



MIT Sloan School of Management

MIT Sloan School Working Paper 4779-10

From Metaphors to Mechanisms: Gender Sorting In(to) an Organizational Hierarchy

Roberto M. Fernandez and Mabel Botelho Abraham

© Roberto M. Fernandez and Mabel Botelho Abraham

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.

This paper also can be downloaded without charge from the
Social Science Research Network Electronic Paper Collection:
<http://ssrn.com/abstract=1589012>

Electronic copy available at: <http://ssrn.com/abstract=1589012>

From Metaphors to Mechanisms: Gender Sorting In(to) an Organizational Hierarchy

Roberto M. Fernandez

Mabel Botelho Abraham

MIT Sloan School of Management

April, 2010

From Metaphors to Mechanisms: Gender Sorting Into an Organizational Hierarchy

Abstract

Numerous studies have examined patterns of gender inequality in organizational advancement, with some showing results indicative of “glass ceilings,” where gender disparities are strong at the upper reaches of the organization, while others suggest “sticky floors,” where the gender differences in advancement occur at the lower levels of the organization. These studies, however, have been less clear on the mechanisms that produce these descriptive patterns. With few exceptions, extant studies have focused on internal promotion practices and have not addressed the extent to which these patterns might reflect gender differences in *external* recruitment and hiring into the organization. We construct queues consisting of both external and internal candidates—the set of candidates under consideration—for over 2,200 job openings during a 27-month period for a large retail bank. We find that women are more likely than men to be hired, and that this pattern holds for jobs up and down the organizational hierarchy. The applicant pools are themselves gendered, however, with women comprising a lower percentage of the applicant pools for high-level jobs, but a greater proportion of the pools for lower-level jobs. Since women are more likely than men to advance from each applicant pool to hire at all levels of the hierarchy, the apparent “glass ceiling” observed among job incumbents is not due to gendered screening practices. Instead, the roots of gender inequality in this firm lie in the initial sorting of applicants into queues.

Keywords: “Glass ceiling,” labor queues, job sex segregation, gender inequality, wage inequality, stratification, hiring processes, employer screening

There is broad interest in the phenomenon of gender inequality in labor market outcomes both in sociology (Morgan 1998; Huffman and Cohen 2004; Barnett et al. 2000) and economics (Albrecht et al. 2003; Dohmen et al. 2004). Numerous studies have adopted the metaphor of a “glass ceiling” when examining whether gender inequality in labor market outcomes is more severe at the top of the reward distribution (Cotter et al. 2001; Morgan 1998; Arulampalam et al. 2007). Most recently, Hassink and Russo (2010) have proposed the idea of a “glass door,” where females are particularly disadvantaged in external recruitment to high-level jobs. These studies have shown mixed results, with some reporting results consistent with the “glass ceiling” notion (Powell and Butterfield, 1994; Yamagata et al. 1997; Cotter et al., 2001), while others have found evidence that suggests “sticky floors,” where women’s disadvantages are most pronounced at lower levels of the reward hierarchy (Booth et al. 2003; Yap and Konrad 2009).

While research reporting general patterns of gender inequality in rewards are relatively plentiful, these studies have been less clear on the mechanisms that produce these descriptive patterns. Numerous scholars have responded to Reskin’s (2003) call to address the mechanisms that produce disparate labor market outcomes. In search of these mechanisms, some scholars have returned to the original conceptualization of the “glass ceiling,” which invokes the idea that hidden organizational barriers are preventing women’s advancement to the top of the corporation (Hymowitz and Schellhardt 1986; Morrison et al. 1987). Several recent studies focus on the organizational nature of barriers to women’s achievement, specifically using the “glass ceiling” idea when studying gender differences in internal advancement. For example, Gorman and Kmec (2009) frame their discussion of female-male differences in promotion to partner in elite law firms as indicative of there being “glass ceilings” in these firms. Using data on a large service organization, Petersen and Saporta (2004) look at female-male differences across a broad range of organizational outcomes, including initial job level, turnover, and wages. The section of the paper devoted to examining sex differences in promotion is titled “The Glass Ceiling” (pp. 887-893). Other organizational studies, too, focus on internal advancement when using the “glass ceiling” idea (Powell and Butterfield 1994; Kalev 2009; Yap and Konrad 2009).

While we applaud the search for specifically organizational mechanisms in this important strand of research, we think there are significant and, to date, unrecognized limitations of this internal focus for research on gender inequality. Research on internal advancement has made the implicit assumption that jobs at higher levels of

the organizational hierarchy are closed to the external labor market. Depending on the setting, this assumption might be warranted. But, to the degree that organizations depart from the Doeringer and Piore (1971) model of the internal labor market, and higher-level jobs are open to competition from the external labor market, then within-firm studies of promotion will obscure this part of the process. Indeed, numerous firm-level studies have questioned whether the internal labor market metaphor is an accurate depiction of the firm (Dohmen et al. 2004; Seltzer and Merrett 2000; Treble et al. 2001). For example, while it might be the case for elite law firms that external hiring tends to be limited to the bottom of the organization (Gorman and Kmec 2009), to the extent that law firms engage in lateral hiring, then studies of internal promotion will misspecify the population at risk of obtaining higher level jobs. If such external market competition is itself gendered, then internal market studies will not reflect the actual mechanisms producing gender stratification, and may yield misleading inferences on the nature of the organizational barriers to women's advancement.

