

NBER WORKING PAPER SERIES

WOMEN HELPING WOMEN? EVIDENCE FROM PRIVATE SECTOR DATA ON
WORKPLACE HIERARCHIES

Astrid Kunze
Amalia R. Miller

Working Paper 20761
<http://www.nber.org/papers/w20761>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2014

We are grateful to Francine Blau, Albrecht Glitz, David Matsa, Øivind Anti Nilsen, Carmit Segal, Geoffrey Warner and various seminar and conference participants for helpful comments and discussions. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2014 by Astrid Kunze and Amalia R. Miller. 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.

Women Helping Women? Evidence from Private Sector Data on Workplace Hierarchies
Astrid Kunze and Amalia R. Miller
NBER Working Paper No. 20761
December 2014
JEL No. J13,J31,J62,J7,K31,M14,M51

ABSTRACT

This paper studies gender spillovers in career advancement using 11 years of employer-employee matched data on the population of white-collar workers at over 4,000 private-sector establishments in Norway. Our data include unusually detailed job information for each worker, which enables us to define seven hierarchical ranks that are consistent across establishments and over time in order to measure promotions (defined as year-to-year rank increases) even for individuals who change employers. We first find that women have significantly lower promotion rates than men across all ranks of the corporate hierarchy, even after controlling for a range of individual characteristics (age, education, tenure, experience) and including fixed effects for current rank, year, industry, and even work establishment. In measuring the effects of female coworkers, we find positive gender spillovers across ranks (flowing from higher-ranking to lower-ranking women) but negative spillovers within ranks. The finding that greater female representation at higher ranks narrows the gender gap in promotion rates at lower ranks suggests that policies that increase female representation in corporate leadership can have spillover benefits to women in lower ranks.

Astrid Kunze
Norwegian School of Economics
Department of Economics
Helleveien 30
N-5045 Bergen
Norway
Astrid.Kunze@nhh.no

Amalia R. Miller
Department of Economics
University of Virginia
P. O. Box 400182
Charlottesville, VA 22904
and NBER
armiller@virginia.edu

A data appendix is available at:
<http://www.nber.org/data-appendix/w20761>

1. Introduction

Although women comprise about 45 percent of the labor force across OECD countries, they continue to earn less than men on average and remain under-represented among business leaders.¹ The OECD average gender wage gap (defined as the difference between median male and female wages, divided by the median male wage) among full-time workers was 15 percent. While 7.7 percent of employed men in OECD countries had managerial responsibilities, only 4.4 percent of employed women did. The highest levels of business leadership are even more male dominated. In June 2014, women held only 4.8 percent of CEO positions in Fortune 500 companies and 5.0 percent of those positions in Fortune 1000 companies.² The gender gap in leadership is even present in Nordic countries. Only 3 percent of CEOs at the 145 largest companies in the Nordic countries are female (Zander, 2014). Across all public limited companies in Norway, women are 6.4 percent of general managers (Statistics Norway, 2014).

In this paper, we use data from Norway to examine a potential mechanism that generates these gender gaps in pay and leadership: gender differences in career progression. We use comprehensive data on a sample of over half a million white-collar worker-year observations across over 4,000 work establishments for the time period 1987-1997 and measure career advancement using promotions to higher organizational ranks. We first quantify gender differences in promotion rates, controlling for a range of factors that might explain women's slower progress up the ranks of corporate hierarchies, including gender differences in education, experience, and tenure, as well as selection into workplaces and childcare obligations. We next examine the role of path dependence from previous male dominance of workplaces. We assess

¹ OECD Statistics are from the Online OECD Employment Database, accessed at <<http://stats.oecd.org/>> on June 12, 2014.

² See Catalyst at <<http://www.catalyst.org/knowledge/women-ceos-fortune-1000>> for the full list of female CEOs in these groups; accessed June 20, 2014.

its importance by measuring the spillover effects of increasing female representation on gender differences in promotion rates.

This analysis is possible because we exploit an unusual dataset that has several advantages over the datasets that have previously been used to study gender differences in promotions. These enable us to advance the literature and reconcile some of the previous conflicting findings.³ First, our data contain detailed job information that enables us to assign workers to one of seven hierarchical levels that are defined consistently across firms and over time. Hence, our sample includes a wide range of employers, unlike many studies of gender gaps in promotions that focus on a single firm (Yap and Konrad, 2009; Ransom and Oaxaca, 2005; Giuliano et al., 2005; Petersen and Saporta, 2004; Jones and Makepeace, 1996; Hersch and Viscusi, 1996; Gerhart and Milkovich, 1989).

Second, for each employer in our sample, we observe the population of workers. We are also able to track workers over time and across establishments in the sample. These features are not typically available in samples based on household or firm surveys that include multiple firms,⁴ and they allow us to control for unobservable differences in promotion rates between workplaces and initial ranks using a wide range of fixed effects. Data on coworkers is what enables our analysis of gender spillovers.

³ Although most studies find gender gaps that favor men, there are several examples of studies that find higher promotion rates for women in the public sector, such as Barnett et al. (2000) on public sector workers in California, Hersch and Viscusi (1996) on workers at a public utility, and Powell and Butterfield (1994) on promotions to senior executive positions in the federal government. Gerhart and Milkovich (1989) finds higher promotion rates for women between 1980 and 1986 in a sample of lower-level workers (who remained at the company for both years) at a single private company. That company was operating under an affirmative action plan at the time, though the plan did not directly apply to the sample. Booth et al. (2003) find similar promotion rates between men and women after controlling for observable factors, but find that women receive lower wage increases from promotions.

⁴ Olson and Becker's (1983) sample consists of two waves of a survey of workers; Blau and DeVaro (2007)'s sample has information on the most recent hires at a cross-section of firms; and Winter-Ebmer and Zweimuller (1997) study white-collar workers from a 1% cross-sectional sample of the Austrian population, defining past promotions implicitly using current rank and education.

We also have detailed information on worker characteristics beyond gender, including age, experience, occupation, tenure at the firm, tenure in the current rank, occupation, part-time status, and children that we can use to examine the role of sex differences in observable characteristics in explaining any gender differences in promotions. Although these controls are valuable in accounting for differences in preferences or abilities that affect human capital investments, it is worth noting they are not able to isolate supply-side explanations entirely. This is because differences in experience, occupation and tenure are all equilibrium outcomes that can be affected by labor demand and opportunities. Also, although finding a persistent gap with these controls can rule out some of the major supply-side explanations for gender differences, it is still possible that unmeasured supply factors, reflecting productivity or preferences, contribute to the remaining “unexplained” gender gap in promotion rates.

The fourth advantage of our data is that it includes individuals who change workplaces; we can track those changes and define promotion outcomes for those workers as well. This allows us to separately study “internal promotions” that depend upon continued employment at the same establishment and “any promotions” that also include promotions involving a change of employer. This distinction may be important, for example, if men are more likely to change jobs if they are not promoted quickly and women are more likely to wait for internal promotions (possibly because they are less geographically mobile).

The main weakness of our data source relative to sources used in some previous studies (such as Blau and DeVaro, 2007, and several of the single-firm analyses) is that we do not observe performance reviews or evaluations for workers. Although it is possible that these (often subjective) evaluations themselves reflect (possibly unconscious) gender bias on the part of

evaluators,⁵ having this information would enable us to distinguish between promotion gaps driven by differences in evaluation scores and those caused by women needing to meet higher performance benchmarks to advance. This latter situation can arise when the positive signal from a high evaluation is perceived as less precise when applied to women: risk averse supervisors may then prefer to promote men over equally qualified women (Aigner and Cain, 1977).⁶ Without this information, the gender differences in promotions we measure potentially include both of these effects.

Our first finding is that women in our sample experience a significantly lower annual likelihood of advancing a rank than do their male counterparts. We observe lower promotion rates for women in the raw data and in regression models with increasing sets of controls for individual characteristics and workplace fixed effects. We also observe a significant promotion gap from each of the six lower ranks in our data. The promotion gaps are present both when promotion is defined exclusively for advances within the same work establishment and for more general advances that also include moves to new work establishments.

In the second part of our analysis, we study gender spillovers in promotion rates by measuring how the female share of the workforce at the same establishment affects the gender gap in promotion. This allows us to focus on workplace, rather than individual, factors that may be related to the demand for female leaders. Rather than focusing exclusively on how the

⁵ Ibarra et al. (2010) discuss gender bias in evaluations in the corporate setting and Goldin and Rouse (2000) find evidence of bias against female candidates for orchestral positions. Bagues and Esteve-Volart (2010) also find evidence of gender bias, but in their setting, evaluators are harsher on candidates of their same sex. Evaluations for promotions into management positions may be further influenced by stereotypical associations between masculinity and leadership (Koenig et al., 2011).

⁶ These concerns, for example, are reflected in a recent McKinsey & Company report that lists “institutional feeling that promoting a woman would be ‘too risky’” as an institutional “mindset” that is a barrier to women’s advancement (Barsh and Yee, 2011). Bjerk (2008) also considers the possibility that women have fewer chances to signal their skill level before entering the labor market or early in their careers.

highest-ranking women (managers, top executives or board members) affect lower ranking workers, our study considers the entire organizational hierarchy of white collar workers and considers changes in the female share of two groups of coworkers: “peers” at the same rank and “bosses” at a higher rank.⁷

There are several reasons why gender spillovers from bosses might be positive, and in fact, positive gender spillovers are a common assumption underlying public policy interventions to promote women in business leadership. However, it is important to note that the spillovers from bosses could also be negative if women who progress block other women from following in their footsteps. Gender spillovers among peers can also be positive or negative in theory.⁸

Our empirical analysis finds significant gender spillover effects from both bosses and peers, but in opposite directions. We find that higher female shares among bosses in the next highest rank are related to increased promotion rates for women relative to men, but that higher female shares among same-ranked coworkers are related to lower rates. These results are consistent across a range of models with different sets of control variables and are robust to alternative definitions of female leadership shares.

Previous studies have measured gender spillovers in hierarchical organizations, but this paper is the first to measure effects of both types of coworkers on promotion rates using private sector data. The only other study we are aware of that also measures effects of demographic interactions from both peers and superiors on promotion rates is Karaca-Mandic et al.’s (2013) study of racial and gender spillovers among the lowest-ranking enlisted members of the US Army. Although that setting is quite different from the white-collar workforce in Norway, it is

⁷ Because these variables are highly correlated with one another, estimated effects of female leadership using observational data without controls for the female shares among peers are likely to suffer from omitted variable bias.

