Title
Gender and Misallocation in the Labor Market

Author
Haegele, Ingrid

Publication Date
2022

Peer reviewed|Thesis/dissertation

Gender and Misallocation in the Labor Market

by

Ingrid Haegele

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Patrick Kline, Chair
Professor David Card
Associate Professor Christopher Walters
Assistant Professor Supreet Kaur

Spring 2022

Gender and Misallocation in the Labor Market

Abstract

Gender and Misallocation in the Labor Market

by

Ingrid Haegele

Doctor of Philosophy in Economics

University of California, Berkeley

Professor Patrick Kline, Chair

Managers are important for firm productivity (Bloom and Van Reenen, 2007, Lazear, Shaw and Stanton, 2015). Consequently, firms must decide how to assign the right workers to higher-level jobs (Rosen, 1982, Holmstrom and Tirole, 1989). However, little is known about how firms' organizational design affects different groups of workers and whether it can explain why women continue to be underrepresented in higher-level positions. In this dissertation, I collect a unique dataset in collaboration with a large multinational firm in order to provide new insights on the key impediments that women face in their career progression.

In the first chapter, I focus on the role of managers for workers' career progression, which is motivated by the fact that most organizations rely on managers to identify talented workers for promotions. However, managers who are evaluated on team performance have an incentive to hoard workers. I find that talent hoarding reduces workers' applications for promotions by more than half, with particularly large effects on women. Marginal female applications who would have not applied under talent hoarding are twice as likely than their male counterparts to land a promotion and three times as likely to outperform their team at a higher level. Talent hoarding thus contributes to misallocation of talent and perpetuates gender inequality in representation and pay.

The second chapter builds on these findings and attempts to identify when and why gender differences in representation first emerge in the leadership hierarchy. The rich personnel records allow me to construct a new and granular measure of job hierarchy that captures career progression throughout the job ladder. I find that the transition to first-level leadership positions represents a key bottleneck in women's career progression. This early leadership gap is not fully explained by employee characteristics and is driven by internal promotions, not differential entry to or exit from the firm. Women who make it to the first-level leadership level are not less likely to get promoted than men, which contrasts the common notion of a glass ceiling.

The third chapter demonstrates why women are less likely to advance to first-level leadership positions than men. Identifying the drivers of female underrepresentation is difficult, because the promotion gap can arise due to both labor supply and labor demand factors.

By combining rich personnel records and the universe of application and hiring decisions at a large multinational firm, I am able to analyze employees' labor supply decisions separately from the firm's labor demand decisions. I find that women at lower hierarchy levels are less likely to apply for promotions to first-level leadership positions than observationally similar men, but do not experience lower hiring likelihoods than men. Using detailed information on every internal job opening in employees' choice sets, I show that preferences for leading a team are a key determinant of the gender gap in applications for promotions. These gender differences in preferences for team leadership are not fully explained by other factors, including correlated job features such as flexibility and skill requirements or the gender composition of the coworkers associated with a job opening.

Together, these findings imply that the organizational design within firms may have large impacts on gender differences in career progression, deterring high-quality women from climbing the job ladder and contributing to misallocation of talent. While all three chapters use the same setting and data, each chapter is intended to be a stand-alone set of research questions, so the respective description is included within each chapter.

To my family,
for always being there.

Contents

List of Figures

List of Tables

Acknowledgments

I am grateful to my advisors David Card and Patrick Kline for their endless encouragement. Their kindness and dedication to their students made my time at Berkeley very special and I thank them for reminding me that I can make a mark. I thank Chris Walters for the many times he showed me how to think about something from a different perspective and whose patience and creativity made this dissertation so much better. I am grateful to Supreet Kaur who has challenged me to discipline my thinking and who has taught me so much.

I was fortunate to be part of a great community at Berkeley. Many people have been extremely generous with their time and advice. Special thanks go to Stefano DellaVigna, Hilary Hoynes, Ulrike Malmendier, Ricardo Perez-Truglia, Jesse Rothstein, Steve Tadelis, and Dmitry Taubinsky. I am especially grateful to Sydnee Caldwell who knows how to make research rigorous, but also really fun. I also thank Camille Fernandez who has been a constant source of support and kindness. During my time at Berkeley, I was lucky to call many different cohorts my friends. Thinking of Berkeley will always mean missing the endless walks, family dinners, camping trips, and coffee breaks with all of you.

This dissertation is shaped by the desire to understand why women continue to be underrepresented in higher-level positions, which led me on a journey of learning about the internal organization of firms and collecting new data that can help shed light on this question. I am indebted to many people who have shared their perspectives and advocated for trying something new. I would like to thank Leon Jacob whose passion for how to improve HR practices is contagious. This dissertation would not have been possible without out the curiosity, trust, and passion of the firm I have collaborated with and I am especially grateful to Julia, Martina, Matthias, and Nina.

Most importantly, I would like to thank my family, whose unconditional love has been a constant companion. As with many other things in life, it is my dad who is the first one to learn about my research. And it is my mom who makes it very clear that life should be enjoyed independent of research. I thank my grandparents, my aunts and uncles, my many cousins, my parents-in-law and sister for making me feel close to home no matter how far away I am. I am especially grateful to Francis, who has helped me with my first assignment in labor economics and has since never stopped. This adventure would have been so much less fun without his endless optimism, love, and humor.

Chapter 1

Talent Hoarding in Organizations

1.1 Introduction

Firms must continually decide how to allocate workers to jobs, a process which has critical implications for productivity (Rosen, 1982, Holmstrom and Tirole, 1989). Because it is difficult to perfectly observe worker ability, most firms rely on managers to identify talented workers who can be promoted to higher-level positions. However, when a talented worker leaves their team for a promotion, team performance suffers. Since managers are rewarded based on team performance and firms cannot perfectly monitor manager actions, the conflicting interests of manager and firm create the potential for moral hazard (Holmstrom, 1979). A growing body of evidence documents that workers in high-level positions have large impacts on firm performance (Bloom and Van Reenen, 2007, Lazear et al., 2015), implying that managers may create significant efficiency costs if they hoard talented workers rather than promote them.

Ample anecdotal and survey evidence points to widespread talent hoarding in organizations. In a global survey, half of organizations report that managers hoard talent by discouraging worker mobility (i4cp, 2016). A US-based survey finds that workers in one-third of firms feel the need to keep internal applications secret from their managers out of fear of retaliation (KornFerry, 2015). In Germany, 83% of the top publicly listed companies cite managerial talent hoarding as a key friction in their organization (hkp, 2021).[1] Despite the apparent prevalence of talent hoarding and its likely detrimental consequences, very little empirical evidence on talent hoarding exists in economics. Studying talent hoarding empirically is challenging. Managers often hoard talent through hidden actions that are difficult to observe, even in rich datasets. Furthermore, identifying the causal impacts of

[1] Talent hoarding also occurs in science and academia. Zuckerman (2021) documents how Katalin Karikó, a seminal developer of mRNA vaccines, experienced talent hoarding when she decided to leave her lab for a new position. Her advisor Robert Suhadolnik vowed "to do whatever he could to stop his protégée from leaving...he made it clear she had two career choices. 'You can work in my lab or go home,' he told her. Suhadolnik followed through on this threat, telling a local immigration office that she was living in the country illegally and should be deported."

talent hoarding requires plausibly exogenous variation in hoarding.

This study provides the first empirical evidence on talent hoarding and its negative impacts on the efficient allocation of talent in organizations. I combine a rich dataset from a large manufacturing firm with a new identification strategy that leverages quasi-random variation in worker exposure to talent hoarding. When managers learn that they will move to a new position on a different team, they no longer have an incentive to hoard workers on their current team. Thus, manager rotations create a temporary window of time for workers in which they will not be subject to talent hoarding, resulting in an increase in workers' applications for promotions of 123%. I show that this increase in applications is consistent with a series of predictions on talent hoarding and that alternative mechanisms, such as loyalty or manager-worker-specific match effects, cannot account for my results by themselves. Talent hoarding deters high-quality workers from applying who would have performed well in higher-level positions, leading to misallocation of talent within the firm. Because women's applications react more to talent hoarding than men's, women experience greater misallocation effects, exacerbating gender disparities in pay and representation at the firm.

I develop a simple conceptual framework that captures managerial moral hazard and provides a formal definition of talent hoarding. Since managers are both tasked with identifying productive workers and are rewarded solely based on team productivity, managers have an incentive to hoard talent by preventing workers from seeking promotions. This framework predicts that managers engage in more hoarding when their workers are more productive and when worker departure is costlier. In addition, by shrinking the pool of workers identified for potential promotion, talent hoarding creates misallocation of talent in the firm. These predictions guide my empirical analysis.

To empirically test for talent hoarding, I collect a unique combination of personnel records and internal job application data from a large manufacturing firm that employs over 200,000 workers. In order to examine internal career progression to higher-level positions, I focus my analysis on the firm's largest internal labor market, consisting of over 30,000 white-collar and management employees in Germany. I demonstrate that the firm is similar to other large firms in terms of its workforce and organizational design. As in many other large firms, employees who want to switch to a new position or to be promoted are required to apply for the position. Since most workers in the same team are at similar hierarchy levels, promotions typically require out-of-team transitions.

The dataset I assemble offers several key advantages that allow me to test the predictions that follow from the conceptual framework. First, by combining personnel records with the universe of application and hiring decisions at the firm, I am able to assess the extent to which talent hoarding deters applications that would have been successful. Second, two novel measures of worker visibility constructed from the firm's internal HR databases allow me to infer managers' propensities to hoard talent by measuring the extent to which they systematically suppress the visibility of workers on their team. Without such data, directly measuring talent hoarding is empirically challenging, because by definition hoarding involves hidden actions. Third, I construct a granular measure of internal job hierarchy that identifies transitions to higher-level positions with more job responsibility. The hierarchy measure enables a direct test of whether talent hoarding causes misallocation of talent by evaluating whether high-quality workers are deterred from moving to higher-level positions in which they would have been more productive.

My research design leverages quasi-random variation in workers' exposure to manager rotations. When a manager learns that they will move to a new position on a different team, they no longer have an incentive to hoard workers on their current team. For workers whose manager will soon rotate, this creates a temporary window of time during which they are not subject to talent hoarding. Therefore, analyzing manager rotations allows me to study the impacts of talent hoarding without requiring direct measurement of talent hoarding behavior. Empirically, rotation effects are large, effectively doubling worker applications in the same quarter. I demonstrate that these effects can be interpreted as reflecting the causal effect of a manager leaving her team. A placebo test shows that manager applications for job rotations only increase worker applications if managers are successful and actually leave the team.

Following the predictions from the conceptual framework, I provide evidence indicating that talent hoarding is a key mechanism underlying the observed impacts of manager rotations. I first show that rotations have larger effects on workers who were previously subject to greater levels of talent hoarding, as captured by three dimensions of heterogeneity: worker quality, the costliness of worker departure, and managers' propensities to hoard talent. I then leverage the rich job application data and show that manager rotations disproportionally increase applications that under talent hoarding carry a greater risk of retaliation, either because managers are likely to find out about the application or because applications are unlikely to be successful. Moreover, I document that manager rotations only affect internal job transitions within the firm that are subject to talent hoarding, but not external job transitions out of the firm, which managers are not able to influence.

A potential threat to the interpretation of the impacts of manager rotations as representing the impacts of talent hoarding is that manager rotations may affect worker applications through additional channels. For instance, workers may refrain from applying for a new position because of loyalty towards their manager or because manager-worker-specific match effects make their current position particularly appealing. In addition, worker applications may result from team-level shocks that are correlated with manager rotations, such as unpleasant working conditions, bad news about the future outlook of the team, or the completion of a major milestone. In a series of tests, I find that these channels alone are not able to account for the observed rotation effects, suggesting that talent hoarding does play a role in deterring worker applications.

My findings indicate that talent hoarding causes misallocation of talent by reducing the quality and performance of promoted workers. To analyze misallocation, I focus on major promotions, such as transitions from individual contributors to team leader positions, that reflect meaningful changes in job responsibility. Manager rotations increase worker applications for major promotions by 123%, indicating that talent hoarding deters a large group of workers from applying for promotions. To quantify how successful deterred applications would have been, I instrument for workers' applications with manager rotations. Marginal applicants, who would not have applied in absence of a manager rotation, are similarly likely as average applicants to land a promotion and are likely to subsequently outperform their teams in higher-level positions. A complier analysis finds that marginal applicants are positively selected in terms of their educational qualifications and past performance. These findings suggest that in addition to reducing the number and the quality of applicants for higher-level positions, talent hoarding lowers team performance at these levels.

I find that talent hoarding has disparate impacts by gender. Talent hoarding deters a larger share of female applicants from applying for major promotions compared to men. Female marginal applicants are twice as likely to land a major promotion than males, implying that talent hoarding is more consequential for women's career progression. Conditional on landing a promotion, women are almost three times as likely as their male counterparts to perform well in their new positions, suggesting that the firm may be failing to realize potential productivity gains by not enabling talented women to progress to higher-level positions. Female marginal applicants are much more qualified than males in terms of their educational qualifications and past performance, indicating that talent hoarding affects women at a higher part of the quality distribution compared to men.

Talent hoarding exacerbates gender inequality in representation and pay. When comparing potential outcomes for marginal applicants, I find that increasing applications through manager rotations is much more beneficial for women than for men, reducing the gender representation gap by 91% and the gender pay gap by 77% within one year. The disparate impacts of talent hoarding by gender are not driven by differential treatment by managers. Rather, a survey of the firm's employees suggests that male and female workers react differently to talent hoarding. In line with the literature on gender differences in preferences (Bertrand, 2011), the survey finds that women in the firm place more value on preserving a good relationship with their manager and rely more on managers' career guidance when making application decisions.

A number of factors suggest that talent hoarding is very likely to manifest similarly in other organizations. The firm I study is similar to other large firms in Germany both in terms of the characteristics of its workforce and its internal organization, in that it is standard that managers are tasked with identifying talented workers, but are neither monitored nor rewarded in doing so (hkp, 2021). Companies across the world report that talent hoarding is commonplace, creates barriers to talent allocation, and occurs through many of the same managerial behaviors that are documented in this study (i4cp, 2016, KornFerry, 2015, Matuson, 2015, Sullivan, 2017).

This study contributes to two broad strands of research on organizations. Prior theoretical research has hypothesized that managers engage in self-interested behavior (Holmstrom and Tirole, 1989), largely focusing on managers' misaligned incentives in the context of biased performance evaluations (Milgrom and Roberts, 1988, Prendergast and Topel, 1996, Fairburn and Malcomson, 2001). While research studying internal labor markets has documented the importance of incentive provision in organizations (Gibbons and Waldman, 1999, Prendergast, 1999), little attention has been paid to how managers' incentive problems may affect the efficiency of job assignments. One notable exception is theoretical work by Friebel and Raith (2013) who show that different organizational designs may change managers' incentives to train subordinates and accurately represent their abilities. My study provides the first empirical demonstration of a costly moral hazard problem among managers that affects the efficient allocation of talent within organizations.

Second, a large empirical literature in economics studies the impacts of managers on firm outcomes. The majority of this literature has focused on upper management, and in particular on documenting the impacts of CEOs on firm performance (Bertrand and Schoar, 2003, Bennedsen, Nielsen, Pérez-González and Wolfenzon, 2007, Malmendier and Tate, 2009, Bennedsen, Perez-Gonzalez and Wolfenzon, 2020). An emerging body of work uses detailed data on managers and workers to show that even managers at lower levels of the firm hierarchy have large impacts on worker outcomes, including worker productivity (Lazear et al., 2015, Frederiksen, Kahn and Lange, 2020, Fenizia, Forthcoming), turnover (Hoffman and Tadelis, 2021), and career progression (Kunze and Miller, 2017, Cullen and Perez-Truglia, 2019, Benson, Li and Shue, 2021). This study adds to these findings by uncovering talent hoarding as an important mechanism that influences managers' impacts on firms and workers. By demonstrating that talent hoarding has meaningful impacts on career progression, this study also contributes to both theoretical and empirical work seeking to understand the dynamics of internal labor markets (Waldman, 1984, Milgrom and Oster, 1987, Baker, Gibbs and Holmstrom, 1994, Benson, Li and Shue, 2019, Huitfeldt, Kostol, Nimczik and Weber, 2021).

The remaining sections are organized as follows. Section 1.2 provides survey evidence and introduces a simple framework that offers a formal definition of talent hoarding. Section 1.3 describes the institutional setting and novel data. Section 1.4 presents the empirical strategy centered around manager rotations. Section 1.5 demonstrates the impacts of talent hoarding on worker applications. Section 1.6 presents results on the efficiency costs of talent hoarding with respect to talent allocation. Section 1.7 provides suggestive evidence unpacking how talent hoarding effects arise. Section 1.8 concludes.

1.2 Background and Conceptual Framework

This section presents anecdotal evidence on the prevalence of talent hoarding. I conduct a large-scale survey at the firm I study, which illustrates how talent hoarding deters workers' career progression. Building on this intuition, a simple conceptual framework offers a formal definition of talent hoarding as well as a set of predictions that guides my empirical analysis.

Talent hoarding is widespread and occurs in a variety of settings. In a survey of 665 global organizations, covering both the private and public sector, half of organizations report that their managers hoard talent by discouraging worker mobility (i4cp, 2016). In Germany, 83% of the top publicly listed companies cite talent hoarding as a crucial friction in their organization (hkp, 2021). Talent hoarding appears to be highly salient to workers. In one-third of US firms, workers feel the need to keep applications secret from their managers out of fear of retaliation (KornFerry, 2015).

News media outlets and industry publications present anecdotal evidence describing how managers hoard talent. A 2015 Forbes article observes that managers who hoard talent "never recommend...people for a promotion in another department." (Matuson, 2015). The industry publication Talent Management & HR writes in 2017 that "hoarding managers, in order to reduce the internal visibility of their top team members, may purposely restrict them from serving on task forces and outside-of-function committees." (Sullivan, 2017). Other talent hoarding strategies are described as underrating potential for higher-level positions or threatening workers who try to leave the team.

1.2.1 Evidence from a Survey within the Firm

To provide the first detailed evidence on the dynamics of talent hoarding in organizations, I conduct a large-scale survey at the firm I study that captures both manager and worker behavior.

All employees in my sample were invited via e-mail by the firm's HR department and were asked to provide their perspectives on the internal labor market at the firm. Employees described challenges regarding their internal career progression both in the form of free-text responses and in multiple-choice answers. An abbreviated version of the survey instrument is presented in Appendix Section B.0.2. The survey received over 15,000 responses, yielding a 50.0% response rate. Respondents are similar to non-respondents in terms of demographics (Appendix Table B.1).

Respondents report a variety of different actions through which managers hoard talent, which include suppressing public signals of worker ability, restricting access to high-visibility projects or training, and explicitly discouraging workers from applying to internal positions. Appendix Table A.2 provides selected quotations. In addition, when asked to state the biggest challenge to their internal career progression, the modal answer (provided by 22% of workers) is managers' limited support for career progression, such as refusal to assist in career planning and denial of requests to participate in development programs that would increase workers' visibility outside the team.

Not only do many employees report managers trying to prevent workers from pursuing promotions, managers' actions appear to strongly influence workers' application decisions. 41% of respondents indicate that they are afraid to apply to internal positions, fearing negative repercussions if managers find out about the application. 25% of workers state that they feel the need to ask for managers' permission before applying for an internal job opening, even though the firm's internal policies are meant to enable workers to initiate an application on their own. 16% of respondents indicate that applying away from the team is seen as disloyal. These findings suggest that fear of retaliation represents a key dimension through which talent hoarding deters worker applications.

In the survey, managers acknowledge the existence of talent hoarding. 32% of managers report that negative repercussions follow when managers find out about internal applications. 28% of managers state that workers should ask their managers for permission before applying. Appendix Table A.3 provides anecdotal evidence of managers' descriptions of talent hoarding, offering direct evidence that misaligned incentives lead managers to hoard talent.

One manager explains, "Managers pursue their own goals and often prevent further development of workers, because they are not rewarded for developing talent.". Another manager reports that "Selfish managers are not willing to promote or recommend subordinates to other areas of the firm, even if that would add value to the firm.".

1.2.2 A Simple Framework of Talent Hoarding

To formally define talent hoarding, consider a firm that employs two types of agents, managers m and workers i. For simplicity, teams are composed of one manager and one worker. Workers are characterized by a marginal productivity α_i drawn from some known distribution G. The firm seeks to efficiently allocate talent to maximize total firm productivity by choosing which workers to promote to managerial positions. Consistent with Rosen (1982), productivity is maximized when high-ability workers are assigned to high-level positions. Thus, in the absence of any constraints, the firm would fill a new managerial vacancy with the most productive worker and would fill the worker's vacated position by hiring a worker from outside of the firm (i.e. a random draw from G).

In practice, firms neither perfectly observe worker productivity, nor do they know which workers would accept a promotion. Accordingly, managers are tasked with identifying high-productivity workers and encouraging them to seek promotions. However, if a high-productivity worker is promoted and leaves their team, that team incurs a productivity loss. Managers, who observe the productivity of their workers, are compensated according to team productivity, creating a conflict of interest between the firm and managers. Talent hoarding is defined as the actions taken by managers that lower the likelihood that a worker applies for and receives a promotions.[2]

The one-period framework proceeds as follows. A managerial vacancy opens exogenously. M managers observe the productivity of the worker on their teams and decide to what extent to engage in talent hoarding. Based on managers' choices, workers decide whether to apply for a promotion. The firm observes noisy signals of worker productivity (e.g. by conducting interviews) and chooses the worker with the highest signal to fill the vacancy. The promoted worker's previous position is replaced with an outside hire. Team productivity is realized and managers are compensated.

Let $\beta \in [0, \infty)$ index the degree of talent hoarding chosen by a manager, with 0 representing no talent hoarding.[3] Let $q(\alpha_i, \beta)$ denote the equilibrium probability that a worker with

[2] Respondents to the survey discussed in Section 1.2.1 report that managers can deter workers from applying by explicitly discouraging or threatening them, underrating worker ability, and restricting access to high-visibility projects or training. In theory, the firm could offer managers a promotion-contingent contract to resolve the misaligned incentives. In practice, firms generally do not compensate managers for promoting their workers, plausibly because of the practical challenges associated with these contracts (discussed in detail in Friebel and Raith (2013)).

[3] Survey responses presented in Appendix Table A.2 indicate that suppression of potential ratings and pressure to refrain from applying are common examples of talent hoarding that can be represented by β. In Section 1.3.3, I construct a direct measure of talent hoarding by comparing the private performance ratings to the public potential ratings that managers assign to workers, where β can be interpreted as reflecting the disparity between these signals.

productivity α_i gets promoted, conditional on applying for a promotion. This promotion probability increases with worker productivity ($\frac{\partial q}{\partial \alpha_i} > 0$), but decreases in the level of talent hoarding ($\frac{\partial q}{\partial \beta} < 0$), and reflects the noisy signal received by the firm. Thus, one can interpret talent hoarding as impacting workers through the likelihood that they get promoted.[4] Furthermore, I assume that the effect of talent hoarding on the conditional promotion probability is larger for more productive workers ($\frac{\partial^2 q}{\partial \beta \partial \alpha_i} < 0$). This assumption relates to the noisy signals of applicant productivity observed by the firm. Intuitively, since low-productivity workers are less likely than high-productivity workers to generate a favorable signal, there is less scope for talent hoarding to lower the likelihood that a low-productivity worker gets chosen among multiple applicants. One example in which this would be the case is if a firm employs a two-part screening strategy in which it chooses a subset of applicants to interview based on the CVs of all applicants. A very low-productivity worker may never clear the bar to be interviewed. Thus, the manager can do little to further lower their hiring likelihood.

Workers decide whether to apply by weighing expected costs and benefits. Let b denote the benefits of a successful application and c denote the costs of applying. Workers apply if

$$q(\alpha_i, \beta)b \geq c + \varepsilon_i \tag{1.1}$$

where $\varepsilon_i \sim \Psi$ captures worker-specific heterogeneity, with Ψ known to the manager. Therefore, from the manager's perspective, the probability that the worker leaves the team is given by

$$p(\alpha_i, \beta) = q(\alpha_i, \beta)\Psi\big(q(\alpha_i, \beta)b - c\big)$$

It follows that talent hoarding reduces the probability that workers leave the team (i.e. $\frac{\partial p}{\partial \beta} < 0$).[5]

Managers optimize by choosing their level of talent hoarding β. If a worker leaves the team for a promotion, the firm hires a worker of unknown ability ($\alpha_j \sim G(.)$ with $E[\alpha_j] = \bar{\alpha}$) from outside the firm. Consequently, a high-productivity worker getting promoted out of the team is likely to decrease team productivity. Without compensation for promoting workers, managers have an incentive to engage in talent hoarding by reducing workers' likelihood of promotion. However, managers incur increasing and convex costs from talent hoarding, which vary in their magnitude by manager according to the parameter $\phi_m > 0$.[6] Thus, managers solve the following problem:

[4] In practice, workers report that managers diminish their visibility, thus lowering their promotion prospects. In theory, talent hoarding can also operate through the cost of applying, which would yield similar predictions.

[5] For simplicity, this framework does not distinguish between different worker types. Talent hoarding may exacerbate between-group differences in promotions if it has a differential effect on workers' application decisions.

[6] The utility costs that managers experience when hoarding talent can be interpreted as consequence of manager altruism, in line with Hoffman, Kahn and Li (2018) who motivate why managers might value making hiring decisions in the interest of the firm and not in their best self-interest. Alternatively, such costs could arise from the probability of detection, for instance in the form of reputation costs.

$$\max_{\beta} \, (1 - p(\alpha_i, \beta))\alpha_i + p(\alpha_i, \beta)\bar{\alpha} - \frac{\phi_m}{2}\big(p(\alpha_i, 0) - p(\alpha_i, \beta)\big)^2$$

This optimization problem yields the following first-order condition, which provides a formal definition of talent hoarding as well as predictions with respect to the realized level of talent hoarding:

$$p(\alpha_i, 0) - p(\alpha_i, \beta^*) = \frac{1}{\phi_m}(\alpha_i - \bar{\alpha})$$

Definition. (Talent hoarding) When worker i's productivity exceeds the expected productivity of an outside hire, the manager optimally hoards talent by choosing $\beta^* > 0$. As a result, the likelihood that a worker leaves the team is lower relative to $\beta = 0$.

Prediction 1. (Worker heterogeneity)
High-productivity workers (i.e. high α_i) experience more talent hoarding.

Prediction 2. (Team heterogeneity)
Workers that are more difficult to replace (i.e. if $\bar{\alpha}$ is lower) experience more talent hoarding.

Prediction 3. (Manager heterogeneity)
Talent hoarding is greater among managers with low utility costs of hoarding (i.e. low ϕ_m).

In addition, workers' decision rule implies that talent hoarding reduces the number of applicants and the quality of the applicant pool, limiting the firm's ability to promote a high-productivity worker and generating misallocation of talent.

Prediction 4. (Number of applicants)
If workers face less talent hoarding, they are more likely to apply for a promotion.

$\Pr[i \text{ applies}|\beta = \beta_1]$¿$\Pr[i \text{ applies}|\beta = \beta_2]$ for $\beta_1 < \beta_2$

Prediction 5. (Composition of applicants)
Talent hoarding has larger impacts for high-productivity workers. Therefore, higher levels of talent hoarding lead to a lower-quality applicant pool:

If $\alpha_1 < \alpha_2$ and $\beta_1 < \beta_2 \implies \frac{\Pr[i \text{ applies}|\alpha_2,\beta_1]}{\Pr[i \text{ applies}|\alpha_1,\beta_1]} > \frac{\Pr[i \text{ applies}|\alpha_2,\beta_2]}{\Pr[i \text{ applies}|\alpha_1,\beta_2]}$

Appendix Section C contains formal derivations of these predictions.

1.3 Setting and Data

My analysis relies on a unique combination of personnel records and internal job application data from a large manufacturing firm. I show that the firm, one of the largest in Europe, is comparable to other firms in its sector. I then use the firm's data to construct a panel dataset that links workers to their managers and includes workers' application and hiring histories at the firm.

1.3.1 Firm Overview

I collect rich data on over 30,000 white-collar and management employees from a large manufacturing firm. This anonymous firm is one of the largest manufacturers in Europe and employs over 200,000 workers around the world, the plurality of which work in Germany.

I restrict my sample to Germany because it represents the largest internal labor market at the firm. The firm operates in many other countries, including the United States. The firm's establishments outside of Germany share many features in common with those in Germany, including organizational design and internal labor market policies (e.g. application systems, widespread use of performance and potential ratings). Because I am interested in career progression to higher-level positions, my analysis focuses on employees in white-collar and management positions (i.e. employees that are either already in or could ultimately be promoted to managerial positions). There are over 200 occupations represented in this sample, ranging from technical positions in engineering to support functions in HR and finance.

Table 1.1 describes my analysis sample, which consists of over 300,000 employee-by-quarter observations from 2015 to 2018.[7] Women represent 21% of employees in the sample, stemming from the underrepresentation of women in technical occupations. Employees' educational qualifications are high, a result of restricting the sample to white-collar and management employees. The average employee holds a Bachelor's degree and 92% of employees work full-time. Tenures at the firm are long, with an average of 13 years, highlighting the importance of internal career progression for employees' long-term income. Managers (i.e. those that lead a team) comprise 19% of the sample.

The demographics of the employees at the firm are comparable to other large manufacturing firms in Germany. In Appendix Table A.1, I compare employees in my sample to those employees in large manufacturing firms in the BiBB, a representative survey of the German workforce conducted in 2018. I find very similar patterns with respect to most employee characteristics (e.g. gender, age, German citizenship, marital and family status).

The firm also resembles other large firms with respect to the design of its internal labor market (hkp, 2021). As in most large German firms, employees who want to switch to a new position or to be promoted are required to apply for the position using a centralized online job portal at the firm. All job openings are posted to the job portal, where openings are visible to all employees. Applications through the portal typically take less than five minutes to complete. While employees can choose to apply to multiple positions at the same time, the median applicant applies to only one position in a given quarter. Callback and hiring decisions are also recorded in the job portal. Appendix Figure A.1 depicts the appearance of the firm's online job portal.

The internal labor market is comprised of competitive openings for new positions, much like the external labor market. Only 25% of applications are successful and the internal labor market is both spatially and interpersonally diffuse. The firm operates in over 50 cities in 250 establishments throughout Germany and one-third of internal applications are for positions in

[7]To maintain confidentiality, I do not disclose the exact number of employees in my sample.

a different city. In 93% (99%) of cases, applicants (applicants' current supervisors) have not previously worked with the hiring manager of the position they are applying for, indicating that application and hiring decisions are distinct in the internal labor market.

Because most teams are small, consisting of one manager and six workers on average, and because most workers in a team are at similar hierarchy levels, workers must leave their team to move up the job ladder. In the data, 97% of applications are to positions outside of a worker's current team. Thus, managers who encourage their workers to pursue promotions lose team members.

1.3.2 Personnel Records and Application Data

To empirically test for talent hoarding, I assemble a unique dataset that combines the firm's internal personnel records from 1998 to 2020 with the universe of application and hiring decisions from 2015 to 2020. The personnel records capture over 30,000 employees, and the application data cover over 16,000 job openings and over 200,000 external and internal applicants. I use a five-step matching algorithm to link personnel records and application data, which matches over 90% of individuals (see Appendix Section B.0.1 for more details). In my main analysis sample, I restrict to employees active at the firm from 2015 to 2018, for whom I can observe outcomes through 2020, yielding a sample of over 300,000 employee-by-quarter observations.

The personnel records contain a large set of employee characteristics, such as age, educational qualifications, tenure, and family status. The records also contain detailed position characteristics, such as position titles, leadership responsibility, and the reporting distance to the CEO. In addition, the records capture workers' assignments to teams and managers. I use this information to characterize manager behavior and construct measures of manager quality based on past outcomes of managers' team members (e.g. promotion, turnover, absenteeism). 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. I supplement this data with payroll data, capturing employees' working hours, earnings, and bonus payments. Finally, I collect information from the firm's talent management system that includes worker evaluations, such as performance and potential ratings, and nominations to succession lists.

I use the richness of these data to account for factors unrelated to talent hoarding that may influence workers' career progression. Unless otherwise noted, all analyses include the following set of controls: worker demographics (female, age, German citizenship, educational qualifications, marital status, family status, parental leave, firm tenure), position characteristics (position type, division, functional area, location, full-time status, hours, number of direct reports), performance and potential ratings, and past mobility at the firm. These characteristics allow me to incorporate key determinants of career progression that are not available in settings studied in prior research.

I use application and hiring decisions to identify the impacts of talent hoarding on applications and promotions. Because all job openings are posted to the centralized portal and all applications and hiring decisions are required to be submitted through the portal, I observe the outcome of each application in terms of rejections, interview callbacks, and subsequent hiring outcomes. These features allow me to construct a panel dataset of employees' application and hiring histories at the firm from 2015 to 2020. Separately measuring applications and hiring outcomes is important given that one would expect workers with a higher application threshold (e.g. because they fear manager retaliation) to have higher hiring likelihoods if they apply. If that is the case, analyzing impacts on promotions alone would underestimate the effects of talent hoarding.

1.3.3 Direct Measures of Talent Hoarding

I infer managers' propensities to hoard talent using two novel measures of worker visibility. Although my main empirical tests for talent hoarding do not rely on directly measuring talent hoarding, doing so offers a useful test of the predictions that emerge from my conceptual framework in Section 1.2.2.

Talent hoarding is difficult to measure because it typically occurs through interpersonal interactions, as indicated by survey responses discussed in Section 1.2.1. The survey suggests that one common way in which managers limit workers' opportunities to leave the team is by suppressing worker visibility. To capture talent hoarding in the data, I identify managers' systematic suppression of worker visibility based on two measures of worker visibility that I collect from the firm's HR databases: potential ratings and nominations to succession lists.