In this paper, we seek to address this limitation of the extant literature. We study the extent to which gender inequality exists across job levels within the organizational hierarchy. In contrast to previous research, we properly specify the candidate pool—both internal and external—at risk of being hired across levels of the organization. We describe where in the organizational hierarchy gender imbalances in job composition occur. Consistent with past research, we find that the percentage of women declines with increasing levels of the organizational hierarchy. Virtually unique among studies in this area (an exception is Fernandez and Mors 2008), we construct queues¹ of both external and internal applicants—the set of candidates under consideration—for over 2,200 job openings during a 27-month period for a large retail bank. We find that women are more likely than men to be hired, and that this pattern holds for jobs up and down the organizational hierarchy. The applicant pools are themselves gendered, however, with women comprising a lower percentage of the applicant pools for high-level jobs, but a greater proportion of the pools for lower-level jobs. Since women are more likely than men to advance from each applicant pool to hire at all levels of the hierarchy, the apparent “glass ceiling” observed among job incumbents is not due to gendered screening practices. Instead, the roots of gender inequality in this firm lie in the initial sorting of applicants into queues.

Data and Setting

¹ For an important paper looking at gender and queues for temporary contract workers, Fernandez-Mateo and King (2010).

We study the process by which jobs are filled in the western U.S. region of a large retail bank. The bank is part of a globally diversified financial services institution. As of December, 1995, the western region bank employed a total of 3,641 workers, 64.5 percent of whom were female. Using the company's internal records, we analyzed data on the hiring process for the western region bank's job openings during the 27-month period from January, 1993 through March, 1995.²

A total of 2,284 job openings were filled over this period. Most important for the purposes of this research, the company's HR department kept databases of job openings (job requisitions submitted by hiring managers) and of applicants to these job openings. By combining these sources, we identified a total of 31,656 applicants to the firm. The information kept on these applicants is sparse, however. Virtually complete information is available to distinguish internal and external applicants (31,627 or 99.9 percent of the 31,656 applications), and whether the applicant was referred to the firm by an employee (99.5 percent). Also important given our purposes is that we were able to identify the gender of the applicant for 93.8 percent of the applications (29,697 of 31,656). The company directly recorded the gender of 2,091 of the 2,098 hires, and we consulted paper employment records to identify the gender of an additional 2,426 applicants as a part of another study (Fernandez and Weinberg 1997). We then hand coded another 25,810 applications based on the applicant's name.

We formed queues of applicants to each of 2,248 job openings across the organizational hierarchy. A total of 13,091 applications were considered for these job openings. These applicants were 54 percent male, and 46 percent female. Of the remaining applicants, 3,757 (65.0 percent male and 35.0 percent female) were in queues into which no one was ultimately hired, either because the job opening was withdrawn or because the job was still open at the end of the hiring window. There were an additional 14,808 applications to the firm, however, which were not matched against any job opening at all. Similar to the organization studied by Fernandez and Mors (2008), the retail bank's database for job requisitions was also used as an internal job posting system. Consequently, internal applicants apply directly and place themselves in queues for particular job openings; only 0.7 percent of these unmatched applications were internal in origin.

In contrast, external applicants are funneled through the company's HR department. Some external applicants asked to be considered for specific jobs, while other applicants applied for employment at the company in

² Ideally, we would have liked to observe the firm's gender composition at the beginning and the end of the hiring window (see Fernandez and Sosa [2005] for an example of a study using this design). However, the earliest time that the records for the overall company are available is 8 months later, December, 1995.

general. The standard operating procedure is for HR recruiters to honor applicants' wishes by matching people to the applicant pool for the specific job opening they requested. While HR successfully matched 51.0 percent (15,318) of external applicants to job openings, the other 49.0 percent (14,699) of externals were not matched against any job opening. These unmatched externals are 48 percent female, slightly greater than the percent female for applicants matched to queues (46.0 percent), but substantially less than the female composition of the overall firm (64.5 percent). Since none of these applications result in a hire, they fall out of the risk set for hiring. Consequently, we set these cases aside for this study.

