

The Old Boys' Club: Schmoozing and the Gender Gap*

Zoë Cullen

Harvard Business School

Ricardo Perez-Truglia

University of California, Berkeley

Abstract

Offices are social places. Employees and managers take breaks together and talk about family and hobbies. In this study, we show that employees' social interactions with their managers can be advantageous for their careers, and that this phenomenon contributes to the gender pay gap. We use administrative and survey data from a large financial institution and exploit quasi-random variation induced by the rotation of managers. We provide evidence that when employees have more face-to-face interactions with their managers, they are promoted at a higher rate. This mechanism could explain a third of the gender gap in promotions at this firm.

JEL Classification: J1, J16, J7, Z1, Z13.

Keywords: gender, career, promotion, social interactions, networking.

*First Draft: December 2019. This Draft: July 2022. Cullen (corresponding author): zcullen@hbs.edu, Rock Center 310, Boston, MA 02163. Perez-Truglia: ricardotruglia@berkeley.edu, 545 Student Services Building #1900, Berkeley, CA 94720. We are thankful for excellent comments from several colleagues and seminar discussants. We are grateful feedback from many colleagues and seminar participants at UCSD, MIT, Berkeley (Econ and Haas), Yale Cowles Inaugural Labor and Public Conference, University of British Columbia, University of Southern California, Cornell, Princeton NLS-E, INSEAD, Columbia University (Micro and Macro), University of Zurich, Georgetown, Tufts University, University of Copenhagen, NBER Summer Institute, RIDGE Public Economics Conference, The World Bank, FAIR Institute, Harvard Business School, American Economic Association Annual Meetings and the Discrimination and Disparities Intercollegiate Seminar. We are especially thankful to Leonardo Bursztyn, Thomas Fujiwara and Amanda Pallais. The collaborating institution provided financial support for the research being conducted. Additionally, Zoe Cullen was a full-time, salaried employee at that institution while the research was being conducted. Dylan Balla-Elliott, Katherine Fang, Anh Nguyen, Giacomo Stazi, Andrew Kao, Aakaash Rao and Jenna Anders provided excellent research assistance.

1 Introduction

Workplaces are social places. Employees and managers often discuss non-work-related topics such as sports, family, and hobbies in personal interactions that can extend beyond office hours – during lunch, smoking, coffee breaks, or the like. Through these interactions, employees form social bonds with their managers. In this study, we explore whether these social interactions influence employees’ careers and whether they can be partially responsible for the gender pay gap.

Women have a harder time than men climbing the corporate ladder. Among U.S. corporations, 48% of entry-level employees are women, but female representation falls to 38% at middle management, 22% at the C-Suite level, and 5% at the CEO level (McKinsey & Company, 2019). The gap in internal promotion rates accounts for the vast majority of the gender pay gap at the population level (Bronson and Thoursie, 2020). While improvements have been made over the last several decades, progress has been agonizingly slow. These gender disparities are not only unfair, but also inefficient, as misallocation of talent slows economic growth (Hsieh et al., 2019).

A growing body of literature has investigated what causes women to lag behind men in the corporate world. One factor that arises in the research is the “old boys’ club”, whereby men have a leg up in promotions because they can schmooze, network, and interact with more powerful men in ways that are less accessible to women.¹ This mechanism can create a self-perpetuating cycle: male managers promote a disproportionate share of male employees, who continue promoting other men.

Abundant anecdotal accounts suggest that the old boys’ club is real (Lang, 2011; Lee, 2014; Elting, 2018). For example, 81% of women say that they feel excluded from relationship-building at work, and many also feel excluded from after-work hours socializing (Gray and Barbara, 2013). Despite this anecdotal evidence, however, little quantitative evidence exists on this subject. Studying social interactions and long-term outcomes presents many challenges, among them are data challenges, as companies do not keep track of social interactions among their employees. And since the choice to socialize is typically endogenous, drawing causal inferences is difficult.

In this study, we provide novel evidence about the role of social interactions. We partnered with a large commercial bank in Southeast Asia (referred to hereinafter as *the firm*) with millions of customers, billions of dollars in assets and revenues, and thousands of employees. The firm is typical in that female representation drops off at higher levels: 75% of entry-level employees are women, falling to 61% among middle managers, 25% at the C-Suite level, and 0% at the CEO and company board levels. Moreover, the gender gaps in pay and promotion rates at the firm are similar to those documented for other corporations in both developed and developing countries.

We leverage rich sources of administrative data spanning four years (2015–2018), 14,638 unique employees, and 1,269 unique managers. These records include employees’ pay grades, the managers to which they were assigned, as well as measures of effort and performance. We also collected survey data to measure other aspects of the employees’ lives, such as whether they smoke or take breaks with their managers.

¹ The term “old boys’ club” was coined in reference to the British elite who attended certain public schools together. In current popular language, the term references the preservation of social elites in general.

In the ideal experiment, we would decide by the flip of a coin which employees get to socialize with their managers. Then, we would measure if the increased socialization has an effect on the employees' subsequent career trajectories. While the ideal experiment would be difficult to implement, we leverage quasi-experimental variation in social interactions driven by smoking habits. We conjectured that an employee who smokes assigned to a manager who also smokes will have increased social interactions with the manager, because of shared smoking breaks. In turn, the increased social interactions with the manager could give the employee a boost in promotions.

Our strategy for causal identification leverages the rotation of managers. In the organization under study, managers rotate across teams and divisions as part of the requirement for managerial promotion. Upon rotation, they assume responsibility for all employees on the team. We leverage the timing of those manager transitions for causal identification in an event-study framework. For example, consider two teams, each managed by a non-smoking manager. One of these teams then transitions from the non-smoking manager to a smoking manager; the other team transitions from the non-smoking manager to a different non-smoking manager. We can compare the outcomes of employees across these two teams before and after the transitions. As both teams undergo a manager transition, this design nets out the effect of the transition itself. Thus, we hypothesize that transitioning to a smoking manager, versus transitioning to a non-smoking manager, will result in better promotion prospects for the smoking employees but will have no effect – or perhaps a negative effect – on the promotion prospects of non-smoking employees.

In our context, a small minority (5%) of women smoke. For this reason, the analysis of smoking habits focuses on male employees and male managers. We analyze a 48-month panel with 2,907 unique employees (33% of whom smoke), 997 unique managers (19% of whom smoke), and 1,798 manager transition events. We conduct a series of empirical tests to confirm testimonies that the manager rotations are as good as random. We show that the type of transition faced by an employee (e.g., from non-smoking to smoking manager) is uncorrelated with the observable characteristic of the employee, as well as with the characteristics of the in-coming and the out-going managers. Most importantly, we use event-study analysis to show that the career progression of smoking and non-smoking employees follow parallel trends leading up to each type of manager transition.

We find a significant smoker-to-smoker advantage in promotions. Relative to smoking employees who transitioned from a non-smoking to a different non-smoking manager, smoking employees who transitioned from a non-smoking to a smoking manager are promoted more quickly.² For instance, at 10 quarters after the transition, the effect is estimated at 0.70 pay grades (p-value=0.002). By contrast, the non-smoking employees experienced similar promotion rates regardless of whether they transitioned from a non-smoking manager to a smoking manager or from a non-smoking manager to another non-smoking manager.

Our preferred explanation for the smoker-to-smoker advantage in promotions centers on the role of social interactions. Smoking employees may use the increased interactions with their smoking managers to gain their managers' favor and use these moments for self-promotion. Further, during these interactions, employees may also learn useful information, such as which tasks or training

² More precisely, we define the smoker-to-smoker advantage as the effect of smoking managers (relative to non-smoking managers) on the careers of smoking employees (relative to non-smoking employees).

are more conducive to promotions. Likewise, managers may learn more about their employees, identifying their effort, accomplishments, and potential. Next, we provide some evidence in favor of this mechanism and rule out some alternative explanations.

We use a survey measure of social interactions between an employee and their manager capturing the share of breaks they take together. We show that there is a smoker-to-smoker advantage in this measure of social interactions. After transitioning from a non-smoking manager to a smoking manager (relative to transitioning to another non-smoking manager), smoking employees spend an additional 24 percentage points (pp) of their breaks with their managers, representing a 63% increase relative to the baseline. In contrast, there is no effect for non-smoking employees.

A natural question is whether the smoker-to-smoker advantage in promotion is due to differences in productivity. For example, smoking managers might be better than non-smoking managers at retaining, motivating, and monitoring smoking employees. However, we do not find any evidence to support this proposition. When we estimate the effects of manager transitions on the probability of staying at the firm, we find point estimates close to zero, statistically insignificant, and precisely estimated. Likewise, we do not find any significant effects on measures of effort (the number of days worked and the number of hours spent in the office) or performance (the employee's own sales revenues).

It is possible that group affinity, instead of face-to-face interactions, explain our findings. To probe this channel, we study the effects of shared traits. A manager has shared traits with an employee if at least one of the following conditions are met: they were born in the same province (true of 16% of pairs), went to the same college (true of 8% of pairs), or were close in age (true of 43% of pairs). We reproduce the analysis of manager transitions, but focus on shared traits instead of co-smoking. That is, we estimate the effect of switching from a manager with whom the employee has no shared traits to one that has shared traits (relative to switching from one manager with no shared traits to a different manager with no shared traits). While these traits all have the potential to create a shared identity, we do not find evidence that they meaningfully increase the employee's promotion probability or interactions with the manager.

We provide additional evidence supporting the channel of face-to-face interactions. We split positions according to whether the employee works in physical proximity to the manager by combining administrative data on office locations with survey data. If driven by face-to-face interactions, the smoker-to-smoker advantage should be stronger when manager and employee pairs work in close physical proximity; by contrast, the effects should be smaller, or even absent, when the manager does not work in physical proximity to the employee. Consistent with the social interactions channel, we find that the smoker-to-smoker advantage is largely concentrated among employees who work in physical proximity to their managers.

Next, we explore whether the social interactions channel contributes to the gender pay gap at this organization. Men may chat with their managers while taking a smoking break; however, there may be plenty of other opportunities for interacting with the manager that have nothing to do with smoking. As a result, it is possible that all men, regardless of their smoking habits, leverage their social interactions with their manager to advance their careers.

We conduct a similar event-study analysis of manager rotations discussed above, with a few

key differences. First, we expand the sample to include female employees and female managers. Second, rather than measuring the smoker-to-smoker advantage, we measure the male-to-male advantage.³ This part of the analysis covers 14,638 employees (65% of whom are female), 1,269 unique managers (49% of whom are female) and 8,670 transition events.

We find a significant male-to-male advantage in promotions. After transitioning from a female manager to a male manager (and relative to transitioning from a female manager to another female manager), male employees are significantly more likely to be promoted. In contrast, this effect is absent for female employees. The male-to-male advantage is similar to the smoker-to-smoker advantage documented above in terms of its timing and magnitude. For instance, at 10 quarters after a male employee transitions from a female to a male manager, the pay grade for male employees is 0.60 points (p-value = 0.003) higher than those of male employees who transitioned from a female manager to a different female manager. This effect is significant in magnitude, roughly equivalent to a 14.6% increase in salary. For reference, our back-of-the-envelope calculations suggest that removing the male-to-male advantage would reduce the gender gap in pay grades by roughly 40%.

When measuring the male-to-male advantage, we provide evidence that the timing of manager rotations is as good as random. For instance, we show that the career progression of the male and female employees follow parallel trends leading up to each type of manager transition. Moreover, thanks to the larger sample size, we can provide an additional robustness check. The baseline analysis focuses on employees who “gain” a male manager (i.e., transitioning from a female manager to a male manager versus transitioning from a female manager to a different female manager). Additionally, we can study employees who “lose” a male manager (i.e., transitioning from a male manager to a female manager versus transitioning from a male manager to a different male manager). The expectation is that the effects of gaining a male manager should be roughly a mirror image of the effects of losing a male manager, in terms of both timing and magnitude. This is a sharp test in the sense that the coefficients are identified by a disjoint set of transition events, and thus there are no mechanical reasons why the results should mirror each other. Indeed, we find that the effects of losing a male manager are roughly a mirror image of the effects of gaining a male manager.

We find that, in addition to its timing and magnitude, the male-to-male advantage shares another core feature with the smoker-to-smoker advantage: all evidence points to social interactions as the underlying mechanism. According to the survey data, a large male-to-male advantage exists in interactions with the manager. After transitioning from a female manager to a male manager (and relative to transitioning from a female manager to another female manager), male employees are significantly more likely to share work breaks with their managers. On the contrary, we do not find any differences among female employees. We find that the male-to-male advantage is concentrated in positions where employees work in close proximity to their managers. And we find no significant male-to-male advantage in effort, performance or retention.

Our study builds on various strands of the literature. Most closely related to our study is the literature on the role of social interactions at work. Despite the universality of socializing in the

³ More precisely, we define the male-to-male advantage as the effect of male managers (relative to female managers) on the careers of male employees (relative to female employees).

workplace, relatively little is known about the returns of these personal interactions and whether these returns differ by gender. Cai and Szeidl (2018) provide experimental evidence that increasing the connections between business owners can increase firm productivity. Evidence also points to spillovers between business school classmates and executives (Shue, 2013; Lerner and Malmendier, 2013; Agarwal et al., 2016; Field et al., 2016).⁴ Evidence indicates that the manager’s social skills affect employee turnover (Hoffman and Tadelis, 2021). In the context of fruit-pickers, managers with fixed pay favor workers with whom they share a connection – to the detriment of firm productivity (Bandiera et al., 2010, 2009). And, in the context of politics, public officials may capitalize on their political and personal networks to gain influence (Cruz and Tolentino, 2019; Xu, 2018; Voth and Xu, 2021)

We contribute to this literature by providing novel evidence about the career and productivity consequences of social interactions in the corporate world. This context has abundant anecdotal evidence about the importance of social interactions and its effects on the gender pay gap, but minimal quantitative evidence probably due to data challenges. For instance, personal interactions are difficult to measure and also consist of sensitive information that firms will not typically want to make public. The lack of quantitative evidence is probably due to challenges with causal identification as well. For example, individuals tend to choose whom they interact with, so getting around that endogeneity challenge is difficult. We address both of these challenges. First, we provide causal evidence based on quasi-experimental variation in the gender and smoking habits of the managers. Furthermore, we collected unique sources of administrative and survey data about social interactions and physical proximity in a real corporation.

Our paper more broadly contributes to the large body of scholarship on the gender wage gap (Goldin, 2014), which demonstrates a consensus that this gap is primarily due to differences in promotion rates (Bertrand, Goldin, and Katz, 2010; Manning and Swaffield, 2008; Goldin, Kerr, Olivetti, and Barth, 2017). According to one account, the gap in internal promotion rates can account for approximately 70% of the gender pay gap by the age of forty-five (Bronson and Thoursie, 2020). Several explanations have been provided for these differences in promotions.⁵ Most related to our study, Kunze and Miller (2017) examine data from a private firm in Norway and find a positive association between the share of male managers at the establishment level and a gender gap in the promotion rate of employees. Other related studies measure the effects of female board members and executives (Bell, 2005; Bertrand et al., 2019; Cardoso and Winter-Ebmer, 2010; Dalvit et al., 2022; Flabbi et al., 2019), female school principals (Grissom, Nicholson-Crotty, and Keiser, 2012; Husain, Matsa, and Miller, 2022), female referees and female committee members (Bagues,

⁴ Another related study is Lleras-Muney et al. (2019), showing that friendships accumulated during high school can have lasting impacts on labor market outcomes. Mengel (2020) uses a laboratory experiment to show that both men and women engage in networking but men develop closer connections.

