

NBER WORKING PAPER SERIES

THE OLD BOYS' CLUB:
SCHMOOZING AND THE GENDER GAP

Zoë B. Cullen
Ricardo Perez-Truglia

Working Paper 26530
<http://www.nber.org/papers/w26530>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2019

We are thankful for excellent comments from several colleagues and seminar discussants. We thank seminar participants at Berkeley-Econ, Berkeley-Haas, UCSD, Massachusetts Institute of Technology Sloan, University of British Columbia, University of Southern California, Tufts University, INSEAD, Columbia University Micro and Macro Labor Conference, Harvard Business School Behavioral Conference, American Economic Association Annual Meetings and the Discrimination and Disparities intercollegiate seminar. We are especially thankful to Leonardo Bursztyn, Thomas Fujiwara and Amanda Pallais. The collaborating institution provided financial support for the research being conducted. Additionally, Zoe Cullen was a full-time, salaried employee at that institution while the research was being conducted. Dylan Balla-Elliott, Katherine Fang, Anh Nguyen, Giacomo Stazi, Andrew Kao, Aakaash Rao and Jenna Anders provided excellent research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Zoë B. Cullen and Ricardo Perez-Truglia. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Old Boys' Club: Schmoozing and the Gender Gap
Zoë B. Cullen and Ricardo Perez-Truglia
NBER Working Paper No. 26530
December 2019, Revised in September 2020
JEL No. J01,J16,J7,Z1,Z13

ABSTRACT

Offices are social places. Employees and managers take coffee breaks together, go to lunch, hang out over drinks, and talk about family and hobbies. In this study, we provide evidence that employees' social interactions with their managers can be advantageous for their careers and that this phenomenon can contribute to the gender pay gap. We use administrative and survey data from a large financial institution. We conduct an event-study analysis of manager rotation to estimate the causal effect of managers' gender on their employees' career progressions. We find that male employees assigned to male managers were promoted faster in the following years than male employees assigned to female managers; female employees, on the contrary, had the same career progression regardless of their managers' gender. These differences were not accompanied by any differences in effort or performance, and they explain a third of the gender gap in promotions at this firm. Then, we provide evidence suggesting that these effects were mediated by the social interactions between male employees and male managers. First, we show that the effects were present only among employees who worked in close proximity to their managers. Second, we show that the effects coincided with an uptick in the share of breaks taken with the managers. Third, we estimate the impact of social interactions on career progression using quasi-random variation induced by smoking habits. When male employees who smoke transitioned to male managers who smoke, they took breaks with their managers more often and were subsequently promoted at higher rates than male smokers who transitioned to non-smoking managers. The boost in socialization and promotion rates closely mirrors the pattern among male employees assigned a male manager.

Zoë B. Cullen
Rock Center 210
Harvard Business School
60 N. Harvard
Boston, MA 02163
zcullen@hbs.edu

Ricardo Perez-Truglia
Haas School of Business
University of California, Berkeley
545 Student Services Building #1900
Berkeley, CA 94720-1900
and NBER
ricardotruglia@berkeley.edu

An Online appendix is available at <http://www.nber.org/data-appendix/w26530>

1 Introduction

Workplaces are social places. Employees and managers often discuss all sorts of non-work related topics, such as sports, family, and movies. These personal interactions extend outside of office hours, such as during lunch, smoking, or coffee breaks. Through these interactions, employees form social bonds with their managers. In this study, we explore whether these social bonds influence employees' careers and whether they can help explain the gender pay gap.

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

A growing literature has investigated what causes women to lag behind men in the corporate world. According to the "old boys' club" hypothesis, this gap arises in part because men can schmooze, network, and interact with more powerful men in ways that are less accessible to women.¹ This mechanism can create a self-perpetuating cycle: male managers promote a disproportionate share of male employees, who continue promoting other men.

Ample anecdotal evidence suggests that the old boys' club is real (Lang, 2011; Lee, 2014; Elting, 2018). For example, 81% of women say that they feel excluded from relationship-building at work, and many also feel excluded from after-work hours socializing (Gray and Barbara, 2013). Some women even believe that being able to use the men's bathroom would give them an advantage at work (Lee, 2014). Despite all the anecdotes, however, there is little quantitative evidence showing that the old boys' club exists. The self-selection of those who engage in social activities creates a number of research challenges to isolating the impact of social interactions. In this study, we propose a quasi-experimental approach for testing this hypothesis and provide novel evidence based on data from a large financial organization.

We partnered with a large commercial bank in Asia (referred to hereinafter as *the firm*) with millions of customers, billions of dollars in assets and in revenues, and thousands of employees. The firm is typical in that female representation drops off at higher levels: 75% of entry-level employees are women, 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 the firm 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. These records include the employees' pay grades, the floor their desks are on, the managers to which they were assigned, as well as measures of effort and performance. We also conducted a series of surveys to measure other

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

aspects of the employees' lives, such as whether they take breaks with their managers, whether they know the manager's favorite sports team, as well as their smoking status to supplement annual health exam reports.

We start by measuring the effect that the manager's gender has on the careers of the employees working under that manager. 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. According to the old boys' club prediction, those assigned to a male manager would benefit the careers of the male employees more than the female employees. Obviously, it would be much too costly and disruptive for any real-world company to randomly shuffle its employees and managers. Instead, we exploit the naturally occurring rotation of managers between teams. These manager transitions are not literally decided by a coin toss, but anecdotal evidence and supporting empirical tests suggests that they can be as good as random.

Our identification strategy is based on event-study analysis. The identification leverages the timing of manager transitions and comparisons between different types of transitions. For example, consider two teams, each managed by a female manager. One of these teams then transitions from the female manager to a male manager, and the other team transitions 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 transition date and each month after the transition. As both teams are affected by a manager transition, this design nets out the effect of the transition. The hypothesis is that transitioning to a male manager, relative to transitioning to the female manager, results in better promotion prospects for the male employees but has no effect (or little effect) on the promotions of female employees.

We focus on manager transitions that are out of the control of the employee. The typical case is a manager rotating laterally to a different team. Our data comprises 8,670 transition events involving 6,021 unique employees and 690 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., transitioning from a female to a male manager) are largely unrelated to the characteristics of the employee, the incoming manager and the outgoing manager.

We find that male employees are promoted more quickly after they transition from a female to a male manager: at 10 quarters after such a manager transition, male employees' pay grades were 0.60 points (p-value = 0.003) higher than those of male employees who transitioned from a female manager to a different female manager. This 0.60 point increase in pay grade is equivalent to a 14.6% increase in salary. By contrast, female employees experienced similar promotion rates regardless of whether they transitioned from a male manager to a female manager or from a male manager to another male manager.

We provide two main robustness checks for our identification strategy. First, we analyze the reverse transition. In the baseline results presented above, we look at employees who "lose" a male manager (i.e., transitioning from a male manager to a female manager versus transitioning from a male manager to a different male manager). Next, we look at employees who "gain" a male manager (i.e., transitioning from a female manager to a male manager versus transitioning from a female manager to a different female manager). The expectation is that the effects of gaining

a male manager should be roughly a mirror image the effects of losing a male manager, in terms of both timing and magnitude. This is a sharp test, in the sense that the coefficients are identified by a disjoint set of transition events and thus there are no “mechanical” reasons why the results should mirror each other. Indeed, we find that the effects of losing a male manager are in the opposite direction of the effects of gaining a male manager, and they 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 that is 0.30 points lower at 10 quarters later (p-value = 0.032), whereas the evolution of pay grades for female employees is unrelated to the manager’s gender.

The second 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. In other words, we would not expect that managers born on an odd date would be beneficial to the careers of their subordinates. We reproduce the whole event-study analysis, but instead of slicing the data based on manager and employee genders, we focus on their birth dates. This test rules out mechanical reasons why our event-study framework would generate spurious effects, and allows us to assess whether our standard errors are adequate. As expected, we find that the estimates are close to zero, statistically insignificant, and precisely estimated.

We define the male-to-male advantage as the effect of male managers (relative to female managers) on the careers of male employees (relative to female employees). Our preferred estimate, based on the transitions in both directions, indicates that the male-to-male advantage in pay grade is highly statistically significant (p-value<0.001) and economically large (0.65 pay grades at 10 quarters after the event). In back of the envelope calculations, we estimate that removing the male-to-male advantage would reduce the gender gap in pay grades by 40%.²

We show that the male-to-male advantage cannot be explained by differences in retention or performance. One potential explanation is that male managers are better at retaining male employees. However, when we estimate the effects of manager transitions on the probability of staying at the firm, we find point estimates that are close to zero, statistically insignificant, and precisely estimated. Another potential explanation is that male employees work harder and more productively under male managers than they would under female managers. For example, male managers might be better than female managers at motivating and monitoring male employees, or male employees may be more responsive to the directions of their male bosses. Contrary to this interpretation, when we estimate the effects of the manager transitions on measures of effort (the number of days worked and the number of hours spent in the office) and performance (the employee’s own sales revenues), we find point estimates that are close to zero, statistically insignificant, and precisely estimated.

Next, we provide evidence on the role of social interactions in the male-to-male advantage in promotions. For example, male employees may use their interactions to gain their managers’ sympathy and favor. Male managers also may learn more about their male employees during the interactions and thus be better able to identify their potential. Such interactions may make the

² For more details, see Section 4.6.

accomplishments and efforts of male employees more noticeable to the manager or may give the male employees opportunities for self-promotion. During interactions with their managers, male employees may learn useful information, such as which tasks or training are more conducive to promotions.

The first test of the social interactions channel exploits the fact that physical proximity is a necessary condition for social interactions. Alternative channels for the male-male advantage, such as statistical discrimination or in-group biases, should not depend on physical proximity. If driven by socialization, the male-to-male advantage should be stronger when manager and employee pairs work in close physical proximity; by contrast, the effects should be smaller, or even null, when the manager does not work in physical proximity to the employee. We categorize positions by the physical proximity with the manager using administrative data on office locations as well as survey data asking employees if their managers work in physical proximity. Consistent with the social interactions channel, we find that the male-to-male advantage is large and statistically significant when the managers and employees work in close physical proximity but approximately zero and statistically insignificant if they do not work in close proximity.

Second, we collect survey data on the frequency of social interactions between employees and their managers. We ask a sub-sample of the firm to report how often they share breaks with their managers. Finding a male-to-male advantage in this form of social interactions would constitute suggestive evidence of the schmoozing mechanism. Indeed, we find that male employees are significantly more likely to share work breaks with their manager after transitioning from a female manager to a male manager (relative to transitioning from a female manager to another female manager). Female employees, on the contrary, are equally likely to spend breaks with male and female managers.

Third, we provide evidence that social interactions in this firm translate into a promotion advantage among male employees. In the ideal experiment, we would flip of a coin to decide which male employees get to socialize more with their male managers. According to the schmoozing channel, the male employees who get to socialize more with their male managers would be promoted faster. While the ideal experiment is not feasible, we exploit quasi-experimental variation based on co-smoking habits. We collected data on the smoking habits of the employees and their managers. 33% of male employees and 37% of male managers smoke.³ We conduct an event-study analysis of the rotation of managers, but this time focusing on their smoking habits instead of their gender. We conjectured that when a male employee who smokes transitions to a male manager who also smokes, they will interact more because of shared smoking breaks. And, according to the schmoozing channel, that increase in social interactions will translate into higher promotion rates.⁴

Consistent with our conjecture, we find a smoker-to-smoker advantage in the frequency of social interactions. After transitioning from a non-smoking manager to a smoking manager (relative to transitioning to another non-smoking manager), smoking employees end up spending more breaks with their new managers; in contrast, there is no effect on non-smoking employees. Indeed, the

³ The smoking rates are negligible among women, and thus we focused this analysis on males only.

⁴ This part of the analysis is based on a sub-sample (males for whom we can infer smoking status) comprising 1,094 unique employees, 250 unique managers and 1,499 unique manager transitions.

magnitude of the smoker-to-smoker advantage in shared breaks is similar to the corresponding male-to-male advantage reported above. Most important, we show that these manager transitions affect promotion rates too: after transitioning from a non-smoking manager to a smoking manager (relative to transitioning to another non-smoking manager), the smoking employees are promoted faster; in comparison, there is no effect on the pay grade of non-smoking employees. This evidence indicates that, consistent with the schmoozing channel, the increased social interactions caused by co-smoking translates into higher promotion rates. By contrast, the increased social interactions do not affect performance (employees' sales revenue) or effort (hours spent working in the office or absenteeism).

This study is related to various strands of literature. Most important, it is related to a literature on the role of social interactions at work. Despite the universality of socializing in the workplace, relatively little is known about the returns of these personal interactions and whether these returns differ by gender. Cai and Szeidl (2018) provide experimental evidence that increasing the connections between business owners can increase firm productivity. There is also evidence of spillovers between business school classmates and executives (Shue, 2013; Lerner and Malmendier, 2013; Agarwal et al., 2016; Field et al., 2016).⁵ There is evidence that the managers' social skills affect employee turnover Hoffman and Tadelis (2020). In the context of fruit-pickers, managers with fixed pay will favor workers with whom they share a connection, to the detriment of firm productivity (Bandiera et al., 2010, 2009). And in the context of politics, public officials may capitalize on their political and personal networks to gain influence (Cruz and Tolentino, 2019; Xu, 2018; Bertrand et al., 2018; Voth and Xu, 2019).

We contribute to this literature by providing novel evidence on the career and productivity consequences of social interactions in the corporate world. This is a context for which there is abundant anecdotal evidence on the importance of social interactions and its effects on the gender pay gap, yet little quantitative evidence. The lack of evidence is probably due to data challenges (e.g., personal interactions are difficult to measure and also sensitive information) as well as challenges with causal identification (e.g., social interactions are highly endogenous). We address both of these challenges. First, we provide causal evidence based on quasi-experimental variation in the gender and smoking habits of the managers. Furthermore, we collected unique sources of administrative and survey data about social interactions and physical proximity in a real corporation that spans culturally distinct regions.

Our paper more broadly contributes to the large literature on the gender wage gap (Goldin, 2014). There is a consensus that the majority of this gap is due to differences in promotion rates (Bertrand, Goldin, and Katz, 2010; Manning and Swaffield, 2008; Goldin, Kerr, Olivetti, and Barth, 2017). By one careful account, the gap in internal promotion rates can account for approximately 70% of the gender pay gap by the age of forty-five (Bronson and Thoursie, 2019). Several explanations have been provided for these differences in promotions.⁶ Most related to our study, there

⁵ Another related study is Lleras-Muney et al. (2019), showing that friendships accumulated during high school can have lasting impacts on labor market outcomes. Also, Mengel (Mengel) use a laboratory experiment to show that men and women both engage in networking but men develop closer connections.

⁶ Some examples include the marriage market incentives (Bursztyan, Fujiwara, and Pallais, 2017), cultural norms

are some studies on the role of the gender of superiors in the education industry, with mixed results. On the one hand, male teachers in public schools are more satisfied with their jobs and more likely to remain working at a school if it is has a male, rather than female, principal (Grissom, Nicholson-Crotty, and Keiser, 2012; Husain, Matsa, and Miller, 2018). On the other hand, female referees and female committee members in academia do not increase the odds of acceptance of female-authored papers or promotion of female candidates (Bagues, Sylos-Labini, and Zinovyeva, 2017; Card, Dellavigna, Funk, and Iriberry, 2019).⁷

We contribute to this literature by showing that the gender of managers can be a major source of gender pay gaps in the corporate world – this mechanism explains around one-third of the gender pay gap in the firm that we study. To the best of our knowledge, Kunze and Miller (2017) provides the only related evidence in the context of a corporation.⁸ The authors use data from a private firm in Norway to measure the association between the gender of managers and the outcomes of their employees. They found that the gender gap in promotions is significantly larger in establishments with a higher share of male superiors. We contribute to this literature in at least three ways. First, we provide causal estimates with the use of quasi-experimental methods. Second, we provide novel evidence about a specific mechanism, social interactions, for which there is abundant anecdotal evidence yet it has been largely ignored in the literature on the gender pay gap. Third, we provide evidence on the limited productivity advantages of socialization between managers and employees despite the fact that they generate career advantages for certain employees.