My primary approach to measuring talent hoarding identifies systematic differences between potential ratings and performance ratings. Performance ratings are meant to provide task-specific feedback to workers about their past performance in their current position. Potential ratings are designed to inform the firm about a worker's future potential for higher-level positions and thus are meant to identify workers who would be likely to perform well if they were promoted. Both ratings are conducted simultaneously by a worker's direct supervisor and are very similar to common worker assessments, such as the nine-box grid, that are used by many organizations across the world (Cappelli and Keller, 2014). An important distinction between performance and potential ratings is the extent of their visibility outside of a worker's team. As in many other firms, performance ratings are private signals to the worker and are not shared with other units in the firm. If a worker applies for a job in a different unit, that unit will not be able to access the worker's past performance ratings. In contrast, potential ratings are public signals of worker talent and are widely circulated within the firm. The firm's HR department regularly circulates lists of high-potential workers to be considered for promotion, making them highly visible.

Intuitively, a manager who wants to hoard talent should give workers lower (public) potential ratings relative to their (private) performance ratings.[8] A survey of the firm's em-

[8]The ideal measure of talent hoarding would compare a private and a public signal of the same type of rating. While the measures required for such a comparison do not exist, I find that 86% of employees who are rated by their manager as having potential for a higher-level position actually receive a high performance

ployees suggests that this method of talent hoarding is commonplace.[9] For a given worker, the difference between performance and potential ratings may reflect worker-specific factors that are unrelated to talent hoarding; however, because managers have discretion in conducting these evaluations, comparing systematic differences between the ratings across workers can capture manager behavior.[10]

My first measure of talent hoarding is defined by the difference between the actual and predicted potential ratings a manager assigns to their workers. Specifically, I take the residuals from the OLS estimation of the following regression of the potential rating given to worker i in quarter t by manager m on their performance rating and other worker characteristics:

$$\text{potential}_{imt} = \beta_1 \text{performance}_{imt} + \beta_X X_{it} + \beta_t + \varepsilon_{imt} \tag{1.2}$$

In Equation 1.2, X_{it} denotes the vector of controls described in Section 1.3.2. I compute the average difference (over workers and quarters) between predicted and actual potential ratings for each manager. I classify a manager as having a high propensity to hoard talent if this difference is in the bottom tercile of the manager distribution (Appendix Figure B.01).

I conduct a number of empirical exercises that support the validity of this measure of talent hoarding. First, the underrating of potential captured by this measure does not represent managers' correct assessment of workers' future performance. When managers who have a high propensity to hoard talent rotate, underrated workers not only experience increases in applications and promotions, they are also likely to perform well at higher levels, demonstrating that their initial low potential rating was inaccurate. Second, the talent hoarding measure is not strongly correlated with managers' ability to assess talent, suggesting that differences between performance and potential ratings are unlikely to be a result of managers' involuntary mistakes. Third, this measure is reasonably stable over time, supporting the systemic notion of talent hoarding the measure is meant to capture. Fourth, this measure of talent hoarding is highly correlated with workers' realized visibility at the firm, confirming that managers' suppression of public signals has a meaningful impact. Appendix Section B.0.3 presents additional details along with multiple additional validity exercises.

My secondary approach to measuring talent hoarding relies on measuring visibility in the form of nominations to succession lists. As in many large organizations, the firm compiles lists of three to five candidates who are potential successors for about one-fifth of the positions in my sample. The lists are assembled by HR employees who search for suitable candidates across the firm. Workers' appearance on such a list represents a measure of their visibility

rating once they get promoted. The strong correlation between a worker's current potential rating and their future performance ratings in higher-level positions suggests that this approach carries much of the information of the ideal comparison.

[9] For instance, one worker states, "Supervisors suppress potential ratings because of fear that employees will leave their current position for a promotion." (Appendix Table A.2).

[10] Illustrating the importance of manager discretion in this setting, Benson et al. (2021) find in contemporaneous work that managers in a large retailer overrate men's potential compared to women's, possibly as a reaction to men's higher turnover risk.

outside of the team. If a manager is successful at hoarding talent, worker visibility should be low, and thus their likelihood of appearing as a nominee on a succession list should also be low.

To construct the measure of talent hoarding based on succession lists, I estimate a version of Equation 1.2 to compute the difference between actual nominations and predicted nominations. I then classify managers as high-propensity and low-propensity talent hoarders, defined as those in the bottom and top terciles of this difference. When testing for talent hoarding, I use this measure as a complement to the primary measure based on potential ratings.

1.3.4 A Granular Measure of Job Hierarchy

In order to assess whether talent hoarding leads to misallocation, I implement a test of whether high-ability workers are underpromoted to high-level positions due to talent hoarding. This test requires a measure of internal job hierarchy. While previous research studying internal hierarchies has typically used occupation or position titles, this approach is not well-suited to testing for misallocation. Since 26% of workers in my sample share the same position title with their supervisor or their supervisor's supervisor, the standard approach would miss granular differences in the job hierarchy and thus likely underestimate the efficiency losses of talent hoarding.

To overcome these challenges, I apply the methods developed in Chapter 2 to define a granular measure of job hierarchy. The key advance provided by Chapter 2 is to form a measure of job hierarchy by combining three distinct dimensions of leadership responsibility: the number of cumulative reports, the reporting distance to the CEO, and the managerial autonomy of a position. The hierarchy ranking is the first principal component of these three dimensions, providing a consistent ordering of all positions at the firm. I define major promotions as increases in the hierarchy index of 20 or more. These transitions typically reflect meaningful increases in leadership responsibility, such as from working as an engineer on a team to managing other engineers. Section 2.3 in Chapter 2 provides further details on the construction and validity of the hierarchy measure. Observed job transitions that represent typical steps in the job ladder are well-described by the hierarchy measure, such that higher levels are associated with more senior positions. The hierarchy measure is strongly correlated with earnings, but is more effective at discerning between hierarchy levels. Because the hierarchy measure is not constructed using pay or salary grades, it also allows me to study how forgone promotions affect pay inequality at the firm.

1.4 Empirical Strategy

My research design leverages quasi-random variation in workers' exposure to manager rotations. When a manager learns that they will move to a new position on a different team, they no longer have an incentive to hoard talented workers on their current team. For workers whose manager will soon rotate, this creates a temporary window of time during which they will not be subject to talent hoarding. Therefore, analyzing manager rotations allows me to study the impacts of talent hoarding without requiring direct measurement of talent hoarding behavior.

In the notation of the conceptual framework presented in Section 1.2.2, when a manager rotates, they temporarily cease to hoard talent by setting $\beta^* = 0$. In practice, the most likely channel through which manager rotations impact worker outcomes in the short-term is by lifting the threat of retaliation. While managers can hoard talent in other ways, such as by underrating worker ability or preventing workers from participating in high-visibility projects or training, the impacts of the cessation of these types of talent hoarding likely require more time to manifest. Even if managers start allowing workers to participate in high-visibility projects, it likely takes time for individuals outside of the team to learn about these workers. Similarly, worker evaluations only occur once or twice per year and therefore manager rotations need not immediately increase potential ratings.

I analyze 1,359 manager rotations by 1,276 unique managers. I define a manager rotation as an instance in which a worker's direct supervisor leaves their team to make an internal job transition within the firm. Restricting attention to internal transitions serves to isolate manager-induced variation that is plausibly orthogonal to worker characteristics and team outcomes. Internal rotations are routine and encouraged by the firm as part of managers' career progression. In contrast, instances in which managers leave the firm through a voluntary exit, layoff, or retirement are likely correlated with other factors that may affect worker outcomes.

During the four-year study period, 20% of managers rotate at least once. Rotations do not occur on a regular schedule, and workers cannot easily predict when managers will leave the team. Appendix Figure A.2 documents large variation in the time that managers spend in a position before rotating. Managers must apply through the firm's application system to rotate. To encourage smooth transitions, the firm's official policy is that managers must inform their teams as soon as possible when they accept a new position. On average, managers learn about their new position two to three quarters before they rotate, at which point they inform their teams about the rotation.[11]

I analyze the effect of manager rotations on worker applications by estimating a linear model for workers' internal application choices using an OLS regression of the following form:

$$\text{Applied}_{it} = \delta_1 \text{Rotation}_{it} + \delta_X X_{it} + \delta_t + u_{it} \tag{1.3}$$

Applied_{it} and Rotiation_{it} are indicators that worker i in quarter t applies for an internal job opening and experiences a manager rotation, respectively. X_{it} includes a broad set of

[11]In the data, over 80% of managers find out about their job transition more than one month in advance, with many applying and accepting offers for new positions at the firm more than six months before actually rotating.

worker and position controls.[12] Section 1.5.2 and Section 1.5.3 provide evidence in support of a causal interpretation of δ_1 and of an interpretation of manager rotations as capturing the impacts of talent hoarding.

1.4.1 Instrumenting for Applications with Manager Rotations

To assess whether talent hoarding causes misallocation of talent, I instrument for worker applications with manager rotations and estimate the marginal probability of landing a promotion, which provides a direct measure of the firm's willingness to promote marginal applicants. Equation 1.4 represents the reduced-form model for the effect of manager rotation on workers' hiring outcomes:

$$\text{Hired}_{it} = \theta_1 \text{Rotation}_{it} + \theta_X X_{it} + \theta_t + \epsilon_{it} \tag{1.4}$$

Hired_{it} is a binary indicator, which is always zero if workers do not apply. The IV-estimator divides the reduced-form effect of manager rotation, θ_1, by the first-stage effect on application choice, δ_1:

$$\beta_{IV} = \frac{\theta_1}{\delta_1} \tag{1.5}$$

I estimate β_{IV} by two-stage least squares, which can be interpreted as the local average treatment effect (LATE), defined as the effect of applications on hiring outcomes for workers induced to apply by manager rotations. Because workers can only get hired if they apply and there are no always takers, LATE equals the treatment effect on the treated (TOT). Interpreting β_{IV} as the LATE requires four assumptions to hold: relevance, independence, exclusion, and monotonicity (Angrist and Imbens, 1995, Angrist, Imbens and Rubin, 1996). I now provide evidence in support of these assumptions.

Relevance.—Section 1.5.1 shows that manager rotations almost double worker applications.

Independence.—A key threat to my research design is that workers who experience a manager rotation might differ in their hiring likelihood from workers who do not. This possibility could arise if manager rotation is not as good as randomly assigned. To evaluate the independence assumption, I test for balance across worker characteristics. Panel B of Appendix Table A.7 shows that workers in teams that experience a manager rotation are observationally similar to those that do not. This similarity holds with respect to worker demographics, such as age, tenure, and marital status, as well as for workers' career trajectories leading up to the manager rotation, such as past earnings growth, absenteeism, applications, and

[12] These controls, described in Section 1.3.2, include female, age, German citizenship, educational qualifications, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, number of direct reports, performance and potential ratings, and past mobility at the firm.

internal job transitions. In unreported results, I find no substantial differences in manager attributes, further supporting the random assignment assumption.[13]

Exclusion.—The exclusion restriction in my setting requires that manager rotation does not affect hiring outcomes other than through workers' decisions to apply. This assumption would be violated if departing managers intercede in workers' subsequent hiring outcomes. Such a violation could occur if departing managers take their workers with them to their new position or replace themselves with workers in their team. However, this occurs in less than 3% of applications, suggesting that departing managers generally do not make subsequent hiring decisions for their subordinates. Alternatively, departing managers may try to influence hiring outcomes by reaching out to hiring managers. I construct a measure of manager ties based on whether they have previously worked with someone in the same team. However, close formal ties between the departing manager and the hiring manager are very rare. For over 99% of applications, the current supervisor has not previously worked with the hiring manager. I test for the influence of more distant formal ties, such as the rotating managers ever having worked in the same location or functional area that the posted job is located in, which increases the probability that the departing manager has ever interacted with the hiring manager. However, I do not find any evidence that relative hiring rates are higher for job openings to which managers have closer ties (Panel A of Appendix Table A.8).

Another violation of the exclusion restriction could occur if manager rotation increases worker qualifications, making them more likely to get hired. If this channel is quantitatively important, one would expect to see larger hiring effects for workers under managers with high managerial quality or who have been exposed to the manager for longer. I measure manager quality using past leave-out means of three team outcomes: promotions, turnover, and absenteeism. Panel B of Appendix Table A.8 documents that the marginal hiring probability is similar for high-quality vs low-quality managers. Panel C of Appendix Table A.8 shows that marginal hiring probabilities are comparable across exposure length, even for workers who have been with the manager for only one quarter, suggesting that such a channel is unlikely to violate the exclusion restriction.[14]

Monotonicity.— I do not find any evidence suggesting that my results are biased by the existence of a large defier population. In my setting, defiers are individuals who would have applied in the absence of a manager rotation, but whose unobserved propensity to apply is reduced by the instrument. Following Arnold, Dobbie and Yang (2018) and Bhuller, Dahl, Løken and Mogstad (2020), I show that the first-stage relationship between applications and manager rotations remains positive for all subgroups of workers defined by eight observable worker characteristics: age, tenure, educational qualifications, marital status, family status, German citizenship, team leadership, and past performance rating (Panel A of Appendix

[13] An alternative approach to illustrate random assignment is presented by the event study in Panel A of Figure 1.1 that shows the absence of pre-trends in worker applications before managers announce their rotation.

[14] Moreover, any estimated gender differences in hiring outcomes in Section 1.6.1 will be unaffected by biases that affect men and women equally.

Table A.9). To further test for the presence of defiers, I use two measures. First, I split my sample into workers who have never applied before and those who have applied before. Columns 1 and 2 of Appendix Table A.9, Panel B show that workers who applied in the past and who are more likely to be defiers do not experience lower or negative application effects when managers rotate. Second, I use the leave-out team mean of past application rates as a predictor for workers' unobserved application propensity. Columns 3 and 4 of Appendix Table A.9, Panel B show that even workers in teams with high application rates experience positive rotation effects.

1.5 Results

This section documents that manager rotations have large effects on worker applications. Several robustness tests support the causal interpretation of these rotation effects. I provide evidence that talent hoarding is a key mechanism underlying the observed impacts of manager rotations.

1.5.1 The Effect of Manager Rotations on Worker Applications

I begin by illustrating the dynamic effects of manager rotations using a quarterly event study around the quarter in which a manager rotates. I estimate a specification with worker and quarter fixed effects, binning event time dummy variables at $t = -8$ and $t = 4$, and clustering standard errors at the worker level. The sample of over 3,000 workers includes those who have not experienced a manager rotation (i.e. never-treated).

Panel A of Figure 1.1 presents quarterly event study coefficients and demonstrates that manager rotations result in an immediate and transitory increase in worker applications. Event time $t = 0$ denotes the quarter in which a manager rotates. Application rates increase up to three quarters before the manager rotation takes place, which is when managers start to inform their teams about their departure. However, before $t = -3$, when managers do not know yet about their job rotation, trends in applications are flat. In the quarter in which a manager rotation occurs, worker applications increase by 2.3 percentage points, almost doubling workers' baseline application rate of 2.7%. As the manager's replacement settles in, application rates return to baseline levels after one quarter.

Since manager rotation appears to have the largest impact on worker applications in the quarter of the rotation, the remainder of the analysis focuses on worker behavior in that quarter. Manager rotations almost double worker applications in the quarter of rotation. Column 1 of Table 1.2 presents OLS estimates for the effect of manager rotations on worker applications in the same quarter based on Equation 1.3. When a manager rotates, applications increase by 2.2 percentage points, representing a 76% increase compared to the baseline application rate of 2.9%.

The effect of manager rotations on worker applications is not a result of managers taking their subordinates with them to the new team or of workers replacing managers in their old position. Rather, 97% of applications in my sample are for positions outside of the worker's current team and not to the manager's new team. Even though manager rotations have large impacts on workers' job transitions within the firm, they have no impact on external transitions out of the firm. Panel B of Figure 1.1 illustrates this finding, indicating that the impacts of manager rotations are confined to worker interactions with the internal labor market. In addition, I find that the characteristics of the incoming manager do not affect the impacts of rotation, suggesting that the impacts are not driven by the incoming manager (Appendix Table A.4). Moreover, the fact that applications start to increase around $t = -3$, long before the firm typically begins to look for a replacement, further suggests that expectations about incoming managers do not play a central role.

1.5.2 Correlated Team-Level Shocks

I interpret manager rotations as reflecting the causal effect of a manager leaving her team. The primary threat to this interpretation is that there are team-level shocks that are correlated with manager rotations, such as unpleasant working conditions, bad news about the future outlook of the team, or the completion of a major milestone. These shocks may induce both managers and workers to apply away from the team. I present five pieces of evidence indicating that common examples of correlated team-level shocks are unlikely to drive the estimated rotation effects.

First, if the observed patterns result from correlated shocks, it should not matter whether a manager who has applied for a rotation actually rotates. Instead, one would expect that worker applications increase even if a manager's application for rotation is unsuccessful. Accordingly, I conduct a placebo test where I examine the effect of managers' unsuccessful applications for a job rotation on worker applications. Panel A of Figure 1.2 shows that while managers' successful job rotations double worker applications, managers' unsuccessful applications have no effect.

Second, in the presence of many common correlated shocks, the effects of a teammate's rotation should be larger than those of a manager's rotation, because teammates are generally more similar to workers than managers. Panel B of Figure 1.2 estimates the impacts of rotations of the most senior teammate and shows that these events have much smaller, rather than larger, impacts on worker applications.[15] This pattern is only consistent with a correlated shock that increases in magnitude with teammates' seniority, which is difficult to reconcile with the nature of many correlated shocks.

Third, in addition to applications, one would expect other team-level outcomes like absenteeism to react to many correlated team-level shocks, like a deterioration in working conditions. Appendix Figure A.3 estimates an event study around manager rotations and rejects economically significant changes in team-level absenteeism rates in the lead-up to a manager rotation. Fourth, many correlated shocks like bad news about the team's long-term future would be expected to have long-lasting impacts on applications, which is not consis-

[15]Similar patterns arise when other types of teammates rotate (e.g. most junior, randomly chosen).

tent with the short-lived nature of the observed increase in applications (Panel A of Figure 1.1).

Fifth, results presented in Section 1.5.3 show that the application effects of manager rotations are much larger for managers who have higher propensities to hoard talent. It is difficult to reconcile these patterns with the expected impacts of many correlated shocks, which would need to increase with managers' talent hoarding propensities to explain the observed patterns.

1.5.3 Interpretation of Manager Rotations As Evidence for Talent Hoarding

I provide evidence indicating that talent hoarding is a key mechanism underlying the observed impacts of manager rotations. Alternative mechanisms, such as loyalty, manager-worker-specific match effects or role-model effects, alone are not able to explain the observed rotation effects.

Short-Lived Impacts.—As soon as a manager learns that they will leave the team, their incentives to hoard talent abate. Worker applications should increase immediately upon learning of a manager's rotation, since the fear of retaliation should then cease. Because all managers face incentives to hoard talent, a new manager taking over a team should also exhibit talent hoarding behavior. Therefore, when measured in relatively high-frequency data, manager rotations should induce a sharp and transitory response in applications. These patterns are borne out in the data, as shown in Panel A of Figure 1.1. Applications increase once managers begin to announce their departure and return to baseline levels within one to two quarters after a new manager has taken over.

Worker Quality.—The conceptual framework in Section 1.2.2 predicts that managers are more likely to hoard high-productivity workers. Therefore, when managers rotate and talent hoarding temporarily subsides, application effects should be larger for the high-productivity workers who experienced more talent hoarding than their low-productivity peers. To test this prediction, I compare the impacts of rotations on applications by worker quality. I construct a measure of worker quality as the predicted values from a regression of applicants' internal hiring probabilities on applicant characteristics. The predicted value of this regression provides an index of worker quality for all workers, weighting worker characteristics by their importance for hiring prospects within the firm.[16] I compare workers in the top and bottom quartile of the index. Panel A of Figure 1.3 shows that high-quality workers experience a 3.4 percentage point increase in applications when a manager rotates, while applications among low-quality workers increase by only 0.9 percentage points, in line with the predicted patterns. Note that baseline application rates (not reported) are very similar across the different worker groups I compare in this section.

[16]This definition of worker quality has the advantage that it reflects the values that the firm places on worker qualifications. Robustness exercises in Appendix Section D.1 find similar patterns when using alternative measures such as educational qualifications or past performance.

Departure Costs.—A second prediction from the conceptual framework is that managers hoard talent more when the costs of worker departures are larger. I test this prediction using two measures of departure costs. The first measure is the number of other workers in the team since a larger team allows for more workers to compensate for a teammate's departure. Panel B of Figure 1.3 presents estimated effects of manager rotations for workers in the bottom and top quartiles of team size. By including detailed worker and position characteristics, this analysis compares workers with similar qualifications in similar positions, but in teams of different sizes. Applications increase by 4.3 percentage points among workers in small teams (one to three teammates) but only by 1.4 percentage points in large teams (10 or more teammates). Measuring departure costs using the average number of days required to fill a worker's vacated position yields similar results. Panel C of Figure 1.3 compares effects for high-cost workers (more than 174 days to fill their position) to low-cost workers (less than 135 days). Applications among workers who are hard to replace increase by 3.3 percentage points, compared to 1.3 percentage points for workers who are easy to replace.

Talent Hoarding Propensity.—The impacts of manager rotations are strongly correlated with direct measures of managers' propensities to hoard talent. Panel D of Figure 1.3 compares the impacts of rotations between managers with a high versus low propensity to hoard talent, using the measure of talent hoarding based on differences between actual and predicted potential ratings (described in detail in Section 1.3.3). Applications increase by 3.7 percentage points when a manager with a high propensity to hoard talent rotates, but only by 1.6 percentage points under managers with a low propensity to hoard talent. An alternative measure of talent hoarding based on nominations to succession lists yields similar results. Panel E of Figure 1.3 shows that rotations of managers with high propensities to hoard talent according to the measure based on succession lists increase applications by 3.2 percentage points, compared to only 1.4 percentage points under low-propensity managers. The conceptual framework also predicts that high-propensity managers have particularly large impacts on more productive workers. Figure 1.4 confirms this prediction, indicating that under both measures of talent hoarding, rotation effects are concentrated among high-quality workers.

Risk of Retaliation—In survey results (Section 1.2.1), workers report fear of manager retaliation as a reason they refrain from applying for internal positions, which represents a form of talent hoarding. If workers are more likely to refrain from applying to positions that carry a greater risk of manager retaliation, applications to these positions should increase the most when managers rotate. I measure the risk of retaliation in two ways. First, applications that are less likely to be successful necessarily carry a greater risk of manager retaliation. I test this prediction by analyzing whether manager rotations have larger effects for applications that typically yield lower success likelihoods. Table 1.2 confirms this prediction and documents that baseline applications for lateral transitions increase by 61%, while small and major promotions increase by 98% and 123%, respectively. Second, managers should be more likely to learn about worker applications to positions that are closer in proximity. I test this prediction by comparing workers' applications to job openings within versus across their current division, functional area, and location. Appendix Table A.6 shows that manager rotations have much larger effects on applications in close proximity to workers' current position with respect to all three dimensions, confirming the prediction.[17]

External Transitions.—While managers may frequently learn about workers' unsuccessful applications to internal positions and possibly retaliate or interfere with those applications, this is not the case for applications to positions outside of the firm. Therefore, manager rotations should only impact internal applications and not external applications outside of the firm. Because I do not observe external applications, I test this prediction by comparing the effect of manager rotations on internal job transitions within the firm and external job transitions out of the firm. Panel B of Figure 1.1 shows that manager rotations only increase worker transitions within the firm and not outside of the firm, even though both types of transitions trend identically in quarters prior to the rotation. Columns 1 and 2 of Appendix Table A.5 confirm this finding quantitatively. Manager rotations increase internal transitions by 1.1 percentage points but have negligible effects on external transitions (0.09 percentage points) compared to the identical baseline rate of 0.7%.

Alternative Mechanism: Loyalty or Match Effects.—It is possible that manager rotations affect worker applications through alternative channels. For instance, workers may refrain from applying for a new position because of loyalty towards their manager or because manager-worker-specific match effects make their current position particularly appealing. While loyalty and match effects are undoubtedly important parts of interpersonal relations in the workplace, I do not find evidence that these mechanisms drive the particular increase in applications around manager rotations.

Loyalty and match effects are typically assumed to compound over time, suggesting that the impacts of rotations increase with exposure time. However, Appendix Figure A.4 finds no evidence that rotation effects vary by the length of time a worker was exposed to the rotating manager. Rotations increase applications even for workers who have been exposed to their manager for one quarter or less. Moreover, one would expect rotations to lead to long-lasting increases in applications since it should take time for workers to become loyal to their new manager or for manager-worker-specific match effects to develop, but the increase in applications is quite transitory.

Under loyalty or match effects, manager rotations increase applications through a decrease in the value of workers' default option and should thus make any type of transition more appealing. Thus, when managers rotate, there should be at least some increase in external job transitions given that internal and external transitions trend identically prior to the rotation. However, Panel B of Figure 1.1 shows that no increase in external transitions occurs.

[17]Manager rotations could lead to an opposite prediction on proximity if they were to operate through the visibility channel instead of retaliation. However, while the suppression of visibility is a common form of talent hoarding, manager rotations do not immediately increase visibility since worker evaluations only occur one to two times a year.

Moreover, under manager-worker-specific match effects, rotations should be more impactful for workers who were hired by that manager, relative to workers who were already on the team when that manager arrived. This is because the manager would have had the ability to select new workers based on match quality. Instead, Columns 3 and 4 of Appendix Table A.5 document that manager rotations have very similar effects on both groups.

Neither mechanism can simultaneously produce the findings that rotations have larger effects among workers who (i) experience a suppression of worker visibility, (ii) are more qualified, (iii) work in small teams, or (iv) work in positions that are hard to replace. In addition, neither mechanism can explain why rotations disproportionally deter applications for positions that are more selective (Table 1.2) or closer in proximity (Appendix Table A.6). Together, these findings indicate that loyalty and match effects alone are unlikely to drive the observed effects of manager rotations.

Alternative Mechanism: Salience and Role-model Effects.—Another channel through which manager rotations could affect applications is through salience or role-model effects, which are important in many settings. For instance, career-driven managers who pursue rotations likely generate information flows that make career planning more salient to workers. Since I find no increase in applications around a manager's unsuccessful application for a rotation, role-model effects in this setting must be limited to only successful rotations in order to drive the observed effects.

Successful manager rotations could increase applications by making career planning more salient or exposing workers to more information about career opportunities. However, the short-lived nature of the observed application effects is difficult to reconcile with an information-based mechanism, because the transfer of information should produce longer-lasting effects.

Role-model effects are often found to be particularly impactful if role models are similar in attributes to affected individuals (e.g. in terms of their gender). However, while only 12% of rotating managers are female, rotation effects are much larger for female workers than for males. In addition, if observing others navigate their career is a key underlying factor of rotation effects, such a role-model effect should not be limited to managers. Since coworkers are more similar to workers than their managers, one would expect even larger effects for observing coworkers rotate. However, Panel B of Figure 1.2 shows that a manager rotation causes four times larger application rates in the same quarter than a coworker rotation, further suggesting that role-model effects alone are unlikely to explain the observed impacts of manager rotations.

External Validity.—The preceding empirical exercises support the interpretation that talent hoarding is behind the observed impacts of manager rotations at this firm. While focusing on a single firm yields large advantages in terms of comprehensive data coverage and is the standard approach in the literature (Baker et al., 1994, Lazear et al., 2015, Cullen and Perez-Truglia, 2019, Hoffman and Tadelis, 2021, Benson et al., 2021), doing so naturally raises questions about external validity. Three features of my environment suggest that the patterns documented here are likely present both in other countries and in other firms. First, because the firm operates in many different countries, the firm's internal personnel records from other countries indicate that talent hoarding is not restricted to the German context. In unreported results, I construct the direct measure of talent hoarding using employee data from the firm's locations in other countries, such as the United States. The observed degree of talent hoarding among employees outside of Germany closely resembles that observed within Germany. Second, the firm is similar to other large firms in Germany in terms of its workforce composition (Appendix Table A.1), as well as its organizational design. Since many other firms in Germany task managers with identifying talented workers and impose little oversight, these firms create the same conditions that give rise to talent hoarding at the firm that I study (hkp, 2021). Third, companies across the world report that talent hoarding is commonplace, creates barriers to talent allocation, and occurs through many of the same managerial behaviors that are documented in this study (i4cp, 2016, KornFerry, 2015, Matuson, 2015, Sullivan, 2017).

1.6 Talent Hoarding, Misallocation, and Gender Inequality

This section demonstrates that talent hoarding has important efficiency costs in the form of misallocation of talent. I find that talent hoarding reduces the quality and subsequent performance of promoted workers. Moreover, misallocation effects are larger for women, implying that talent hoarding exacerbates gender inequality with respect to representation and pay.

To test for misallocation, I evaluate whether qualified workers are deterred from moving to higher-level positions in which they could be more productive. This notion of misallocation is in line with the literature on optimal hierarchies (Rosen, 1982), which holds that firms must promote high-ability workers to high-level positions to efficiently allocate talent. I analyze transitions to higher-level positions by focusing on major promotions, described in Section 1.3.4. These promotions represent critical decisions for the firm's allocation of talent because they reflect meaningful changes in job responsibility, such as transitions from individual contributors to managers.

I begin by analyzing the extent to which talent hoarding deters applications for major promotions. Column 1 of Table 1.3 shows that rotations increase applications for major promotions by 0.65 percentage points, corresponding to a 123% increase. This finding suggests that talent hoarding deters a large group of workers from applying for major promotions, shrinking applicant pools and potentially limiting the ability of the firm to fill high-level positions with high-ability workers.

While these findings are suggestive of the negative impacts of talent hoarding on the efficient allocation of talent, impacts may be modest if marginal applicants are unlikely to be successful. To evaluate whether this is the case, I use manager rotation as an instrument for worker applications to estimate the success probability of deterred (marginal) applications.

Column 2 of Table 1.3 reports 2SLS estimates of Equation 1.5. I find that marginal applications have a 15.11% hiring probability, which is similarly high as the baseline success rate of 17.00%. This finding implies that a substantial share of marginal applicants forgoes high-stakes applications that would have been successful, indicating that talent hoarding creates misallocation. Effects are not limited to major promotions as documented by robustness exercises in Appendix Section D.1, which examine other transition types (e.g. small and very large promotions).

Going a step further, I leverage data on employees' performance ratings to show that deterred applicants would have been more productive at higher-level positions. Performance ratings are designed to provide task-specific feedback on whether a worker has accomplished her tasks in the past evaluation cycle. Since most workers in a team perform very similar tasks, performance ratings are particularly well-suited for drawing comparisons across workers within teams. I assess whether talent hoarding causes forgone performance at higher levels by estimating a 2SLS regression in which the outcome is defined as a worker landing a major promotion *and* performing better than the leave-out average performance of the new team one year later. Column 3 of Table 1.3 presents the results of this analysis and shows that 8.40% of marginal applicants land a promotion and outperform their teammates, strongly suggesting that the firm would have foregone higher performance at higher-level jobs had the marginal workers not been promoted.

I provide more detail on the characteristics of deterred applicants using a complier analysis based on Abadie (2003).[18] Table 1.4 compares average characteristics across the entire employee population (Column 1), always takers who apply even in absence of manager rotations (Column 2), and marginal applicants who only apply if a manager rotates (Column 3), and shows that marginal applicants are positively selected. For instance, while 48.6% of always takers hold a graduate degree, this is true for 63.3% of marginal applicants. Similarly, 56.9% of always takers received high performance ratings prior to applying, compared to 65.2% of marginal applicants, and 5.7% of marginal applicants (but only 2.4% of always takers) have been nominated to a succession list at the firm.

Taken together, these results are consistent with the predictions of the conceptual framework in Section 1.2.2: talent hoarding reduces the share of high-quality workers in the applicant pool, limiting the firm's ability to efficiently fill high-level positions. Not only are marginal applicants well-qualified and likely to be successful in their applications, they would also perform well in high-level positions.

[18] Under standard IV assumptions discussed in Section 1.4, complier characteristics can be estimated as $E[X_{it}|\text{Compliers}]$ for some characteristic X_{it}. I calculate average complier characteristics and standard errors by performing 2SLS using the first-stage Equation 1.3 and a reduced-form equation replacing the outcome variable in Equation 1.4 with $X_{it}A_{it}$, where X_{it} corresponds to a characteristic of individual i and A_{it} is a binary indicator for i applying in quarter t. I compute characteristics for always takers, who apply even in the presence of talent hoarding, by estimating an OLS regression of $X_{it}A_{it}(1 - Z_{it})$ on $A_{it}(1 - Z_{it})$, which allows me to estimate $E[X_{it}|\text{Always takers}]$.

1.6.1 Gender Differences in Misallocation

Survey responses discussed in Section 1.2.1 indicate that managers hoard talent through direct interpersonal interactions. Such behavior raises the possibility that talent hoarding may be particularly impactful for workers who depend more heavily on managerial support or are more sensitive to confrontation with their manager. Motivated by previous work on gender differences in preferences (Bertrand, 2011), I test whether talent hoarding has differential effects by gender.

I first analyze the effects of manager rotations on worker applications separately for men and women. Column 1 of Table 1.5 shows that men's applications increase by 0.55 percentage points (a 98% increase), while manager rotations increase women's application rate by 1.05 percentage points, a 244% increase (Column 2). These findings reveal that talent hoarding deters a larger share of female applicants from applying for major promotions compared to males.

To evaluate whether talent hoarding is more detrimental for women than for men, I compare hiring probabilities for men and women at the margin, which can be interpreted as a Becker outcome test. Columns 3 and 4 of Table 1.5 report the 2SLS estimates of Equation 1.5, which are separately estimated by gender and capture the probability of landing a major promotion for marginal applicants. Both men and women experience positive and statistically significant marginal hiring probabilities of 12.78% and 25.78%, respectively. However, the hiring probability for women is twice that for men, implying that talent hoarding is more consequential for women's career progression.

Differences in promotion likelihoods may also reflect differential labor demand, such as affirmative action policies. To verify that this is not the case, I test whether women would also be more likely to perform well in high-level positions. Columns 5 and 6 of Table 1.5 present the 2SLS estimate for landing a major promotion and performing better than the average in the new team one year later. Women exhibit a marginal probability of outperforming their teammates of 17.66%, which is almost three times higher than men's marginal probability of 6.38%. This finding suggests that marginal female applicants would not only be more likely to land promotions, they would also be more likely to perform well in these positions.