⁵ Some examples include the marriage market incentives (Bursztyn, Fujiwara, and Pallais, 2017), cultural norms (Bursztyn, Fujiwara, and Pallais, 2017; Alesina, Giuliano, and Nunn, 2013; Jayachandran, 2021), recognition for group work (Sarsons, 2017; Isaksson, 2019; Sarsons et al., 2021), differences in aspirations and performance (Azmat and Ferrer, 2017), the child penalty (Schönberg and Ludsteck, 2014; Bertrand et al., 2010; Kleven et al., 2019; Kuziemko et al., 2018), preference for flexible hours (Wasserman, 2022), and household work more generally (Cortés and Pan, 2019).

Sylos-Labini, and Zinovyeva, 2017; Card, Dellavigna, Funk, and Iriberry, 2020; Kim, 2020).⁶

We contribute to this literature in several ways. First, we provide causal estimates with the use of quasi-experimental methods, and we examine a range of career and performance outcomes. Second, we provide novel evidence of a specific mechanism, social interactions, which has been largely ignored in the literature on the gender pay gap.

The rest of the paper proceeds as follows. Section 2 describes the institutional context and the data. Sections 3 and 4 present the results for the smoker-to-smoker and male-to-male advantage, respectively. Section 5 provides additional discussion; the last section concludes.

2 Institutional Context and Data

2.1 The Firm

We collaborated with a private commercial bank in Southeast Asia. To keep the identity of the firm confidential, we refrain from providing exact information about its characteristics. This bank has millions of customers, billions of dollars in assets and in revenues, and thousands of employees.

Although we do not claim that the firm is representative of all firms in the world, we believe that this context is not an outlier. Since we study smoking habits, one important feature of the context is the smoking rates. In our sample, 33% of men and 5% of women smoke. According to 2016 estimates from the World Bank, these rates of smoking are almost identical to the world average of 35% among men and 6% among women.⁷

Because we study the gender pay gap, it is natural to compare the size of the gap at this organization versus that in other contexts. As a benchmark, the gender pay gap at the firm (23%) is close to the average of similar-sized firms in the financial sector in the United States (31%).⁸ The firm is typical in that men and women in a given position get paid similarly. The bulk of the gender pay gap thus is due to differences in positions among men and women. For example, 75% of firm employees at the entry-level are female, falling to 61% in middle management, 25% at the C-Suite level, and 0% at the CEO level. Data for U.S. corporations suggest a similar drop from 48% of female employees in entry-level positions to 38% in middle management, 22% in C-Suite positions, and 5% in CEO positions (McKinsey & Company, 2019). The content of the country (in Southeast Asia) is also comparable to that of United States in some dimensions, such as the gap in labor force participation.⁹ There are some notable differences, however. For instance, a majority

⁶ Other related studies look at the gender of peers instead of the gender of managers (Dahl et al., 2021; Hill, 2017; Stoddard et al., 2020), or the role of demographics other than gender, such as race (Mas and Moretti, 2009; Bandiera et al., 2010; Giuliano et al., 2011; Hjort, 2014; Glover et al., 2017).

⁷ These data are collected by the World Health Organization and made available by the World Bank in its Global Health Observatory Data Repository (last accessed on May 2021). Among OECD member countries, however, there is more balance in the smoking rates by gender: 29% of men smoke versus 18% of women.

⁸ Results based on wage rates for men and women working in the financial sector in firms with over 1,000 employees, as reported in Yildirmaz et al. (2019).

⁹ According to the World Bank Databank and International Labour Organization ILOSTAT database, the gender gap in labor force participation in the country of study (8.5%) is comparable to the corresponding U.S. gap (13.2%).

(64%) of employees in our sample is female; in comparison, in 2017, the U.S. workforce of Bank of America, for example, was 53% female.¹⁰

And, to the extent that we study social interactions in the office, the workplace context deserves some attention too. Arguably, we study the context of typical office jobs. Some employees work in one of the two headquarter offices, while the rest of the employees are dispersed in hundreds of branches around the country. Employees are expected to be in their offices from 9am to 5pm, and quite often stay for longer. As far as work breaks and social interactions are concerned, the firm maintains a general policy of managerial discretion during the work day, allowing the manager to decide whether and how to coordinate lunch breaks and other midday breaks. The social norms at the firm level include taking time out of the workday to celebrate employee birthdays, the firm's birthday, international women's day, and the new year. The American phrase "sharing is caring" is a popular phrase used by employees to capture the general sentiment that personal news is welcome.

2.2 Smoking

We collaborated with different divisions of the organization to create a centralized and anonymous database of employee characteristics and outcomes. One feature of the data that is critical for the analysis is the smoking status. We measure the smoking status of employees and their managers by combining multiple sources of data. First, we use data from the 2017 annual health exam, which included a question on smoking status.¹¹ These data provide a snapshot of smoking status for a large cross-section of employees as of September 2017. To complement the annual health exam data, we use two supplemental surveys.

First, we conducted a 2-minute survey that was exclusively dedicated to smoking habits.¹² We asked respondents about their own smoking status and the smoking status of current and past co-workers, including those who left the bank prior to the annual health exam. We invited 6,022 employees via email on February 2018 and received a response rate of 39%.¹³

The second survey was about a different topic (relationship with managers), but included a few questions related to smoking.¹⁴ We asked respondents whether they smoked and also about the smoking status of their current and past managers. We distributed the survey to employees in the largest division: sales and distribution, which comprises 62% of the firm's employees.¹⁵ We invited 4,847 employees by email to complete the survey in December 2017. A total of 3,345 employees completed the survey, implying a 69% response rate.¹⁶

¹⁰ Source: <https://about.bankofamerica.com/en-us/what-guides-us/our-global-workforce.html>.

¹¹ For some employees, their health review took place onsite during the workday, while other employees completed it online.

¹² A sample of the survey instrument is attached as Appendix A.

¹³ While survey participation was voluntary, the invitation mentioned cash prizes to be raffled off to survey respondents.

¹⁴ A copy of the survey instrument is included as Appendix B.

¹⁵ We were able to coordinate a detailed survey with the Sales and Distribution division because of the strong relationship we built with the head of that division.

¹⁶ The head of the division requested full participation from employees and gave permission to conduct the survey during work hours. We emphasized that answers to these survey questions would not be revealed to co-workers or managers.

We combine the different sources of data as follows. If an employee appears in the 2017 annual health exam data, we use their response to assign smoking status. For employees who do not appear in the annual health exam data, we impute their smoking status using data from the two complementary surveys.¹⁷ Using this method, we assign smoking status to 57% of employees who worked at the firm in the period of study. Of those, 59% are classified using their annual health exam, and the remaining 41% are classified using the complementary surveys.¹⁸

2.3 Pay Grade

Our main outcome variable is pay grade, which ranges from 41 to 66, and is observed on a monthly basis from January 2015 to December 2018. Pay grade is arguably the best measure of vertical career progression in the organization – indeed, employees commonly mention pay grades in conversations with coworkers to refer to their ranking. Typically an increase in pay grade is associated with a promotion.¹⁹ Consistent with anecdotal evidence, the changes in pay grade suggest that there is ample opportunity for upward mobility in the firm, with most employees experiencing at least one pay grade increase during the four-year period.²⁰ As in university contexts, the firm conducts merit reviews on a regular cycle. While this pattern changes from position to position, the typical employee faces a promotion opportunity every 1.5 years. In any given year, roughly two-thirds of employees face a promotion opportunity, and roughly half of those opportunities are converted into promotions. As a result, the average employee experiences an annual increase of 0.34 pay grades per year.

Due to the sensitive nature of the data, we do not have the full compensation details for the whole sample; however, we do have the compensation details for one specific point in time (March 2017). According to that data, pay grades explain the vast majority of variation in salaries.²¹ To aid in the interpretation of the findings, a 1-point increase in pay grade is associated with a 25% increase in salary (see Appendix C.1 for more details).

Through their merit review input, managers can have a direct influence on the career progression of employees on their teams. While this setting is highly competitive, employees are not necessarily competing with their teammates. There are no limits on the number of employees on a team that can be promoted, and different employees from the same team may seek promotions into different positions. Employees compete for promotions with employees from other teams in the firm, and as

¹⁷ More precisely, we classify an individual as a smoker if over one-third of the survey reports flag the individual as a smoker. This one-third threshold is arbitrary but largely inconsequential for the categorization.

¹⁸ Since some employees appear both on the annual health exam data and the complementary survey data, we can use the overlapping responses to test the validity of the complementary surveys. As expected, we find the two sources of data to be highly consistent with each other: the smoking status according to the complementary surveys coincides with the health records 82% of the time.

¹⁹ Conditional on an increase in pay grade, there is an 84% chance of a change in position title; in comparison, there is a 1% chance of a change in position title when there is no pay grade increase.

²⁰ More precisely, among the 7,622 employees who worked at the bank during the full sample period of four years, 50% experienced at least one pay grade increase, and 16% experienced more than one increase.

²¹ Typically, HR personnel carry out the precise salary negotiation within the range determined by pay grade, using market benchmark data.

the company routinely hires new employees, they implicitly compete with outside candidates.²²

2.4 Manager Assignments

We use longitudinal data from the firm’s organizational chart to link each employee to a manager in each month that the employee appears in the sample. The employee-manager assignment is constructed using a simple, two-step algorithm: identify the employee’s team, and then identify the “director” of that team.²³

To validate our manager assignments, we created a tailored question for the survey about managers (introduced in Section 2.2 above). We asked employees to identify the managers who “have directly influenced your key performance indicator and pay grade.” Respondents were shown six candidates identified from the organizational chart as being potentially the employee’s current manager or past manager. Most importantly, respondents were given the option to list an additional name if they did not see their current manager listed. Our manager assignment is highly consistent with the employees’ self-reports. For instance, in the month of the survey, December 2017, 91% of the managers we assigned using the organization chart were self-reported by the employee as a manager in the survey.²⁴

The managers tend to be significantly above their subordinates in the firm’s hierarchy. For example, the modal (mean) pay grade difference between managers and their employees is 5 (5.3). While the relationship with the manager may differ across different types of positions, the anecdotal accounts suggest that managers typically have a lot of influence on the careers of the employees. Most importantly, a manager provides key input in deciding whether to promote an employee. Even if the employee is not promoted, the manager provides input that influences that employee’s raises and bonuses. The manager also has discretion to distribute the workload across team members. Even if the work hours are rigid, such as for tellers, the manager still has latitude to approve leaves of absences or late days.

2.5 Manager Transitions

We focus on manager transitions that result from the reassignment of managers across teams as part of managerial rotations. The most typical case occurs when managers rotate laterally across different teams, but also include instances when the team’s manager is an existing employee who was promoted to a higher position, or a new hire. We identify these exogenous transition events in the data by observing that the new manager assumes responsibility for all employees on the

²² Indeed, during the period of study the firm experienced both high employee turnover (12.5% yearly) and net growth in the number of employees (5.9% yearly).

²³ An employee may have more than one “superior” to whom they report, so one challenge with this exercise is to identify the most relevant one. And in cases where the team has no directors listed in the organizational chart, we assign the team to the director listed at the next highest level in the organizational chart hierarchy.

²⁴ This comparison is restricted to pairs in the administrative organization chart that remain together for one year or more. When we include all pairs, even those who have been together for just one month, we still find substantial overlap: 78% of the assignments coincide with the employee’s self-report.

team. Each team has a unit number, and the manager is reassigned to a new unit number. In other words, the whole team, rather than a specific employee, experiences the manager transition. We exclude managers who are temporary replacements by excluding transitions whereby the new manager remained with the team for less than one quarter. As a result, new managers stay with the team for at least one quarter and may stay for up to several years. We also exclude events that are most likely team reorganizations, by excluding transitions in which more than half of the team members changed around the transition event.

Testimonies from executives and employees from Human Resources (HR) suggest that manager transition assignments can be considered as good as random. As part of corporate strategy, managers are expected to gain experience in all areas of banking. For this reason, managers are transitioned to other departments to gain exposure to new people and activities; for example, a manager from HR may move to a team in IT, and vice versa. By the time they reach the position of senior vice-president, most managers will have directed teams in most divisions. Opportunities for transferring happen on short notice. When managers quit or request a transfer, they are required to give thirty days' notice, and the set of candidates available to fill the role in time is (anecdotally) small. This shortage also helps explain why banks rotate managers quickly from distant divisions and why job postings for every managerial level of the bank can be found on the internal and external company dashboards. While all of these accounts can be persuasive, instead of relying on them, we use a variety of strategies to assess whether the manager transitions are truly exogenous.

Appendix C.2 provides a number of additional descriptive statistics about the transition events, which we summarize here. Manager transitions are distributed roughly uniformly over the four sample years. Each event will affect, on average, 6 employees. Around 44% of employees experience at least one event at some point in the four-year period, and 16% experience two or more events. We show that the sample of employees who experience a manager transition is quite representative of the whole firm in observable characteristics. Moreover, and consistent with the anecdotal evidence on the random nature of the transitions, we show that the characteristics of employees and managers are approximately balanced across the different types of manager transitions.

When interpreting the effects of manager transitions, one note is important to keep in mind: our estimates measure a reduced form effect of an increased but not necessarily permanent exposure to a manager of a given type. For the sake of simplicity, consider an employee who transitioned from a non-smoking manager to a smoking manager. As time passes, this employee may end up with a non-smoking manager again. For example, the employee may stay on the team but the team may be re-assigned to a non-smoking manager. Alternatively, the employee may move laterally or vertically to a different position that happens to be supervised by a non-smoking manager. In all of these cases, our estimates would under-estimate the effect of the manager's smoking status: if the employee were to stay with a smoking manager forever, the effects would presumably be even stronger. In practice, however, the attenuation bias is probably minor, as the transitions are, on average, quite persistent. During the 2.5-year window following the transition, smoking employees who gained a smoking manager spent on average 1.6 years more under a smoking manager (for details, see Appendix E.1). Likewise, during the 2.5-year window following the transition, male employees who gained a male manager spent, on average, an extra 1.5 years under a male manager

(for details see Appendix F.4).

2.6 Social Interactions with the Manager

The survey about managers opened by asking respondents to identify their current and previous managers. For each of the managers the respondent chose, we asked a series of questions about the respondent's relationship with that manager. The modal respondents reported information on their last three managers, resulting in a total of 9,068 employee-manager pairs.

The key question was about the frequency of social interactions with the manager: “out of 10 work breaks (including lunch or random breaks), how many would include [*Manager's Name*]?”²⁵ We construct a simple outcome that equals the fraction of breaks shared with the manager. This outcome ranges from 0, corresponding to no breaks shared with the manager, to 1, corresponding to all breaks shared with the manager.²⁶ Although the share of breaks taken with the manager is probably not a perfect measure of social interactions, in Appendix C.3 we provide some suggestive evidence that this outcome contains meaningful variation. First, we show that employees who spend more breaks with their managers are more likely to know more about their manager's personal life. Second, we show that spending breaks with the manager is predictive of future promotions.