⁸ Section 4 discusses theories that would generate positive and negative spillovers from peers or bosses.

interesting that the pattern of gender spillovers observed in that study is the same: positive effects of female superiors but negative effects of female peers.⁹

Papers that have focused exclusively on “downward” flowing gender spillovers on promotion rates at lower ranks tend not to find significant effects. Blau and DeVaro (2007) use cross-sectional survey data on recent hires in U.S. firms. They find female workers have lower promotion rates and the sex of the worker’s immediate supervisor does not interact significantly with the worker’s own sex in affecting promotion probability. Giuliano et al. (2005) also find no differential effects of supervisor sex on promotion rates for female employees (though they do find a slight reduction in quit rates); they argue this may be because their sample is “relatively youthful and predominantly female” (p. 3). Using data on Danish companies, Smith et al. (2013) find lower rates of promotion from VP to CEO positions for women (including both internal and external hires); female leadership at the hiring company does not mediate this effect.

In contrast with these results focused on promotions for individual workers, studies that estimate effects of female leaders on the representation of women at lower ranks in their organizations do tend to find positive spillovers. These include Cohen et al.’s (1998) analysis of 333 savings and loan banks in California and Kurtulus and Tomaskovic-Devey’s (2012) study of the more comprehensive data on the demographic composition of workforces reported by private sector employers to the U.S. Equal Employment Opportunity Commission (EEOC). Matsa and Miller (2011) and Bell (2005) find evidence of women helping women at the highest levels of major US corporations: increases in female board representation are followed in later years by greater female representation at the CEO and top executive level and smaller gender pay gaps among top executives. However, because the spillover effects found in these papers are not based

⁹ Ciliberto et al.’s (2013) study of peer effects in fertility decisions also finds evidence of heterogeneous effects, with negative spillovers dominating within some types and positive spillovers across types.

on longitudinal data on individual workers, the estimates are not able to separately identify increases in promotion rates from changes in hiring pattern. Indeed, Cohen et al. (1998) argue that a possible source of the spillovers they find is that female leaders are better able to identify and recruit talented female workers to work at their companies.

By studying the effects of female representation across the organizational hierarchy on promotion rates, this paper also contributes to a broader literature on the effects of female business leaders on lower-ranking workers in their organizations. Recent papers have considered compensation levels and gender pay gaps (Flabbi et al., 2014; Tate and Yang, 2013; Cardoso and Winter-Ebmer, 2010) as well as employment and downsizing rates (Matsa and Miller 2013, 2014).

The remainder of this paper is organized as follows. Section 2 describes the data and the empirical relevance of gender differences in promotions. Section 3 presents our approach to estimating promotion regressions and the results from our examination of gender differences in promotion rates. Section 4 presents our approach to measure gender spillovers and results from that analysis. Section 5 concludes.

2. Data and Descriptive Statistics

This paper uses longitudinal linked employer-employee data on private sector workers in Norway based on register data from Statistics Norway and data on job types from an establishment-level survey compiled by Statistics Norway and the Confederation of Norwegian Enterprise (NHO, Næringslivets Hovedorganisasjon), the primary employer association in Norway. The goal of the data collection was to obtain an overview of earnings among white-collar workers; response rates among surveyed establishments (with at least 5 such workers) were very high because employers were legally bound to collect and report their data. Our

sample covers the time period 1987 to 1997, which is the longest time period for which all of our key variables are available, and includes workers born between 1936 and 1969.¹⁰

Because we rely on the NHO survey, our estimation sample is limited to workers at firms that are NHO members. Hence, the sample over-represents manufacturing and under-represents finance and public services. See Table 1A in the Appendix for the distribution of workers across industry groups in our sample of establishments and in the full working population of Norway in 1997. Establishments in the NHO sample also tend to be larger and older than average and their employees tend to be more educated and to earn higher wages. Nevertheless, the sample has broad coverage. In 1997, NHO employees represented 37 percent of all male workers and 27 percent of all female workers in Norway (Appendix Table 1A).

We also restrict our data to white-collar workers to ensure a relatively comparable set of jobs with substantial presence of both male and female workers. Women comprise 28 percent of our sample that includes all ranks and 30 percent of the sample of workers below the top rank, who have a chance of promotion. Like any employment-based sample, our sample is limited to current workers and excludes people who are unemployed or out of the labor market. Workers on temporary leave from a job, such as disability or parental leave, however, are included in the data.

The rest of this section describes the key outcome measures and explanatory variables used in the analysis.

¹⁰ Before 1987, only limited register information is available from Statistics Norway. After 1997, the NHO ceased collecting the data. For more details, see the yearly publication *Lønnsstatistikk for funksjonærer*. The birth cohorts 1936 to 1969 are those who can potentially be followed during the 10 year observation window.

2.1 Data on Promotions

Jobs in the NHO data are categorized into intersecting occupational groups and hierarchical ranks. The occupational groups are: (A) technical; (B) production supervision; (C) administrative, including clerical, accounting, and shipping; (D) sales; (E) storage; and (F) other. Within each occupational group, up to seven “task levels” are distinguished, ranging from unskilled and routine tasks in the lower levels up to high-skill and leadership tasks in the higher levels. (Appendix Table 2A contains sample job titles and occupations within each level.) These task levels and job descriptions are used to construct our seven hierarchical ranks. We define promotions as year-to-year changes in job category that entail an increase in hierarchical rank. Increases in rank can occur within an occupational group or can involve an occupational transition. Promotion rates may differ depending on starting occupational group, which implies that sex-based occupational segregation (Blau et al., 2012) can potentially be a source of differential promotion rates. In the analysis that follows, we eliminate this channel out by controlling for occupational group fixed effects.

Because not all establishments have positions in all seven ranks and many workers stay at their same establishment from one year to the next, a promotion across multiple ranks may have different meanings in different establishments.¹¹ We therefore focus on binary measures of progress instead of comparing the size of the rank change. Because these ranks are defined consistently across all establishments in the sample and over time, we are able to track career progress for individuals who change jobs within the sample. Our primary measure exploits this

¹¹ For example, rank 5 is far less commonly observed than rank 4 or 6. Smaller establishments that offer no jobs at this rank will have a “standard” promotion of two ranks for workers going from rank 4. In general, workers at establishments with deeper hierarchies will have more promotion opportunities for internal promotion. In addition to defining promotions as binary rank increases, we also address this issue in our analysis by including promotions that involve job changes in our outcome variable and by controlling for establishment fixed effects.

information and includes all promotions, but for completeness, we also study internal (within establishment) promotions and rates of staying at the same establishment.

Table 1 summarizes the distribution of workers across ranks in our sample, overall and separately by sex. Figure 1 plots the female share of workers in each rank. Women are vastly over-represented in the lowest rank, which is over 80 percent female. The next two ranks are closer to representative (women comprise 25 percent of rank 2 and 36 percent of rank 3; both comparable to their overall share of 28 percent of the sample). Rank 4 has even lower female presence (17 percent), and men dominate the three highest ranks, comprising over 90 percent of each. We also report a summary of different plant-year level measures of female representation in Appendix Table 3A. Across all years in our data, women are never more than 6 percent of the top three ranks, on average, even as their overall share of the average workplace increases from 25 to 33 percent.

Table 1 also reports mean values of our annual promotion rates in the sample (7.1 percent) and separately for men (7.3 percent) and women (6.6 percent). The vast majority of men and women (87 percent) stay in the same rank from year to year. But promotion rates vary dramatically across ranks, as shown in Table 2, dropping from 13 to 1.7 percent from ranks 1 to 6. This means that the overall gender difference in Table 1 understates the difference in promotion opportunities for similarly situated men and women. This variation is illustrated graphically in Figure 1, which plots the ratio of female to male promotion rates separately by rank. With the exception of the small fifth rank, this ratio is always substantially less than 1, ranging from 0.53 to 0.69.

The raw gender gap in internal promotions in Table 1 is substantially smaller than the gap for any promotions, suggesting that differential mobility may be part of the promotion gap, but

the sex ratios for internal promotion rates in Figure 1 fall into a similar range from 0.56 to 0.71 for ranks other than 5. The average shares of men and women staying at the same work establishment (conditional on remaining in the sample) are identical at 74 percent across all ranks (Table 1), though the shares staying tend to be somewhat higher for men, conditional on rank, and a large gap emerges at the highest rank (Table 2 and Figure 1). Women's greater likelihood of leaving their firm (or, as discussed in Section 2.3, of departing the sample entirely) from this most elevated rank leads to even lower rates of female representation in leadership than would be predicted from gender differences in promotions alone.

2.2 Explanatory Variables

The summary statistics in Table 2 demonstrate the importance of controlling for starting rank in attempting to measure gender differences in promotion opportunities. But starting rank is not the only observable dimension over which men and women in our sample differ. The summary statistics in Table 1 show gender differences in age, formal schooling, work experience, tenure at the current workplace, rank-specific tenure (time since initial entry or since the last promotion, whichever is more recent), with men having a higher average value for each of these variables. We also define a part-time status variable as having usual weekly work hours of less than 37 over the year. This variable is defined based on contracted hours, which means that leave taking does not reduce workers' average weekly hours. Nevertheless, women are significantly more likely to be working part-time (26 percent do so) than men (only 7.8 percent) in our sample.

Each of these variables may be related to worker productivity and may therefore affect promotion rates, so we use them as controls in our empirical analyses. Because we are concerned that the variables related to recent labor supply decisions such as experience, tenure and part-time status may themselves be endogenous, and affected by expectations about promotion

opportunities, we also present baseline models that exclude these variables. We include all of these controls in our main specification to provide a stronger test for how much of the gender gap in promotions is “unexplained” by observable work-related differences. We also estimate expanded versions of the model in Section 3 that add controls for family status. Specifically, we use birth register data to create indicator variables for whether or not the worker has any children under the age of 17 and under the age of 7.

In Section 4, we analyze how the female share at a workplace affects its promotion rates by adding explanatory variables that measure coworker demographics. The first is the female share among peers – defined here as workers at the same rank, work establishment and year. This variable has a mean of 30 percent for all workers, with a mean of 18 percent for men and 58 percent for women (Table 1), consistent with some sex segregation across workplaces and ranks. The second coworker variable is the female share among “bosses” working at the next highest rank at the same work establishment and year. This variable has a lower mean of 17 percent, reflecting women’s under-representation at higher ranks. The gender gap in this variable is smaller than for the peer variable; the mean values are 14 percent for men and 25 percent for women (Table 1). Across all observations in our sample, over 25 percent of workers have zero female bosses and only 1 percent of workers have all female bosses.