While we do not think that this firm's hiring practices are atypical of other large firms (cf. Kalleberg et al. 1996), given our decision to study this one setting, we make no claims regarding generalizability. Our main goal in adopting this empirically grounded, case-study approach is to elucidate the workings of hitherto unaddressed mechanisms that affect the sex composition of jobs at different level of the organizational hierarchy. While we would expect that there will be some contingency in the ways these processes work in different settings, it is impossible to distinguish among the mechanisms at play without the unique, fine-grained data we analyze here. To our knowledge, this is only the second empirical study³ which measures the gender composition of the applicants at risk of being hired for specific job openings, for both internal and external candidates. The theoretical significance of this case is that it provides a window through which one can view the operations of a set of processes that are normally hidden from view. Thus, the insights gleaned from this case study can be used to guide broader-gauge research designed to represent wider populations of organizations.

Analysis

Table 1 shows the average annual salary distribution for job incumbents by level of the organization as of December 31, 1995.⁴ The organization has 21 levels, with levels 13 – 21 being salaried jobs and exempt from U.S. laws governing overtime, and levels 1 -12 being hourly jobs subject to overtime regulation. The first three columns show the mean, minimum and maximum average annual salary for people working at each level. With the exception of levels 10 and 11, the average annual salary increases directly with level. The aggregate correlation between level

³ The only other study of which we are aware is Fernandez and Mors (2008).

⁴ As we mentioned above, while not ideal (see fn. 2), this is the earliest point for which we have data available to offer a snapshot of the entire western region bank. In addition, because of privacy concerns, these data were supplied to us in an anonymized form, so it is not possible to match these individuals to those in the jobs and applicant databases. Despite these limitations, these data form a valuable baseline for subsequent analyses of the hiring process.

(measured as 1 – 21) and average wage is 0.79. Also noteworthy, however, is the fact that there is substantial overlap of the salary ranges across levels (see the second and third columns).

Most importantly, the fourth column shows that females are heavily concentrated in the lower ranks of the firm, and that with few exceptions the percentage female declines as one looks up the hierarchy. This pattern of differential sorting has important consequences for wage inequality. First, the correlation across the 21 levels between percent female and average annual salary is -0.84.⁵ Second, this firm exhibits the familiar pattern of wage inequality and sex segregation. Ignoring job level, women overall earn 58 percent of what men earn: the average annual salary for females is \$33,650 compared with \$57,572 for males. Consistent with many previous studies (e.g. Petersen and Morgan, 1995) the wage gap virtually disappears when women and men in the same level are compared: within level, women on average earn 97 percent of what men earn. Thus, a key component of wage inequality can be traced to allocative processes that sort men and women into different levels of the hierarchy (Petersen and Saporta 2004). It is precisely this process that we seek to shed light on with this study.

Looking more closely at the percent female column, we see two places where the fall-off is most precipitous. First, among non-exempt workers the drop from level 10 to levels 11 and 12 (i.e., 66.7 to 21.4, and 25.0 percent) is fairly dramatic. As the number of cases column shows, levels 10 – 12 are sparsely populated. HR personnel told us that these levels are being phased out, but have historically been used to reward employees with long tenure in the company. The data on tenure with the firm is consistent with this account. For both men and women, the incumbents of levels 10 -12 are those with the longest tenure among non-exempt employees in the firm.⁶ Moreover, very few jobs at this level were filled during the hiring window, either internally or externally (see the last column of Table 1).⁷

For the exempt jobs, the break in percent female occurs at the very top of the organization, i.e., between level 18 and above (percent female at level 18 = 25.0, 19 = 4.5, 20 = 0.0 and 21 = 0.0). The fact that there are no women at all in the top two levels of the organization might be interpreted as evidence that the firm has a “glass

⁵ The individual-level correlation between female (coded 1) and annual salary is -0.376.

⁶ Specifically, there is a sharp jump between the average tenure for level 9 (7.1 and 7.2 years for females and males) and the average years of tenure of those in levels 10, 11 and 12 (respectively, female = 12.0, 10.7, and 15.8; males = 11.2, 12.5, and 15.5).

⁷ The data on hiring that we analyze in this paper is based on job openings that are filled over the course of the study window. Thus, these data reflect both new people being hired into the organization (externals) as well as internal transfers and promotions. While it would be desirable to distinguish the latter two categories, the HR system on which we base our analyses tracked the filling of open slots, but did not keep information on which job internals had moved from. Identifying information which would distinguish transfers from promotions would require linking to the larger HR database. Because the larger HR data has been anonymized (see fn. 4), it is not possible for us to link these data sources.

ceiling” – a hidden barrier that women cannot penetrate--at level 19. It is important to recognize that such an inference would be based on very few cases, however, since there is only one male incumbent in each of the top two salary bands (for a discussion of a similar pattern, see Petersen and Saporta 2004:889-92). Having such a small number of cases in the top levels of the organization is challenging in another way. The fact is that no job openings at all were filled by either internal or external hires in levels 19 and above (last column of Table 1). Thus we cannot offer any insight into how jobs at these levels are filled.

However, the data in the last column show that we do have a sufficient number of hires to study how gender sorting works for exempt jobs at level 17 and below. In order to preserve the 8 hires at level 18, we will combine levels 17 and 18 for subsequent analyses. Following a similar logic, we also combine levels 1 – 3 and levels 9 – 12 and use the resulting 12 level variable to categorize job levels for the hiring analyses to follow.