2.7 Proximity to the Manager

To further investigate the social interaction mechanism, we split positions according to whether the employee works in physical proximity to the manager. For this classification, we combine two sources of data. For employees working in the headquarters (45% of the sample), we use card swipe data provided by the security division. These data include information about the floor on which the employee works, which we use to calculate the share of employees of each position who work on the same floor as their managers. If the share is above the median across all positions, we categorize the position as high-proximity; and if the share is below the median, we categorize the position as low-proximity. As a result, roughly half of the employees are categorized as high-proximity and the other half as low-proximity. In the high-proximity positions, 80% of employees work on the same floor as their manager, compared to only 8% among the low-proximity positions.²⁷

The card swipe data are not available for positions in the sales and distribution division, which are located outside of headquarters. We use a question from the manager relationship survey to

²⁵ To minimize the incentive for respondents to under-report so as to appear more focused and productive, we ask the question about a share of 10 breaks, rather than about the overall number of breaks. The downside is that we do not have a measure of the overall number of minutes spent together in a given week.

²⁶ The survey then asked a second question on social interactions with the manager, but those responses are not useful for the analysis because of the lack of significant variation. Specifically, we asked: “Of the last 10 emails you sent to [*Manager's Name*], how many included some part that was personal?” On average, employees report that just 5% of their emails are of personal nature. We suspect employees may have under-reported this type of behavior thinking that perhaps it would be seen as a violation of the firm's email policy.

²⁷ These statistics apply to the full sample of employees. For the subsample of employees used for the smoker-to-smoker analysis, these patterns look similar: 94% of the employees in the high-proximity group share the same floor as their manager, compared to only 37% in the low-proximity group.

fill this gap in the data. For each manager that the respondent listed in the survey, we asked the following question: “How often are (or were) you physically working near <manager name> (i.e., same floor and area)?” Respondents could choose from the following options: “Every day or most days (4–6 times per week),” “Some days (2–3 times per week),” or “Infrequently.” To split positions in high versus low proximity, we follow a similar procedure as the one described above for the headquarter offices. For each position, we calculate the average response to this survey question. Then, we categorize a position as high-proximity if the score is above the median across all positions, and low-proximity otherwise. Due to the availability of survey responses, we were able to categorize a majority (62%) of the positions in the sales and distribution division. By construction, roughly half these employees are categorized as high-proximity and the other half as low-proximity.²⁸

In total, we were able to successfully classify proximity for a large majority (82.5%) of the employees.²⁹ An example of a high-proximity position is the customer support specialist, who sits in a specific location near the manager. An example of a low-proximity position is the sales and quality development director, who usually travels between branches and reports back to the manager by phone or email.

2.8 Effort, Performance and Retention

To provide evidence of the underlying mechanisms, we rely on administrative data to measure effects on other employee outcomes. We have two measures of effort. The first measure is based on administrative records from the Human Resources (HR) division on employee absenteeism, including vacation days, sick days, and parental leave, among others. We proxy for the number of days worked by taking the maximum number of workdays in a month and then subtracting the days the employee was absent (according to these records). The second measure of employee effort is only available for employees in the headquarters offices (45% of the sample). These employees have to clock in and out using an electronic card-swipe system that is strictly enforced by security personnel. With data provided by the security division, we can use these time stamps to calculate the average number of hours an employee spent in the office on a given month. And we have one measure of performance, which is available only for employees with a sales role (42% of the sample). The bank uses an official formula to aggregate an employee’s sales across all products (e.g., credit cards, loans, mortgages). We use this formula to construct a monthly sales performance index.

To study effects on retention, we use the HR records: for employees who left the firm at some point, these records include the exact date when they left. Using this dependent variable presents a unique challenge for the event-study analysis. By construction, employees do not experience

²⁸ To create the average score by position, we coded the survey options with 0 for “Infrequently,” 1 for “Some days...”, and 2 for “Every day...” The average scores are 1.80 versus 1.48 in high-proximity versus low-proximity positions, respectively.

²⁹ We classify 100% of workers in headquarters and 62% of workers in the sales and distribution division. We consider a worker to be in headquarters or in sales and distribution if, over the course of the panel, they are only ever in that group.

manager transitions after they leave the company, which means that while we can estimate the post-transition coefficients, we cannot estimate the pre-transition coefficients. When using this outcome, we follow an approach similar in spirit to Kleven et al. (2019), by assigning “hypothetical” events to individuals who left the firm. More precisely, we take advantage of the fact that after an employee leaves the firm, the employee’s former team remains. Thus, we take the transition events experienced by the team and assign them to the employee, even if the employee no longer works for the firm.

2.9 Relevant Samples and Descriptive Statistics

We study the monthly evolution of pay grades from January of 2015 to December of 2018. Importantly, the analysis of smoker-to-smoker advantage is limited to a subset of the sample: male employees and male managers for whom we can assign smoking status. This subsample includes 2,907 unique employees and 997 unique managers; 33% of these employees (and 19% of unique managers) are smokers. There are a total of 1,798 manager transition events, involving 497 unique manager transitions, 1,226 unique employees, and 273 unique managers. The analysis of the male-to-male advantage uses the full monthly panel spanning 14,638 unique employees and 1,269 unique managers; 65% of these employees (and 49% of managers) are female. This sample includes 8,670 manager transition events, including 877 unique manager transitions, 6,021 unique employees, and 690 unique managers.

3 Results: Smoker-to-Smoker Advantage

3.1 Econometric Model

We introduce the econometric specification for the event-study analysis. Let $y_{i,t}$ be a generic outcome, where the subscripts i and t denote employees and time, respectively. The main outcome in our analysis is the employee’s pay grade, but we also consider other outcomes such as the employee’s effort, performance, and retention. Let S_i be an indicator variable that takes the value 1 if the employee smokes and the value 0 if the employee does not smoke. Likewise, we use superscripts S and N to refer to employees that smoke and do not smoke, respectively.

Let $D_{i,t}^j$ denote the traditional event-study variables that indicate the periods leading up to and following a transition event of type j . For example, $D_{i,t+e}^j$ is an indicator variable that equals 1 if individual i experiences an event of type j in period $t+e$, and 0 otherwise. $J_S = \{N2S, N2N, S2N, S2S\}$ denotes the set of types of manager transitions, where $N2S$ stands for a transition from a non-smoking manager to a manager who smokes, $N2S$ corresponds to a transition from a non-smoking manager to a different non-smoking manager, and so on. And let the set $\mathcal{E} = \{-30, -29, \dots, -4, 0, +1, \dots, +30\}$ be the event-study window, spanning from 30 months before the event to 30 months after the event.³⁰ Note that the omitted categories in \mathcal{E} represent the quarter prior to the event (i.e.,

³⁰ The width of this time window was chosen due to the length of our panel data. Following Stevenson and Wolfers

-3, -2, and -1 months). In the event-study graphs, we aggregate these monthly coefficients to the quarterly level for ease of presentation.³¹

The baseline econometric specification is the following:

$$y_{i,t} = \sum_{j \in J_S} \sum_{e \in \mathcal{E}} \beta_{j,e}^S \cdot S_i \cdot D_{i,t+e}^j + \sum_{j \in J_S} \sum_{e \in \mathcal{E}} \beta_{j,e}^N \cdot (1 - S_i) \cdot D_{i,t+e}^j + \gamma_i + \eta_{i,t} + \delta_t^S + \delta_t^N + \varepsilon_{i,t} \quad (1)$$

Note that we interact the event-study dummies with the dummy for whether the employee smokes (S_i) to estimate event-time coefficients for smokers ($\beta_{j,e}^S$) and non-smokers ($\beta_{j,e}^N$) separately. This baseline specification includes employee fixed effects (γ_i), manager fixed effects ($\eta_{i,t}$), and a separate set of month effects for smokers and non-smokers (δ_t^S and δ_t^N). All regressions use two-way clustering of the standard errors at the team and manager levels.

To isolate the impact of a change in manager smoking status from a change in manager more generally, we always compare between employees who are experiencing manager transitions. For example, we compare the effects of transitioning from a non-smoking manager to a smoking manager versus the effects of transitioning from a non-smoking manager to a different non-smoking manager. In the case of smoking employees, the object of interest is $\beta_{N2S,e}^S - \beta_{N2N,e}^S$ where e indicates the time since (or until) the transition date. In the case of non-smoking employees, the corresponding object of interest is $\beta_{N2S,e}^N - \beta_{N2N,e}^N$. Hereinafter, we refer to these objects as the *single-difference*, because they correspond to the difference between two types of transitions.

What we capture with the single-difference estimates is the impact of receiving a smoking manager relative to the impact of receiving a new non-smoking manager. However, we are ultimately interested in whether the effects of manager smoking status differ for smoking and non-smoking employees. For example, smoking managers who increase pay grades for smoking and non-smoking employees alike would not constitute evidence of a smoker-to-smoker advantage. Thus, we must take the difference of the single-difference estimates between smoking and non-smoking employees: $(\beta_{N2S,e}^S - \beta_{N2N,e}^S) - (\beta_{N2S,e}^N - \beta_{N2N,e}^N)$. A positive difference would be consistent with a smoker-to-smoker advantage. We refer to these estimates as the *double-differences*, because they take differences first with respect to types of transitions and second with respect to the employee's own smoking status.

3.2 Effects on Pay Grade

Figure 1 presents the main evidence of the smoker-to-smoker advantage in pay grades. These results are based on the comparison between employees who transitioned from a non-smoking to a smoking manager, relative to employees who transitioned from a non-smoking manager to another non-smoking manager. In each panel, the x-axis corresponds to the distance to the transition event, from 10 quarters leading up to a manager transition to 10 quarters after the manager transition. The quarter before the event (-1) corresponds to the omitted category, and thus the corresponding coefficient is always zero by construction.

(2006), we add absorbing dummies for the extreme categories of ≤ -31 and $\geq +31$ months.

³¹ In all the event-study graphs, the period "0" denotes the exact month of the transition event.

When inspecting Figure 1, the following must be kept in mind: this context has ample upward mobility, meaning that employee pay grades increase over time. The event-study coefficients refer to differences across transition types. As a result, a coefficient of zero in the post-transition period does not imply that employees remain in the same pay grade; rather, it indicates similar growth of pay grades across employees transitioning from non-smoking to smoking managers versus employees transitioning from non-smoking to non-smoking managers.

Panel (a) of Figure 1 corresponds to the single-difference estimates. Coefficients for smoking employees are shown as purple triangles, while the coefficients for non-smoking employees appear as orange diamonds. For instance, the purple triangles compare the pay grades of smoking employees who transition from a manager who does not smoke to a manager who does smoke (relative to transitioning from a manager who does not smoke to another manager who does not smoke).

Rather than relying exclusively on testimony that the manager rotations are as good as random, we test that assumption through the event-study framework. More precisely, we measure the evolution of the outcome before the date of the transition to confirm that, prior to the transitions, employees were on the same pay-grade trajectories. Panel (a) of Figure 1 shows that, prior to the event date, the coefficients for the smoking employees (denoted by the purple triangles) are statistically indistinguishable from the coefficients for the non-smoking employees (denoted by the orange diamonds). This evidence indicates that the assumption about parallel trends holds. In contrast, after the transition date, the evolution of pay grades starts to gradually diverge between smoking and non-smoking employees. At 10 quarters after transitioning to a smoking manager (relative to transitioning to another non-smoking manager), the pay grades of smoker employees increase by an additional 0.70 points (p -value=0.002). In contrast, the corresponding point estimate is close to zero (0.07) and statistically insignificant (p -value=0.722) for the non-smoking employees.

The estimates from panel (a) of Figure 1 suggest that the effects of the manager transition build slowly over time. The timing of the effects is largely consistent with the timing of promotions and transitions. While there can be differences across positions and units, employees typically face a promotion opportunity every 1.5 years, in either May or October of each year. By contrast, manager rotations can happen any month of the year. Indeed, these transition events are roughly uniformly distributed over the course of the year. Thus, while some employees may happen to be up for promotion right after the manager switch, other employees may need to wait months or even over a year until the next promotion opportunity arises. As a result, we would expect the effects on pay grade to accumulate with the time since the transition event, while more and more employees face promotion opportunities.

The event-study results have to be interpreted carefully. Consider, for example, the finding that smoking employees do better under smoking managers than non-smoking managers. Our preferred interpretation is that smoking managers are favorable to smoking employees; however, an alternative explanation could be that non-smoking managers are unfavorable to non-smoking employees. Likewise, the lack of negative effects on non-smoking employees suggests that the additional promotions of smoking employees are not crowding out promotions of non-smoking employees working on the same team; however, they probably crowd out promotions of employees working in different teams or external hires.

For a more direct measurement of the smoker-to-smoker advantage, panel (b) of Figure 1 presents the double-differences estimates: i.e., the coefficients for smoking employees (purple triangles from panel (a) of Figure 1) minus the corresponding coefficients non-smoking employees (orange diamonds from panel (a) of Figure 1). At 10 quarters after the transition, the smoker-to-smoker advantage is estimated at 0.63 pay grades (p-value= 0.035). This 0.63 effect on pay grade is roughly equivalent to a 15% effect on base salary.³² An alternative way of illustrating the magnitude of this effect of 0.63 pay grades is to compare it to the average change in pay grades. In 10 quarters following a manager transition, the average employee experiences a gain of 0.96 pay grades. The 0.63 effect is then equivalent to 65% ($= \frac{0.63}{0.96}$) of that baseline.

As an additional robustness check, Appendix D presents the results from a placebo exercise. We reproduce the event-study analysis, this time focusing on a characteristic that we know ex ante should not be relevant for promotions: whether the managers and employees have an odd or an even birthday. For instance, we would not expect that managers with an odd birthday would be beneficial to the careers of employees with an odd birthday. We reproduce the whole event-study analysis, but instead of slicing the data based on smoking status, we focus on whether their birth dates were odd or even. This test rules out mechanical reasons why our event-study framework would generate spurious effects, and allows us to assess whether our standard errors are adequate. As expected, we find that the estimates are close to zero, statistically insignificant, and precisely estimated.

3.3 Effects on the Time Spent with the Manager

Our preferred explanation is that smoking employees use the social interactions with their smoking managers to get a boost in promotions. To assess whether this mechanism is at play, we assess whether, after transitioning to a smoking manager, smoking employees share more breaks with their managers.

Because the sample sizes are much smaller with this survey outcome than with the pay grade data, we need to use a stylized version of the event-study framework.³³ We follow the same notation from Section 3.1 above, but with a few differences. The first difference is that, instead of an employee-month pair, each observation corresponds to an employee-manager pair: i denotes the employee and m the manager, respectively. Let $Share_{i,m}$ be the share of breaks that employee i took with manager m . Consider the following regression:

$$Share_{i,m} = \sum_{j \in J_S} \beta_{j,post}^S \cdot S_i \cdot D_{i,m}^j + \sum_{j \in J_S} \beta_{j,post}^N \cdot (1 - S_i) \cdot D_{i,m}^j + \sum_{j \in J_S} \beta_{j,pre}^S \cdot S_i \cdot D_{i,m+1}^j + \sum_{j \in J_S} \beta_{j,pre}^N \cdot (1 - S_i) \cdot D_{i,m+1}^j + X_{i,m} \gamma + \epsilon_{i,m} \quad (2)$$

³² A single pay-grade increase is associated with a log increase of 0.227 (Appendix C.1), and thus a 0.63 pay-grade increase should be equivalent to a 15% ($= e^{0.63 \cdot 0.227} - 1$) higher salary.