These findings suggest that talent hoarding may deter higher-quality women compared to men. I again conduct a complier analysis, now separately by gender. Since men and women work in very different positions at the firm, I adjust population characteristics using the same set of baseline controls used in computing complier and always taker characteristics. Table 1.6 documents that marginal female applicants who are deterred from applying by talent hoarding are strongly positively selected compared to always takers and average workers. While 74.7% of marginal female applicants hold a graduate degree, this is only true for 39.0% of always takers. Similarly, 73.6% (73.4%) of women at the margin received high performance (potential) ratings in the past relative to 53.8% (43.1%) of always takers. 8.7% of female marginal applicants have been nominated to succession lists (i.e. suitable successors for high-level positions), while the firm only nominated 2.9% of always takers. For men, the extent of this positive selection is substantially less pronounced, suggesting that talent hoarding affects women at a higher part of the quality distribution compared to men. Together, these results indicate that talent hoarding has more severe misallocation effects for women.

1.6.2 Impacts on Gender Inequality in Pay and Representation

The preceding results documenting disparate impacts of talent hoarding by gender suggest that talent hoarding may exacerbate gender inequality at the firm. To quantify the differential effect of talent hoarding, I make use of the potential outcomes framework that follows from the interpretation of β_{IV} as the LATE for marginal applicants. I compare the average potential outcomes for compliers in the treated state (i.e. when marginal applicants apply because talent hoarding ceases) and the untreated state (i.e. when marginal applicants do not apply due to talent hoarding).

Within the potential outcomes framework, I compare the gender gap in log real annual earnings and hierarchy levels one year later between the untreated state and the treated state, allowing me to evaluate the effect of applications on gender disparities. The advantage of this framework is that it allows a comparison of the same set of individuals across two different potential outcomes, avoiding potential composition bias. Limiting attention to marginal applicants is not restrictive, since they represent the group of workers whose outcomes differ because of talent hoarding. To test for gender differences in these worker outcomes, I follow the literature assessing gender pay gaps (Blau and Kahn, 2017) and estimate worker outcomes separately by gender for marginal applicants in each potential outcomes state, using the same set of controls X_{it} as in previous models.

Panel A of Figure 1.5 presents estimates of hierarchy levels in quarter $t + 4$ for marginal applicants in both potential outcomes states by gender. Outcomes are reported in terms of percentiles at the firm. Both men and women experience higher hierarchy levels if they choose to apply; however, the larger gains realized by women lead to a reduction in the representation gap by 91%. This reduction in gender differences with respect to hierarchy levels translates into a substantial reduction in the earnings gap. Panel B presents log annual real earnings across treatment states in percentiles. Applying substantially increases worker earnings four quarters later and appears to reduce gender disparities in pay by 77%. This finding suggests that talent hoarding exacerbates gender inequality with respect to pay and representation in the firm, highlighting the negative consequences of talent hoarding with respect to both efficacy and equity in the internal labor market.

1.7 Unpacking Talent Hoarding: Suggestive Evidence

The documented costs of talent hoarding, particularly for women, raise the question of which factors underlie these impacts. Do talent hoarding effects depend on manager and worker characteristics? For instance, do managers treat women differently, or are there differences in how women and men react to talent hoarding? This section investigates the role of manager and worker characteristics and what this implies for how organizations may react to talent hoarding.

Previous research has highlighted the importance of manager gender as a key correlate of manager behavior, particularly when trying to explain gender differences in worker outcomes (e.g. Kunze and Miller, 2017, Cullen and Perez-Truglia, 2019). Accordingly, I begin by comparing the first-stage effects of manager rotations on worker applications by manager gender. Columns 1 and 2 of Table 1.7, Panel A show that there is no statistically detectable difference between the impacts of rotations by male and by female managers. Similarly, Columns 3 and 4 show that there are no substantial differences by whether managers and workers have opposite genders. In unreported results, I find that this pattern persists when conducting estimation separately for male versus female workers.

Talent hoarding behavior may also differ due to other manager characteristics. Table 1.7 examines heterogeneity in rotation effects by key manager characteristics, such as age (Columns 5 and 6 of Panel A), experience (Columns 1 and 2 of Panel B), and managerial quality (Columns 3 to 6 of Panel B). Besides managers' own performance ratings which they receive from their direct supervisor, I measure manager quality using the leave-out mean absenteeism rate that the manager's team had in the past. Comparing the bottom and top quartiles of these measures, I do not find differential effects by manager characteristics. In unreported results, I also find that leave-out mean team turnover rates in the past and other manager evaluations by their supervisor that complement performance ratings (e.g. rating of problem-solving ability) do not predict the magnitude of rotation effects.

I find similar patterns using suppression of worker visibility as a direct measure of managers' talent hoarding propensities. An F-test of all manager characteristics included in the logit regressions presented in Columns 1 and 2 of Table 1.8—which besides manager gender include age, marital and family status, tenure at the firm, division, function, and location—rejects their joint significance in explaining the propensity to hoard talent. While the survey results support the interpretation of manager differences in talent hoarding as differences in self-interest or altruism, future research that elicits managers' personality traits that are not contained in personnel records may help to provide additional information on the reasons why some managers hoard more than others.

The finding that talent hoarding behavior is not correlated with manager or team-level characteristics that are easily observable has important implications for how to detect talent hoarding. My analysis suggests that it is not possible to accurately predict a manager's propensity to hoard talent using observable characteristics. Without a rigorous data collection and analysis effort it is difficult for firms to pin down which managers are hoarding talent because forgone promotions can be attributed to a range of different factors – including managers hoarding talent, workers not wanting to pursue promotions, or hiring managers choosing other candidates.[19]

[19] In line with infrastructure constraints, the firm in this study is not able to assemble and analyze a dataset that would allow them to detect talent hoarding. Consequently, the findings in this research study represent the first empirical test of talent hoarding at the firm that is difficult for practitioners to implement.

Workers seem not to be able to find out whether a manager that they have not worked with before is a talent hoarder. When assessing the number of workers who apply for an internal job opening, the talent hoarding propensity of the team's manager carries no predictive power for worker application decisions. This finding is in line with the general pattern that internal applicants only in rare cases have worked with anyone in the team before applying to a job opening, reducing worker ability to gather information about a manager's talent hoarding propensity. These results echo findings from the literature on asymmetric information in firms (Kahn and Lange, 2014) and help to rationalize why talent hoarding may persist, even if firms know that misaligned incentives exist.

In addition to manager characteristics, the impacts of talent hoarding likely differ by workers. Since talent hoarding has very different effects by gender, a first-order question is whether managers hoard women more or whether women react to talent hoarding more. Because the rotation effects on applications depend both on managers and workers, they cannot distinguish worker effects from manager effects. To isolate manager behavior, I use the two measures of managerial talent hoarding described in Section 1.3.3. Columns 3 and 4 of Table 1.8 document that worker gender does not significantly affect managers' talent hoarding behavior. I also find no gender difference in workers' likelihood of reporting fearing manager retaliation, which is a key dimension of talent hoarding (Figure 1.6). These findings suggest that managers do not hoard women more than men.

Instead, employees' survey responses suggest that men and women react differently to talent hoarding. In the survey, women are 29% more likely than similar men to mention the importance of manager support for their career development (Figure 1.6), suggesting that they rely more on managers' career guidance. In addition, women are 19% more likely to rank a good relationship with their supervisor as the most important feature of their job (Figure 1.6), indicating that women seem to place more value on preserving a good relationship with their manager. These findings are in line with the fact that talent hoarding typically occurs through direct interpersonal interactions and are consistent with a large body of research on gender differences in preferences (Bertrand, 2011). Taken together, my results suggest that despite being gender-neutral, talent hoarding produces disparate effects due to workers' underlying gender differences. Mitigating talent hoarding may thus allow organizations to reduce gender inequality.

1.8 Conclusion

This paper provides the first empirical evidence that talent hoarding is an important source of frictions in organizations. Using novel personnel records and internal application data from a large manufacturing firm, I show that talent hoarding leads to misallocation of talent and perpetuates gender inequality at the firm. While my results provide the first detailed insights on talent hoarding, additional evidence suggests that such talent hoarding behavior is endemic. In a survey of the top publicly listed companies in Germany, 83% cite talent hoarding as a crucial friction in their organization (hkp, 2021). Firm surveys in other countries, such as the United States, document that talent hoarding is not limited to German organizations (i4cp, 2016, KornFerry, 2015).

Because talent hoarding arises due to misaligned incentives, a natural solution would be to more closely align the incentives of managers with those of the firm. Surveys of German firms suggest that accomplishing this realignment through financial incentives is infeasible (hkp, 2021). However, policies that increase application rates — such as implementing regular application schedules and having other organizational agents, such as the HR department, directly invite workers to apply for positions — could reduce the scope for managers to engage in talent hoarding. The employee survey I conduct shows that for such policies to be effective, the firm must be able to deter managers from retaliating against workers, for instance by assuring full confidentiality for applicants.

While the costs of these policies are likely to be non-negligible, their potential benefits are substantial given the potential gains for firms. A key contribution of this analysis has been to provide the first empirical evidence showing that talent hoarding has meaningful efficiency costs in the form of talent misallocation. It is likely that organizations suffer additional efficiency costs through other channels. The employee survey suggests that talent hoarding leads managers to deter workers from pursuing training programs and from participating in high-profile projects in order to suppress worker visibility. These actions likely cause substantial underinvestment in human capital. In addition, workers who report being subject to talent hoarding are 30% more likely to report having searched for external jobs, indicating that talent hoarding may create unwanted turnover of high-quality workers the firm would like to retain. These findings suggest that the estimates in this study represent a lower bound for the efficiency costs of talent hoarding, which likely exacerbate gender inequality in the labor market.

1.9 Figures and Tables

Figure 1.1: Effect of Manager Rotations on Applications and Job Transitions

Panel A. Application Probability

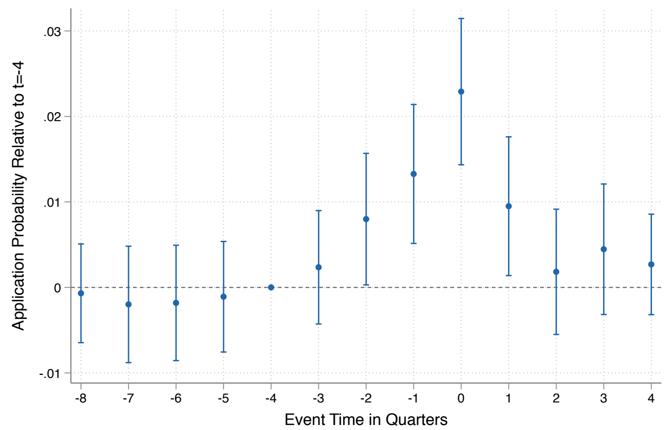

Panel B. Cumulative Transition Probability

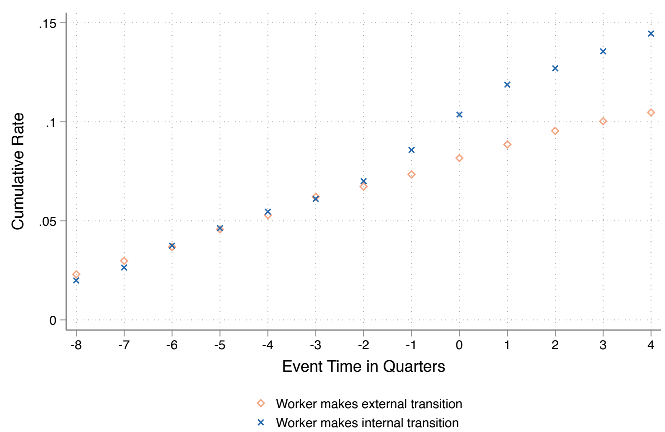

Notes: This figure depicts internal applications (Panel A) and job transitions (Panel B) around a manager rotation. Panel A presents estimates from an event study regression, in which the outcome is an indicator that the worker applied in a quarter and event time is defined relative to the occurrence of a manager rotation. The specification includes worker and quarter fixed effects. I bin event time dummy variables at $t = -8$ and $t = 4$ and cluster standard errors at the worker level. The mean application rate as of $t = -4$ is 0.027. The sample of 3,xxx workers includes those who have not experienced a manager rotation. Panel B plots the cumulative share of workers who have exited the team via internal (i.e. within the firm) and external (i.e. out of the firm) transitions around a manager rotation. Workers are assigned to their team as of ten quarters before the team experiences a manager rotation.

Figure 1.2: Manager Rotation Placebo Test

Panel A. Results by Outcome of Manager Application

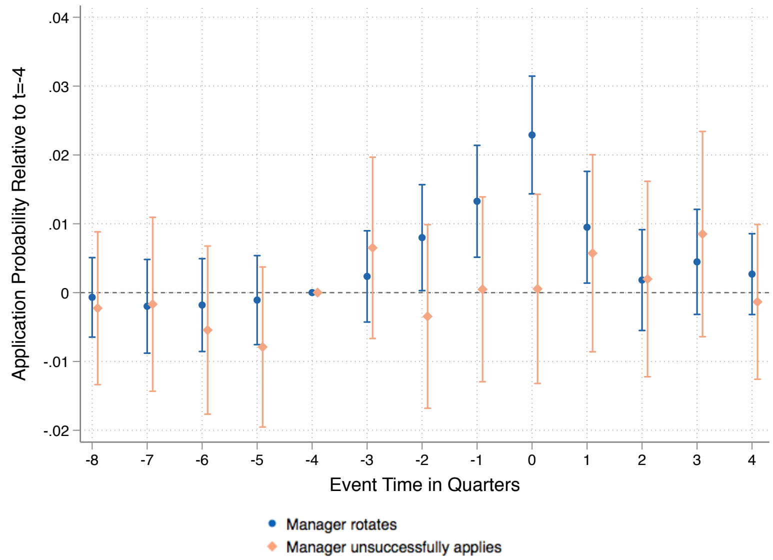

Panel B. Results by Type of Rotating Teammate

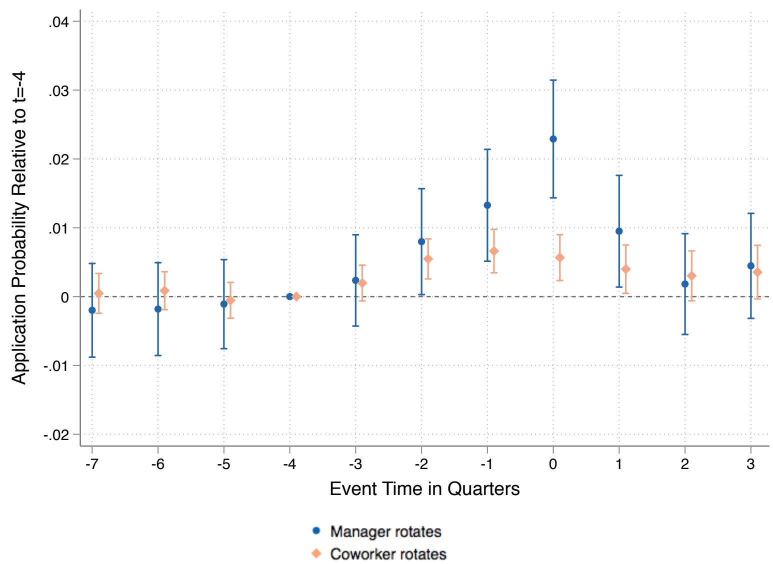

Notes: This figure presents placebo tests for manager rotations. Estimates stem from event study regressions, in which the outcome is an indicator that the worker applied in a quarter and event time is defined relative to the occurrence of a rotation event. The specification includes worker and quarter fixed effects. I bin event time dummy variables at $t = -8$ and $t = 4$ and cluster standard errors at the worker level. The mean application rate as of $t = -4$ is 0.027. The sample includes those who have not experienced a manager rotation. Panel A compares a successful manager rotation (in blue, N=3, xxx) to a placebo event, in which a manager applied for an internal job rotation, but did not land the position and stayed in the team (in orange, N= 1,xxx). Panel B compares a manager rotation (in blue, N=3, xxx) to the rotation of the most senior coworker (in orange, N= 2x,xxx).

Figure 1.3: Heterogeneity in Application Effects by Predicted Level of Talent Hoarding

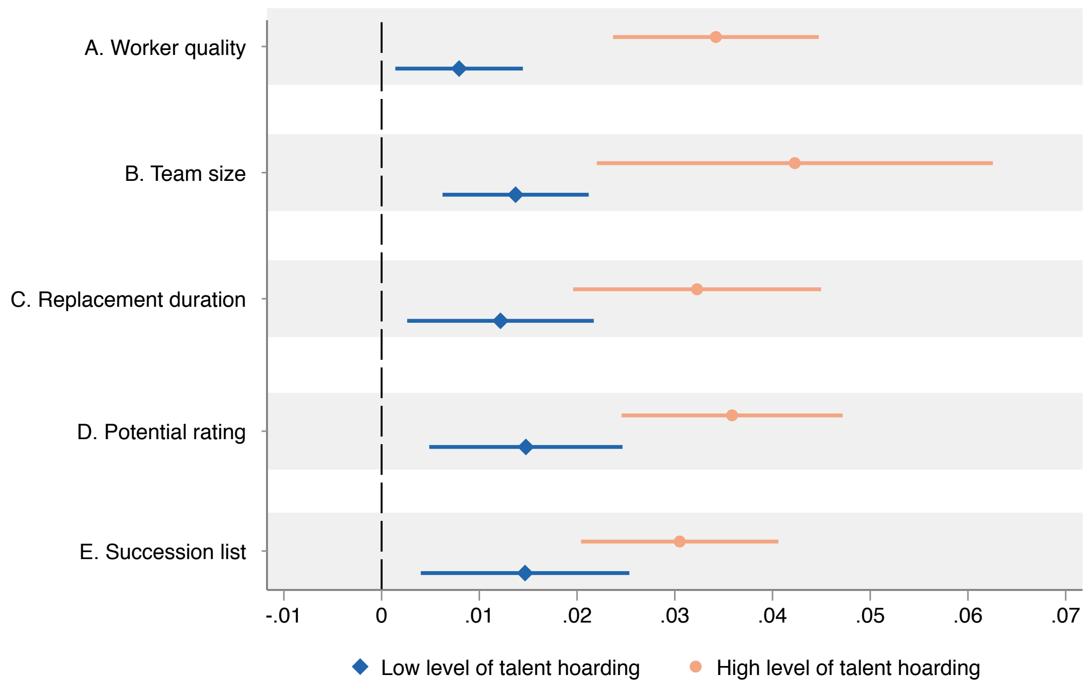

Notes: This figure demonstrates larger impacts of manager rotations on worker applications for subgroups that are expected to experience *high* instead of *low* predicted levels of talent hoarding. Each coefficient stems from a separate regression based on Equation 1.3 using robust standard errors. Panels A, B, and C focus on workers in the bottom and top quartile of the respective measure to distinguish between *high* and *low* levels of hoarding. Panel A uses a quality index, constructed using the predicted value from an OLS regression of workers' internal hiring probability on worker characteristics. Panel B uses team size (*high*: <4 teammates, *low*: >9 teammates). Panel C uses the average number of days it takes to replace a position (*high*: >174 days, *low*: <135 days). Panels D and E compare rotations of manager types with *high* versus *low* propensity to hoard based on measures of worker visibility. Panel D uses managers' mean deviations between actual and predicted potential ratings. Panel E uses managers' mean deviations in subordinates' probability to be nominated to succession lists. Baseline application rates are very similar across subgroups and are not separately reported. Controls: Female, age, German citizenship, educational qualifications, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects. N=3xx,xxx.

Figure 1.4: Heterogeneity in Application Effects by Hoarding Propensity and Worker Quality

Panel A. Potential Ratings as Measure of Talent Hoarding Propensity

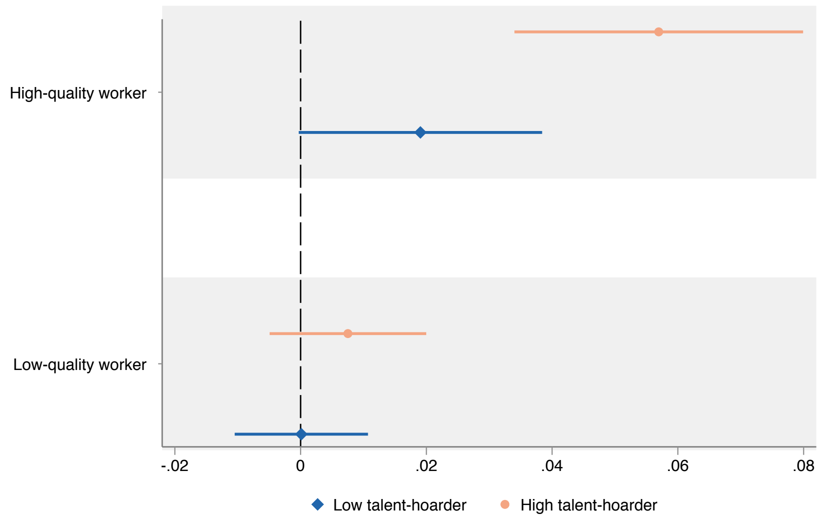

Panel B. Succession Lists as Measure of Talent Hoarding Propensity

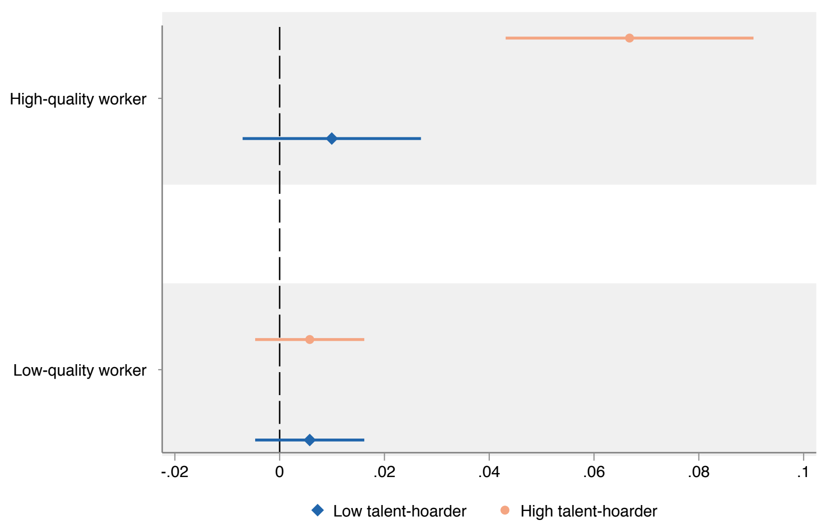

Notes: This figure demonstrates larger impacts of working under a manager with high talent hoarding propensity for high-quality workers than low-quality workers. Each coefficient stems from a separate regression based on Equation 1.3 using robust standard errors. High-quality (low-quality) workers represent workers in the top (bottom) quartile of a quality index I construct. Panel A uses the mean deviation between actual and predicted potential ratings to compare rotations of manager types with *high* versus *low* propensity to hoard. Panel B uses mean deviations in subordinates' probability to be nominated to succession lists. Controls: Female, age, German citizenship, educational qualifications, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects. N=3xx,xxx.

Figure 1.5: Potential Outcomes by Application Status for Marginal Applicants

Panel A. Hierarchy Level at $t+4$ Panel B. Log Earnings at $t+4$

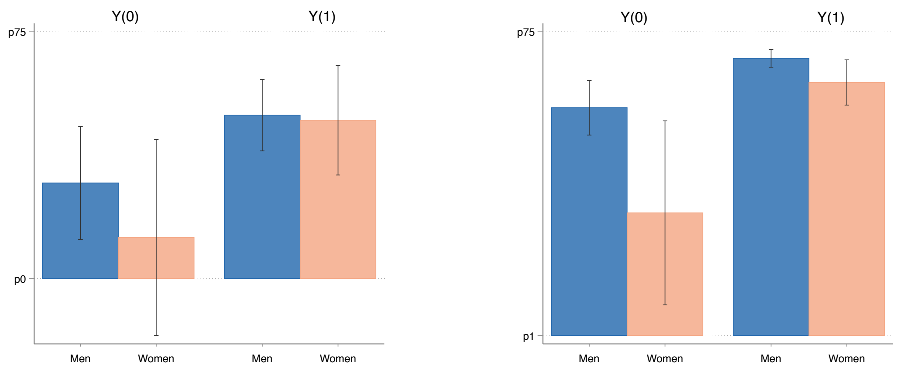

Notes: This figure depicts potential outcomes in quarter $t+4$ measured in percentiles. The left two bars in each panel represent the potential outcomes for marginal applicants had they not applied, labeled $Y(0)$. The right two bars represent the potential outcome for marginal applicants had they applied, labeled $Y(1)$. Panel A presents workers' hierarchy level in quarter $t+4$. Panel B presents workers' log real annual earnings in quarter $t+4$. Controls: Female, age, German citizenship, educational qualifications, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects. Standard errors are robust. N=3xx,xxx.

Figure 1.6: Gender Differences in Survey Responses

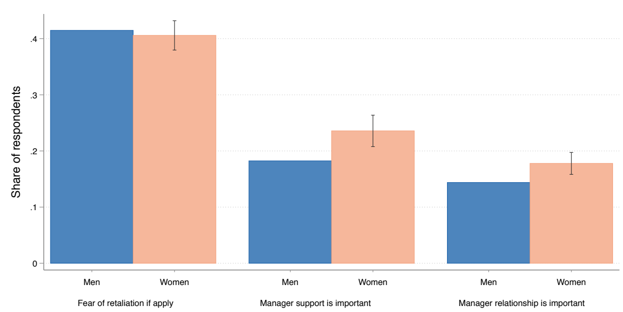

Notes: This figure depicts survey responses separately estimated by gender using robust standard errors. Bars 1 and 2 represent the share of respondents who report fearing retaliation if managers find out about internal applications. Bars 3 and 4 represent the share of respondents who report manager support as critical for their career development. Bars 5 and 6 report the share of respondents who rank a good relationship with their manager as most important job feature. While men and women are similarly likely to report fearing retaliation, women are 29% more likely to value manager support and 19% more likely to value a good relationship with their manager than similar men. Controls: Age, tenure, schooling, nationality, children, functional area, location, full-time, hours, and team leadership. N=1x,xxx.

Table 1.1: Summary Statistics of Analysis Sample

	Mean	Std. deviation	p25	p75
Demographics				
Female	0.21	0.40	0.00	0.00
German citizen	0.89	0.31	1.00	1.00
Age (years)	43.41	10.03	35.00	51.50
Tenure at firm (years)	13.34	9.65	5.00	19.25
Schooling (years)	15.81	2.74	12.00	18.00
Married	0.62	0.49	0.00	1.00
Children	0.75	0.43	0.00	1.00
On parental leave	0.03	0.17	0.00	0.00
Position Characteristics				
Technical position	0.63	0.48	0.00	1.00
Full-time	0.92	0.27	1.00	1.00
Weekly hours	41.15	4.56	40.75	43.50
Team leadership	0.19	0.39	0.00	0.00
Number of direct reports	5.00	3.89	2.00	7.00
Career Progression				
High performance rating	0.54	0.50	0.00	1.0
High potential rating	0.27	0.44	0.00	1.00
Time in position (quarters)	13.34	9.74	5.00	21.00
Internal application	0.03	0.17	0.00	0.00
Internal job transition	0.01	0.09	0.00	0.00
Observations		3xx ,xxx		

Notes: This table reports summary statistics for the quarterly analysis sample. This sample consists of over 300,000 employee-by-quarter observations from 2015 to 2018. A technical position is defined as a job related to engineering, IT, quality management, or production. The number of direct reports is only calculated for employees with team leadership. A high performance rating is defined as *sometimes exceeds expectations* or *often exceeds expectations*. High potential rating refers to supervisors' assessment that workers have future potential for higher-level positions. Internal application and job transition rates are at the quarterly level.

Table 1.2: Application Effects of Manager Rotations by Position Selectivity

	Application for			
	Any position	Lateral transition	Small promotion	Major promotion
	(1)	(2)	(3)	(4)
Manager Rotation	0.0224	0.0030	0.0115	0.0065
	(0.003)	(0.001)	(0.002)	(0.001)
Outcome Mean	0.0290	0.0049	0.0118	0.0053
Size of Effect in %	76	61	98	123
Observations	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx

Notes: This table illustrates the effect of manager rotations on worker applications in the same quarter. Column 1 shows the effect of manager rotations for any positions, Column 2 represents lateral transitions. Column 3 focuses on small promotions, which are transitions defined by a cutoff of 10 with respect to the increase in hierarchy index. Column 4 focuses on applications for major promotions, which are defined as an increase in the hierarchy index of 20 or more and represent large career jumps, such as transitions from individual contributors to team leader positions. Controls: Female, age, German citizenship, educational qualification, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects. Robust standard errors in parentheses.

Table 1.3: Misallocation Effects of Talent Hoarding

	OLS	IV	IV
	Applied for major promotion	Hired for major promotion	Perform>team average if land major promotion
	(1)	(2)	(3)
Manager Rotation	0.0065	-	-
	(0.001)	-	-
Applied	-	0.1511	0.0840
	-	(0.034)	(0.026)
Outcome Mean	0.0053	0.0009	0.3673
Observations	3xx,xxx	3xx,xxx	3xx,xxx

Notes: This table reports the effects of manager rotations on workers' career progression. Each coefficient is based on a separate regression. Column 1 reports the first-stage effect of manager rotation on applications for major promotions based on Equation 1.3. Column 2 reports the estimate from a two-stages least squares regression on landing a major promotion that instruments for applying with manager rotation based on Equation 1.5 which represents the LATE. Column 3 estimates a similar two-stages least squares regression, but uses an indicator for landing a major promotion *and* performing better than the leave-out team average one year later as outcome variable. Controls: Female, age, German citizenship, educational qualification, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects. Robust standard errors in parentheses.

Table 1.4: Characteristics of Marginal Applicants (in %)

	All workers (1)	Always takers (2)	Marginal applicants (3)
German citizen	89.8	87.2	86.4
Age >=40yrs	60.4	39.2	50.6
Married	61.7	54.8	49.5
Children	73.3	68.6	63.6
Tenure at firm <2yrs	37.5	53.4	49.2
Tenure at firm 2-5yrs	40.5	38.0	38.5
Tenure at firm >=5yrs	21.9	8.6	12.3
Graduate degree	47.6	48.6	63.3
Full-time	92.5	94.4	97.1
High performance	54.0	56.9	65.2
High potential	28.2	44.0	43.0
Technical position	63.2	56.7	59.1
Low-level position	68.9	73.6	77.7
First-level leadership position	11.5	9.7	7.4
Time in position <2yrs	37.1	38.3	39.3
Time in position 2-5yrs	36.2	40.9	42.1
Time in position >=5yrs	26.7	20.8	18.5
Nominated to succession list	1.6	2.4	5.7
Applied 12 months before	2.6	11.2	2.6

Notes: This table illustrates results from a complier analysis as described in Section 1.6. Each number is based on a separate regression including controls and represents an adjusted mean (in %). Column 1 shows means for all workers, Column 2 represents always takers, and Column 3 represents marginal applicants, who only apply if managers rotate and talent hoarding temporarily abates. Each number represents the share of workers in a given group that exhibit the respective characteristic. A technical position is defined as a job related to engineering, IT, quality management, or production. Low-level positions are defined as positions at low hierarchy levels without leadership responsibility (i.e. individual contributors). First-level leadership represents positions with limited leadership responsibility, such as team leaders. Controls: Female, age, German citizenship, educational qualification, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects.

Table 1.5: Misallocation Effects of Talent Hoarding by Gender

	OLS Applied for major promotion		IV Hired for major promotion		IV Perform>team average if land major promotion	
	Men	Women	Men	Women	Men	Women
	(1)	(2)	(3)	(4)	(5)	(6)
Manager Rotation	0.0055	0.0105	-	-	-	-
	(0.001)	(0.003)	-	-	-	-
Applied	-	-	0.1278	0.2578	0.0638	0.1766
	-	-	(0.036)	(0.106)	(0.026)	(0.086)
Outcome Mean	0.0056	0.0041	0.0009	0.0009	0.3817	0.3572
Observations	3xx,xxx	8x,xxx	3xx,xxx	8x,xxx	3xx,xxx	8x,xxx

Notes: This table reports the effects of manager rotations on workers' career progression by gender. Each coefficient is based on a separate regression. Columns 1 and 2 report first-stage effects of manager rotation on applications for major promotions based on Equation 1.3. Columns 3 and 4 report estimates from a two-stages least squares regression on landing a major promotion that instruments for applying with manager rotation based on Equation 1.5. Columns 5 and 6 estimate a similar two-stages least squares regression, but use an indicator for landing a major promotion *and* performing better than the leave-out team average one year later as outcome variable. Controls: Age, German citizenship, educational qualification, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects. Robust standard errors in parentheses.

Table 1.6: Characteristics of Marginal Applicants by Gender (in %)

	Men			Women		
	All workers	Always takers	Marginal applicants	All workers	Always takers	Marginal applicants
	(1)	(2)	(3)	(4)	(5)	(6)
German citizen	90.7	87.9	85.1	84.4	84.8	91.4
Age >=40yrs	63.6	43.2	50.9	59.4	25.6	49.2
Married	64.5	60.1	51.4	55.8	36.9	43.0
Children	76.0	72.9	60.0	68.8	54.3	80.2
Tenure at firm <2yrs	36.1	51.0	51.6	35.2	61.3	37.9
Tenure at firm 2-5yrs	41.1	40.2	38.0	41.5	30.7	41.8
Tenure at firm >=5yrs	22.8	8.7	10.4	23.3	8.0	20.3
Graduate degree	51.5	51.4	61.0	48.4	39.0	74.7
Full-time	97.4	97.9	100.0	73.2	82.4	79.9
High performance	56.1	57.8	63.6	55.4	53.8	73.6
High potential	28.7	44.3	36.1	27.7	43.1	73.4
Technical position	71.0	64.9	66.3	45.5	29.1	32.3
Low-level position	67.8	72.5	76.4	68.8	77.4	83.6
First-level leadership position	12.8	10.8	8.0	12.4	5.8	4.9
Time in position <2yrs	36.7	37.7	43.8	37.3	40.7	20.0
Time in position 2-5yrs	36.2	40.0	37.2	36.2	43.8	63.0
Time in position >=5yrs	27.1	22.4	19.0	26.4	15.5	17.0
Nominated to succession list	1.6	2.3	5.0	1.6	2.9	8.7
Applied 12 months before	2.6	11.7	2.4	2.4	9.4	4.3

Notes: This table illustrates results from a complier analysis by gender, as described in Section 1.6. Each number is based on a separate regression including controls and represents an adjusted mean (in %). Columns 1 and 4 show means for all workers, Columns 2 and 5 represent always takers, and Columns 3 and 6 reflect marginal applicants, who only apply if managers rotate and talent hoarding temporarily abates. A technical position is defined as a job related to engineering, IT, quality management, or production. Low-level positions are defined as positions at low hierarchy levels without leadership responsibility (i.e. individual contributors). First-level leadership represents positions with limited leadership responsibility, such as team leaders. Controls: Age, German citizenship, educational qualification, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects.