Two features of our measure of female bosses are worth noting. First, the measure is based only on establishment level rates; we are not able to identify workers’ immediate supervisors (as in Karaca-Mandic et al., 2013, and in Blau and DeVaro, 2007). If the effects of female bosses run exclusively within the chain of authority, then our estimates, based on a noisy measure, will tend to understate the importance of female bosses. To the extent that promotion decisions involve multiple decision-makers and that mentoring and role model effects occur

outside of the immediate chain of command, our approach will capture meaningful spillover effects. Second, our main measure of female bosses is rank-specific, and focuses on female representation at the next higher rank. This approach has the advantage of measuring effects of female leadership deeper within organizational hierarchies. In our robustness analysis of Section 4.3, we also explore the possibility that female shares two ranks higher or among the establishment's top leadership, defined based on rank, earnings or leadership job tasks, affect gender gaps in promotion rates.

2.3 Empirical Relevance of Gender Gaps in Promotion

Before turning to the econometric analysis of gender gaps in promotion, this section examines the empirical relevance of promotions, defined in this paper as increases in rank, in affecting two other gender gaps of interest: the gender gap in business leadership and the gap in pay. The first step of this examination is to quantify the contribution of sex differences in promotion rates into the outcome of sex differences in the distribution of workers across the hierarchical ranks. The second step is to show that these gendered distributions across ranks contribute substantially to gender pay differences in our sample.

In quantifying the role of promotions in gendered rank distributions, we use a simple framework to relate the three major factors that contribute to unequal distributions of workers across ranks by sex. The simple framework relates the distribution of workers across ranks in period t (the vector \underline{v}^t) to the basic inputs: the transition rates between ranks (\underline{M} ; where element M_{ij} is the year-to-year probability of going to rank i from rank j), the exit flow from each rank (\underline{e}), and the flow of new entrants into the ranks (\underline{u}). The exit flows can be used to construct a matrix \underline{R} that reflects that probability of staying in the sample into the next period, conditional on

current rank. The matrix $\underline{\mathbf{R}}$ has zeros in the off-diagonal elements, and values of $(1-e_i)$ for the diagonal elements R_{ii} .

The input flows can be combined with the current period's rank distribution to generate next period's rank distribution according to this equation:

$$(1) \quad \underline{\mathbf{v}}^{t+1} = (\underline{\mathbf{R}} \underline{\mathbf{M}}) \underline{\mathbf{v}}^t + \underline{\mathbf{u}}$$

Assuming that the three inputs are stable over time, the steady state distribution of workers across ranks will be:

$$(2) \quad \underline{\mathbf{v}}^\infty = (1 - \underline{\mathbf{R}} \underline{\mathbf{M}})^{-1} \underline{\mathbf{u}}$$

In this exercise, we use data from our sample on each of the inputs to predict the steady state distribution of workers across ranks. We can then examine the potential importance of each input – promotion rates, entry ranks, and rank-specific exit rates – by using the model to predict how eliminating gender differences in that input affects the steady state distribution of men and women across ranks.

There are various reasons why this model will not apply exactly on our sample: 1) the inputs are not stable over time, as female labor force participation is increasing over the period; and, therefore, 2) our data may not reflect a steady state. Also, we introduce noise in this exercise by not using the full transition matrix observed in our data but instead focusing on promotions alone and treating all promotions as increases of a single rank. Indeed, the steady state model (using average entry, exit and promotion rates for men and women observed over the entire period) predicts higher levels of female leadership than those in our data. The share of female workers in the top three ranks is 6.3 percent in the sample; the steady state predicted value is 14.4 percent, which suggests that our time period is one of transition toward greater gender equality (consistent with trends in Appendix Table 3A). Similarly, the ratio of the female share

in the top three ranks to the female share in the bottom three ranks is 0.16 in the data while the predicted steady state ratio is 0.33.

Varying each of the inputs in turn to equalize promotion, entry and exit rates between the sexes produces higher steady state predictions for female leadership shares. The female share in the top three ranks increases a similar amount under each change, going to 20 percent under equal promotion rates for men and women, to 21 percent under equal entry ranks, and to 18 percent under equal rank-specific exit rates. These changes also increase the predicted ratio of female shares in top and bottom ranks, but the effects differ by input. Equalizing entry ranks has the largest effect, driving the ratio to 0.67. Equalizing promotions increases the ratio to 0.48, but equalizing exits increases it to only 0.39 (this is because the change also lowers the chance that women in lower ranks exit the sample).

Although these values are unlikely to be accurate predictions of the exact impact of reducing the gender gap in promotions, they do support the quantitative importance of gender gaps in promotions in generating the observed gender gaps in leadership.

The second step of this examination is to measure the importance of gendered rank distributions in producing gender gaps in pay. This step fits with previous evidence that much of the gender pay gap is driven by women and men working in different jobs (e.g., Lazear and Rosen, 1990). For this step, we construct an hourly wage variable by taking the ratio of monthly earnings to normal hours of work.¹² Tables 1 and 2 present summary statistics for this variable, overall and by sex. The summary statistics in Table 2 (and the wage ratio curve in Figure 1) show initial evidence that rank is an important factor in predicting wages and creating gender

¹² The variable for normal work hours is based on contracted hours and excludes overtime hours. Earnings are from work and benefit claims that exclude overtime payments. Together, these allow us to measure hourly earnings from regular work.

wage gaps. The ratio of female average pay to male average pay across all ranks is 0.76. The ratios within each rank are closer to equality, ranging from 0.88 to 0.98.

We formalize this step by regressing the natural log of hourly pay on a *Female* indicator variable and four sets of control variables (described in Section 2.2) to measure the “unexplained” gender pay gap. For each set of controls, we estimate versions of the model with and without fixed effects for our 7 hierarchical ranks. The results of these regressions, reported in Figure 2, show that ranks play a substantial role in producing gender pay gaps. In the first column, with no other controls, the raw gap decreases by 59 percent (from 0.27 to 0.11) with the inclusion of rank effects. This pattern is repeated across the other models, in which the gender pay gaps are reduced by 55 to 57 percent with the rank controls.

The results in this section provide empirical support for the relevance of our primary outcome of interest. In particular, studying promotion rates can help researchers learn about a key mechanism that generates gender differences in leadership and pay.

3. Gender Differences in Promotion Rates

This section presents our empirical analysis of sex differences in promotions up the hierarchical ranks of the workplaces in our sample. As discussed in Francesconi (2001) and elsewhere, there are several theories of promotions. Internal promotion schemes can be viewed as a way to reward performance and thereby induce effort among lower-ranking workers. Internal promotions can also reflect workers accumulating human capital (some of it firm-specific) in lower tasks that increases their productivity at higher tasks or firms learning over time about the quality of their matches with individual workers. External promotions are more like hiring decisions and initial task assignments, but outside offers can also be affected by wages and promotion chances at the current employer.

Because, as discussed in the previous sections, men and women differ in many observable characteristics that may affect promotion rates under the various theories, we estimate regression models of promotion outcomes using various sets of controls and fixed effects and focus on measuring the unexplained or residual systematic variation in promotion rates related to gender.

Although some of this residual variation may be attributable to discrimination on the part of employers or higher ranking coworkers (based on discriminatory tastes, as in Arrow, 1973, or statistical beliefs about future commitment to the workplace as in Lazear and Rosen, 1990), our approach is not able to rule out supply-side differences in preferences and behavior between male and female workers.¹³ In fact, the notion of a sharp distinction between supply-side and demand-side explanations for gender differences may be misleading because it ignores the interactions between the two sides of the market. For example, if women tend to specialize in childcare and home production and invest less in their careers than men, profit-maximizing employers may react by offering them fewer employment and advancement opportunities and lower salaries. At the same time, if women face worse labor market opportunities than men, on average, that can justify and perpetuate a division of labor within households that has women specializing in home production. It is also worth noting that the extent to which fertility and household responsibilities affect workplace productivity also depends in part on the structure of jobs and workplaces.

3.1 Overall Gender Gaps in Promotions

Our basic estimation equation takes the following form:

$$(3) \quad Promotion_{ijt} = \beta_F Female_i + \beta_X X_{ijt} + \varepsilon_{ijt}$$

¹³ Examples of preference differences include tastes for risk (Eckel and Grossman 2008), competition (Niederle and Vesterlund 2007) or valuing money and career over relationships (Fortin 2008). For recent surveys, see Bertrand (2010) and Croson and Gneezy (2009).

The unit of observation is an individual i observed in year t working at establishment j . Where $Promotion_{ijt}$ is an indicator for that person being observed at higher rank in year $t+1$ than in year t , $Female_i$ is an indicator for the worker being female, X_{ijt} is the set of covariates, and ε_{ijt} is a random error term. We estimate all models using OLS and account for potential correlations in promotion rates among peers by clustering standard errors at the level of the plant-rank-year group.¹⁴ In this first phase of estimation, the X_{ijt} vectors in different models include different sets of individual characteristics and fixed effects. Because our outcome variable is defined based on *changes* in rank, this specification has the same spirit as a first-differences specification for rank, which accounts for unobservable differences that affect starting rank. One difference is that our outcome variable focuses exclusively on rank increases and ignores their magnitude. Another difference is that, in our model, the controls for unchanging factors (gender, education) and for those that identically increment by one for all of the workers in our sample (age, experience, tenure) account for how those factors affect the rate of change of an individual's rank rather than how they affect its current level. The level effects of these factors are impossible to estimate in a first-difference model because the variable for their year-to-year changes would have the same value for all observations (zero for the fixed and one for the incrementing factors).

In the basic version of the model, we control for year, industry, occupation, age, schooling and rank fixed effects. In this specification, female workers have a 3.5 percentage point lower chance of promotion, as reported in Column 1 of the top panel in Table 3. The

¹⁴ The marginal effects from Probit models corresponding to the specification in Column (2) of Table 3 are very similar to the OLS estimates: -0.027 (s.e. of 0.001) for all promotions, -0.020 (s.e. of 0.001) for internal promotions, and -0.0002 (s.e. of 0.002) for staying at the plant. In separate estimation, we also confirmed that the gender gaps in promotions (all and internal) remain statistically significant at the 1 percent level when errors are clustered at the plant-year or plant level. These alternative schemes account for possible correlations in promotion rates within establishments at a point in time or over time.

promotion outcome in this panel exploits the fact that we can track workers across establishments and includes any promotions, even those that include a workplace transition.

This 3.5 percentage point gender gap represents a 52 percent reduction relative to the average promotion rate for male workers. This is *larger* than the raw sex difference in promotion rates in Table 1, in large part because women are far more likely to be working in the lowest rank where promotion rates are highest (Table 2). The importance of controlling for current level may explain some of the discrepancies across other studies of sex differences in promotion rates: studies without such controls that compare men at higher starting ranks than women may not find significant gender differences even if they are present conditional on the same starting rank.