³³ Two reasons account for the smaller sample size. First, we collected survey data on a minority of employees. Second, even among surveyed employees, we measure their social interactions only at a handful of points in time (as opposed to the monthly data for four years from the administrative records).

$D_{i,m}^j$ is a dummy variable that equals 1 if individual i experiences an event of type j from manager $m - 1$ to manager m . The coefficients $\beta_{j,post}^S$ and $\beta_{j,post}^N$ are intended to capture the change in social interactions after the employee transitions to the new manager. In turn, $D_{i,m+1}^j$ is a dummy variable that equals 1 if individual i will experience an event of type j from manager m to manager $m + 1$. The coefficients next to these variables ($\beta_{j,pre}^S$ and $\beta_{j,pre}^N$) are intended to provide the usual tests for pre-trends: they measure whether future manager transitions affect the employee’s social interactions with the current manager. Additionally, the regression includes a set of basic controls ($X_{i,m}$): unit size, manager’s pay grade, dummies for whether the manager or employee smokes, as well as position title dummies. To make the effect sizes more intuitive, we follow an approach similar to Hastings et al. (2021) by shifting the coefficients by the baseline mean. And, for reference, we also report the baseline levels in each graph.

Figure 2 presents the results from the stylized event-study analysis. Both panels (a) and (b) correspond to the comparison between employees who transitioned from a non-smoking manager to a smoking manager, relative to transitioning from a non-smoking manager to a different non-smoking manager. Panel (a) presents the effects for smoking employees. According to the baseline levels, smoking employees do socialize with their non-smoking managers. The “post-transition” coefficient indicates that, after transitioning from a non-smoking manager to a smoking manager – and relative to transitioning from a non-smoking manager to a non-smoking manager – smoking employees interact more with their managers. More precisely, gaining a smoking manager increases the share of breaks that smoking employees take with their managers from 38 pp to 62 pp. This 24 pp increase is both statistically significant (p-value=0.002) and economically significant, amounting to a 63% increase with respect to the baseline level. In comparison, the coefficient “pre-transition” implies a placebo effect that is close to zero (-2 pp) and statistically insignificant (p-value=0.846).

Panel (b) of Figure 2 is identical to panel (a), except that it presents the results for the non-smoking employees instead of the smoking employees. The baseline level indicates that non-smoking employees socialize with their non-smoking managers. However, the “post-transition” coefficient indicates that gaining a smoking manager does not have a significant effect on the time that non-smoking employees spend with their managers. More precisely, the effect for non-smoking employees is close to zero (-3 pp), and statistically insignificant (p-value=0.625). Moreover, the “pre-transition” coefficient implies a placebo effect that is close to zero (1 pp) and statistically insignificant (p-value=0.845).

3.4 Effects on Attrition, Effort, and Performance

A question that naturally arises is whether the smoker-to-smoker advantage in promotions is due to differences in productivity. Smoking employees may reach higher positions under smoking managers because they work longer hours, perform better, or are more likely to stay at the firm when working under a smoking manager than their non-smoking counterparts.³⁴ To probe this

³⁴ For instance, smoking managers may be better able to discern the best candidates to promote due to information gleaned during breaks (Brogaard et al., 2014).

channel, we measure the effects of manager transitions on effort, performance, and retention.

These results are presented in Figure 3. Each panel of Figure 3 is equivalent to panel (b) of Figure 1, except that it uses a dependent variable other than pay grade. Panels (a) and (b) of Figure 3 correspond to the effects on the two measures of effort, while panel (c) corresponds to sales performance and panel (d) corresponds to retention. To make the comparison of coefficients more intuitive across different panels with different dependent variables, we follow an approach similar to Hastings et al. (2021): in each panel, we normalize the range of the y-axis to be approximately twice the within-individual standard deviation of the corresponding dependent variable. For example, the within-individual standard deviation in pay grade is about 0.5, so in the event-study graphs for that dependent variable the y-axis ranges from -1 to 1.³⁵

The results from Figure 3 show no evidence of differences in effort, performance, or retention. In panel (a) of Figure 3, the dependent variable is (the logarithm of) the monthly number of days worked. All the coefficients are close to zero, statistically insignificant, and precisely estimated. For example, the smoker-to-smoker advantage at 10 quarters after the transition is close to zero (0.015 log points) and statistically insignificant (p -value=0.707). This effect on days worked of (roughly a 1.5% increase) is not only statistically insignificant, but also economically insignificant relative to the corresponding magnitude of the smoker-to-smoker advantage (a 15% salary increase, from panel (a) of Figure 1). And, consistent with the assumption of balanced pre-trends, the coefficients preceding the transition date are close to zero, precisely estimated, and statistically insignificant.

In panel (b) of Figure 3, the dependent variable is (the logarithm of) the average number of hours spent in the office, according to the card swipe data. Again, we find no systematic evidence of a smoker-to-smoker advantage on the time spent in the office. However, these results must be taken with a grain of salt, as this outcome is defined for a subsample (employees in headquarters), the sample size is quite limited and thus so is the statistical power. In panel (c) of Figure 3, the dependent variable is the sales revenue index.³⁶ The point estimates are again close to zero, statistically insignificant, and precisely estimated. However, because these results are based on a subsample (employees with a sales role), they must be taken with a grain of salt due to the limited statistical power.

In panel (d) of Figure 3, the dependent variable is a dummy for whether the employee has left the firm. The evidence indicates a lack of smoker-to-smoker advantage on retention: the post-transition coefficients are close to zero, precisely estimated, and statistically insignificant. For example, at 10 quarters after the event, the smoker-to-smoker coefficient for attrition is close to zero (-0.010), statistically insignificant (p -value = 0.887), and precisely estimated. On average, 35% of workers experiencing a transition event have left the firm after 10 quarters. Relative to that baseline, the estimated smoker-to-smoker advantage of -1 pp is economically insignificant. One caveat with the results on retention, however, is that there are some significant pre-transition coefficients. Given that these differences in pre-trends do not show up in any other part of the analysis, our best guess is that this result is spurious.

³⁵ To allow for familiar scales, we use round numbers. For example, the within-individual standard deviation of pay grade is 0.517; so instead of using a range from -1.034 to 1.034, we use a range from -1 to 1.

³⁶ As this outcome has bunching at zero, we cannot use the logarithmic transformation for it.

3.5 The Affinity Channel

Perhaps the smoker-to-smoker advantage is simply due to affinity. That is, smoking managers may promote smoking employees not because of the additional time shared with them, but simply because of their inclination to favor employees with whom they have things in common. To probe this channel, we measure the effects of manager transitions in which the employee gains (or loses) a shared trait with the manager.

Because smoking status is not required for this analysis, and to maximize statistical power, we use the full sample (i.e., the same sample used for the analysis of the male-to-male advantage). We construct an indicator variable for whether the employee and manager have at least one shared trait: either they were born in the same province (true of 16% of pairs), went to the same college (true of 8% of pairs), or are close in age (true of 43% of pairs). As reported in Appendix E.2, we find that sharing one of these traits with the manager does not increase the share of breaks taken with the manager. As a result, if having a shared trait with the manager affected the employee's career progression, that would constitute evidence that the affinity channel, and not just social interactions, is at play.

The effects of shared traits with the manager on the employee's evolution of pay grade are presented in Figure 4. Panel (a) corresponds to the effects of gaining a shared trait with the manager: i.e., switching from a manager with whom the employee has no traits in common to one with whom the employee shares a trait, relative to switching from one manager with no traits in common to a different manager with no traits in common.³⁷ The pre-transition coefficients are close to zero and statistically insignificant, which supports the view that the timing of manager transitions is as good as random. The post-transition coefficients are positive but smaller in magnitude, and statistically less significant, than the corresponding coefficients for the smoker-to-smoker advantage. For instance, at 10 quarters after the transition, we find an increase in pay grade (0.04 pay grades) that is statistically insignificant (p-value = 0.504). Moreover, the magnitude of this affinity coefficient (0.04 pay grades) is much smaller than the corresponding coefficient for the smoker-to-smoker advantage (0.63 pay grades, from panel (b) of Figure 1). Thus, while we cannot rule out that affinity played some role, the evidence suggests that this channel is far from fully explaining the smoker-to-smoker advantage.

In panel (b) of Figure 4, we look at transitions in the opposite direction: i.e., comparing an employee who transitioned from a manager with a shared trait to a manager without shared traits, relative to transitioning from a manager with a shared trait to a different manager with a shared trait. Because these employees are losing the shared trait with the manager, the affinity channel would predict a drop in pay grade. Consistent with the results from panel (a), the results from panel (b) suggest that the affinity channel is not playing a major role. More specifically, the post-transition coefficients are close to zero and mostly statistically insignificant.³⁸ Again, our preferred interpre-

³⁷ Due to the nature of the transitions, we cannot estimate the single-difference estimates. Thus, we estimate the double-differences estimates instead.

³⁸ Some of the post-transition coefficients are statistically significant, but they are always economically small and have the opposite sign as the one predicted by the affinity channel. For instance, at 10 quarters after the transition, the coefficient is statistically significant (p-value=0.036) but is positive and quite small in magnitude (0.16 pay grades).

tation is that while the affinity channel may explain some of the smoker-to-smoker advantage, it is nowhere near explaining it entirely.

3.6 Heterogeneity by Proximity to the Manager

To provide further evidence on the social interactions channel, we explore the heterogeneity by proximity to the manager. If face-to-face interactions with the manager are driving the smoker-to-smoker advantage, we would expect this advantage to be more pronounced in positions with high proximity to the manager, where such face-to-face interactions are more likely to happen (Bandiera et al., 2009).

The results are presented in Figure 5. This figure reproduces the results from panel (b) of Figure 1, except that rather than having a single set of event dummies, we split this set in two, one set for high-proximity positions and another for low-proximity positions. The coefficients for the high-proximity positions are presented in panel (a) of Figure 5, while the coefficients for the low-proximity positions are reported in panel (b) of Figure 5. The results suggest that, consistent with the importance of face-to-face interactions, the smoker-to-smoker advantage is nearly three times as large in the high-proximity positions as in the low-proximity. For example, the estimated smoker-to-smoker advantage after 10 quarters is 1.02 paygrades ($p=0.017$) in the high-proximity sample; in contrast, the corresponding effect for the low proximity positions is only 0.34 paygrades, and statistically insignificant ($p=0.410$). However, this difference (1.02 versus 0.34 pay grades) has to be taken with a grain of salt because it is imprecisely estimated and thus statistically insignificant ($p\text{-value}=0.269$).

4 Results: Male-to-Male Advantage

The evidence presented above suggests that social interactions with the manager may give some employees a leg up in promotions. In this section, we explore whether this mechanism could contribute to the gender pay gap. Intuitively, smoking together on work breaks would be one opportunity men have to interact with each other, but other opportunities may arise that do not involve smoking. As a result, male employees – even the ones who do not smoke – may be more likely to be promoted than female employees.

4.1 Effects on Pay Grade

To estimate the male-to-male advantage, we conduct a similar event-study analysis of manager rotations discussed in Section 3 above, but with two key differences. First, we expand the sample to include not only male employees and managers, but also female employees and managers. Second, rather than measuring the smoker-to-smoker advantage, we measure the male-to-male advantage. More precisely, we use a specification identical to equation (1), except that the smoking status is

replaced everywhere by the gender status.³⁹

Figure 6 presents the main evidence on the effects of the manager's gender on employees' pay grades. This figure presents the single-difference coefficients, with the effects for female employees presented as red squares and the effects for male employees presented as blue circles. Panel (a) of Figure 6 corresponds to the effects of gaining a male manager: i.e., the comparison between transitioning from a female manager to a male manager, relative to a transition from a female manager to a different female manager.⁴⁰

Panel (a) of Figure 6 shows that, in the 10 quarters prior to the transition, the coefficients for male employees and the coefficients for female employees track each other closely. This lack of difference in pre-trends supports the view that the timing of manager transitions may be as good as random. While the pay grades of male and female employees track each other closely before the manager transition, after the transition they diverge. The pay grades of male employees rise more quickly when transitioning to a male manager compared to a female manager. At 10 quarters after the manager transition, the relative gain for male employees is 0.60 pay grades (p-value < 0.001), which is roughly equivalent to a 15% gain in salary.⁴¹

On the other hand, panel (a) of Figure 6 shows that the pay grades of female employees evolve similarly regardless of whether they transitioned to a female manager or a male manager. At 10 quarters after the transition, the corresponding coefficient for female employees is close to zero (-0.043 pay grades), statistically insignificant (p-value = 0.736), and precisely estimated. Moreover, this coefficient of -0.043 points for female employees is statistically different from the corresponding coefficient of 0.60 for male employees (p-value < 0.001). The lack of effects on female employees suggests that the additional promotions of male employees are not crowding out promotions of female employees working on the same team; however, they are probably crowding out promotions of female and male employees working on different teams as well as external hires.

For a more direct measurement of the male-to-male advantage, Figure 7 presents the double-differences estimates. For instance, panel (a) of Figure 7 corresponds to the difference between the male and female coefficients from panel (a) of Figure 6. Panel (a) of Figure 7 shows that, at 10 quarters after the transition, the male-to-male advantage amounts to 0.65 pay grades, which is not only highly statistically significant (p-value < 0.001) but also economically large.

The event-study coefficients need to be interpreted carefully. Take, for example, the finding that male employees do better under male managers than under female managers, which could be interpreted as evidence that male managers are favorable to male employees, that female managers are unfavorable to male employees, or a combination of both. Ideally, we would tease these two

³⁹ Instead of S_i , the alternative specification uses the dummy F_i that takes the value 1 if the employee is female and 0 if the employee is male. The set of manager transition is $J_G = \{F2M, F2F, M2F, M2M\}$, where $F2M$ denotes a transition from a female manager to a male manager, $F2F$ denotes a transition from one female manager to another female manager, and so on. Last, instead of superscripts S and N , we use superscripts F and M to refer to female and male employees, respectively.

⁴⁰ These coefficients refer to differences across transition types. For reference, Appendix F.2 reports the raw coefficients $\beta_{j,s}^M$ and $\beta_{j,s}^F$.

⁴¹ A single pay-grade increase is associated with a log increase of 0.227 (Appendix C.1), and thus a 0.60 pay grade increase should be equivalent to a salary that is 15% ($= e^{0.60 \cdot 0.227} - 1$) higher.

findings apart by comparing them to a third group: “gender-neutral” managers. In some settings, it is possible to have a gender-neutral condition. For instance, one could remove all the information about gender from a job application. In a real-world, face-to-face context like ours, a gender-neutral condition is infeasible. Based on our knowledge of the institutional context and the fact that there is a significant gender gap in promotions that favor men, our preferred interpretation is that male managers are favorable to the careers of male employees.⁴²

4.2 Reverse Transitions

The sample sizes are substantially larger for the analysis of male-to-male advantage, which allows for sharper robustness checks. Panel (a) of Figure 6 presents the results for gaining a male manager: i.e., transitioning from a female manager to a male manager versus transitioning from a female manager to a different female manager. In turn, panel (b) presents the case of losing a male manager: i.e., transitioning from a male manager to a female manager versus transitioning from a male manager to a different male manager.⁴³ The expectation is that the effects of gaining a male manager should be roughly a mirror image of the effects of losing a male manager, in terms of both timing and magnitude. This is a sharp test in that the coefficients are identified by a disjoint set of transition events, and thus there are no mechanical reasons why the results should mirror each other.⁴⁴

The comparison between panels (a) and (b) of Figure 6 indicates that, as expected, the effects of losing a male manager are approximately a mirror image of the effects of gaining a male manager. Female employees do not fare better or worse by gaining a male manager (panel (a)) or losing a male manager (panel (b)). By contrast, male employees end up with higher pay grades after gaining a male manager (panel (a)), and end up with lower pay grades after losing a male manager (panel (b)). The timing of these effects is similar across panels in that they slowly accrue over time.

