

The Old Boys' Club

Schmoozing and the Gender Gap

Zoë Cullen*

Ricardo Perez-Truglia

Harvard Business School

University of California, Los Angeles

This Draft: December 2019.

Abstract

The *old boys' club* refers to the alleged advantage that male employees have over their female counterparts in interacting with powerful men. For example, male employees may schmooze with their managers in ways that female employees cannot. We study this phenomenon using data from a large financial institution. We use an event study analysis of manager rotation to estimate the causal effect of managers' gender on their employees' career progression. We find that when male employees are assigned to male managers, they are promoted faster in the following years than they would have been if they were assigned to female managers. Female employees, on the contrary, have the same career progression regardless of the manager's gender. These differences in career progression cannot be explained by differences in effort or output. This male-to-male advantage can explain a third of the gender gap in promotions. Moreover, we provide suggestive evidence that these manager effects are due to socialization between male employees and male managers. We show that these manager effects are present only if the employee works in close proximity to the manager. We use survey data to show that, after transitioning to a male manager, male employees spend more time with their managers. Finally, we study a shock to socialization within males, based on the anecdotal evidence that employees who smoke tend to spend more time together. We find that when male employees who smoke switch to male managers who smoke, they spend more of their breaks with their managers and are promoted faster in the following years. Moreover, the effects of these smoking manager switches are similar in timing and magnitude to the effects of the gender manager switches.

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

Keywords: gender, career, promotion, socialization, networking.

*Corresponding author: zcullen@hbs.edu, Rock Center 310, Boston, MA 02163. We are thankful for excellent comments from several colleagues and seminar discussants at USC, UCSD, Berkeley-Haas, HBS and UBC. 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 and Anh Nguyen provided excellent research assistance.

1 Introduction

Women have a harder time than men climbing the corporate ladder. Among U.S. corporations, 48% of entry-level employees are women, but the female representation falls to 38% at middle-management, 22% at the C-Suite level and 5% at the CEO level (McKinsey & Company, 2019). The improvement over the last several decades has been agonizingly slow. Not only is this unfair, it is inefficient; the economy is missing out on women who would make great managers (Hsieh et al., 2019).

The lack of female representation is due to multiple factors. We focus on the hypothesis that male employees enjoy the privileges of the *old boys' club*.¹ In the business context, this idiom refers to male employees having an advantage over their female counterparts in their relationships with male managers. For example, men can schmooze, network, and interact with more powerful men in ways that are less accessible to women. This mechanism can create a self-perpetuating cycle: male managers will promote a disproportionate share of male employees, who will continue promoting other men. Ample anecdotal evidence suggests that men can provide and take advantage of their male relationships in the corporate world in ways that women cannot (Lang, 2011; Lee, 2014; Elting, 2018). However, there is little quantitative evidence on the matter. In this study, we provide novel evidence from a natural experiment in a large financial organization.

In an ideal experiment, we would randomize employees to male and female managers and then measure the effects on their career progression in subsequent years. Such an experiment is difficult to implement in practice, however, so we turn to the next best thing: exploiting quasi-experimental variation in manager assignment generated by the rotation of managers within the organization. Our identification strategy relies on the timing of manager switches and on the comparison between different types of switches. For example, consider two teams, each managed by a female manager. One of these teams then switches from the female manager to a male manager, and the other team switches from the female manager to a different female manager. We can compare the outcomes of the male employees each month leading up to the manager switch date and each month after the switch. As both teams are affected by a manager switch, this design nets out the effect of the transition. The old boys' club hypothesis is that, relative to switching to a different female manager, switching to a male manager benefits the careers of the male employees on the team; and the corresponding effects on the careers of female employees should be zero, or at least weaker than for male employees.

¹This term originated from the fact that much of the British elite attended certain public schools as boys. In popular language, however, the term has come to be used in reference to the preservation of social elites in general.

We collaborated with a large commercial bank in Asia that has millions of customers, billions of dollars in assets and in revenues, and thousands of employees. This organization is typical in that female representation drops off at higher levels of the corporation: 75% of entry-level employees are women, which falls to 61% in middle management, 25% at the C-Suite level, and 0% at the CEO and company board levels. Indeed, the gender gaps in pay and promotion rates at this corporation are similar to those documented for other corporations in both developed and developing countries.

We have rich sources of administrative data spanning four years (2015-2018) and 14,736 unique employees, 1,269 of whom had a managerial role at some point. We focus on manager switches that are largely out of the control of the employee. The typical case is a manager rotating laterally to a different team.² Our data comprises 10,101 events involving 6,536 unique employees and 751 unique managers. Events are uniformly distributed across the four years, and they affect employees at every level. Whether the employee has an event and the type of event (e.g., switching from a female to a male manager) are generally uncorrelated to the characteristics of the employees and the managers involved.

We show that male employees are promoted faster after they transition from a female to a male manager: at 10 quarters after such a manager transition, male employees increased their pay grades by an additional 0.53 points (p-value = 0.005), roughly equivalent to 13% higher pay, compared to male employees who transitioned from a female manager to a different female manager. On the contrary, female employees had the same career progression regardless of whether they transitioned from a female manager to a male manager or from a female manager to another female manager. The triple-difference is consistent with a male-to-male advantage: male managers (relative to female managers) improve the career progression of male employees (relative to female employees). The triple-difference in paygrades (0.50 as of 10 quarters after the event, p-value=0.003) is not only highly statistically significant, but also economically large: removing this advantage would reduce the gender gap in pay grades by 38%.³

We provide two main robustness checks for our identification strategy. First, we analyze the opposite types of transitions. In the baseline results, we look at employees who “gain” a male manager (i.e., switching from a female manager to a male manager versus switching from a female manager to a different female manager). In this robustness check, we look at employees who “lose” a male manager (i.e., switching from a male manager to a female

²A team also may be assigned a replacement manager (an external hire or internal rotation) when the current manager leaves for another company.

³The average difference between male and female pay grades is 0.85 points. Recall that we estimate that male employees who transitioned from a female to male manager increased their pay grades by an additional 0.5 points after 10 quarters. Since 66% of managers are male, removing the male-to-male advantage would reduce the pay grade gap by $0.5 \cdot 0.66 = 0.33$ points, which is equivalent to 38.8% of the 0.85 gap.

manager versus switching from a male manager to a different male manager). The expectation is that the effects of gaining a male manager should mirror 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 switch 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 in the opposite direction of the effects of gaining a male manager, but are similar in terms of timing and magnitude. Male employees who transition to a female manager (relative to transitioning to another male manager) end up with a pay grade 0.38 points lower 10 quarters later, whereas the evolution of pay grade for female employees is unrelated to the manager’s gender.

The second robustness test introduces placebo events. We reproduce the whole analysis, but instead of focusing on gender as the relevant characteristic of managers and employees, we focus on a characteristic that we know *ex ante* should not be relevant: whether someone was born on an even or odd date. For example, we compare switches from an odd-birthday manager to an even-birthday manager among odd-birthday and even-birthday employees. As expected, we find that all effects are close to zero, statistically insignificant, and precisely estimated.

Next, we provide evidence about the mechanisms underlying the male-to-male advantage. We start with the role of attrition. When we estimate the effects of the manager switches on the probability of leaving the firm, we find point estimates that are close to zero, statistically insignificant, and precisely estimated. Similarly, we do not find any significant effects on the probability of internal transfers. This evidence rules out the explanation that male managers are better at retaining male employees.

Another potential mechanism is that male employees work harder and are more productive under male managers than they would be under female managers. For example, male managers might be better than female managers at motivating and monitoring male employees. To test this hypothesis, we exploit data on two measures of effort (the number of days worked and the number of hours spent in the office) and one measure of output (sales revenues). Contrary to this mechanism, we find no significant evidence of a male-to-male advantage in any of these measures of effort or output.

The timing of the male-to-male advantage is informative about the plausible mechanisms: only after the first year do we begin to see a gap in the promotion rates between men and women. This gap cannot be entirely due to the spacing between promotion events, as some males became eligible for promotion right after the event. The delayed timing highlights the potential role of slower mechanisms. For example, socialization involves developing affinities through spending time together during breaks or other activities, which requires a meaningful

amount of time to develop. The timing of the male-to-male advantage is less consistent with other mechanisms, such as statistical discrimination or chauvinism, that are more likely to show up immediately after an event, when the first promotions of male employees arise.

We provide three additional pieces of evidence in favor of the socialization mechanism. First, we show that the advantage conferred to male employees by male managers stems from positions where the manager and employee work in close proximity, and this advantage is not significant in positions where the manager and employee often work apart. Second, we use survey data on the socialization between managers and their employees to show that, for a male employee, having a male manager increases the share of breaks taken with that manager. Third, we show that within the sample of male employees and male managers, a shock to socialization increases the subsequent promotion rates.

The socialization shock we study is the transition from non-smoking managers to smoking managers, focusing exclusively on the sample of male employees and male managers. In our context, smokers tend to take smoking breaks together, and thus for an employee who smokes, having a manager who also smokes increases their socialization. We use data on the smoking status of a subsample of employees and managers from the annual health exam and supplement it with survey data. In this sample, 33% of male employees smoke and 37% of male managers smoke. We reproduce the event-study framework on the effects of manager gender, but we focus on smoking status instead of gender.⁴ We show that transitioning from a non-smoking to a smoking manager (relative to transitioning from a non-smoking manager to another non-smoking manager) increases how often the smoking employees socialize with their managers but has no effect on the socialization of non-smoking employees. Then, we show that these manager switches affect promotion rates: transitioning from a non-smoking to a smoking manager (relative to transitioning from a non-smoking to another non-smoking manager) increases the subsequent promotion rates of smoking employees but does not affect the promotions of non-smoking employees. Moreover, the effects of the transitions from non-smoking to smoking managers are similar in magnitude and timing to the corresponding effects of the transitions from female to male managers.

Although we offer evidence from a specific corporation, our methodology is not specific to this organization. Rotation of managers is common across other corporations, so the event-study analysis could be applied to other places. The types of datasets that we use in our analysis, such as position titles and pay grades, should also be available for other corporations. We hope our methods will be applied to data from other companies in different industries and different countries to identify the contexts in which the male-to-male advantage is most

⁴The event-study analysis spans 365 manager switches involving 917 unique employees and 227 unique managers.

pervasive.

This paper is related to various strands of literature. It is of course related to a large literature on the gender wage gap (Goldin, 2014). There is a consensus that the vast majority of this gap is due to differences in promotion rates (Bertrand et al., 2010; Manning and Swaffield, 2008). By one careful account, the gap in internal promotion rates can account for approximately 70% of the gender pay gap by age 45 (Bronson and Thoursie, 2019). Several explanations have been provided for those gaps, such as marriage market incentives (Bursztyn, Thomas, and Pallais, 2017), cultural norms (Bursztyn, Gonzalez, and Yanagizawa-Drott, 2018), recognition for group work (Sarsons, 2019, 2017), differences in effort and performance (Azmat and Ferrer, 2017), and the child penalty (Bertrand et al., 2010; Kleven et al., 2019; Kuziemko et al., 2018).⁵ Our contribution to this literature is twofold. First, we contribute by identifying and quantifying a new channel that, in the organization at hand, can explain over one-third of the gender pay gap. Second, we provide evidence on a mechanism, socialization at work, that has been largely overlooked in the literature.

This study is also related to literature on the importance of managers for the outcomes of employees (Lazear et al., 2015). More specifically, there is a set of studies that measures whether the gender of leaders is important in corporations. Evidence shows that firms with female leadership or female representation on the board are more likely to have female executive representation and better gender wage policies (Bell, 2005; Cardoso and Winter-Ebmer, 2010; Dalvit et al., 2018; Flabbi et al., 2019). Other studies look at the correlation between the gender of managers and the outcomes of male and female employees, such as their salaries and job satisfaction (Grissom, Nicholson-Crotty, and Keiser, 2012; Srivastava and Sherman, 2015; Lucifora and Vigani, 2016; Kunze and Miller, 2017; Halldén, Säve-Söderbergh, and Rosén, 2018; van Hek and van der Lippe, 2019).⁶ We advance this literature in two ways. First, whereas most studies of corporate jobs are based on correlations, we provide causal identification using quasi-experimental methods. Second, the existing studies look at the contemporaneous association between the gender of managers and the salaries of their employees.⁷ Our evidence shows that focusing on these short-term effects misses the development of a gender gap due to differential rates of promotion.