Table A1 in the Appendix shows distributions of average wages and sex composition using the recoded levels. Whereas the correlation between sex composition and job level was -0.84 across the 21 categories, the corresponding correlation for the 12 categories is -0.86. The average within-job level wage gap remains unchanged: for both the 21 and 12 category classifications, when women and men are in the same category, women on average earn 97 percent of what men earn. By these standards at least, the simplified recoding does a good job of preserving the relationship between job-level average annual salary and percent female.

A key component of our argument is that it is important to properly specify the candidate pool—both internal and external—before making inferences about gender differences the risk of being hired into jobs across levels of the organization. As we argued above, to the degree that the organization departs from the internal labor market model—i.e., jobs being filled by externals predominantly through “ports of entry” at the bottom of the organization, and higher level openings filled through internal promotions—then the focus on promotions as the gender sorting mechanism responsible for gender inequality by organizational rank can be misleading. Several firm-level studies have appeared which question the internal labor market model, specifically arguing that there is little evidence for “ports of entry” at the bottom of the organization (Dohmen et al. 2004; Seltzer and Merrett 2000; Treble et al. 2001).

Table 2 shows the percent external among hires and candidates filled by level of the organization. Overall, 60.1 percent of jobs are filled by external hires, but this varies by level. Broadly consistent with the internal labor market model, the highest percentage of external hires is found in the lowest rung of the organization (75.2), and the

lowest percentage of external hires is found in the top tier (37.1). However, we would be hard pressed to interpret this pattern as evidence for there being distinct “ports of entry” to the firm: a non-trivial percentage of the job openings at each level are filled by externals. In addition, the second column, which lists the external/internal composition of the candidate pools, shows that a substantial percentage of the pools at all job levels are composed of externals. Apparently, externals are standing in line for jobs at all levels of the organization. These findings suggest that an external focus is needed in order to correctly specify the process by which people are sorted into job openings at different levels of the organizational hierarchy.

Hiring Analyses

The first column of Table 3 shows the percent female for job incumbents by the recoded level of the organization. The major feature of Table 1 is in evidence here as well: percent female tends to decline as job level increases (correlation = -0.86). Using the 1 -12 coding of level, the average level for male job incumbents is 8.6, compared with 6.1 for female job incumbents ($p < .0001$; F-test = 491.64, with 1 and 3,639 d.f.).

The second column shows the data on the sex composition of job openings that have been filled during our study window by level.⁸ As noted above (see fn. 7) these jobs have been filled by either internal or external applicants. (We will analyze the two sources separately below). The sex composition of hires and job incumbents track closely: they are correlated 0.88. On average, males are being hired into jobs that are 1.9 levels higher than the levels into which females are being hired. Using the 1 -12 coding of level, the average job level that males are filling is 7.6 compared with 5.7 for females ($p < .0001$; F-test = 156.198, with 1 and 2,239 d.f.).

The fact that the sex composition of hires by level lines up closely with that of job incumbents might lead some to conclude that females are suffering from a disadvantage in allocation to jobs (for a similar pattern, see Hassink and Russo 2010). This conclusion, however, would be premature. In considering questions of allocation, both the first and second columns of Table 2 present data that have been selected on the dependent variable (see Fernandez and Weinberg [1997] for a discussion of the limitations of what they call “start-with-hire” studies). That is, the data in the first and second columns look only at the survivors of the allocative process, and do not take into account the sex composition of the pools of candidates from which these hires have been selected.

⁸ Our data on job openings are based on job requisitions that have been submitted by hiring managers. These requisitions, however, are sometimes written to span levels (e.g., “9-11”). The job levels we report in this paper use the top of the range, in the example of levels “9-11”, we code this as level “11.” None of the results change if we use the bottom of the range.

A major contribution of this study, therefore, lies in the data reported in columns 4 – 6 of Table 3. The third column shows the sex composition of the candidate pool by level of the organization’s hierarchy. Several patterns are noteworthy in these data. First, perhaps not surprisingly given that the reward structure, the higher level jobs attract more candidates than the lower level jobs (cf. the number of candidates for the exempt and non-exempt jobs in the column labeled “Candidates”). The fourth column shows that the degree of competition for these jobs increases dramatically with the level of the job. Hiring rates decline from 43.8 and 47.5 percent for the lowest two job levels, to 5.8 and 5.5 percent for the top two job levels.

Also of interest is the relationship between the job level and the sex composition of candidates. Figure 1 graphs the first three columns of Table 3, i.e., the percent female for job incumbents, hires and candidates for each job level. Overall, the line for the candidates closely mirrors the pattern for the hires (correlation = 0.90) and job incumbents (correlation = 0.89): female candidates are disproportionately concentrated in queues for jobs in the lower ranks of the organization, while male candidates are more prevalent in the candidate pools for the upper level jobs.