Although we offer evidence from a specific firm, the firm’s establishments span regions with distinct cultural heritage, allowing us to examine heterogeneity by gender norms. Moreover, our methodology is not specific to our setting. The rotation of managers is a common practice in large organizations, and the data on pay grades, assignments, and demographics could be obtained for most firms. Thus, our research design can be applied in other firms from different industries and countries to identify the contexts in which the male-to-male advantage is most pervasive. Our study already provides suggestive evidence that the male-to-male advantage may be exacerbated in some occupations (i.e., in which the manager and the employee work in close proximity with each other) and regions (i.e., where stronger gender norms prevail).

The rest of the paper proceeds as follows. Section 2 summarizes 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 the results. Section 6 concludes.

(Bursztyn, Gonzalez, and Yanagizawa-Drott, 2018; Alesina, Giuliano, and Nunn, 2013; Jayachandran, 2020), recognition for group work (Sarsons, 2017; Isaksson, 2019; Sarsons et al., 2019), differences in aspirations and performance (Azmat and Ferrer, 2017), the child penalty (Schönberg and Ludsteck, 2014; Bertrand et al., 2010; Kleven et al., 2019; Kuziemko et al., 2018), preference for flexible hours (Wasserman, 2018) and household work more generally (Cortés and Pan, 2019).

⁷ Other related studies look at gender roles among peers instead of among managers (Dahl et al., 2018; Hill, 2017; Karpowitz et al., 2020) and at the role of other demographics besides gender, such as race (Mas and Moretti, 2009; Bandiera et al., 2010; Giuliano et al., 2011; Hjort, 2014; Glover et al., 2017).

⁸ A related literature studies whether female representation at the very top of the firm, such as owners, CEOs, and chairs, can affect the female employees working at those companies (Bell, 2005; Bertrand et al., 2019; Cardoso and Winter-Ebmer, 2010; Dalvit et al., 2018; Flabbi et al., 2019).

2 Research Design

2.1 Conceptual Framework

Our analysis revolves around the effects of manager characteristics on the subsequent career progressions of their employees. For example, we want to measure whether male employees fare better after transitioning from a female to a male manager and whether employees who smoke are promoted faster when they transition from a non-smoking to a smoking manager. To estimate these manager effects, ideally we would randomize employees to their managers. As this type of experiment is not feasible, we instead exploit naturally occurring variations in manager assignments generated by the rotation of managers within the organization. Rather than assuming that these natural manager transitions are as good as random changes, we test that assumption using an event-study analysis. The formal econometric framework for the event-study analysis is provided below.

2.2 Effects of Manager's Gender

Let $y_{i,t}$ be a generic outcome, where the subscripts i and t denote employees and time, respectively. The main outcome in our analysis is the employee's pay grade, but we also consider other outcomes such as firm exit, effort, and performance.

The transition between two managers can result in one of four different types of gender transitions. Let J_G denote the set of these types: $J_G = \{F2M, F2F, M2F, M2M\}$, where $F2M$ denotes a transition from a female manager to a male manager, $F2F$ denotes a transition from one female manager to another female manager, and so on. Let $D_{i,t}^j$ denote the traditional event-study variables that indicate the periods leading up to and following a transition event. For example, $D_{i,t+s}^j$ is an indicator variable that equals 1 if individual i experiences an event of type j in period $t+s$.

The event-study regression relates the outcome variable to the event-study dummies:

$$y_{i,t} = \sum_{j \in J_G} \sum_{s \in S} \beta_{j,s}^F \cdot F_i \cdot D_{i,t+s}^j + \sum_{j \in J_G} \sum_{s \in S} \beta_{j,s}^M \cdot (1 - F_i) \cdot D_{i,t+s}^j + \gamma_i + \eta_{i,t} + \delta_t^F + \delta_t^M + \varepsilon_{i,t} \quad (1)$$

Note that we interact the event-study dummies with a gender indicator (F_i) to estimate event-time coefficients for men ($\beta_{j,t}^M$) and women ($\beta_{j,t}^F$) separately. The set S , the event-study window, spans from 30 months before the event to 30 months after the event (this time window is due to the length of our panel data). We include the usual absorbing dummies at extremes of ≤ -31 and $\geq +31$ months (Stevenson and Wolfers, 2006). In the event-study graphs, we aggregate these monthly coefficients to the quarterly level for ease of presentation. The omitted categories in S are the three months prior to the event (i.e., -3, -2, and -1 months), aiding our visual depiction at the quarterly frequency. This baseline specification includes employee fixed effects (γ_i), manager fixed effects ($\eta_{i,t}$) and gender-specific month effects (δ_t^M and δ_t^F). In this study, 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 always compare employees undergoing manager transitions where one of those transitions results in a change of manager gender and the other does not. For example, we compare the effects of transitioning from a female manager to a male manager versus the effects of transitioning from a female manager to a different female manager. In the case of male employees, the object of interest is $\beta_{F2M,s}^M - \beta_{F2F,s}^M$, where s indicates the time since (or until) the transition date. In the case female employees, the corresponding object of interest is $\beta_{F2M,s}^F - \beta_{F2F,s}^F$. Hereinafter, we refer to these objects as the *single-differences*, because they are differences between types of transitions.

What we capture with the single-difference estimates is the impact of receiving a male manager relative to the impact of receiving a new female manager. However, we are ultimately interested in whether the effects of manager gender differ for male and female employees. For example, if male managers increase pay grades for male and female employees alike, that would not constitute evidence of a male-to-male advantage. Thus, we must take the difference of the single-difference estimates between male and female employees: $(\beta_{F2M,s}^M - \beta_{F2F,s}^M) - (\beta_{F2M,s}^F - \beta_{F2F,s}^F)$. A positive difference would be consistent with a male-to-male advantage. We refer to these estimates as the *double-differences*, because they take differences first with respect to types of transitions and second with respect to the employee's own gender.

The key assumption is that, prior to the transitions, male and female employees were on the same pay-grade trajectories. The event-study framework provides a natural test of the identifying assumption: we can assess the evolution of the outcome in each month before the date of the transition to confirm if the trends were truly parallel before the event date.

The manager transitions provide an additional validation check, based on the principle that transitions in the opposite direction should result in approximately the opposite effects. In the previous example, we discussed the effects of “gaining” a male manager (i.e., what happens when an employee transitions from a female manager to a male manager, relative to what would have happened if the employee transitioned from a female manager to another female manager). Likewise, we can measure the effects of “losing” a male manager (i.e., what happens when an employee transitions from a male manager to a female manager, relative to what would have happened if the employee transitioned from a male manager to another male manager). The expectation is that the effects of gaining a male manager should roughly mirror the effects of losing a male manager in terms of timing and magnitude. Because these coefficients are identified by a disjointed set of transition events, there are no mechanical reasons why the results should mirror each other.

To maximize statistical power, we estimate the average male-to-male advantage using all four types of gender transitions. That is, we average the double-difference estimates from “gaining” a male manager and the (negative of) the double-difference estimates from “losing” a male manager: $\frac{1}{2}\{(\beta_{F2M,s}^M - \beta_{F2F,s}^M) - (\beta_{F2M,s}^F - \beta_{F2F,s}^F) - [(\beta_{M2F,s}^M - \beta_{M2M,s}^M) - (\beta_{M2F,s}^F - \beta_{M2M,s}^F)]\}$. We refer to this object as the *dual-double-difference*.

When interpreting the event-study results, there are a few caveats to keep in mind. First, our estimates measure a reduced form effect of an increased but likely transitory exposure to a given managerial gender. As time goes by, many reasons explain why an employee ends up with a manager of a different gender. For example, the employee may be promoted to a different position and assigned a manager of a different gender, 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. In practice, this is a minor concern, as we find gender transitions to be persistent over time.

A second caveat is that our framework cannot disentangle whether male managers are favorable to male employees or female managers are dis-favorable to male employees. Indeed, this challenge is not unique to our methodology or to our context. Even in a randomized controlled trial, we could compare male managers versus female managers only, because there are no gender-neutral managers to compare against. Likewise, the favorable or dis-favorable conditions may be due to the behavior of the employee, the behavior of the manager, or both. For example, male employees may do better under male managers because the managers treat them better or because the employees are more attentive to them.

2.3 Placebo: Effects of Manager’s Birthday-Evenness

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. This placebo test is designed to rule out mechanical reasons why our event-study framework would generate spurious effects. This placebo analysis can also be used to assess whether our standard errors are adequate: e.g., if we found statistically significant coefficients, it would suggest that the inference is misleading. Let O_i be an indicator variable that equals 1 if the employee was born on an even day and 0 otherwise. The regression of interest is identical to the main specification from equation (1), except that gender is replaced everywhere by the birthday-evenness:

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

where J_E is the set of manager transitions $J_E = \{E2O, E2E, O2E, O2O\}$: $E2O$ denotes a transition from a manager with an even birthday to a manager with an odd birthday, and so on. We identify analogous single-difference, double-difference, and dual-double-difference estimates for these placebo events. For example, the following single-difference estimate measures how the odd-birthday employee reacts to gaining an odd-birthday manager (i.e., transitioning from an even-birthday manager to an odd-birthday manager, relative to transitioning from an even-birthday manager to another even-birthday manager): $\beta_{E2O,s}^O - \beta_{E2E,s}^O$. We use the following double-difference estimate to measure the odd-to-odd advantage: $(\beta_{E2O,s}^O - \beta_{E2E,s}^O) - (\beta_{E2O,s}^E - \beta_{E2E,s}^E)$.

In the results section, we present an additional placebo test measuring if the transitions in the manager’s gender affect even-birthday and odd-birthday employees differentially.

2.4 Effect of Manager’s Smoking Habits

We also directly evaluate a non-gender shock to social interactions. Intuitively, we begin by restricting to the sample to male employees and male managers. We then compare two teams, each 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 the transitions for smoking employees and for non-smoking employees separately. The prediction is that transitioning to the smoking manager should benefit the subsequent career of the smoking employees, whereas it should not affect, or less prominently affect, 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:

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

where J_S is the set of the types of manager transitions $J_S = \{N2S, N2N, S2N, S2S\}$. For instance, $N2S$ denotes a transition from a non-smoking manager to a manager who smokes. Again, we define analogous single-difference, double-difference, and dual-double-difference estimates for these manager transitions. For example, the following single-difference estimate measures how smoker employees react to gaining a smoker manager (i.e., transitioning from a non-smoker manager to a smoker manager, relative to transitioning from a non-smoker manager to another non-smoker manager): $\beta_{N2S,s}^S - \beta_{N2N,s}^S$. Likewise, we can use the following double-difference estimate to measure the smoker-to-smoker advantage: $(\beta_{N2S,s}^S - \beta_{N2N,s}^S) - (\beta_{N2S,s}^N - \beta_{N2N,s}^N)$.

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 spans several culturally distinct regions and has millions of customers, billions of dollars in assets and in revenues, and thousands of employees. Although we do not claim that this firm is representative of all firms in the world, we have evidence that this is not an extreme context. The firm may be unusual for the financial sector in that a majority (64%) of its employees are female. Besides that, the gender gaps at this organization are 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%).⁹ The firm is typical in that men and women in a given position get paid similarly. The bulk of the gender pay gap thus is due to differences in positions among men and women. For example, 75% of firm employees at the entry-level are female, and that fraction falls to 61% in middle management, 25%

⁹ Results based on wage rates for men and women working in the financial sector in firms with over 1,000 employees, as reported in Yildirmaz et al. (2019).

at the C-Suite level, and 0% at the CEO level. Data for U.S. corporations suggest a similar drop from 48% of female employees in entry-level positions to 38% in middle management, 22% in C-Suite positions, and 5% in CEO positions (McKinsey & Company, 2019).

When looking at the regions where the firm operates as a whole, the gender gaps are similar to those in the United States. For example, the gender gap in labor force participation (8.5%) is similar to the one in the United States (13.2%).¹⁰ According to survey data, the gender norms also are not unusual. For example, data from the 2006 World Value Survey suggest that 12% of women in the firm's country describe work as unimportant or of little importance, and the respective share is 19% in the United States. In Section 4.5, we leverage variation in gender norms across culturally distinct regions where the firm operates to examine the mediating role of local norms.

3.2 Administrative Data: Pay Grade

We collaborated with different divisions 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 whom have been assigned to a manager role at some point. Finally, 64% of the employees are female, and 49% of the managers are female.

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, and the metric over which managers have the most direct input.¹¹ 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 associated with a promotion, or increase in responsibilities. Conditional on an increase in pay grade, there is an 84% chance of a change in position title; in comparison, there is a 1% chance of a change in position title when there is no pay grade increase. Variation in pay grade 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 during the full sample period of four years, 50% experienced at least one pay grade increase, and 16% experienced more than one increase.

Due to the sensitive nature of the data, we do not have the exact compensation details for the whole sample of employees. However, for a different project on a different topic (Cullen and Perez-Truglia, 2018), we have a cross-section of the pay grades and base salaries of employees in a given month (March of 2017). According to these data, there is a strong linear relationship between the logarithm of salary and the pay grade (results presented in Appendix A.1). The slope of the relationship (0.227) indicates that a 1-point increase in pay grade is associated with a 25% increase in salary ($= e^{0.227} - 1$). The R^2 of the regression (0.83) also indicates that pay grade explains the vast majority of variation in salaries.

¹⁰ Labor force participation data come from the World Bank Databank and International Labour Organization ILOSTAT database. These 2017 figures are the most recent for which male and female labor force participation data are available in both countries.

¹¹ HR personnel often carry out the precise salary negotiation within the range determined by pay grade, using market benchmark data.

Although the setting involves employees competing for promotions, employees are not necessarily competing with their teammates. There are no limits on the number of employees on a team that can be promoted, and different employees from the same team may seek promotions into different positions. Indeed, these employees compete for promotions with employees from other teams in the firm, and as the company routinely hires new employees, they also implicitly compete with outside candidates.¹²

3.3 Other Outcomes: Attrition, Effort, and Performance

We know the dates when the employees join and exit the company, which allows us to construct a dummy variable for employee attrition. We also have some measures of effort and performance. The first measure of effort is the number of days worked. We construct this measure using data from the human resources divisions on absences. We subtract the number of absences, including parental leave, sick days, and vacation days, from the total number of workdays in the month. We use an additional measure of effort to complement the administrative data: the number of hours spent in the office. However, we measure this outcome only for employees working in the headquarters offices (29% of the sample), as those employees clock in and out using an electronic card-swipe system that 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 performance 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 across all products (e.g., credit cards, loans, mortgages) by mapping each product to the expected revenue generated from the sale. We use these data to construct a monthly *sales performance index*.