Finally, despite the universality of socializing in the workplace, little is known about

⁵Although large, the child penalty cannot fully account for the gender gaps in career progression and pay. For example, Bronson and Thoursie (2019) show that the promotion gap is significant even before the birth of the first child, and is even present between men and women who never have children.

⁶Other studies look at the gender of peers instead of the gender of managers (Dahl et al., 2018; Hill, 2017). Also, some studies focus on other manager characteristics such as race and friendship links (Mas and Moretti, 2009; Bandiera et al., 2010; Hjort, 2014; Glover et al., 2017; Giuliano et al., 2011).

⁷For example, Grissom et al. (2012) and Kunze and Miller (2017) measure whether male or female managers are associated with higher salaries of male and female employees in that same year or the subsequent year.

the returns to these personal interactions and whether those returns differ by gender. Some evidence in the context of politics suggests that public officials can capitalize on their political and personal networks to gain influence (Cruz and Tolentino, 2019; Bertrand et al., 2018; Voth and Xu, 2019).⁸ We contribute to this literature by providing novel evidence that, even in the corporate world, socializing with superiors may help with career advancement.

The rest of the paper proceeds as follows. Section 2 presents an overview of the research design and our econometric specification. Section 3 presents the institutional context for this study and describes the data. Sections 4 and 5 present our empirical results. Section 6 concludes.

2 Research Design

In this section, we discuss the conceptual framework and present the econometric specifications.

2.1 Switches in Manager’s Gender

While we cannot explicitly randomize employees to male and female managers, we can exploit quasi-experimental variation in manager assignments generated by the rotation of managers within the organization.

Our identification strategy relies on the timing of manager switches, as well as on the comparison between different types of switches. Consider for example two teams that are led by a female manager. One of those teams transitions from a female manager to a male manager, while the other team transitions from a female manager to a different female manager. The old boys’ club prediction is that, relative to switching to the new female manager, switching to the male manager should be beneficial for the subsequent career of the male employees; in contrast, the differential effects of this transition should be absent, or at least less pronounced, for the female employees.

To formalize this framework, we present our econometric specification:

⁸There are some related findings in business contexts. For example, Field et al. (2016) show that going through business training with a female friend increases the likelihood that a female participant engages in future business activity. And Mengel (2015) find that, in a laboratory experiment, there is similar networking by men and women but men promote closer connections.

$$\begin{aligned}
y_{i,t} = & \sum_{j \in J_G} \sum_{s \in S} \beta_{j,s}^F \cdot F_i \cdot D_{i,t+s}^j + \sum_j \sum_s \beta_{j,s}^M \cdot (1 - F_i) \cdot D_{i,t+s}^j \\
& + \gamma_i + \delta_t^F + \delta_t^M + \epsilon_{i,t}
\end{aligned} \tag{1}$$

Let $y_{i,t}$ be a generic outcome, where i denote employees and t denote time. The main outcome in our analysis is the employee's pay grade, but we also consider other outcomes such as effort, output and attrition. Let J_G be the set of types of manager switches $J_G = \{F2M, M2F, F2F, M2M\}$, where $F2M$ denotes a transition from a female manager to a male manager, $F2F$ denotes a transition from a female manager to another female manager, and so on. The variables $D_{i,t+s}^j$'s are the traditional event-study dummies. The dummy $D_{i,t+s}^j$ takes the value 1 if individual i experience an event of type j in period $t + s$. The set S , the event-study window, spans from 30 months before the event to 30 months after the event, with the usual absorbing dummies at extremes of -31 and +31 months (Stevenson and Wolfers, 2006).⁹ The omitted categories in S are the 3 months prior to the event (i.e., -3, -2 and -1 months). In the event-study graphs, we aggregate these monthly coefficients to the quarterly level, for ease of presentation.

We interact event-study lags and leads with a dummy for female employee (F_i) to estimate event time coefficients for men (β_{jt}^M) and women (β_{jt}^F). This baseline specification includes employee fixed effects (γ_i) as well as gender specific month effects, (δ_t^M and δ_t^F). In the analysis we always use two-way clustering of the standard errors at the team and manager levels.

To isolate the impact of a change in manager gender from a change in manager more generally, we focus on the difference between outcomes for employees whose manager changes gender, with the outcomes of those whose manager changes to someone of the same gender. What we capture with the difference-in-difference, $\beta_{F2M,t}^M - \beta_{F2F,t}^M$, is the impact of receiving a male manager relative to the impact of receiving a new female manager. When we see a positive effect of switching from female to male manager for male employees, we detect either a gender bias, or an overall boost that may come with male managers for all employees, male and female alike. By presenting the triple difference, we compare the net effect of the same switches for men and women, $(\beta_{F2M,t}^M - \beta_{F2F,t}^M) - (\beta_{F2M,t}^F - \beta_{F2F,t}^F)$. If this triple-difference is positive, it would be consistent with a male-to-male advantage.

Underlying our interpretation of the event-study coefficients is an important identifying assumption: the relative trajectory of outcomes of employees is orthogonal to the type of manager assigned to the unit. An example of a violation would be if the divergent success

⁹We do not have power to extend our event-study window further due to the length of our panel and frequency of events.

of male employees relative to their female teammates led the female manager to recruit a male manager to take her place. The event-study framework provides a natural test of the identifying assumption: we can look at the evolution of the outcome in each of the months before the date of the switch, to confirm if the trends were truly parallel.

This design of manager switches provides a natural robustness check based on the principle that switches in the opposite direction should result in approximately the opposite effects. More precisely, the comparison described above refers to a situation in which a team is “gaining” a male manager: i.e., switching from a female manager to a male manager versus switching from a female manager to a different female manager. Instead, we could look at teams who are “losing” a male manager: i.e., switching from a male manager to a female manager versus switching from a male manager to a different male manager. The expectation is that the effects of gaining a male manager should be the mirror image of the effects of losing a male manager, in terms of both timing and order of magnitude. Note that since these coefficients are identified by totally disjoint set of switch events, there are no “mechanical” reasons why the results should be a mirror image of each other.

When interpreting the event-study results, there are a few caveats to keep in mind. First, it is important to note that our estimates measure a reduced form effect of an increased but likely transitory exposure to a given managerial gender. As time goes by, there are all sort of reasons why the employee may end up with a manager of a different gender. For example, the employee may be promoted to a different position and be assigned a manager of a different gender as a result, or the employee may move laterally to another team with a manager of a different gender. In this sense, our estimates will under-estimate the effect of the manager’s gender: if the employee were to stay with the new manager gender forever, the effects would presumably be even stronger.¹⁰ Second, this framework cannot distinguish between disentangle if the male managers are favoring male employees or the female managers are hurting male employees. Indeed, this would be true even if we were able to randomize the gender of the manager, and is related to the lack of gender-neutral managers. We will return to this issue in the results section. The last caveat is that our effects may be attributed to the manager, the employee, or both: e.g., male managers may treat male employees differently, or it may be the male employees are the ones who are reacting differently to male managers. For example, it is possible that female managers treat male employees just like the male managers but the male employees are more willing to take orders from the male managers than from the female managers (Giuliano et al., 2011).

¹⁰In Appendix C.3 we quantify the persistence of the gender switches.

2.2 Placebos for Manager's Switches

As a robustness check, we reproduce the analysis but, instead of focusing on gender as the relevant characteristic of managers and employees, we focus on a characteristic that we know ex-ante should not be relevant: whether someone was born on an even or odd date. The specification is identical to the main specification in Equation 1, except that gender is replaced everywhere by the evenness of birth date:

$$y_{i,t} = \sum_{j \in J_E} \sum_{s \in S} \beta_{j,s}^O \cdot O_i \cdot D_{i,t+s}^j + \sum_j \sum_s \beta_{j,s}^E \cdot (1 - O_i) \cdot D_{i,t+s}^j + \gamma_i + \delta_t^E + \delta_t^O + \epsilon_{i,t} \quad (2)$$

Where J_E is the set of manager switches $J_E = \{O2E, E2O, O2O, E2E\}$: $O2E$ denotes a transition from a manager with an odd birthday to one with an even birthday, and so on.

2.3 Effect of Manager's Smoker Status

We also directly evaluate a non-gender shock to socialization. Intuitively, we begin by restricting to the sample of male employees and male managers. We compare two teams who are led by a non-smoking manager. One team transitions to a smoking manager, and the other team transitions to a different non-smoking manager. We compare the differential effects of this transitions for smoking employees and non-smoking employees separately. The prediction is that switching to the smoking manager should benefit the subsequent career of the smoking employees, while it should not affect, or affect less prominently, the careers of the non-smoking employees.

We use a variant of the same specification to identify the smoker events, based on the restricted sample of male employees and male managers. Again, the event study specification is identical to that in Equation 1, except that the gender status is replaced everywhere by the smoker status:

$$y_{i,t} = \sum_{j \in J_S} \sum_{s \in S} \beta_{j,s}^S \cdot S_i \cdot D_{i,t+s}^j + \sum_j \sum_s \beta_{j,s}^N \cdot (1 - S_i) \cdot D_{i,t+s}^j + \gamma_i + \delta_t^S + \delta_t^N + \epsilon_{i,t} \quad (3)$$

Where J_S is the set of the types of manager switches $J_S = \{N2S, S2N, N2N, S2S\}$: $N2S$ denotes a transition from a non-smoker manager to a smoker manager, and so on.

3 Institutional Context and Data

3.1 Institutional Context

We collaborated with a private commercial bank in Asia. To keep the identity of the firm secret, 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.

While we do not want to claim that the evidence is representative of the world, it is still useful to understand whether this context is out of the ordinary. The firm may be unusual for the financial sector in that a majority (65%) of its employees are female. Besides that, however, the gender gaps at this organization are quite average by U.S. standards. The gender pay gap at this firm (26%) is close to the average of similar sized firms in the financial sector in the United States (31%)¹¹. Moreover, the firm is also typical in that female and male employees within a given position get paid about the same. The bulk of the gender pay gap, thus, is due to difference in the positions that female and male employees hold. For example, while 75% of employees at entry-level are female, that fraction falls to 61% at middle management, 25% at C-Suite and 0% at CEO. Sata for U.S. corporations suggest a similar drop, from 48% of female employees at entry-level to 38% at middle management, 22% at C-Suite and 5% at CEO (McKinsey & Company, 2019).

When looking at the firm's country as a whole, the gender gaps are not different from those in the United States either. For example, the gender gap in labor force participation (8.5%) is similar to the one in the United States (13.2%).¹² According to survey data, the gender norms are not out of the ordinary either. For example, data from the 2006 World Value Survey suggest that 12% of women describe work as unimportant or of little importance, while the respective share is 19% in the United States.

3.2 Administrative Data and Outcome Variables

We collaborated with the different units of the organization to create a centralized and anonymous database of every employee in the firm. We constructed a monthly panel spanning four years, from January of 2015 to December of 2018. This panel includes 14,736 unique employees, 1,269 of which are assigned to a manager role at some point. 64% of the employees are female, while 49% of the managers are female.

¹¹Results based on wage rates for men and women working in the financial sector in firms with over 1,000 employees reported in Yildirmaz et al. (2019).

¹²Labor force participation data come from the World Bank Databank, and the International Labour Organization ILOSTAT database. These are 2017 figures, as that is the most recent year for which male and female labor force participation data are available in both countries.