While all three lines follow the same broad pattern, it is critical to note that the line for candidates is *always* below that for hires. This means that compared to their representation in the candidate pool, females are *overrepresented* among hires. This is true for the firm overall, where women constitute 65.6 percent of hires, but are only 45.9 percent of the total candidate pool. Most importantly, this pattern repeats itself at every level of the organization: the percentage of females hired is *greater* than the percentage of females in the candidate pools for jobs up and down the organization’s hierarchy. This latter pattern implies that hire rates for females are *greater* than hire rates for males, irrespective of organizational rank (see Figure 2 and the last two columns of Table 3).

Table 4 elaborates further on this finding investigating the between- and within- job level patterns of sex differences in getting hired. In order to address the question of interactions between sex composition and internal vs. external on chances of being hired, the analyses are presented both combined for internals and externals (Models 1 – 3), and separately for externals (Models 4 – 6) and internals (Models 7 – 9). For each population, we exclude cases that are missing on any variable included in subsequent models in order to ensure that the results within each population are based on the same cases.⁹

⁹ The substantive results do not change when the models are re-estimated including all the non-missing cases for each of the separate models.

Model 1 presents the coefficients of a logit regression model with sex predicting hires, combining external and internal candidates. The coefficient for sex in Model 1 shows a strong, statistically significant effect on the odds of getting hired: overall, the odds of females getting hired are 2.673 times the odds of males getting hired.

Model 1, however, does not distinguish within- and between- job-level variation in gender composition and hire. However, jobs at different levels differ in many ways, and these factors—whether observed or unobserved—can also affect hiring managers' choices. In order to make proper inferences about the sources of variation in outcome, it is important to separate the within- and between-job level variation in the gender of the hire. Toward this end, we estimate a logit model predicting hire with fixed effects for job-levels (Model 2). This estimation strategy, therefore, purges all between-job level factors—both observed and unobserved—associated with job levels. The sex coefficient in Model 2 is attenuated compared to Model 1 (1.964 vs. 2.673), indicating that some part of the observed preference for females in hiring is due to the kinds of jobs for which they are being considered. Nevertheless, even after controlling level job-level factors, females still show substantially higher odds of getting hired than do males.

The analyses to this point do not address the possibility that females' odds of being hired might differ for high and low level jobs. Model 3 adds an interaction term between sex and the average annual salary of the job level in the logit model with job-level fixed effects. (Because the average annual salary measure does not vary within job level, the fixed effects absorb the main effect of salary level, which drops out of the analysis). The interaction term, however, picks up within-job level variation in each sex. In Model 3, therefore, the main effects measure the sex difference for the jobs coded zero on salary (i.e., the lowest level jobs), and the interaction term measures the extent to which sex differences in the hire outcome change as average job-level salary increases.¹⁰

The main effect of sex in Model 3 shows that females are twice as likely as males to be hired in the lowest job level (odds ratio = 2.036). The interaction term, while negative is very close to 1 and is not statistically reliable ($p > .541$). The .998 interaction effect implies that the observed preference for females within job level erodes by 0.2 percent per \$1,000 increase in average job-level annual salary. By this metric, the observed female advantage is reduced to 1.690 at the very top of the organization (job level with average annual salary = \$109,000). Thus, we

¹⁰ In order to facilitate the interpretation of the interaction term, we rescaled job-level annual salary by subtracting the lowest observed job-level wage (\$16.03 in thousands of dollars) from the average salary for each job level. Thus, the rescaled salary variable varies from 0 to \$93.43.

have no statistically reliable evidence of a decline in the female advantage as one looks up the organizational hierarchy.

The second and third panels of Table 4 replicate the analysis for externals (Models 4 – 6) and internals (Models 7 – 9). For internals, the story is simple: we have no statistically reliable evidence of sex differences in hiring outcomes neither overall (Model 7) nor within job-levels (Model 8), and these findings hold irrespective of the job level considered (Model 9). Thus, internal allocative processes of promotion and transfer are not significant contributors to the sorting of men and women across the organizational hierarchy. In contrast, external candidates show quite strong patterns of sex differences in allocation to jobs. Females are advantaged both overall (Model 4) and within job levels (Model 5). However, Model 6 shows no statistically significant evidence of the female advantage eroding across the organizational hierarchy. These patterns strongly suggest that the pooled results (reported in Models 1 – 3 and Figures 1 and 2) are driven by the evident female advantage in hiring among external candidates.

Conclusion

The “glass ceiling” has been embraced as a useful metaphor guiding much research on gender inequality. A number of scholars have moved beyond the descriptive use of the idea, to study the mechanisms producing barriers to women’s advancement in organizational hierarchies. In this regard, a particular focus of research has quite reasonably been the study of gender sorting in internal promotions. The patterns we have documented here, however, suggest that such a narrowing of focus can obscure important processes by which males and females are sorted into jobs. We have argued that, to the degree that firms are open to external recruitment throughout the hierarchy, a focus on internal processes will misspecify the ways in which men and women come to be sorted into jobs of different levels and rewards.