3.4 Manager Assignments

Because a single employee may consider more than one person to be his or her manager, we identify the most relevant manager as the one who has the most power over the employee. We use longitudinal data from the firm's organizational chart to link each employee to a manager in each month that the employee appears in the sample. The employee-manager assignment is constructed using a simple, two-step algorithm: identify the employee's team, and then identify the "director" of that team.¹³ To validate our manager assignments, we conducted a survey of the sales and distribution division (described in Section 3.6). We asked employees to identify the 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 also are reported by the employees to be their managers.¹⁴

¹² More specifically, there is both a high employee turnover (12.5% yearly) and growth in the number of employees (5.9% yearly).

¹³ In cases where the team has no directors listed in the organizational chart, we assign the team to the director listed at the next highest level in the organizational chart hierarchy.

¹⁴ Our comparison is restricted to pairs in the administrative organization chart that remain together for one year or more. When we include all pairs, even those who have been together for just one month, we still find substantial

The managers tend to be significantly above their subordinates in the firm's hierarchy. For example, the modal (mean) distance between managers and their employees is 5 (5.3) pay grades. The manager typically can influence the careers and daily lives of the employees in various ways. Most importantly, the manager provides key input in decisions to promote employees. Even if the employee is not promoted, the manager still provides input that influences employee raises and bonuses. The manager also has discretion to distribute workload across team members. 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.

3.5 Manager Transitions

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 may be promoted to a higher position and thus assigned to a different team with a different manager, or an employee who dislikes his or her manager may ask to be transferred to another team. In this study, we focus on manager transitions that are outside the control of the employee. The most typical case occurs when managers rotate laterally across different teams, but also include instances when the team's manager is promoted to a higher position or accepts a position at a different firm and a replacement is necessary.

In identifying these exogenous transition events in the data, we impose a few conditions. We require that the new manager must assume responsibility for all employees in the team. In other words, the whole team, rather than a specific employee, experiences the manager transition. We also exclude managers who are temporary replacements by requiring the new manager to remain with the team for at least one quarter, and we exclude events in which more than half of the team members changed at the transition event. In the results section, we show that the results are robust under different criteria for the definition of the transition events.

We use the event-study framework to assess whether the manager transitions are exogenous. Anecdotal evidence suggests that this exogeneity is plausible. As part of corporate strategy, managers are expected to gain experience in all areas of banking. For this reason, managers are transitioned across teams within divisions and across divisions to gain exposure to new people and activities; for example, a manager from HR may move to a team in IT and vice versa. By the time they reach the position of senior vice president, most managers will have directed teams in most divisions. When managers quit or request a transfer, they are required to give thirty days' notice, and the set of candidates available to fill the role in time is (anecdotally) very small and sometimes empty. This shortage contributes helps explain why banks reward managers who are willing to transfer quickly from distant divisions and why job postings for every managerial level of the bank can be found on the internal and external company dashboards.

Over the span of our data, we identify 8,670 events involving 6,021 unique employees and 690 unique managers. These events are distributed uniformly over the four-year panel (e.g., see Appendix Figure A.2.i). 41% of employees experience at least one event at some point in the

overlap: 78% of the managers we identify also are listed by the employee.

four-year period, but only 13% experience two or more events. Each event will affect on average 9.75 employees, and the inter-quartile range of events affects teams of 3 and 10 employees. In Appendix A.2 we show that the sample of employees who experience a manager transition (41%) is quite representative of the whole firm in observable characteristics. Moreover, we show that the characteristics of employees and managers are similar across the different types of manager transitions.

When we define placebo events or smoker events, the manager transitions are the same, but we categorize those events differently, basing them on manager's birthday-evenness or smoking habits instead of gender. By construction, the number of placebo events equals the number of gender events. Because the smoker analysis is based on a subsample (male employees and managers for whom we could infer smoking status), the number of smoking events is smaller than the number of gender events. As for gender events, we find that the placebo events and smoker events are largely homogeneous over time and across individuals (see also Appendix A.2).

3.6 Survey Data: Relationship with Managers

We collected self-reported data on manager assignments to validate our method of identifying managers through the administrative data. To obtain data on relationship between employees and their managers, we distributed a survey to the employees in the largest division: sales and distribution. Sales and distribution comprise 62% of the firms employees, and 100% of employees outside of headquarters.¹⁵ Appendix B includes a sample of the survey instrument. The survey asks respondents to list managers who “directly influenced your key performance indicator and pay grade either in your current position or past positions”. They could select up to six managers.¹⁶ The rest of the survey asked a series of questions (described in the following sections) for each manager listed by the respondent.

We invited 4,847 employees by email to complete the survey in December 2017. Appendix B includes a sample of the emailed invitation. The head of the sales and distribution division requested full participation from employees and gave permission to conduct the survey during work hours. We emphasized that answers to these survey questions would not be revealed to co-workers or managers. A total of 3,345 employees completed the survey, implying an 70% response rate. The modal respondents reported information on their last three managers. The final dataset contained 9,068 employee-manager pairs.

3.7 Proximity to the Manager

To investigate the social interaction mechanism, we split positions by whether the employee works in physical proximity to the manager. For employees working in the headquarters offices, we use card swipe data provided by the security division. These data include information about the floor

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

¹⁶ If they had more than six managers to list, we asked employees to prioritize the most important ones since 2015.

where the employee works, which we use to calculate the share of employees of each position who work on the same floor as their managers. We split these positions by whether the position averages exceed or fall below the median. As a result, roughly half of the employees are categorized as higher-proximity and the other half as lower-proximity. In the higher-proximity positions, 80% of employees work on the same floor as their manager, compared to only 8% among the lower-proximity positions.

Security data are not available for positions outside headquarters. Thus, we included a question in the manager relationship survey to supplement these data. The question was repeated for each manager whom the employee identified in the survey. We asked “How often are (or were) you physically working near <manager name> (i.e. same floor and area)?”. Respondents could choose from the following options: “Every day or most days (4-6 times per week)”, “Some days (2-3 times per week)”, or “Infrequently”. Similar to the procedure for the swipe data, we calculate the average proximity of each position and then split positions by whether their average exceeds or falls below the median. Using this method, we categorize 62% of the position titles in the sales and distribution division for which survey data were collected. By construction, half of these employees are categorized as higher-proximity and the other half as lower-proximity.¹⁷

3.8 Frequency of Social Interactions

A third goal of the survey is to measure social interactions between employees and managers. For each manager listed by the employee, we ask, “Out of 10 work breaks (including lunch or random breaks), how many would include [*Manager’s Name*]?”¹⁸ We construct a simple variable that equals the fraction of breaks shared with the manager.¹⁹ To assess whether employees and managers discuss personal matters, we ask respondents to share their favorite sport teams and to guess the favorite sport team of their managers. For the pairs of employees and managers who responded to the survey, we measure the accuracy of the employee’s answers to this question.

3.9 Smoking Habits

We measure the smoker status of employees and their managers in two ways. We use data on smoking status from the 2017 annual health exam that occurs onsite during the workday and a corresponding online workplace health survey with the same questions and framing. To complement the previously described data comprising snapshots of employees working in September 2017, we use two additional supplemental surveys.

¹⁷ In the higher-proximity positions, 88% of employees report working with their manager every day or most days, compared to only 65% of employees in the lower-proximity positions.

¹⁸ We ask the question about a share of 10 breaks, rather than asking about the overall number of breaks, to minimize the 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.

¹⁹ The survey also asks about an alternative form of social interactions with the manager: “Of the last 10 emails you sent to [*Manager’s Name*], how many included some part that was personal?” However, there is too little variation in this outcome to be useful for the analysis: the average share of personal emails is just 5%. We suspect employees may have under-reported this type of behavior for fear of violating company policy.

Section 3.6 describes the survey of manager relationship, which includes a question about whether the employee and their current and past managers smoke. Additionally, we deployed a 2-minute survey exclusively about smoking. Appendix C includes a sample of this survey. This survey asks about the respondent’s own smoking status and the smoking status of current and past co-workers, including those who left the bank prior to the annual health exam. We emailed invitations to the survey on February 2018, and the invitation included information about cash prizes to be raffled to 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 survey data. Using this method, we assign smoking status to 57% of employees from the main sample.²⁰ Some employees appear on both the annual health exam data and the crowdsourced survey data. We can use that overlap to validate the crowdsourced data. As expected, we find that the two sources of data are highly consistent with each other: the crowdsourced measure of smoker status coincides with the health records 82% of the time.²¹

4 Results: Effects of Manager’s Gender

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

4.1 Descriptive Analysis

Before diving into the event-study analysis, we provide some simple descriptive evidence on the association between past exposure to male managers and the employee’s subsequent promotions. Let $\Delta P_{i,t}$ be employee i ’s change in pay grade from t at 10 quarters later. Let $S_{i,t-1}$ indicate the employee’s recent exposure to male managers (i.e., the fraction of the past year that employee i was assigned to a male manager). Consider the following regression:

$$\Delta P_{i,t} = \alpha_0^M \cdot (1 - F_i) + \alpha_1^M \cdot S_{i,t-1} \cdot (1 - F_i) + \alpha_0^F \cdot F_i + \alpha_1^F \cdot S_{i,t-1} \cdot F_i + \beta \cdot T_{i,t} + \rho_{P_{i,t}} + \varepsilon_{i,t} \quad (4)$$

Note that we interact $S_{i,t-1}$ with a gender indicator (F_i) to estimate the relationship separately for male and female employees. The regression includes basic control variables: the employee’s tenure ($T_{i,t}$) and, to flexibly compare employees who started at the same level, fixed effects for initial pay grade ($\rho_{P_{i,t}}$).

Figure 1 presents the results in binned scatterplot form. The x-axis indicates if the employee is assigned to a female (towards the left) or male (towards the right) manager. The y-axis indi-

²⁰ Of those, 59% (33% of the sample) are classified using their annual health exam, and the remaining 41% are classified using crowdsourced data.

²¹ More precisely, we classify an individual as a smoker if over one-third of the crowdsourced survey reports flag the individual as a smoker. This one-third threshold is arbitrary but largely inconsequential for the categorization. Moreover, in Appendix A.16 we show the results are robust to the use of alternative thresholds.

cates the change in pay grade 10 quarters later. This figure suggests that women are promoted at roughly similar rates under male and female managers ($\alpha_F = 0.056$, p-value<0.001). In contrast, male employees are promoted substantially faster under male managers than they are under female managers ($\alpha_F = 0.380$, p-value<0.001). More precisely, Figure 1 shows that when employees are assigned mostly (i.e., above 75% of the time) to female managers, they tend to be promoted at the same rate, regardless of whether they are female or male. The gender gap is small (0.022 pay grades) and statistically insignificant (p-value=0.403). On the contrary, when employees are assigned mostly (i.e., above 75% of the time) to male managers, then the male employees are promoted 0.30 pay grades higher than female employees (p-value<0.001).

The evidence from Figure 1 suggests that female and male employees receive equal treatment under female managers, but male managers promote their male employees faster than their female employees. This evidence, however, is subject to the usual concerns with causal inference. For example, it is possible that the share of male managers correlates with manager, employee, or position characteristics that are favorable to the promotion of male employees. In the following sections, we address these causality concerns with the event-study analysis of manager transitions.

4.2 Event-Study Analysis

We start by comparing the pay grade effects from transitioning from a female to male manager relative to transitioning from a female manager to another female manager. Figure 2.a presents the results based on the econometric framework described in Section 2. This event-study graph shows the evolution of pay grades in each of the 10 quarters leading up to a manager transition and the 10 quarters after the manager transition. We present coefficients for female employees (red squares) and male employees (blue circles) separately. The quarter before the event (-1) corresponds to the omitted category, and thus the corresponding coefficient is always zero by construction.

When inspecting Figure 2.a, note that these coefficients refer to differences across transition types.²² As a result, a coefficient of zero in the post-treatment period does not imply that employees remain in the same pay grade; rather, it indicates similar growth rates of pay grades across employees transitioning from female to male managers versus employees transitioning from female to female managers. This context has ample upward mobility, meaning that employee pay grades increase over time.

Figure 2.a shows that, in the 10 quarters prior to the transition, the coefficients are similar in magnitude between male employees and female employees, confirming that female and male employees share similar trends prior to the manager transition.²³ On the contrary, the evolution of

²² We focus on the single-difference estimates to isolate the effects of the change of gender from the effects of changing manager *per se*. For reference, Appendix A.3 reports the raw coefficients $\beta_{j,s}^M$ and $\beta_{j,s}^F$, that is, without “differencing” between transition types.

²³ The most important fact is that there are no systematic trends. Having said that, some of the pre-treatment gender differences are statistically significant (most notably, at quarter 4 before the transition the difference p-value is 0.022). This is probably spurious: given the large number of falsification coefficients presented in the paper, a minority of them are expected to be statistically significant just by chance. In any case, as discussed below, the pre-treatment coefficients are consistently close to zero for the transitions in the opposite direction (from male to female managers)

pay grades diverge substantially between male and female employees after the transition date. On the one hand, male employees advance further in the organization after being assigned to a male manager, relative to how they would have fared if they instead were assigned to female managers. At 10 quarters after the transition, pay grades among men exceed those among women by 0.65 points (p-value < 0.001), roughly equivalent to a salary that is 15% higher,²⁴ when transitioning from a female manager to a male manager (relative to transitioning from a female manager to a different female manager). An alternative way of illustrating the magnitude of this effect is to compare it to a baseline: 10 quarters after experiencing a manager transition employees gain an average of 0.96 pay grades (for details, see Appendix A.1).

On the other hand, Figure 2.a shows that female employees do not advance similarly after being assigned to male managers, relative to being assigned to female managers. Female employees have pay grades that are 0.043 points (p-value = 0.736) lower at 10 quarters after transitioning from a female to a male manager (relative to transitioning to a different female manager).²⁵ Moreover, this coefficient of -0.043 points for female employees is statistically significantly different from the corresponding coefficient of 0.60 for male employees (p-value < 0.001).

Now, we assess the robustness of the identification strategy by analyzing the manager transitions in the opposite direction. Figure 2.b is equivalent to Figure 2.a, except that it corresponds to the opposite type of transition (comparing a transition from a male manager to a female manager minus the transition from a male manager to a different male manager). Keep in mind that the coefficients are identified by a disjointed set of transition events, and thus there are no “mechanical” reasons why the results should mirror each other. A comparison of Figures 2.a and 2.b indicates 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 timing and order of magnitude. For example, Figure 2.a indicates that male employees gain 0.60 points (p-value < 0.001) at 10 quarters after gaining a male manager. In turn, Figure 2.b indicates that male employees lose 0.30 points (p-value = 0.031) at 10 quarters after losing a male manager.

Figure 3 presents the double-difference estimates described in Section 2. Intuitively, the coefficients from Figure 3.a correspond to the difference between the male and female coefficients from Figure 2.a. Figure 3.a shows that at 10 quarters after the transition, the male-to-male advantage amounts to 0.65 pay grades, which is not only highly statistically significant (p-value < 0.001) but also economically large. Figure 3.b is equivalent to Figure 3.a, except that it corresponds to the transitions in the opposite direction. According to Figure 3.b, there is a statistically significant (p-value = 0.001) male-to-male advantage of 0.44 pay grades at 10 quarters after the transition. This point estimate of 0.44 is smaller in magnitude than the corresponding estimate of 0.65 from Figure 3.a, but we cannot reject the null hypothesis that these two coefficients are equal (p-value = 0.343).

In Figure 3.c, we present the dual-double-difference estimates. Intuitively, Figure 3.c corre-

and in the dual-double-differences specification.