For a more quantitative comparison, we can turn to the double-difference estimates from Figure 7. Panel (a) of Figure 7 shows that at 10 quarters after the transition, the male-to-male advantage amounts to 0.65 pay grades (p-value<0.001). Panel (b) of Figure 7 is equivalent to panel (a), except that it corresponds to the transitions in the opposite direction. According to panel (b), at 10 quarters after the transition, there is a male-to-male advantage of 0.44 pay grades (p-value<0.001). While these point estimates (0.65 versus 0.44) differ somewhat, we cannot reject the null hypothesis that they are equal (p-value= 0.343).

To maximize statistical power, we can estimate the male-to-male advantage based on all four types of manager transitions. Let the *dual-double-differences* be the combination of the double-

⁴² For more supporting evidence, see Appendix F.1.

⁴³ Male employees transitioning from a male manager to a different male manager may lose some of the relational capital they accrued by schmoozing with their original manager. However, male employees transitioning from a male manager to a female manager experience the same loss. So, as long as we take the difference between the two transition types, the effects of losing the relational capital should be differenced out.

⁴⁴ We can conduct a similar robustness check for the smoker-to-smoker analysis by comparing the effects of gaining a smoking manager to the effects of losing a smoking manager. Unfortunately, due to the sample size, we are under-powered for such analysis. The results are reported in Appendix E.3.

differences estimates from gaining a male manager and the (negative of) the double-differences estimates from losing a male manager.⁴⁵ The dual-double-differences estimates are presented in panel (c) of Figure 7. The estimated male-to-male advantage amounts to 0.54 pay grades at 10 quarters after the transition (p-value<0.001).

4.3 Timing of the Effects

Because the dual-double-differences estimates are quite precisely estimated, we can provide a more detailed analysis of the timing of the effects. As discussed in Section 3.2, due to the timing of promotion opportunities and manager transitions, we would expect the effects on pay grade to accrue slowly with the time as more and more employees face promotion opportunities.

We can use the estimates from panel (c) of Figure 7 to further explore the timing of the effects. The male-to-male advantage already appears in the very first quarter after the manager switch: the coefficient corresponding to +1 quarters is 0.10, and statistically significant (p-value=0.006). The male-to-male advantage then grows smoothly over time, presumably because more employees face promotion opportunities. To better illustrate this explanation, we can compare the size of the male-to-male advantage relative to the average change in pay grade at each time horizon. The average pay grade change in each of the first eight quarters after a manager transition were: 0.05, 0.15, 0.25, 0.34, 0.47, 0.56, 0.67, and 0.75. The estimated male-to-male advantage at the corresponding time horizons were, respectively: 0.10, 0.10, 0.12, 0.16, 0.20, 0.21, 0.30, and 0.38 (each of them statistically significant, with p-values of 0.006, 0.032, 0.022, 0.012, 0.011, 0.016, 0.001, and <0.001). In other words, the male-to-male advantage rises in proportion to the growth of pay grades. For instance, the male-to-male advantage roughly doubles from the end of the first year to the end of the second year, from 0.16 to 0.38. Likewise, the average pay grade change doubles in the same period, from 0.34 to 0.75.

4.4 Effects on the Time Spent with the Manager

The male-to-male and smoker-to-smoker advantages are not only similar in terms of timing and magnitude, but the bulk of evidence also suggests that they are both driven by social interactions. We start with the most direct test of this channel, based on the survey data on socialization with the manager.

The results for the stylized event-study analysis are shown in panels (c) and (d) of Figure 2. Both of these panels correspond to the comparison between employees who transitioned from a female manager to a male manager, relative to transitioning from a female manager to a different female manager. Panel (c) presents the estimates for male employees. According to the baseline levels, male employees do socialize with their female managers. The coefficient labeled “post-transition” indicates that, after transitioning from a female manager to a male manager – and relative to transitioning from a female manager to a female manager – male employees interact more with their managers. Gaining a male manager increases the share of breaks that male employees spend with

⁴⁵ More precisely: $\frac{1}{2} \{ (\beta_{F2M,e}^M - \beta_{F2F,e}^M) - (\beta_{F2M,e}^F - \beta_{F2F,e}^F) - [(\beta_{M2F,e}^M - \beta_{M2M,e}^M) - (\beta_{M2F,e}^F - \beta_{M2M,e}^F)] \}$.

their managers, from 46.7 pp to 61.2 pp. This 14.5 pp increase is both statistically significant (p -value=0.017) and economically significant, as it is equivalent to a 31% increase relative to the baseline. This figure also reports the corresponding event-study falsification test. The coefficient labeled “pre-transition” is close to zero (0.2 pp) and statistically insignificant (p -value=0.987).

The corresponding results for female employees are presented in panel (d) of Figure 2. According to the baseline levels, female employees do socialize with their female managers. For this group, however, there is no robust evidence that the share of breaks with the manager changed when transitioning from a female manager to a male manager, relative to transitioning from a female manager to a different female manager. The “post-transition” coefficient suggests a small (-8 pp) and statistically significant (p -value = 0.037) reduction in the share of breaks taken with the manager. While at first glance this finding could constitute evidence that female employees socialize less with male managers than with female managers, it must be taken with a grain of salt as it fails the falsification test: the “pre-transition” coefficient (-11 pp) is statistically significant (p -value=0.080) and similar in magnitude to the “post-transition” coefficient (-8 pp).

4.5 Heterogeneity by Proximity to the Manager

If socializing with the manager plays an important role, then we should observe stronger effects for employees whose jobs require frequent face-to-face interactions with the manager and potentially no effect when the manager and employee are physically separated.

Figure 8 presents the heterogeneity results. To maximize statistical power, we estimate the same dual-double-differences model from Figure 7. However, rather than having a single set of event dummies, we split this set in two: one set for high-proximity positions and another for low-proximity positions. Panel (a) of Figure 8 presents coefficients from high-proximity events, and panel (b) of Figure 8 presents coefficients from low-proximity events. Panel (a) of Figure 8 shows a significant male-to-male advantage when the employee works in high proximity to the manager. Panel (b) of Figure 8 further shows that the male-to-male advantage is close to zero and statistically insignificant when the employee works in a low-proximity environment. For example, panel (a) of Figure 8 indicates that at 10 quarters after the transition, the male-to-male advantage in pay grade is 0.76 (p -value<0.001) in the high-proximity group, compared to 0.21 (p -value=0.178) in the low-proximity group, and their difference is statistically significant (p -value= 0.013).

4.6 Additional Robustness Checks

Some additional robustness checks are presented in Appendix F and summarized below. The male-to-male advantage in promotions may be due to underlying differences in productivity. Contrary to this interpretation, however, Appendix F.3 shows that the male-to-male advantage in promotions is not accompanied by any differences in effort, performance, or retention. And Appendix F.5 shows that the results are similar under alternative definitions of manager transition events, such as dropping the largest transition events.

4.7 Comparison of Male-to-Male and Smoker-to-Smoker Advantages

The male-to-male advantage shares a number of important features with the smoker-to-smoker advantage, leading to our preferred interpretation: both advantages are driven in large part by social interactions. We summarize support for this hypothesis below.

The smoker-to-smoker and male-to-male advantages in promotions have similar timing, with effects building up over the course of two and a half years after the transition. In terms of magnitude, the two effects are also quite comparable. For example, at 10 quarters after the transition, the smoker-to-smoker advantage (0.63 pay grades, from panel (b) of Figure 1) is close to the corresponding male-to-male advantage (0.65 pay grades, from panel (a) of Figure 7), and this difference is statistically insignificant (p -value=0.956). Both the smoker-to-smoker and male-to-male advantages are driven primarily by employees who work in close proximity to the manager. In neither case do we see significant differences in effort, performance, or retention. And last, but not least, in both cases we observe comparable effects on the share of breaks taken with the manager. More precisely, the effect of smoking managers on smoker employees (25 pp, from panel (a) of Figure 2) is comparable in magnitude to the corresponding effect of male managers on male employees (14.5 pp, from panel (c) of Figure 2). Indeed, we cannot reject the null hypothesis that these two effects (25 pp and 14.5 pp) are equal (p -value=0.360).

Given that, in this context, men are more likely to be smokers than women (33% versus 5%, respectively), part of the male-to-male advantage that we measure arises mechanically from the smoker-to-smoker advantage.⁴⁶ In Appendix F.6, however, we show that only a small fraction of the male-to-male advantage can be directly attributed to the smoker-to-smoker advantage. That finding indicates that while smoking breaks provide one excuse for male employees to interact with male managers, men must have other opportunities to interact with their male managers that have nothing to do with smoking per se.

5 Discussion

5.1 Contribution to the Gender Pay Gap in Promotions

To assess the economic magnitude of the male-to-male advantage, we compute what would happen to the overall gender gap if we were to remove the male-to-male advantage. Our setting's unconditional gender gap is approximately 0.90 pay grades.⁴⁷ According to the dual-double-differences estimates, the male-to-male advantage amounts to 0.54 pay grades. As 66% of male employees have male managers, if the male-to-male advantage were to be removed, the gender pay gap would be reduced by 0.36 pay grades ($= 0.54 \cdot 0.66$). In other words, removing the male-to-male advan-

⁴⁶ Intuitively, because male managers are more likely to smoke than female managers, when a male employee transitions from a female manager to a male manager, he is also more likely to be transitioning from a non-smoking to a smoking manager than when transitioning to another female (almost certainly another non-smoker).

⁴⁷ This figure is estimated using a cross section of the bank in the last period of our sample (December, 2018).

tage would reduce the gender pay gap by 40% (from 0.90 to 0.54 pay grades).⁴⁸

To further contextualize the magnitude of these findings, we turn to a result that is well established in the literature: the so-called “motherhood penalty” (Schönberg and Ludsteck, 2014; Kleven et al., 2019). From the administrative HR data, we are able to identify workers who take maternity leave at any point in our sample.⁴⁹ Looking at the cross-section of employees in December 2018, we find that the gap between women who never took maternity leave and men is 0.83 pay grades. In comparison, the gap between women who took maternity leave and men is 1.09 pay grades. The difference between these two gaps is statistically significant ($p\text{-value} < 0.001$) and suggests that 23.9% ($= \frac{1.09 - 0.83}{1.09}$) of the gender pay gap could be attributed to the child penalty. While the 23.9% difference is not a proper causal estimate of the motherhood penalty, this simple calculation falls within the range of causal estimates provided in the literature.⁵⁰ For instance, Kleven et al. (2019) estimates that the motherhood penalty generates a long-run gender gap in earnings of around 20%. Overall, this finding constitutes suggestive evidence that, in the firm, the male-to-male advantage could be as quantitatively important as the motherhood penalty.

5.2 Comparison to other Studies

In terms of context, our results are most directly comparable to those of Kunze and Miller (2017), who study the effects of manager gender in a sample of private firms in Norway. Consistent with our results, they find a positive association between the share of male managers at the establishment level and a gender gap in the promotion rate of employees.⁵¹ While our study and Kunze and Miller (2017) primarily analyze managers, another strand of the literature focuses on the role of gender among the very top leaders, such as boardroom members and CEOs. On the one hand, female representation in the boardroom does not seem to significantly benefit the average female employee at the firm (Bertrand et al., 2019). On the other hand, evidence indicates that female leaders in executive positions, such as CEOs, benefit female employees (Cardoso and Winter-Ebmer, 2010), although those benefits seem to be concentrated in women at the top of pay distribution (Bell, 2005; Dalvit et al., 2022; Flabbi et al., 2019). Our social interaction mechanisms provide a potential explanation for the differences across these studies. Boardroom members do not have face-to-face interactions with the average employee at the firm, which may foreclose the male-to-male advan-

⁴⁸ One caveat with this interpretation is that if some effects were due to a negative effect of female managers on male employees, then the effects on the gender pay gap would be smaller. Thus, the 40% reported here may be considered an upper bound. In the extreme case where all effects are due to the negative effects of female managers on male employees, removing these manager effects should actually increase the gender pay gap, as male employees' pay grades would increase and female employees' would remain unaffected.

⁴⁹ Women are entitled to six months of maternity with partial pay, and in our sample, the average leave is 4.5 months.

⁵⁰ We measure maternity leave using the firm's HR records. This is a rough estimate; for example, some employees may have had children before they joined the company.

⁵¹ There are some differences in interpretation, however. Kunze and Miller (2017) note that their results could be driven by women helping women or men helping men. While they tend to prefer the interpretation that female managers are helping female employees, we prefer the interpretation that male managers favor male employees. This interpretation is supported by the evidence that female employees do equally well regardless of whether they transition to a female manager or a male manager.

tage. Executive leaders may interact face-to-face with other executives but not with employees at the bottom of the wage distribution, which could explain why female executives benefit women at the top of the pay distribution but do not benefit other women.

Yet another strand of the literature looks at the effects of gender of superiors in the education sector. The evidence indicates that male teachers in public schools are more satisfied with their jobs and more likely to remain working at a school if they have a male, rather than female, principal (Grissom et al., 2012; Husain et al., 2022). Also consistent with our evidence, those studies show that female teachers show similar job satisfaction and turnover rates whether working in schools run by female principals or male principals. Similarly, in the context of the academy, female referees and female committee members do not seem to benefit female candidates (Bagues et al., 2017; Card et al., 2020; Kim, 2020). However, in contrast to our evidence, male referees and male committee members do not seem to benefit male candidates. Again, our social interaction mechanism can provide a potential explanation for the differences across these studies. School principals have to engage in face-to-face interactions with teachers, but committee members and referees do not get to interact face-to-face with the candidates, which may explain why the gender of the superiors matters in one case but not the other.

We can also offer a quantitative comparison between our findings and the most closely related study. However, we must take this comparison with a grain of salt due to obvious differences in context and research design. With that caveat in mind, Kunze and Miller (2017) report a gender gap in promotion rates of 3.3 pp (page 772), a gap 2 pp larger in establishments with 100% male superiors, relative to establishments with 0% male superiors (column (1) of Table 2). We can use the estimates from Kunze and Miller (2017) for back-of-the-envelope counterfactual analysis. Because 83% of managers in their sample are male, the male-to-male advantage would account for 50% ($= \frac{0.83 \cdot 2}{3.3}$) of the overall gender gap. This percentage is in the same order of magnitude as the estimates from our study, according to which, in our context, the male-to-male advantage accounts for 40% of the gender gap in pay grades.