Table 1.7: Application Effects of Manager Rotation by Manager Characteristics

Dependent variable: Workers' internal applications

Panel A: Manager Attributes

	Manager gender		Manager vs worker gender		Manager age	
	Male	Female	Opposite	Same	Old	Young
	(1)	(2)	(3)	(4)	(5)	(6)
Manager Rotation	0.0211	0.0315	0.0268	0.0210	0.0226	0.0234
	(0.003)	(0.009)	(0.006)	(0.003)	(0.003)	(0.006)
Outcome Mean	0.028	0.028	0.028	0.028	0.028	0.028
Adj R-squared	0.013	0.012	0.012	0.012	0.012	0.013
P-value of t-test	0.2612		0.6329		0.8956	
Observations	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx

Panel B: Manager Quality

	Experience as manager		Manager performance		Team absenteeism	
	High	Low	High	Low	High	Low
	(1)	(2)	(3)	(4)	(5)	(6)
Manager Rotation	0.0193	0.0238	0.0206	0.0181	0.0207	0.0201
	(0.005)	(0.005)	(0.004)	(0.005)	(0.005)	(0.005)
Outcome Mean	0.028	0.028	0.028	0.028	0.028	0.028
Adj R-squared	0.012	0.013	0.013	0.012	0.012	0.012
P-value of t-test	0.5487		0.7240		0.9353	
Observations	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx

Notes: This table illustrates heterogeneity in the effect of rotations on worker applications by the characteristics of the rotating manager. Each coefficient stems from a separate regression based on Equation 1.3, where the rotation event is restricted to managers with a given characteristic. Panel A compares application effects by manager gender (Columns 1 and 2), manager vs worker gender (Columns 3 and 4), and manager age split at the sample median of 40 years (Columns 5 and 6). Panel B compares application effects by manager quality as measured by experience as manager at the firm (Columns 1 and 2), managers' own performance rating (Columns 3 and 4), as well as absenteeism rates (Columns 5 and 6). All splits in Panel B are with respect to the top and bottom quartile of the respective measure. In both panels, I find no statistical difference in the effect between each pair-wise comparison as indicated by the p-value of the corresponding t-test. Controls: Female, age, German citizenship, educational qualification, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects.

Table 1.8: Impact of Manager and Worker Gender on Talent Hoarding Propensity

	Manager-level		Worker-level	
	Potential rating	Succession nomination	Potential rating	Succession nomination
	(1)	(2)	(3)	(4)
Female	-0.003	-0.118	-0.013	0.054
	(0.217)	(0.228)	(0.012)	(0.036)
Outcome Mean	0.2875	0.3135	0.3768	0.0160
Av ME for Women	-0.0007	-0.0239	-0.0019	0.0010
Gender Gap in %	-0	-8	-1	6
Prob > chi2	0.398	0.341	0.000	0.000
Observations	1,xxx	1,xxx	3xx,xxx	3xx,xxx

Notes: This table examines the impact of gender on managers' decisions to make worker talent visible. Each column is based on a separate logit regression where the regressor of interest is whether the manager is female (Columns 1 and 2) or the worker is female (Columns 3 and 4). Columns 1 and 2 are at the manager-level and estimate the propensity that managers manipulate worker visibility through suppressing potential ratings or nominations to succession lists. Controls include manager age, marital and family status, experience at the firm, division, functional area, and location. Columns 3 and 4 are at the worker-level and estimate the propensity that workers are made visible through rating them as potentials or nominating them to a succession list. Controls include worker age, tenure, schooling, nationality, married, kids, parental leave, position title, division, function, location, full time, hours, leadership, direct reports, past mobility, quarter, and performance. Robust standard errors in parentheses.

Chapter 2

The Broken Rung: Gender and the Leadership Gap

2.1 Introduction

The previous chapter demonstrated that talent hoarding causes high-quality women to forgo high-stakes applications for higher-level positions in which they would have performed well, exacerbating gender inequality in representation and pay at the firm. Besides talent hoarding, there are many other factors that may differentially affect men and women's career progression in organizations. Up to date, remarkably little empirical evidence exists on when and why gender gaps in career progression first emerge in the leadership hierarchy. Answering these questions requires the ability to distinguish between different levels of leadership. However, most datasets only contain coarse measures of job hierarchy.

Due to these data limitations, most attention has been devoted to the fact that women are less likely to hold top leadership positions than men. While 47% of S&P 500 workers are female, women only make up 6% of CEOs (Catalyst, 2021). Accordingly, most attempts to increase female leadership have focused on the very top of the job ladder, as exemplified by the increasing number of countries that have established female quotas for corporate boards. However, recent work indicates that increased representation in top positions does not necessarily trickle down to lower rungs of the job ladder (Bertrand, Black, Jensen and Lleras-Muney, 2018, Maida and Weber, 2020), bringing into question whether addressing representation in top positions is sufficient to increase gender equality throughout the leadership hierarchy. This lack of evidence highlights the importance of identifying when and why gender differences first occur and which areas of the job hierarchy would benefit most from increased attention.

To identify when gender differences in career progression first occur, this study collects novel personnel records in collaboration with a large manufacturing firm. The data contain an unusually rich set of job characteristics, allowing me to construct a new and granular measure of job hierarchy that captures career progression along the job ladder, not just at the top. I find that the transition to first-level leadership positions represents a key bottleneck for women's career progression. This early leadership gap occurs across both male-dominated and female-dominated areas of the firm and is not fully explained by employee characteristics, such as working hours or family status. My results demonstrate that the key driver of this bottleneck are gender differences in internal promotions, not differential entry to or exit from the firm. Women who make it to first-level leadership positions are not less likely to get promoted than men, rejecting the common notion of a glass ceiling at the firm. My findings suggest that analyzing why women are less likely to take on first leadership levels is of first-order importance in order to improve overall gender equality.

The data in this study come from a large multinational firm that employs over 200,000 workers and is one of the largest manufacturers in Europe. To examine internal career progression to higher-level positions, I focus my analysis on the firm's largest internal labor market, consisting of over 30,000 white-collar and management employees in Germany. One advantage of my setting is that the firm employs over 200 occupations, capturing a broad set of positions characterized by female shares that vary between 9% in engineering and 69% in HR. The firm's workforce is comparable to that of other large German firms in terms of demographics and female representation.

A key empirical challenge in studying leadership progression is the complexity of the job ladder. As in most organizations, there are several levels of leadership positions at the firm, which differ in the amount of authority and autonomy over decision-making. The usual first step towards becoming a senior leader (e.g. department head) is to take on a first-level leadership position. First-level leadership positions involve responsibility for managing a small team or project and are distinct from lower-level positions, which do not involve any leadership responsibilities. Identifying a position's leadership level and comparing it to other positions is difficult. Position titles do not always reflect the level of leadership a position entails. In addition, similar position titles may not be comparable across functional areas. For example, many senior marketing positions require leading a team, while senior engineering positions are more likely to have responsibility over a project or product, but not a team.

The difficulty of measuring complex job hierarchies has represented a key impediment for previous research, leading researchers to focus on relatively narrow settings to study gender differences in career progression, such as supermarket workers (Ransom and Oaxaca, 2005), lawyers (Azmat and Ferrer, 2017), central bankers (Hospido, Laeven and Lamo, 2019), and academics (Bosquet, Combes and Garcia-Penalosa, 2019). Focusing on a narrow setting has the advantage of circumventing the need to define a consistent job hierarchy measure that tracks leadership levels across many different occupations and functional areas, but precludes a broader analysis of job mobility.[1] In my data, 15% of promotions occur across functional areas (e.g. HR to IT), and multiple career paths exist within each functional area (e.g.

[1]In the literature on internal labor markets, a common approach is to infer the job hierarchy from flows between position titles when focusing on small labor markets with few different occupations (Baker et al., 1994, Huitfeldt et al., 2021). However, given the complex structure of the internal labor market at the large manufacturer, such approaches are not well suited to this study's setting.

recruiters vs. talent management specialists in HR), indicating that a broader analysis of job mobility is necessary to fully capture employees' career progression.

A key advantage of this study is that it applies a granular measure of job hierarchy that can be used to compare leadership levels across different career paths. To construct this hierarchy measure, I leverage the detailed personnel records to combine three key dimensions of leadership responsibility that directly capture employees' authority and autonomy over decision-making and which are also comparable across occupations: the cumulative number of direct reports, the reporting distance to the CEO, and the extent of managerial autonomy. My hierarchy ranking is the first principal component of these three dimensions, providing a consistent ordering of all positions at the firm.

I use the continuous hierarchy ranking to identify points along the leadership hierarchy that represent bottlenecks for female representation. Female shares drop substantially around the transition to first-level leadership positions, from 22% to 12%. In stark contrast to the common notion of a glass ceiling at higher-level leadership positions, there appear to be no large bottlenecks at higher levels. Female representation falls only modestly after the transition to first-level leadership positions, from 12% to 7% at the highest levels. This pattern exists across different areas of the firm, irrespective of a functional area's gender composition. First-level leadership positions appear to represent the key bottleneck for female representation, both in male-dominated and female-dominated functional areas.

The bottleneck at first-level leadership positions is driven by gender differences in internal promotions. Women a lower-levels of the hierarchy are 68% less likely to move to first-level leadership positions. This promotion gap is persists across different employee groups and cannot be fully explained by differences in hours worked, family demands, or educational qualifications. However, employees at first-level leadership positions do not exhibit gender differences in subsequent promotions. While previous research has almost exclusively focused on the possibility that a glass ceiling at higher levels is a key determinant of gender differences (Blau and Kahn, 2017), my findings echo a growing narrative among practitioners that a broken rung at the beginning of the career ladder represents a major impediment for women's career progression (McKinsey and LeanIn.Org (2021)). In contrast to internal promotions, differential entry to and exit from the firm do not play critical roles in explaining female underrepresentation at higher levels. Even though women are more likely than men to enter the firm at higher hierarchy levels, the vast majority of positions at the firm are filled by internal candidates, highlighting the importance of internal promotions. In addition, women are less likely than observationally similar men to exit the firm.

This study contributes to two strands of literature in economics. First, a large literature has documented substantial gender differences in labor market outcomes, particularly pay (Goldin, 2014, Blau and Kahn, 2017). Even though gender differences in career progression appear to be an important contributor to disparities in pay (Bronson and Thoursie, 2019), most of the literature has focused on analyzing gender pay gaps, with little consensus on differences in career progression. Previous work on gender differences in representation has mostly focused on the highest levels of the job ladder (Bertrand et al., 2018, Maida and Weber, 2020), likely due to data limitations precluding the analysis of the full job hierarchy. Understanding how representation evolves along the entire job ladder, however, seems particularly important given that previous work on gender pay gaps has documented that critical differences emerge early in employees' careers (Bertrand, Goldin and Katz, 2010, Goldin, Kerr, Olivetti and Barth, 2017). By collecting a new dataset, this study demonstrates that a key gender promotion gap occurs early on in employees' careers, at the transition to first-level leadership positions. Second, by constructing a novel measure on internal job hierarchy, this paper also contributes to the large body of work that analyzes the functioning of the internal labor market (Waldman, 1984, Milgrom and Oster, 1987, Baker et al., 1994, Benson et al., 2019, Huitfeldt et al., 2021).

The rest of this chapter proceeds as follows. Section 2.2 introduces the setting and data. Section 2.3 describes the construction of the new measure of job hierarchy, which allows me to make granular distinctions between different leadership levels. Section 2.4 uses the hierarchy measure to provide new facts about gender and the leadership hierarchy. Section 2.5 discusses my findings in light of the common notion of a glass ceiling and concludes.

2.2 Setting and Data

This paper analyzes gender differences in career progression in a large multinational firm that employs over 200 typical occupations in both male-dominated and female-dominated functional areas. I collect rich personnel records that contain detailed information on positions' leadership responsibility.

2.2.1 Firm Overview

The large multinational firm that I study employs over 200,000 workers around the world. To analyze career progression to higher-level positions, I focus my sample on on all 30,000 white-collar and management employees, who are either already in or could ultimately attain management positions at the firm. I further restrict my sample to all employees based on Germany which represents the largest internal labor market at the firm. The median employee in my sample holds an engineering position, but as a large manufacturer, the firm also employs many female-leaning occupations, such as marketing, finance, and HR, allowing me to assess gender disparities across both male-leaning and female-leaning areas.

Germany is similar to other Western countries, such as the United States, in terms of observed gender disparities in the workplace. In 2019, the gender pay gap for full-time employees was 14% in Germany and 18% in the United States (OECD, 2022). Both in Germany and the United States, women are underrepresented on corporate boards and hold only 29% and 24% of seats, respectively (Deloitte, 2021). Gender role attitudes on women's labor force participation, as measured by the World Value survey, are also similar in Germany compared to the United States (Fortin, 2005). One notable exception are norms towards working mothers, which are typically more conservative in Germany than in the

United States. To account for this potential difference of the German context, I restrict supplementary results to only employees without children.

Table 2.3 provides summary statistics for my main analysis sample which consists of over 400,000 employee-by-quarter observations from 2015 to 2019.[2] Employee tenures at the firm tend to be long, with an average tenure of 13 years, allowing me to follow employees' internal career progression over time. Because I restrict to white-collar and management employees with regular employment contracts (as opposed to those with marginal employment such as mini jobs), employee qualifications in my sample is high. The average employee holds a Bachelor's degree and 92% of employees work full-time.

Women at the firm are underrepresented in higher-level positions. In my overall sample, women represent 21% of employees, which is consistent with the high share of technical occupations that this manufacturing firm employs. However, while top leadership positions, such as senior executives, corporate board members, C-suite positions, are held by 8% of men, this is only the case for 4% of women. Similarly, 21% of men but only 10% of women have responsibility over a team. Even conditional on leading a team, women have significantly fewer direct reports.

The demographics of the employees at the firm are comparable to other large manufacturing firms in Germany. In Appendix Table A.1, I compare employees in my sample to those employees in large manufacturing firms in the BiBB, a representative survey of the German workforce conducted in 2018. I find very similar patterns with respect to most employee characteristics (e.g. gender, age, German citizenship, martial and family status). In addition, the BiBB illustrates that the gender leadership gap in the firm I study aligns with broader patterns of female underrepresentation in Germany, suggesting that this setting is fairly typical for German firms.

2.2.2 Data

I collect the firm's internal personnel records, which provide detailed information on demographics and position characteristics for all employees in my sample. I collect detailed demographic information from the personnel records, including gender, age, citizenship, educational qualifications (such as highest degree, major, and institution), marital status, family status, and parental leave history at the firm. The records also contain detailed position characteristics, such as occupation and position title, functional area (e.g. marketing versus engineering), business unit, location, leadership responsibility, and the reporting distance to the CEO. I supplement these data with payroll information, capturing employees working hours, earnings, and bonus payments. Finally, I collect information on worker evaluations, such as performance and potential ratings.

For my empirical analysis, I construct an employee-by-quarter dataset spanning 2015 to 2019. I restrict my sample to only white-collar and management employees who are regular employees at the firm (e.g. excluding marginal employment such as mini jobs). This dataset

[2] To maintain confidentiality, I do not disclose the exact number of employees in my sample.

allows me to construct a granular measure of job hierarchy and to examine gender differences in internal promotions. I collapse the data to a quarterly level. My main analysis sample contains over 400,000 employee by quarter observations and covers over 30,000 unique white-collar and management employees.

2.3 Constructing a New Measure of Job Hierarchy

The goal of this study is to identify when and why gender differences occur in the leadership hierarchy. Detecting when bottlenecks for female representation arise requires a granular hierarchy measure that makes fine distinctions between leadership levels.

Classic theories on the functioning of the firm characterize higher hierarchy levels as exhibiting more authority or a larger span of control (Rosen, 1982). Managers are typically thought of as those with supervisory, coordination, and arbitration functions (Holmstrom and Tirole, 1989). In practice, capturing granular differences in managerial or leadership responsibility has been empirically challenging. Most common datasets, such as matched employer-employee data, do not contain information on the degree of leadership over a team or the extent of autonomy employees have in their decision-making. Even studies that draw on internal records from firms rarely use direct information on the extent of leadership responsibility an employee has.

Previous studies have typically followed one of two strategies. A popular approach to infer a position's hierarchy level, particularly in settings in which little other hierarchy information is available, uses pay-based measures, such as individual salaries or the salary band a position is assigned to (Bronson and Thoursie, 2019, Cullen and Perez-Truglia, 2019, Benson et al., 2021). For the objective of this study, however, such pay-based measures are not well suited. Salary bands or job grades are often relatively coarse, particularly for lower-level positions which represent the focus of this study. Salary bands are not designed to distinguish differences in the authority or autonomy an employee has, but often group positions according to factors such as market wages. Both the individual salary and the salary band an employee is assigned to are likely influenced by factors that are unrelated to the job hierarchy, such as candidates' negotiation success or whether outside offers are matched. Previous work has documented substantial gender differences in negotiation (Bertrand, 2011), suggesting that gender gaps in pay-based hierarchy measures may be distorted.

Another common strategy to infer job hierarchy is based on flows between occupation or positions titles. This approach has been particular popular in settings in which internal labor markets are relatively homogeneous and are comprised by a limited number of positions. For instance, the firms studied in Huitfeldt et al. (2021) have an average of ten position titles, and the firms in Baker et al. (1994) (a medium-sized service-sector firm) and Ransom and Oaxaca (2005) (a supermarket) have a relatively small set of possible career trajectories. In the firm that I study, the internal labor market consists of over 200 different occupations and multiple non-intersecting career paths, making it difficult to construct a universal hierarchy ranking based on position titles.[3] Moreover, occupation and position titles are often both noisy and relatively coarse.[4] In my sample, 26% of employees share a position title with either their supervisor or their supervisor's supervisor. Assigning them the same measure of job hierarchy would underestimate differences in leadership responsibility and might lead to biased estimates of the gender leadership gap.

To make granular distinctions between leadership levels across different occupations, I construct a new measure of internal job hierarchy that directly captures three key dimensions of leadership responsibility: the cumulative number of direct reports, the reporting distance to the CEO, and the positions' managerial autonomy. I choose the cumulative number of direct reports to distinguish between responsibility over individual contributors versus over team leaders. The reporting distance to the CEO is constructed by linking reporting relationships between supervisors, resulting in 8 different levels which resemble the logic of an organizational chart. The information I use to capture managerial autonomy distinguishes between five different levels, ranging from employees with neither autonomy over working hours nor decision-making, to employees with autonomy over hours but not decision-making, to three groups of employees with increasing autonomy over profit and loss. The key advantage of these inputs is that they directly capture the extent of leadership responsibility that an employee exhibits and that they can be easily compared across different occupations.[5] Another advantage is that these inputs represent common elements of firms' internal personal records and are available in many firms.

The hierarchy ranking I construct is the first principal component of the three dimensions of leadership responsibility, which explains 61% of variation and provides a consistent order of all positions at the firm. The loadings on the first principal component are very granular, capturing over 600 different values. All three inputs are similarly important and load on the first component as follows: $0.5591 \times \{\text{cumulative reports}\} + 0.6336 \times \{\text{managerial autonomy}\} + 0.5348 \times \{\text{reporting distance to the CEO}\}$. The resulting one-dimensional hierarchy ranking ranges from 0 to 100. This approach assigns the lowest hierarchy rankings to entry-level positions (e.g. junior engineering positions), while the position of the CEO receives the highest ranking of 100. Figure 2.1 shows the pyramidal structure of the hierarchy, where 69% of employees are situated at positions with a hierarchy ranking of 20 or below. I use the hierarchy ranking to capture granular dynamics along the leadership hierarchy, allowing me to identify where exactly gender differences in representation occur. For simplicity, I provide descriptive statistics by grouping the continuous hierarchy measure into deciles.

[3] The existence of multiple career paths implies that institutional structures, such as organizational charts, are not sufficient to create a hierarchy measure that is comparable across different parts of the firm.

[4] Note that this caveat is not specific to the firm I study. The occupation titles in my data coincide with the occupation titles the firm reports to the German social security administration, providing the basis for the widely-used German matched employer-employee data.

[5] Note that these inputs are computed for each individual, not each position, since position titles are not sufficiently granular. Nevertheless, the firm's internal records suggest that the extent of leadership responsibility is typically not an outcome of individual negotiation. Instead, the three inputs I use are determined before job openings are advertised.

To evaluate the fit of the hierarchy measure, I begin by comparing it to employee earnings. Figure 2.2 demonstrates that the hierarchy measure, which itself is not based on pay, is strongly correlated with earnings, suggesting that it captures meaningful differences between positions. However, Figure 2.2 also shows that the hierarchy measure is more effective at discerning between hierarchy levels at the bottom of the hierarchy relative to pay. Given that firms do not immediately adjust not surprising, but highlights that pay-based measures are not well-suited to studying differences at lower-levels of the hierarchy However, particularly at the bottom of the hierarchy, there are many positions for which earnings are similar, but the hierarchy ranking substantially differs, suggesting that hierarchy measures solely based on earnings likely underestimates differences in hierarchy.

The hierarchy ranking captures systematic differences. Table 2.2 documents that hierarchy levels, which are constructed by grouping the hierarchy ranking into deciles, differ substantially in terms of characteristics not used to construct the hierarchy measure. As hierarchy levels increase, bonus payments represent larger shares of employees' total compensation (Column 1). Higher shares of bonus payments are usually associated with higher-level positions that have large autonomy. While the vast majority of positions with an hierarchy ranking of 20 or less do not entail leadership of a team, the number of direct reports substantially increases as hierarchy levels rise (Column 2). Positions at higher hierarchy levels are filled by employees with more work experience (Column 3) and higher educational qualifications (Column 4). These differences echo the differences in how positions at different hierarchy levels are advertised with respect to stated job requirements (Table 3.2).

Increases in hierarchy levels represent typical steps in the job ladder, as illustrated by the transition matrix for employees who switch positions. Table 2.1 indicates that employees are most likely to move to adjacent hierarchy levels. Figure 2.3 shows that hierarchy levels correlate well with position title not used to construct the hierarchy measure. Even though position titles do not always reflect leadership responsibility, the share of position titles that reference leadership (e.g. "team lead", "head of department") rises as hierarchy levels increases. This finding suggests that the ordering of positions that the hierarchy measure induces is sensible, further supporting the validity of the hierarchy measure. However, Figure 2.3 also illustrates that using position titles alone fails to distinguish between leadership levels, as many position titles do not sufficiently reflect how much leadership responsibility a position entails. For instance, some positions at high hierarchy levels are labeled as engineering or specialist positions, even though they entail leadership over a department.

While the continuous hierarchy ranking allows me to identify when in the leadership hierarchy gender differences occur, I use data-driven insights to group hierarchy rankings into relevant parts of the leadership hierarchy in order to make my main analysis tractable. I use information on the job characteristics at different levels of their hierarchy ranking that are based on job features recorded in the firm's administrative data or based on how jobs are described in job openings or by employees who currently fill them. For instance, while the majority of employees with a hierarchy ranking of 20 or less have no responsibility over a team, the average employee with a hierarchy ranking between 30 and 40 (40 and 50) has three (six) direct reports (Table 2.2). Motivated by the data, I define positions with a hierarchy ranking between 0 and 20 as lower-level positions, given that these positions generally lack leadership responsibilities (e.g. entry-level engineering position). Positions with an index value between 20 and 40 are defined as first-level leadership positions (e.g. project or team lead), and those above 40 are senior leadership positions (e.g. department or divisional head), with the CEO having a value of 100.

First-level leadership positions represent the first step along the leadership pipeline towards higher-level leadership positions. They differ substantially from lower-level positions in their job attributes (e.g. salary, required work experience, and educational qualifications). These differences are reflected in how these positions are advertised in job postings, which likely influences employees' application choices. Table 3.2 presents characteristics of job attributes in 11,xxx job ads that have been posted to the firm's online job portal from 2015 to 2019, and allows a comparison of lower-level positions (Column 1) to first-level leadership positions (Column 2). Lower-level positions offer significantly lower pay compared to first-level leadership positions. They also require less schooling and work experience. First-level leadership positions, however, are more likely to require frequent business travel and negotiations on-the-job. They are more likely to require English proficiency and strong communication and analytical skills. Leadership positions are almost four times more likely to be described as strategic positions at the firm (i.e. of high importance), and their job ads provide more detail when describing the position and stating job requirements.

To capture employees' likelihood of moving along the leadership pipeline, I define different types of promotions based on the increase in the hierarchy ranking that is associated with an internal job transition. One can conceptualize small promotions as those carrying small increases in hierarchy but no substantial increases in leadership responsibilities. A key outcome of interest regarding employees' career progression is the occurrence of major promotions, such as transitions from lower-level positions without leadership responsibility to first-level leadership positions that require responsibilities over projects or teams. In my data, large increases in leadership responsibility typically occur with increases in the hierarchy ranking of 20 or more. I therefore define major promotions as increases in the hierarchy ranking of 20 or more. For employees at low levels (20 or below), over 89% of major promotions defined in this way result in a transition to a first-level leadership position (e.g. promotion from specialist to team leader). Thus, this definition of major promotions effectively captures key transitions along the leadership pipeline.[6]

2.4 First-Level Leadership Positions as Bottleneck

I document that gender differences in representation arise at the transition to first-level leadership positions and are driven by differences in internal promotions. Women who make it to first-level leadership positions are not less likely than men to advance to subsequent hierarchy levels.

[6]To show robustness to alternative definitions of promotions, I present results using different promotion types along with my baseline estimates.

2.4.1 Female Representation Along the Job Hierarchy

The granular hierarchy ranking allows me to identify when in the job hierarchygender differences in representation first occur. Figure 2.4 documents the female share across the hierarchy ranking. On average, women fill 22% of lower-level positions at the firm, but only 12% of first-level leadership positions and 7% of top leadership positions.[7] This pattern signifies that the sharpest reduction in female representation occurs early in the leadership pipeline, around the transition to first-level leadership positions. Beyond this level, female representation falls only gradually at higher hierarchy levels. This findings echoes general cross-firm patterns in Sweden (Bronson and Thoursie, 2019) and the United States (McKinsey and LeanIn.Org, 2021 and rejects the notion that a glass ceiling at higher leadership levels remains the key barrier to equal gender representation in leadership.

First-level leadership positions represent the key bottleneck for female representation in both male- and female-dominated functional areas. Figure 2.4 shows that the pattern exists across both technical functional areas (e.g. engineering, IT, and production-related positions) and non-technical areas (e.g. HR, marketing, finance, and purchasing), despite substantial variation in the share of female employees. Even though almost half of the lower-level positions in non-technical areas are filled by women, the majority of first-level leadership positions in these areas (as well as higher-level positions) are nonetheless held by men. This finding suggest that there may be underlying causes that deter women from progressing to leadership positions that even apply if these positions are in more female-leaning settings involve leadership over mostly female teams.

2.4.2 Gender Differences in Internal Promotions

In order to test whether the transition to first-level leadership positions indeed represents a bottleneck for womens career progression, I analyze gender differences in career progression separately for employees at lower-levels (hierarchy ranking of 20 or lower) and employees at first leadership levels (hierarchy ranking between 20 and 40).

I estimate gender differences in experiencing a major promotion using a logit regression of an indicator for getting promoted in a given quarter on gender, quarter fixed effects and varying sets of employee demographics and position controls.

$$\Pr(\text{Promoted}_{it} = 1) = \Lambda(\theta_1 \text{Female}_i + \theta_t + \theta_X X_{it}) \tag{2.1}$$

Women in low-level positions are substantially less likely to move to first-level leadership positions than men. Column 1 of Table 2.4, Panel A documents the results from a logit regression of an indicator of experiencing a major promotion in a given quarter on worker gender and quarter fixed effects. Major promotions are defined as an increase in hierarchy ranking of 20. For employees at low levels, a major promotion represents the transition to a first-level leadership position, for instance from a specialist to a team leader position.

[7] I define lower-level positions as positions with a hierarchy ranking of 20 or less. First-level leadership positions are defined as having a hierarchy ranking between 20 and 40.

Women are 68% less likely than men in low-level positions to move to first-level leadership positions in a given quarter. Even after controlling for detailed employee demographics and position controls, this promotion gap persists, as illustrated by Columns 2 and Column 3, respectively. These findings demonstrate that men and women at low-level positions differ substantially in their probability to climb the first rung of the job ladder, even when controlling for observable differences in employee qualifications and position characteristics.

However, among employees at first-level or senior leadership positions, there do not exist any gender differences in subsequent career progression. Panel B of Table 2.4 documents that women who made it to the first leadership level are not less likely to experience a major promotion than men. This finding is robust to including varying sets of controls and contrasts with the more common idea that there is a glass ceiling that is the primary bottleneck for gender equality. Instead, I find that the transition to first-level leadership positions represents the key bottleneck for womens career progression to top leadership positions.

The gender gap in promotions at lower-levels is robust to alternative definitions of promotions. Column 1 of Table 2.7 focuses on small promotions as defined by increases in the hierarchy ranking between 5 and 20. Column 2 of Table 2.7 uses transitions to becoming any type of team leader for the first time. Column 3 of Table 2.7 restricts to transitions that induce an increase of more than five direct report. I find meaningful and statistically significant gender differences in each type of promotion definition, indicating that gender differences in early promotions go beyond my preferred definition of promotion. However, the results in Table 2.7 also demonstrate that gender differences will likely be understated depending on the coarseness of the hierarchy definitions in use.

2.4.3 Gender Differences in Exit from the Firm

A common hypothesis for women's underrepresentation at higher-levels conjectures a leaky pipeline. Consequently, differential exit at higher hierarchy levels may contribute to decreasing female shares. I begin to test this hypothesis by examining gender differences in the likely to exit the firm in a given quarter, separately for employees at low levels and at the first leadership level. I analyze gender differences in firm exit by estimating a logit regression based on Equation 2.2 that uses an indicator for whether an employ exits the firm in a given quarter as outcome variable.

$$\Pr(\text{Exit}_{it} = 1) = \Lambda(\theta_1 \text{Female}_i + \theta_t + \theta_X X_{it}) \tag{2.2}$$

Table 2.5 shows the coefficients from a logit regression of an indicator for exiting the firm on a coefficient for female, quarter fixed effects and an increasing set of worker controls. Panel B of Table 2.5 shows that women at first-level leadership positions are not more likely than men to exit the firm. This finding indicates that differential exit at first-level leadership positions is not a key determinant of the decrease in female shares at higher leadership levels. When taking into account underlying differences between position types, Column 3 of Table 2.5, Panel A documents that also at lower levels women are less likely to exit the firm.

2.4.4 Gender Differences in Entry to the Firm

Another potential channel that may exacerbate the decrease in female shares at higher leadership levels is differential entry to the firm. I use a logit regression based on Equation 2.3 to test whether women are less likely to enter the firm than men in a given quarter.

$$\Pr(\text{Entry}_{it} = 1) = \Lambda(\theta_1 \text{Female}_i + \theta_t + \theta_X X_{it}) \tag{2.3}$$

Panel B of Table 2.6 documents that if anything, women are more likely to enter the firm at first-level leadership positions than men. This finding suggests that differential entry is not a key driver of the bottleneck at first-level leadership positions. In addition, since 87% of high-level positions at the firm are filled by internal candidates, entry into the firm is much less important for representation at higher-levels than to internal promotions. At lower levels, women are less likely to enter the firm than similar men (Panel A of Table 2.6). In unreported results I use information on all over 200,000 external applicants who have applied to the firm between 2015 and 2019 and find that conditional on applying to low level positions, women are not less likely to be hired into the firm than men.

2.5 Conclusion

This chapter provides new evidence on the determinants of female underrepresentation in leadership positions. I use novel personnel data from a large manufacturing firm to document the existence of a major bottleneck in women's career progression at the transition to first-level leadership positions.Conditional on holding a first-level leadership position, women face similar career trajectories as men.

My results contrast the common notion that a glass ceiling at higher-level leadership positions is the key barrier to gender equality. This finding highlights the importance of focusing on gender differences that arise early on in the leadership pipeline. If gender differences in early promotions represent a common bottleneck, policies that provide early exposure to leadership and encourage women to try out first-level leadership, such as job rotations or mentoring programs, could be effective tools to bridge the leadership gap.

These findings raise an outstanding question is: does fixing the broken rung increase overall gender equality? The absence of gender differences at higher levels raises the question whether the observed gender gaps at lower levels are driven by heterogeneity (some women are more similar to men and these are the women who fill leadership positions) versus state-dependence (when women are exposed to first leadership experience it changes their application behavior). This distinction is an important component of future work, as it determines whether policy interventions targeted at increasing female shares in first-level leadership positions can be effective for increasing female representation throughout the leadership pipeline.