²⁴ A single pay-grade increase is associated with a log increase of 0.227 (Appendix A.1), and thus a 0.65 pay grade increase should be equivalent to a salary that is 15% ($= e^{0.65 \cdot 0.227} - 1$) higher.

²⁵ Although this evidence suggests that promotions among male employees do not crowd out promotions among female teammates, it also does not imply that male employees do not crowd out anyone. Indeed, male employees are probably crowding out other employees in the same position but on different teams, as well as external hires.

sponds to the average male-to-male advantage implied by Figures 3.a and 3.b. The estimated male-to-male advantage amounts to 0.54 pay grades at 10 quarters after the transition (p-value<0.001). Unsurprisingly, this point estimate is in the middle of the corresponding point estimates from Figures 3.a and 3.b. However, these estimates combine their variation and are thus more precisely estimated than the corresponding coefficients from Figures 3.a and 3.b on their own. As a result, we use the dual specification to maximize statistical power, such as when measuring the heterogeneity of the effects.

Given that we have the most statistical power for the dual-double-difference specification, we can use it to explore the timing of the effects. First of all, notice that there is a significant jump in the first quarter after the manager switch: the coefficient corresponding to +1 quarters after the switch is 0.10, and statistically significant (p-value=0.006). Note also that the male-to-male advantage grows smoothly over time, which happens mechanically because while some employees may happen to be up for promotion right after the manager switch, most employees are months or sometimes years away from their next promotion opportunity at the time of the switch. Just like in academia, this company reviews promotions at the end of the year, and depending on the position some employees may be considered up for promotion every other year or so.

To illustrate this better, we can compare the size of the male-to-male advantage relative to the average change in pay grade at each time horizon. In the first four quarters after a manager transition, the male-to-male advantages are estimated at 0.10, 0.10, 0.12, and 0.16 (each of them statistically significant, with p-values of 0.006, 0.032, 0.022, and 0.012). The average pay grade change in each of the first four quarters after a manager transition were 0.05, 0.15, 0.25, and 0.34, respectively. The male-to-male advantage grows stronger during the second year: in the fifth through eight quarters after the transition, the male-to-male advantages are estimated at 0.20, 0.21, 0.30, and 0.38 (each of them statistically significant, with p-values of 0.011, 0.016, 0.001, and <0.001). For comparison, the average pay grade change in the fifth through eight quarters after a manager transition were 0.47, 0.56, 0.67, and 0.75. The male-to-male advantage seems to taper off in the third year: the point estimates for the ninth and tenth quarters are 0.48 and 0.54 (both p-values < 0.001), with their difference being small and statistically insignificant. For comparison, the average pay grade change in the ninth and tenth quarters after a manager transition were 0.84 and 0.96.

In the appendix, we report some additional robustness checks. In Appendix A.4, we measure the persistence of gender transitions. Appendix A.5 shows that the results are similar under alternative specifications; including controls for employee characteristics and specifications with and without manager fixed effects. Appendix A.6 shows that the results are robust under alternative definition of events, such as excluding the largest events. Appendix A.7 shows that the results are robust if we restrict the sample to employees who joined the firm before the start of the panel. In Appendix A.8, we show that the results are robust if we focus on the employees' first transition event only.

4.3 Placebo Analysis: Birthday-Evenness

As a placebo test, 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 provides a useful sanity check. First, it helps rule out mechanical reasons why our event-study framework would generate spurious effects. Second, this placebo analysis can be used to assess whether our standard errors are conservative enough.

Figure A.9.i is equivalent to Figure 2, but it is based on birthday-evenness instead of gender. Figure 4.a compares transitions from an even-birthday manager to an odd-birthday manager versus transitions from an even-birthday manager to another even-birthday manager. We directly present double-difference coefficients for odd-birthday employees relative to even-birthday employees.

As expected, Figure 4.a shows no significant difference between the two types of transition, either before or after the event. For instance, at 10 quarters after transitioning from an even-birthday to an odd-birthday manager (relative to another even-birthday manager), the difference between the pay grades of odd-birthday and even-birthday employees is close to zero (0.06), statistically insignificant (p -value=0.518), and precisely estimated. Moreover, we can reject the null hypothesis that this coefficient for odd-birthday employees the same as the corresponding coefficient of 0.65 estimated for male employees in Figure 2.a (p -value=0.001). Moreover, Figure 4.b shows that the results are virtually the same if we use the transitions in the opposite direction (i.e., odd-to-even instead of even-to-odd). For the sake of brevity, we report the single-difference and dual-double-difference estimates in Appendix A.9.

In Appendix A.10, we also show that the results are robust to an alternative placebo specification that combines the gender of the manager with the birthday-evenness of the employees. We take the same gender transitions of the managers from the previous section and show that, despite strong heterogeneity with respect to the gender of the employee, there is no significant heterogeneity with respect to the birthday-evenness of the employee.

Ideally we could replicate the results using an alternative characteristic, such as race or ethnicity, that would provide another shared demographic trait against which to benchmark the gender results. Unfortunately, in our context, racial and ethnic diversity are too limited for such a benchmark.

4.4 Effects on Attrition, Effort and Performance

Male employees may reach higher positions under male managers because they are less likely to leave the firm, work longer hours, or perform better than their female counterparts. Alternatively, the male-to-male advantage may be the consequence of favoritism without accompanying productivity justifications. To probe these factors, we measure the effects of manager transitions on additional outcomes. Figure 5 presents the results under the dual-double-difference specification, which combines all transition types and thus maximizes statistical power. Each panel of Figure 5 is equivalent to Figure 3.c, except it uses a different dependent variable instead of pay grade. As we use different dependent variables, we follow Hastings et al. (2019) by setting the scale of each graph at

approximately twice the within-individual standard deviation.²⁶ For example, the within-individual standard deviation in pay grade is about 0.5, so in the event-study graphs for that dependent variable the y-axis ranges from -1 to 1.²⁷ This hopefully allows for a more intuitive comparison between event-study graphs that involve different outcomes.

Figure 5.a shows the effects on the probability of leaving the firm (i.e., a dummy variable that equals 1 for every month after the employee leaves the firm). When using this specific dependent variable, there is an extra challenge for the event-study analysis. By construction, employees do not experience manager transitions after they leave the company. We can still estimate the post-treatment coefficients, but we cannot estimate the pre-treatment coefficients. We address this common challenge in event-study analysis by using the standard approach of assigning hypothetical events to individuals who left the firm (Kleven et al., 2019). To do this, we take advantage of the fact that after an employee leaves the firm, the employee's former team still exists. Thus, we take the transition events experienced by the team and assign them to the employee, even if the employee no longer works for the firm.

Figure 5.a shows that, consistent with the assumption of balanced pre-trends, the coefficients preceding the transition date are close to zero, precisely estimated, and statistically insignificant. The evidence also indicates a lack of male-to-male advantage on attrition: the post-event coefficients are also close to zero, precisely estimated, and statistically insignificant. For example, at 10 quarters after the event, the male-to-male coefficient for attrition is close to zero (-0.3 percentage points), statistically insignificant (p-value = 0.667), and precisely estimated. On average, the probability of leaving the firm at 10 quarters after an event is 35 percentage points. Thus, the estimated effect of less than one percentage point is quite small relative to that baseline.

Next, we assess whether there is a male-to-male advantage in employee effort or performance. For example, male managers may be better role models than female managers for male employees (Kofoed and McGovney, 2019), or perhaps male managers are better than female managers at communicating with or monitoring male employees. Figure 5.b shows the event-study graph with the (logarithm of) the monthly number of days worked as the dependent variable. The coefficients are close to zero, statistically insignificant, and precisely estimated. For example, the male-to-male advantage at 10 quarters after the transition is close to zero (0.012 log points), statistically insignificant (p-value=0.313), and precisely estimated. We can interpret the magnitude as a percentage increase of <1% in the days worked. This difference is tiny compared with the magnitude of the male-to-male advantage in pay grades reported in Figure 3.c, which is roughly equivalent to a 13% salary difference.²⁸

Figure 5.c presents the results for the other measure of effort: (the logarithm of) the average number of hours spent in the office, according to security log data for employees working at headquarters (43% of the sample).²⁹ Again, we find no male-to-male advantage on time spent in the

²⁶ Hastings et al. (2019) perform a similar normalization but use the inter-quartile range instead.

²⁷ To allow for familiar scales, we use round numbers. For example, the within-individual standard deviation of pay grade is 0.479, so instead of using a range from -0.958 to 0.958, we use a range from -1 to 1.

²⁸ A single pay-grade increase is associated with a log increase of 0.227 (Appendix A.1), and thus a 0.54 pay-grade increase should be equivalent to a salary that is 13% ($= e^{0.54 \cdot 0.227} - 1$) higher.

²⁹ In Appendix A.13, we report the effects on paygrade for this same subsample. The effects are less precisely estimated

office. The point estimates are close to zero, statistically insignificant, and precisely estimated. For example, at 10 quarters after the transition, the male-to-male advantage is small (relative to the within-individual standard deviation) and statistically insignificant (p-value = 0.822).

Figure 5.d presents the effects on sales revenue for the subsample of employees who have a sales role (42% of the sample).³⁰ The point estimates are again close to zero, statistically insignificant, and precisely estimated.³¹ For instance, at 10 quarters after the transition, the male-to-male advantage is small (relative to the within-individual standard deviation) and statistically insignificant (p-value = 0.790).

In sum, the analysis presented in this section indicates that the higher promotion rates that male employees enjoy under male managers are not accompanied by any differences in attrition, effort, or performance, compared to female employees.

A related mechanism could be that male managers are better able to discern the best candidates to promote by virtue of their social interactions with their male employees (Brogaard et al., 2014).³² If male managers were systematically better at allocating tasks or roles to male employees through an information advantage, in most circumstances this would translate into higher sales revenue, the firm's bottom line. We do not find this to be true even 10 quarters after the manager switch. In Appendix A.11.1, we focus on a sample of employees, junior sales associates, for whom their post-promotion productivity is revealed ex-post through their sales performance. We show that, across men and women, the employees who are promoted by male managers are on average significantly less productive than those promoted by female managers in their subsequent sales roles, while overall rates of promotion do not differ significantly.

In the appendix, we present some additional robustness checks. For instance, the results presented here are based on the dual-double-difference specification. In Appendix A.12, we show that the results are robust when looking at two directions of the transitions (i.e., gaining and losing a male manager) separately.

4.5 Heterogeneity by Gender Norms

The effects of manager gender could be mediated by social norms about gender roles (Jayachandran, 2020). For example, in more chauvinistic contexts, male managers may be more prone to spending time, becoming friends with and promoting their male employees. While we do not have data on the gender norms at the individual level, we take advantage of geographic variation.

The employees of the firm work in hundreds of geographically dispersed branches and two corporate towers, one in the "northern region" and one in the "southern region". There are sharp

but still follow the basic patterns from the whole sample.

³⁰ In Appendix A.13, we report the effects on paygrade for this same subsample. The effects are less precisely estimated but still follow the basic patterns from the whole sample.

³¹ As shown in Appendix A.11, the results are robust to using the inverse hyperbolic sine transformation. As this outcome equals zero a non-trivial fraction of the time, we cannot use the logarithm of sales revenues as a dependent variable. We use the inverse hyperbolic sine transformation instead, which can be interpreted like a log transformed variable, as $\text{arcsinh}(x) \rightarrow \ln(2x) = \ln(2) + \ln(x)$ rapidly.

³² According to the asymmetric information story, the marginal employee who is promoted by a male manager should ex-post be more productive than the marginal employee promoted by a female manager.

differences in the cultural and institutional past between the northern and southern regions that could generate persistent differences in gender norms. Anecdotally, southerners are more westernized due to early European colonization while northerners are more communist-influenced due to Chinese rule. Based on these roots and prior research on persistent gender differences between East and West Germany by Boelmann, Raute, and Schönberg (2020), we expect the southern regions to have stronger and more unequal gender norms around roles at work.

Indeed, these anecdotal accounts are supported by different sources of data. First, the gender gap in the labor force participation is three times as large in the southern regions (5 percentage points) as in the northern regions (16 percentage points).³³ The anecdotal accounts are also supported by survey measures of gender norms. The most recent wave of the World Values Survey covered the country where the firm is located and included a relevant question on the role of gender in business leadership. Respondents were asked whether they agree with the statement “Men make better business executives than women do” in a scale from “strongly disagree” to “strongly agree.” Among respondents in the top income quartile (which is the most relevant population for the employees in our firm), 52% of men agreed with the statement that they men are better business executives than women. This is roughly 12 percentage points higher (or 31% higher) than the share of men in the northern regions who agreed with this statement (difference p-value = 0.016).³⁴

For the purpose of the heterogeneity analysis, we split units between the northern and southern regions using data on the birthplace of employees. A unit is categorized as *northern* if the modal employee in that unit was born in one of the northern provinces under greater communist-influence. Under this categorization, 68% of employees work in northern units and the remaining 32% work in the southern units. Employees are broadly similar between the northern and southern units in terms of their observable characteristics. A slightly greater share of the workers in the northern units are male (32% relative to 30%) or have a college degree (86% relative to 78%), and the average worker in a northern unit works in a slightly larger unit (100 workers relative to 77) – for more details, see Appendix A.17. Using the firm’s pay grade data, we find that the gender pay gaps is 39% higher in the southern units than in the northern units (1.1 pay grade gap in the south versus 0.8 pay grade gap in the north, p-value of the difference=0.016). The sign of this difference is consistent with the anecdotal accounts of the difference in gender norms between the north and the south.

Figure 6 presents the heterogeneity of results between northern and southern units. To maximize statistical power, we estimate the same dual-double-difference model from Figure 3 but split the set of event dummies in two: one set for the northern units and another set for the southern units. Figure 6.a presents coefficients for the southern units and Figure 6.b presents coefficients for the northern units. Figure 6.a shows a significant male-to-male advantage in the southern units (where there are stronger gender norms), while Figure 6.b shows that the male-to-male advantage is smaller and less statistically significant in the northern units. For example, at 10 quarters after the transition the male-to-male advantage in pay grade is 0.69 (p-value<0.001) in the south (Figure 6.a), compared with 0.44 (p-value=0.001) in the north (Figure 6.b). This specific difference

³³ These figures were calculated from the most recently available census data (from 2009).

³⁴ The difference is similar, but less precisely estimated, among female respondents. For more details, see Appendix A.17

must be taken with a grain of salt because, although large, it is not precisely estimated and thus is statistically insignificant (p-value= 0.269). The difference between Figure 6.a and Figure 6.b are consistent in direction and magnitude for all the time horizons.³⁵ In sum, the evidence suggest that while the effects are present in both the northern and southern units, they are more pronounced in the southern units, where there are stronger gender norms.

4.6 Interpreting the Magnitude of the Effects

Next, we discuss the economic magnitude of the male-to-male advantage. Under the assumption that our findings are due to a positive effect of male managers on male employees, we compute what would happen to the overall gender gap if we were to remove this male-to-male advantage.³⁶ The unconditional gender gap in pay grade in our setting is approximately 0.90 pay grades.³⁷ As 66% of male employees have male managers, the gender pay gap would be reduced by 0.36 pay grades ($= 0.54 \cdot 0.66$) if the male-to-male advantage were removed. That is, removing the male-to-male advantage would reduce the gender pay gap by 40% (from 0.90 to 0.54 pay grades).^{38,39}

We can also compare our findings to the results from related studies. However, we must take these comparisons with a grain of salt due to obvious differences in context and research design (e.g., we rely on quasi-experimental methods). The closest related study is Kunze and Miller (2017), which is based on data on white-collar employees from a private firm in Norway. Consistent with our findings, they find that the gender gap in promotions is higher in establishments where the share of male superiors is higher.⁴⁰ We also provide a quantitative comparison to Kunze and Miller (2017). They report a gender gap in promotion rates of 3.3 percentage points (page 772). That gap is 2 percentage points larger in establishments with 100% male superiors, relative to establishments with 0% male superiors (column (1) of Table 2). We can use the estimates from Kunze and Miller (2017) in an equivalent counterfactual analysis as the above. Since 83% of managers in their sample are male, this gap can explain $50\% = \frac{0.83 \cdot 2}{3.3}$ of the gender gap in promotions. This is in the same

³⁵ For example, looking at the effects 4 quarters after the transition, the male-to-male advantage is estimated at 0.29 pay grades in the south (p-value=0.010) vs. 0.09 pay grades in the north (p-value=0.205).