5.3 Social Interactions between Women

We find that men fare better under male managers than under female managers. One of the puzzling findings is that women do not seem to fare better under female managers than under male managers. This finding echoes results from other studies. For example, as discussed above, male teachers fare better under male than female principals, but female teachers fare similarly under female than under male principals (Grissom et al., 2012; Husain et al., 2022).

This finding can be interpreted in at least two ways. First, relative to women, men may spend more time interacting and socializing with their managers. In this regard, we suspect that cultural and institutional factors may be significant mediators. Men may have some unique opportunities to chat with each other that are less accessible to women, yet women may not have the equivalent opportunities to socialize with each other that are less accessible to men. For example, in this context, smoking breaks are socially acceptable and, furthermore, areas are officially designated for such a purpose; meanwhile, these norms and institutions are not as overt for female-centered

activities. Another reason men may have more time to socialize with each other could be that, in this context, gendered norms about parental duties may limit the time that women have to interact with each other during work or after hours (Juhn and Rubinstein, 2020; Cubas et al., 2019). To the extent the cultural and institutional factors may be different across countries and industries (Jayachandran, 2021), we would expect that the findings may differ somewhat across the different contexts.

Second, regardless of time spent with their managers, female employees may not leverage those face-to-face interactions to advance their careers as much as male employees do. Indeed, some related evidence is consistent with this view. For example, relative to women, men are more assertive in claiming credit for group work and also more aggressive in self-promotion (Sarsons et al., 2021; Isaksson, 2019; Coffman et al., 2021; Exley and Kessler, 2022). Additionally, evidence suggests that men may receive favorable treatment by managers by getting assigned tasks that are more conducive to promotions (Lehmann, 2013; Babcock et al., 2017).

Due to data limitations, we cannot disentangle these two explanations. For instance, we know little about how employees use time on their shared breaks – whether to promote themselves or to relax. Likewise, while we measure shared breaks with the managers, the social interactions may extend to work hours, the coffee shop, or the pub. Based on our survey evidence, our preferred interpretation is that both channels play some role. On the one hand, we observe that male employees increase their allocation of break time with the manager when he is male, and male managers shift the allocation of shared breaks toward their male employees. This constitutes suggestive evidence in support of the first explanation – that men interact more with their managers. On the other hand, we observe that female employees do spend a significant share of their break time with their managers, regardless of whether they are male or female. This constitutes suggestive evidence in support of the second explanation – that women do interact with their managers but do not leverage those interactions to advance their careers.

6 Conclusions

We presented evidence that social interactions with the manager may give some employees a boost in promotions. Furthermore, we documented that this mechanism may contribute significantly to the gender pay gap in promotions. We conclude by discussing some implications of these findings and avenues for future research.

Our identification strategy can be applied to other contexts. The rotation of managers is a common practice in large organizations, and the data on pay grades, assignments, and demographics could be obtained for most firms. Thus, our research design can be applied in other firms from different industries and countries to identify the contexts in which the male-to-male advantage is most pervasive. Indeed, our study already provides suggestive evidence that the male-to-male advantage may be exacerbated in some occupations, such as those in which the manager and the employee work in close proximity to each other. And we suspect that some cultural and institutional factors may mediate the size of the effects as well.

Our results can hopefully inspire future research to more fully elucidate the role of face-to-face interactions and networking for career advancement. For instance, our findings provide one potential explanation to a puzzle in the literature on working from home. In their seminal experiment, Bloom et al. (2014) found that while productivity rose among employees assigned to work from home, increasing by 13% over the nine months of the experiment, those working from home also experienced declines in rates of promotion of about 50%. Our evidence provides a natural explanation for this puzzle: though they are less productive, the in-person employees have a leg up in promotions due to their face-to-face interactions with managers.

Our findings have implications for policies aimed at reducing gender gaps in pay and leadership. Companies may be able to curb favoritism by changing their promotion review systems. For example, involving multiple managers in promotion decisions may make it more difficult for employees, male or female, to socialize their way into promotions. Companies also could standardize the review process to use objective indicators, such as revenue generated and hours worked. Another strategy to curb these gender gaps may be leveling the opportunities for employees to socialize and connect with their managers. For example, companies could promote gender-neutral social activities. While 81% of women say that they feel excluded from relationship-building at work, 92% of men believe that they are not excluding women (Gray and Barbara, 2013). This gap reveals a large blind spot. In light of this evidence, firms could try raising awareness of the differential access to social interactions at work. If male employees and male managers were aware of the gender disparity in socialization, some of them might want to take action to fix it. Testing the effects of these policy proposals offers avenues for future research.

References

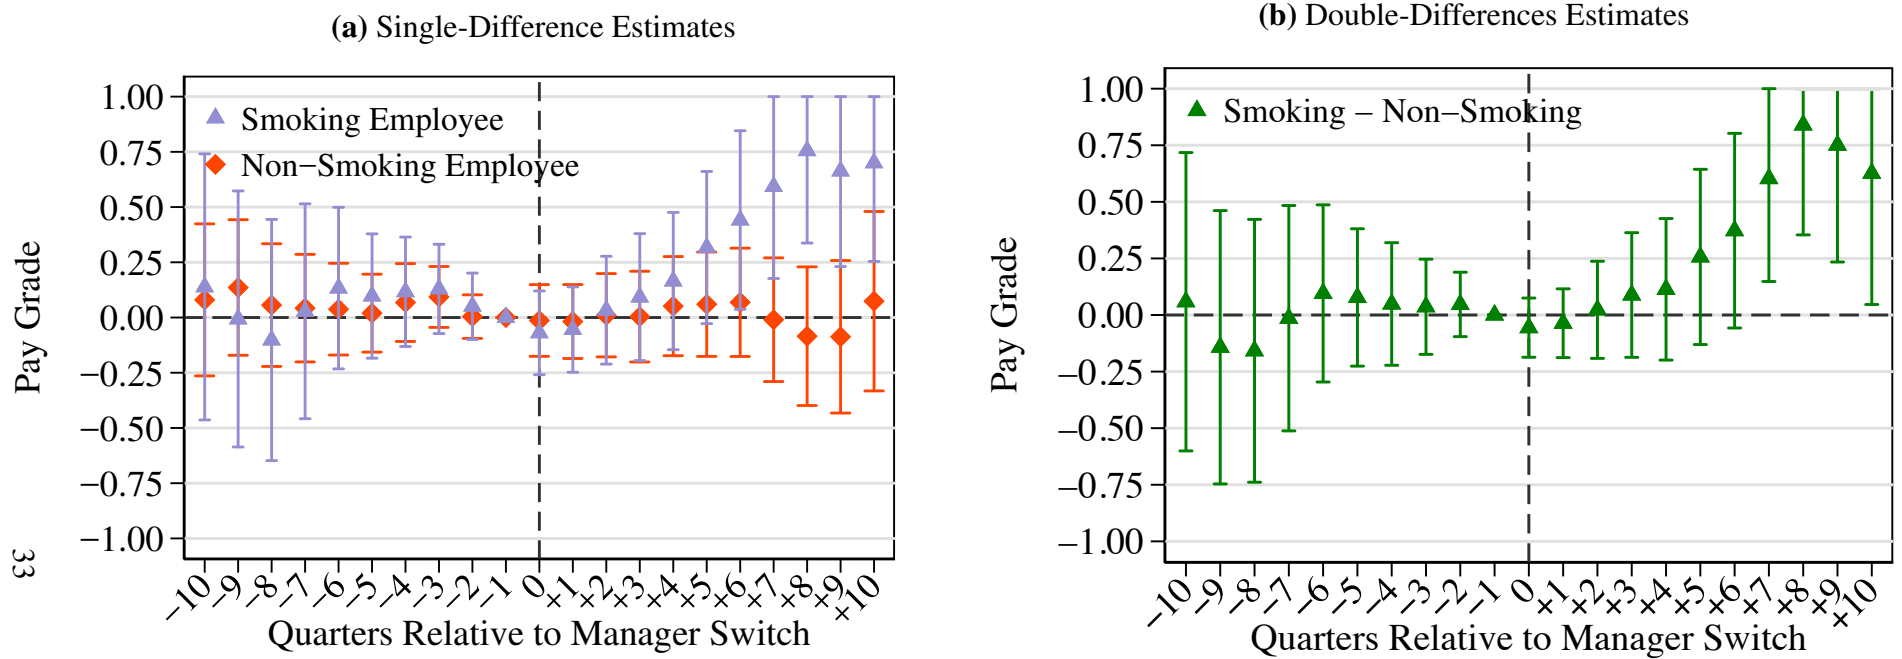
- Agarwal, S., W. Qian, D. M. Reeb, and T. F. Sing (2016). Playing the boys game: Golf buddies and board diversity. *American Economic Review* 106(5), 272–276.
- Alesina, A., P. Giuliano, and N. Nunn (2013). On the Origins of Gender Roles: Women and the Plough. *Quarterly Journal of Economics* 128(2), 469–530.
- Azmat, G. and R. Ferrer (2017). Gender Gaps in Performance: Evidence from Young Lawyers. *Journal of Political Economy* 125(5), 1306–1355.
- Babcock, L., M. P. Recalde, and L. Vesterlund (2017). Gender Differences in the Allocation of Low-Promotability Tasks: The Role of Backlash. *American Economic Review, Papers and Proceedings* 107(5), 131–135.
- Bagues, M., M. Sylos-Labini, and N. Zinovyeva (2017). Does the Gender Composition of Scientific Committees Matter? *American Economic Review* 107(4), 1207–1238.
- Bandiera, O., I. Barankaray, and I. Rasul (2010). Social Incentives in the Workplace. *Review of Economic Studies* 77(2), 417–458.
- Bandiera, O., I. Barankay, and I. Rasul (2009). Social Connections and Incentives in the Workplace: Evidence From Personnel Data. *Econometrica* 77(4), 1047–1094.
- Bell, L. A. (2005). Women-Led Firms and the Gender Gap in Top Executive Jobs. *IZA Discussion Paper No. 1689*.

- Bertrand, M., S. E. Black, S. Jensen, and A. Lleras-Muney (2019). Breaking the Glass Ceiling? The Effect of Board Quotas on Female Labour Market Outcomes in Norway. *Review of Economic Studies* 86(1), 191–239.
- Bertrand, M., C. Goldin, and L. F. Katz (2010). Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors. *American Economic Journal: Applied Economics* 2(3), 228–255.
- Bloom, N., J. Liang, J. Roberts, and Z. J. Ying (2014, 11). Does Working from Home Work? Evidence from a Chinese Experiment. *The Quarterly Journal of Economics* 130(1), 165–218.
- Brogaard, J., J. Engelberg, and C. A. Parsons (2014). Networks and productivity: Causal evidence from editor rotations. *Journal of Financial Economics* 111(1), 251–270.
- Bronson, M. A. and P. S. Thoursie (2020). The Wage Growth and Within-Firm Mobility of Men and Women: New Evidence and Theory. *Working Paper*.
- Bursztyn, L., T. Fujiwara, and A. Pallais (2017). Acting Wife: Marriage Market Incentives. *American Economic Review* 107(11), 3288–3319.
- Cai, J. and A. Szeidl (2018). Interfirm relationships and business performance. *Quarterly Journal of Economics* 133(3), 1229–1282.
- Card, D., S. Dellavigna, P. Funk, and N. Iriberry (2020). Are Referees and Editors in Economics Gender Neutral? *The Quarterly Journal of Economics* 135(1), 269–327.
- Cardoso, A. R. and R. Winter-Ebmer (2010). Female-Led Firms and Gender Wage Policies. *Industrial and Labor Relations Review* 64(1), 143–163.
- Coffman, K., C. B. Flikkema, and O. Shurchkov (2021). Gender Stereotypes in Deliberation and Team Decisions. *Games and Economic Behavior* 129, 329–349.
- Cortés, P. and J. Pan (2019). When time binds: Substitutes for household production, returns to working long hours, and the skilled gender wage gap. *Journal of Labor Economics* 37(2), 351–398.
- Cruz, C. and C. Tolentino (2019). Gender, Social Recognition, and Political Influence. *Working Paper*.
- Cubas, G., C. Juhn, and P. Silos (2019). Coordinated Work Schedules and the Gender Wage Gap. *NBER Working Paper No. 26548*.
- Cullen, Z. and R. Perez-Truglia (2022). How Much Does Your Boss Make? The Effects of Salary Comparisons. *Journal of Political Economy* 130(3), 766–822.
- Dahl, G. B., A. Kotsadam, and D.-O. Rooth (2021). Does Integration Change Gender Attitudes? The Effect of Randomly Assigning Women to Traditionally Male Teams. *Quarterly Journal of Economics* 136(2), 987–1030.
- Dalvit, N., A. Patel, and J. Tan (2022). Intra-firm Hierarchies and Gender Gaps. *Labour Economics*.
- Elting, L. (2018). How To Navigate A Boys’ Club Culture. *Forbes*, July 27 2018.
- Exley, C. and J. Kessler (2022). The Gender Gap in Self-Promotion. *Quarterly Journal of Economics*.
- Field, E., S. Jayachandran, R. Pande, and N. Rigol (2016). Friendship at Work: Can Peer Effects Catalyze Female Entrepreneurship? *American Economic Journal: Economic Policy* 8(2), 125–153.
- Flabbi, L., M. Macis, A. Moro, and F. Schivardi (2019). Do Female Executives Make a Difference? The Impact of Female Leadership on Gender Gaps and Firm Performance. *The Economic Journal* 129(622), 2390–2423.
- Giuliano, L., D. I. Levine, and J. Leonard (2011). Racial Bias in the Manager-Employee Relationship. *Journal of Human Resources* 46(1), 26–52.