Across the four models in Table 3, we find consistent statistically and economically significant gender differences in promotion rates. Adding quadratic controls for experience and tenure, a linear control for rank-specific tenure, and an indicator for part-time status (defined as having usual weekly work hours below 37) in Column 2 reduces the *Female* coefficient to 2.9 percentage points, a 40 percent reduction from the male rate. Effects of age, schooling and civilian work experience are separately identified in our sample because of variation in time out of the labor market and time in military service. The full set of estimates for the control variables, other than fixed effects, are reported in Appendix Table 4A. The large and significant negative effect of part-time status on promotions explains much of the decline in the gender gap between Columns 1 and 2 of Table 3. Although it is illegal for employers in Norway to discriminate against part-time workers, it appears that their career progress is still slower.

The next column adds plant-level fixed effects and contains our preferred model. This change increases the size of the *Female* coefficient to -0.033 .¹⁵ This increase suggests that the lower promotion rates for women are *not* explained by women sorting into workplaces, possibly those with family friendly policies, that offer fewer promotion opportunities for *all* of their workers. Instead, it suggests that women tend to work at plants with higher promotion rates.

Finally, in the fourth column, we add controls for family status using two indicator variables: one for having any children under the age of 17 and another for any children under the age of 7. We also interact these variables with the *Female* indicator to allow children to have different effects on career progress for men and women. For men, fatherhood is associated with a greater chance of promotion, and there is no significant difference in the effects of younger children (estimates are in Column 5 of Appendix Table 4A). For women, children have a negative effect on promotion rates and that effect is even more negative if they are younger. This result is consistent with studies of wage growth that find flattening of wage profiles after motherhood (e.g., Miller, 2011; Ejrnæs and Kunze, 2013).¹⁶ The *Female* coefficient in Column 4 of Table 3 should be interpreted as the gender difference in promotions for workers with no children. It is smaller than the overall estimate of -0.033 in the previous column, but still economically and statistically significant (-0.025).

¹⁵ Using firm fixed effects instead produces nearly identical results; see Appendix Table 4A, Columns 2 and 3. This suggests that average promotion rates are similar across establishments at multi-establishment firms.

¹⁶ In a separate regression model, we also account for changes in Norwegian parental leave policy that extended the total leave period (from 165 to 294 days at 100 percent salary replacement, or from 210 to 364 days at 80 percent, subject to a salary cap) and that introduced paternity leave (reserving 4 weeks to the father). In the regression we add interaction terms between a Post-1992 indicator (after the policy change) and the two indicators for having young children. These interactions are statistically significant (and the policy is associated with a small decline in promotions for fathers of young children and an increase in promotions for mothers of young children), but they leave the other coefficient estimates unchanged.

In addition to these main specifications, we also estimated an exploratory model with an additional indicator for earning a high wage (in the top 30 percent) for one's rank, establishment and year. This variable is meant to proxy for the productivity of the worker relative to other contenders for promotion;¹⁷ it is a strong predictor of promotion within the next year (coefficient of 0.074, Column 7 of Appendix Table 4A). The mean value of this variable is substantially higher for men than for women (Table 1), which suggests that it might explain part of the gender gap in promotions. Indeed, including the variable reduces the *Female* coefficient somewhat (to -0.023), but it remains highly significant. We omit this variable from our preferred models, however, because it combines true productivity with evaluation and compensation decisions by employers that may themselves contain gender biases. Including it as a control therefore runs the risk of over-controlling for differences in how men and women are treated in the workplace.

In the second panel of Table 3, we repeat the analysis for promotion outcomes within the same establishment. Individuals who leave the establishment are assigned a zero value for this variable, so the sample size is unchanged between the two panels of the table. For internal promotions, we again find negative and significant *Female* coefficients across the different specifications, though the magnitudes are smaller, ranging from -0.028 to -0.020 . These values represent 34 to 47 percent lower internal promotion rates for women relative to the mean rate for men of 0.058 (Table 1). The estimates are similar for the two measures of promotions, but the higher estimates for the main measure suggests that part of the gender gap is from men being more likely to change workplaces in order to obtain a promotion.

We explored overall sex differences in turnover rates in the bottom panel of Table 3 using the same sample as the previous panels. The outcome variable is an indicator for staying at

¹⁷ Because our rank levels are broad enough to include multiple job titles and sub-ranks, high wages might also indicate having achieved a higher sub-rank within the rank.

the same plant in the next year. We find no significant gender differences in this outcome across the specifications in Table 3.¹⁸ The lack of overall differences in workplace transitions within the sample does not preclude a role for differential mobility in generating gender gaps in promotions. First, our sample is limited to individuals who are in our sample of private sector workers in two adjacent years, so we ignore differences in exit rates from the population of interest. Second, average transition rates may be the same for men and women, but the reasons for job changes may still differ systematically (for example, if men leave for better pay and higher status and women leave for greater workplace flexibility). Finally, if men who receive competing job offers are more likely to receive attractive counter-offers from their current employers, the sexes could end up with similar job change rates as men become more likely to be promoted internally.

3.2 Gender Gaps in Promotions by Rank

A key question about the average differences in promotion rates is if the differences are present across all ranks or if they are mainly concentrated at the lowest ranks, corresponding to a “sticky floor” story, or at the highest ranks, because of a “glass ceiling” beyond which promotions become more difficult or impossible for women.¹⁹ Different papers have found evidence of gender gaps in promotion at different points in organizational hierarchies: Smith et al. (2013), find gaps at the highest level for promotions from VP to CEO but not for promotions to VP, and single-firm studies such as Yap and Konrad (2009), Jones and Makepeace (1996), and Petersen

¹⁸ This also holds for the additional specifications shown in Appendix Table 6A, with one exception. In the model with firm fixed effects in Column 3, the *Female* coefficient is positive and significant. This difference is not robust to including plant fixed effects.

¹⁹ Bjerk (2008) and Yap and Konrad (2009) discuss the different career stages at which gender differences may emerge. In addition to the “glass ceiling” and “sticky floor”, there can also be a “mid-level bottleneck”. This concept of a sticky floor differs somewhat from that Booth et al. (2003), which instead relates to wages within ranks.

and Saporta (2004), find gender gaps at the lower and middle ranks. This paper is the first multi-firm analysis that measures promotion gaps separately at all hierarchical levels.

The results in Table 4 are from an expanded version of Equation 3 that includes a full set of interaction terms between *Female* and the 6 rank indicators (for each of the levels from which promotion is possible). We report estimates from our preferred specification in the previous table (Table 3, Column 3) with controls for industry, occupation, rank, year, age, schooling, experience, tenure, and part-time status, as well as establishment fixed effects. Across all starting ranks, we find that women are less likely than men to be promoted to a higher rank by the next year. This is true for either measure of promotions (Columns 1 and 2). This finding indicates that the sex differences in promotion rates at the highest ranks found in Smith et al. (2013) and for entry-level workers in the single-firm studies are present across all hierarchical ranks.

The outcome in the final column is staying at the same plant. There are no significant gender differences in the lower ranks, but a significant (at the 10 percent level) gap emerges at the sixth rank.

4. Gender Spillovers in Promotion Rates

Having documented substantial and significant gender gaps in promotion rates in Section 3, this section turns to the second phase of our analysis and asks if the gender mix of coworkers at workplaces affects these gaps. Specifically, we ask if increases in the shares of female bosses and female peers affect gender differences in promotion rates.

The prior literature suggests several reasons why gender spillovers from bosses might be positive. Positive spillovers will occur if higher-ranking women serve as mentors, role models,

and advocates for their lower-ranking coworkers (Athey et al., 2000).²⁰ Spillovers from bosses will also be positive if taste-based discrimination (against people of the opposite sex) or statistical discrimination (for example, if people are better at evaluating the work performance of others who are more similar to them, as in Aigner and Cain, 1977) is important for promotions. Increasing female representation at higher ranks may also improve promotion rates for lower-ranking women by weakening the associations of leadership with masculinity (Koenig et al., 2011) and providing reassurance to decision-makers that women are capable of performing well at higher ranks.²¹ Finding positive gender spillovers from female bosses, under any of these theories, would suggest that men’s historical dominance of business leadership can be self-perpetuating in equilibrium, but may potentially be disrupted by increasing female shares at higher ranks.

Although policymakers and advocates for increasing female business leadership often assume that gender spillovers from bosses are positive, it is worth noting that this need not be the case. The spillovers from female bosses can be negative if the “queen bee” phenomenon is common. This happens when a woman who achieves career success in a male-dominated field blocks other women from advancing (Staines et al., 1974). Although women are often expected to be more favorable judges of other women’s work,²² there is evidence of the opposite pattern in some male-dominated fields, which could also lead to negative spillovers, as in Bagues and

²⁰ An effective mentor can improve a worker’s performance by providing useful information and guidance, and can increase their chances of promotion with or without improved performance. Ibarra et al. (2010) distinguish between “mentors” and “sponsors”: the former provide advice and emotional support and can be at any rank, while the latter are relatively senior and act as advocates. They suggest that one reason that women with mentors report lower promotion rates than men with mentors is that women’s mentors tend to be less powerful in their organizations and less likely to act as sponsors.

²¹ It is also possible that sexual harassment against female workers is less frequent in settings with more female bosses.

²² There is also experimental evidence from cognitive psychology that unconscious biases against women in performance ratings are less severe when women comprise a larger share of the team (summarized in Valian, 1998, pp. 139-44).

Esteve-Volart (2010). The authors hypothesize this occurs because female evaluators have a “self-enhancement” motive to identify with the male majority.

Gender spillovers among peers can also be positive or negative. For example, women may be more willing to compete with other women (as in Gneezy et al., 2003), which could lead to higher overall female performance and higher promotion rates. However, unlike the laboratory setting with individual tasks, the workplace often involves both individual and cooperative tasks. While it is possible that women collaborate more with other women, it is also possible that women feel that their closest competitors for promotion are other women. This could happen, for example, in response to tokenism at workplaces or informal limits on the numbers of women in higher ranks (Kanter, 1977). In that case, women might be less cooperative with one another than with men and even engage in sabotage (Lazear, 1999), which could lower their relative promotion rates. This form of gender discrimination in promotions, through which female workers are competing most intensively with their female peers to be promoted into positions occupied by female bosses, would imply negative gender spillovers within ranks but positive spillovers from higher-ranking to lower-ranking women. Without this type of competition, there is no reason to expect a mechanical relationship between female shares at the current or higher rank and the likelihood that any individual woman is promoted. Another reason for negative spillovers from same-sex peers could be that more within-gender group socializing leads to more shirking from work; in contrast with the previous mechanism, this would imply positive (or zero) gender spillovers across ranks.