Our main outcome variable is pay grade. This outcome ranges from 41 to 66, and is the best measure of the vertical career progression in the organization. Indeed, employees commonly use pay grades as a measure of their rank in the firm in conversations with other employees. An increase in pay grade is typically associated with a promotion: conditional on an increase in pay grade, there is an 84% chance of a change in the position title; in comparison, there is less than 1% chance of a change in position title when there is no pay grade increase. Variation in the pay grade outcome suggests that, consistent with anecdotal evidence, there is ample opportunity for upward mobility in the firm. Among the 7,622 employees who worked at the bank for the full period of four years, 50% increase their pay grade at least once and 16% increase it more than once.

Due to the sensitive nature of the data, we do not have the exact compensation details for the employees for the whole sample. However, for a different project on a different topic (Cullen and Perez-Truglia, 2018), we have a cross-section on the pay grade and the base salary of the employees at a given month (March of 2017). We can use that data to show that the pay grade is highly correlated to the base salary. The results are presented in Figure 1. Note that there is a robust, linear relationship between the logarithm of salary and the pay grade. The slope of the relationship (0.227) indicates that a 1-point increase in pay grade is associated to a 25% increase in salary ($= e^{0.227} - 1$).

We also have some measures of effort and output. The first measure of effort is the number of days worked. We construct this measure using data from the Human Resources divisions on approved absences. We subtract the number of approved leave days (i.e. parental leave, sick days, vacation days, etc) from the total number of work days in the month. In this measure, we are unable to observe unapproved absences; if an employee simply doesn't show up to work and the absence is never reported to HR, we would not observe this absence in the administrative data. We have an additional measure of effort that we can use to complement the administrative data: the number of hours spent in the office. We can only measure this outcome for employees working in the headquarter offices (29% of the sample). Those employees must clock-in and -out from the office using an electronic card-swipe system, which is strictly enforced by security personnel. We use these time stamps to calculate the average number of *hours in the office*. Finally, our measure of output is based on the 38% of employees who have a sales role. We measure sales performance based on their sales revenues. The bank uses an official formula to aggregate an employee's sales over all products (e.g., credit cards, loans, mortgages). We use this data to construct a *sales performance index* on a monthly basis.

3.3 Manager Assignments

An employee’s manager is another employee from the same firm who has a higher pay grade. The manager is typically in a position the employee would like to be promoted to in the future. The manager is someone that the employee may need to report to daily, and someone who has power over the employee. While a single employee may have more than one person that he or she may call manager, we are interested in identifying the most relevant manager: i.e., the one who has the most power over the employee.

We use longitudinal data from the firm’s organizational chart in order to link each employee to a manager in each month the employee appears in the sample. The employee-manager assignment was constructed using a simple algorithm that boils down to two steps: identifying the team that the employee is a member of, and then identifying the “director” of that team.¹³

To validate our manager assignments, we conducted a survey of the Sales and Distribution division (described in Section 3.5 below) where employees are asked to self-report managers who “have directly influenced your key performance indicator and pay grade.” In the month of the survey, December 2017, 91% of the managers we identify using the organization chart are also listed by the employee as someone who directly influences their key performance indicator and pay grade.¹⁴

The managers tend to be significantly above their subordinates in the firm’s hierarchy. For example, the modal (mean) distance between the manager and its employees is 5 (5.3) pay grades. The manager typically can influence the careers and daily lives of the employees in various ways. Most important, the manager provides key input in the decision to promote the employee. And even if the employee is not promoted, the manager still provides input that influences the raises and bonuses that the employee may receive. The manager also has discretion to distribute workload across team members however they see fit. Even if the work hours are rigid, such as for a tellers, the manager still has latitude to approve leaves of absences or late days.

3.4 Manager Switches

Employees can change managers over time for a variety of reasons. Some of those reasons are under the control of the employee and thus likely endogenous. For example, employees

¹³In some cases, the team does not have any directors designed in the organizational chart. In that case, the team is assigned to the director one layer higher in the organizational chart hierarchy.

¹⁴Our comparison is restricted to pairs in the administrative organization chart that remain together for one year or more. If instead we include all pairs, even those who have been together for just one month, we still find substantial overlap: 78% of the managers we identified are also listed by the employee.

may be promoted to a higher position and as a result be assigned to a different team with a different manager. Or an employee who disliked his or her manager may ask to be transferred to another team. Instead, we focus on manager switches that are, arguably, outside of the control of the employee. For example, a team may lose its manager because the manager was promoted to a higher position. The most typical case is that managers rotate laterally between different teams. Other times, a manager leaves the company, is promoted to a higher team, or must transfer to another team due to personal issues. In those cases, the firm is forced to find a new manager for that position.

In identifying these exogenous events in the data, we impose a few conditions. We require that the new manager must assume responsibility for all the employees in the unit. In other words, the unit experiences a manager switch, not an individual employee. Also, we exclude managers that are temporary replacements, requiring them to remain with the unit for at least one quarter. In the results section we discuss a series of robustness checks regarding the definition of the switch events.

We do not need to *assume* that these manager events are exogenous. Instead, we test that assumption directly in our event-study analysis. If manager switches are orthogonal to the outcomes of employees, in the period leading up to the event itself the outcomes of employees will follow the same trajectory irrespective of their future event type. However, we do have anecdotal evidence suggesting that the gender of the manager in these switches is as good as random. As part of corporate strategy, managers are expected to gain experience in all aspects of the bank. For this reason, managers move across teams within divisions and across divisions; for example a manager from HR can 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 across major divisions. When managers quit or request a transfer, they are required to give 30 days notice, and the set of candidates available to fill the role in time is (anecdotally) a very small set and sometimes an empty set. This shortage contributes to the reason that the bank rewards managers willing to transfer quickly from distant divisions, and the reason that job postings for every managerial level of the bank can be found on the internal and external dash board. In Table 1 we show the reasons (quit/hire, lateral transfer or promotion) for the transition of the incoming and outgoing manager. The most typical case for manager switches is that both are rotated laterally.

Over the span of our data, there are 10,101 events involving 6,536 unique employees and 706 unique managers. In Figure 2 we show that these events are distributed uniformly over the four year panel. 44% of employees will experience at least one event in this window, but only 33.6% experience two or more events. Given the distribution of team sizes, an event will affect on average five workers, and the interquartile range of events affect teams with

between three and ten workers. In the results section we discuss robustness checks under alternative treatment of the event data.

One relevant question is whether the sample of individuals who move (44%) is representative of the whole organization. Columns (1) and (2) of Table 2 compare the characteristics of employees and individuals who do and do not experience at least one switch. The samples are almost identical in age and education. The sample of individuals with events is a bit more female and has lower pay grade than the ones who do not experience events. This simply reflects the fact that there is more turnover and rotation at the bottom of the hierarchy, which happens to be less male and have lower pay grades, than the top of the firm.

Another question is whether the characteristics of the employees and managers are similar across the different types of manager switches. This is not necessary for the identification strategy: the critical condition is that the evolution of the outcomes are parallel, not that the levels are the same. However, the comparisons in levels can give a sense of how plausible the parallel trends are. Columns (3) and (4) of Table 2 compare between switches from female to male managers and switches from female to female managers, while columns (5) and (6) compare between switches from male to female managers and switches from male to male managers. There are a few differences here and there, but for the most part the characteristics of employees, incoming and outgoing managers are quite similar across the pairs of event types.

When we define placebo events, or smoker events, the definition of manager switches remains largely the same. We merely categorize the events based on different traits about the manager and employee, such as the evenness of the birthday or their smoker status rather than their gender. By construction, the number of placebo events is equal to the number of gender events. Due to the fact that the smoker analysis is based on a subsample (males with data on smoker status), the number of smoker events is smaller than the number of gender events. Like the gender events, we find that the placebo events and smoker events are largely homogeneous over time and across individuals – for more details, see Appendix C.

3.5 Survey on Relationship with Managers

To obtain data on relationship between employees and their managers, we were able to distribute a survey to the employees in the largest division: Sales and Distribution. A sample of the survey instrument is included in Appendix A, and described below.

The survey starts by asking respondents to list all the managers that “directly influenced your key performance indicator and pay grade either in your current position or past positions”. They could select up to 6 managers. If they had more than 6 to list, we asked employees to prioritize the most important and recent ones since 2015. The rest of the survey

repeated a series of questions for each of the managers listed by the respondent.

The goals of the survey were several. The first goal was to get the employees to self-identify managers that could influence their career outcomes, which we can use to validate our method of identifying managers. A second goal of the survey was to identify whether the employee works in close proximity to the manager. We wanted this variable to split positions into those with low vs. high proximity to the manager. The question uses the exact name of the managers the employee listed at the top of the survey and asks, “how often are (or were) you physically working near *[Manager’s Name]*?” with a multiple choice answer, 4-6 days per week, 2-3 days per week, or less frequently.

A third goal of the survey was to measure the socialization between the employee and the manager. We asked, “out of 10 work breaks (including lunch or random breaks), how many would include *[Manager’s Name]*?”. We can construct a simple variable which is the fraction of the breaks shared with the manager.¹⁵ The survey also asked an alternative socialization question: “Of the last 10 emails you sent to *[Manager’s Name]*, how many included some part that was personal?.”

We also included questions intended to measure whether the employee knows the preferences of the manager. With that goal in mind, we asked the respondent to provide his or her own favorite sport team, and to guess what the favorite sport team of the manager is. For the pairs of employees and managers who responded to the survey, we can measure if the employee identified the sports team of the manager accurately.

Last, this survey also included a couple of questions on the smoker status of the employees and their managers. These questions were intended to supplement other measures of smoker status, as described in the next section.

We invited 4,847 employees to complete the survey in December 2017. The employees were invited by email. We include a sample of the invitation email in Appendix A. The head of the Sales and Distribution Division requested full participation and gave permission to carry out the survey during work hours. We emphasize that answers to these survey questions would not be revealed to co-workers or managers. A total of 3,345 employees completed the survey, amounting to a 89% response rate. The median respondent completed the survey in 12 minutes. Respondents reported answers for up to their last six managers – the modal respondents reported information on their last three managers. This results in a dataset with 9,068 employee-manager pairs.

¹⁵We chose to ask the question this way, as a share of 10 breaks, rather than to report the overall number of breaks, to ensure that there would not be a incentive to under-report so as to appear more focused and productive. The downside is that we do not have a measure of the overall number of minutes spent together in a given week.

3.6 Data on Smoker Status

We measure the smoker status of employees and their managers in two ways. We collect self-reports of current smoking status and smoking history during the 2017 annual health exam that occurs on-site during the work day and a corresponding online workplace health survey with the same questions and framing. To complement the above data, which were snapshots of employees working in September 2017, we conducted additional surveys to crowdsource the smoker status of past and present employees.

In addition to the survey described in the previous section, which included two questions about smoker status, we deployed a very short survey exclusively about smoker status. A sample of the survey instrument is included in Appendix B. The 2-minute survey asks respondents about the smoker status of their past co-workers in order to collect the smoker status of employees who had quit the bank prior to the annual health exam. The invitations to the survey were distributed via email on February 2018, and offered cash prizes to be raffled among survey respondents. We invited a total of 6,022 employees and had a response rate of 39%.

If an employee appears in the 2017 annual health exam data, we use his or her response to assign the smoker status. For employees who do not appear in the annual health exam data, we impute their smoker status using the crowdsourced data: we classify them as smokers if over one-third of the crowdsourced reports flag the subject as a smoker.¹⁶ Indeed, we can validate the crowdsourced measure by using the subsample of respondents for which we have data on both the annual health exam and the crowdsourcing survey. Given our definition, the crowdsourced measure of smoker status coincides with the health records 82% of the time.

When we construct the sample for our smoking event study analysis, we are able to assign smoking status to 57% of the main sample. 59% of these employees (33% of the sample) self report their smoking status; the remaining 41% have their smoking status crowdsourced.

4 Results: Effects of Manager’s Gender

In this section, we document the effects of manager gender on the employee’s careers.