- Glover, D., A. Pallais, and W. Pariente (2017). Discrimination as a Self-Fulfilling Prophecy: Evidence from French Grocery Store. *Quarterly Journal of Economics* 132(3), 1219–1260.
- Goldin, C. (2014). A Grand Gender Convergence: Its Last Chapter. *American Economic Review* 104(4), 1091–1119.
- Goldin, C., S. P. Kerr, C. Olivetti, and E. Barth (2017). The expanding gender earnings gap: Evidence from the LEHD-2000 census. *American Economic Review* 107(5), 110–114.
- Gray, J. and A. Barbara (2013). *Work With Me: Eight Blind Spots Between Men and Women in Business*. Palgrave Macmillan.
- Grissom, J. A., J. Nicholson-Crotty, and L. Keiser (2012). Does My Boss’s Gender Matter? Explaining Job Satisfaction and Employee Turnover in the Public Sector. *Journal of Public Administration Research and Theory* 22(4), 649–673.
- Hastings, J. M., R. Kessler, and J. Shapiro (2021). The Effect of SNAP on the Composition of Purchased Foods: Evidence and Implications. *American Economic Journal: Economic Policy* 13(3), 277–315.
- Hill, A. J. (2017). The Positive Influence of Female College Students on their Male Peers. *Labour Economics* 44(7), 151–160.
- Hjort, J. (2014). Ethnic Divisions and Production in Firms. *Quarterly Journal of Economics* 4(4), 1899–1946.
- Hoffman, M. and S. Tadelis (2021). People Management Skills, Employee Attrition, and Manager Rewards: An Empirical Analysis. *Journal of Political Economy* 129(1), 243–285.
- Hsieh, C.-T., E. Hurst, C. I. Jones, and P. J. Klenow (2019). The Allocation of Talent and U.S. Economic Growth. *Econometrica* 87(5), 1439–1474.
- Husain, A., D. A. Matsa, and A. R. Miller (2022). Do Male Workers Prefer Male Leaders? An Analysis of Principals’ Effects on Teacher Retention. *Journal of Human Resources*.
- Isaksson, S. (2019). It Takes Two: Gender Differences in Group Work. *Working Paper*.
- Jayachandran, S. (2021). Social Norms as a Barrier to Women’s Employment in Developing Countries. *IMF Economic Review* 69(3), 576–595.
- Juhn, C. and Y. Rubinstein (2020). Task Requirements, Hours Requirements, and the Gender Gap in Jobs and Pay. *Working Paper*.
- Kim, M. (2020). Opening the door for others? Female leadership and gender disparities in academia. *Working Paper*.
- Kleven, H., C. Landais, and J. E. Sogaard (2019). Children and Gender Inequality: Evidence from Denmark. *American Economic Journal: Applied Economics* 11(4), 121.
- Kunze, A. and A. R. Miller (2017). Women Helping Women? Evidence from Private Sector Data on Workplace Hierarchies. *The Review of Economics and Statistics* 99(5), 769–775.
- Kuziemko, I., J. Pan, J. Shen, and E. Washington (2018). The Mommy Effect: Do Women Anticipate the Employment Effects of Motherhood? *NBER Working Paper No. 24740*.
- Lang, I. H. (2011). Co-Opt the Old Boys’ Club. *Harvard Business Review*, November 2 2011.
- Lee, A. S. (2014). How to break up the old boys’ club in your office. *Quartz*, April 8 2014.
- Lehmann, J.-Y. (2013). Job Assignment and Promotion Under Statistical Discrimination: Evidence from the Early Careers of Lawyers. *Working Paper*.
- Lerner, J. and U. Malmendier (2013). With a little help from my (random) friends: Success and failure in

- post-business school entrepreneurship. *Review of Financial Studies* 26(10), 2411–2452.
- Lleras-Muney, A., M. Miller, S. Sheng, and V. Sovero (2019). Party On: The Labor Market Returns to Social Networks and Socializing. *NBER Working Paper No. 27337*.
- Manning, A. and J. Swaffield (2008). The Gender Gap in Early-Career Wage Growth. *Economic Journal* 118(530), 983–1024.
- Mas, A. and E. Moretti (2009). Peers at Work. *American Economic Review* 99(1), 112–145.
- McKinsey & Company (2019). Women in the Workplace. Technical report.
- Mengel, F. (2020). Gender Differences in Networking. *Economic Journal* 130(630), 1842–1873.
- Sarsons, H. (2017). Interpreting Signals in the Labor Market: Evidence from Medical Referrals. *Working Paper*.
- Sarsons, H., K. Gërxhani, E. Reuben, and A. Schram (2021). Gender Differences in Recognition for Group Work. *Journal of Political Economy* 129(1), 101–147.
- Schönberg, U. and J. Ludsteck (2014). Expansions in maternity leave coverage and mothers' labor market outcomes after childbirth. *Journal of Labor Economics* 32(3), 469–505.
- Shue, K. (2013). Executive networks and firm policies: Evidence from the random assignment of MBA peers. *Review of Financial Studies* 26(6), 1401–1442.
- Stevenson, B. and J. Wolfers (2006). Bargaining in the Shadow of the Law: Divorce Laws and Family Distress. *The Quarterly Journal of Economics* 121(1), 267–288.
- Stoddard, O., C. Karpowitz, and J. Preece (2020). Strength in Numbers: A Field Experiment in Gender, Influence, and Group Dynamics. *IZA Discussion Paper No. 13741*, 1–80.
- Voth, J. and G. Xu (2021). Patronage for Productivity: Selection and Performance in the Age of Sail. *Working Paper*.
- Wasserman, M. (2022). Hours Constraints, Occupational Choice, and Gender: Evidence from Medical Residents. *Review of Economic Studies*.
- Xu, G. (2018). The Costs of Patronage: Evidence from the British Empire. *American Economic Review* 108(11), 3170–3198.
- Yildirmaz, A., C. Ryan, and J. Nezaj (2019). 2019 State of the Workforce Report: Pay, Promotions and Retention. Technical report, ADP Research Institute.

Figure 1: Smoker-to-Smoker Advantage in Pay Grades




Notes: See Section 3.1 for details about the regression specification. All coefficients are estimated from the same regression that includes 94,728 observations of 2,907 employees (966 Smoking & 1,941 Non-Smoking). The dependent variable is the pay grade of the employee. 912 employees (275 Smoking & 637 Non-Smoking) experience events. There are 287 transitions from a non-smoking manager to a smoking manager and 939 from a non-smoking manager to another non-smoking manager. The estimates shown in the graph are based on the coefficients of the event-study variables. In panel (a), the orange diamonds correspond to the coefficient for non-smoking employees, while the purple triangles correspond to the coefficients for smoking employees. Panel (a) corresponds to the difference between transitions from a non-smoking manager to a smoking manager versus transitions from an non-smoking manager to another non-smoking manager. In panel (b), the green triangles correspond to the difference between the coefficient for smoking employees and non-smoking employees. The estimates shown in Panel (b) are the double-differences estimates $(\beta_{N2S}^S - \beta_{N2N}^S) - (\beta_{N2S}^N - \beta_{N2N}^N)$. The 95% confidence intervals are presented in brackets, with two-way clustering by manager and employee. Results are based on the subset of male employees and male managers for which smoking status is available. The coefficient for period “0” corresponds to the exact month of the transition.

Figure 2: Smoker-to-Smoker Advantage in Social Interactions with the Manager


(a) Non-Smoking to Smoking Mgr.
minus Non-Smoking to Non-Smoking Mgr.




(b) Non-Smoking to Smoking Mgr.
minus Non-Smoking to Non-Smoking Mgr.



(c) Female to Male Manager
minus Female to Female Manager

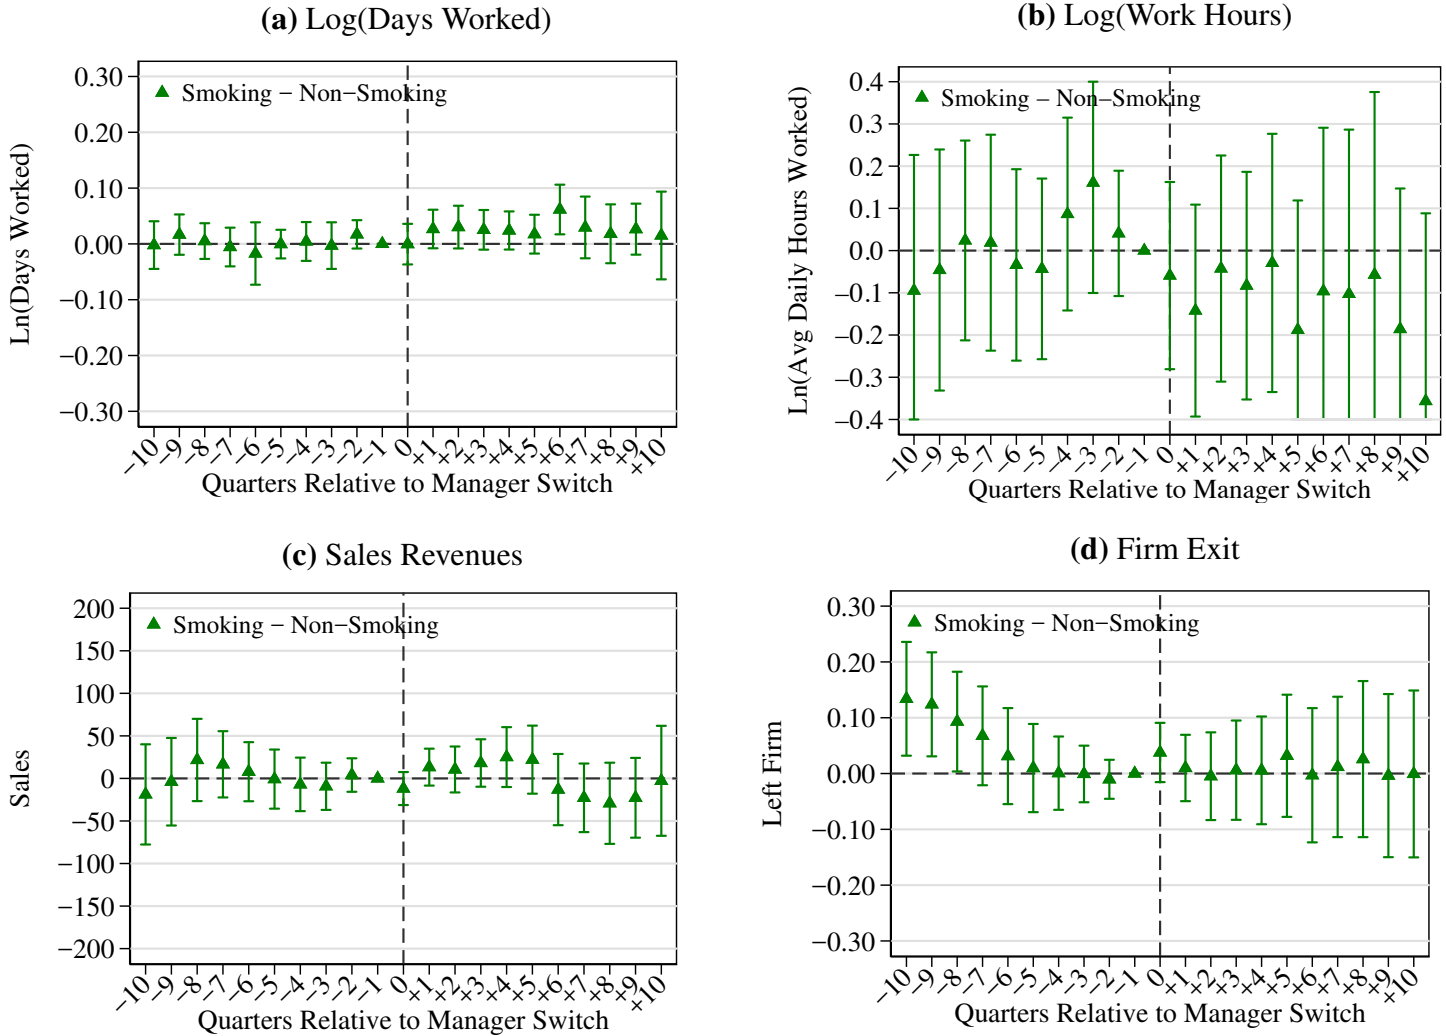


(d) Female to Male Manager
minus Female to Female Manager



Notes: Regression results with the share of breaks. See Section 3.3 for the full econometric specification. The baseline corresponds to the average in the control group: e.g., in panel (a) the baseline is the mean outcome among smoking employees who transitioned from a non-smoking manager to a non-smoking manager. Panels (a) and (b): These two panels are based on one regression which includes 1,287 observations of 699 workers (176 smoker & 523 Non-smoker). 193 employees (50 Smoking & 143 Non-Smoking) of these workers experience a transition event. There are 49 transitions from a non-smoking manager to a smoking manager and 157 from a non-smoking manager to another non-smoking manager. The within-individual standard deviation of this outcome is 0.174. Results for these panels are based on the subset of male employees and male managers for which smoking status is available. Panel (c) and (d): These two panels are based on one regression which includes 4,843 observations of 2,638 workers (698 Male & 1,940 Female). 411 employees (82 Male & 329 Female) of these workers experience a transition event. There are 235 transitions from a female manager to a male manager and 241 from a female manager to another female manager. The 95% confidence intervals are presented in brackets, with two-way clustering by manager and employee.


Figure 3: Smoker-to-Smoker Advantage in Effort, Performance and Retention




Notes: See Section 3.1 for details about the regression specification. In panel (a), the dependent variable is the logarithm of the total number of days worked in the month (inferred from data on approved leaves of absence). The within-employee standard deviation of the dependent variable is 0.123. All coefficients were estimated from a single regression including 89,223 observations of 2,769 employees. In panel (b), the dependent variable is the logarithm of the average number of hours worked in a given month (inferred from data on swipes in and out of the building, and available for headquarter employees only). The within-employee standard deviation of the dependent variable is 0.255. 95 CI are trimmed at -0.4 and 0.4. All coefficients were estimated from a single regression including 33,512 observations of 1,480 employees. In panel (c) the dependent variable is the sales revenue (available for employees with sales roles only) normalized to have mean 100. The within-employee standard deviation of the dependent variable is 82.6. All coefficients were estimated from a single regression including 89,863 observations of 3,195 employees. In panel (d), the dependent variable is an indicator that takes the value 1 in every month after the employee left the firm (these results include additional events after the employees left the firm). The within-employee standard deviation of the dependent variable is 0.184. All coefficients were estimated from a single regression including 114,679 observations of 3,006 employees. In all panels, The 95% confidence intervals are presented in brackets, with two-way clustering by manager and employee. Results are based on the subset of male employees and male managers for which smoking status is available. The coefficient for period “0” corresponds to the exact month of the transition.