4.1 Empirical Approach and Identification

Motivated by the theoretical ambiguity of these effects, we measure gender spillovers on promotion rates by adding variables and interaction terms to our empirical model as follows:

$$(4) \quad Promotion_{ijt} = \beta_F Female_i + \beta_X X_{ijt} + \beta_P FemaleSharePeers_{ijt} + \\ \beta_{PF} FemaleSharePeers_{ijt} * Female_i + \beta_P FemaleShareBosses_{ijt} + \\ \beta_{PF} FemaleShareBosses_{ijt} * Female_i + \varepsilon_{ijt}$$

Because the measures of female shares among peers and bosses are correlated with one another, it is important to include them together in the regression models.

Adding these spillover variables to the model potentially introduces new identification challenges related to the estimation of social interactions or peer effects (e.g., Manski, 1993). These variables are direct measures of “contextual effects” (related to peer characteristics rather than the “endogenous effects” of peer behavior). It is important to be clear that their estimated effects will incorporate both the direct effect of the characteristics as well as any effects that are caused by the behavior of these peers that is related to their gender. In that sense, we are not aiming to isolate the purely contextual effects in our estimates.

The two other identification challenges related to estimating spillover effects derive from the possibility of common shocks (unobservable to researchers) affecting all coworkers and endogenous group formation.²³ It is not obvious what direction we should expect these biases to have. For the estimated effects of peers at the same rank, we might expect positive bias in the *FemaleSharePeers*Female* interaction if women are drawn to establishments where they face higher chances of promotion; those workplaces would have higher shares of female peers and higher promotion rates for women. By contrast, if female bosses attract more female workers at lower levels, that would not bias the estimated effects of female bosses on promotion rates unless the female bosses were able to attract higher *quality* (along unmeasured dimensions) female candidates. In that case, there would be a positive correlation between female bosses and female

²³ Because we are not including endogenous peer effects (promotions and retention) as covariates in our model, our estimates will not suffer from the “reflection problem”; only common shocks and systematic worker sorting that are correlated with the peer demographics and other controls will introduce bias.

promotion rates that is based in part on “selection” into the workplace rather than events after the worker joins the establishment. This type of selection could still be interpreted as a form of gender spillover, but it would differ from the channels related to human capital investments, productivity or discrimination.

In the next section, we can address the major concerns about omitted variable biases, potentially related to non-random sorting of workers to workplaces, by including different sets of control variables and fixed effects in our regression models. Our main models in Section 4.2 include a full set of establishment, rank and year fixed effects and Section 4.3 shows the robustness of the estimates to including additional controls and interaction terms.²⁴

4.2 Main Results for Gender Spillovers

Our main estimates of gender spillovers from the model in Equation 4 are reported in Table 5. As in Table 3, the top panel presents results for any promotions, the middle panel is for internal promotions, and the bottom panel is for staying at the same plant. The three columns include the same controls and fixed effects as Columns 2 to 4 of Table 3: the first has controls for industry, occupation, rank, year, age, schooling, experience, tenure and part-time status; the second adds plant fixed effects; and the third adds the two indicators for having children (under the ages of 17 and 7) and their interactions with the *Female* indicator.

We find significant gender spillovers from peers and bosses that act in opposite directions. Having more female workers in the next highest rank tends to *reduce* the gender gap in promotion rates by raising women’s relative promotion chances, but having more female

²⁴ Given our empirical setting, we are clearly not able to exploit experimental variation in female representation for identification. An approach that is sometimes used in the peer effects literature to generate quasi-random variation in peer groups using observational data is to focus on individuals who stay in a group, but whose peer characteristics are changed by others who enter or leave the group. Unfortunately, given the endogenous nature of the decisions to leave a workplace and which new workplace to enter, this strategy is not appealing in our setting.

coworkers in the same rank tends to *increase* it. The coefficient on the *Female*FemalePeerShare* interaction ranges from -0.060 to -0.062 , and is substantially larger than the *Female*FemaleBossShare* interaction of 0.018 to 0.020 .

Adding plant fixed effects has only minor effects on the estimates, which provides empirical evidence against non-random selection into workplaces being an important source of bias for our estimated spillover effects.²⁵ The coefficient of 0.019 from our main specification in Column 2 implies that a standard deviation increase in the female boss share (an increase of 0.24) would decrease the gender gap in promotions by 13.8 percent (from a base of 0.033 ; see Table 8). The model in Column 2 also implies that increasing the female share among peers by a standard deviation (or 0.29) would increase the gender gap in promotions by 52.7 percent. These patterns are repeated in the middle panel of the table for internal promotions.

The positive interaction effect from female bosses is consistent with women helping women below them in the hierarchy through mentoring or taste-based discrimination by supervisors. The finding provides further support for the theory that demand-side factors play a role in creating the observed gender gaps in promotion rates. It suggests that one reason for women's slow progress to the top of corporate hierarchies is the historical male domination of those ranks. In that sense, it provides empirical support for the presence of a key mechanism underlying public policies aimed at promoting women in business leadership positions – that female leaders tend to help lower-ranking women in their organizations.

However, the finding of negative gender interactions among women at the same rank complicates this story of women helping women. Because many of the theories of positive gender spillovers (mentoring, sponsorship, role models, access to powerful networks) are more important for bosses than for peers, the absence of positive spillovers may not be surprising. As

²⁵ We report estimates from models with additional sets of fixed effects in the next subsection.

discussed above, the negative effects of female peers may come from greater competition (and less cooperation) among peers of the same sex (which would also generate the positive effect on male peers) or from women in lower ranks facing greater competition for scarce sex-specific resources such as mentors and sponsors. Interestingly, Karaca-Mandic et al. (2013) also find negative gender spillovers among peers in the lower ranks of the US Army.

Although female bosses are associated with higher promotion rates for female workers, they appear to have no significant effects on promotion rates for males. This gender difference provides further evidence against the presence of an omitted variable that is related to promotion rates overall (such as corporate growth and expansion) and to female boss shares (for example, if growing firms are able to increase diversity at higher ranks more rapidly). The gender differences are larger for the effects of female peers, where the estimates show positive and significant effects on male promotion rates across the models. These positive effects for men suggest that the negative effects on women are not driven by an omitted variable that is related to the female share in a rank and overall lower promotion opportunities, which would be the case if women tended to sort into “dead-end” job groups identifiable by their work establishment and hierarchical rank combination. Furthermore, the fact that we find opposite effects of female peers and bosses also rules out stories of sorting or workplace-specific omitted variables (such as female-friendliness) that create spurious correlations between female representation at all levels, including bosses and peers, and promotion outcomes for women.

Finally, the results in the bottom panel of Table 5 show no significant gender spillovers in the probability of staying at the same workplace, after plant fixed effects are included in the model (Columns 2 and 3). The significant spillover coefficients in Column 1 seem to suggest that there is some sorting out of workplaces based on gender compositions of peers and bosses.

However, their dramatic reduction in size and significance in the latter two columns indicates that this sorting is captured with establishment fixed effects. As discussed above, these same fixed effects leave the spillover effects on promotions unchanged.

Finally, in an unreported regression model, we explored the interactions between plant mobility and gender spillovers in promotion rates. Using the main promotion outcome, we added a control variable to capture changing plants during the time period and interactions between this variable and the *Female*, *FemaleBossShare*, *FemalePeerShare*, *Female*FemaleBossShare*, and *Female*FemalePeerShare* variables. The last two interaction terms show that the positive gender spillovers from female bosses are indistinguishable between workers who stay at the same plant and those who change plants (the point estimates is 0.002, with a standard error of 0.007, relative to a boss spillover of 0.018, standard error of 0.006), but the negative gender spillovers from female peers are about 50% larger for workers who change plants (the point estimate on the interaction is -0.023 , with a standard error of 0.008, relative to a peer spillover of -0.053 , standard error of 0.006, for those who stay at the same plant). This suggests that, while female peers do not affect the relative mobility rates of male and female workers, they may affect the reasons or circumstances associated with changing plants.

4.3 Robustness Checks for Gender Spillover Results

This subsection examines the robustness of the main gender spillover results for the main promotion outcome in Section 4.2.²⁶ Table 6 expands the main model in Column 2 of Table 5 with plant fixed effects but without fertility controls by adding various sets of additional controls

²⁶ Results for the other outcomes are similarly robust. As in Section 3.1 above, the OLS estimates in Section 4.2 (Table 1, Column 1) are nearly identical to the average marginal effects from corresponding Probit models. The main relationships are also statistically significant under alternative clustering schemes (at the establishment-year or establishment level).

and fixed effects. The goal of this exercise is to further explore the potential role of selection into workplaces and omitted variables on the main spillover effects from bosses and peers.

The first column addresses the concern that female dominated ranks may differ in size from male dominated ranks and that these size differences may affect promotion rates. Although there promotions are not zero-sum in the sense that only one person can advance up from a particular rank at each establishment in a given year, there may be restrictions on the number of available slots at the higher ranks. When those limits are more binding, competition among peers may be more intense. We are not able to observe the number of candidates eligible for promotion or the number of available slots, but we proxy for each using the size of the current rank (at the establishment and year) and the next highest rank (at the same establishment and year). Adding these controls leaves the main results unchanged: it reduces the estimated gender spillover from bosses by 0.003 to 0.016 and the estimated gender spillover from peers by 0.004 to -0.056 .

Column 2 adds a full set of interactions between the *Female* and rank indicators. This is to address the concern that differences in the female shares of peers and bosses may be related to workers' ranks, which would mean that the estimated interactions between the coworker variables and the *Female* indicator reflect in part the cross-rank differences in gender gaps in promotion found in Table 4. Column 3 addresses the concern that within-establishment variation across ranks may be related to both female shares and promotion changes by adding a full set of establishment-by-rank fixed effects. In both cases, the estimated gender spillovers remain statistically significant and qualitatively unchanged.

We next consider the possibility that our results are driven by sorting *within* ranks and establishments that leaves women disproportionately working in occupational groups with limited promotion opportunities at their workplaces. We assess the importance of this

explanation by adding a control to our model for the share of workers in the next higher rank in the organization from the worker's current occupational group (recall that these groups span multiple ranks). Including this variable does not alter the estimated gender gap in promotion rates overall or (as shown in Column 4) the estimated gender spillovers in promotions from peers or bosses. Controlling for the female share in the rank *below* (unreported) also leaves the estimates unchanged; the female share in the lower rank does not affect promotion rates for men or gender differences in promotion rates.

The last two columns of Table 6 are motivated by concerns about time-varying omitted variables that are not addressed by the establishment and rank fixed effects. The first is that selection into establishments is based on the current reputation of the establishment (which is related to current promotion rates) and that this variable (unobserved by researchers) varies over time. The second is that the relationship between ranks and promotion rates is changing over time. These time-varying factors can introduce bias in our models because female representation in the sample, and in the higher ranks, are both increasing over the period (Appendix Table 3A). We address the first point in Column 5 with the inclusion of establishment-by-year fixed effects and the second point in Column 6 with rank-by-year fixed effects. Estimates with these additional controls confirm the results from Table 5.