¹⁶This one-third threshold is arbitrary, but largely inconsequential. 21% of the sample flips their smoke status when we raise the threshold to require all reports indicate the person is a smoker and 9% flip status when we lower the threshold to any smoker report. As shown in Appendix F, our results hold even under those extreme thresholds.

4.1 Effects on Pay Grade

Figure 3.a corresponds to the difference-in-difference estimates described in Section 2. This figure corresponds to the effects on pay grade from switching from a female to male manager relative to the effect of switching from a female manager to another female manager. This event-study graph shows the evolution of the pay grade for female employees (red circles) and male employees (blue squares) separately in each of the 10 quarters leading up to a manager switch and 10 quarters after the manager switch. The quarter before the event (-1) corresponds to omitted category and thus it is normalized to zero by construction.

Figure 3.a shows that, in the 10 quarters prior to the switch, the evolution of pay grades is identical between male employees and female employees. On the contrary, the path of pay grades starts to diverge a few quarters after the switch: relative to female employees, the pay grades of the male employees grow faster.

At 10 quarters after such a manager transition, male employees increased their pay grades by an additional 0.53 points (p-value = 0.005), roughly equivalent to 13% higher pay, compared to male employees who transitioned from a female manager to a different female manager. In contrast, female employees have the same career progression regardless of the manager's gender. Indeed, this finding echoes the results of earlier studies that show female referees and committee members do not serve to increase the odds of acceptance or promotion for females under evaluation [Bagues et al. \(2017\)](#); [Card et al. \(2019\)](#). This finding is also consistent with the results from [van Hek and van der Lippe \(2019\)](#) and [Srivastava and Sherman \(2015\)](#) that earnings are similar for female employees under female managers and female employees under male managers.¹⁷

The flat evolution of coefficients identified for female employees does not imply that female employees are not getting promoted, or are stuck in the same pay grade. The average pay grades of females and males are both going up over time. The female employee coefficients are simply capturing the fact that the growth rates of pay grades for females are similar across those who experience transitions between female and male managers, and those who experience transitions between female and new female managers.

Note that, if promotions were a tournament at the team level, we would expect the positive effect on male employees would mechanically create a negative effect on the female employees, due to crowding out. However, anecdotal evidence suggests that promotions are not a tournament at the team level. In other words, employees are not competing with their teammates for a fixed number of promotions. Consistent with the anecdotal evidence, we do not see any evidence that the male employees are crowding out their female teammates.

¹⁷Keep in mind that the results from those studies are not directly comparable, because we use a quasi-experimental specification and because we look at effects in a longer-term.

However, this evidence does not imply that the male employees are not crowding out anyone: they are probably crowding out other employees in the same position but in different teams and/or external hires.¹⁸

Figure 4.a corresponds to the triple-difference estimate described in Section 2: i.e., the coefficients from Figure 3.a for male employees minus the coefficients from the same figure but for the female employees. The triple-difference is consistent with a male-to-male advantage: male managers (relative to female managers) improve the career progression of male employees (relative to female employees). At 10 quarters after the transition, the male-to-male advantage amounts to 0.50 paygrades, which is not only highly statistically significant (p-value=0.003), but also economically significant.

We provide two main robustness checks for our identification strategy. The first test consists of analyzing the opposite types of transitions. Figure 3.b is equivalent to Figure 3.a, only that it corresponds to the opposite type of transition: i.e., comparing a switch from a male manager to a female manager minus the switch from a male manager to a different male manager. Comparing Figures 3.a and 3.b indicate that, as expected, the effects of “losing” a male manager are the opposite of the effects of “gaining” a male manager, both in terms of the timing as well as the magnitude. Figure 3.b shows that, again, the evolution of pay grades is identical between male employees and female employees in the 10 quarters prior to the switch. In contrast, the path of pay grades starts to diverge some quarters after the switch date. More precisely, the path of pay grades of female employees is not affected by the switch, while the male employees grow slower after the switch.

We can provide a more quantitative comparison between the transitions in opposite directions. Figure 4.b, like Figure 4.a, presents the triple-differences estimates, but using the transitions in the opposite direction. The male-to-male advantage 10 quarters after the switch is 0.50 looking at employees “gaining” a male manager (Figures 4.a), while it is -0.38 looking at employees “losing” a male manager (Figures 4 over the same time frame.b). Indeed, we cannot reject the null hypothesis that these two effects are equal in absolute value (p-value=0.588) 10 quarters after the event.

4.2 Placebo Events

Another robustness test is based on placebo events. We reproduce the whole analysis but, instead of focusing on gender as the relevant characteristic of managers and employees, we focus on a characteristic that we know *ex ante* should not be relevant: whether someone was born on an even or odd date. This placebo provided a sanity check for a number of

¹⁸The external hire channel may be important because of the significant attrition (12.5% per year) and growth (5.9% yearly growth rate in the overall number of employees).

reasons. First, it is intended to show that there are no mechanical reasons why our event-study framework would generate significant effects. Second, this falsification test provides a nice check that our standard errors are not prone to over-rejection.

Figure 3.c shows that there is no meaningful difference between switching to an odd or even birthday manager. For employees with an odd or even birthday, pay grade in the period leading up to the event and the subsequent 10 quarters follows the same trajectories. Thus, when we plot the difference between effects for workers with odd and even birthdays (i.e. the even-to-even advantage) in Figure 4.c, the results are reassuringly close to zero, statistically insignificant and precisely estimated. The point estimate after 10 quarters of -0.12 is close to zero, statistically insignificant (p-value=0.22) and precisely estimated. This even-to-even advantage of -.12 is smaller than the male-to-male advantage of .5 presented above – moreover, the difference is statistically different (p-value=0.036). Last, in Appendix D.3 we show that the results are virtually the same if we use the even-to-odd transitions instead of the odd-to-even transitions.

4.3 Effects on Attrition

Next, we provide evidence about the mechanisms underlying the male-to-male advantage. One natural mechanism could be attrition: i.e., male employees reach higher positions under male managers because they are less likely to leave the firm. To probe this channel, we measure the effects of the manager switches on the probability of leaving the firm.

When using this specific outcome, there is an extra challenge for the event-study analysis. By construction, employees do not experience event switches after they leave the company. As a result, while we can still estimate the post-treatment coefficients, but it is not possible to estimate the pre-treatment coefficients. This is a typical challenge in event-study analysis, which we address using a standard tool (Kleven et al., 2019): creating hypothetical events for individuals after they quit. We do this in a straightforward way: we assign them the events experienced by the team in which they were working prior to leaving the firm.

The results are presented in Figure 5.a. Looking at 10 quarters after the event, the effects on cumulative attrition are close to zero (-2.8 percentage points) and statistically insignificant (p-value = 0.33). On average, 10 quarters after an event, 35% of the employees leave the firm. Thus the estimated effect is quite small even relative to the baseline rate.

Another form of attrition consists of internal transfers: i.e., employees who are transferred laterally to a different team. In Appendix Figure D.4.1, we measure the effects of the switch events on this alternative outcome. We find that the effects on internal transfers, like the effects on firm exits, are small and statistically insignificant. For example, the estimated male-to-male advantage in internal transfer 10 quarters after the event is small (-1.1 percentage

points) and statistically insignificant (p-value=0.76).

4.4 Effects on Effort and Output

A natural question is whether male employees are promoted faster under male managers because they work harder or because they are more productive. For example, it is possible that male employees may see role models in their male managers but not in their female managers (Kofoed and McGovney, 2019). Or perhaps male managers are better able to communicate with the male employees, or better at monitoring them. To probe these explanations, we estimate the effects on effort and output. To allow for intuitive comparison of results in different units, we follow Hastings et al. (2019) by setting the scale of each graph at approximately twice the within-individual standard deviation.¹⁹

Absenteeism is a measure of performance that the bank closely tracks and considers an important factor in team productivity. For this reason we have excellent measures of leave days, including half days, along with the explanation for these leaves. Even if a manager excuses a leave, security will register the absence. We subtract all excused and unexcused absences when calculating the total number of work days for an employee in a given month. In Figure 5.b we replicate the event study analysis, focusing on the female to male manager transition, with log monthly work days as the dependent variable. We find no difference between male and female employees leading up to the event, nor after the event. The point estimate for the male-to-male advantage after 10 quarters is small (0.016) and statistically insignificant (p-value = 0.39). We can interpret the magnitude as a percentage increase in the days worked of 1.6%, which likely falls short from justifying the 13% pay gap that we document above.

For employees that work office jobs within the bank’s headquarters, we observe a proxy for effort: the number of hours the employee works in the office as measured by the time-stamped swipes in and out of the lobby. In Figure 5.c we carry out the event study analysis with the log average daily hours worked in a given month as the outcome of interest. We find no evidence of differential changes in behavior between the men and women who experience their manager change from female to male, compared to female to female. At precisely 10 quarters after the event the male-to-male advantage is 0.21 and statistically insignificant (p-value = 0.15). While statistically insignificant, the point estimate is still large. However, this is an outlier: for all the other quarters the coefficients are close to zero and statistically

¹⁹Hastings et al. (2019) perform a similar normalization, but use the inter-quartile range instead. To allow for familiar scales, we choose a “round” number that is close to twice the within-individual standard deviation. For example, the within-individual standard deviation of pay grade is 0.479, so the y-axis on this graph ranges from -1 to 1 instead of -0.958 to 0.958.

insignificant. For instance, the coefficients for +8 and +9 quarters are 0.067 (p-value=0.43) and -0.011 (p-value=0.91).

Next, we measure the effects on sales performance. This outcome is available for the subsample of employees who have any sort of sales role – about one-third of male employees have this role. Since we are looking at a subsample, it is important to verify that the effects on paygrade are present in this subsample too, which is the case – results reported in Appendix Figure F.3.

The outcome is the sales performance index, which aggregates the revenues across different products. To avoid sharing confidential information, and for ease of interpretation, we normalize this index to have a mean of 100. The within-individual standard deviation of this outcome is 95.

In Figure 5.d, we display the results of an event study analysis with sales revenue as the outcome. The point estimates are mostly close to zero and negative. This suggests that, if anything, male employees are doing worse in terms of performance after switching from female to male managers (relative to switching to female to another female managers). Hence, our analysis of sales supports the hypothesis that the observed male-to-male advantage is not due to higher productivity.

While the estimates of these coefficients in some periods are negative and statistically significant, we cannot conclusively say that the switch to male manager causes male employees to decrease their sales. First, the point estimates are not precisely estimated and do not follow a consistent path in the post-switch period. Most important, this finding is not robust to the symmetry test: when we consider switches from a male manager to a female manager, instead of finding positive effects on male employees we find the effects to be precisely estimated around zero (results reported in Appendix D.4).

4.5 Counterfactual Analysis

Next, we discuss the economic magnitude of the male-to-male advantage. A first natural benchmark would be the pay grade gap. The average difference in pay grade between male and female employees is 0.85. For male employees, the advantage of being under a male manager instead of a female manager is 0.5 pay grades (after 10 quarters). Relative to the overall gender gap, this advantage is quite substantial ($58.8\% = \frac{0.50}{0.85}$).

Moreover, we would like to say something about what would happen to the overall gender gap if we were to remove the male-to-male advantage. For this counterfactual analysis, we need to come back to a caveat discussed in Section 2: our specification cannot distinguish whether the male-to-male advantage is driven by whether male employees get favorable treatment from male managers, unfavorable treatment by female managers, or a combination of

both. Indeed, this challenge is not unique to our methodology or to our context. For example, even in a randomized controlled trial, we can only compare male managers versus female managers because there is no gender-neutral managers to compare against.

There is, however, some suggestive evidence that the male-to-male advantage documented above is due mostly to favorable treatment by male managers. First of all, promotion rates are higher for male employees than for female employees. Thus, the favorable treatment by male managers provides a likely explanation. Of course, it could still be possible that there is unfavorable treatment by female managers and that there are other factors that favor men and bend the gender gap in the opposite direction.