³⁶ As discussed in Section 2, our specification cannot distinguish whether the male-to-male advantage is driven by favorable treatment from male managers, unfavorable treatment by female managers, or a combination of both. The descriptive analysis presented in Section 4.1, however, suggests that the favorable treatment by male managers is a more likely explanation.

³⁷ This figure is estimated using a cross section of the bank in the last period of our sample (December, 2018).

³⁸ This exercise combines all types of manager transition. The resulting magnitude would be bit larger if we used the transitions starting with a male manager only (it would explain 48% of the gender gap, instead of 40%) or the transitions starting with a female manager only (it would explain 32%).

³⁹ One caveat with this interpretation is that if some effects were due to a negative effect of female managers on male employees, then the effects on the gender pay gap would be smaller and the 40% reported here could be considered an upper bound. In the extreme case where all effects are due to negative effects of female managers on male employees, then removing these manager effects should actually increase the gender pay gap, as male employees' pay grades would increase and female employees' would remain unaffected. As discussed throughout, we did not find any supporting empirical evidence of a negative bias of female managers against male employees.

⁴⁰ While their preferred interpretation is that the difference is due to female managers helping female employees, they also describe the indeterminacy between women helping women or men helping men given the absence of a gender-neutral benchmark. Our evidence instead suggests that male managers help male employees.

order of magnitude as our baseline estimate of 40% presented above.

Our finding that women do not benefit from having female managers echoes results from earlier studies in non-corporate contexts: female referees and female committee members do not increase the odds of acceptance of female-authored papers or promotion of female candidates (Bagues et al., 2017; Card et al., 2019); and female teachers in public schools show similar job satisfaction and turnover rates whether working in schools run by female principals or male principals (Grissom et al., 2012; Husain et al., 2018).

Last, to contextualize this effect size further, we turn to a result that is well established in the literature: the so-called “motherhood penalty”. From the administrative HR data, we are able to identify workers who take maternity leave at any point in our sample. Women are entitled to six months of maternity with partial pay, and in our sample, the average leave is 22 weeks (4.5 months). Looking at the same December 2018 cross section, we find that the gap between men and women who never take maternity leave is 0.83 pay grades. When this gap is measured instead using women who ever take maternity leave, the gap is 31% larger or 1.09 pay grades. The difference is highly statistically significant with a p-value of $p < 0.001$.⁴¹ While we do not have an instrument to causally estimate the impact of childbirth and maternity leave on the pay gap, the unconditional difference in means is similar to the carefully estimated gender gap in Kleven et al. (2019), which grows between 20% to 30% over four years after birth. Overall, this constitutes suggestive evidence that the male-to-male advantage in this firm may be in the same order of magnitude as the motherhood penalty (Schönberg and Ludsteck, 2014; Kleven et al., 2019).

5 Results: Social Interactions Channel

We use social interactions as an all-encompassing term to refer to a family of mechanisms featuring face-to-face, personal interactions between employees and their managers. For example, male managers may become emotionally attached to male employees over time and thus feel increasing pressure to promote them. Perhaps male employees use the interactions to gain the manager’s sympathy and schmooze their way into promotions. Socializing with the manager may make the accomplishments and efforts of employees more salient to the manager, thus making those employees more likely to be rewarded with a promotion. With more frequent interactions, male managers might better identify potential among their male employees (Brogaard et al., 2014). Male employees also may use the time spent with their manager to claim credit and engage in self-promotion (Sarsons et al., 2019; Isaksson, 2019; Coffman et al., 2019). Male employees may get favorable treatment from managers by getting assigned tasks that are more conducive to promotions (Lehmann, 2013; Babcock et al., 2017). It is also possible that male managers are more willing to work alongside with and train their male subordinates, compared with their female subordinates (Ranganathan, 2019).

⁴¹ We measure maternity leave using the HR records of the firm. This is a rough estimate. For example, it is possible that some employees had children before they joined the company. It is also possible that employees became parents without formally logging parental leave with HR.

In each of the following sections, we provide suggestive evidence that the male-to-male pay grade advantage operates at least partially through the social interactions channel, and that these social interactions are not accompanied by higher employee productivity or the managers' enhanced ability to discern and promote more productive employees.

5.1 Heterogeneity by Proximity to the Manager

The first test of the social interactions channel exploits heterogeneity according to proximity to the manager. If socializing with the manager plays an important role, then we should observe stronger effects for employees whose jobs require frequent face-to-face interactions with the manager (Bandiera et al., 2009).

Recall from Section 3.7 that we use a combination of administrative and survey data to split positions into higher and lower proximity to the manager. An example of a high-proximity position is customer support specialist, who normally sit in a specific location near the manager. An example of a low-proximity position is the sales and quality development director, who usually travels between branches and reports back to the manager by phone or email. We were able to successfully classify the proximity for a large majority (88.2%) of the employees in the sample.⁴²

Figure 7 presents the heterogeneity results. To maximize statistical power, we estimate the same dual-double-difference model from Figure 3. However, rather than having a single set of event dummies, we split this set in two: one set for high-proximity positions and another for low-proximity positions. Figure 7.a presents coefficients from high-proximity events, and Figure 7.b presents coefficients from low-proximity events. Figure 7.a shows a significant male-to-male advantage when the employee works in high proximity to the manager. Figure 7.b further shows that the male-to-male advantage is close to zero and statistically insignificant when the employee works in a low proximity environment. For example, Figure 7.a indicates that at 10 quarters after the transition, the male-to-male advantage in pay grade is 0.76 (p-value<0.001) in the high-proximity group, compared with 0.21 (p-value=0.178) in the low-proximity group, and their difference is statistically significant (p-value= 0.013).

In the appendix, we present some additional robustness checks. In the baseline results presented above, we estimate the dual-double-difference estimator that combines all types of transitions and thus maximizes the statistical power. In Appendix A.14, we show that the results are robust when looking at transitions in each direction (i.e., a “gain” or a “loss” of a male manager). Similarly, to maximize power, our measure of proximity combines administrative data and survey data. In Appendix A.14, we show that the results are robust even when looking at the administrative and survey measures of proximity separately.

In Appendix A.14.1 we compare the observable characteristics across those low and high proximity groups and show that, although not large, there are some systematic differences. One potential concern is that the differences in results between high and low proximity groups stem partially from differences along those other characteristics. We provide evidence against this concern. We repli-

⁴² In Appendix A.13, we report the effects on paygrade for this same subsample. The effects are almost identical as for the whole sample.

cate the analysis using propensity score matching to maintain balance along critical observables, including position pay grade, position share male, and position share sales roles and average sales revenue.⁴³ We show that the results are practically identical after re-weighting.

5.2 Effects on the Time Spent with the Manager

If the social interactions channel plays a role, we should observe that male employees interact more with their managers after transitioning to a male manager (relative to transitioning to another female manager). To test this hypothesis, we use our survey measure of social interactions: the fraction of the last ten breaks that the employee took that was shared with his or her manager.

Although the share of breaks taken with the manager is probably not the perfect measure of social interactions, we start by providing some suggestive evidence that it contains some meaningful variation. First, we show that employees who spend more breaks with their managers get to know them better. Among the 3,072 employee-manager pairs for whom both the manager and employee responded to our survey (so that we can determine if the employee guessed the manager's preference correctly), we find that spending more breaks with the manager is positively associated with an accurate guess about their favorite sports team. The association is highly statistically significant (p -value < 0.001) and large in magnitude: increasing the share of breaks taken with the manager from 0% to 100% is associated with a 44% increase in the probability of correctly guessing the manager's favorite team (from 25 to 36 percentage points). Moreover, we show that our measure of shared breaks is correlated to promotion rates. Among the 5,047 employee-manager pairs for which the employee answered our survey, we find that spending breaks with the manager is positively associated with promotions. This correlation is not only statistically significant (p -value = 0.014), but also economically significant: increasing shared breaks from 0% to 100% is associated with an additional increase of 0.1 pay grade.⁴⁴

Next, we assess whether male employees change the shared breaks with their managers after transitioning to a male manager. Ideally, we would implement the same quarterly event-study analysis that we employ for the outcomes measured with administrative data. Unfortunately, due to the smaller sample size, that is not feasible for this survey outcome. For instance, although the analysis of pay grades is based on 374,913 observations (employee-month pairs), the dataset on share of breaks has only 9,068 observations (employee-manager pairs).⁴⁵ Instead, we use a stylized version of the event-study framework tailored to the smaller survey dataset.

We follow the same notation from Section 2.2, with a few differences. The first difference is that, instead of the employee-level pair, observations are denoted by employee-manager pair, where i denotes the employee and m the manager, respectively. Let $Share_{i,m}$ be the share of breaks that

⁴³ We selected these observables based on a study by Bandiera et al. (2009) showing that managers exhibit favoritism under low powered incentives in particular, so we are careful to re-weight so that the incentive schemes and masculinity across the two groups are balanced.

⁴⁴ For a binned scatterplot of these two relationships, see Appendix Figure A.15.ii.

⁴⁵ The smaller sample size is due to two reasons. First, we collected survey data on a minority of employees. Second, even among surveyed employees, we measure their social interactions only at a handful of points in time (as opposed to the monthly data for four years from the administrative records).

employee i took with manager m . Consider the following regression:

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

$D_{i,m}^j$ is a dummy variable that equals 1 if individual i experiences an event of type j from manager $m - 1$ to manager m . As in Section 2.2, we interact these dummies with gender indicators to allow the effects to be gender-specific. The coefficients $\beta_{j,post}^F$ and $\beta_{j,post}^M$ are intended to capture the change in social interactions after the employee transitions to the new manager. In turn, $D_{i,m+1}^j$ is a dummy variable that equals 1 if individual i experiences an event of type j from manager m to manager $m + 1$. The coefficients next to these variables ($\beta_{j,pre}^F$ and $\beta_{j,pre}^M$) are intended to provide the usual tests for pre-trends: they measure whether future manager transitions affect the employee's social interactions with the current manager. Additionally, the regression includes gender-specific time effects (δ_m^F and δ_m^M) and a set of basic controls ($X_{i,m}$): unit size, manager's pay grade, and position title dummies.

Figure 8 presents the results from the stylized event-study analysis. Figure 8.a presents the results for the gender manager transitions. The findings suggest that social interactions may play a role in the male-to-male promotion advantage. The coefficients for the male employees are consistent in sign with the effects on pay grades reported in the previous section. The male coefficient labeled "after transition" corresponds to the effects following a transition. For male employees, the share of breaks taken with the manager increases by 15 percentage points (p-value=0.017) after transitioning from a female manager to a male manager, relative to transitioning from a female manager to another female manager. This coefficient is statistically and economically significant: it is almost as large as the within-employee standard deviation of the dependent variable (17.4 percentage points). The corresponding falsification test is reported in Figure 8.a as the coefficient labeled "before transition". As expected, the falsification coefficients is close to zero (0.2 percentage points) and statistically insignificant (p-value=0.987).

For female employees, in contrast, there is no robust evidence that the share of breaks with the manager changed as a result of a change in gender of the manager.⁴⁶ This evidence also aligns with the lack of female-to-female advantage in promotions. One possible interpretation is that female managers socialize equally with female and male employees, thus offering no advantages to one gender. Another possible interpretation is that even if male and female employees spend equal time with their manager, gender differences may still occur in their ability to convert those interactions into a higher promotion probability. That is, female employees may be less successful than male employees at taking advantage of opportunities to schmooze with managers.

We also validate this research design by estimating the stylized event-study with our placebo

⁴⁶ Figure 8.a shows that even though for female employees the "after transition" coefficient (-8 percentage points) is statistically significant (p-value = 0.037), it is probably spurious because it is almost identical to the corresponding falsification coefficient (-11 percentage points, p-value=0.080).

events. Figure 8.b presents the results. As expected, both even-birthday and odd-birthday employees are equally likely to share breaks with their manager after transitioning from an even-birthday manager to an odd-birthday manager (relative to transitioning from an even-birthday manager to another even-birthday manager).

5.3 Co-Smoking Shocks to Social Interactions

For a final test of the social interactions channel, ideally we would flip a coin to determine the frequency of social interactions among male employees and male managers. According to the schmoozing channel, the male employees assigned to socialize more with their male managers should be promoted faster. Although this ideal experiment is not feasible, we exploit quasi-experimental variation based on the transitions between non-smoker and smoker managers.

We conjecture that for an employee who smokes, having a manager who also smokes can increase the frequency of their social interactions due to shared smoking breaks. We start by using the survey measure of shared breaks to test this conjecture. Figure 8.c is equivalent to Figure 8.a except focusing on smoking status rather than gender. The results from Figure 8.c confirm the conjecture that sharing a smoking habit constitutes a significant shock to social interactions between an employee and manager. The “after transition” coefficient indicates that male employees who smoke increase the share of breaks taken with their managers by 25 percentage points (p -value = 0.002) after transitioning from a non-smoking male manager to a smoking male manager (relative to transitioning from a non-smoking male manager to another non-smoking male manager). In contrast, the corresponding coefficient for non-smoking employees is close to zero (-3 percentage points) and statistically insignificant (p -value=0.702). Moreover, the falsification coefficients, labeled “before transition”, are close to zero and statistically insignificant, both for the smoker and non-smoker employees.

These results confirm that a shared smoking habit increases socialization between an employee and manager. According to the social interactions channel, this increased socialization should result in higher promotion rates for those employees. To test this hypothesis, we estimate the event-study effects of smoker-manager transitions on pay grades. Figure 9 presents the results, which are identical to Figure 2, except focusing on manager smoking status rather than manager gender. Note that the event-study coefficients from Figure 9 (on smoker transitions) are substantially less precisely estimated than the corresponding coefficients from Figure 2 (on gender transitions), due to differences in sample sizes. The smoker analysis is limited to male employees and managers, who constitute less than half the sample, and is further limited to employees and managers for whom data on smoking status is available. Thus, the analysis of smoker transitions are based on a sample size (94,750 observations) that is roughly a quarter of the sample size used for gender transitions (380,964 observations).

Figure 9.a compares the pay grades of male employees who transition from a male manager who does not smoke to a male manager who smokes (relative to transitioning from a male manager who does not smoke to another male manager who does not smoke). Prior to the event date, the coefficients for the smoking employees (denoted by the violet triangles) are statistically indistin-

guishable from the coefficients for the non-smoking employees (denoted by the orange diamonds). This evidence indicates that the assumption about parallel trends holds. In contrast, after the transition date, the evolution of pay grades starts to gradually diverge between smoking and non-smoking employees. At 10 quarters after transitioning to a smoker manager (relative to transitioning to another non-smoker manager), the pay grades of smoker employees increase by an additional 0.70 points (p-value=0.002). In contrast, the corresponding point estimate is close to zero (0.07) and statistically insignificant (p-value=0.722) for the non-smoking employees.