The second set of robustness checks, in Table 7, explores how the estimated effect of female bosses on the gender gap in promotions varies under three alternative measures of the female boss share. In Columns 1 and 2, we use a measure based on the female share of workers two ranks higher than the worker's current rank. This incorporates the possibility that individuals in the higher rank are involved in deciding on promotions to the rank below them. This type of gender spillover, if present, also has the appealing feature of potentially generating a cascade of

effects down the hierarchy. Such effects are not implied if gender spillovers are limited to effects from the immediate rank above onto promotions into that rank. These spillovers appear to be positive and substantial in our data (larger than the estimates from our primary measure) and they are present even after controlling for our primary measure (Column 2).

Column 3 uses an establishment-wide measure of female leadership based on the female share in the top three ranks (averaging about 4 to 5 percent; Appendix Table 3A) and uses a sample limited to workers in ranks 1 to 4. Column 4 identifies leaders based on the top earners in the establishment and year (using a top 20 percent cutoff). Column 5 defines bosses to include all workers whose detailed job descriptions contain leadership tasks. The female boss share is lowest under this definition (Appendix Table 3A). Across each of these three alternative measures, the estimated effect of female bosses on the promotion rates of women relative to men are positive, substantial, and statistically significant. These results confirm the robustness of the overall finding of positive gender spillovers from female bosses to female workers and suggest that various types of female bosses can each contribute to the career progression of female employees.

Table 8 provides some perspective on the magnitude of the estimated effects of female coworkers relative to sample variation in those measures and to the overall gender gap in promotions in our data. The first row reports the coefficients from our preferred estimation models for each measure of female shares, the second notes the location (table and column number) of the estimate earlier in the paper, and the third reports the sample standard deviation of the female share measure. Finally, the last row reports the predicted effect of a standard deviation increase in the female share of peers or bosses (for each of the alternative measures) on the gender gap in promotions. The negative gender spillovers from peers are largest, accounting

for a 52.7 percent increase in the gap. The positive gender spillovers from bosses are smaller, but still substantial, with a standard deviation increase in female leadership narrowing the gap by 11.6 to 21.1 percent for the first four measures and by a more modest 4.1 percent for the final and most stringent measure (with the lowest levels of female representation, as shown in Appendix Table 3A). This pattern confirms the quantitative importance of female leadership within organizations, and not just at the highest levels, to ensure female progress up the ranks.

5. Conclusion

This paper uses an 11-year employer-employee matched panel on white-collar workers at over 4,000 workplaces in Norway to study gender differences in promotion rates across seven hierarchical ranks. We find that women are significantly less likely than men to be promoted to a higher rank in the organizational hierarchy, conditional on a wide range of individual characteristics and workplace, time and rank fixed effects. This gender gap in promotions is present whether we include or exclude promotions that involve a change in workplace. It is also present across all six of the lower ranks from which workers can be promoted, which indicates that both “sticky floors” and “glass ceilings” are contributing to women’s severe underrepresentation at the highest levels of corporate hierarchies.

Using data on the gender composition of workplaces, we next show how the promotion gap is affected by the presence of female bosses and female peers. Several theories predict the increasing female leadership can improve outcomes for lower-ranking women: lower-ranking women may gain access to better mentors and role modeling and to powerful professional networks. Consistent with this prediction, we find that higher shares of female workers at the next highest rank are associated with significantly smaller gender gaps in promotion. This finding is not coming from higher relative promotion rates for women overall at firms that hire

more women in all ranks (because they are more “female friendly” in general). In fact, we find that having a greater share of female coworkers at the same establishment and rank actually depresses promotion rates for women, relative to men. This negative effect may be caused by greater competition within gender (either between men or women) among peers, either for promotions or for gender-specific mentoring and support.

The finding that greater female representation at higher ranks narrows the gender gap in promotion rates at lower ranks supports the theory that men’s dominance in corporate leadership continues to present barriers to women’s advancement in corporate hierarchies. It also suggests that policies that promote greater female representation in corporate leadership will have spillover benefits to women in lower ranks. There are, however, reasons to believe that our estimates may understate the true importance of gender spillovers from female leaders to other female workers. Specifically, because our empirical setup is focused on measuring spillovers within organizations, our estimates will fail to capture any spillover benefits from female leaders to women in other firms or establishments. These can be important if mentoring occurs across these boundaries (possibly through professional or social groups outside of the workplace) or one of the ways in which women help other women progress in their careers is by hiring them through external recruitment from other organizations.

References

- Aigner, Dennis J., and Glen G. Cain (1977). “Statistical Theories of Discrimination in Labor Markets.” *Industrial and Labor Relations Review* 30 (2): 175–87.
- Arrow, Kenneth (1973). “The Theory of Discrimination,” in *Discrimination in Labor Markets*, eds. Orley A. Ashenfelter and Albert Rees. Princeton, NJ: Princeton University Press.
- Athey, Susan, Christopher Avery, and Peter Zemsky (2000). “Mentoring and Diversity.” *American Economic Review*, 90 (4): 765-786.
- Bagues, Manuel F. and Berta Esteve-Volart (2010). “Can Gender Parity Break the Glass Ceiling? Evidence from a Repeated Randomized Experiment.” *Review of Economic Studies* 77(4): 1301–1328.

- Barnett, William P., James N. Baron, and Toby E. Stuart (2000). "Avenues of Attainment: Occupational Demography and Organizational Careers in the California Civil Service." *American Journal of Sociology* 106(1): 88–144.
- Barsh, Joanna and Lareina Lee (2011). *Unlocking the Full Potential of Women in the US Economy*. McKinsey & Company Report.
- Bell, Linda A. (2005). "Women-Led Firms and the Gender Gap in Top Executive Jobs." IZA Working Paper.
- Bertrand, Marianne (2010). "New Perspectives on Gender." in Orley Ashenfelter and David Card eds., *Handbook of Labor Economics*, Volume 4B. Elsevier.
- Bjerk, David (2008). "Glass Ceiling or Sticky Floors? Statistical Discrimination in a Dynamic Model of Hiring and Promotion." *Economic Journal* 118(530): 961–82.
- Blau, Francine D., Peter Brummund, and Albert Yung-Hsu Liu (2012). "Trends in Occupational Segregation by Gender 1970-2009: Adjusting for the Impact of Changes in the Occupational Coding System." NBER Working Paper No. 17993.
- Blau, Francine D. and Jed L. DeVaro (2007). "New Evidence on Gender Differences in Promotion Rates: An Empirical Analysis of a Sample of New Hires." *Industrial Relations* 46(3): 511-50.
- Booth, Alison L., Marco Francesconi, and Jeff Frank (2003). "A Sticky Floors Model of Promotion, Pay, and Gender." *European Economic Review* 47(2): 295-322.
- Cardoso, Ana and Rudolf Winter-Ebmer (2010). "Mentoring and Segregation: Female-Led Firms and Gender Wage Policies." *Industrial and Labor Relations Review* 64(1):143-163.
- Ciliberto, Federico, Amalia Miller, Helena Skyt Nielsen and Marianne Simonsen (2013). "Playing the Fertility Game at Work: An Equilibrium Model of Peer Effects." Working Paper.
- Cohen, Lisa E., Joseph P. Broschak, and Heather A. Haveman (1998). "And Then There Were More? The Effect of Organizational Sex Composition on the Hiring and Promotion of Managers." *American Sociological Review* 63: 711-727.
- Croson, Rachel and Uri Gneezy (2009). "Gender Differences in Preferences." *Journal of Economic Literature*, 47(2): 1-27.
- Eckel, Catherine C., and Philip J. Grossmann (2008). "Men, Women and Risk Aversion: Experimental Evidence." in Plott, C. and V. Smith (eds.), *Handbook of Experimental Results*.
- Ejrnæs, M. and A. Kunze (2013): "Work and Wage Dynamics around Childbirth." *Scandinavian Journal of Economics* 115(3): 856-877.
- Flabbi, Luca, Mario Macis, Andrea Moro, and Fabiano Schivardi (2014). "Do Female Executives Make a Difference? The Impact of Female Leadership on Gender Gaps and Firm Performance." IZA DP No. 8602.
- Fortin, Nicole M. (2008). "The Gender Wage Gap among Young Adults in the United States The Importance of Money versus People." *Journal of Human Resources* 43(4): 884-918.
- Francesconi, Marco (2001). "Determinants and Consequences of Promotions in Britain." *Oxford Bulletin of Economics and Statistics*, 63(3): 279-310.
- Gerhart, Barry A., and George T. Milkovich (1989). "Salaries, Salary Growth, and Promotions of Men and Women in a Large, Private Firm." In *Pay Equity: Empirical Inquiries*, eds. Robert T. Michael, Heidi I. Hartmann, and Brigid O'Farrell, pp. 23–43. Washington, DC: National Academy Press.