One challenge with the event-study analysis is that it compares the evolution of the pay grades over time but conceals the information about the levels. As complementary evidence, we use a different approach that does not difference out the information about levels. Intuitively, we would like to compare similar male and female employees who recently worked mainly under male managers with those who have recently worked mainly under female managers. Consider the following regression:

$$\begin{aligned} \Delta P_{i,t} = & \alpha_0^M \cdot (1 - F_i) + \alpha_1^M \cdot S_{i,t-1} \cdot (1 - F_i) + \alpha_0^F \cdot F_i + \alpha_1^F \cdot S_{i,t-1} \cdot F_i \\ & + \beta T_{i,t} + \rho_{P_{i,t}} + \epsilon_{it} \end{aligned} \tag{4}$$

Let $\Delta P_{i,t}$ be the change in pay grade from t to 10 quarters later. $S_{i,t-1}$ is the share of time between $t - 1$ and t that the employee spent working under a male manager. The regression includes some additional controls: $T_{i,t}$ is the employee's tenure and $\rho_{P_{i,t}}$ is a fixed effect for current pay grade, which allows us to flexibly compare between employees starting out from the same paygrade.

According to the event-study results from above, we would expect $\alpha_1^M > 0$ (i.e., male employees do better under male managers) and $\alpha_1^F = 0$ (i.e., female employees do equally well under male or female managers). Unlike the event-study analysis, this simpler framework is not equipped to deal with causal inference. On the other hand, this simpler framework has the advantage that it can disentangle between favorable treatment of male managers versus unfavorable treatment of female managers. If the effects were entirely due to unfavorable treatment by female managers, we would expect male employees to do as well as female employees under male managers, but worse than female employees under female managers. If the effects were entirely due to favorable treatment by male managers, we would expect male employees to do as well as female employees under female managers, but better than female employees under male managers.

The results are presented in binned scatterplot form in Figure 6.²⁰ The main results are consistent with the event study analysis: 10 quarters later, male employees reach an extra 0.5 promotions under male managers relative to under female managers; and the female employees are promoted at similar rates under female and male managers. The magnitude of this simplified analysis need not necessarily line up with the event study, but the fact that they do is reassuring.

Figure 6 shows that under female managers (i.e., towards the left of the x-axis), female and male employees tend to get promoted at the same rate. Under male managers (i.e., towards the right of the x-axis) the female employees are still promoted at roughly the same rate, but the male employees are promoted more. In particular, we estimate $\alpha_1^F = 0.059$ (p-value<0.001) and $\alpha_1^M = 0.35$ (p-value<0.001). Moreover, the difference between these two coefficients is highly significant (p-value<0.001). In sum, these patterns are consistent with the interpretation that the male-to-male advantage arises from favorable treatment by male managers rather than unfavorable treatment by female managers.

Under the assumption that our findings are due to a positive effect of male managers on male employees, we can finally compute what would happen to the overall gender gap if we were to remove this male-to-male advantage. Since 66% of male employees have male managers, the average pay grade of male employees under these would be reduced by 0.33 ($= 0.50 \cdot 0.66$) if the male-to-male advantage were removed. In turn, this would reduce the gender pay gap by 38.8% (from 0.85 to 0.52 pay grades). As a result, the male-to-male advantage can explain 38.8% of the gender gap at this organization. However, if some of the effects were due to a negative effect of female managers on male employees, then the effects on the gender pay gap would be smaller. In the extreme case where all the effects were due to negative effects of female managers on male employees, then removing these manager effects would actually increase the gender pay gap, as male employees would see their pay grades go up while the female employees would remain unaffected. In this sense, the 38.8% reported above serves as an upper bound.

When we calculate the male-to-male advantage from the sample of employees who start out with a female manager, and transition to either male or female managers, we find a male-to-male advantage of 0.5. However, when we use the reverse transitions, among those that start with a male manager, the magnitude is somewhat smaller (0.38). The results of this counterfactual analysis would be somewhat smaller if we used the sample who started with a male manager and transitioned away: instead of explaining 38.8% of the gender gap, it would explain 29.5%.

²⁰As a benchmark, Appendix D.1 presents the same analysis using promotions in the following 5 quarters instead of 10 quarters.

5 Results: Socialization Mechanism

In this section, we discuss some potential mechanisms that may mediate the male-to-male advantage documented in the previous section, and we provide suggestive evidence about the role of socialization.

5.1 Timing of the Effects

The male-to-male advantage is probably due to a number of channels. One natural family of explanations, which have nothing to do with socialization, are based on pre-existing in-group biases. For example, male managers may have biased beliefs about the productivity of other men, they may care more about the well-being of male employees, or they may seek to pay the favor forward for the men who have helped them in the past.

There is a second family of channels, which involve socialization. For example, male managers may become emotionally attached to the male employees over time, and thus feel increasing pressure to promote their male employees. Perhaps male employees use the face-to-face interactions to gain the manager's sympathy and schmooze their way into promotions. Socializing with the manager may make the accomplishment and effort of the employee more salient to the manager, thus making the employee more likely to be rewarded with a promotion. Male employees may use the time spent with their manager to claim credit and engage in self-promotion (Isaksson, 2019; Coffman et al., 2019; Exley and Kessler, 2019; Sarsons, 2019). Male employees may get favorable treatment from the managers by getting assigned tasks that are more conducive to promotions (Lehmann, 2013). It is also possible that male managers are more willing to get more involved in working alongside with, or training, their male subordinates (Ranganathan, 2019).

A first way of disentangling between these two families of explanations is by considering the predictions made about the timing of the effects. The first family of channels, which do not involve socialization, predict that the male-to-male advantage should manifest right after the manager's assignment: some male employees are coming up for promotion right after the manager switch, and if the male manager is biased then those biases should affect those immediate promotion decisions. Indeed, these channels may even predict that the male-to-male advantage would diminish over time, as managers collect more information about their employees and thus correct some of their initial biases. On the contrary, the family of mechanisms related to socialization predict that the effects should take time to materialize and should also build up over time. For example, if male manager favor male employees due to an emotional attachment, one would expect that such emotional attachment would take time to develop.

While the two families of explanations could both play a role, the timing of the male-to-male advantage presented in the previous section is more consistent with the socialization channel. In the first one to four quarters after the manager switch, the estimated male-to-male advantage is close to zero (0.054), statistically insignificant and precisely estimated (p-value = 0.55). Only after the first year the male-to-male advantage begins to appear and builds up over the following quarters; by quarter 10, the estimated advantage is 0.50 (p-value = 0.003). Below we provide more direct evidence in support of the socialization channels.

5.2 Heterogeneity by Proximity to the Manager

The first test of the socialization channel exploits heterogeneity by the proximity to the manager. If socializing with the manager plays an important role, then we should observe stronger effects when the employee has more frequent face-to-face interactions with the manager.

Recall the survey question about proximity with one’s manager described in Section 3.5. We use responses to that question to split positions in two groups: higher and lower access to the manager. We average the response to this question for all employees in a position to create a position-level value. Using this method, we can categorize 62% of the position titles in the Sales and Distribution division (i.e., for which the survey data was collected). We then split these position into high and low access positions according to the median value across positions.²¹ An example of a position that was classified as high-access is a customer support specialist who will be present in one place to receive calls, with his or her manager close by. On the other hand, a sample position that was classified as low-access is a sales and quality development director, who usually travels between branches and reports back to his or her manager about the local situation.

We compare the advantageous effect of switching to a male manager for male and female employees when the position requires employees and managers to work in close proximity (above median levels of interaction) and low proximity (below the median). Since this analysis is restricted to a subsample from the Sales and Distribution division with classified positions (63% of the employees in the main sample), we verify that the effects of our main event study analysis with pay grade are still present in this subsample (results reported in Appendix F.3). Additionally, in addition to restricting to 64% of the sample, we will split the remaining sample in half, resulting in estimates that are less precisely estimated and thus must be interpreted with extra care.

The results are presented in Figure 7. We find that not only is the effect more pronounced among those with high access, but it is insignificant among those with low access.

²¹In positions in the higher-access group, 88% of employees report working with their manager “everyday or most days”, compared to only 65% of the employees in the low access group.

In the high access group, 10 quarters later the male-to-male advantage in pay grade is 0.87 (p-value=0.002). In comparison, the corresponding coefficient for the low-access group is negative (-0.62) and statistically insignificant (p-value=0.21). Moreover, we reject the null hypothesis that these two coefficients are equal (p-value= 0.006).

5.3 Effects on Time Spent with Manager

To explore the role of socialization further, we measure the effects of the manager switches on the socialization between managers and employees during breaks.

As a proxy for socialization, we use responses to the survey question about the share of the last ten breaks that the employee took that was shared with his or her manager.²² We must keep a couple of important caveats in mind, however. First, in addition to the amount of time spent with the manager, there may still be gender differences in the ability to convert the social interactions into promotions. Second, our measure of interaction is just a proxy, which may understate or overstate the degree to which employees socialize with their managers. For example, it is possible that some employees and managers are on the same break “schedule” but do not share meaningful social interaction during these breaks.²³

We can provide suggestive evidence to indicate that employees who share more breaks with their managers have better knowledge about the manager’s personal interests. In this same survey, we asked employees both what their preferred sports team is, and to guess their manager’s preferred sports team. For managers that themselves responded to the survey, we can compare their stated sports preference to their employees’ guesses and use this as a rough proxy for socializing. We find that employees who take more breaks with their manager are better able to successfully guess their manager’s favorite sports team. In particular, an increase from 0% to 100% in the share of breaks taken with the manager is associated with an 11 percentage point increase in the likelihood of accurately guessing the manager’s preferences (p-value < 0.001).²⁴

The survey dataset on share of breaks taken with the manager has substantially fewer observations (9,068 employee-manager pairs) relative to the the administrative dataset (374,913

²²The survey included an alternative measure of social interactions with the manager: the share of emails exchanged with the manager that were of a personal nature. In practice, however, there is not enough variation in this outcome (the average share of personal emails is just 5%) for the event-study analysis. This average response may indicate that personal emails are quite rare, which would be consistent with the bank policy that prohibits employees from using their work email addresses for personal communications. However, the rate of personal emails may be higher than reported because employees may be under-reporting the extent of this behavior.

²³Similarly, it is possible that a manager and an employee have an informal social relationship but which is reflected in time spent at night or during weekends rather than during the workdays.

²⁴For more details, see Appendix D.2.

employee-month pairs).²⁵ Thus, the survey data is not equipped to the type of detailed event-study analysis presented for the administrative data. Thus, we use a miniaturized version of the event-study framework tailored to this survey dataset. Let subscripts i and m denote a pair of an employee and one of his or her managers, respectively. Let $Share_{i,m}$ be the share of breaks that employee i took with manager m . Consider the following econometric specification:

$$\begin{aligned}
Share_{i,m} = & \sum_{j \in J} \beta_{j,m}^F \cdot F_i \cdot D_{i,m}^j + \sum_{j \in J} \beta_{j,m}^M \cdot (1 - F_i) \cdot D_{i,m}^j + \\
& \sum_{j \in J} \beta_{j,m+1}^F \cdot F_i \cdot D_{i,m+1}^j + \sum_{j \in J} \beta_{j,m+1}^M \cdot (1 - F_i) \cdot D_{i,m+1}^j + \\
& \delta_m^F + \delta_m^M + \gamma X_{i,m} + \epsilon_{i,m}
\end{aligned} \tag{5}$$

The set $J \in \{J_G, J_E, J_S\}$ refers again to the types of manager switches: in the gender analysis, for example, $J_G = \{F2M, M2F, F2F, M2M\}$. The variable $D_{i,m}^j$ is a dummy variable that takes the value 1 if individual i experiences an event of type j from manager $m - 1$ to manager m . As usual, we interact these dummies with gender dummies, to allow the effects to be gender-specific. The coefficients next to these variables ($\beta_{j,m}^F$ and $\beta_{j,m}^M$) are intended to capture the effect of the manager switch on the socialization of the employee’s with the manager that the employee switched to. The variable $D_{i,m+1}^j$ is a dummy variable that takes the value 1 if individual i experiences an event of type j from manager m to manager $m + 1$. The coefficients next to these variables ($\beta_{j,m+1}^F$ and $\beta_{j,m+1}^M$) are intended to provide the usual tests for pre-trends: they measure whether future manager switches affect the employee’s socialization with the current manager. Additionally, the regression includes gender-specific time effects (δ_m^F and δ_m^M) and a set of basic controls ($X_{i,m}$): unit size, manager’s paygrade, and a vector of position title dummies.