Although data limitations prevent us from offering much insight into gender inequality at the very top of the firm, the balance of the organization that we study here bears several of the well known earmarks of gender inequality. Over all, women’s salaries are 58 percent of men’s salaries. In addition, the percentage of females declines as one looks up the organizational hierarchy. Moreover, there is also sex inequality in initial placement with the firm since male new hires are placed into higher level jobs than are women. Taken in isolation, these facts would seem to suggest that women likely face barriers to their achievement in this organization.

While gender inequality is clearly evident in this firm, the mechanism which accounts for this pattern is not the usual suspect of internal promotion barriers. Although we cannot separate the promotion from transfers in these data, we find no evidence of women being disadvantaged by internal processes. Instead, we find that external recruitment plays an important role in our analyses. We see that jobs across the spectrum are often being filled by external candidates. Contrary to Hassink and Russo (2010), we find no evidence of “glass doors” where women are being excluded in the allocation process from the external market. Although we do not have any controls for the quality of the candidates, the fact is that women are *more* likely than men to be hired from the external labor market, and this pattern holds for jobs up and down the organization’s hierarchy. Virtually unique among studies in this area, we have shown that the candidate pools from which managers are choosing hires are gendered, with women comprising a lower percentage of the applicant pools for high-level jobs, but a greater proportion of the pools for lower-level jobs. Since women are more likely than men to advance from each applicant pool to hire at all levels of the hierarchy, the apparent “glass ceiling” observed among job incumbents cannot be attributed to hiring manager’s gendered screening practices. Instead, we must look to factors that lead men and women to sort into different candidate pools. At least in this firm, the roots of gender inequality lie in the initial sorting of applicants to queues.

References

- Albrecht, James, Anders Björklund, and Susan Vroman. 2003. “Is There a Glass Ceiling in Sweden?.” *Journal of Labor Economics* 21:145-177.
- Arulampalam, Wiji, Alison L. Booth, and Mark L. Bryan. 2007. “Is There a Glass Ceiling over Europe? Exploring the Gender Pay Gap across the Wage Distribution.” *Industrial and Labor Relations Review* 60:163-186.
- Barnett, William P., James N. Baron, and Toby E. Stuart. 2000. “Avenues of Attainment: Occupational Demography and Organizational Careers in the California Civil Service.” *The American Journal of Sociology* 106:88-144.
- Booth, Alison L., Marco Francesconi, and Jeff Frank. 2003. “A sticky floors model of promotion, pay, and gender.” *European Economic Review* 47:295-322.
- Cotter, David A., Joan M. Hermsen, Seth Ovardia, and Reeve Vanneman. 2001. “The Glass Ceiling Effect.” *Social Forces* 80:655-681.
- Doeringer, P. B., M. J. Piore, and others. 1971. *Internal labor markets and manpower analysis*. Heath Lexington Books.
- Dohmen, Thomas J., Ben Kriechel, and Gerard A. Pfann. 2004. “Monkey Bars and Ladders: The Importance of Lateral and Vertical Job Mobility in Internal Labor Market Careers.” *Journal of Population Economics* 17:193-228.
- Fernandez, Roberto M., and Marie Louise Mors. 2008. “Competing for jobs: Labor queues and gender sorting in the hiring process.” *Social Science Research* 37:1061-1080.
- Fernandez, Roberto M., and M. Lourdes Sosa. 2005. “Gendering the Job: Networks and Recruitment at a Call Center.” *The American Journal of Sociology* 111:859-904.
- Fernandez, Roberto M., and Nancy Weinberg. 1997. “Sifting and Sorting: Personal Contacts and Hiring in a Retail Bank.” *American Sociological Review* 62:883-902.