Figure 2.1: Distribution of Hierarchy Index

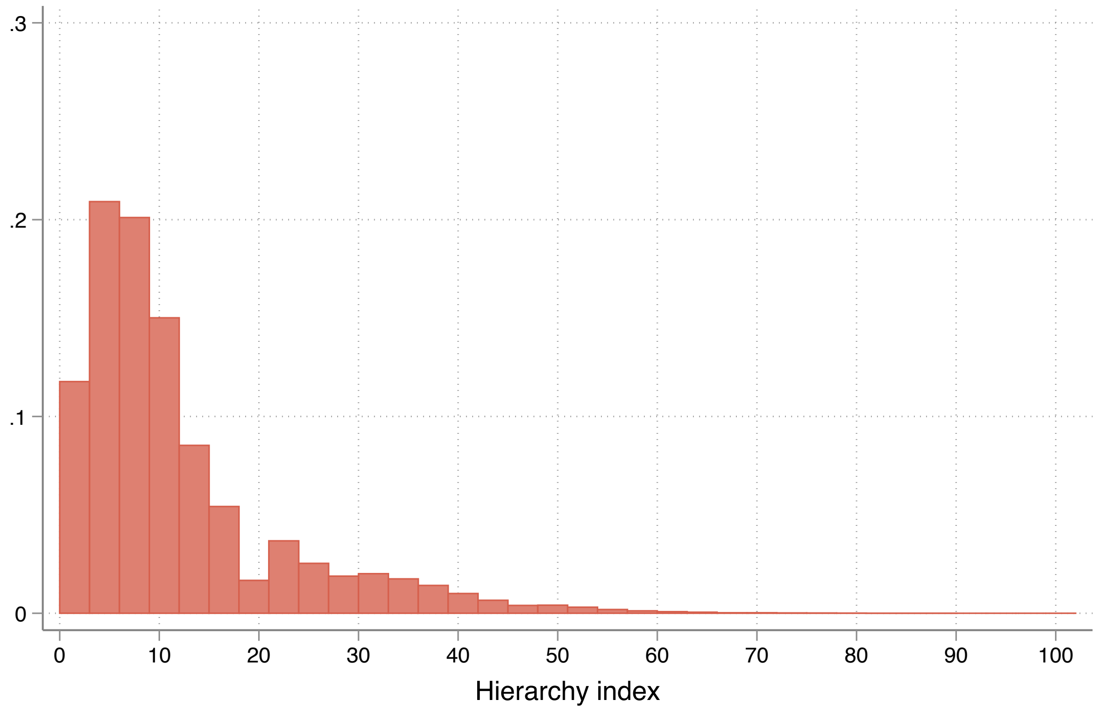

Notes: This figure displays the distribution of the continuous index of job hierarchy in my sample. Low values represent low-level positions, such as entry-level engineering jobs. The highest value of 100 represents the CEO. The majority of workers (69%) are situated at positions with a hierarchy index of 20 or below. The total number of observations is 4xx,xxx.

Figure 2.2: Fit of Hierarchy Index Compared to Earnings

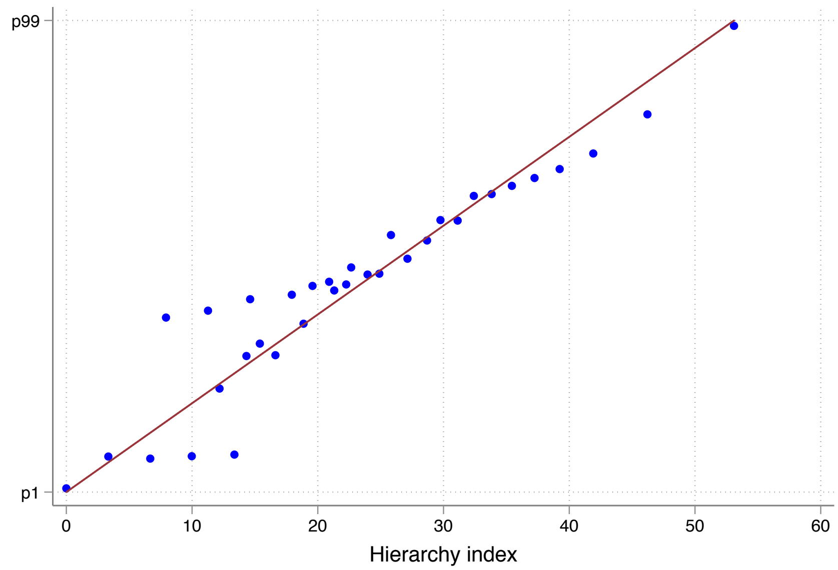

Notes: This figure displays the correlation between the continuous hierarchy index and the percentile of employees' log real annual earnings. The total number of observations is 4xx,xxx.

Figure 2.3: Composition of Position Titles by Hierarchy Level

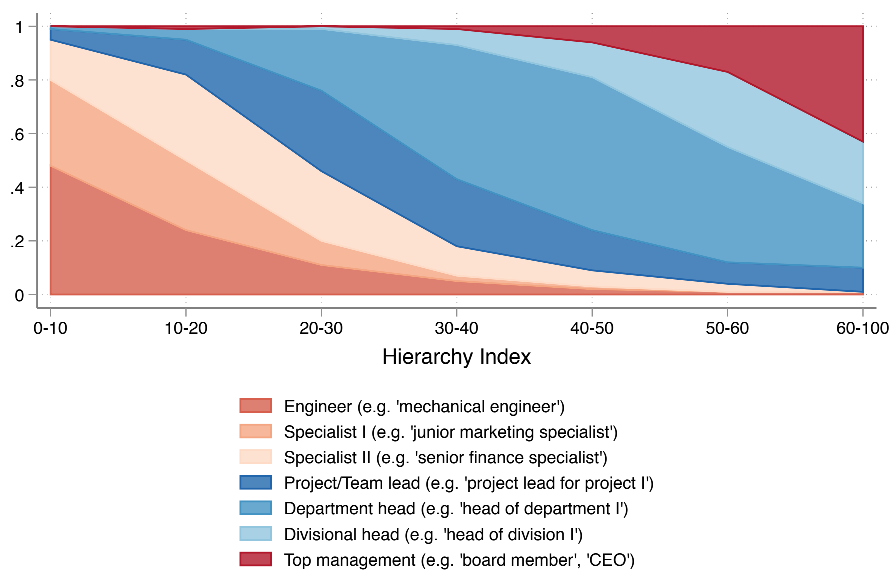

Notes: This figure displays the composition of position titles by deciles of the hierarchy index. I extract key terms from position titles that likely indicate the type of leadership responsibility that the position entails (e.g. "engineer", junior specialist", "head of department"). The total number of observations is 4xx,xxx.

Figure 2.4: Female Share by Decile of Hierarchy Index

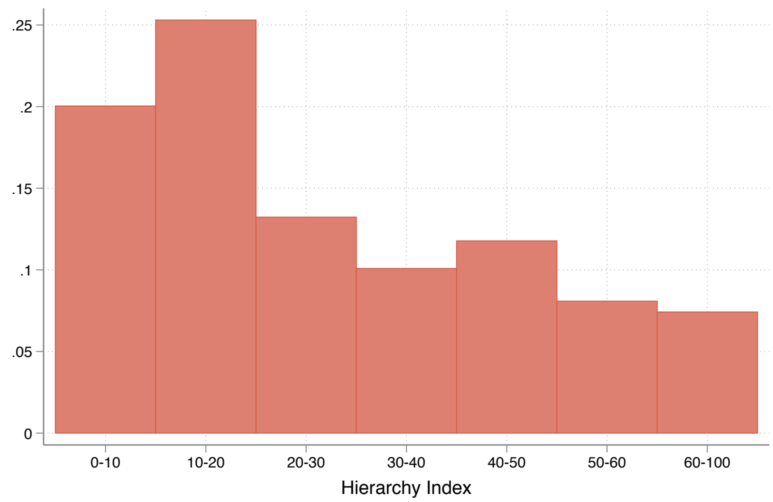

Panel A. Sample Average

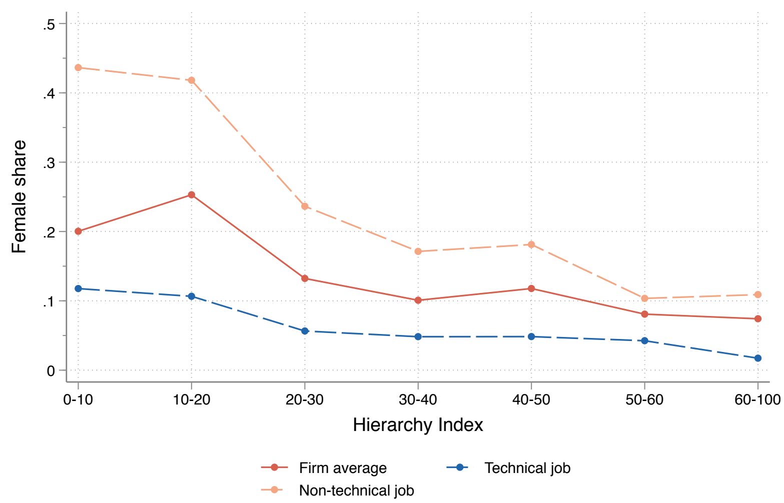

Panel B. Across Different Functional Areas

Notes: This figure illustrates the share of female employees at a given decile of the hierarchy index, which is constructed via PCA and combines three dimension of leadership responsibility: cumulative number of reports, reporting distance to the CEO, and managerial autonomy. Values below 20 represent low-level positions without leadership responsibility. Positions with an index between 20 and 40 represent first-level leadership positions, such as team or project leaders, followed by higher-level leadership positions. Panel A reports the sample average. Panel B distinguishes different functional areas at the firm. Technical areas include engineering, IT, quality management, and production-related positions, with non-technical jobs capturing the remainder of the sample. N=4xx,xxx.

Table 2.1: Transition Matrix

	Hierarchy Index$_{t+1}$						
Index$_t$	0-10	10-20	20-30	30-40	40-50	50-60	60-100
0-10	83.1	16.0	0.6	0.2	0.0	0.0	0.0
10-20	9.4	73.0	14.3	2.7	0.3	0.3	0.1
20-30	0.0	12.9	65.6	16.4	4.3	0.8	0.0
30-40	0.0	3.1	8.5	65.9	19.4	3.1	0.0
40-50	0.0	0.0	12.5	20.0	55.0	7.5	5.0
50-60	0.0	0.0	10.0	0.0	40.0	50.0	0.0
60-100	0.0	0.0	0.0	25.0	0.0	50.0	25.0

Notes: This table represents the transition matrix for employees who switch positions between year t and the subsequent year $t+1$. Each coefficient represents the share of employees (in %) who start out at a hierarchy decile in year t and transition to a respective decile in year $t+1$. For instance, 83.1% of employees from the first hierarchy decile transition to a position in the same decile. The total number of observations is 2,xxx.

Table 2.2: Characteristics by Hierarchy Index

Hierarchy index	Fraction bonus/pay (1)	Number reports (2)	Tenure (years) (3)	Share >= BA (4)	External hire (5)
0-10	0.11	0.03	11.7	0.42	0.024
10-20	0.13	0.32	13.5	0.56	0.017
20-30	0.16	3.04	16.2	0.72	0.007
30-40	0.22	5.61	17.6	0.77	0.003
40-50	0.30	7.74	17.3	0.79	0.003
50-60	0.43	9.69	16.8	0.79	0.002
60-100	0.45	14.98	17.2	0.78	0.001
N	4xx,xxx				

Notes: This table reports average characteristics by decile of hierarchy index: bonus pay as fraction of total pay (Column 1), number of direct reports (Column 2), firm tenure in years (Column 3), share with at least a Bachelor's degree (Column 4), share who are hired from externally in a given year (Column 5). Note that for the construction of the hierarchy measure the number of all reports is used, not the number of direct reports. The total number of observations is 4xx,xxx.

Table 2.3: Summary Statistics

	Men (1)	Women (2)	Δ (3)
Demographics			
German nationality	0.90	0.86	0.04***
Age (years)	44.2	40.5	3.7***
Schooling (years)	16.05	14.91	1.14***
Firm tenure (years)	13.6	12.2	1.4***
Married	0.65	0.51	0.14***
Children	0.77	0.65	0.12***
On parental leave	0.02	0.08	-0.06***
Position			
Limited contract	0.01	0.03	-0.02***
Full-time	0.97	0.73	0.24***
Weekly hours	42.0	37.8	4.3***
Technical position	0.71	0.33	0.39***
Salary decile	5.9	3.8	2.2***
Team leadership	0.21	0.10	0.11***
Direct reports	6.09	5.10	0.99***
Cumulative reports	25.36	9.01	16.36***
Top management	0.08	0.04	0.04***
Number of teammates	9.1	8.6	0.5***
Evaluations			
High performance	0.56	0.47	0.10***
High potential	0.27	0.26	0.02***
Observations	3xx,xxx	1xx,xxx	4xx,xxx

Notes: This table reports summary statistics for my quarterly analysis sample from 2015 to 2019. Column 1 restricts to male employees, Column 2 restricts to female employees, and Column 3 represents the differences between men and women. Technical positions refer to positions in engineering, IT, or production-related areas. Salary deciles are computed based on the distribution of employees' real annual earnings. Top management positions refer to senior executives, board members, and the CEO. Performance and potential ratings are employee evaluations conducted by the direct supervisor. *p< 0.10, **p< 0.05, ***p< 0.001.

Table 2.4: Gender Differences in Promotions

Panel A: Employees at Low Levels

	(1)	(2)	(3)
Female	-0.665**	-0.732**	-0.743*
	(0.323)	(0.370)	(0.395)
Demographic Controls	-	X	X
Position Controls	-	-	X
Outcome Mean	0.0003	0.0003	0.0003
Observations	2xx,xxx	2xx,xxx	2xx,xxx

Panel B: Employees at First-Leadership Levels

	(1)	(2)	(3)
Female	0.071	0.012	0.096
	(0.321)	(0.355)	(0.372)
Demographic Controls	-	X	X
Position Controls	-	-	X
Outcome Mean	0.0013	0.0013	0.0013
Observations	5x,xxx	5x,xxx	5x,xxx

Notes: This table documents gender differences in major promotions from 2015 to 2019. Panel A restricts to employees in low-level positions without leadership responsibility (hierarchy ranking of 20 or lower). Panel B restricts to employees at first-level leadership positions (hierarchy ranking between 20 and 40). Major promotions represent transitions to higher-level positions. Each coefficient stems from a separate logit regression of an indicator for getting promoted on worker gender, quarter fixed effects, and a varying set of controls. Demographic controls: Age, degree, nationality, marital and family status, parental leave, tenure. Position controls: Division, functional area, full-time status, weekly hours, performance rating. Robust standard errors in parentheses. *p< 0.10, **p< 0.05, ***p< 0.001.

Table 2.5: Gender Differences in Exit

Panel A: Employees at Low Levels			
	(1)	(2)	(3)
Female	0.154***	0.087	-0.212***
	(0.049)	(0.054)	(0.064)
Demographic Controls	-	X	X
Position Controls	-	-	X
Outcome Mean	0.0067	0.0067	0.0067
Observations	2xx,xxx	2xx,xxx	2xx,xxx
Panel B: Employees at First-Leadership Levels			
	(1)	(2)	(3)
Female	-0.097	-0.078	-0.211
	(0.173)	(0.190)	(0.205)
Demographic Controls	-	X	X
Position Controls	-	-	X
Outcome Mean	0.0053	0.0053	0.0053
Observations	5x,xxx	5x,xxx	5x,xxx

Notes: This table documents gender differences in firm exit from 2015 to 2019. Panel A restricts to employees in low-level positions without leadership responsibility (hierarchy ranking of 20 or lower). Panel B restricts to employees at first-level leadership positions (hierarchy ranking between 20 and 40). Each coefficient stems from a separate logit regression of an indicator for exiting the firm on worker gender, quarter fixed effects, and a varying set of controls. Demographic controls: Age, degree, nationality, marital and family status, parental leave, tenure. Position controls: Division, functional area, full-time status, weekly hours, performance rating. Robust standard errors in parentheses. *p< 0.10, **p< 0.05, ***p< 0.001.

Table 2.6: Gender Differences in Entry

Panel A: Employees at Low Levels			
	(1)	(2)	(3)
Female	-0.047	-0.111***	-0.102**
	(0.038)	(0.042)	(0.047)
Demographic Controls	-	X	X
Position Controls	-	-	X
Outcome Mean	0.0122	0.0122	0.0122
Observations	2xx,xxx	2xx,xxx	2xx,xxx
Panel B: Employees at First-Leadership Levels			
	(1)	(2)	(3)
Female	1.050***	0.293	0.294
	(0.152)	(0.181)	(0.207)
Demographic Controls	-	X	X
Position Controls	-	-	X
Outcome Mean	0.0033	0.0033	0.0033
Observations	5x,xxx	5x,xxx	5x,xxx

Notes: This table documents gender differences in entry to the firm from 2015 to 2019. Panel A restricts to employees in low-level positions without leadership responsibility (hierarchy ranking of 20 or lower). Panel B restricts to employees at first-level leadership positions (hierarchy ranking between 20 and 40). Each coefficient stems from a separate logit regression of an indicator for entering the firm on worker gender, quarter fixed effects, and a varying set of controls. Demographic controls: Age, degree, nationality, marital and family status, parental leave, tenure. Position controls: Division, functional area, full-time status, weekly hours, performance rating. Robust standard errors in parentheses. *p< 0.10, **p< 0.05, ***p< 0.001.

Table 2.7: Gender Differences in Promotions

Panel: Employees at Low Levels			
	Small Promotion (1)	Become Team Lead (2)	More Direct Reports (3)
Female	-0.086**	-0.301***	-0.688***
	(0.041)	(0.082)	(0.182)
Demographic Controls	-	X	X
Position Controls	-	-	X
Outcome Mean	0.0184	0.0052	0.0017
Observations	2xx,xxx	2xx,xxx	2xx,xxx

Notes: This table documents gender differences in promotions from 2015 to 2019 for employees at low-level positions using alternative promotion definitions. Column 1 focuses on small promotions which are defined as increases in the hierarchy ranking between 5 and 20. Column 2 focuses on becoming a team leader for the firs time. Column 3 restricts to increases in the number of direct reports of more than 5. Each coefficient stems from a separate logit regression of an indicator for getting promoted on worker gender, quarter fixed effects, and a varying set of controls. Controls: Age, degree, nationality, marital and family status, parental leave, tenure, division, functional area, full-time status, weekly hours, performance rating. Robust standard errors in parentheses. *p< 0.10, **p< 0.05, ***p< 0.001.

Chapter 3

Gender, Leadership, and Differences in Job Applications

3.1 Introduction

The central question raised by the previous chapter is, why are women less likely to advance to first-level leadership positions than men? Identifying the drivers of female underrepresentation is difficult, because the promotion gap can can arise due to both labor supply and labor demand factors. In addition, taking on a leadership position for the first time typically comes with several changes in terms of job features and work environments, making it difficult to pin down what exactly it is that makes leadership positions less appealing to women. Previous research has mostly focused on job attributes that are easily measurable, such as pay or hours of a position. Responsibility over a team, which is a very common feature of first-level leadership positions, however, has been understudied, likely because of a lack of available data.

This study provides new insights on the drivers of the gender leadership gap by combining rich personnel records and the universe of application and hiring decisions at a large multinational firm. The data allow me to analyze employees' labor supply decisions separately from the firm's labor demand decisions. I find that women at lower hierarchy levels are less likely to apply for promotions to first-level leadership positions than observationally similar men, but do not experience lower hiring likelihoods than men. Using detailed information on every internal job opening in employees' choice sets, I show that preferences for leading a team are a key determinant of the gender gap in applications for promotions. These gender differences in preferences for team leadership are not fully explained by other factors, including correlated job features such as flexibility and skill requirements or the gender composition of the coworkers associated with a job opening.

Using the universe of application and hiring decisions at the firm, I distinguish labor supply factors (i.e. application decisions) from labor demand factors (i.e. hiring decisions). I find that women at lower-level positions are substantially less likely to apply for promotions to first leadership levels compared to the base application rate, even when using very detailed controls including employee qualifications and performance. Previous literature has demonstrated the important role of employee characteristics, such as family status and hours worked, for understanding gender differences in labor market outcomes (Bertrand et al., 2010, Goldin, 2014, Kleven, Landais and Søgaard, 2019, Wiswall and Zafar, 2018, Mas and Pallais, 2017). The richness of the firm records allows me to compare application decisions across different groups of workers to account for such factors.

In contrast to applications for promotions, women are not less likely to apply for lateral transitions than men, even if these require switching divisions, functional areas, or locations. This pattern suggests that women in lower-level positions are not generally averse to applying for positions that cause a substantial change in their work environment. In addition, I find no evidence that women are less likely than men to get hired for internal job openings, conditional on applying. This finding underscores the potential importance of gender differences in employees' self-selection, echoing a recent focus in economics on supply-side factors as determinants of gender gaps (Bertrand, 2018).

Why are first-level leadership positions less appealing to women at lower levels than to men? Answering this question is complicated, since leadership responsibility is often correlated with other job attributes, such as hours, pay, and skill requirements that may be less appealing to women. Transitioning to first-level leadership positions often comes with changes in employees' work environment, such as a higher share of male coworkers, which may affect women's likelihood of applying for these positions independently of the required leadership responsibility. Moreover, moving from a lower-level position to a first-level leadership position typically entails multiple changes to job responsibilities. For instance, first-level leadership positions may require employees to both manage a team and assume responsibility over a project. It is therefore difficult to pinpoint which factors drive the application gap.

Responses to a large-scale survey of the firm's employees suggest that responsibility for a team is a particularly salient dimension of leadership that is less appealing to women in lower-level positions. The survey of employees received a 50.0% response rate, yielding over 15,000 responses. When asked where employees would like to see themselves with respect to their career progression, women at low levels are 32% less likely to report preferences for leading a team compared to the baseline mean. This gender gap in stated preferences is not only large in magnitude, it also appears predictive of employees' application behavior. Both men and women who state preferences for leading a team are 30% more likely to report having searched for internal promotions. In contrast, men and women who already hold leadership positions do not differ in their reported preferences for leading a team. These results mirror the finding that gender gaps in promotions and applications are absent once employees begin holding leadership positions.

To test the hypothesis that responsibility for a team is a key determinant of the gender gap in applications, I estimate employees' revealed preferences for leading a team. I leverage a unique feature of my data, which is that I observe all vacancies to which employees could potentially apply. I construct a dataset at the employee-by-vacancy level that allows me to control for detailed job features of both employees' current position as well as those of every job opening in employees' application choice set. Since not all first-level leadership positions involve responsibility for a team, I am able to estimate the extent to which gender differences in revealed preferences for team responsibility explain the gender gap in applications.

The requirement to assume responsibility for a team can explain the entire gap in applications for promotions to first-level leadership positions. Controlling for a broad range of job and employee characteristics, men are 91% more likely to apply for first-level leadership positions if they require leading a team; however, leading a team does not make women more likely to apply. The effect of team leadership is not explained by alternative channels, such as stated job flexibility, how selective the job opening appears, or the female composition of coworkers. Moreover, male and female survey respondents are equally informed about job postings. Together, these findings imply the existence of a gender gap in employees' revealed preferences for leading a team.

This study contributes to three strands of literature in economics. First, a large literature has documented substantial gender differences in labor market outcomes, particularly pay (Goldin, 2014, Blau and Kahn, 2017). Even though gender differences in career progression appear to be an important contributor to disparities in pay (Bronson and Thoursie, 2019), most of the literature has focused on analyzing gender pay gaps, with little consensus on differences in career progression.

Second, by analyzing gender differences in applications as a key driver of the promotion gap, this study adds to a growing body of research on the determinants of gender differences in application behavior. Most previous studies analyze application gaps in relatively narrow settings and have largely focused on individuals' personal traits or on information provided in the application process (Flory, Leibbrandt and List, 2015, Coffman, Collis and Kulkarni, 2022, Gee, 2019, Del Carpio and Guadalupe, 2021, Abraham and Stein, 2020, Delfino, 2021, Cortés, Pan, Pilossoph and Zafar, 2021). The study closest to this paper is Fluchtmann, Glenny, Harmon and Maibom (2021), who analyze gender differences in applications of UI recipients in Denmark with respect to the relative coarse vacancy characteristics, such as industry or and wage. A key innovation of this study is to combine detailed information on employees' full application choice sets with realized application and hiring outcomes, allowing me to isolate the role of key job features, such as team leadership, from other potentially correlated job attributes and employee characteristics.

Third, by documenting the important role of gender differences in preferences for team leadership, this study contributes to literature on gender differences in preferences for job characteristics (Mas and Pallais, 2017, Wiswall and Zafar, 2018, Wasserman, 2022). My results also speak to a related body of work that has explored other explanations for the gender leadership gap, such as employee aspirations (Azmat and Ferrer, 2017), behavioral attributes Eckel and Grossman, 2008, Gneezy, Niederle and Rustichini, 2003, Niederle and Vesterlund, 2007), self-stereotyping (Coffman, 2014), expected backlash (Chakraborty and Serra, 2021), negative beliefs about women's leadership ability (Beaman, Chattopadhyay, Duflo, Pande and Topalova, 2009, Macchiavello, Menzel, Rabbani and Woodruff, 2020), differential recognition for group work (Sarsons, 2017, Sarsons, Gërxhani, Reuben and Schram, 2021), child penalties (Bertrand et al., 2010, Kleven et al., 2019), and differential impacts of managers (Kunze and Miller, 2017, Cullen and Perez-Truglia, 2019, Benson et al., 2021). This study suggests that gender differences are particularly large among employees without prior exposure to leadership responsibility, but that women who transition to the first-level leadership level face similar career outcomes than men, underscoring the importance of focusing on employees' early career progression in order to identify a critical root causes for the persistent gender disparities in labor market outcomes.

The rest of the chapter proceeds as follows. Section 3.2 introduces the setting and Section 3.3 discusses the unique dataset I collect. Section 3.5 documents a large gender gap in applications for promotions among employees at lower-levels. Section 3.5 demonstrates that gender differences in preferences for team leadership are a key driver of this application gap. Section 3.6 concludes.

3.2 Setting

This paper uses rich personnel records from a large multinational firm that employs over 200,000 workers around the world and represents one of the largest manufacturers in Europe. To maintain confidentiality, I refrain from providing details that could be used to identify the firm. As a large manufacturer, the firm's internal labor market consists of over 200 different occupations. The majority of positions are in technical areas, such as engineering or production, which are traditionally male-dominated. However, the firm also employs more female-leaning occupations, such as marketing, finance, and HR, allowing me to assess gender disparities across both male-leaning and female-leaning areas.

Since the goal of this study is to analyze career progression to higher-level positions, I restrict my analysis to white-collar and management employees at the firm (i.e. employees that are either already in or could ultimately attain management positions). While the firm operates in many different countries, including the United States, Germany represents the largest internal labor market at the firm. I therefore focus my analysis on all 30,000 white-collar and management employees who are based in Germany.

Table 3.1 provides summary statistics for my main analysis sample which consists of over 400,000 employee-by-quarter observations from 2015 to 2019. Women represent 21% of employees in the sample, which is consistent with the underrepresentation of women in technical occupations. Employee tenures at the firm tend to be long, with an average tenure of 13 years, allowing me to follow employees' internal career progression over time. Because I restrict to white-collar and management employees with regular employment contracts (as opposed to those with marginal employment such as mini jobs), employee qualifications in my sample is high. The average employee holds a Bachelor's degree and 92% of employees work full-time.

The demographics of the employees at the firm are comparable to other large manufacturing firms in Germany. In Appendix Table A.1, I compare employees in my sample to those employees in large manufacturing firms in the BiBB, a representative survey of the German workforce conducted in 2018. I find very similar patterns with respect to most employee characteristics (e.g. gender, age, German citizenship, martial and family status). In addition, the BiBB illustrates that the gender leadership gap in the firm I study aligns with broader patterns of female underrepresentation in Germany, suggesting that this setting is fairly typical for German firms.

Like many other large organizations, the firm requires employees to actively apply using the centralized online job portal in order to make internal job transitions, including promotions. Such active application systems are very common among large organizations (hkp, 2021). Employees can access every job opening at the firm through a centralized online job portal (see Appendix Figure A.1 for an illustration of the job portal). Employees are required to submit their application through the online portal, which typically takes less than five minutes to complete. This institutional feature enables me to directly measure employees' labor supply, in terms of internal applications, and the firm's labor demand, in form of final callback and hiring decisions.

Quarterly application rates for internal positions are 3%. While employees can choose to apply to multiple positions at the same time, the median applicant applies to only one internal position in a given quarter. Only 25% of applications are successful, with lower success rates for applications to higher-level positions. Similar to the external labor market, 93% of applicants have not previously worked with the hiring manager of the position they are applying for. In addition, the firm operates in over 50 cities in 250 establishments throughout Germany and one-third of internal applications are for positions in a different city. Consequently, the internal labor market is both spatially and interpersonally diffuse and application decisions are typically made under uncertainty.

The usual first step towards become a senior leader is to take on a first-level leadership position that typically involves leadership responsibility in the form of managing a team or being responsible for a project or product. These positions are distinct from lower-level positions as individual contributors that do not involve any leadership responsibility. Job openings usually indicate the type of leadership responsibility a position entails (e.g. number of direct reports, type of project), making leadership responsibility a very salient job characteristic at the time that application decisions are made. Focusing on first-level leadership positions has several advantages. Because the majority of senior leadership positions, which also involve greater pay, require previous leadership experience, first-level leadership positions represent an important prerequisite for employees who want to climb the job ladder.

3.3 Novel Data on Internal Career Progression

I assemble a unique dataset that combines personnel records and job application data at the firm, allowing me to isolate the role of application differences for the gender leadership gap.

3.3.1 Personnel Records and Job Application Data

I collect the firm's internal personnel records from 1998 to 2020, which provide detailed information on demographics and position characteristics for all employees in my sample. The richness of these data allow me to account for key differences between men and women that may influence workers' career progression independent of leadership responsibility. I collect detailed demographic information from the personnel records, including gender, age, citizenship, educational qualifications (such as highest degree, major, and institution), marital status, family status, and parental leave history at the firm. The records also contain detailed position characteristics, such as occupation and position title, functional area (e.g. marketing versus engineering), business unit, location, leadership responsibility, and the reporting distance to the CEO. I supplement these data with payroll information, capturing employees working hours, earnings, and bonus payments. Finally, I collect information on worker evaluations, such as performance and potential ratings.

To isolate the role of labor supply for the gender leadership gap, I collect data on the universe of application and hiring decisions from 2015 to 2020 at the firm. Because the firm requires all application and hiring decisions to be submitted through a centralized online portal, I am able to separately measure labor supply factors (i.e. applications) and labor demand factors (i.e. interview and hiring decisions). I observe the exact timing and identity of each application at the firm, covering both applications from existing employees and from external applicants. In total, the application data cover over 16,000 job openings and over 200,000 external and internal applicants. Because I also observe the outcome of each application in terms of rejections, interview callbacks, and subsequent hiring outcomes, I am able to construct a panel dataset of employees application and hiring histories at the firm from 2015 to 2020. The data allow me to disentangle whether an employee did not switch jobs because they did not apply or because they did not get hired conditional on applying.

Since the firm requires all job openings to be posted to a centralized job portal, I am able to collect the original job posting for the universe of job openings at the firm from 2015 to 2020. Extracting job characteristics from job postings has several key advantages. First, the job postings capture relevant job features that are usually not contained in personnel records, including possible arrangements with respect to job flexibility or job tasks such as frequent negotiations. Second, I am able to control for those job features that are salient to applicants at the time of application, which aids the interpretation of observed application decisions as revealed preferences. Third, this approach also captures how jobs are described, for instance whether language is used that sounds particularly competitive.

From the raw job posting, I extract advertised position characteristics, including position title, business unit, location, pay, part-time options, hours, and responsibility for a team. I also observe the job requirements that are stated in the posting, such as frequent business travel, educational qualification, work experience, and communication skills. For validity exercises, I also characterize job postings by the type of language they include (e.g. male-leaning vs female-leaning). Table 3.2 provides an overview about the type of job characteristics that I extract from the job postings.

To measure characteristics of job openings beyond those that are explicitly stated in the job description (e.g. how challenging a job seems), I draw on characteristics of the applicant pool to each opening. I use the number of external applicants (who were not employed by the firm at the time of application) to measure the relative attractiveness of a job opening conditional on the job's other features. I also use information about other applicants' quality, such as their educational qualifications, to characterize job openings.

3.3.2 Data Linkages and Sample Construction

For my empirical analysis, I construct two primary analysis samples.

First, to test the role of applications for the gender leadership gap, I construct an employee-by-quarter dataset spanning 2015 to 2019. I restrict my sample to only white-collar and management employees who are regular employees at the firm (e.g. excluding marginal employment such as mini jobs). I combine the personnel records with the job application data using a five-step matching algorithm, which matches over 90% of individuals (see Appendix Section B.0.1 for more details). I then collapse the data to a quarterly level. My main analysis sample contains over 400,000 employee by quarter observations and covers over 30,000 unique white-collar and management employees.

Second, I create an employee-by-vacancy dataset from 2015 to 2019 that combines each employee in my main analysis sample with every available job opening they could have applied to. I refine these choices based on observed application patterns, dropping combinations that never occur in the data.[1] To test whether men and women at lower-levels differ in their preferences for team leadership,I restrict my analysis to employees who are at lower-level positions, yielding over 2 million employee-by-vacancy observations, each corresponding to a potential application choice. In a given quarter, the average employee has 35 job openings in their final application choice set. Because I both observe employees' complete choice sets and their realized choices as well as detailed job features of each opening, I am able to estimate employees' revealed preferences for job features.

3.3.3 Employee Survey

To capture employees' perceptions with respect to career progression in general and leadership responsibility in particular, I designed and conducted a large-scale survey at the firm. All employees in my main analysis sample were invited via e-mail by the firm's human resources department and were asked to provide their perspectives on the firm's internal labor market. The survey received over 15,000 responses, yielding a 50.0% response rate. Respondents are similar to non-respondents in terms of demographics (Appendix Table B.1). I find no evidence for differential selection into response by gender (Appendix Table B.2). Employees described challenges regarding their internal career progression both in the form

[1]For instance, I drop combinations between employee location and vacancy location for which applications never occur in my data.

of free-text responses and in multiple-choice answers. The median response time was 13 minutes. For my main analysis, I only keep respondents who took at least five minutes to respond and have no missing observations.