The results from the miniaturized version of the event-study specification are presented in Figure 8.a. The coefficient labeled “after switch” corresponds to the differential effect of switching from a female manager to a male manager (relative to switching from a female manager to a different female manager). We find that the male-to-male advantage, caused by a switch from female to male managers increases the share of breaks taken with the manager by 0.23 (p-value = 0.001). On the contrary, there is no effect for female employees. The coefficient labeled “before switch” corresponds to the falsification test from the event-study

²⁵We do not have data on all employees, and even for the employees we have data on, we only have data on less than a handful points in time (as opposed to the monthly data for four years, in the administrative records).

framework: i.e., it is the “effect” of a switch that will happen in the future but has not happened yet. As expected, we find these falsification coefficients to be close to zero and statistically insignificant.

We can again validate this design by rerunning this miniaturized event study with our placebo events. These results are presented in Figure 8.b. The estimated even-to-even advantage following a move to a manager with an even-birthday compared to a manager with an odd-birthday is null and precisely estimated ($p = 0.62$).

5.4 Smoker Events

Our last piece of evidence on the socialization mechanism is based on a shock that increases socialization while holding gender constant. Anecdotally, smokers tend to take smoker breaks together. As a result, for an employee who smokes, having a manager that smokes can increase socialization with the manager.

When we construct the sample for our smoking event study analysis, we are able to assign smoking status to 57% of the main sample. 59% of these employees (33% of the sample) self report their smoking status; the remaining 41% have their smoking status crowdsourced.

We start by using the survey data to test the conjecture that smoking employees socialize more with smoking managers. In Figure 8.c we show that the impact of moving from a non-smoking male manager to a smoking male manager is a rise of approximately 0.27 (p -value = 0.008) in the share of breaks taken with the manager when the employee is a smoker. Thus, the male-to-male advantage of 0.23 (p value = 0.001) is similar in magnitude to the smoker-to-smoker advantage. On the contrary, for non-smokers the smoker status of the manager does not affect the share of breaks taken with their manager.

The results described above confirm that the smoker events are significant shocks to socialization. If these positive effects on socialization coincide with positive effects on promotion rates, that would provide support for the socialization mechanism. Figure 4.d shows the event-study graphs for pay grade. We compare the outcomes of male employees who transition male-non-smoker to male-smoker versus the outcomes of male employees who transition male-non-smoker to male-non-smoker. We split male employees by their smoker status, to compare the effect on smoker vs. non-smoker employees. Before proceeding with the remaining results, it is important to note that the analysis of the smoker events is based on a subsample (57%) of the data used for the gender events: we focus on the male employees and managers, who constitute less than half the sample, and we further focus on the subset of those for whom we observe the smoker status. As a result, this event-study analysis is based on 365 manager switches involving 917 unique employees and 227 unique managers. The smaller samples result mechanically in event-study coefficients that are less precisely

estimated, yet they are still precise enough for the purposes of our analysis.

Figure 4.d shows that, for smoker employees, switching to smoker manager causes pay grade to increase by an additional 0.76 points ($p\text{-value} < 0.001$) after 10 quarters. This smoker-to-smoker advantage is in the same order of magnitude as the corresponding male-to-male advantage (0.50 points). Indeed, we cannot reject the null hypotheses that these two effects (0.76 and 0.50) are equal to each other ($p\text{-value} = 0.31$).

Under the assumption that the smoker-to-smoker advantage works purely through socialization, the similarities between the smoker-to-smoker and male-to-male advantages would suggest that the male-to-male advantage also works mostly through socialization.

We can also examine the reverse smoker-status transitions: among everyone who starts with a smoking manager, we can compare those who transition to a non-smoking manager with those that transition to another smoking manager. Unfortunately, we have significantly less power because smoking managers are a minority and transitioning from one smoking manager to another smoking manager is rare. Nevertheless, we include the analysis in Appendix D.3. The point estimates have the expected timing and direction, but are somewhat smaller in magnitude, less precisely estimated and thus statistically insignificant.

In our particular setting there are gender differences in smoker status: 33% of men smoke while less than 5% of women smoke. As a result, some of the male-to-male advantage can be a byproduct of the smoker-to-smoker advantage. However, due to the fact that smokers (both managers and employees) are a minority, this contribution is minor. In Appendix G we control for the smoker events in the gender event-studies, and find that the male-to-male advantage is only slightly smaller. Last, Appendix F shows that the results are robust to changes in the specification, such as using a different criteria to code the smoker status data.

6 Conclusions

We test the old boys' club hypothesis using data from a real-world corporation. At this institution, manager rotations across teams are common and create transitions in the gender of the manager that are largely out of the employee's control. We use an event study analysis over a four-year period and find that male employees who are assigned a male manager are promoted at a faster rate than their female counterparts, with no observable change in performance. Women, in turn, are promoted at the same rate whether assigned to a male manager or female manager. The male-to-male advantage can explain one-third of the gender gap in pay grades.

We provide suggestive evidence that socialization plays a role in how the male-to-male advantage develops. The advantage develops slowly over the course of two years, consistent

with the time-intensive process of relationship building. This advantage is concentrated in positions where male managers and employees spend time in close proximity. When a male employee transitions from a female manager to a male manager, the share of breaks taken together rises significantly. We also show that, even among male employees and male managers, a shock to socialization through co-smoking has similar effects on promotion rates.

We developed our identification strategy so that it can also be applied to other contexts. The rotation of managers is common in large corporations. And the data necessary for the analysis, such as paygrades, demographics and manager assignments, are probably available for most large organizations. We hope this methodology will be applied to other firms, countries, and industries, which will help to generalize the findings and identify where these biases are more or less problematic.

Last, our findings have implications for policies aimed at reducing gender gaps in pay and leadership. Companies can reduce 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 schmooze their way into promotions. Companies also could standardize the review process to use objective indicators, such as sales revenues and hours worked. Another strategy to curb these gender gaps is to level the opportunities for employees to socialize and connect with their managers. For example, companies could promote gender-neutral social activities. Whether these policies are effective at limiting gender gaps is an open empirical question that offers several avenues for future research.

References


- Azmat, G. and R. Ferrer (2017). Gender Gaps in Performance: Evidence from Young Lawyers. *Journal of Political Economy* 125(5), 1306–1355.
- 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. Barankay, and I. Rasul (2010). Social incentives in the workplace. *Review of Economic Studies* 77(2), 417–458.
- Bell, L. A. (2005). Women-Led Firms and the Gender Gap in Top Executive Jobs. *IZA Discussion Paper No. 1689*.
- Bertrand, M., R. Burgess, A. Chawla, and G. Xu (2018). Social Proximity and Bureaucrat Performance: Evidence from India. *NBER Working Paper No. 25389*.
- 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.
- Bronson, M. A. and P. S. Thoursie (2019). The Wage Growth and Within-Firm Mobility of Men and Women: New Evidence and Theory. *Working Paper*.
- Bursztyn, L., A. Gonzalez, and D. Yanagizawa-Drott (2018). Misperceived Social Norms: Female Labor Force Participation in Saudi Arabia. *NBER Working Paper No. 24736*.
- Bursztyn, L., F. Thomas, and A. Pallais (2017). 'Acting Wife': Marriage Market Incentives. *American Economic Review* 107(11), 3288–3319.
- Card, D., S. Dellavigna, P. Funk, and N. Iriberry (2019). Are Referees and Editors in Economics Gender Neutral? *The Quarterly Journal of Economics*, forthcoming.
- 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 (2019). Gender Stereotypes in Deliberation and Team Decisions. *Working Paper*.
- Cruz, C. and C. Tolentino (2019). Gender, Social Recognition, and Political Influence. *Working Paper*.
- Cullen, Z. and R. Perez-Truglia (2018). How Much Does Your Boss Make? The Effects of Salary Comparisons. *NBER Working Paper No. 24841*.

- Dahl, G. B., A. Kotsadam, and D.-O. Rooth (2018). Does Integration Change Gender Attitudes? The Effect of Randomly Assigning Women to Traditionally Male Teams. *Working Paper*.
- Dalvit, N., A. Patel, and J. Tan (2018). Intra-firm Hierarchies and Gender Gaps. *Working Paper*.
- Elting, L. (2018). How To Navigate A Boys' Club Culture. *Forbes*, July 27 2018.
- Exley, C. and J. Kessler (2019). The Gender Gap in Self-Promotion. *NBER Working Paper 26345*.
- Field, E., S. Jayachandran, R. Pande, and N. Rigol (2016). Friendship at Work: Can Peer Effects Catalyze Female Entrepreneurship? Faculty Research Working Paper Series. *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, 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.
- 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.
- Halldén, K., J. Säve-Söderbergh, and Å. Rosén (2018). Gender of the immediate manager and women's wages: The importance of managerial position. *Social Science Research* 72, 115–133.
- Hastings, J. M., R. Kessler, and J. Shapiro (2019). The Effect of SNAP on the Composition of Purchased Foods: Evidence and Implications. *Working Paper*.
- Hill, A. J. (2017). The Positive Influence of Female College Students on their Male Peers. *Labour Economics* 44(July 2016), 151–160.
- Hjort, J. (2014). Ethnic Divisions and Production in Firms. *Quarterly Journal of Economics* 4(4), 1899–1946.
- 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.

- Isaksson, S. (2019). It Takes Two: Gender Differences in Group Work. *Working Paper*, 1–49.
- Kleven, H., C. Landais, and J. E. Sogaard (2019). Children and Gender Inequality: Evidence from Denmark. *American Economic Journal: Applied Economics* 11(4), 121.
- Kofoed, M. S. and E. McGovney (2019). The Effect of Same-Gender or Same-Race Role Models on Occupation Choice. *The Journal of Human Resources* 54(2), 430–467.
- 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 Series*, 24740.
- Lang, I. H. (2011). Co-Opt the Old Boys’ Club. *Harvard Business Review*, November 2 2011.
- Lazear, E. P., K. L. Shaw, and C. T. Stanton (2015). The Value of Bosses. *Journal of Labor Economics* 33(4), 823–861.
- Lee, A. S. (2014). How to break up the old boys’ club in your office. *Quartz*, April 8 2014.
- Lee, D. S. (2009). Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects. *Review of Economic Studies* 76(3), 1071–1102.
- Lehmann, J.-Y. (2013). Job Assignment and Promotion Under Statistical Discrimination: Evidence from the Early Careers of Lawyers. *Working Paper*.
- Lucifora, C. and D. Vigani (2016). What If Your Boss Is a Woman? Work Organization , Work-Life Balance and Gender Discrimination at the Workplace. *IZA Discussion Paper Series*.
- 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. (2015). Gender Differences in Networking. *Working Paper*.
- Ranganathan, A. (2019). Getting their Hands Dirty: How Female Managers Motivate Female Worker Productivity through Subordinate Scut Work. *Working Paper*.
- Sarsons, H. (2017). Interpreting Signals in the Labor Market: Evidence from Medical Referrals. *Working Paper*.