We also examine the reverse smoker-status transitions: among everyone who starts with a smoking manager, we compare those who transition to a non-smoking manager versus those who transition to another smoking manager. Unfortunately, these types of transitions are much less common, resulting in estimates that are highly imprecisely estimated. Figure 9.b presents the results. The point estimates have the expected sign, indicating that smoker employees are less likely to be promoted as a result of losing their smoking manager. However, the point estimates are somewhat smaller in magnitude, less precisely estimated, and thus statistically insignificant.

For a more direct measurement of the smoker-to-smoker advantage, Figure 10 presents the double-difference estimates. Figure 10.a corresponds to the difference of coefficients between smoking and non-smoking employees from Figure 9.a. At its peak in the 8th quarter after the transition, the smoker-to-smoker advantage is estimated at 0.84 pay grades (p-value < 0.001) and remains 0.63 pay grades (p-value= 0.035) 10 quarters after the transition. Figure 10.b corresponds to the manager transitions in the opposite direction; the smoker-to-smoker advantage of 0.37 pay grades is statistically insignificant at its peak 7 quarters out, with a p-value of 0.117. Figure 10.c presents the dual-double-difference estimates, which combine the transitions in both directions. Due to noisy estimates in the late quarters of the manager transitions with an outgoing smoking manager, the estimated smoker-to-smoker advantage at 10 quarters is somewhat understated (0.30 pay grades, p-value = 0.145). However, in quarters 7, 8, and 9, the smoker-to-smoker advantage is large (0.49, 0.49, and 0.44) and statistically significant (p-values 0.003, 0.013, and 0.029).

We also present productivity results for smoking events in Appendix A.12. We find no evidence of differential performance (sales revenue) or effort (hours worked in the office, retention or absenteeism) after the manager switch. If the smoking managers were better at allocating tasks or roles to smokers through an information advantage gained during smoking breaks, in most circumstances this would translate into higher sales revenue within 10 quarters after the manager switch. However, we find no evidence of allocation improvements; mean sales revenue across smokers and non-smokers remains constant across manager switch events and the variance remains constant as well.

The surprising similarly-sized boost in social interactions arising from a match in smoking status and a match between males allows for a simple comparison of subsequent patterns of employee promotions. Both the timing and magnitude of the smoker-to-smoker advantage are similar to those of the male-to-male advantage reported in the previous section. The smoker-to-smoker advantage in the share of breaks together (a 25 percentage point increase, from Figure 8.c is comparable in magnitude to the corresponding male-to-male advantage (15 percentage points, from Figure 8.a) – we cannot reject the null hypothesis that these two effects are equal (p-value=0.360). At 10 quarters

after the transition, the smoker-to-smoker advantage (0.63 pay grades, from Figure 10.a) is statistically indistinguishable from the corresponding male-to-male advantage (0.65 pay grades, from Figure 3.a), with a p-value of the difference 0.956.

In the appendix, we present some additional robustness checks. Given the differences by gender in rates of smoking (33% of men smoke, and less than 5% of women smoke), a natural question is whether the male-to-male advantage arises purely because of co-smoking between men. In Appendix A.15, we show that only a small fraction of the male-to-male advantage can be attributed directly to the smoker-to-smoker advantage. Smoking-breaks are likely just one of several social outlets men share. Appendix A.16 further shows that the results are robust using a different criteria to code the smoker status.

6 Conclusions

We test the old boys' club hypothesis using data from a real-world corporation. At this firm, 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 of these manager transitions to show that male employees are promoted at a faster rate when assigned to male manager than when assigned to a female manager. Women, in turn, are promoted at the same rate whether they are assigned to a male or female manager. The magnitude of this male-to-male advantage in promotions explains over one-third of the gender gap in pay grades, but it cannot be explained by gender differences in attrition, effort, or performance.

We provide evidence that social interactions play a role in the male-to-male advantage in promotions. The effects of male managers on male employees is only evident in positions where managers and employees work in close physical proximity, which is a necessary condition for the social interactions channel. We show that the male-to-male advantage in promotions coincides with an increase in the frequency of social interactions with the manager. Furthermore, we provide evidence that social interactions are important even among male employees: when male employees who smoke transitioned to male managers who smoke, they took breaks with their managers more often and were subsequently promoted at higher rates.

Thus far we have evidence that some contextual features may be important for the male-to-male advantage. In the case where the employee and manager are physically distant, such as in traveling sales roles or when the manager is assigned to sit on a different floor than the employee, the male-to-male advantage disappears. We also find that two regions of the country with distinct cultural norms and attitudes toward gender equality predicts differences in the male-to-male advantage at work. In the region with long-standing communist rule and more equal gender roles, the male-to-male advantage is weaker than in regions colonized early by Europe, where people tend to believe men are better equipped for executive roles.

Our identification strategy can be applied to other contexts. The rotation of managers is common in large corporations, and the data necessary for the analysis, such as pay grades, demographics, and manager assignments, are likely collected and preserved by 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 the male-to-male advantage is most pervasive and why.

References

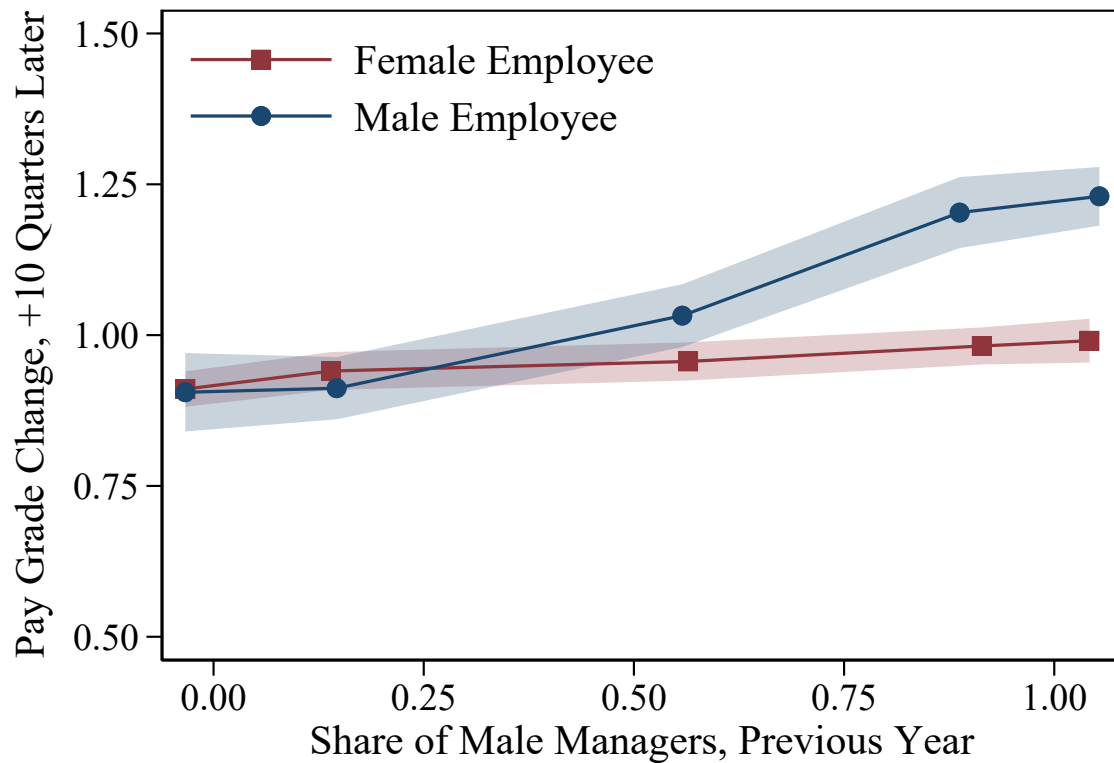
- Agarwal, S., W. Qian, D. M. Reeb, and T. F. Sing (2016). Playing the boys game: Golf buddies and board diversity. *American Economic Review* 106(5), 272–276.
- Alesina, A., P. Giuliano, and N. Nunn (2013). On the Origins of Gender Roles: Women and the Plough. *Quarterly Journal of Economics* 128(2), 469–530.
- Azmat, G. and R. Ferrer (2017). Gender Gaps in Performance: Evidence from Young Lawyers. *Journal of Political Economy* 125(5), 1306–1355.
- Babcock, L., M. P. Recalde, and L. Vesterlund (2017). Gender Differences in the Allocation of Low-Promotability Tasks: The Role of Backlash. *American Economic Review, Papers and Proceedings* 107(5), 131–135.
- Bagues, M., M. Sylos-Labini, and N. Zinovyeva (2017). Does the Gender Composition of Scientific Committees Matter? *American Economic Review* 107(4), 1207–1238.
- Bandiera, O., I. Barankaray, and I. Rasul (2010). Social Incentives in the Workplace. *Review of Economic Studies* 77(2), 417–458.
- Bandiera, O., I. Barankay, and I. Rasul (2009). Social Connections and Incentives in the Workplace: Evidence From Personnel Data. *Econometrica* 77(4), 1047–1094.
- Bell, L. A. (2005). Women-Led Firms and the Gender Gap in Top Executive Jobs. *IZA Discussion Paper No. 1689*.
- Bertrand, M., S. E. Black, S. Jensen, and A. Lleras-Muney (2019). Breaking the Glass Ceiling? The Effect of Board Quotas on Female Labour Market Outcomes in Norway. *Review of Economic Studies* 86(1), 191–239.
- Bertrand, M., 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.
- Boelmann, B., A. Raute, and U. Schönberg (2020). Wind of Change? Cultural Determinants of Maternal Labor Supply. *Mimeo*.
- Brogaard, J., J. Engelberg, and C. A. Parsons (2014). Networks and productivity: Causal evidence from editor rotations. *Journal of Financial Economics* 111(1), 251–270.
- Bronson, M. A. and P. S. Thoursie (2019). The Wage Growth and Within-Firm Mobility of Men and Women: New Evidence and Theory. *Working Paper*.
- Bursztyn, L., T. Fujiwara, and A. Pallais (2017). Acting Wife: Marriage Market Incentives. *American Economic Review* 107(11), 3288–3319.
- Bursztyn, L., A. Gonzalez, and D. Yanagizawa-Drott (2018). Misperceived Social Norms: Female Labor Force Participation in Saudi Arabia. *NBER Working Paper No. 24736*.
- Cai, J. and A. Szeidl (2018). Interfirm relationships and business performance. *Quarterly Journal of Economics* 133(3), 1229–1282.

- Card, D., S. Dellavigna, P. Funk, and N. Iriberry (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*.
- Cortés, P. and J. Pan (2019). When time binds: Substitutes for household production, returns to working long hours, and the skilled gender wage gap. *Journal of Labor Economics* 37(2), 351–398.
- Cruz, C. and C. Tolentino (2019). Gender, Social Recognition, and Political Influence. *Working Paper*.
- 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. *NBER Working Paper No. 24351*.
- 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.
- Field, E., S. Jayachandran, R. Pande, and N. Rigol (2016). Friendship at Work: Can Peer Effects Catalyze Female Entrepreneurship? *American Economic Journal: Economic Policy* 8(2), 125–153.
- Flabbi, L., M. Macis, A. Moro, and F. Schivardi (2019). Do Female Executives Make a Difference? The Impact of Female Leadership on Gender Gaps and Firm Performance. *The Economic Journal* 129(622), 2390–2423.
- Giuliano, L., D. I. Levine, and J. Leonard (2011). Racial Bias in the Manager-Employee Relationship. *Journal of Human Resources* 46, 26–52.
- Glover, D., A. Pallais, and W. Pariente (2017). Discrimination as a Self-Fulfilling Prophecy: Evidence from French Grocery Store. *Quarterly Journal of Economics* 132(3), 1219–1260.
- Goldin, C. (2014). A Grand Gender Convergence: Its Last Chapter. *American Economic Review* 104(4), 1091–1119.
- Goldin, C., S. P. Kerr, C. Olivetti, and E. Barth (2017). The expanding gender earnings gap: Evidence from the LEHD-2000 census. *American Economic Review* 107(5), 110–114.
- Gray, J. and A. Barbara (2013). *Work With Me: Eight Blind Spots Between Men and Women in Business*. Palgrave Macmillan.
- Grissom, J. A., J. Nicholson-Crotty, and L. Keiser (2012). Does My Boss’s Gender Matter? Explaining Job Satisfaction and Employee Turnover in the Public Sector. *Journal of Public Administration Research and Theory* 22(4), 649–673.
- Hastings, J. M., R. Kessler, and J. Shapiro (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.
- Hoffman, M. and S. Tadelis (2020). People Management Skills, Employee Attrition, and Manager Rewards: An Empirical Analysis.

- Hsieh, C.-T., E. Hurst, C. I. Jones, and P. J. Klenow (2019). The Allocation of Talent and U.S. Economic Growth. *Econometrica* 87(5), 1439–1474.
- Husain, A., D. A. Matsa, and A. R. Miller (2018). Do Male Workers Prefer Male Leaders? An Analysis of Principals' Effects on Teacher Retention. *NBER Working Paper No. 25263*.
- Isaksson, S. (2019). It Takes Two: Gender Differences in Group Work. *Working Paper*.
- Jayachandran, S. (2020). Social Norms as a Barrier to Women's Employment in Developing Countries. *Working Papers*.
- Karpowitz, C., J. Preece, and O. Stoddard (2020). Strength in Numbers : A Field Experiment in Gender , Influence , and Group Dynamics. *Mimeo*, 1–80.
- 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.
- Lee, A. S. (2014). How to break up the old boys' club in your office. *Quartz*, April 8 2014.
- Lehmann, J.-Y. (2013). Job Assignment and Promotion Under Statistical Discrimination: Evidence from the Early Careers of Lawyers. *Working Paper*.
- Lerner, J. and U. Malmendier (2013). With a little help from my (random) friends: Success and failure in post-business school entrepreneurship. *Review of Financial Studies* 26(10), 2411–2452.
- Lleras-Muney, A., M. Miller, S. Sheng, and V. Sovero (2019). Party On: The Labor Market Returns to Social Networks and Socializing. pp. 1–69.
- 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. Gender Differences in Networking. *Economic Journal*.
- 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., K. Gërkhani, E. Reuben, and A. Schram (2019). Gender Differences in Recognition for Group Work. *Journal of Political Economy*, forthcoming.
- Schönberg, U. and J. Ludsteck (2014). Expansions in maternity leave coverage and mothers' labor market outcomes after childbirth. *Journal of Labor Economics* 32(3), 469–505.
- Shue, K. (2013). Executive networks and firm policies: Evidence from the random assignment of MBA peers. *Review of Financial Studies* 26(6), 1401–1442.
- Stevenson, B. and J. Wolfers (2006). Bargaining in the Shadow of the Law: Divorce Laws and Family