- Giuliano, Laura, David I. Levine, and Jonathan Leonard. 2005. "Do Race, Gender, and Age Differences Affect Manager-Employee Relations? An Analysis of Quits, Dismissals, and Promotions at a Large Retail Firm." Paper 18. Center for Responsible Business, University of California, Berkeley.
- Gneezy, U., Niederle, M. and Rustichini, A. (2003). "Performance in Competitive Environments: Gender Differences." *Quarterly Journal of Economics*, 118 (3): 1049–1074.
- Goldin, Claudia, and Cecilia Rouse (2000). "Orchestrating Impartiality: The Impact of 'Blind' Auditions on Female Musicians." *American Economic Review* 90(4): 715–742.
- Hersch, Joni, and W. Viscusi (1996). "Gender Differences in Promotions and Wages." *Industrial Relations: A Journal of Economy and Society* 35(4): 461-472.
- Hunnes, Arngrim, Jarle Møen, and Kjell G. Salvanes (2009). "Wage Structure and Labor Mobility in Norway, 1980-97," in *The Structure of Wages: An International Comparison*, Edward P. Lazear and Katherine Shaw, eds., NBER and University of Chicago Press.
- Ibarra, Herminia, Nancy M. Carter, and Christine Silva (September, 2010). "Why Men Still Get More Promotions Than Women." *Harvard Business Review*.
- Jones, David R., and Gerald H. Makepeace (1996). "Equal Worth, Equal Opportunities: Pay and Promotion in an Internal Labour Market." *Economic Journal*: 401-409.
- Kanter, Rosabeth. M. (1977). *Men and Women of the Corporation*. New York, NY: Basic Books.
- Karaca-Mandic, Pinar, David Powell, and Nicole Maestas (2013). "Peer Groups and Employment Outcomes: Evidence Based on Conditional Random Assignment in the United States Army." Working Paper.
- Koenig, Anne M., Alice H. Eagly, Abigail A. Mitchell, and Tiina Ristikari (2011). "Are Leader Stereotypes Masculine? A Meta-Analysis of Three Research Paradigms." *Psychological Bulletin* 137(4): 616–642.
- Kurtulus, Fidan Ana, and Donald Tomaskovic-Devey (2012). "Do Female Top Managers Help Women to Advance? A Panel Study using EEO-1 Records." *Annals of the American Academy of Political and Social Science* 639: 173-197.
- Lønnsstatistikk for funksjonærer (1997). *utarbeidet av Næringslivets Hovedorganisasjon område for arbeidsmarked og Statistikk*. ISSN 0333-2608. Available annually.
- Lazear, Edward P. (1989). "Pay Equality and Industrial Politics." *Journal of Political Economy* 97(3): 561-580.
- Lazear, Edward P. and Sherwin Rosen (1990). "Male-Female Wage Differentials in Job Ladders." *Journal of Labor Economics*: S106-S123.
- Manski, Charles F. (1993). "Identification of Endogenous Social Effects: The Reflection Problem." *Review of Economic Studies* 60 (3), 531-542.
- Matsa, David A. and Amalia R. Miller (2011). "Chipping Away at the Glass Ceiling: Gender Spillovers in Corporate Leadership." *American Economic Review*, 101(3): 635–39.
- Matsa, David A. and Amalia R. Miller (2013). "A Female Style in Corporate Leadership? Evidence from Quotas." *American Economic Journal: Applied Economics* 5(3):136–169.
- Matsa, David A. and Amalia R. Miller (2014). "Workforce Reductions at Women-Owned Businesses in the United States." *Industrial and Labor Relations Review*, 67(2): 422–452.
- Miller, Amalia R. (2011). "The Effects of Motherhood Timing on Career Path." *Journal of Population Economics* 24(3): 1071-1100.
- Olson, Craig A. and Brian E. Becker (1983). "Sex Discrimination in the Promotion Process." *Industrial and Labor Relations Review* 36(4): 624–41.

- Petersen, Trond and Ishak Saporta (2004). "The Opportunity Structure for Discrimination." *American Journal of Sociology* 109(4): 852-901.
- Powell, Gary N. and D. Anthony Butterfield (1994). "Investigating the 'Glass Ceiling' Phenomenon: An Empirical Study of Actual Promotions to Top Management." *Academy of Management Journal*, 37(1): 68–86.
- Ransom, Michael, and Ronald L. Oaxaca (2005). "Intrafirm Mobility and Sex Differences in Pay." *Industrial and Labor Relations Review* 58(2): 219–37.
- Smith, Nina, Valdemar Smith, and Mette Verner (2013). "Why Are So Few Females Promoted into CEO and Vice-President Positions? Danish Empirical Evidence 1997-2007." *Industrial and Labor Relations Review* 66(2): Article 4.
- Staines, Graham L., Carol Tavris, and Toby E. Jayarante (1974). "The Queen Bee Syndrome." *Psychology Today* 7:55-60.
- Tate, Geoffrey, and Liu Yang (2013). "Female Leadership and Gender Equity: Evidence from Plant Closure." forthcoming in the *Journal of Financial Economics*.
- Valian, Virginia (1998). *Why So Slow? The Advancement of Women*. Cambridge, MA: MIT Press.
- Winter-Ebmer, Rudolf and Josef Zweimuller (1997). "Unequal Assignment and Unequal Promotion in Job Ladders." *Journal of Labor Economics* 15(1): 43–71.
- Yap, Margaret, and Alison M. Konrad (2009). "Gender and Racial Differentials in Promotions: Is There a Sticky Floor, a Mid-Level Bottleneck, or a Glass Ceiling?" *Relations Industrielles/Industrial Relations* 64(4): 593-619.
- Zander, Christina (May 21, 2014). "Even Scandinavia Has a CEO Gender Gap." *The Wall Street Journal*.

Table 1. Summary Statistics

	Male Mean	Female Mean	Total Mean	Total Std. Dev.	Total Obs.
Promotion	0.073	0.066	0.071		597,552
Internal Promotion	0.059	0.056	0.058		597,552
Stay at Plant	0.74	0.74	0.74		597,552
Stay at Rank	0.87	0.87	0.87		597,552
Female	0	1	0.3		597,552
Female Boss Share	0.14	0.25	0.17	0.24	597,552
Female Peer Share	0.18	0.58	0.3	0.29	597,552
Age	40.2	37.1	39.3	9.04	597,552
Years Schooling	12.6	11.7	12.3	2.62	597,552
Years Work Experience	18.4	14.5	17.3	7.33	597,552
Years Tenure	7.41	6.52	7.15	5.28	597,552
Rank Specific Tenure	4.29	4.04	4.22	2.76	597,552
Working Part-Time	0.076	0.26	0.13		597,552
Any Children	0.7	0.56	0.66		597,552
Children under 7	0.33	0.27	0.31		597,552
High Wage for Rank	0.27	0.13	0.23		597,552
Rank = 1	0.024	0.27	0.094		744,531
Rank = 2	0.18	0.15	0.17		744,531
Rank = 3	0.28	0.4	0.31		744,531
Rank = 4	0.29	0.14	0.25		744,531
Rank = 5	0.079	0.012	0.06		744,531
Rank = 6	0.13	0.025	0.1		744,531
Rank = 7	0.015	0.0011	0.011		744,531
Hourly Wage (NOK)	127.7	96.6	118.9	36.5	744,531

Notes: Data from NHO (1987-1997) and Statistics Norway. Sample includes workers in ranks 1 to 6 for all variables except rank indicators and average wage, which also include rank 7. Hourly wages are the ratio of earnings (measured in real Norwegian Kroner, NOK) and hours worked. Female boss share is defined for workers 1 rank higher.

Table 2. Promotion, Mobility and Wages by Rank and Sex

	Rank	All Mean	Male Mean	Female Mean
Promotion	1	0.13	0.18	0.11
	2	0.084	0.091	0.063
	3	0.064	0.077	0.041
	4	0.054	0.057	0.039
	5	0.069	0.069	0.072
	6	0.017	0.017	0.01
	7	0	0	0
	Total	0.066	0.067	0.063
Internal Promotion	1	0.11	0.15	0.098
	2	0.068	0.073	0.052
	3	0.052	0.062	0.035
	4	0.043	0.045	0.032
	5	0.057	0.056	0.061
	6	0.013	0.013	0.0077
	7	0	0	0
	Total	0.053	0.053	0.054
Stay at Plant	1	0.73	0.75	0.72
	2	0.69	0.7	0.67
	3	0.74	0.74	0.74
	4	0.73	0.73	0.72
	5	0.74	0.75	0.72
	6	0.72	0.73	0.69
	7	0.66	0.66	0.5
	Total	0.73	0.73	0.72
Hourly Wage	1	82.6	89.5	81
	2	98.3	100.9	90.4
	3	102.4	106.9	94.5
	4	129.9	131.7	120.7
	5	163.6	163.7	161.2
	6	173.5	174.4	162.1
	7	221.7	221.9	214.7
	Total	118.9	127.7	96.6

Notes: Data from NHO (1987-1997) and Statistics Norway. Sample includes workers in all ranks. Hourly wages are the ratio of earnings (measured in real Norwegian Kroner, NOK) and hours worked.

Table 3. Gender Differences in Promotions and Mobility

	(1)	(2)	(3)	(4)
	Promotion			
Female	-0.035***	-0.029***	-0.033***	-0.025***
	[0.0012]	[0.0012]	[0.0013]	[0.0015]
N	597,552	597,552	597,552	597,552
R ²	0.046	0.049	0.072	0.072
	Internal Promotion			
Female	-0.028***	-0.022***	-0.026***	-0.020***
	[0.0011]	[0.0012]	[0.0012]	[0.0014]
N	597,552	597,552	597,552	597,552
R ²	0.038	0.04	0.061	0.061
	Stay at Plant			
Female	0.0020	0.00013	0.0020	0.0030
	[0.0020]	[0.0020]	[0.0017]	[0.0020]
N	597,552	597,552	597,552	597,552
R ²	0.32	0.32	0.44	0.44
Fixed Effects for Industry, Occupation, Rank, Year, Age, Schooling	Yes	Yes	Yes	Yes
Controls for Experience, Tenure, Part-Time Status	No	Yes	Yes	Yes
Plant Fixed Effects	No	No	Yes	Yes
Children Fixed Effects	No	No	No	Yes

Notes: OLS estimates with standard errors clustered at the plant-rank-year level. Experience and overall Tenure controls are quadratic; Rank specific tenure is linear; Part-time status is <37 hours per week.

Children effects are indicators for any children under age 17, any children under age 7, and interactions of each with the Female indicator. Controls are defined in period t for promotion and mobility outcomes between periods t and t+1. Occupation groups are listed in Appendix Table 1A.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

Table 4. Gender Differences in Promotions and Mobility by Rank

	(1)	(2)	(3)
	Promotion	Internal Promotion	Stay at Plant
Female and Rank = 1	-0.048*** [0.0071]	-0.038*** [0.0064]	0.0045 [0.010]
Female and Rank = 2	-0.016*** [0.0026]	-0.012*** [0.0024]	0.00038 [0.0038]
Female and Rank = 3	-0.044*** [0.0018]	-0.035*** [0.0017]	0.0012 [0.0025]
Female and Rank = 4	-0.024*** [0.0024]	-0.019*** [0.0022]	0.0014 [0.0035]
Female and Rank = 5	-0.020*** [0.0065]	-0.015** [0.0062]	0.0086 [0.0081]
Female and Rank = 6	-0.0086*** [0.0028]	-0.0061** [0.0024]	0.013* [0.0066]
N	597,552	597,552	597,552
R ²	0.072	0.061	0.44

Notes: Estimates from OLS models with standard errors clustered at the plant-rank-year level. Controls as in Table 3, Column (3).