- Sarsons, H. (2019). Gender Differences in Recognition for Group Work. *Journal of Political Economy*, forthcoming.
- Srivastava, S. B. and E. L. Sherman (2015). Agents of Change or Cogs in the Machine? Re-examining the Influence of Female Managers on the Gender Wage Gap. *American Journal of Sociology* 120(6), 1778–1808.
- 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.
- van Hek, M. and T. van der Lippe (2019). Are Female Managers Agents of Change or Cogs in the Machine? An Assessment with Three-Level Manager-Employee Linked Data. *European Sociological Review* 35(3), 316–331.
- Voth, J. and G. Xu (2019). Patronage for Productivity: Selection and Performance in the Age of Sail. *Working Paper*.
- 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: Relationship between Pay Grade and Salary



Notes: The above presents a binned scatter plot of log base salary against pay grade in March, 2017. We use this cross section of the bank's employees as we have access to their base salary from related work (Cullen and Perez-Truglia, 2018).


Figure 2: Descriptive Statistics about the Manager Switch Events



Notes: These figures display counts for events as defined in Section 3.4. Similar tables for smoker events and for placebo events are available in Figure G.1 and Figure G.2.


Figure 3: Difference in Difference Effects on Pay Grade

a. Female to Male
minus Female to Female




This regression includes 374,913 observations of 14,332 workers (5,071 Male & 9,261 Female). 3,160 of these workers experience a switch event (819 Male & 2,341 Female). There are 1846 transitions from a female manager to a male manager, 2120 from one female manager to another female manager. The within individual standard deviation of pay grade is 0.479. 95% C.I. displayed.

b. Male to Female
minus Male to Male




This regression includes 374,913 observations of 14,332 workers (5,071 Male & 9,261 Female). 4,489 of these workers experience a switch event (1,458 Male & 3,031 Female). There are 1745 transitions from a male manager to a female manager, 4291 from one male manager to another male manager. The within individual standard deviation of pay grade is 0.479. 95% C.I. displayed.

c. Odd to Even
minus Odd to Odd



This regression includes 374,913 observations of 14,332 workers (6,945 Odd BD & 7,387 Even BD). 4,705 of these workers experience a switch event (2,312 Odd BD & 2,393 Even BD). There are 3433 transitions from an odd-birthday manager to an even-birthday manager, 3086 from one odd-birthday manager to another odd-birthday manager. The within individual standard deviation of pay grade is 0.479. 95% C.I. displayed.

d. Non-Smoker to Smoker
minus Non-Smoker to Non-Smoker




This regression includes 93,534 observations of 2,894 workers (957 Smoking & 1,937 Non-Smoking). 912 of these workers experience a switch event (275 Smoking & 637 Non-Smoking). There are 287 transitions from a non-smoking manager to a smoking manager, 939 from one non-smoking manager to another non-smoking manager. The within individual standard deviation of pay grade is 0.519. 95% C.I. displayed, trimmed at -1 and 1.

Notes: The coefficients plotted above are difference-in-difference estimates of the form $\beta_{DD,t}^g = \beta_{\text{Gender Switch},t}^g - \beta_{\text{Same Gender},t}^g$ where $g \in \{\text{Male, Female}\}$ indexes the gender of the employee and the subscript indexes the switch event type and time since the event. See 2 for a formal discussion of the event study specification and Figure 4 for triple difference versions of the event study analysis. Coefficients in panels a) and b) are identified from the same regression. All specifications include employee, time, and time-by-(gender(a,b) birthday(c) smoker-status(d)) fixed effects. Additionally, all event-study coefficients are estimated by month, and are linearly smoothed to create quarter estimates; events outside the window are “absorbed” by indicators $\mathbf{1}(\text{Event Time} \geq 31 \text{ months})$ and $\mathbf{1}(\text{Event Time} \leq -31 \text{ months})$. See the notes below each panel for details about sample size, event counts, etc. Results for employees experiencing even to odd manager switches and smoker to non-smoker manager switches are available in Appendix Figure G.5.


Figure 4: Triple Differences Effects on Pay Grade

a. Female to Male
minus Female to Female



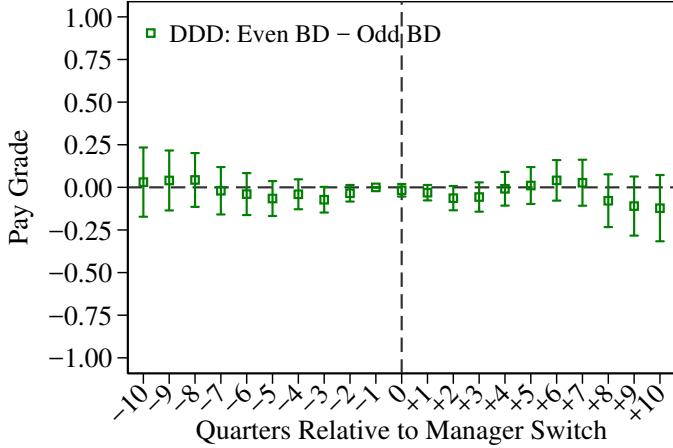
This regression includes 374,913 observations of 14,332 workers (5,071 Male & 9,261 Female). 3,160 of these workers experience a switch event (819 Male & 2,341 Female). There are 1846 transitions from a female manager to a male manager, 2120 from one female manager to another female manager. The within individual standard deviation of pay grade is 0.479. 95% C.I. displayed.

b. Male to Female
minus Male to Male




This regression includes 374,913 observations of 14,332 workers (5,071 Male & 9,261 Female). 4,489 of these workers experience a switch event (1,458 Male & 3,031 Female). There are 1745 transitions from a male manager to a female manager, 4291 from one male manager to another male manager. The within individual standard deviation of pay grade is 0.479. 95% C.I. displayed.

c. Odd to Even
minus Odd to Odd



This regression includes 381,028 observations of 14,645 workers (7,538 Even BD & 7,107 Odd BD). 4,705 of these workers experience a switch event (2,393 Even BD & 2,312 Odd BD). There are 3433 transitions from an odd-birthday manager to an even-birthday manager, 3086 from one odd-birthday manager to another odd-birthday manager. The within individual standard deviation of pay grade is 0.374. 95% C.I. displayed.

d. Non-Smoker to Smoker
minus Non-Smoker to Non-Smoker




This regression includes 93,534 observations of 2,894 workers (957 Smoking & 1,937 Non-Smoking). 912 of these workers experience a switch event (275 Smoking & 637 Non-Smoking). There are 287 transitions from a non-smoking manager to a smoking manager, 939 from one non-smoking manager to another non-smoking manager. The within individual standard deviation of pay grade is 0.519. 95% C.I. displayed, trimmed at -1 and 1.

Notes: The coefficients plotted above are triple difference estimates of the form $\beta_{DDD,t} = (\beta_{F2M,t}^M - \beta_{F2F,t}^M) - (\beta_{F2M,t}^F - \beta_{F2F,t}^F)$ where β^M and β^F are effects for male and female workers, respectively and $F2M, F2F$ are manager switch events from female to male managers and from one female manager to another, respectively. See 2 for a formal discussion of the event study specification and Figure 3 for difference-in-difference versions of the event study analysis. Coefficients in panels a) and b) are identified from the same regression. All specifications include employee, time, and time-by-(gender(a,b) birthday(c) smoker-status(d)) fixed effects. Additionally, all event-study coefficients are estimated by month, and are linearly smoothed to create quarter estimates; events outside the window are “absorbed” by indicators $\mathbf{1}(\text{Event Time} \geq 31 \text{ months})$ and $\mathbf{1}(\text{Event Time} \leq -31 \text{ months})$. See the notes below each panel for details about sample size, event counts, etc. Results for employees experiencing even to odd manager switches and smoker to non-smoker switches are available in Figure G.5.


Figure 5: Triple Difference Effects on Other Outcomes: Female to Male versus Female to Female

a. Firm Exit




This regression includes 524,785 observations of 15,827 workers (5,532 Male & 10,295 Female). 3,209 of these workers experience a switch event (828 Male & 2,381 Female). There are 1872 transitions from a female manager to a male manager, 2143 from one female manager to another female manager. The within individual standard deviation of left firm is 0.177. 95% C.I. displayed.

b. Log(Days Worked)




This regression includes 355,344 observations of 14,162 workers (4,913 Male & 9,249 Female). 2,971 of these workers experience a switch event (762 Male & 2,209 Female). There are 1682 transitions from a female manager to a male manager, 1994 from one female manager to another female manager. The within individual standard deviation of ln(days worked) is 0.166. 95% C.I. displayed, trimmed at $-.3$ and $.3$.

c. Log(Work Hours)



This regression includes 104,231 observations of 4,876 workers (1,881 Male & 2,995 Female). 982 of these workers experience a switch event (285 Male & 697 Female). There are 386 transitions from a female manager to a male manager, 801 from one female manager to another female manager. The within individual standard deviation of ln(avg daily hours worked) is 0.208. 95% C.I. displayed, trimmed at $-.4$ and $.4$.

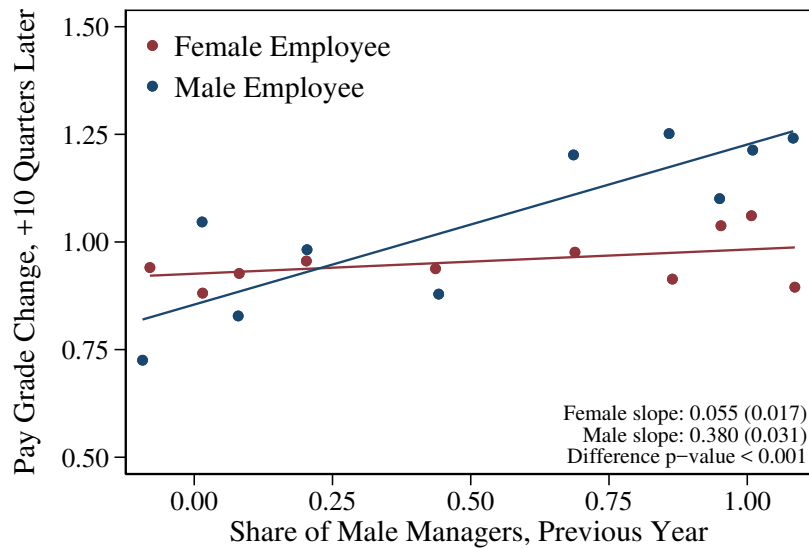
d. Sales Revenues



This regression includes 380,964 observations of 14,638 workers (5,193 Male & 9,445 Female). 1,956 of these workers experience a switch event (463 Male & 1,493 Female). There are 1398 transitions from a female manager to a male manager, 1036 from one female manager to another female manager. The within individual standard deviation of sales is 95.1. 95% C.I. displayed, trimmed at -200 and 200 .

Notes: The coefficients plotted above are triple difference estimates of the form $\beta_{DDD,t} = (\beta_{F2M,t}^M - \beta_{F2F,t}^M) - (\beta_{F2M,t}^F - \beta_{F2F,t}^F)$ where β^M and β^F are effects for male and female workers, respectively and $F2M, F2F$ are manager switch events from female to male managers and from one female manager to another, respectively. See Section 2 for a formal discussion of the event study specification. All specifications include employee, time, and time-by-gender fixed effects. Additionally, all event-study coefficients are estimated by month, and are linearly smoothed to create quarter estimates; events outside the window are “absorbed” by indicators $\mathbf{1}(\text{Event Time} \geq 31 \text{ months})$ and $\mathbf{1}(\text{Event Time} \leq -31 \text{ months})$. See the notes below each panel for details about sample size, event counts, etc. See Section 4.4 for further discussion of the construction of these outcome variables. Results for employees experiencing male to female manager switches and non-smoker to smoker switches are available in Figure G.12 and Figure G.13, respectively.