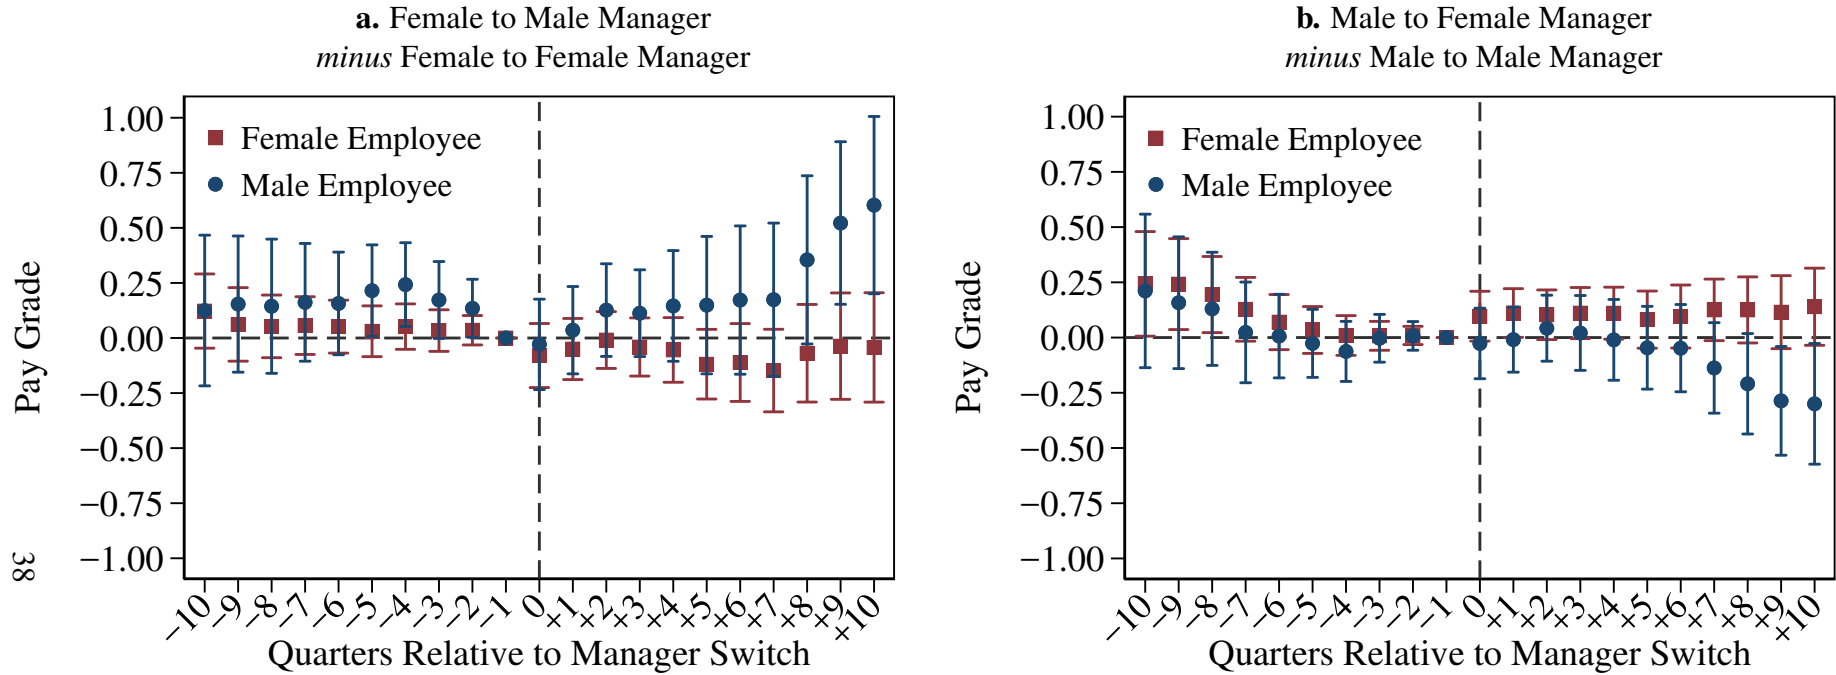
- Distress. *The Quarterly Journal of Economics* 121(1), 267–288.
- Voth, J. and G. Xu (2019). Patronage for Productivity: Selection and Performance in the Age of Sail. *Working Paper*.
- Wasserman, M. (2018). Hours Constraints, Occupational Choice, and Gender: Evidence from Medical Residents. *Working Paper*.
- Xu, G. (2018). The Costs of Patronage: Evidence from the British Empire. *American Economic Review* 108(11), 3170–3198.
- Yildirmaz, A., C. Ryan, and J. Nezaj (2019). 2019 State of the Workforce Report: Pay, Promotions and Retention. Technical report, ADP Research Institute.

Figure 1: Link between Past Exposure to Male Managers and Future Pay Grade Changes




Notes: See Section 4.1 for details about the regression specification. This binned scatterplot shows the relationship between the share of male managers in the previous year and the change in pay grade at 10 quarters later. Results based on employees who are in the panel for at least 14 quarters (so that we can compute the left-hand-side and right-hand-side variables without truncation). The red squares correspond to the female employees while the blue circles correspond to the male employees. The analysis uses the following control variables: the employee's seniority, an indicator variable for the employee's gender and initial pay grade fixed effects. The 95% confidence intervals are represented by the shaded areas.

Figure 2: Effects of Manager’s Gender on Pay Grade: Single-Differences Estimates



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


Figure 3: Effects of Manager’s Gender on Pay Grade: Double-Differences Estimates



39

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

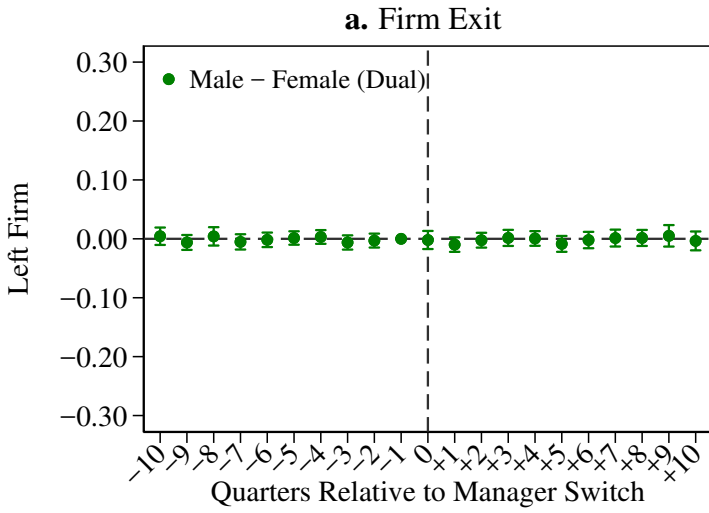
Figure 4: Placebo: Double-Difference Estimates



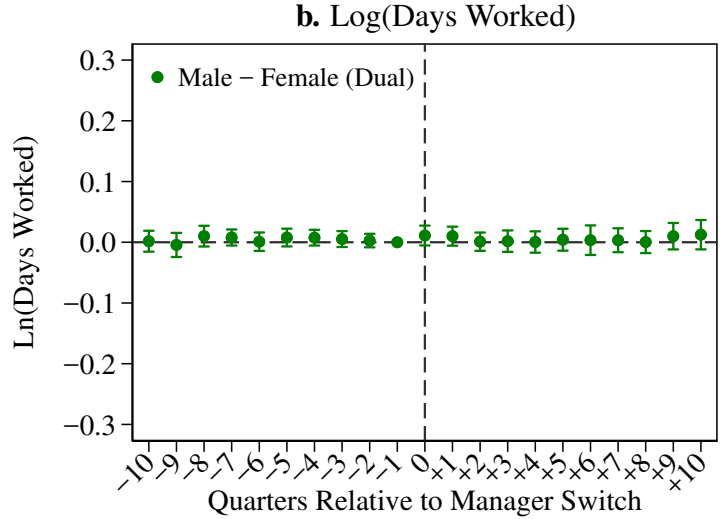
40

Notes: All coefficients were estimated from a single regression including 380,964 observations of 14,638 employees (7,533 Even BD & 7,105 Odd BD). Panel (a): 4,536 employees (2,385 Even BD & 2,151 Odd BD) experience events: 3,014 transitions from a even-birthday manager to a odd-birthday manager and 3,131 from a even-birthday manager to another even-birthday manager. Panel (b): 4,244 employees (2,155 Even BD & 2,089 Odd BD) experience events: 2,922 transitions from a odd-birthday manager to a even-birthday manager and 2,453 from a odd-birthday manager to another odd-birthday manager. Panel (c): Panel (c) corresponds to the average between the coefficients from panel (a) and the (negative value of) the coefficients from panel (b). This “symmetric” double-differences estimates is then $\frac{1}{2}\{(\beta_{E2O,t}^E - \beta_{E2E,t}^E) - (\beta_{E2O,t}^O - \beta_{E2O,t}^O) - [(\beta_{O2E,t}^E - \beta_{O2O,t}^E) - (\beta_{O2E,t}^O - \beta_{O2O,t}^O)]\}$. The 95% confidence intervals are presented in brackets, with two-way clustering by manager and employee. The within-employee standard deviation of the dependent variable is 0.475.

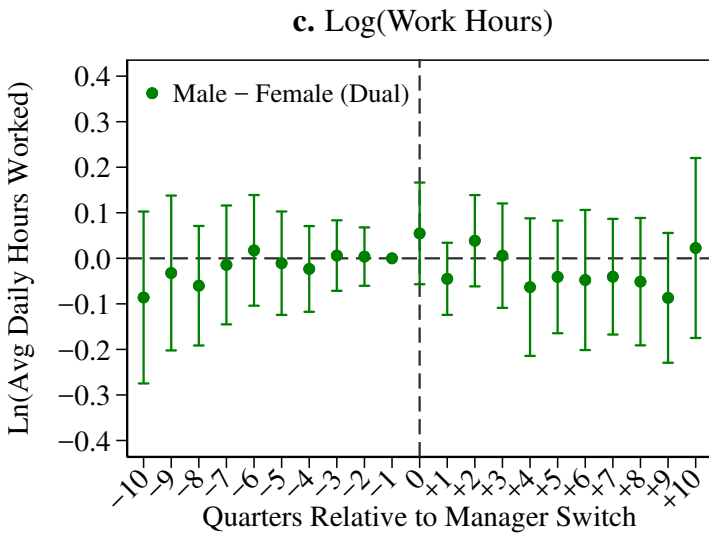
Figure 5: Dual-Double-Differences Estimates: Additional Outcomes



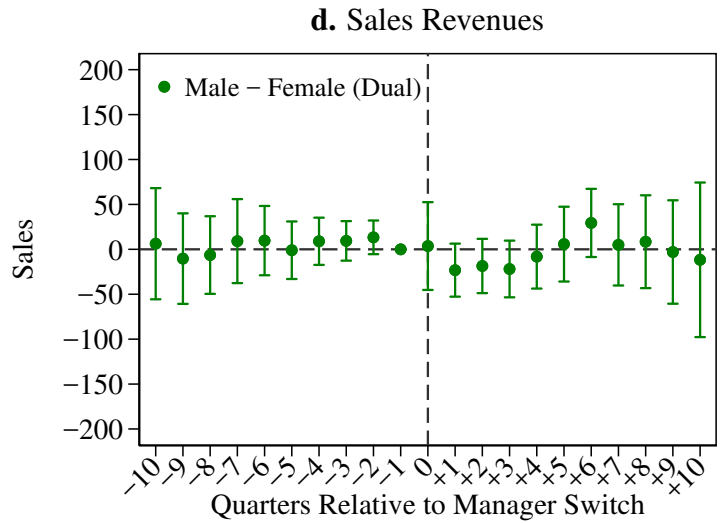
All coefficients were estimated from a single regression including 359,225 observations of 14,601 employees (5,157 Male & 9,444 Female). 6,579 employees (2,046 Male & 4,533 Female) experience events: 1,865 transitions from a female manager to a male manager (F2M): 2,106 F2F, 1,770 M2F, 4,243 M2M. The within-employee standard deviation of the dependent variable is 0.177.



All coefficients were estimated from a single regression including 352,282 observations of 14,154 employees (4,913 Male & 9,241 Female). 5,647 employees (1,697 Male & 3,950 Female) experience events: 1,261 transitions from a female manager to a male manager (F2M): 1,766 F2F, 1,490 M2F, 3,388 M2M. The within-employee standard deviation of the dependent variable is 0.138.



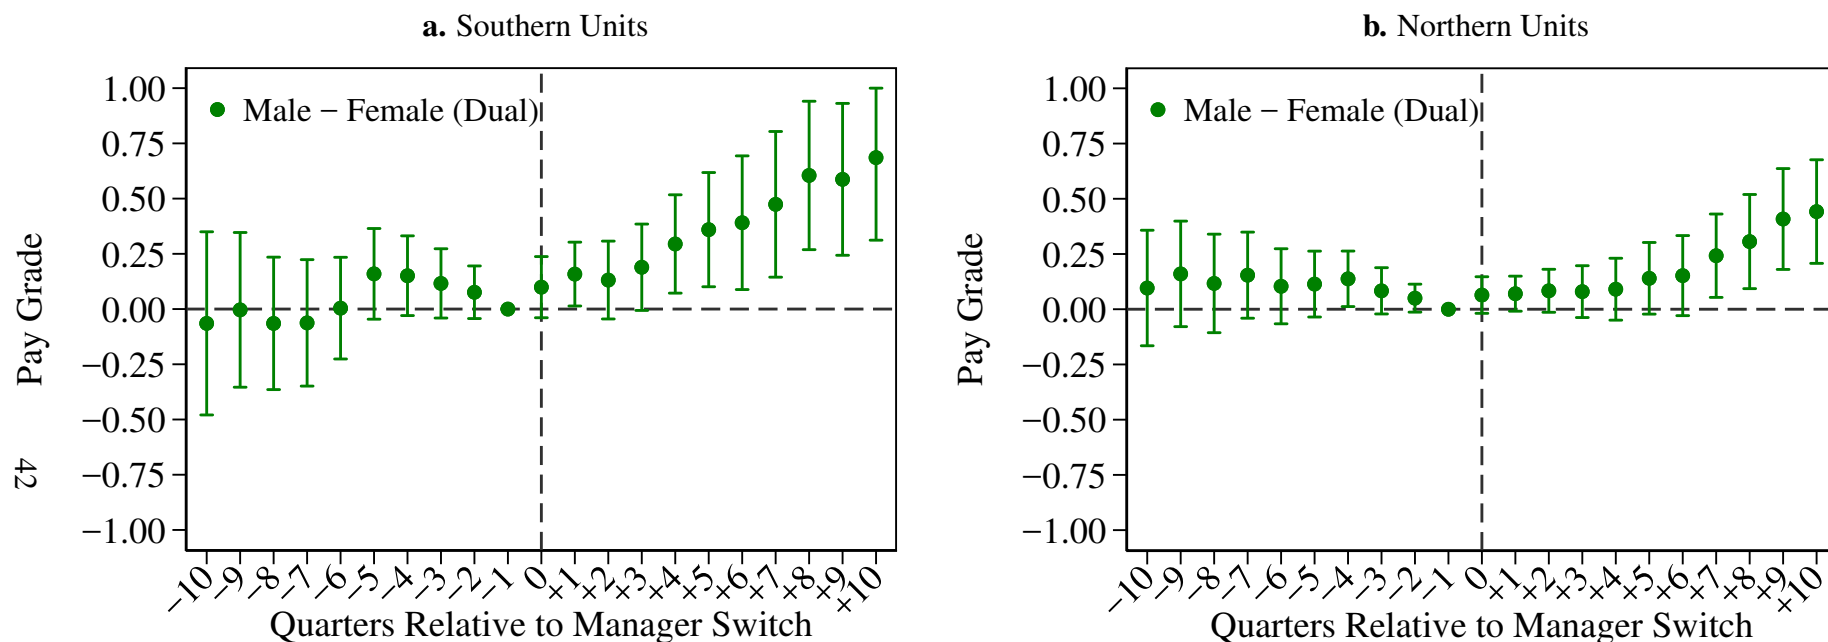
All coefficients were estimated from a single regression including 104,215 observations of 4,875 employees (1,881 Male & 2,994 Female). 1,677 employees (581 Male & 1,096 Female) experience events: 370 transitions from a female manager to a male manager (F2M): 690 F2F, 548 M2F, 588 M2M. The within-employee standard deviation of the dependent variable is 0.208.