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

Table 5. Gender Spillovers in Promotions and Mobility

	(1)	(2)	(3)
	Promotion		
Female * Female Boss Share	0.020*** [0.0044]	0.019*** [0.0047]	0.018*** [0.0047]
Female * Female Peer Share	-0.060*** [0.0054]	-0.060*** [0.0057]	-0.062*** [0.0057]
Female Boss Share	0.0048 [0.0037]	-0.0058 [0.0049]	-0.0054 [0.0049]
Female Peer Share	0.050*** [0.0048]	0.036*** [0.0059]	0.037*** [0.0059]
N	597,552	597,552	597,552
R ²	0.05	0.073	0.073
	Internal Promotion		
Female * Female Boss Share	0.022*** [0.0040]	0.019*** [0.0043]	0.019*** [0.0043]
Female * Female Peer Share	-0.055*** [0.0051]	-0.052*** [0.0053]	-0.053*** [0.0053]
Female Boss Share	0.0019 [0.0033]	-0.0068 [0.0044]	-0.0065 [0.0044]
Female Peer Share	0.045*** [0.0045]	0.031*** [0.0056]	0.031*** [0.0056]
N	597,552	597,552	597,552
R ²	0.041	0.061	0.061
	Stay at Plant		
Female * Female Boss Share	0.015* [0.0080]	0.0044 [0.0068]	0.0046 [0.0068]
Female * Female Peer Share	-0.045*** [0.010]	0.012 [0.0086]	0.013 [0.0086]
Female Boss Share	-0.016** [0.0082]	0.0016 [0.0080]	0.0017 [0.0080]
Female Peer Share	0.031*** [0.0093]	-0.0032 [0.0098]	-0.0038 [0.0098]
N	597,552	597,552	597,552
R ²	0.32	0.44	0.44
Plant FE	No	Yes	Yes
Children	No	No	Yes

Notes: OLS estimates with standard errors clustered at the plant-rank-year level. Basic controls in all columns are as in Table 3, Column (2). Female Boss Share is the female share of workers in the next highest rank at the same plant-year. Peer share is for workers at the same rank in the same plant-year.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

Table 6. Robustness Checks for Gender Spillovers in Promotions

	(1)	(2)	(3)	(4)	(5)	(6)
	Promotion					
Female * Female Boss Share	0.016*** [0.0047]	0.010** [0.0053]	0.014*** [0.0047]	0.025** [0.0102]	0.020*** [0.0045]	0.020*** [0.0046]
Female * Female Peer Share	-0.056*** [0.0057]	-0.048*** [0.0065]	-0.061*** [0.0063]	-0.079*** [0.0094]	-0.068*** [0.0055]	-0.060*** [0.0057]
Female Boss Share	-0.0056 [0.0048]	-0.0026 [0.0049]	-0.0188*** [0.0048]	-0.0103 [0.0109]	0.00094 [0.0048]	-0.0054 [0.0048]
Female Peer Share	0.026*** [0.0058]	0.029*** [0.0063]	-0.026*** [0.0078]	0.047*** [0.0106]	0.055*** [0.0058]	0.038*** [0.0059]
Size of Ranks (current and next higher)	Yes	No	No	No	No	No
Rank-Sex FE	No	Yes	No	No	No	No
Rank-Plant FE	No	No	Yes	No	No	No
Occupation Group Share in Next Rank	No	No	No	Yes	No	No
Plant-Year FE	No	No	No	No	Yes	No
Rank-Year FE	No	No	No	No	No	Yes
N	597,463	597,552	597,552	318,760	597,552	597,552
R ²	0.074	0.073	0.083	0.075	0.15	0.075

Notes: OLS estimates with standard errors clustered at the plant-rank-year level. Basic controls in all columns are as in Table 5, Column (2).

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

Table 7. Gender Spillovers in Promotions with Alternative Definitions of Female Boss Share

	(1)	(2)	(3)	(4)	(5)
	Promotion				
Female * Female Boss (+2) Share	0.041*** [0.0079]	0.033*** [0.0082]			
Female * Female Boss (+1) Share		0.021*** [0.0064]			
Female * Female Boss (Top 3) Share			0.051*** [0.010]		
Female * Female Boss (Top Earners) Share				0.069*** [0.0071]	
Female * Female Boss (Leaders) Share					0.021** [0.010]
Female * Female Peer Share	-0.073*** [0.0071]	-0.076*** [0.0073]	-0.061*** [0.0064]	-0.064*** [0.0057]	-0.085*** [0.0085]
N	453,924	453,924	471,561	583,859	325,320
R ²	0.078	0.078	0.074	0.07	0.07

Notes: OLS estimates with standard errors clustered at the plant-rank-year level. Controls are as in Column (3) of Table 3. Female Boss (+2) is the share 2 ranks higher; Female Boss (Top 3) is the share in ranks 5-7; Top earners are in the top fifth for their establishment-year; Female Boss (Leader) applies to higher ranks with leadership tasks.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

Table 8: Predicted Effects of Increasing Female Peer and Boss Shares on Women's Relative Promotion Rates

	(1)	(2)	(3)	(4)	(5)	(6)
	Peers	Bosses	Bosses	Bosses	Bosses	Bosses
Female share measure:	Same Rank	Rank +1	Rank +2	Rank 5-7	Top Earn	Leader
Estimated coefficient for female share	-0.060	0.019	0.041	0.051	0.069	0.021
Standard deviation across individuals	0.290	0.240	0.170	0.075	0.090	0.064
Predicted change in the gender promotion gap from a standard deviation increase in female representation	52.7%	-13.8%	-21.1%	-11.6%	-18.8%	-4.1%

Notes: The coefficient estimates for the female share in Columns 1 and 2 are from Table 5, Column 2. The remaining coefficients are from Table 7. Column 3 here is from Column 1; Column 4 is from Column 3; Column 5 is from Column 4 and Column 6 is from Column 5.

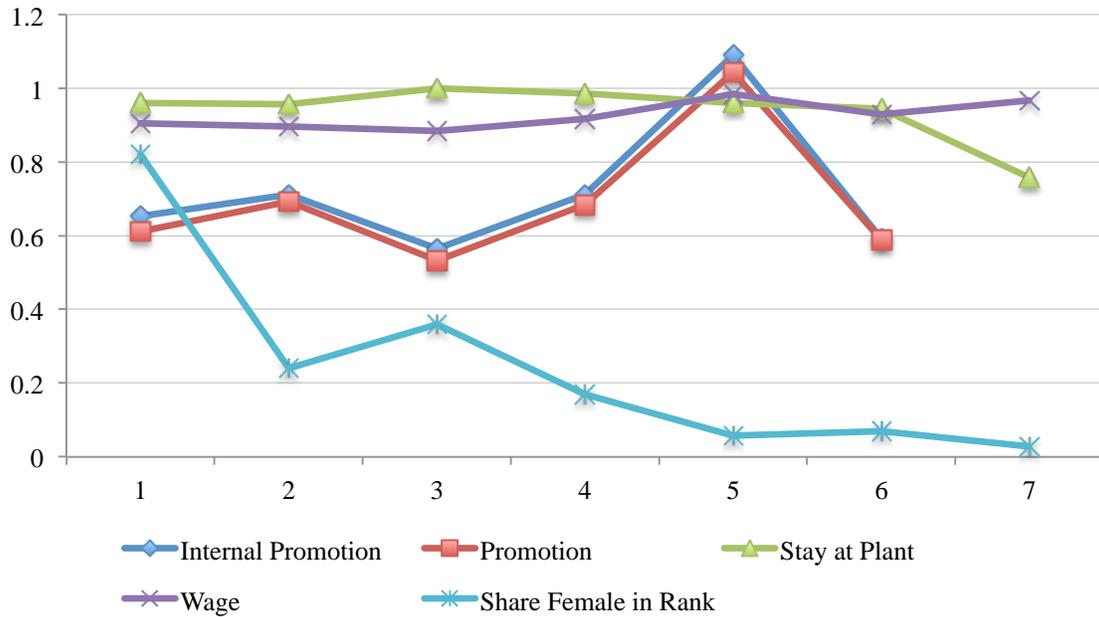


Figure 1: Sex Differences by Rank. The lines for internal promotion, promotion, stay at plant, and wage each plot the female to male ratio of the mean values of the outcome within each rank. The line for female share in rank plots the female share of workers in that rank. The female share of the entire sample is 28 percent.

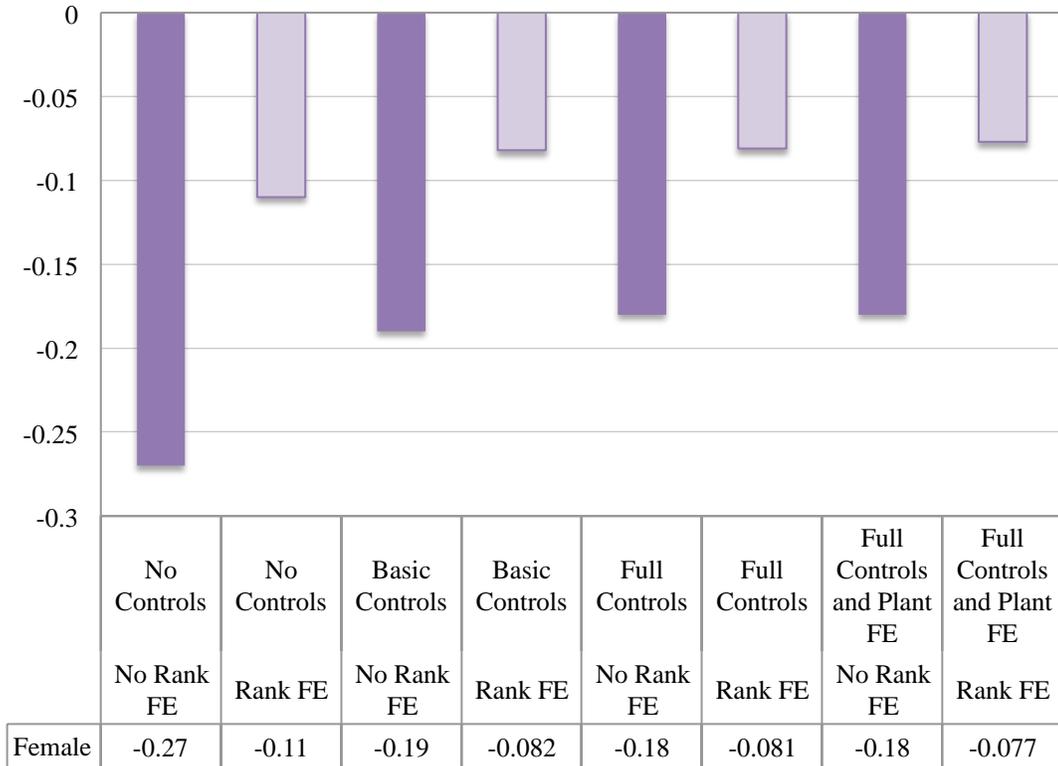


Figure 2: Gender Pay Gap Estimates with and without Rank Fixed Effects. Each bar is a point estimate for the *Female* coefficient from an OLS regression of the log of wage rate (computed by dividing average monthly earnings by normal hours) on *Female* without (dark bars) and with (light bars) rank fixed effects. The first pair of lines has no additional controls. The second pair includes fixed effects for industry, occupation, year, age and schooling. The third pair also includes controls for tenure (overall and rank-specific), experience and part-time status. The last pair has plant fixed effects. All reported coefficients are significant at the 1 percent level.