Figure 6: Association between Past Exposure to Male Managers and Future Changes in Pay Grade




Notes: We present binned scatter plots with linear trend lines of the change in pay grade over the following 10 quarters against the share of managers in the previous year that are male. This restricts our panel to workers are in the panel for at least 14 quarters. This is because we require that they have already been at the bank for one year so that the share of male manager is always computed over exactly 12 months. This is simply the share of months in the last year that an employee worked under a male manager. The change in pay grade outcome then mechanically requires that a worker remain in the panel for an additional 10 quarters, so that we can observe the change in pay grade at the end of that period. We control for tenure and include pay grade fixed effects, as we may expect promotions to be more or less likely based on the current pay grade. The difference p-value corresponds to the null hypothesis that the male and female slopes are the same.

Figure 7: Triple Differences Effects on Pay Grade by Proximity to the Manager: Female to Male versus Female to Female


PHYSICAL PROXIMITY TO MANAGER

a. Closer



This regression includes 367,581 observations of 14,288 workers (4,999 Male & 9,289 Female). 1,065 of these workers experience a switch event (343 Male & 722 Female). There are 679 transitions from a female manager to a male manager, 591 from one female manager to another female manager. The within individual standard deviation of pay grade is 0.475. 95% C.I. displayed, trimmed at -1 and 1.

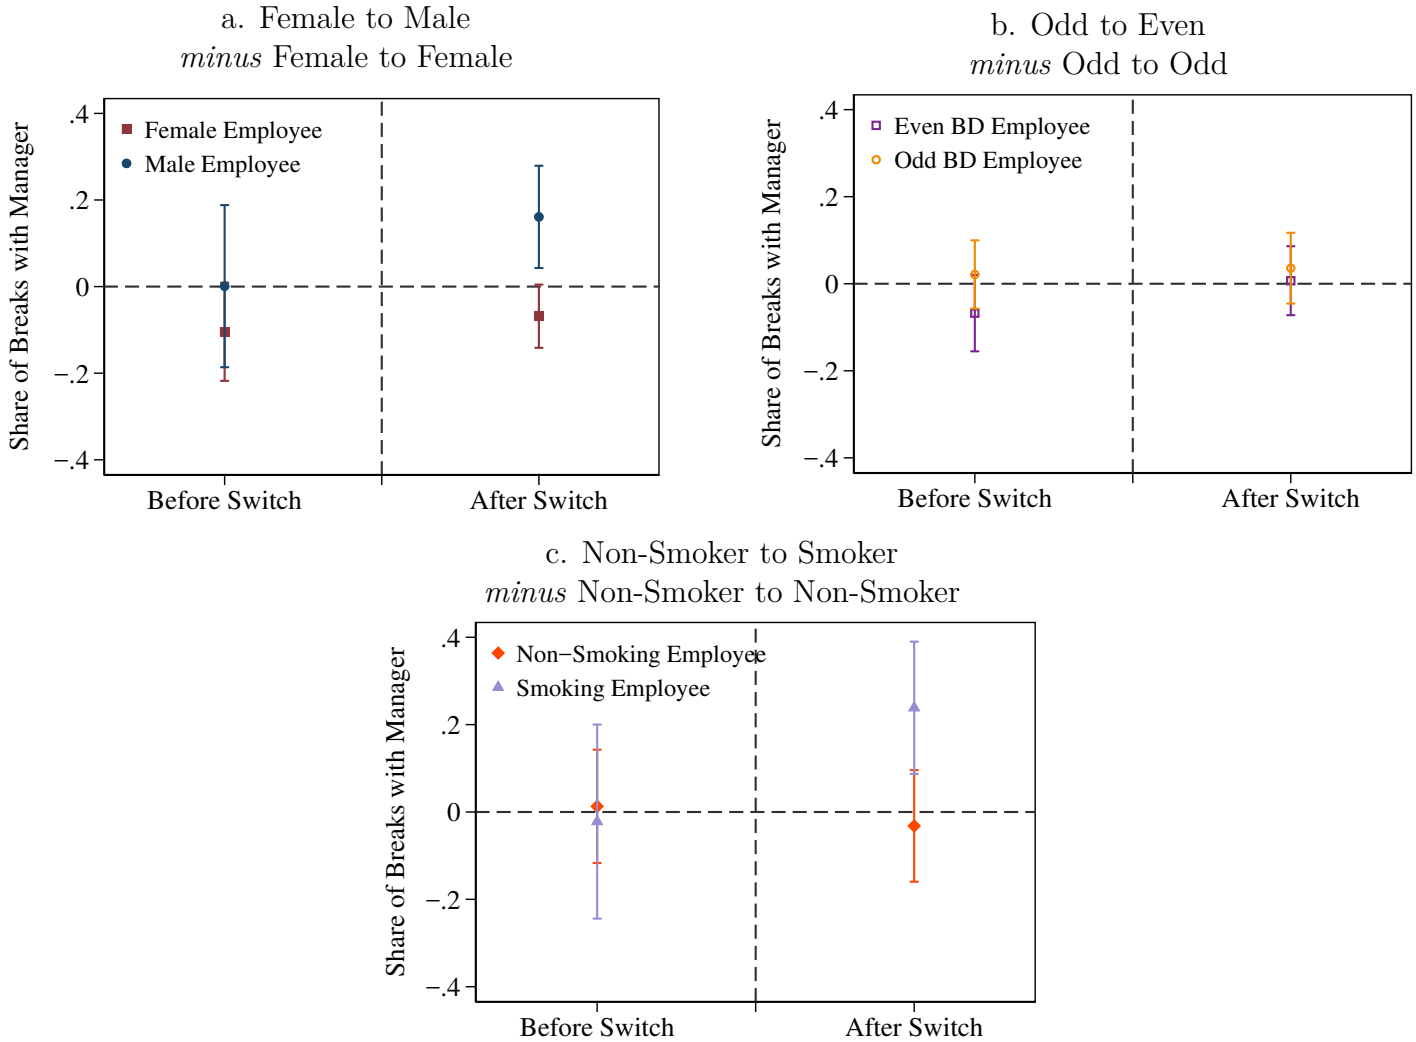
b. Farther



This regression includes 367,581 observations of 14,288 workers (4,999 Male & 9,289 Female). 974 of these workers experience a switch event (149 Male & 825 Female). There are 703 transitions from a female manager to a male manager, 474 from one female manager to another female manager. The within individual standard deviation of pay grade is 0.475. 95% C.I. displayed, trimmed at -1 and 1.

Notes: The coefficients plotted above are triple difference estimates of the form $\beta_{DDD,t} = (\beta_{F2M,t}^M - \beta_{F2F,t}^M) - (\beta_{F2M,t}^F - \beta_{F2F,t}^F)$ where β^M and β^F are effects for male and female workers, respectively and $F2M, F2F$ are manager switch events from female to male managers and from one female manager to another, respectively. See Section 2 for a formal discussion of the event study specification. All specifications include employee, time, and time-by-gender fixed effects. Additionally, all event-study coefficients are estimated by month, and are linearly smoothed to create quarter estimates; events outside the window are “absorbed” by indicators $\mathbf{1}(\text{Event Time} \geq 31 \text{ months})$ and $\mathbf{1}(\text{Event Time} \leq -31 \text{ months})$. We code positions as “closer” (“farther”) if the average worker in that position self-reports spending more than (less than) 4.5 days a week working in close physical proximity with their manager. We then categorize events as high or low proximity based on the position of the worker in the month they experience a switch event. We estimate the high and low proximity event coefficients on the same regression. To normalize across groups, we depart from our usual practice of normalizing with respect to an omitted period $t \in \{-1, -2, -3\}$. Rather, we explicitly estimate coefficients for the preperiod $t \in \{-1, -2, -3\}$ separately for the high and low proximity groups, and explicitly difference the smoothed estimate for $q = -1$ out of each panel. This is analogous to including group fixed effects in a linear regression; we simply subtract some constant so that the reference period in each panel is equal to 0.

Figure 8: Effects of Manager Switches on the Share of Breaks Taken with the Manager



Notes: 95% C.I displayed in brackets. Regression results with the share of breaks. The within-individual standard deviation of this outcome is 0.174. See Section 5.3 for full econometric specification. Panel (a): This regression includes 4,843 observations of 2,638 workers (698 Male & 1,940 Female). 430 of these workers experience a switch event (83 Male & 347 Female). There are 254 transitions from a female manager to a male manager, 243 from one female manager to another female manager. Panel (b): This regression includes 4,947 observations of 2,648 workers (1,296 Odd BD & 1,352 Even BD). 670 of these workers experience a switch event (330 Odd BD & 340 Even BD). There are 403 transitions from a odd-birthday manager to a even-birthday manager, 321 from one odd-birthday manager to another odd-birthday manager. Panel (c): This regression includes 1,287 observations of 699 workers (176 smoker & 523 Non-smoker). 196 of these workers experience a switch event (51 smoker & 145 Non-smoker). There are 50 transitions from a non-smoker manager to a smoker manager, 160 from one non-smoker manager to another non-smoker manager.

Table 1: Categorization of the Reasons Behind the Manager Switches

Outgoing	Incoming Manager		
	New Hire	Promotion	Lateral Move
Quit	29	53	89
Promotion	28	79	56
Lateral Move	91	154	324

Notes: We say that an outgoing (incoming) manager quit (was hired) if they quit (were hired) in the six months after (before) the switch. Similarly, we code a switch a promotion if there is a change in pay grade in the three months before or after the event.

Table 2: Characteristics of the Managers and Employees, by Type of Manager Switch

EMPLOYEES	Had Event?		Female to ...		Male to ...	
	No	Yes	Female	Male	Female	Male
Unique Employees	8,200	6,536	1,759	1685	1,671	3,305
Pay Grade	49.065 (2.74)	48.822 (2.56)	49.066 (2.52)	48.586 (2.52)	49.129 (2.69)	48.679 (2.52)
Male (%)	0.371 (0.48)	0.292 (0.45)	0.239 (0.43)	0.273 (0.45)	0.249 (0.43)	0.344 (0.48)
Age	29.844 (5.46)	30.084 (5.30)	30.223 (5.58)	29.537 (5.03)	30.690 (5.39)	30.005 (5.21)
College (%)	0.851 (0.36)	0.853 (0.35)	0.870 (0.34)	0.845 (0.36)	0.864 (0.34)	0.843 (0.36)
MANAGERS (INCOMING)	Had Event?		Female to ...		Male to ...	
	No	Yes	Female	Male	Female	Male
Unique Incoming Managers	480	706	215	162	177	265
Pay Grade	53.478 (2.10)	53.640 (2.14)	53.882 (2.18)	54.031 (2.07)	53.243 (2.07)	53.597 (2.21)
Male (%)	0.454 (0.50)	0.548 (0.50)	0.000 (0.00)	1.000 (0.00)	0.000 (0.00)	1.000 (0.00)
Age	36.856 (5.30)	35.437 (4.34)	35.445 (4.44)	35.714 (4.61)	35.328 (3.79)	35.427 (4.39)
College (%)	0.958 (0.20)	0.928 (0.26)	0.937 (0.24)	0.901 (0.30)	0.917 (0.28)	0.941 (0.24)
MANAGERS (OUTGOING)	Had Event?		Female to ...		Male to ...	
	No	Yes	Female	Male	Female	Male
Unique Outgoing Managers	536	650	206	120	137	269
Pay Grade	53.205 (1.87)	53.866 (2.24)	53.845 (2.13)	53.754 (2.31)	54.448 (2.39)	53.666 (2.20)
Male (%)	0.427 (0.49)	0.592 (0.49)	0.000 (0.00)	0.000 (0.00)	1.000 (0.00)	1.000 (0.00)
Age	36.330 (4.59)	36.013 (4.44)	36.122 (4.42)	36.115 (5.01)	36.241 (4.59)	35.813 (4.16)
College (%)	0.946 (0.23)	0.930 (0.26)	0.933 (0.25)	0.925 (0.26)	0.917 (0.28)	0.936 (0.24)

Notes: Average characteristics, with standard deviations in parentheses. Note that some employees/managers experience multiple events. For employees/managers who experience an event, we compute the average characteristic in the month of the corresponding event. For those who never experience an event (first column) we show the average over their whole tenure.