All coefficients were estimated from a single regression including 136,341 observations of 6,244 employees (1,814 Male & 4,430 Female). 2,444 employees (611 Male & 1,833 Female) experience events: 581 transitions from a female manager to a male manager (F2M): 572 F2F, 542 M2F, 1,701 M2M. The within-employee standard deviation of the dependent variable is 95.1.

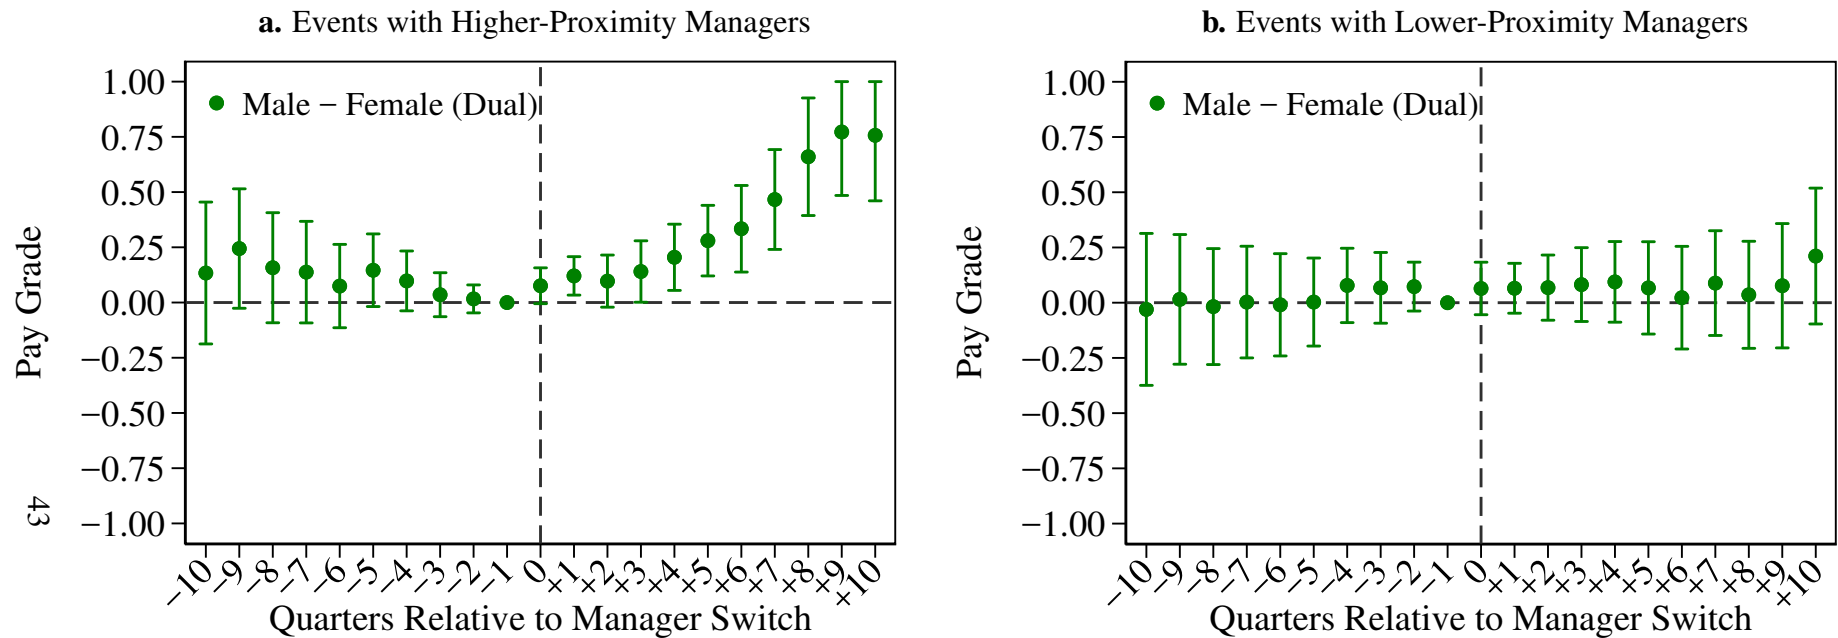
Notes: See Section 2 for details about the regression specification. These results are based on the symmetric specification reported in panel (c) of Figure 3, which combines data on the four types of gender transitions. The only difference is that in this figure, instead of pay grade, we use different dependent variables: in panel (a) the dependent variable is an indicator that takes the value 1 in every month after the employee left the firm (these results include additional events after the employees left the firm); in panel (b) the dependent variable is the logarithm of the total number of days worked in the month (inferred from data on approved leaves of absence); in panel (c) the dependent variable is the logarithm of the average number of hours worked in a given month (inferred from data on swipes in and out of the building, and available for headquarter employees only); in panel (d) the dependent variable is the sales revenue (available for employees with sales roles only) normalized to have mean 100. The 95% confidence intervals are presented in brackets, with two-way clustering by manager and employee.

Figure 6: Effects of Manager Gender on Pay Grade: Heterogeneity by North/South Birthplace (Dual-Double-Differences Estimates)




Notes: See Section 2 for details about the regression specification. These results use the symmetric specification reported in panel (c) of Figure 3, based on the four types of gender transitions. We split the events in two subsets: South and North, based on the birthplace of the modal worker in the unit. All coefficients are estimated from the same regression with 380,959 observations of 14,638 employees (5,193 Male & 9,445 Female). 1,890 employees (566 Male & 1,324 Female) in predominantly-southern units (panel (a)) experience events. There are 496 transitions from a female manager to a male manager (F2M): 482 F2F, 399 M2F, and 1,284 M2M. 4,225 employees (1,297 Male & 2,928 Female) in predominantly-northern units (panel (b)) experience events: 893 transitions from a female manager to a male manager (F2M): 1,396 F2F, 1,130 M2F, and 2,424 M2M. The 95% confidence intervals are presented in brackets, with two-way clustering by manager and employee. Confidence intervals in panel (a) are trimmed at +1.

Figure 7: Effects of Manager Gender on Pay Grade: Heterogeneity by Proximity to the Manager (Dual-Double-Differences Estimates)



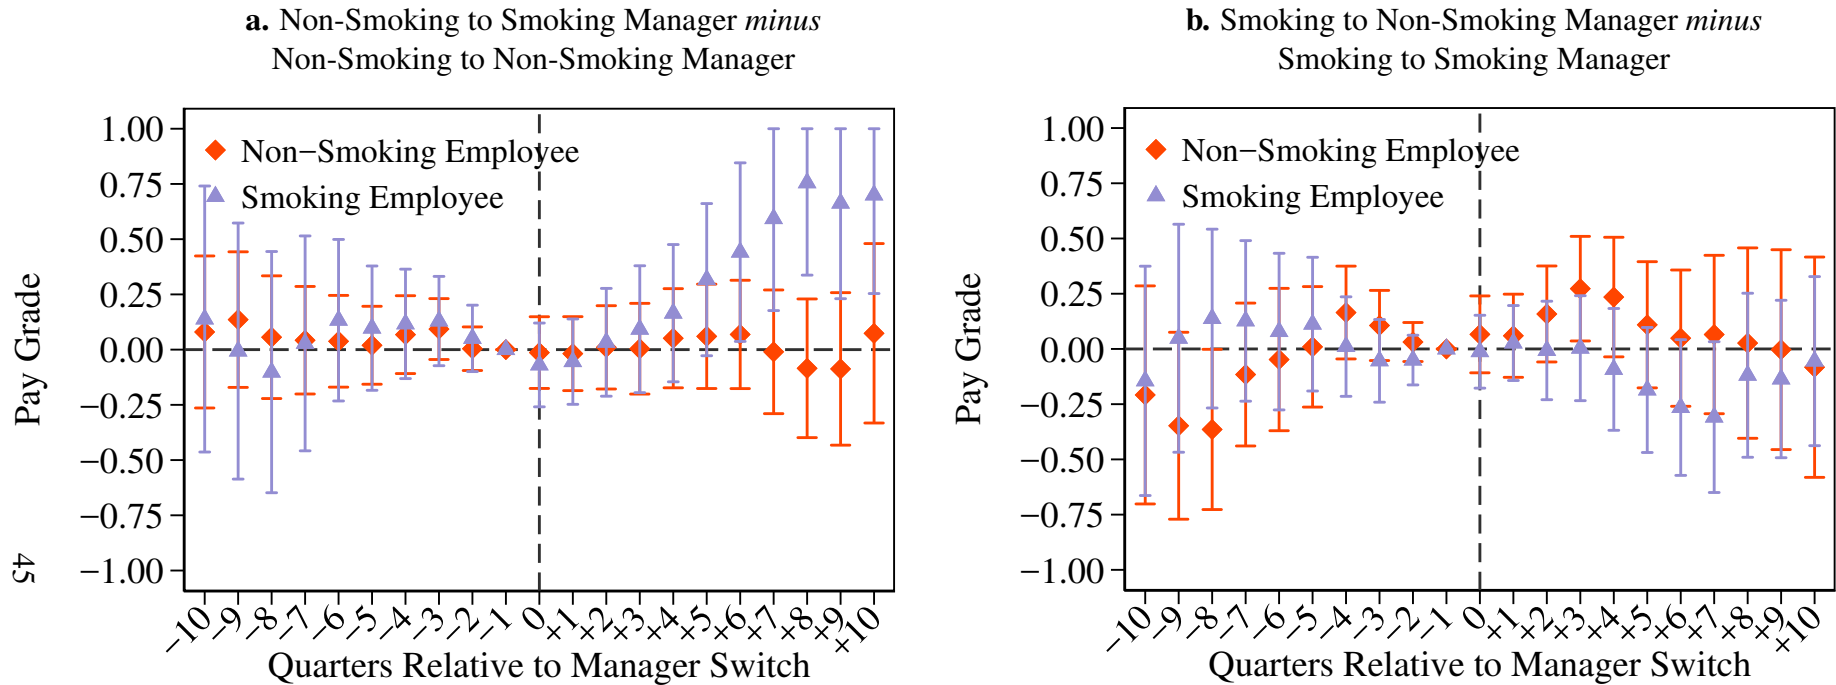
Notes: See Section 2 for details about the regression specification. These results use the symmetric specification reported in panel (c) of Figure 3, based on the four types of gender transitions. The only difference is that we split the events in two subsets: high and low proximity events, based on whether the position of the employee in the month of the event was of higher or lower proximity to the manager. All coefficients are estimated from the same regression with 360,239 observations of 13,814 employees (4,912 Male & 8,902 Female). The higher-proximity events (panel (a)) affect 2,983 employees (1,043 Male & 1,940 Female), with 743 transitions from a female manager to a male manager and 617 transitions from a female manager to a male manager (F2M): 1,075 F2F, 754 M2F, 1,508 M2M. The lower-proximity events (panel (b)) affect 3,056 employees (783 Male & 2,273 Female), with 743 transitions from a female manager to a male manager and 762 transitions from a female manager to a male manager (F2M): 751 F2F, 742 M2F, 2,182 M2M. The 95% confidence intervals are presented in brackets, with two-way clustering by manager and employee.

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




Notes: Regression results with the share of breaks. See Section 5.2 for full econometric specification. Panel (a): This regression includes 4,843 observations of 2,638 workers (698 Male & 1,940 Female). 411 employees (82 Male & 329 Female) of these workers experience a transition event. There are 235 transitions from a female manager to a male manager and 241 from a female manager to another female manager. Panel (b): This regression includes 4,947 observations of 2,648 employees (1,352 Even BD & 1,296 Odd BD). 813 employees (429 Even BD & 384 Odd BD) experience events. 418 transitions from a even-birthday manager to a odd-birthday manager and 494 from a even-birthday manager to another even-birthday manager. Panel (c): This regression includes 1,287 observations of 699 workers (176 smoker & 523 Non-smoker). 193 employees (50 Smoking & 143 Non-Smoking) of these workers experience a transition event. There are 49 transitions from a non-smoking manager to a smoking manager and 157 from a non-smoking manager to another non-smoking manager. The within-individual standard deviation of this outcome is 0.174. The 95% confidence intervals are presented in brackets, with two-way clustering by manager and employee.

Figure 9: Effects of Manager’s Smoking Habits on Pay Grade: Single-Differences Estimates



Notes: See Section 2 for details about the regression specification. All coefficients are estimated from the same regression that includes 94,728 observations of 2,907 employees (966 Smoking & 1,941 Non-Smoking). The dependent variable is the pay grade of the employee. The estimates shown in the graph are based on the coefficients of the event-study variables. The orange diamonds correspond to the coefficient for non-smoking employees, while the lavender triangles correspond to the coefficients for smoking employees. Panel (a) corresponds to the difference between transitions from a non-smoker manager to a smoker manager versus transitions from an non-smoker manager to another non-smoker manager. 912 employees (275 Smoking & 637 Non-Smoking) experience events. There are 287 transitions from a non-smoking manager to a smoking manager and 939 from a non-smoking manager to another non-smoking manager. Panel (b) corresponds to the difference between transitions from a smoker manager to a non-smoker manager versus transitions from a smoker manager to another smoker manager. 464 employees (198 Smoking & 266 Non-Smoking) experience events, comprised of 296 transitions from smoker manager to non-smoker manager and 276 from a smoker manager to another smoker manager. The 95% confidence intervals are presented in brackets, with two-way clustering by manager and employee.

Figure 10: Effects of Manager’s Smoking Habits on Pay Grade: Double-Differences Estimates



46

Notes: See Section 2 for details about the regression specification. All coefficients are estimated from the same regression that includes 94,750 observations of 2,907 employees (966 Smoking & 1,941 Non-Smoking). The dependent variable is the pay grade of the employee. The estimates shown in the graph are based on the coefficients of the event-study variables. The green triangles correspond to the difference between the coefficient for smoking employees and non-smoking employees. Panel (a) corresponds to the difference between transitions from a non-smoker manager to a smoker manager and transitions from a non-smoker manager to another non-smoker manager (as in panel (a) of Figure 9). The estimates shown in Panel (a) are the double-differences estimates $(\beta_{N2S}^S - \beta_{N2N}^S) - (\beta_{N2S}^N - \beta_{N2N}^N)$. Panel (b) corresponds to the difference between transitions from a smoker manager to a non-smoker manager and transitions from a smoker manager to another smoker manager (as in panel (b) of Figure 9). The estimates shown in Panel (b) are the double-differences estimates $(\beta_{S2N}^S - \beta_{S2S}^S) - (\beta_{S2N}^N - \beta_{S2S}^N)$. Panel (c) corresponds to the average between the coefficients from panel (a) and the (negative value of) the coefficients from panel (b). The dual-double-differences estimates shown in (c) are then $\frac{1}{2}\{(\beta_{N2S}^S - \beta_{N2N}^S) - (\beta_{N2S}^N - \beta_{N2N}^N) - [(\beta_{S2N}^S - \beta_{S2S}^S) - (\beta_{S2N}^N - \beta_{S2S}^N)]\}$. The 95% confidence intervals are presented in brackets, with two-way clustering by manager and employee.