3.4 Gender Differences in Applications

Many firms rely on employees to actively apply for internal job switches in order to allocate talent (hkp, 2021). An emerging consensus in economics, however, underscores that supply-side factors are an important channel through which gender differences in labor market outcomes manifest (Bertrand, 2018, Coffman, 2014), suggesting that active application systems may exacerbate gender differences in the workplace. An advantage of this study is that key labor-supply factors (e.g. application decisions for internal job transitions) are observable independent of labor-demand factors (e.g. callback and hiring decisions). By leveraging the detailed information on every job openings' characteristics, I am also able to analyze different types of applications (e.g. for lateral transitions vs promotions).

I begin by analyzing gender differences in the quarterly likelihood to apply for an internal job opening at the firm.[2] Specifically, I estimate gender differences in applications using a logit regression of an indicator for applying in a given quarter on gender and quarter fixed effects.

$$\Pr(\text{Applied}_{it} = 1) = \Lambda(\theta_1 \text{Female}_i + \theta_X X_{it} + \theta_t)$$

Prior research has highlighted the importance of factors such as employees' family status and hours requirements for gender differences in labor market outcomes. In addition, because of common workplace segregation, female employees may experience different application opportunities than male employees who work in different areas at the firm. To account for such factors, I include a broad set of controls X_{it}, which capture worker demographics (age, German citizenship, educational qualification, marital status, family status, parental leave, firm tenure), position characteristics (position type, division, functional area, location, full-time, hours, number of direct reports), employee evaluations (performance and potential ratings), and past mobility at the firm.

I find that women at low hierarchy levels are substantially less likely to apply for promotions to first-level leadership positions than men, but no gender differences exist with respect to applications for lateral transitions. Column 1 of Table 3.3, Panel A shows that the gender gap in applications for all transition types is 6.8%. Column 2 indicates that there is no gender gap in applications for lateral transitions. Women are 11.1% less likely to apply for small promotions (Column 3) and 26.0% less likely to apply for major promotions (Column 4) relative to base application rates.

These gender gaps in applications for promotions are large in magnitude and robust across different specifications. The large magnitudes are striking given that these estimates

[2] This approach is motivated by the fact that the median applicant in my sample applies to only one job opening in a given quarter.

control for employees' current position characteristics and detailed qualification measures including past performance and potential ratings. My findings are not driven by specific sets of controls and are not fully explained by factors commonly cited in relation to gender-specific preferences, such as parenthood (Bertrand et al. (2010), Kleven et al. (2019)) or flexibility in hours (Goldin (2014), Wiswall and Zafar (2018), Mas and Pallais (2017)). Gender gaps exist even among employees who work full-time or do not have children (Panel A of Table 3.5). Results are similar when estimating the gender gap only within the set of workers who applied for at least one job within a given quarter, suggesting that differences in access to the job portal do not account for the observed differences (Panel B of Table 3.5).

Gender gaps are absent for lateral transitions, even those that require substantial changes in an employee's work environment. Panel B of Table 3.3 indicates no significant gender differences in applications for lateral transition that require a switch to a different division (Column 1), functional area (Column 2), or location less than 100 kilometers away (Column 3). Although women are less likely to apply for lateral transitions that are more than 100 kilometers away (Column 4), collectively these findings suggest that women are not less likely to apply for promotions because they entail changes in work environments. In addition, since applications at the firm are competitive and hiring likelihoods are small even for lateral applications, the absence of any gender gap in applications for lateral transitions suggests that factors such as fear of rejection or preferences for risk and competition are not likely to be the sole drivers of the observed gender gap in applications for promotions.

While women at lower levels are less likely to apply for promotions, they are not less likely to get hired conditional on applying. Panel A of Table 3.6 presents results from a logit regression of an indicator for getting hired on gender and the controls used in previous specifications. Women who apply for major promotions are 35% more likely to land the position. For women at first-level leadership positions, the hiring advantage conditional on applying is even larger (Panel B of Table 3.6). Even though women are likely positively selected into leadership positions, complicating the direct comparison of application and hiring margins, these results suggest that the hiring stage is not a key bottleneck for women's career progression. This finding is supported by evidence from Chapter 1 that finds that female marginal applicants have twice as large marginal hiring probabilities than their male counterparts when using an instrumental variable approach that circumvents potential selection bias.

For completeness, I also analyze gender differences in application and hiring decisions at higher hierarchy levels. While women at lower-levels are substantially less likely to apply for promotions, Table 3.4 demonstrates that women at first-level leadership positions are not less likely than similar men to apply for promotions. These results indicate that identifying why women are less likely to sort into leadership positions than men is critical.

3.5 Gender Differences in Preferences for Team Leadership

Why are women at lower-levels less likely to apply for promotions but not for lateral transitions? Responses to a survey of the firm's employees suggest that responsibility for a team is a particularly salient dimension of higher-level positions that is less appealing to women in lower-level positions. Revealed preference estimation using data on employees' full application choice sets underscores this finding, demonstrating that gender differences in preferences for team leadership account for the observed application gap.

3.5.1 Survey Evidence on Reported Preferences

The employee survey reveals large gender differences in preferences for leading a team among employees without leadership responsibility. Figure 3.1 presents answers to a question asking where employees would like to see themselves in five years with respect to their career. For men at low levels (Panel A), the most frequent career aspiration is to have responsibility for a team (26%), followed by more pay (21%), working in a similar position as the employee's current position (19%), more challenges (16%), more job security (11%), and more flexibility (7%). Women, however, are 32% less likely to report preferences for leading a team compared to the baseline mean. There do not exist similarly large gender gaps with respect to preferences for more pay, a similar position, or a more challenging position.

In contrast, among employees who already hold leadership positions, there is no gender gap in preferences for more team leadership responsibilities. Panel B of Figure 3.1 documents very similar response patterns for men and women who hold leadership positions. The absence of a gender gap in preferences for team leadership at higher levels mirrors the results in Section 3.4 finding no differences in applications for promotions among employees who have attained leadership positions.

Differences in preferences for team leadership likely translate into differences in applications. For all employees, reported preferences for team leadership are highly predictive of job search. Both men and women who report wanting responsibility for a team are 30% more likely to also report to search for an internal promotion. The importance of team leadership is supported by qualitative responses to the survey. Responsibility for a team is described in open-ended responses by many employees as a particularly salient feature of leadership positions. When respondents are prompted to describe their reasoning when making internal application decisions, a common response is that employees assess whether a position involves responsibility for a team.

3.5.2 Gender Differences in Revealed Preferences

Empirically detecting why women are less likely to apply for first-level leadership positions is difficult, because higher-level positions can differ from lower-level positions in many ways, such as pay, hours, and responsibility for a team. Transitioning to first-level leadership positions may also induce changes in employees' work environment, such as a higher share of male coworkers, which may affect women's likelihood of applying for these positions independently of the required leadership responsibility. To isolate the role of responsibility over a team for explaining the observed application gap, I use detailed information on the trade-offs that employees make between their application choices, allowing me to infer their revealed preferences for team leadership while controlling for other relevant job features.

Two unique features of my data make this revealed-preference approach possible. First, because the data contain all vacancies from 2015 to 2020 to which employees could potentially apply as well as employees' realized application choices, I am able to infer preferences over positions. The revealed preference logic suggests that employees, who choose to apply for a position even though alternative positions are available at the same time, prefer this position over the alternatives. Second, my data contain detailed information on each vacancy's job features, which may be correlated with responsibility for a team (e.g. pay, required work experience). To isolate preferences for team leadership, I leverage the fact that not all first-level leadership positions require the responsibility for a team.[3]

The revealed preference interpretation assumes that employees are informed about application choices. This assumption is supported by the fact that the firm's internal policy requires every job opening to be posted to a centralized online job portal that is accessible to all employees. Several pieces of evidence from the employee survey further support the assumption that women and men have similar access to information. First, men and women are equally likely to actively search for job openings (Column 1 of Table 3.7, Panel A) and to get approached by others with recommendations about job openings (Column 2 of Table 3.7, Panel A). Second, the type of individual who makes these recommendations (e.g. supervisor vs. teammate) is similar for men and women (Panel A of Figure 3.2). Third, when asked what would be most supportive for their career progression, only 4.5% of male and 4.5% of female respondents name better access to information about job openings.

To characterize employee's choice sets, I construct a dataset at the employee-by-vacancy level that combines each employee with every available job opening they could have applied to. I refine these choices based on observed application patterns, dropping combinations that never occur in the data. I restrict my analysis to employees who are at lower-level positions, yielding over 2 million employee-by-vacancy observations, each corresponding to a potential application choice. In a given quarter, the average employee has 35 job openings in their final application choice set. For each vacancy I observe detailed information about the job features advertised in the job ad, such as position title, division, functional area, location, type of leadership responsibility, required degree and work experience, and pay. I also observe detailed information on employees' current positions and on their demographics. In addition, the data contain an indicator for employees' realized application choices.

[3]While 79% of first-level leadership positions require team leadership, the remainder are characterized by leadership responsibilities with respect to a project. In my data, such variation in the type of leadership responsibilities exists even within functional areas and occupation groups.

I begin by establishing that the gender gap in applications for first-level leadership positions documented in Section 3.4 persists when controlling for detailed features of each vacancy to which employees could have applied. I estimate a linear probability model for the likelihood that employee i chooses to apply to first-level leadership job j in quarter t, controlling for employees' current job features and demographics X_{it} and detailed features of the job opening W_j:

$$\Pr(\text{Applied}_{ijt} = 1) = \gamma_1 \text{Female}_i + \gamma_X X_{it} + \gamma_w W_j + \gamma_t$$

I cluster standard errors at the employee level and scale the outcome by 100.

The results from this linear probability model are presented in Table 3.8. In my preferred specification (Column 3), women are 38% less likely to apply for a first-level leadership position in their application choice set relative to the outcome mean. This finding suggests that the gender gap in applications for first-level leadership positions documented in Section 3.4 is not fully explained by the composition of employees' application choice sets or by other job features that are correlated with first-level leadership positions (e.g. higher pay, lower hours). Instead, stark gender differences in revealed preferences for first-level leadership positions appear to exist.

To isolate the importance of team leadership as suggested by the responses from the employee survey, I estimate the linear probability including an indicator for whether a job opening requires team leadership and its interaction with gender.

$$\begin{aligned}\Pr(\text{Applied}_{ijt} = 1) = {} & \gamma_1 \text{Female}_i + \gamma_2 \text{Leading team}_j + \gamma_3 (\text{Female}_i \times \text{Leading team}_j) \\ & + \gamma_X X_{it} + \gamma_w W_j + \gamma_t\end{aligned}$$

Table 3.9 shows that the requirement to assume responsibility for a team explains the entire application gap. Controlling for a broad range of job and employee characteristics, men are 91% more likely to apply for leadership positions if they require leading a team; however, leading a team does not make women more likely to apply.

3.5.3 Alternative Channels

While these results suggest that preferences for team leadership are key drivers of gender gaps in applications, it remains possible that differences in applications are explained by other attributes that result from organizations' choice of job architecture. Positions that require responsibility for a team may differ in many respects from those that do not, or might be advertised differently. For instance, team leadership positions might offer less flexibility. If men and women differ in their preferences for job flexibility, only controlling for flexibility will not absorb this differential effect. Positions with team leadership could also differ in the work environment they are in (e.g. the share of female colleagues or supervisors is likely lower since most team leaders are male).

I test whether such alternative channels can explain why positions with team leadership are less appealing to women by explicitly controlling for potential gender differences in preferences for alternative features. Each coefficient in Table 3.10 stems from a separate regression in which I include an interaction of the alternative feature and worker gender. The objective of this exercise is to test whether adding these interactions substantially reduces the coefficient of interest on the interaction of gender and team leadership.

Panel A of Table 3.10 shows that controlling for gender differences with respect to other key job features does not explain the team leadership gap. Column 1 controls for potential gender differences with respect to job flexibility in terms of full-time status, number of weekly hours, and the requirement for frequent business travel. Columns 2 and 3 control for stated job requirements with respect to on-the-job negotiation and previous work experience, respectively. Neither regression substantially alters the coefficient on gender and team leadership. Panel B of Table 3.10 uses characteristics from the applicant pool to measure characteristics of the position that are not explicitly stated in the job ad. Intuitively, the number and quality of applicants may contain information on how selective a job opening may appear to prospective applicants. Controlling for the total number of applicants (Column 1), a high share of applicants with a graduate degree (Column 2), and applicants with high prior performance ratings (Column 3) does not substantially alter the estimated gender gap in team leadership. Panel C of Table 3.10 finds that even after controlling for the gender of the direct supervisor (Column 1), gender representation among coworkers (Column 2), and gender representation among the broader organizational unit (Column 3), the effect of team leadership persists.

Taken together, these results suggest that preferences for leading a team are an important driver of gender differences in applications, and that the observed relationship between application behavior and leadership responsibility is not explained by other attributes of the position or the work environment that are likely correlated with team leadership.

3.6 Conclusion

Understanding the root causes of female underrepresentation is critical for identifying effective policy remedies. This chapter shows gender differences in applications for first-level leadership positions are an important driver of female underrepresentation. The gender application gap arises due to gender differences in preferences for team leadership, highlighting the important role played by organizational job architecture.

The large gender gap in preferences for team leadership raises the question, what is driving these differences in preferences? The results in this study suggest that employee demographics and other features of leadership positions alone cannot fully explain the gender differences in preferences for team leadership. In ongoing work, a follow-up survey with employees at the firm aims to elicit whether gender differences in the perception of leadership responsibility represent a relevant underlying mechanism. Identifying why first-level leadership positions are less appealing to (qualified) women can thus help improve gender representation and talent allocation in organizations.

Figure 3.1: Reported Preferences for Team Leadership

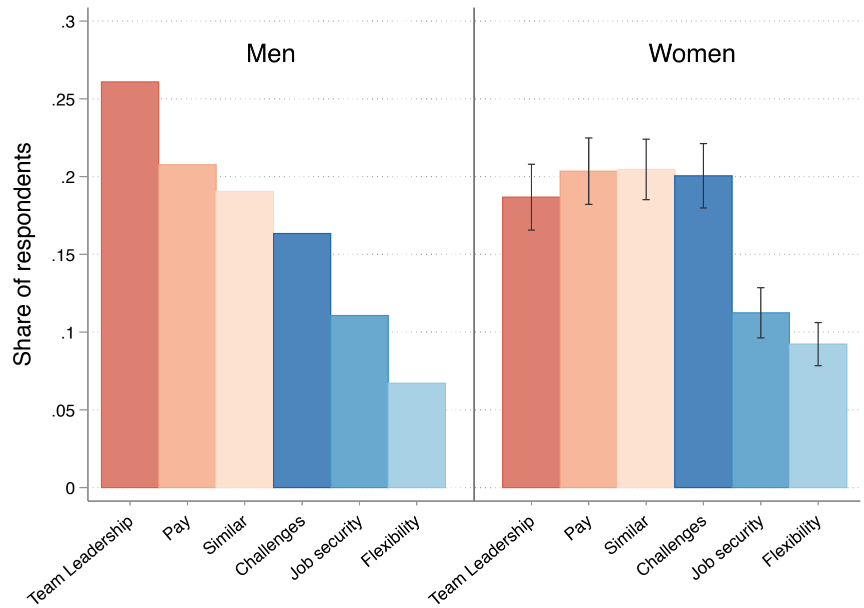

Panel A. Employees w/o Leadership Responsibility

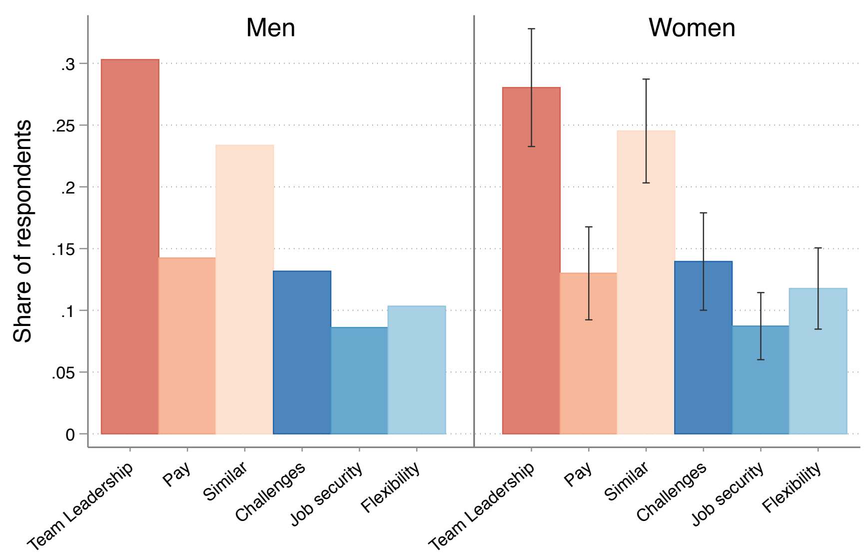

Panel B. Employees with Leadership Responsibility

Notes: This figure documents responses from the employee survey. Respondents were asked to choose one out of following answers to the question where they would like to see themselves with respect to their career in five years: more team leadership responsibility, more pay, staying in a similar job, more challenges, more job security, more job flexibility. Panel A restricts to employees at low levels without leadership responsibility (N=1x,xxx). Panel B restricts to employees in leadership positions (N=4,xxx). Controls: Age, tenure, children, nationality, full time, hours, location, functional area, and job switch in past 12 months.

Figure 3.2: Type of Individual who Provided Job Recommendations

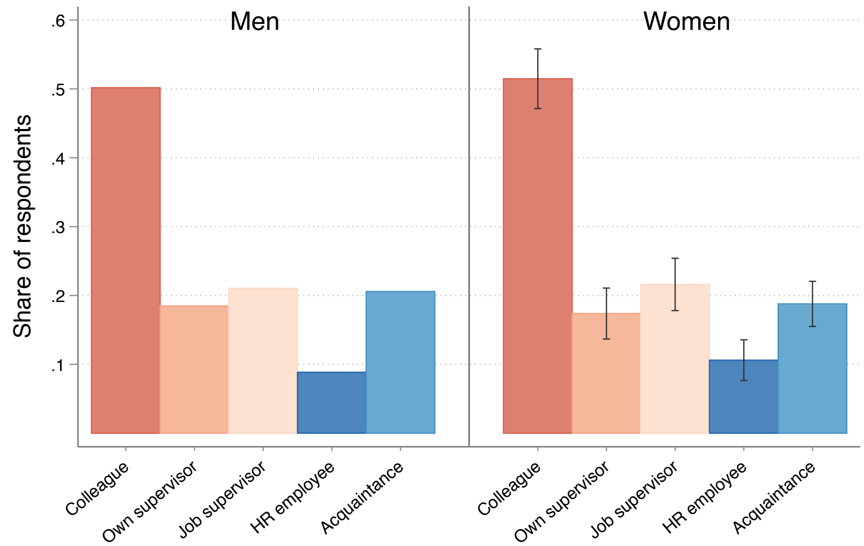

Panel A. Employees w/o Leadership Responsibility

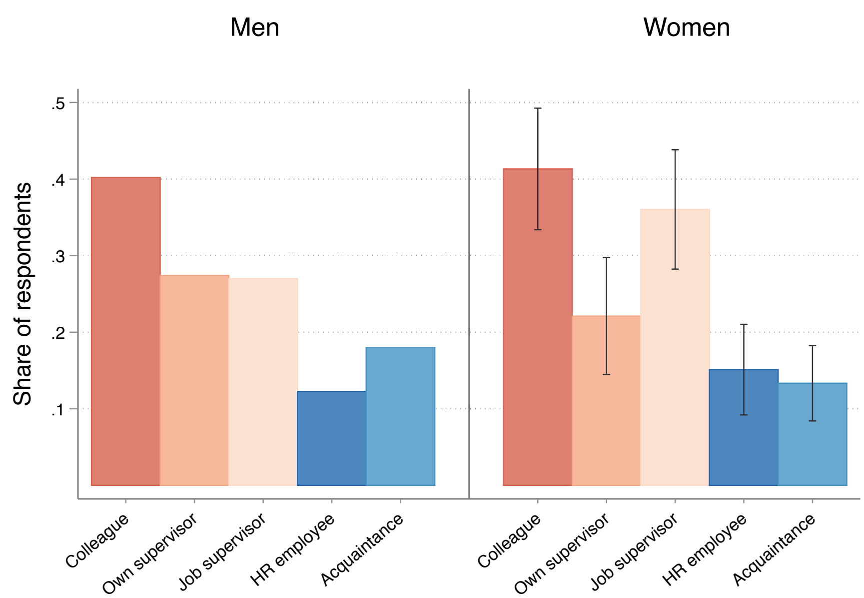

Panel B. Employees with Leadership Responsibility

Notes: This figure documents responses from the employee survey. Respondents who report having received recommendations about job opportunities were asked to identify which type of individual had approached them in the past 12 months. Panel A restricts to employees at low levels without leadership responsibility (N=1x,xxx). Panel B restricts to employees in leadership positions (N=4,xxx). Controls: Age, tenure, children, nationality, full time, hours, location, functional area, and job switch in past 12 months.

Table 3.1: Summary Statistics

	Men (1)	Women (2)	Δ (3)
Demographics			
German nationality	0.90	0.86	0.04***
Age (years)	44.2	40.5	3.7***
Schooling (years)	16.05	14.91	1.14***
Firm tenure (years)	13.6	12.2	1.4***
Married	0.65	0.51	0.14***
Children	0.77	0.65	0.12***
On parental leave	0.02	0.08	-0.06***
Position			
Limited contract	0.01	0.03	-0.02***
Full-time	0.97	0.73	0.24***
Weekly hours	42.0	37.8	4.3***
Technical position	0.71	0.33	0.39***
Salary decile	5.9	3.8	2.2***
Team leadership	0.21	0.10	0.11***
Direct reports	6.09	5.10	0.99***
Cumulative reports	25.36	9.01	16.36***
Top management	0.08	0.04	0.04***
Number of teammates	9.1	8.6	0.5***
Evaluations			
High performance	0.56	0.47	0.10***
High potential	0.27	0.26	0.02***
Observations	3xx,xxx	1xx,xxx	4xx,xxx

Notes: This table reports summary statistics for my quarterly analysis sample from 2015 to 2019. Column 1 restricts to male employees, Column 2 restricts to female employees, and Column 3 represents the differences between men and women. Technical positions refer to positions in engineering, IT, or production-related areas. Salary deciles are computed based on the distribution of employees' real annual earnings. Top management positions refer to senior executives, board members, and the CEO. Performance and potential ratings are employee evaluations conducted by the direct supervisor. *p< 0.10, **p< 0.05, ***p< 0.001.

Table 3.2: Characteristics of Job Posting by Leadership Level

	Low-level position (1)	First-level leadership (2)	Δ (3)
Attributes			
Salary (decile)	3.62	5.86	-2.24***
Technical role	0.68	0.49	0.19***
Strategic position	0.05	0.16	-0.11***
Part-time possible	0.10	0.10	0.00
Weekly hours	41.21	41.81	-0.59***
Based at Top 5 location	0.68	0.73	-0.05***
Frequent business travel	0.23	0.35	-0.12***
Frequent negotiations	0.06	0.13	-0.07***
Requirements			
At least Bachelor's degree	0.61	0.72	-0.11***
Work experience			
Few years	0.10	0.06	0.04***
Several years	0.46	0.62	-0.16***
Many years	0.03	0.06	-0.03***
English proficiency	0.85	0.92	-0.07***
Communication skills	0.52	0.61	-0.10***
Analytical skills	0.19	0.22	-0.03***
Length of Ad			
Description of position (decile)	4.84	5.69	-0.84***
Stated requirements (decile)	4.92	5.62	-0.70***
Observations	5,xxx	6,xxx	11,xxx

Notes: This table illustrates how positions are advertised in all job postings at the firm from 2015 to 2019 that received at least one application from an employee in my sample. Column 1 restricts to low-level positions without leadership responsibility, while Column 2 focuses on first-level leadership positions. Column 3 reports the difference between Column 1 and 2. Salary is expressed in terms of the salary decile at the firm. Technical roles refer to positions in engineering, IT, and production-related areas. Strategic positions are jobs which are described as strategic positions in the job ad (i.e. of high importance). Being based at a top 5 locations refers to the five biggest locations of the firm. The length of the job ad is measured as the length of the string corresponding to the description of the position and the listings of requirements and is measured in deciles with respect to all job openings at the firm. *p < 0.10, **p< 0.05, ***p< 0.001.

Table 3.3: Application Likelihood for Employees at Low Levels

Panel A: Applications by Hierarchy Increase of Transition

	Any Application (1)	Lateral Switch (2)	Small Promotion (3)	Major Promotion (4)
Female	-0.061**	0.049	-0.084*	-0.284***
	(0.030)	(0.054)	(0.043)	(0.076)
Outcome Mean	0.0283	0.0080	0.0121	0.0055
Av ME for Women	-0.0019	0.0004	-0.0013	-0.0014
Gender Gap in %	-6.8	5.5	-11.1	-26.0
Observations	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx

Panel B: Applications by Destination for Lateral Transitions

	Division Different (1)	Function Different (2)	Location <100 km (3)	Location >=100 km (4)
Female	-0.028	-0.027	0.071	-0.201***
	(0.042)	(0.047)	(0.080)	(0.074)
Outcome Mean	0.0132	0.0104	0.0042	0.0047
Av ME for Women	-0.0004	-0.0003	0.0003	-0.0011
Gender Gap in %	-3.4	-3.3	6.6	-23.2
Observations	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx

Notes: This table documents gender differences in applications between 2015 and 2019 for employees at low-level positions. Panel A distinguishes transitions by their induced increase in hierarchy. Major promotions represent transitions to first-level leadership positions. Panel B focuses on lateral switches and distinguishes transition by their destination. Gender gaps in % are computed by dividing the average marginal effect for women based on the logit coefficient by the outcome mean. Each coefficient stems from a separate logit regression of an indicator for applying for a given transition type on gender and a large set of controls. Controls: Age, German citizenship, educational qualification, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, number of direct reports, performance and potential rating, time on position, and quarter fixed effects. Standard errors in parentheses. *p< 0.10, **p< 0.05, ***p< 0.001.

Table 3.4: Application Likelihood for Employees at First Leadership Level

Panel A: Applications by Hierarchy Increase

	Any Application (1)	Lateral Switch (2)	Small Promotion (3)	Major Promotion (4)
Female	0.032	-0.140	0.177	0.081
	(0.084)	(0.120)	(0.145)	(0.192)
Outcome Mean	0.0285	0.0154	0.0100	0.0049
Av ME for Women	0.0011	-0.0023	0.0020	0.0005
Gender Gap in %	3.8	-15.0	20.0	11.2
Observations	5x,xxx	5x,xxx	5x,xxx	5x,xxx

Panel B: Applications by Destination for Lateral Switches only

	Division Different (1)	Function Different (2)	Location <100 km (3)	Location >=100 km (4)
Female	0.004	0.018	-0.155	-0.095
	(0.109)	(0.115)	(0.241)	(0.176)
Outcome Mean	0.0168	0.0159	0.0052	0.0069
Av ME for Women	0.0001	0.0003	-0.0006	-0.0007
Gender Gap in %	0.4	2.1	-11.7	-9.8
Observations	5x,xxx	5x,xxx	5x,xxx	5x,xxx

Notes: This table documents gender differences in applications from 2015 to 2019 for employees at first-level leadership positions. Panel A distinguishes transitions by the the increase in hierarchy the transition induces. Major promotions represent transitions to higher-level leadership positions. Panel B focuses on lateral switches and distinguishes transition by destination type. Gender gaps in % are computed by dividing the average marginal effect for women based on the logit coefficient by the outcome mean. Each coefficient stems from a separate logit regression of an indicator for applying for a given transition type on worker gender and a large set of controls.Controls: Age, degree, nationality, marital and family status, parental leave, job title, division, functional area, location, hours, reports, performance and potential rating, time on position, and quarters. Standard errors in parentheses. *p< 0.10, **p< 0.05, ***p< 0.001.

Table 3.5: Robustness in Applications of Employees at Low Levels

Panel A: Applications for Major Promotions				
	Full-time (1)	Part-time (2)	Children (3)	No Children (4)
Female	-0.098**	-0.131	-0.175***	-0.128**
	(0.044)	(0.212)	(0.057)	(0.064)
Outcome Mean	0.0124	0.0089	0.0124	0.0089
Av ME for Women	-0.0018	-0.0014	-0.0025	-0.0026
Gender Gap in %	-14.6	-15.3	-20.4	-28.9
Pseudo R-squared	3xx,xxx	2x,xxx	2xx,xxx	9x,xxx

Panel B: Application Gaps if Applied				
	Any Application (1)	Lateral Switch (2)	Small Promotion (3)	Major Promotion (4)
Female	0.178***	-0.041	-0.304***	-0.183***
	(0.066)	(0.061)	(0.082)	(0.060)
Outcome Mean	0.2840	0.4267	0.1939	0.6221
Av ME for Women	0.0330	-0.0092	-0.0372	-0.0417
Gender Gap in %	11.6	-2.1	-19.2	-6.7
Pseudo R-squared	9,xxx	9,xxx	9,xxx	9,xxx

Notes: This table documents robustness with respect to gender differences in applications from 2015 to 2019 for employees at low-level positions. Panel A estimates gender gaps in applications for promotions separately for employees who work full-time vs. part-time (Columns 1 and 2) and employees who have children vs. do not have children (Columns 3 and 4). Panel B restricts to employees who applied at least once in a given quarter. Controls: Age, German citizenship, educational qualification, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, number of direct reports, performance and potential rating, time on position, and quarter fixed effects. Standard errors in parentheses. *p< 0.10, **p< 0.05, ***p< 0.001.

Table 3.6: Hiring Likelihood for Employees

Panel A: Employees at Low Levels				
	Any Application (1)	Lateral Switch (2)	Small Promotion (3)	Major Promotion (4)
Female	0.190***	0.098	0.261**	0.421**
	(0.066)	(0.126)	(0.102)	(0.184)
Outcome Mean	0.2690	0.3439	0.2138	0.1565
Av ME for Women	0.0353	0.0194	0.0439	0.0562
Gender Gap in %	13.1	5.6	20.5	35.9
Observations	9,xxx	2,xxx	4,xxx	1,xxx
Panel B: Employees at First Leadership Level				
	Any Application (1)	Lateral Switch (2)	Small Promotion (3)	Major Promotion (4)
Female	0.624***	0.653**	1.412***	1.879
	(0.183)	(0.275)	(0.380)	(1.259)
Outcome Mean	0.2748	0.3352	0.1890	0.1625
Av ME for Women	0.1255	0.1327	0.2099	0.2170
Gender Gap in %	45.7	39.6	111.1	133.5
Observations	1,xxx	1,xxx	1,xxx	1,xxx

Notes: This table documents gender differences in hiring likelihoods for employees who applied for an internal position from 2015 to 2019. Panel A restricts to employees at low-level positions. Panel B restricts to employees in first-level leadership positions. Each coefficient stems from a separate logit regression of an indicator for getting hired to a given position type on worker gender. Controls: Age, German citizenship, educational qualification, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, number of direct reports, performance and potential rating, time on position, and quarter fixed effects.Standard errors in parentheses. *p< 0.10, **p< 0.05, ***p< 0.001.

Table 3.7: Access to Information about Job Openings

Panel A: Employees w/o Leadership Responsibility

	Searched actively	Received recommendation
	(1)	(2)
Female	-0.009	0.002
	(0.0127)	(0.0125)
Outcome mean	0.407	0.326
Gender Gap in %	-2	1
Observations	10,xxx	10,xxx

Panel B: Employees with Leadership Responsibility

	Searched actively	Received recommendation
	(1)	(2)
Female	-0.013	0.046*
	(0.0258)	(0.0261)
Outcome mean	0.379	0.403
Gender Gap in %	-3	11
Observations	4,xxx	4,xxx

b

Notes: This table documents responses from the employee survey. Respondents were asked whether they actively searched for internal job openings in the past 12 months (Column 1) and if someone within the firm has approached them with information or recommendations about job openings in the past 12 months (Column 2). Panel A restricts to employees without leadership responsibility. Panel B restricts to employees with leadership responsibility. Controls: Age, tenure, kids, nationality, full-time, hours, location, functional area, and job switch in past 12 months. Standard errors in parentheses. *p $<$ 0.10, **p$<$ 0.05, ***p$<$ 0.001.

Table 3.8: Application Likelihood for First Leadership Level

	Applied for Job at First Leadership Level		
	(1)	(2)	(3)
Female	-0.0050**	-0.0051**	-0.0044*
	(0.0020)	(0.0026)	(0.0026)
Employee characteristics	-	X	X
Current job features	-	X	X
Opening job features	-	-	X
Outcome Mean	0.0117	0.0117	0.0117
Observations	2,xxx,xxx	2,xxx,xxx	2,xxx,xxx

Notes: This table documents gender differences in applications from 2015 to 2019 while controlling for employees' application choice sets. I restrict to employees at low-level positions without leadership responsibility. Each coefficient stems from a separate linear probability model where the outcome of interest is an indicator for applying to a first-level leadership position in employees' application choice set, scaled by 100. Column 1 provides the raw estimates. Column 2 adds employee demographics and characteristics of employees' current positions. Column 3 adds detailed characteristics of the job vacancy. Employee controls: Age, tenure, martial and family status, German nationality, full-time, hours, location, division, functional area, and job switch in past 12 months. Vacancy controls: location, division, functional area, job requirements stated in job ad, hours, salary, negotiation on the job. Standard errors are clustered at the employee level. *p< 0.10, **p< 0.05, ***p< 0.001.

Table 3.9: Impact of Team Leadership on Application Likelihood for First Leadership Level

	Applied for Job at First Leadership Level		
	(1)	(2)	(3)
Female	0.0003	0.0006	0.0009
	(0.0018)	(0.0028)	(0.0028)
Leading team	0.0161***	0.0170***	0.0106***
	(0.0027)	(0.0030)	(0.0026)
Female × Leading team	-0.0114***	-0.0129***	-0.0122***
	(0.0041)	(0.0045)	(0.0047)
Employee characteristics	-	X	X
Current job features	-	X	X
Opening job features	-	-	X
Outcome Mean	0.0117	0.0117	0.0117
Observations	2,xxx,xxx	2,xxx,xxx	2,xxx,xxx

Notes: This table documents gender differences in applications from 2015 to 2019 for positions in employees' choice sets. I restrict to employees at low-level positions without leadership responsibility. Each coefficient stems from a separate linear probability model where the outcome of interest is an indicator for applying to a first-level leadership position, scaled by 100. Column 1 provides the raw estimates. Column 2 adds employee demographics and characteristics of employees' current positions. Column 3 adds detailed characteristics of the job vacancy. The coefficient on leading a team indicates employees' revealed preferences for team leadership. Employee controls: Age, tenure, martial and family status, German nationality, full-time, hours, location, division, functional area, and job switch in past 12 months. Vacancy controls: location, division, functional area, job requirements stated in job ad, hours, salary, negotiation on the job. Standard errors are clustered at the employee level. *p< 0.10, **p< 0.05, ***p< 0.001.