- Fernandez-Mateo, Isabel. 2010. "Gender Sorting and Competition for Projects in the Staffing Sector." Unpublished paper.
- Gorman, Elizabeth H., and Julie A. Kmec. 2009. "Hierarchical Rank and Women's Organizational Mobility: Glass Ceilings in Corporate Law Firms." *American Journal of Sociology* 114:1428-1474.
- Hassink, Wolter H.J., and Giovanni Russo. 2010. "The Glass Door: The Gender Composition of Newly-Hired Workers Across Hierarchical Job Levels." IZA Discussion paper No. 4858.
- Huffman, ML, and PN Cohen. 2004. "Occupational segregation and the gender gap in workplace authority: National versus local labor markets." *Sociological Forum* 19:121-147.
- Hymowitz, Carol, and Timothy Schellhardt. 1986. "The glass ceiling: Why women can't seem to break the invisible barrier that blocks them from the top jobs." *The Wall Street Journal*.
- Kalleberg, Arne L., David Knoke, Dr. Peter V. Marsden, and Joe L. Spaeth. 1996. *Organizations in America: Analysing Their Structures and Human Resource Practices*. Sage Publications, Inc.
- Kalev, Alexandra. 2009. "Cracking the Glass Cages? Restructuring and Ascriptive Inequality at Work." *American Journal of Sociology* 114:1591-1643.
- Morgan, Laurie A. 1998. "Glass-Ceiling Effect or Cohort Effect? A Longitudinal Study of the Gender Earnings Gap for Engineers, 1982 to 1989." *American Sociological Review* 63:479-493.
- Morrison, Ann M., Randall P. White, Ellen Van Velsor, and Center for Creative Leadership. 1987. *Breaking the Glass Ceiling: Can Women Reach the Top of America's Largest Corporations?* Addison Wesley Longman Publishing Co.
- Petersen, Trond, and Laurie A. Morgan. 1995. "Separate and Unequal: Occupation-Establishment Sex Segregation and the Gender Wage Gap." *The American Journal of Sociology* 101:329-365.
- Petersen, Trond, and Ishak Saporta. 2004. "The Opportunity Structure for Discrimination." *The American Journal of Sociology* 109:852-901.
- Powell, Gary N., and D. Anthony Butterfield. 1994. "Investigating the "Glass Ceiling" Phenomenon: An Empirical Study of Actual Promotions to Top Management." *The Academy of Management Journal* 37:68-86.
- Reskin, BF. 2003. "Including mechanisms in our models of ascriptive inequality." *American Sociological Review* 68:1-21.
- Seltzer, Andrew, and David Merrett. 2000. "Personnel Policies at the Union Bank of Australia: Evidence from the 1888-1900 Entry Cohorts." *Journal of Labor Economics* 18:573-613.
- Treble, John, Edwin van Gameren, Sarah Bridges, and Tim Barmby. 2001. "The internal economics of the firm: further evidence from personnel data." *Labour Economics* 8:531-552.
- Yamagata, Hisashi, Kuang S. Yeh, Shelby Stewman, and Hiroko Dodge. 1997. "Sex Segregation and Glass Ceilings: A Comparative Statics Model of Women's Career Opportunities in the Federal Government Over a Quarter Century." *The American Journal of Sociology* 103:566-632.
- Yap, M., and A. Konrad. 2009. "Gender and Racial Differentials in Promotions: Is There a Sticky Floor, a Mid-Level Bottleneck, or a Glass Ceiling?." *Relations Industrielles* 64:593-619.

Table 1. Annual Salary and Percent Female for Job Incumbents by Level of the Organization (As of December 31, 1995)

Level	<u>Mean</u>	<u>Minimum</u>	<u>Maximum</u>	<u>Percent Female</u>	<u>Number of Cases</u>	<u>Jobs Filled in Study Window*</u>
<u>Non-Exempt:</u>						
1	\$13,873	\$12,480	\$29,567	82.6	23	15
2	\$15,573	\$12,954	\$23,088	75.4	69	176
3	\$17,061	\$13,992	\$24,338	72.2	79	15
4	\$18,868	\$13,744	\$29,121	70.8	219	198
5	\$20,705	\$13,631	\$28,621	79.5	185	238
6	\$21,955	\$17,138	\$32,320	86.0	707	217
7	\$26,157	\$17,848	\$38,080	80.5	292	148
8	\$30,764	\$12,514	\$44,418	83.7	208	155
9	\$32,944	\$22,033	\$46,720	68.0	125	45
10	\$43,135	\$33,967	\$50,000	66.7	18	2
11	\$43,086	\$33,650	\$50,600	21.4	14	1
12	\$53,313	\$47,500	\$55,500	25.0	8	0
<u>Exempt:</u>						
13	\$30,283	\$24,000	\$43,200	58.0	50	49
14	\$36,446	\$25,272	\$80,000	62.0	395	460
15	\$47,682	\$31,500	\$130,000	54.3	440	320
16	\$65,912	\$40,100	\$96,000	38.8	418	147
17	\$90,762	\$58,550	\$160,000	37.6	255	54
18	\$134,757	\$83,346	\$176,885	25.0	112	8
19	\$182,831	\$131,410	\$240,000	4.5	22	0
20	\$240,000	\$240,000	\$240,000	0.0	1	0
21	\$299,657	\$299,657	\$299,657	0.0	1	0
Total	\$42,139	\$12,480	\$299,657	64.5	3641	2248

Table 2. Percent External by Recoded Level of the Organization (Number of cases in parentheses)

Level	<u>Hires</u>	<u>Candidates</u>
<u>Non-Exempt:</u>		
Levels 1 -3	.752 (155)	.892 (428)
4	.631 (125)	.823 (354)
5	.748 (178)	.894 (785)
6	.585 (127)	.809 (496)
7	.708 (104)	.873 (604)
8	.429 (66)	.783 (656)
Levels 9 -12	.458 (22)	.824 (300)
<u>Exempt:</u>		
13	.531 (26)	.896 (398)
14	.623 (286)	.898 (2138)
15	.600 (192)	.909 (1907)
16	.442 (65)	.933 (2470)
17 +	.371 (23)	.927 (1087)
Total	.610 (1369)	.891 (11623)