Figure 4: Pay Grade Advantage from Having a Shared Trait with the Manager

(a) No Shared Traits to Shared Traits *minus*
No Shared Traits to No Shared Traits

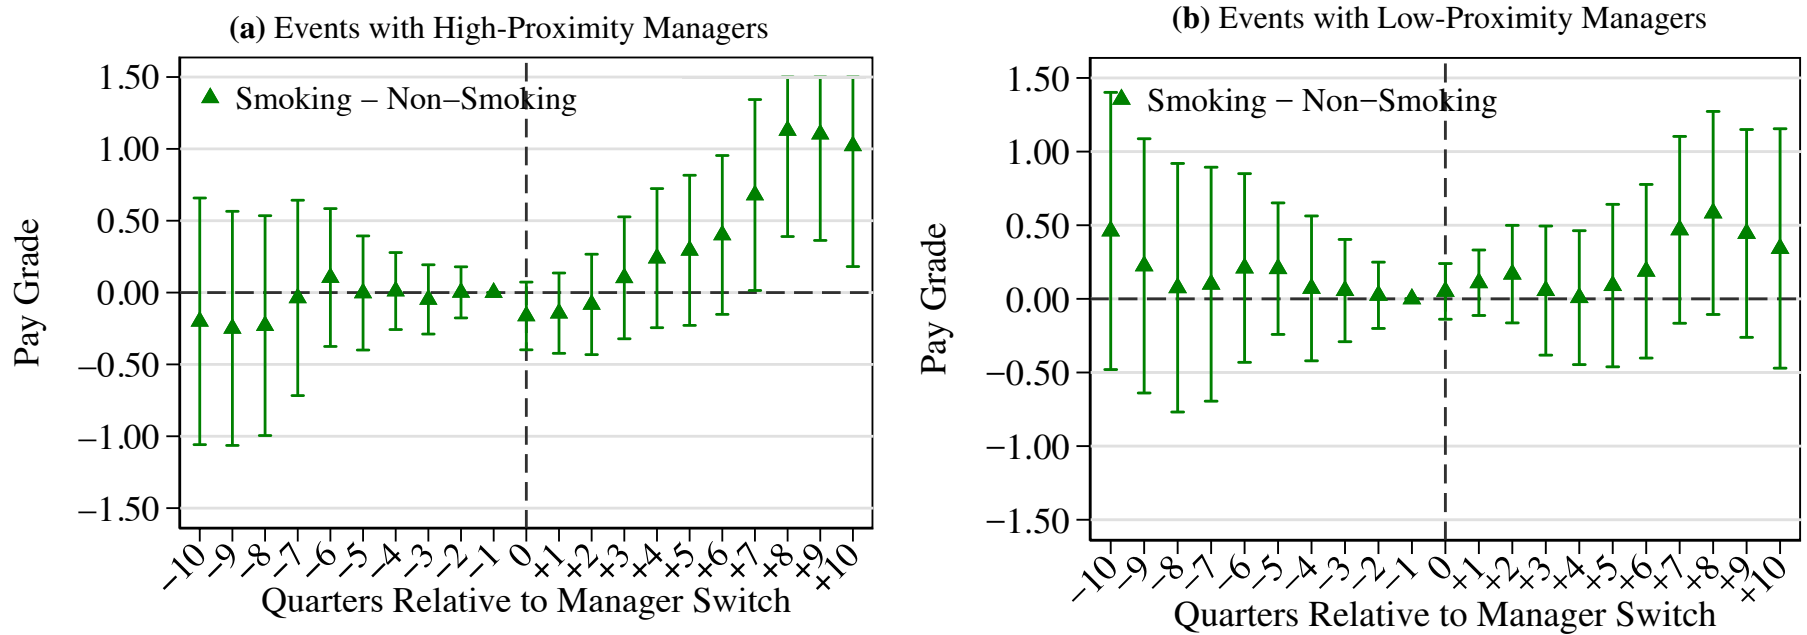


(b) Shared Traits to No Shared Traits *minus*
Shared Traits to Shared Traits



Notes: All coefficients estimated from a regression with 380,126 observations of 14,635 employees. Panel (a) 3,818 employees experience events; 1,902 transition from a manager with whom they have no traits in common to one with whom they share at least one trait, and 2,257 transition between two managers with whom they have no traits in common. Panel (b) 3,013 employees experience events; 1,150 transition from a manager with whom they share at least one trait one with whom they have no traits in common, and 2,044 transition between two managers with whom they share at least one trait. The 95% confidence intervals are presented in brackets, with two-way clustering by manager and employee. Results are based on the subset of male employees and male managers for which smoking status is available. The coefficient for period “0” corresponds to the exact month of the transition.

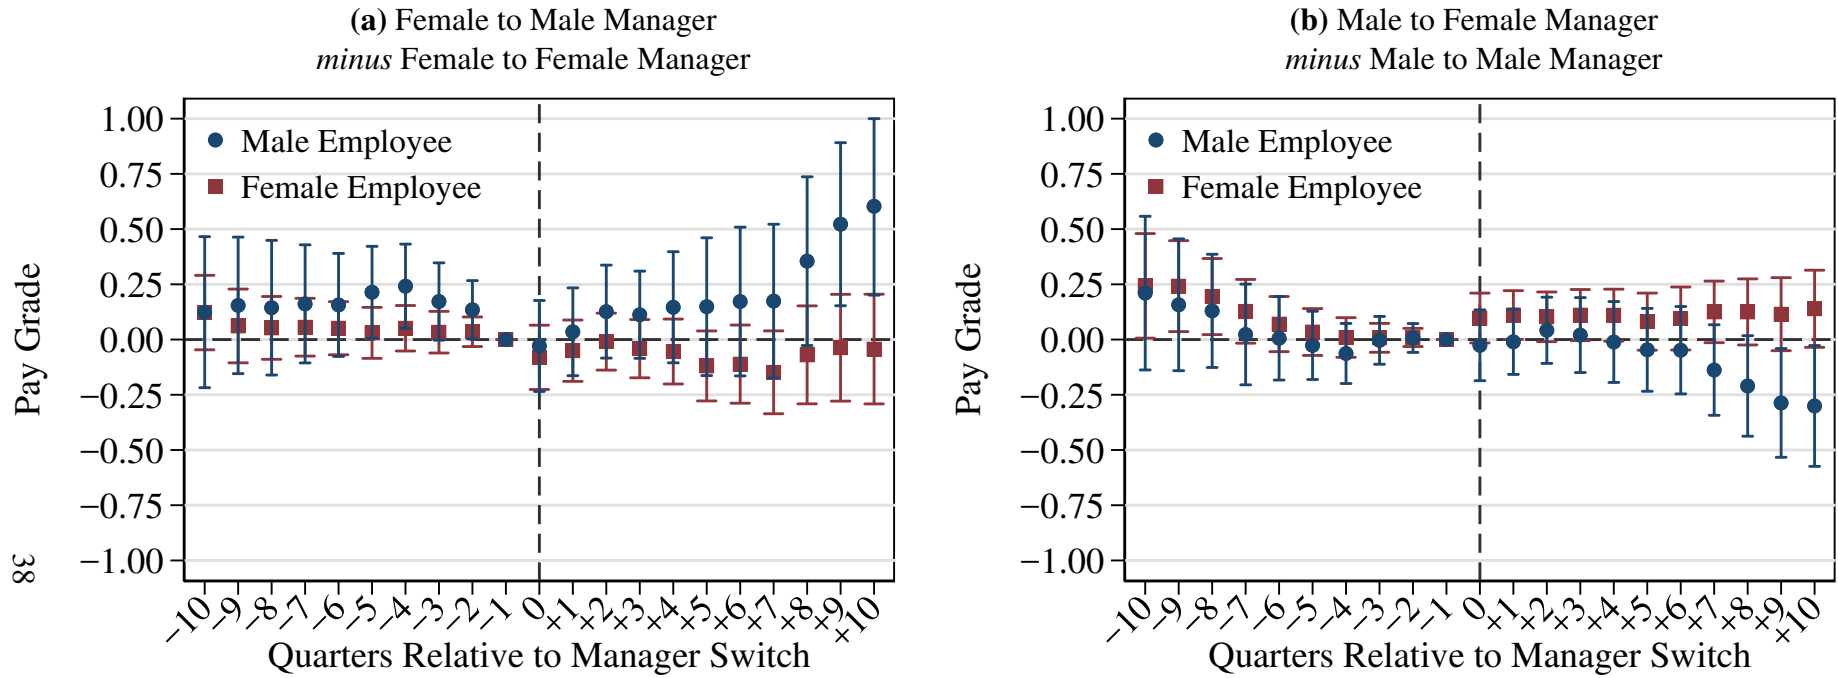
Figure 5: Smoker-to-Smoker Advantage in Pay Grades: Heterogeneity by Proximity to the Manager



37


Notes: See Section 3.1 for details about the regression specification. These results use the specification reported in panel (a) of Figure 1, based on the four types of gender transitions. The only difference is that we split the events in two subsets: high and low proximity events, based on whether the position of the employee in the month of the event was of high or low proximity to the manager. All coefficients were estimated from a single regression including 88,373 observations of 2,829 employees (947 Smoking & 1,882 Non-Smoking). The high-proximity events (panel (a)) affect 395 employees (100 Smoking & 295 Non-Smoking, with 138 transitions from a non-smoking manager to a smoking manager and 351 from a non-smoking manager to another non-smoking manager). The low-proximity events (panel (b)) affect 510 employees (161 Smoking & 349 Non-Smoking), with 133 transitions from a non-smoking manager to a smoking manager and 560 from a non-smoking manager to another non-smoking manager. The 95% confidence intervals are presented in brackets, with two-way clustering by manager and employee. Results are based on the subset of male employees and male managers for which smoking status is available. The coefficient for period “0” corresponds to the exact month of the transition.

Figure 6: Male-to-Male Advantage in Pay Grades: Single-Difference Estimates



Notes: See Section 3.1 for details about the regression specification. Each panel plots single-difference estimates. Panel (a) plots $\beta_{F2M,e}^g - \beta_{F2F,e}^g$ and panel (b) plots $\beta_{M2F,e}^g - \beta_{M2M,e}^g$, where $g \in \{\text{Male, Female}\}$ indexes the gender of the employee and the subscript indexes the transition event type and time since the event. All coefficients were estimated from a single regression including 380,959 observations of 14,638 employees (5,193 Male & 9,445 Female). Panel (a) corresponds to the difference between transitions from a female manager to a male manager and transitions from a female manager to another female manager. 2,712 employees (729 Male & 1,983 Female) experience events: 1,417 transitions from a female manager to a male manager and 1,916 from a female manager to another female manager. Panel (b) corresponds to the difference between transitions from a male manager to a female manager and transitions from a male manager to another male manager. 4,157 employees (1,309 Male & 2,848 Female) experience events: 1,571 transitions from a male manager to a female manager and 3,766 from a male manager to another male manager. The 95% confidence intervals are presented in brackets, with two-way clustering by manager and employee. The within-employee standard deviation of the dependent variable is 0.475. The coefficient for period “0” corresponds to the exact month of the transition.

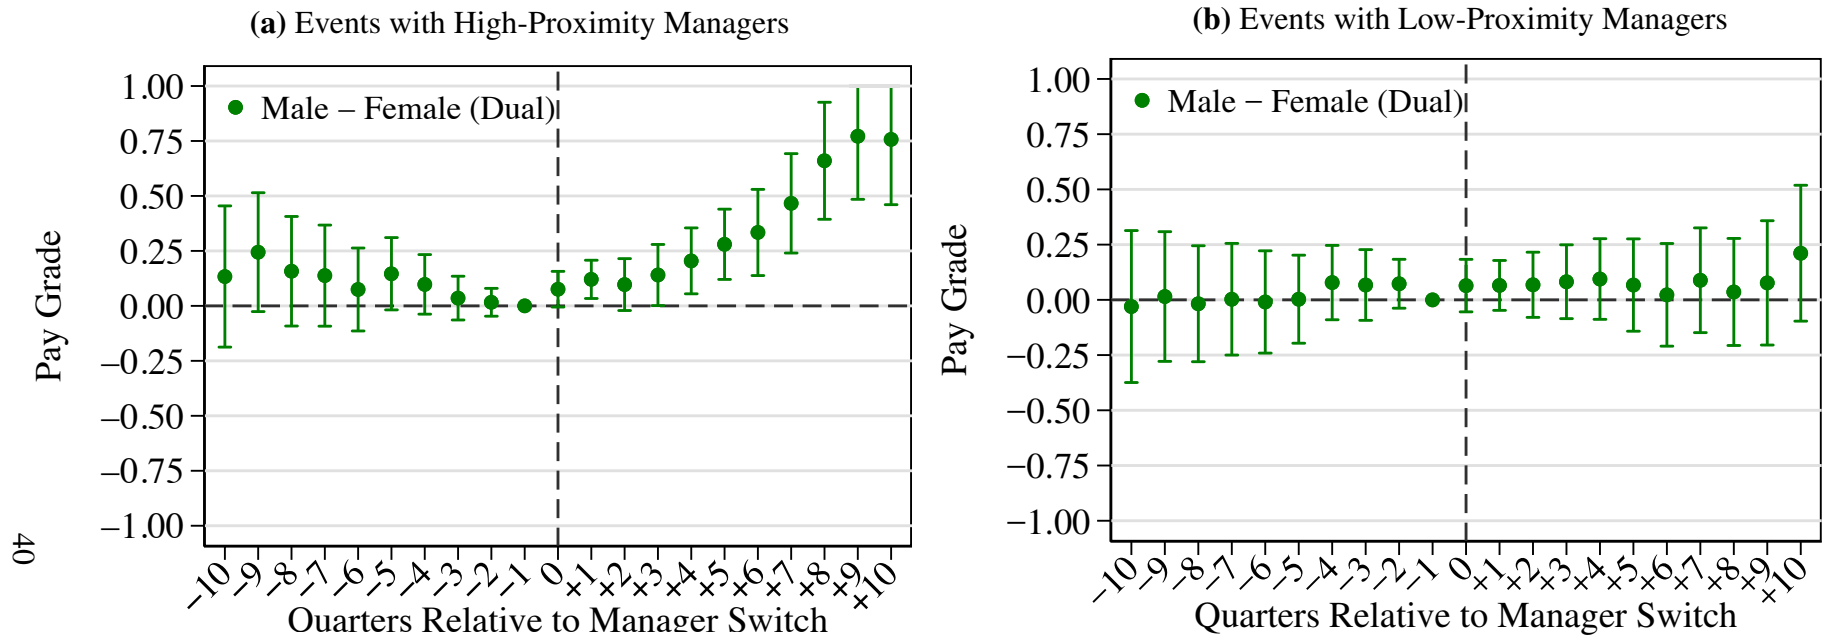
Figure 7: Male-to-Male Advantage in Pay Grades: Double-Differences Estimates



39

Notes: See Section 3.1 for details about the regression specification. All coefficients are estimated from the same regression that includes 380,964 observations of 14,638 workers (5,193 Male & 9,445 Female). The dependent variable is the pay grade of the employee. The estimates shown in the graph are based on the coefficients of the event-study variables. The coefficients shown in panel (a) correspond to the double-differences $(\beta_{F2M,t}^M - \beta_{F2F,t}^M) - (\beta_{F2M,t}^F - \beta_{F2F,t}^F)$ where β^M and β^F are effects for male and female workers, respectively and $F2M, F2F$ are manager transition events from female to male managers and from one female manager to another, respectively. Panel (b) is equivalent to panel (a), but based on the comparison between transitions from a male manager to a female manager and from a male manager to another male manager: $(\beta_{M2F,t}^M - \beta_{M2M,t}^M) - (\beta_{M2F,t}^F - \beta_{M2M,t}^F)$. Panel (c) corresponds to the average between the coefficients from panel (a) and the (negative value of) the coefficients from panel (b). This “symmetric” double-differences estimates is then $\frac{1}{2}\{(\beta_{F2M,t}^M - \beta_{F2F,t}^M) - (\beta_{F2M,t}^F - \beta_{F2F,t}^F) - [(\beta_{M2F,t}^M - \beta_{M2M,t}^M) - (\beta_{M2F,t}^F - \beta_{M2M,t}^F)]\}$. The 95% confidence intervals are presented in brackets, with two-way clustering by manager and employee. The coefficient for period “0” corresponds to the exact month of the transition.

Figure 8: Male-to-Male Advantage in Pay Grades: Heterogeneity by Proximity to the Manager



Notes: See Section 3.1 for details about the regression specification. Each panel plots dual-double-differences estimates. These results use the symmetric specification reported in panel (c) of Figure 7, based on the four types of gender transitions. The only difference is that we split the events in two subsets: high and low proximity events, based on whether the position of the employee in the month of the event was of high or low proximity to the manager. All coefficients are estimated from the same regression with 360,239 observations of 13,814 employees (4,912 Male & 8,902 Female). The high-proximity events (panel (a)) affect 2,983 employees (1,043 Male & 1,940 Female), with 617 transitions from a female manager to a male manager (F2M): 1,075 F2F, 754 M2F, 1,508 M2M. The low-proximity events (panel (b)) affect 3,056 employees (783 Male & 2,273 Female), with 762 transitions from a female manager to a male manager (F2M): 751 F2F, 742 M2F, 2,182 M2M. The 95% confidence intervals are presented in brackets, with two-way clustering by manager and employee. The coefficient for period “0” corresponds to the exact month of the transition.