Table 3.10: Application Likelihood for First Leadership Level: Controlling for Other Features

Panel A: Other features of job opening

	Job flexibility (1)	Negotiate on-the-job (2)	Work experience (3)
Female × Leading team	-0.0115**	-0.0122***	-0.0121**
	(0.0049)	(0.0047)	(0.0047)

Panel B: Selectivity of applicant pool for job opening

	Total applicants (1)	Graduate degree (2)	High performance (3)
Female × Leading team	-0.0121**	-0.0132***	-0.0132***
	(0.0047)	(0.0045)	(0.0047)

Panel C: Work environment of job opening

	Female supervisor (1)	≥ 1 female coworker (2)	Fem share of unit (3)
Female × Leading team	-0.0120***	-0.0097**	-0.0127***
	(0.0047)	(0.0049)	(0.0048)

Notes: This table documents gender differences in applications from 2015 to 2019 for positions in employees' choice sets. I restrict to employees at low-level positions without leadership responsibility. Each coefficient stems from a separate linear probability model where the outcome is an indicator for applying to a first-level leadership position, scaled by 100. In addition to controlling for employees' current characteristics and vacancy features, each regression tests whether an alternative channel might be driving the team leadership result by including the interaction of the respective variable with gender. Employee controls: Age, tenure, martial and family status, German nationality, full-time, hours, location, division, functional area, and job switch in past 12 months. Vacancy controls: location, division, functional area, job requirements stated in job ad, hours, salary, negotiation on the job. Standard errors are clustered at the employee level. *p< 0.10, **p< 0.05, ***p< 0.001.

Bibliography

Abadie, Alberto, "Semiparametric Instrumental Variable Estimation of Treatment Response Models," *Journal of Econometrics*, 2003, *113* (2), 231–263.

Abraham, Lisa and Alison Stein, "Words Matter: Experimental Evidence from Job Applications," Working Paper 2020.

Angrist, Joshua D and Guido W Imbens, "Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity," *Journal of the American Statistical Association*, 1995, *90* (430), 431–442.

_ , _ , and Donald B Rubin, "Identification of Causal Effects Using Instrumental Variables," *Journal of the American Statistical Association*, 1996, *91* (434), 444–455.

Arnold, David, Will Dobbie, and Crystal S Yang, "Racial Bias in Bail Decisions," *The Quarterly Journal of Economics*, 2018, *133* (4), 1885–1932.

Azmat, Ghazala and Rosa Ferrer, "Gender Gaps in Performance: Evidence from Young Lawyers," *Journal of Political Economy*, 2017, *125* (5), 1306 – 1355.

_ , Vicente Cuñat, and Emeric Henry, "Gender Promotion Gaps: Career Aspirations and Workplace Discrimination," *Available at SSRN 3518420*, 2020.

Baker, George, 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.

Barbanchon, Thomas Le, Roland Rathelot, and Alexandra Roulet, "Gender Differences in Job Search: Trading off Commute Against Wage," *The Quarterly Journal of Economics*, 2021, *136* (1), 381–426.

Beaman, Lori, Raghabendra Chattopadhyay, Esther Duflo, Rohini Pande, and Petia Topalova, "Powerful Women: Does Exposure Reduce Bias?," *The Quarterly journal of economics*, 2009, *124* (4), 1497–1540.

Bennedsen, Morten, Francisco Perez-Gonzalez, and Daniel Wolfenzon, "Do CEOs Matter? Evidence from Hospitalization Events," *The Journal of Finance*, 2020, *75* (4), 1877–1911.

_ , Kasper Meisner Nielsen, Francisco Pérez-González, and Daniel Wolfenzon, "Inside the Family Firm: The Role of Families in Succession Decisions and Performance," *The Quarterly Journal of Economics*, 2007, *122* (2), 647–691.

Benson, Alan, Danielle Li, and Kelly Shue, "Promotions and the Peter Principle," *The Quarterly Journal of Economics*, 08 2019, *134* (4), 2085–2134.

_ , _ , and _ , ""Potential" and the Gender Promotion Gap," Working Paper 2021.

Bertrand, Marianne, "New Perspectives on Gender," in "Handbook of Labor Economics," Vol. 4, Elsevier, 2011, pp. 1543–1590.

_ , "Coase Lecture: The Glass Ceiling," *Economica*, 2018, *85* (338), 205–231.

_ and Antoinette Schoar, "Managing with Style: The Effect of Managers on Firm Policies," *The Quarterly journal of economics*, 2003, *118* (4), 1169–1208.

_ , Claudia Goldin, and Lawrence F. Katz, "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors," *American Economic Journal: Applied Economics*, July 2010, *2* (3), 228–55.

_ , Sandra E Black, Sissel Jensen, and Adriana Lleras-Muney, "Breaking the Glass Ceiling? The Effect of Board Quotas on Female Labour Market Outcomes in Norway," *The Review of Economic Studies*, 06 2018, *86* (1), 191–239.

Bhuller, Manudeep, Gordon B Dahl, Katrine V Løken, and Magne Mogstad, "Incarceration, Recidivism, and Employment," *Journal of Political Economy*, 2020, *128* (4), 1269–1324.

Blau, Francine D and Lawrence M Kahn, "The Gender Wage Gap: Extent, Trends, and Explanations," *Journal of Economic Literature*, 2017, *55* (3), 789–865.

Bloom, Nicholas and John Van Reenen, "Measuring and Explaining Management Practices Across Firms and Countries," *The Quarterly Journal of Economics*, 11 2007, *122* (4), 1351–1408.

Bosquet, Clement, Pierre-Philippe Combes, and Cecilia Garcia-Penalosa, "Gender and Promotions: Evidence from Academic Economists in France," *The Scandinavian Journal of Economics*, 2019, *121* (3), 1020–1053.

Bronson, Mary Ann and Peter S Thoursie, "The Wage Growth and Within-Firm Mobility of Men and Women: New Evidence and Theory," 2019.

Cappelli, Peter and JR Keller, "Talent Management: Conceptual Approaches and Practical Challenges," *Annu. Rev. Organ. Psychol. Organ. Behav.*, 2014, *1* (1), 305–331.

Carpio, Lucia Del and Maria Guadalupe, "More Women in Tech? Evidence from a Field Experiment Addressing Social Identity," *Management Science*, 2021.

Catalyst, "Pyramid: Women in the United States Workforce," Industry Report 2021.

Chakraborty, Priyanka and Danila Serra, "Gender and Leadership in Organizations: Promotions, Demotions and Angry Workers," Working Papers 20210104-001, Texas A&M University, Department of Economics January 2021.

Coffman, Katherine Baldiga, "Evidence on Self-Stereotyping and the Contribution of Ideas," *T*he Quarterly Journal of Economics, 09 2014, *1*29 (4), 1625–1660.

Coffman, Katherine, Manuela Collis, and Leena Kulkarni, "Whether to Apply?," Working Paper 2022.

Cortés, Patricia, Jessica Pan, Laura Pilossoph, and Basit Zafar, "Gender Differences in Job Search and the Earnings Gap: Evidence from Business Majors," Technical Report, National Bureau of Economic Research 2021.

Cullen, Zoë B and Ricardo Perez-Truglia, "The Old Boys' Club: Schmoozing and the Gender Gap," Technical Report, National Bureau of Economic Research 2019.

Delfino, Alexia, "Breaking Gender Barriers: Experimental Evidence on Men in Pink-Collar Jobs," 2021.

Deloitte, "Women in the Boardroom: A Global Perspective," Industry Report 2021.

Eckel, Catherine C and Philip J Grossman, "Men, Women and Risk Aversion: Experimental Evidence," *H*andbook of experimental economics results, 2008, *1*, 1061–1073.

Fairburn, James A. and James M. Malcomson, "Performance, Promotion, and the Peter Principle," *T*he Review of Economic Studies, 2001, *6*8 (1), 45–66.

Fenizia, Alessandra, "Managers and Productivity in the Public Sector," Technical Report, Econometrica Forthcoming.

Flory, Jeffrey A, Andreas Leibbrandt, and John A List, "Do Competitive Workplaces Deter Female Workers? A Large-Scale Natural Field Experiment on Job Entry Decisions," *T*he Review of Economic Studies, 2015, *8*2 (1), 122–155.

Fluchtmann, Jonas, Anita Marie Glenny, Nikolaj Harmon, and Jonas Maibom, "The Gender Application Gap: Do men and women apply for the same jobs?," Technical Report, IZA Discussion Papers 2021.

Fortin, Nicole M, "Gender role attitudes and the labour-market outcomes of women across OECD countries," *o*xford review of Economic Policy, 2005, *2*1 (3), 416–438.

Frederiksen, Anders, Lisa B Kahn, and Fabian Lange, "Supervisors and Performance Management Systems," *J*ournal of Political Economy, 2020, *1*28 (6), 2123–2187.

Friebel, Guido and Michael Raith, "Managers, Training, and Internal Labor Markets," Simon School Working Paper No. FR 13-31 2013.

Gee, Laura K, "The More You Know: Information Effects on Job Application rRtes in a Large Field Experiment," *M*anagement Science, 2019, *6*5 (5), 2077–2094.

Gibbons, Robert and Michael Waldman, "Careers in Organizations: Theory and Evidence," *H*andbook of Labor Economics, 1999, *3*, 2373–2437.

Gneezy, Uri, Muriel Niederle, and Aldo Rustichini, "Performance in Competitive Environments: Gender Differences," *T*he Quarterly Journal of Economics, 08 2003, *1*18 (3), 1049–1074.

Goldin, Claudia, "A Grand Gender Convergence: Its Last Chapter," *A*merican Economic Review, 2014, *1*04 (4), 1091–1119.

_ , Sari Pekkala Kerr, Claudia Olivetti, and Erling Barth, "The expanding gender earnings gap: Evidence from the LEHD-2000 Census," *A*merican Economic Review, 2017, *1*07 (5), 110–14.

hkp, "Next Generation Talent Management Session on Talent Hoarding," Industry Report 2021.

Hoffman, Mitchell and Steven Tadelis, "People Management Skills, Employee Attrition, and Manager Rewards: An Empirical Analysis," *J*ournal of Political Economy, 2021, *1*29 (1), 243–285.

_ , Lisa B Kahn, and Danielle Li, "Discretion in Hiring," *T*he Quarterly Journal of Economics, 2018, *1*33 (2), 765–800.

Holmstrom, Bengt, "Moral Hazard and Observability," *T*he Bell Journal of Economics, 1979, *1*0 (1), 74–91.

Holmstrom, Bengt R. and Jean Tirole, "The Theory of The Firm," in "in," Vol. 1 of *Handbook of Industrial Organization*, Elsevier, 1989, pp. 61–133.

Hospido, Laura, Luc Laeven, and Ana Lamo, "The Gender Promotion Gap: Evidence from Central Banking," *T*he Review of Economics and Statistics, 2019, pp. 1–45.

Huitfeldt, Ingrid, Andreas Kostol, Jan Nimczik, and Andrea Weber, "Internal Labor Markets: A Worker Flow Approach," 2021.

i4cp, "Talent Mobility Matters," Industry Report 2016.

Kahn, Lisa B. and Fabian Lange, "Employer Learning, Productivity, and the Earnings Distribution: Evidence from Performance Measures," *T*he Review of Economic Studies, 2014, *8*1 (4), 1575–1613.

Kleven, Henrik, Camille Landais, and Jakob Egholt Søgaard, "Children and Gender Inequality: Evidence from Denmark," *American Economic Journal: Applied Economics*, 2019, *11* (4), 181–209.

KornFerry, "Executive Survey on Internal Talent Mobility," Industry Report 2015.

Kunze, Astrid and Amalia R. Miller, "Women Helping Women? Evidence from Private Sector Data on Workplace Hierarchies," *The Review of Economics and Statistics*, 12 2017, *99* (5), 769–775.

Lazear, Edward P, Kathryn L Shaw, and Christopher T Stanton, "The Value of Bosses," *Journal of Labor Economics*, 2015, *33* (4), 823–861.

Macchiavello, Rocco, Andreas Menzel, Atonu Rabbani, and Christopher Woodruff, "Challenges of Change: An Experiment Promoting Women to Managerial Roles in the Bangladeshi Garment Sector," Technical Report, National Bureau of Economic Research 2020.

Maida, Agata and Andrea Weber, "Female Leadership and Gender Gap within Firms: Evidence from an Italian Board Reform," *ILR Review*, 2020, *0* (0), 0019793920961995.

Malmendier, Ulrike and Geoffrey Tate, "Superstar CEOs," *The Quarterly Journal of Economics*, 2009, *124* (4), 1593–1638.

Mas, Alexandre and Amanda Pallais, "Valuing Alternative Work Arrangements," *American Economic Review*, 2017, *107* (12), 3722–59.

Matuson, Roberta, "5 Signs You've Become A Talent Hoarder," 2015.

McCue, Kristin, "Promotions and Wage Growth," *Journal of Labor Economics*, 1996, *14* (2), 175–209.

McKinsey and LeanIn.Org, "Women in the Workplace," Industry Report 2021.

Milgrom, Paul and John Roberts, "An Economic Approach to Influence Activities in Organizations," *American Journal of Sociology*, 1988, *94*, S154–S179.

_ and Sharon Oster, "Job Discrimination, Market Forces, and the Invisibility Hypothesis," *The Quarterly Journal of Economics*, 1987, *102* (3), 453–476.

Niederle, Muriel and Lise Vesterlund, "Do Women Shy Away from Competition? Do Men Compete Too Much?," *The Quarterly Journal of Economics*, 2007, *122* (3), 1067–1101.

OECD, "Gender Wage Gap (Indicator)," Accessed on 18 March 2022 2022.

Prendergast, Canice, "The Provision of Incentives in Firms," *Journal of Economic Literature*, 1999, *37* (1), 7–63.

_ and Robert H Topel, "Favoritism in Organizations," *J*ournal of Political Economy, 1996, *1*04 (5), 958–978.

Ransom, Michael and Ronald L Oaxaca, "Intrafirm Mobility and Sex Differences in Pay," *I*LR Review, 2005, *5*8 (2), 219–237.

Rosen, Sherwin, "Authority, Control, and the Distribution of Earnings," *T*he Bell Journal of Economics, 1982, pp. 311–323.

Sarsons, Heather, "Recognition for Group wWrk: Gender Differences in Academia," *A*merican Economic Review, 2017, *1*07 (5), 141–45.

_ , Klarita Gërxhani, Ernesto Reuben, and Arthur Schram, "Gender Differences in Recognition for Group Work," *J*ournal of Political Economy, 2021, *1*29 (1), 101–147.

Sullivan, John, "Managers Are Hoarding Talent: A Hidden Problem That's Likely Hurting Your Firm," 2017.

Waldman, Michael, "Job Assignments, Signalling, and Efficiency," *T*he RAND Journal of Economics, 1984, *1*5 (2), 255–267.

Wasserman, Melanie, "Hours Constraints, Occupational Choice, and Gender: Evidence from Medical Residents," *F*orthcoming, 2022.

Wiswall, Matthew and Basit Zafar, "Preference for the Workplace, Investment in Human Capital, and Gender," *T*he Quarterly Journal of Economics, 2018, *1*33 (1), 457–507.

Zuckerman, Gregory, *A Shot to Save the World: The Inside Story of The Life-Or-Death Race For a Covid-19 Vaccine*, Penguin Books Ltd, 2021.

Appendices

Appendix A

Supplementary Tables and Figures

A.1 Appendix Tables

Table A.1: Comparison of Analysis Sample to Representative Survey of German Workforce

	Analysis sample (1)	Large manufacturers (2)	German workforce (3)
Female	0.21	0.21	0.45
German citizen	0.89	0.85	0.88
Age (years)	43.3	44.4	43.9
Tenure (years)	14	11	7
Schooling (years)	16	12	12
Married	0.61	0.60	0.54
Children	0.54	0.63	0.62
Weekly hours	41	41	38
Manager	0.18	0.26	0.29
Observations	2x,xxx	1,848	13,791

Notes: This table compares average employee characteristics of my analysis sample in 2018 to representative survey measures for employees in Germany in 2018. Column 1 reports statistics for my analysis sample. Column 2 uses data from the BiBB survey on white-collar employees at manufacturing firms with at least 100 employees in West Germany. Column 3 uses data from the BiBB survey for all labor force participants in West Germany. Columns 2 and 3 are weighted to achieve representativeness of the German population.

Table A.2: Worker Quotations from Employee Survey

Form of talent hoarding	Quotations by workers
Informal underrating	“[My] former boss made me look bad at the potential new superiors behind my back in order to keep me on the team.”
Formal underrating	“Supervisors suppress potential ratings because of fear that employees will leave their current position for a promotion.” “Out of fear that workers will leave the team, supervisors tend to underrate employees.” “My supervisor would certainly find out if I applied to another position at the firm and that would have a negative impact on my assessment.”
Suppressed development	“[Supervisors] keep employees in their positions by preventing further development, rejecting training courses, and increasing workloads to prevent capacity for new tasks and development measures.” “Career development is not supported by direct supervisors, instead it is actively blocked with the goal of keeping people in their current positions.”
Soft pressure	“I decided not to apply internally ... because the message communicated to me in the employee dialogue was that I can’t leave the team within three years of joining.” “My boss strongly hinted that in order for my success until now to be considered I need to follow through on my project until the very end otherwise my efforts will not be fully taken into account.”
Threats	“... my supervisor communicated very openly in a workshop that if one of his employees applied and was not hired, his career in his current department would also be at an end.”
Retaliation	“The position should fit exactly so that negative effects of applying regarding the current manager are worthwhile.” “Fear of negative reactions from the supervisor: I have seen it from many colleagues who openly stated that they were applying. These colleagues then received no further training and no more interesting projects. The supervisor had written them off. This made the months leading up to the final change of job very difficult.”

Notes: This table displays worker quotations regarding talent hoarding at the firm. The quotations are based on workers’ free-text responses to questions on internal career development at the firm.

Table A.3: Manager Quotations from Employee Survey

Statement	Quotations by managers
Acknowledgment	"Many managers are not necessarily interested in developing workers or helping them to get a better job within the firm, because they would lose a good worker. Switches to other areas at the firm are not encouraged, even if it would have been the right move for the worker." "If you are good at what you are doing, it is very unlikely that you will be suggested for a higher-level position, especially if that position is in a different team." "Supervisors have the policy to do whatever it takes to keep people in their team, even if this means ignoring the development of the team members."
Misaligned incentives	"Managers have no interest in developing talents because they have no direct benefit from it." "Managers pursue their own goals and often prevent further development of workers, because they are not rewarded for developing talent." "Selfish managers are not willing to promote or recommend subordinates to other areas of the firm, even if that would add value to the firm." "Regarding the development of your subordinates, there often is a conflict of interests for managers, since the employee then usually leaves the team and the position will not be approved to be refilled again." "[Middle] Managers can't improve the situation themselves since they depend on the upper management. That is why they block the development of their subordinates."
Departure costs	"I observe that managers are not interested in letting good workers leave their current position or develop them, otherwise they would have to fill the current position again." "Managers don't actively support workers in switching positions because the vacant position is usually not (immediately) filled again."
Improvement	"Make sure that supervisors at all levels are incentivized to get their employees to the next stage of their career." "Positions are cut if an employee takes on a new job or even dies, so managers are afraid that their team will shrink and they don't know how to do all the work with the rest of the team. Managers would be more open to talk to employees about developing them if they would know that they can fill the vacant position again."

Notes: This table displays manager quotations regarding talent hoarding at the firm. The quotations are based on managers' free-text responses to questions on internal career development at the firm.

Table A.4: Effect of Manager Rotations on Applications by Incoming Manager's Characteristics

	Characteristics of incoming manager					
	Female	Male	Old	Young	Married	Unmarried
	(1)	(2)	(3)	(4)	(5)	(6)
Manger Rotation	0.0109	0.0149	0.0144	0.0148	0.0140	0.0161
	(0.004)	(0.002)	(0.002)	(0.003)	(0.002)	(0.003)
Outcome Mean	0.012	0.012	0.012	0.012	0.012	0.012
P-value of t-test	0.3707		0.9207		0.5523	
Observations	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx

	Characteristics of incoming manager					
	Firm tenure		Division		Functional area	
	Long	Short	Same	Different	Same	Different
	(1)	(2)	(3)	(4)	(5)	(6)
Manger Rotation	0.0149	0.0126	0.0140	0.0158***	0.0132	0.0149
	(0.002)	(0.004)	(0.002)	(0.003)	(0.002)	(0.002)
Outcome Mean	0.012	0.012	0.012	0.012	0.012	0.012
P-value of t-test	0.6220		0.6257		0.5582	
Observations	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx

Note: This table tests the influence of incoming managers' characteristics for the impacts of manager rotations on worker applications. Each coefficient stems from a separate regression based on Equation 1.3, where I restrict the rotation event to transitions to incoming managers with the respective characteristics. Worker controls: Female, age, German citizenship, educational qualification, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects. Robust standard errors in parentheses.

Table A.5: Effect of Manager Rotations on Job Transitions and Applications

	Job Transition		Internal Application	
	Internal	External	Incumbent	Newly hired
	(1)	(2)	(3)	(4)
Manager Rotation	0.0109	0.0009	0.0146	0.0175
	(0.002)	(0.001)	(0.002)	(0.002)
Outcome Mean	0.007	0.007	0.027	0.030
P-value of t-test	0.0000		0.4018	
Observations	3xx,xxx	3xx,xxx	1xx,xxx	1xx,xxx

Notes: This table provides evidence in support of talent hoarding as underlying mechanism. Columns 1 and 2 illustrate the effect of manager rotations on workers' job transitions within the firm (Column 1) and out of the firm (Column 2). Consistent with talent hoarding, only internal transitions are affected by rotations. Columns 3 and 4 document the application effects of manager rotations for workers who have been in the team before the manager arrived (Column 3) and workers who in the past were hired by the rotating manager (Column 4). I do not find that rotations have larger effects for workers who the manager was able to select herself, which is what one would expect under manager-worker-specific match effects. Controls: Female, age, German citizenship, educational qualification, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects. Robust standard errors in parentheses.

Table A.6: Effect of Manager Rotations by Destination of Application

	Division		Functional area		Location	
	Same	Different	Same	Different	Same	Different
	(1)	(2)	(3)	(4)	(5)	(6)
Manager Rotation	0.01557	0.00716	0.01407	0.00818	0.01769	0.00438
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)
Outcome Mean	0.01438	0.01434	0.01575	0.01310	0.01777	0.00963
P-value of t-test	0.0021		0.0307		0.0000	
Observations	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx

Notes: This table illustrates heterogeneity in the effect of manager rotations on worker applications by the destination of applications. Under talent hoarding, workers should disproportionally hold off on applications about which managers are likely to find out (e.g. because of proximity to the hiring manager). I assess three dimension of proximity compared to worker's current position: same versus different division (Columns 1 and 2), same versus different functional area (Columns 3 and 4), and same versus different location (Columns 5 and 6). Each coefficient stems from a separate regression based on Equation 1.3. Controls: Female, age, German citizenship, educational qualification, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects. Robust standard errors in parentheses.

Table A.7: Testing for Instrument Relevance and Independence

Panel A: Application Effects of Manager Rotation

	All	All	Men	Women
	(1)	(2)	(3)	(4)
Manager Rotation	0.0222	0.0224	0.0227	0.0203
	(0.003)	(0.003)	(0.003)	(0.006)
Outcome Mean	0.029	0.029	0.029	0.027
Quarter Fixed Effects	X	X	X	X
Other Controls	-	X	X	X
Observations	3xx,xxx	3xx,xxx	3xx,xxx	8x,xxx

Panel B: Balance Across Worker Characteristics by Manager Rotation

	No Rotation	Rotation	Difference (in %)
	(1)	(2)	(3)
German citizen	0.90	0.88	2.22
Age (years)	43.42	42.83	1.35
Tenure at firm (years)	13.35	12.75	4.49
Schooling (years)	15.81	15.98	-1.08
Married	0.62	0.61	1.61
Children	0.75	0.73	2.67
On parental leave	0.03	0.03	0.00
Full-time	0.92	0.93	-1.09
Weekly hours	41.15	41.12	0.00
Number of teammates	9.01	8.96	0.56
Quarters worked with manager	9.68	9.32	3.72
Performance rating	2.72	2.70	0.74
Past earnings growth	0.05	0.05	0.00
Past share absent	0.09	0.09	0.00
Past share applied	0.03	0.03	0.00
Past share internal switch	0.01	0.01	0.00
Observations			3xx,xxx

Notes: This table illustrates relevance (Panel A) and independence (Panel B) of manager rotation as instrument for worker applications. Panel A reports the first-stage effect of manager rotations on applications based on Equation 1.3. Columns 1 and 2 contain the full sample, Columns 3 and 4 focus on male and female workers, respectively. I estimate δ_1 by OLS regression and assess instrument relevance by testing the hypothesis that δ_1 is significantly different from zero. Controls: Female, age, German citizenship, educational qualification, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects. Robust standard errors in parentheses. Panel B compares average characteristics of workers, who do not experience a manager rotation (Column 1), and workers, who do experience a manager rotation (Column 2), in a given quarter. Column 3 represents the differences between Column 1 and Column 2 as % share of Column 1. Variables that refer to the past represent worker characteristics 12 months before.

Table A.8: Testing for Instrument Exclusion

Panel A: IV Estimates on Hiring by Rotating Manager's Ties to Destination

	Rotating manager has worked in the job opening's ...					
	Division		Functional area		Location	
	Ever	Never	Ever	Never	Ever	Never
	(1)	(2)	(3)	(4)	(5)	(6)
Applied	0.488	0.277	0.548	0.317	0.445	0.590
	(0.066)	(0.084)	(0.074)	(0.093)	(0.062)	(0.181)
Outcome Mean	0.005	0.002	0.004	0.002	0.005	0.002
Observations	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx

Panel B: IV Estimates on Hiring by Rotating Manager's Quality

	Promotion		Turnover		Absenteeism	
	High	Low	High	Low	High	Low
	(1)	(2)	(3)	(4)	(5)	(6)
Applied	0.479	0.471	0.518	0.463	0.486	0.537
	(0.111)	(0.063)	(0.137)	(0.060)	(0.113)	(0.111)
Outcome Mean	0.007	0.007	0.007	0.007	0.007	0.007
Observations	3xx,xxx	3xx,xxx	3 xx,xxx	3xx,xxx	3xx,xxx	3xx,xxx

Panel C: IV Estimates on Hiring by Rotating Manager's Exposure to Workers

	Exposure length in quarters					
	$\leq$1	2-3	4-6	7-8	9-10	11-14
	(1)	(2)	(3)	(4)	(5)	(6)
Applied	0.4429	0.3177	0.5593	0.4367	0.5135	0.4227
	(0.199)	(0.140)	(0.193)	(0.138)	(0.211)	(0.167)
Outcome Mean	0.009	0.008	0.008	0.008	0.008	0.007
Observations	3x,xxx	6x,xxx	7x,xxx	3x,xxx	3x,xxx	4x,xxx

Notes: This table finds no violation of the exclusion restriction regarding manager rotation as instrument for worker applications. Each panel presents two-stages least squares estimates of applying on getting hired based on Equation 1.5. Panel A estimates hiring outcomes for applications to which rotating managers have varying degrees of formal ties. Each column uses a different split of applications based on whether the rotating manager has ever worked in the same area as the job opening the worker applies to, where area is defined as division (Columns 1 and 2), functional area (Columns 3 and 4), and location (Columns 5 and 6). Panel B estimates hiring outcomes by rotating managers' quality using past leave-out team-level means for three outcomes: promotions (Columns 1 and 2), turnover (Columns 3 and 4), and absenteeism (Columns 5 and 6). Each column uses a different manager type to define the rotation event. Panel C estimates hiring outcomes by manager's length of exposure to workers. Each column uses a different split of workers based on the number of quarters the worker has been with the manager. Controls: Female, age, German citizenship, educational qualification, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects. Robust standard errors in parentheses.

Table A.9: Testing for Instrument Monotonicity

Panel A: Application Effects of Manager Rotations by Subgroup

Sample (1)	Observations (2)	Baseline application rate (3)	First-stage effect (4)	Standard error (5)
Age $\leq$ 40 years	1xx,xxx	0.0389	0.0244	(0.005)
Age > 40 years	2xx,xxx	0.0204	0.0213	(0.003)
Tenure $\leq$ 5 years	1xx,xxx	0.0348	0.0282	(0.006)
Tenure > 5 years	2xx,xxx	0.0255	0.0226	(0.003)
Schooling $\leq$ 13 years	1xx,xxx	0.0203	0.0130	(0.005)
Schooling > 13 years	2xx,xxx	0.0305	0.0258	(0.003)
Married	2xx,xxx	0.0259	0.0203	(0.003)
Not married	1xx,xxx	0.0322	0.0268	(0.005)
Parent	2xx,xxx	0.0265	0.0220	(0.003)
Non-parent	1xx,xxx	0.0337	0.0246	(0.005)
German citizen	3xx,xxx	0.0272	0.0232	(0.003)
Non-German citizen	4x,xxx	0.0338	0.0194	(0.009)
Team leadership	8x,xxx	0.0242	0.0164	(0.006)
No team leadership	2xx,xxx	0.0287	0.0242	(0.003)
High performance	2xx,xxx	0.0283	0.0284	(0.004)
Low performance	1xx,xxx	0.0289	0.0172	(0.004)

Panel B: Application Effects of Manager Rotations by Predicted Application Propensity

	Never applied (1)	Applied before (2)	Low-application team (3)	High-application team (4)
Manager Rotation	0.0183	0.0407	0.0216	0.0145
	(0.003)	(0.013)	(0.003)	(0.008)
Outcome Mean	0.017	0.067	0.019	0.067
Observations	3xx,xxx	4x,xxx	2xx,xxx	8x,xxx

Notes: This table finds no violation of the monotonicity assumption regarding manager rotation as instrument for worker applications. Panel A presents first-stage effects of manager rotation on applying for several subpopulations of interest, as indicated by Column 1. Estimation is based on Equation 1.3 and conducted separately in each subpopulation. Column 2 contains the number of observations, Column 3 presents baseline application rates, Column 4 provides first-stage effects, and Column 5 contains robust standard errors. I find that manager rotation has a positive and statistically significant first-stage effect (Column 4) for each subpopulation. Panel B presents first-stage effects of manager rotation on applying by workers' predicted application propensity. I use two approaches to predict workers' unobserved application propensity. Columns 1 and 2 split the sample by workers' own past application activity. Columns 3 and 4 split the sample by whether teams' leave-out application rates in the past were high or low. I find a significant and positive first-stage effect even for the subsets of workers who previously had a high propensity to apply (Columns 2 and 4). Controls: Female, age, German citizenship, educational qualification, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects. Robust standard errors in parentheses.

A.2 Appendix Figures

Figure A.1: Example of Firm's Internal Job Portal

Panel A. Search Interface

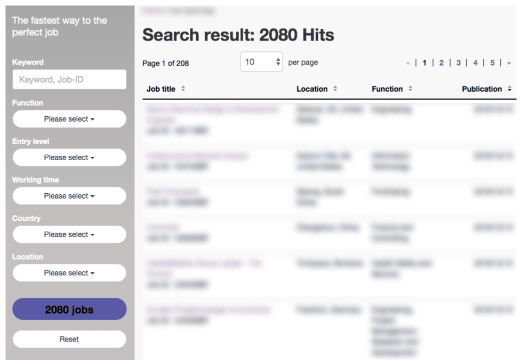

Panel B. Typical Job Ad

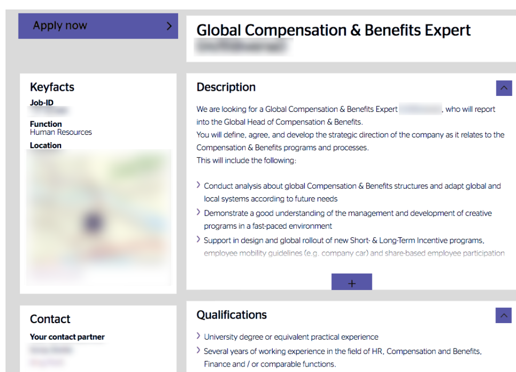

Panel C. Application Interface

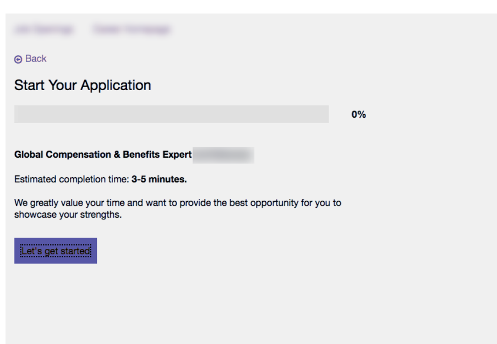

Notes: This figure provides an stylized example of the firm's internal job portal. Panel A displays the search interface, Panel B illustrates a typical job ad, and Panel C presents the application interface through which employees submit internal applications.

Figure A.2: Number of Manager Rotations by Length in Position

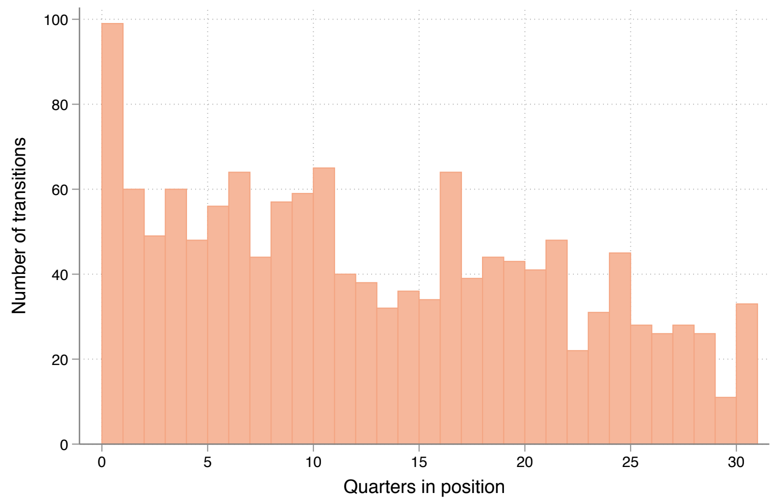

Notes: This figure illustrates the variation in the timing of manager rotations, as measured by the number of quarters a manager has been in their position at the time of rotation. The number of observations is 1,359, representing the total number of all internal job transitions in my sample.

Figure A.3: Effect of Manager Rotations on Team-Level Absenteeism Rates

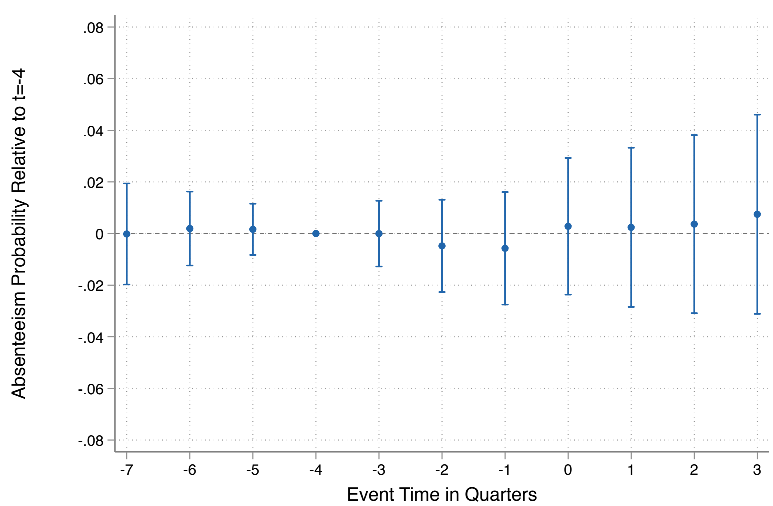

Notes: This figure documents team-level absenteeism rates around a manager rotation. Estimates stem from an event study regression, in which the outcome is the team average of absenteeism rates in a given quarter and event time is defined relative to the occurrence of a manager rotation. The specification includes team and quarter fixed effects. I bin event time dummy variables at $t = -8$ and $t = 4$ and cluster standard errors at the team level. The mean absenteeism rate as of $t = -4$ is 0.085. I find no evidence that manager rotations are preceded by changes in absenteeism. The sample of 6,xx teams includes those who have not experienced a manager rotation (i.e. never-treated).

Figure A.4: Application Effects of Manager Rotations by Exposure Length

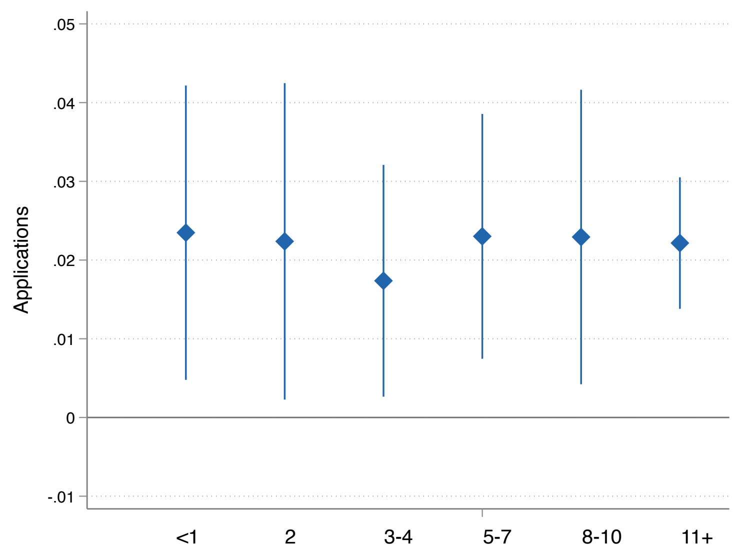

Notes: This figure assesses heterogeneity in the impact of manager rotation on applications by workers' length of exposure to the rotating manager. Each coefficient stems from a separate regression based on Equation (1.3) using robust standard errors. Worker subgroups are defined by the number of quarters a worker has worked under the manager. Baseline application rates are 0.031 (<1), 0.026 (2), 0.027 (3-4), 0.029 (5-7), 0.030 (8-10) and 0.026 (11+). The total number of observations are 3x,xxx (<1), 3x,xxx (2), 6x,xxx (3-4), 5x,xxx (5-7), 3x,xxx (8-10) and 1xx,xxx (11+). Controls: Female, age, German citizenship, educational qualifications, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects.

Appendix B

Data Appendix

This section provides additional information on the data I assemble. I illustrate how I merge personnel records and job application data. I present details on the survey I conduct at the firm. I then discuss the validity of the direct measures of managers' propensities to hoard talent. Finally, I provide details on the granular measure of internal job hierarchy used to test for misallocation.

B.0.1 Merging Personnel Records and Application Data

In order to relate internal application and hiring decisions to employees' career progression at the firm, I use a five-step matching algorithm to merge personnel records to application data. The algorithm uses exact matches based on names and date of birth, which matches over 90% of individuals.

Before beginning the matching process, I prepare names and birth dates to be matched. I standardize the names in both data sources. In order to be able to use date of birth in the matching algorithm, I need to impute birth dates for applicants since it is not contained in the main application data. I therefore parse 180,000 applicant CVs to isolate the date of birth for applicants mentioned on the CV and add it to the application data using a unique applicant identifier. This procedure allows me to capture birth dates for 33% of applicants in my sample. The remaining applicants will be merged using exact matches based on their names in later steps of the algorithm.

The matching algorithm follows an iterative process over five steps. In each step, individuals are only considered for matching if they have not been matched in previous rounds. The steps are constructed as follows:

- Step 1: Exact match on last name, first name, and birth date
- Step 2: Exact match on last name, first name, and year of birth
- Step 3: Exact match on last name, first three letters of first name, and birth date
- Step 4: Exact match on last name, first three letters of first name, and year of birth

– Step 5: Exact match on last name and first name

Note that Step 2 allows for the fact that there are different norms of whether to list month or day first when stating birth dates on a CV. Step 3 and 4 accommodate that some applicants add their middle name to their first name. I only keep exact matches. Disambiguous matches are resolved by using additional information on occupations, locations, and employees' work history. Since it is not clear how many applicants should be matched to the personnel records in the first place, a simple comparison of applicants and employees is not possible. Instead, I evaluate the match rate by checking how many applicants who got hired and therefore should appear in the personnel records are matched by the algorithm. This test represents a valid alternative to assess match quality, since the algorithm does not contain any specific treatment of hired applicants compared to applicants who are not hired. I find a match rate of over 90%. The match rate does not differ by gender.

B.0.2 Employee Survey

All employees in my sample were invited via e-mail by the firm's human resources department and were asked to provide their perspectives on the internal labor market at the firm. The survey received over 15,000 responses, yielding a 50.0% response rate. Respondents are similar to non-respondents in terms of demographics (Appendix Table B.1). I find no evidence for differential selection into response by gender (Appendix Table B.2).

Employees described challenges regarding their internal career progression both in the form of free-text responses and in multiple-choice answers. The median response time was 13 minutes. For my main analysis, I only keep respondents who took at least five minutes to respond and have no missing observations.

The translation of the relevant questions for this study is presented in abbreviated format:

A. Please rate following six statements. "Actively applying for positions at [Company] ..." {I strongly agree, I agree, Undecided, I do not agree, I totally do not agree}

A.1 "... increases future promotion chances."
A.2 "... does not matter since jobs are only posted proforma."
A.3 "... would cause negative consequences by my current supervisor."
A.4 "... is seen as disloyal to my current team."
A.5 "... is appropriate once employees are unsatisfied with their job."
A.6 "... should only be done after checking in with one's direct supervisor."

B. Which job characteristics are most important to you? Please select the two most important characteristics from the following list. {Potential for training, Potential for promotion, Pay, Flexible hours, Location, Meaningful tasks, Familiar tasks, Challenging tasks, Good relationship with colleagues, Good relationship with supervisor}

Table B.1: Comparison of Analysis Sample to Respondents of Employee Survey

	Sample (1)	Survey (2)
Demographics		
Female	0.22	0.24
German citizen	0.91	0.94
Age $<$30 yrs	0.09	0.12
Age 30-39 yrs	0.29	0.32
Age 40-49 yrs	0.26	0.26
Age $\geq$50 yrs	0.35	0.28
Tenure $\leq$2 yrs	0.14	0.14
Tenure 3-5 yrs	0.16	0.18
Tenure 6-9 yrs	0.15	0.17
Any children	0.57	0.58
Position Characteristics		
Weekly hours	37	40
Location small	0.16	0.15
Location medium	0.12	0.17
Location large	0.72	0.68
Engineering	0.45	0.41
Finance	0.05	0.05
Marketing and Sales	0.08	0.07
Observations	3x,xxx	1x,xxx

Notes: This table compares average characteristics of the analysis sample (Column 1) to the subset of employees who responded to the employee survey (Column 2). The sample of survey respondents is restricted to only contain responses who took at least five minutes to respond and have no missing observations.

C. At the end of this survey, we are interested in your personal opinion about current challenges and potential improvements with respect to careers at [Company].

C.1 What were the reasons why you decided in the past not to apply for internal job openings at [Company]? {free-text response}

C.2 What are the main challenges that you have encountered in your career development at [Company]? {free-text response}

C.3 What are some of the ways that [Company] could be helpful to you as you are planning your career? {free-text response}

Table B.2: Selection into Survey Response by Gender

	Survey response before reminder
	(1)
Female	0.023
	(0.0503)
Outcome mean	0.607
Av ME for Women	0.005
Gender Gap in %	0.9
Observations	1x,xxx

Notes: This table tests differential selection into survey responses by gender. Estimates stem from a logit regression of completing the survey before reminders were sent out on gender and employee controls. Controls: Age, tenure, nationality, children, team leadership, full-time, hours, location, functional area. Robust standard errors in parentheses.

B.0.3 Validity of Direct Measures of Talent Hoarding

This section presents validity exercises for my primary measure of talent hoarding, which infers managers' propensities to hoard talent based on the systematic suppression of potential ratings.

One way through which managers can hoard talent is by giving workers lower public potential ratings relative to their private performance ratings. Section 1.3.3 describes the construction of the measure. Even though performance and potential ratings are designed to capture different objects, potential ratings are highly predictive of future performance ratings. Among employees in my sample who are rated by their manager as having potential for higher-level positions, 86% actually receive a high performance rating once they get promoted to a higher-level position, motivating the comparison of performance and potential ratings. Moreover, managers' mean deviation between actual and predicted potential ratings is reasonably stable over time, supporting the systemic notion of talent hoarding the measure is meant to capture. When using earlier years to estimate a manager's mean deviation, its correlation with the manager's deviation based on later years is 0.64.

I find strong evidence against the possibility that the low potential ratings the talent hoarding measure identifies as underrating result from managers' accurate assessment of worker potential. When managers with high propensities to hoard talent rotate, underrated workers not only experience increases in applications and promotions, but are also likely to perform well in higher-level positions, demonstrating that the low potential rating was inaccurate.[1] In addition, while low potential ratings could in theory stem from the fact that managers have an incentive to hire low-potential workers to avoid the possibility of losing

[1] My 2SLS results demonstrate that these workers face a marginal probability of 0.15 (p-value 0.001) to land a position and a marginal probability of 0.08 (p-value 0.018) to perform in higher-level positions.

talent, this is not confirmed in the data. I find that talent hoarding effects occur both for incumbent workers and workers who are newly hired by a rotating manager (Columns 3 and 4 of Appendix Table A.5) .

The measure of talent hoarding is highly correlated with workers' realized visibility at the firm, confirming that managers' suppression of public signals has a meaningful impact. I measure worker visibility by collecting data on workers' nominations to succession lists. As in many large organizations, the firm compiles lists of three to five candidates as potential successors for about one-fifth of positions in my sample. The lists are assembled by HR employees who search for suitable candidates across the firm. Workers' appearance on such a list represents a measure of their visibility outside of the team. If a manager is successful at hoarding talent, worker visibility should be low, and thus their likelihood of appearing as a nominee on a succession list should also be low. I estimate a version of Equation 1.2 to compute the difference between actual nominations and predicted nominations, then classify managers as high- and low-propensity talent hoarders, defined as those in the bottom and top terciles of this difference.

The underrating of potential the measure captures does not appear to result from managers' involuntary mistakes. The deviation between performance and potential ratings signals is not driven by managers' ability to assess talent, as measured by the experience of leading a team, or other key manager attributes, such as gender, age, and experience. While Table 1.8 only reports the coefficient on gender for brevity, the coefficients on age and experience are not statistically significant. A F-test of all manager characteristics included in this logit regression—which besides gender include age, marital and family status, experience at the firm, division, function, and location— further rejects their joint significance. Moreover, survey evidence from employees in my sample documents that managers purposefully underrate potential (Appendix Table A.2).

I do not find evidence for alternative channels that would explain why managers suppress potential ratings. For instance, managers may be reluctant to rate a worker as high potential (despite high performance) if the worker has expressed disinterest in promotions. In contrast, under talent hoarding managers have less incentives to suppress potential ratings for workers who are less likely to leave the team. I use two different measures for workers' willingness to switch jobs to distinguish between these competing explanations: (i) workers' past internal applications and (ii) workers' consent that the firm can include the worker in their internal recruiting pool.[2] Appendix Table B.01 demonstrates that managers are more likely to suppress potential ratings for workers who have signaled their willingness to switch position using either measure, which is in line with talent hoarding as underlying mechanism, but contrasts with alternative explanations.

A placebo test provides further evidence against the importance of alternative channels. In the internal labor market, managers are generally likely to learn about workers' unsuccess-

[2]This internal feature is only used very recently at the firm and not available for the entire dataset. However, since it represents a public signal of the desire to switch positions, I use it in the available sample for robustness tests.

Table B.01: Impact on Talent Hoarding Propensities

	Underrated Potential		Undernominated for Succession	
	(1)	(2)	(3)	(4)
Applied in past 12 months	0.328		0.157	
	(0.015)		(0.039)	
Consent to be recruited		0.623		0.270
		(0.014)		(0.042)
Outcome Mean	0.3768	0.3768	0.0160	0.0160
Av ME for Women	0.0494	0.0913	0.0031	0.0057
Gender Gap in %	13	24	20	36
Observations	3xx,xxx	1xx,xxx	3xx,xxx	1xx,xxx

Notes: This table provides a robustness test for the direct measures of talent hoarding by examining the impact on managers' decisions to make worker talent visible. Each column is based on a separate logit regression at the worker level where the regressor of interest is whether the worker has applied internally in the past 12 months (Columns 1 and 3) or the worker has given their consent to be included in the firm's internal recruiting pool (Columns 2 and 4). Columns 1 and 2 estimate the propensity that managers manipulate worker visibility through suppressing potential ratings, Columns 3 and 4 focus on nominations to succession lists. Controls: Female, age, German citizenship, educational qualifications, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects. Robust standard errors in parentheses.

ful applications and have some power to diminish workers' application probability. However, it is unlikely that managers will be able to observe or intervene with respect to applications outside of the firm. Consequently, under talent hoarding, comparing managers with high and low propensities to hoard talent should lead to differential effects with respect to workers' internal career progression, but to zero impacts on transitions to jobs outside of the firm. However, if managers' suppression of potential ratings is not a sign for talent hoarding, but reflects other types of manager-specific behavior that affect workers, we would not necessarily expect the effect on external transitions to be zero and the effect under both high-propensity and low-propensity talent hoarders to be similar. I use this intuition to conduct a placebo test, comparing rotation effects for managers with high versus low propensities to hoard for internal applications, internal job transitions within the firm, and external job transitions out of the firm. Panel C of Appendix Figure B.02 documents that I find a zero effect on external transitions for both managers with high and low propensities to hoard talents, which contrasts my findings on internal applications (Panel A) and internal job transitions (Panel B). See Section D.1 for additional tests, which verify that my results are robust to choosing different cutoffs for the measure of talent hoarding.

Figure B.01: Mean Deviation Between Actual and Predicted Potential Ratings as Measure for Managers' Propensities to Hoard Talent

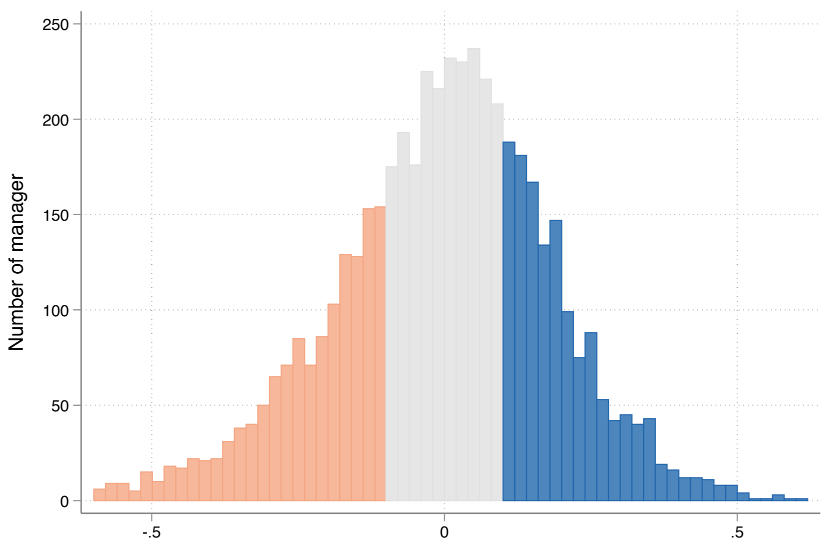

Notes: This figure depicts the mean deviation by manager between actual and predicted potential ratings, which captures systematic discrepancies in worker visibility, and serves as a measure of managers' propensities to hoard talent. The y-axis reports the number of unique managers with a given deviation. Each year, managers simultaneously conduct a performance rating (i.e. private signal of worker talent only shared with worker) and potential ratings (i.e. public signal of worker talent that is widely circulated) for each worker in their team. I assess managers' systematic underreporting of public potential ratings by comparing managers' actual potential rating to the predicted potential rating based on managers' own assessment of worker performance and worker characteristics. Values below zero represent managers who on average lower their public signal below their private signal of worker talent, which captures one likely dimension of talent hoarding. I use the bottom (marked in orange) and top (marked in blue) tercile of this distribution to classify managers as high-propensity versus low-propensity to hoard talents. The total number of observations is 7,xxx.

Figure B.02: Placebo Test for Talent Hoarding Measure

Panel A. Internal Applications

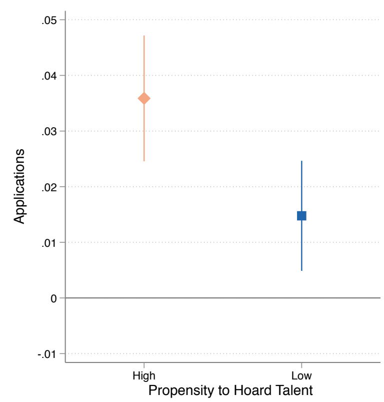

Panel B. Internal Transitions

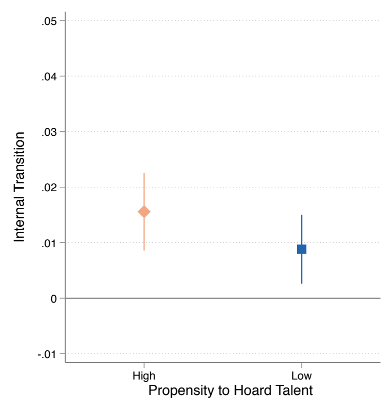

Panel C. External Transitions

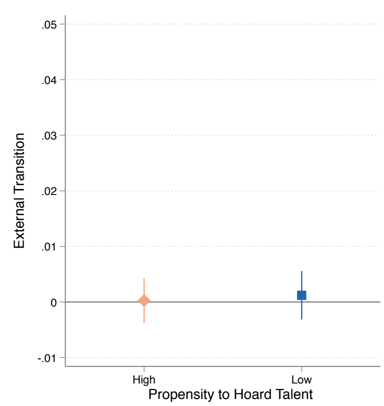

Notes: This figure provides a placebo test for my primary measure of managers' propensities to hoard tlanet based on the mean deviation between actual and predicted potential ratings. Each coefficient stems from a separate regression based on Equation 1.3 using robust standard errors. Each panel compares rotations of managers with high versus low propensity to hoard. The outcomes of interest are internal applications (Panel A), internal job transitions within the firm (Panel B), and external job transitions out of the firm (Panel C). Controls: Female, age, German citizenship, educational qualifications, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects. N=3xx,xxx.

Appendix C

Theoretical Appendix

In this section, I provide the formal derivations for predictions 4 and 5 referenced in Section 1.2.2.

Prediction 4. If $\beta_1 < \beta_2 \implies$ Pr[i applies$|\beta = \beta_1$]¿Pr[i applies$|\beta = \beta_2$]
This prediction implies that greater levels of talent hoarding reduce the number of workers who apply for a promotion.

Workers apply if $q(\alpha_i, \beta)b \geq c + \varepsilon_i$, where $\varepsilon_i \sim \Psi$ captures worker-specific heterogeneity. A worker's probability to apply can be expressed as Pr[i applies$|\beta_m$] $= \Psi\big(\mathrm{q}(\alpha_i,\beta_m)b - c\big)$. Because workers' promotion probability is decreasing in talent hoarding ($\frac{\partial q}{\partial \beta} < 0$), if $\beta_1 < \beta_2$:

$$
\begin{aligned}
q(\alpha_i, \beta_1) &> q(\alpha_i, \beta_2) \\
\Psi\big(q(\alpha_i, \beta_1)b - c\big) &> \Psi\big(q(\alpha_i, \beta_2)b - c\big) \\
\implies \Pr[i \text{ applies}|\beta = \beta_1] &> \Pr[i \text{ applies}|\beta = \beta_2]
\end{aligned}
$$

Prediction 5. If $\alpha_1 < \alpha_2$ and $\beta_1 < \beta_2 \implies \frac{\Pr[i \text{ applies}|\alpha_2,\beta_1]}{\Pr[i \text{ applies}|\alpha_1,\beta_1]} > \frac{\Pr[i \text{ applies}|\alpha_2,\beta_2]}{\Pr[i \text{ applies}|\alpha_1,\beta_2]}$
This prediction implies that greater levels of talent hoarding change the composition of applicants, causing a lower share of workers with high productivity in the applicant pool.

Let $r(\alpha_i, \beta_m)$ be Pr[i applies$|\alpha_i, \beta_m$]$= \Psi\big(\mathrm{q}(\alpha_i,\beta_m)b - c\big)$. We have assumed that $\frac{\partial^2 q}{\partial \beta \partial \alpha} < 0$, $\frac{\partial q}{\partial \beta} < 0$, $\frac{\partial q}{\partial \alpha} > 0$. We want to show that $\frac{\partial}{\partial \beta}\frac{r(\alpha_1,\beta)}{r(\alpha_2,\beta)} < 0$ for $\alpha_2 > \alpha_1$.

$$
\begin{aligned}
\frac{\partial}{\partial \beta}\frac{r(\alpha_2, \beta)}{r(\alpha_1, \beta)} &= \frac{\partial}{\partial \beta}\frac{\Psi(q(\alpha_2, \beta))}{\Psi(q(\alpha_1, \beta))} \\
&= \frac{\Psi(q(\alpha_2, \beta))\psi(q(\alpha_1, \beta))\frac{\partial q(\alpha_1,\beta)}{\partial \beta} - \Psi(q(\alpha_1, \beta))\psi(q(\alpha_2, \beta))\frac{\partial q(\alpha_2,\beta)}{\partial \beta}}{\big[\Psi(q(\alpha_1, \beta)\big]^2}
\end{aligned}
$$

Omitting the denominator since $\big[\Psi(q(\alpha_1, \beta)\big]^2 > 0$ leaves to show that

$$
\Psi(q(\alpha_1, \beta))\psi(q(\alpha_2, \beta))\frac{\partial q(\alpha_2, \beta)}{\partial \beta} < \Psi(q(\alpha_2, \beta))\psi(q(\alpha_1, \beta))\frac{\partial q(\alpha_1, \beta)}{\partial \beta}
$$

Rearranging leads to following expression

$$\underbrace{\frac{\frac{\partial q(\alpha_2,\beta)}{\partial \beta}}{\frac{\partial q(\alpha_1,\beta)}{\partial \beta}}}_{<0} < \underbrace{\frac{\Psi(q(\alpha_2,\beta))\psi(q(\alpha_1,\beta))}{\Psi(q(\alpha_1,\beta))\psi(q(\alpha_2,\beta))}}_{>0}$$

Since the left-hand side of the equation is below zero and the right-hand side of the equation is above zero, it holds that $\frac{\partial}{\partial \beta}\frac{r(\alpha_1,\beta)}{r(\alpha_2,\beta)} < 0$ for $\alpha_2 > \alpha_1$.

Appendix D

Robustness Results

D.1 Supplementary Results for Chapter 1

This section presents supplementary results that demonstrate the robustness of the main findings in Chapter 1.

I first verify that misallocation effects are not limited to major promotions. I show that similar patterns arise when considering other types of promotions. Appendix Table D.11 presents two-stage least squares results for any type of promotion (Column 1), small promotions (Column 2), and very large promotions (Column 3), which complement my preferred outcome of major promotions.[1] For each of these different promotion types, I find that marginal applicants, who only apply in the event of a manager rotation, face economically meaningful and statistically significant marginal probabilities with respect to landing higher-level positions and performing well in them.

Next, I test the robustness of the heterogeneity analyses with respect to the effect of manager rotations on worker applications, conducted in Section 1.5.3. I verify that my finding, which states that rotations of managers with higher propensities to hoard talent lead to bigger application effects, is not sensitive to specific cutoff choices I made when constructing the measure of talent hoarding. While my preferred approach compares manager rotations of managers in the bottom and top tercile of the mean deviation between actual and predicted potential ratings, Appendix Table D.12, Panel A, Columns 3 and 4 document very similar results when using bottom and top quartiles. Similarly, instead of assessing the effect of different manager rotations in the joint sample, Columns 1 and 2 of Appendix Table D.12, Panel A depicts similar results when splitting the sample of workers by whether their manager is in the bottom and top tercile. I conduct the same robustness tests for nominations to succession lists as a measure of talent hoarding (Columns 5 to 8 of Appendix Table D.12, Panel A) and find very similar patterns.

I use a similar approach to verify the robustness of the measure of worker quality. While

[1] While major promotions are defined as an increase of 20 or more in the hierarchy index, the cutoffs for small and very large promotions are 10 and 30, respectively.

my preferred approach uses the bottom and top quartile of my quality index, Panel B of Appendix Table D.12 documents very similar results when using bottom and top halves (Columns 1 and 2) and terciles (Columns 3 and 4) of my quality index. I also use an alternative measure to distinguish between high-quality and low-quality workers using years of education (Columns 5 and 6) and past performance ratings (Columns 7 and 8). The resulting application patterns are very similar to my preferred measure of worker quality.

Finally, I document that my complier analysis in Section 1.6 is not sensitive to the controls that I use for covariate adjustments. Appendix Table D.13 presents results from the complier analysis without the use of any controls. Appendix Table D.13 presents very similar patters with respect to the positive selection of marginal applicants compared to Table 1.4 which uses covariate adjustment, documenting the robustness of this finding.

Table D.11: Misallocation Effects of Talent Hoarding by Promotion Types

Panel A: 2SLS Results for Landing a Promotion

	Any promotion (1)	Small promotion (2)	Large promotion (3)
Applied	0.3519	0.2240	0.0725
	(0.053)	(0.043)	(0.017
Outcome Mean	0.0047	0.0024	0.0004
Observations	3xx,xxx	3xx,xxx	3xx,xxx

Panel B: 2SLS Results for Landing a Promotion and Outperforming the Team

	Any promotion (1)	Small promotion (2)	Large promotion (3)
Applied	0.4077	0.1154	0.0373
	(0.058)	(0.032)	(0.017)
Outcome Mean	0.365	0.365	0.365
Observations	3xx,xxx	3xx,xxx	3xx,xxx

Notes: This table documents the robustness of my main misallocation effects by evaluating different types of promotions. Column 1 refers to any type of promotion, Column 2 refers to small promotions (increase of 10 in hierarchy index), while Column 3 refers to very large promotions (increase of 30 in hierarchy index). Panel A reports estimates from two-stage least squares regressions on landing a promotion, where applying is instrumented for by manager rotation based on Equation 1.5. Panel B reports estimates from similar two-stage least squares regressions, but for the outcome of landing a promotion and performing better than the leave-out team average one year later. Controls: Female, age, German citizenship, educational qualification, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects. Robust standard errors in parentheses.

Table D.12: Heterogeneity in Application Effects by Talent Hoarding Levels

Panel A: Robustness in Measure of Managerial Talent Hoarding Propensity

	Public Signal		Public Signal		Succession List		Succession List	
	Bottom $\frac{1}{3}$	Top $\frac{1}{3}$	Bottom $\frac{1}{4}$	Top $\frac{1}{4}$	Bottom $\frac{1}{3}$	Top $\frac{1}{3}$	Bottom $\frac{1}{4}$	Top $\frac{1}{4}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Manager Rotation	0.0345	0.0187	0.0367	0.0137	0.0295	0.0150	0.0341	0.0171
	(0.006)	(0.005)	(0.007)	(0.005)	(0.005)	(0.005)	(0.006)	(0.006)
P-value of t-test	0.0361		0.0068		0.0442		0.0467	
Observations	1xx,xxx	1xx,xxx	3xx,xxx	3xx,xxx	1xx,xxx	1xx,xxx	3xx,xxx	3xx,xxx

Panel B: Robustness in Measure of Worker Quality

	Quality Index		Quality Index		Education		Performance	
	Top $\frac{1}{2}$	Bottom $\frac{1}{2}$	Top $\frac{1}{3}$	Bottom $\frac{1}{3}$	Top $\frac{1}{2}$	Bottom $\frac{1}{2}$	Top $\frac{1}{2}$	Bottom $\frac{1}{2}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Manager Rotation	0.0315	0.0116	0.0346	0.0101	0.0280	0.0177	0.0284	0.0172
	(0.004)	(0.003)	(0.006)	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)
P-value of t-test	0.0001		0.0002		0.0522		0.0484	
Observations	1xx,xxx	1xx,xxx	1xx,xxx	1xx,xxx	1xx,xxx	1xx,xxx	1xx,xxx	1xx,xxx

Notes: This table documents the robustness of the heterogeneity analysis with respect to the application effects of manager rotations. Each coefficient stems from a separate regression based on Equation 1.3. Panel A documents the robustness of the measure of talent hoarding based on deviations between managers' actual and predicted potential ratings. Columns 1 and 2 use the bottom and top tercile of this deviation, but conduct estimation in two separate samples of workers based on whether their manager is in the bottom or top tercile. Columns 3 and 4 use the same joint sample approach as my preferred measure, but split managers based on top and bottom quartiles instead of terciles. Columns 5 and 6 split workers into two samples based on whether their managers are in the top or bottom tercile of nominations to succession lists. Columns 7 and 8 use the same joint sample approach as my preferred measure, but split managers based on top and bottom quartiles instead of terciles. Panel B documents the robustness of my measure of worker quality. Columns 1 and 2 split workers into high-quality versus low-quality using the median of the quality index, while Columns 3 and 4 use top and bottom terciles. Columns 5 and 6 distinguish high-quality versus low-quality workers based on the median education level (having 18 or more years of schooling versus less than 18 years of schooling). Columns 7 and 8 make the quality distinction based on the median past performance rating. Controls: Female, age, German citizenship, educational qualification, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects. Robust standard errors in parentheses.

Table D.13: Characteristics of Marginal Applicants Without Covariate Adjustment (in %)

	All workers (1)	Always taker (2)	Marginal applicants (3)
German citizen	89.8	86.8	84.4
Age >=40yrs	60.4	38.7	50.6
Married	61.7	54.6	49.6
Children	73.3	68.1	63.5
Tenure at firm <2yrs	37.5	53.9	50.0
Tenure at firm 2-5yrs	40.5	37.7	38.2
Tenure at firm >=5yrs	21.9	8.4	11.9
Graduate degree	47.6	48.4	63.9
Full-time	92.5	94.4	97.5
High performance	54.0	56.7	63.5
High potential	28.2	44.2	43.4
Technical position	63.2	56.3	65.1
Low-level position	68.9	73.7	80.2
First-level leadership position	11.5	9.6	7.0
Time in position <2yrs	37.1	38.9	39.8
Time in position 2-5yrs	36.2	40.6	42.3
Time in position >=5yrs	26.7	20.6	17.9
Indicated desire to switch position	46.7	76.3	67.4
Nominated to succession list	1.6	2.4	5.6
Applied 12 months before	2.6	11.1	2.9

Notes: This table illustrates results from a complier analysis as described in Section 1.6. Each number is based on a separate regression *without* controls. Column 1 shows means for all workers, Column 2 represents always takers, and Column 3 represents marginal applicants, who only apply if managers rotate and talent hoarding temporarily subsides. Each number represents the share of workers in a given group that exhibit the respective characteristic (in %). A technical position is defined as a job related to engineering, IT, quality management, or production. Low-level positions are defined as positions at low hierarchy levels without leadership responsibility (i.e. individual contributors). First-level leadership represent positions with limited leadership responsibility, such as team leaders. I measure workers' indicated desire to switch position based on a recent internal feature at the firm that elicits workers' consent that the firm can include the worker in their internal recruiting pool. Controls: Female, age, German citizenship, educational qualification, marital status, family status, parental leave, firm tenure, position type, division, functional area, location, full-time, hours, team leadership, number of direct reports, performance and potential rating, time on position, and quarter fixed effects. N=3xx,xxx.