Table 3. Percent Female and Hires Per Candidate by Recoded Level of the Organization (Number of cases in parentheses)

Level	Job Incumbents ^a	Hires ^b	Candidates ^b	Hires Per Candidate		
	<u>Percent Female</u>	<u>Percent Female</u>	<u>Percent Female</u>	<u>All</u>	<u>Female</u>	<u>Male</u>
<u>Non-Exempt:</u>						
Levels 1 -3	74.9 (171)	76.7 (206)	69.8 (470)	.438 (470)	.482 (328)	.338 (142)
4	70.8 (219)	78.8 (198)	69.3 (417)	.475 (417)	.540 (289)	.328 (128)
5	79.5 (185)	72.3 (238)	69.1 (863)	.276 (863)	.289 (596)	.247 (267)
6	86.0 (707)	80.2 (217)	70.0 (593)	.366 (593)	.419 (415)	.242 (178)
7	80.5 (292)	77.6 (147)	77.0 (675)	.218 (675)	.219 (520)	.213 (155)
8	83.7 (208)	75.8 (153)	55.9 (800)	.191 (800)	.259 (447)	.105 (353)
Levels 9 - 12	61.8 (165)	72.9 (48)	55.7 (357)	.134 (357)	.176 (199)	.082 (158)
<u>Exempt:</u>						
13	58.0 (50)	75.0 (48)	45.8 (406)	.118 (406)	.193 (186)	.054 (220)
14	62.0 (395)	60.6 (457)	44.7 (2221)	.206 (2221)	.279 (993)	.147 (1228)
15	54.3 (440)	45.6 (320)	33.4 (2002)	.160 (2002)	.219 (668)	.130 (1334)
16	38.8 (418)	42.2 (147)	30.1 (2544)	.058 (2544)	.081 (767)	.048 (1777)
17 +	32.0 (391)	37.1 (62)	28.2 (1132)	.055 (1132)	.072 (319)	.048 (813)
Total	64.5 (3641)	65.6 (2241)	45.9 (12480)	.180 (12480)	.256 (5727)	.114 (6753)

^a As of 12/31/95.

^b Includes jobs filled by external or internal candidates from 1/1/93 – 4/30/95.

Table 4. Logistic Regression Models Predicting Hire (Coefficients are Odds Ratios; z-values in parentheses)

	All			Externals			Internals		
	Model [1]	Model [2] ^a	Model [3] ^a	Model [4]	Model [5] ^a	Model [6] ^a	Model [7]	Model [8] ^a	Model [9] ^a
Female Candidate	2.673 (20.16)	1.964 (12.87)	2.036 (8.99)	2.636 (16.14)	1.772 (8.80)	1.960 (6.81)	1.177* (0.29)	1.038* (0.29)	.986* (-0.07)
Female^x Annual Salary of Level^b			.998* (-0.61)			.995* (-1.36)			1.002* (0.32)
Likelihood Ratio X^2 Degrees of Freedom	427.30 1	1,214.94 12	1,215.32 13	272.79 1	989.07 12	990.96 13	1.96 1	179.36 12	179.46 13
Number of Cases	12,480	12,480	12,480	11,081	11,081	11,081	1,399	1,399	1,399

* Not significant at $p < .05$ level.

^a Logit model, with fixed effects for 12 levels.

^b Annual salary of the top level of the salary range (in thousands of dollars).

Appendix 1. Annual Salary and Percent Female for Job Incumbents by Recoded Level of the Organization

Level	<u>Mean</u>	<u>Minimum</u>	<u>Maximun</u>	<u>Percent Female</u>	<u>Number of Cases</u>
<u>Non-Exempt:</u>					
Levels 1 -3	\$16,032	\$12,480	\$29,567	74.9	171
4	\$18,868	\$13,744	\$29,121	70.8	219
5	\$20,705	\$13,631	\$28,621	79.5	185
6	\$21,955	\$17,138	\$32,320	86.0	707
7	\$26,157	\$17,848	\$38,080	80.5	292
8	\$30,764	\$12,514	\$44,418	83.7	208
Levels 9 -12	\$35,904	\$22,033	\$55,500	61.8	165
<u>Exempt:</u>					
13	\$30,283	\$24,000	\$43,200	58.0	50
14	\$36,446	\$25,272	\$80,000	62.0	395
15	\$47,682	\$31,500	\$130,000	54.3	440
16	\$65,912	\$40,100	\$96,000	38.8	418
17 +	\$109,460	\$58,550	\$299,657	32.0	391
Total	\$42,139	\$12,480	\$299,657	64.5	3,641

Figure 1. Percent Female for Job Incumbents, All Hires, and All Candidate Pool

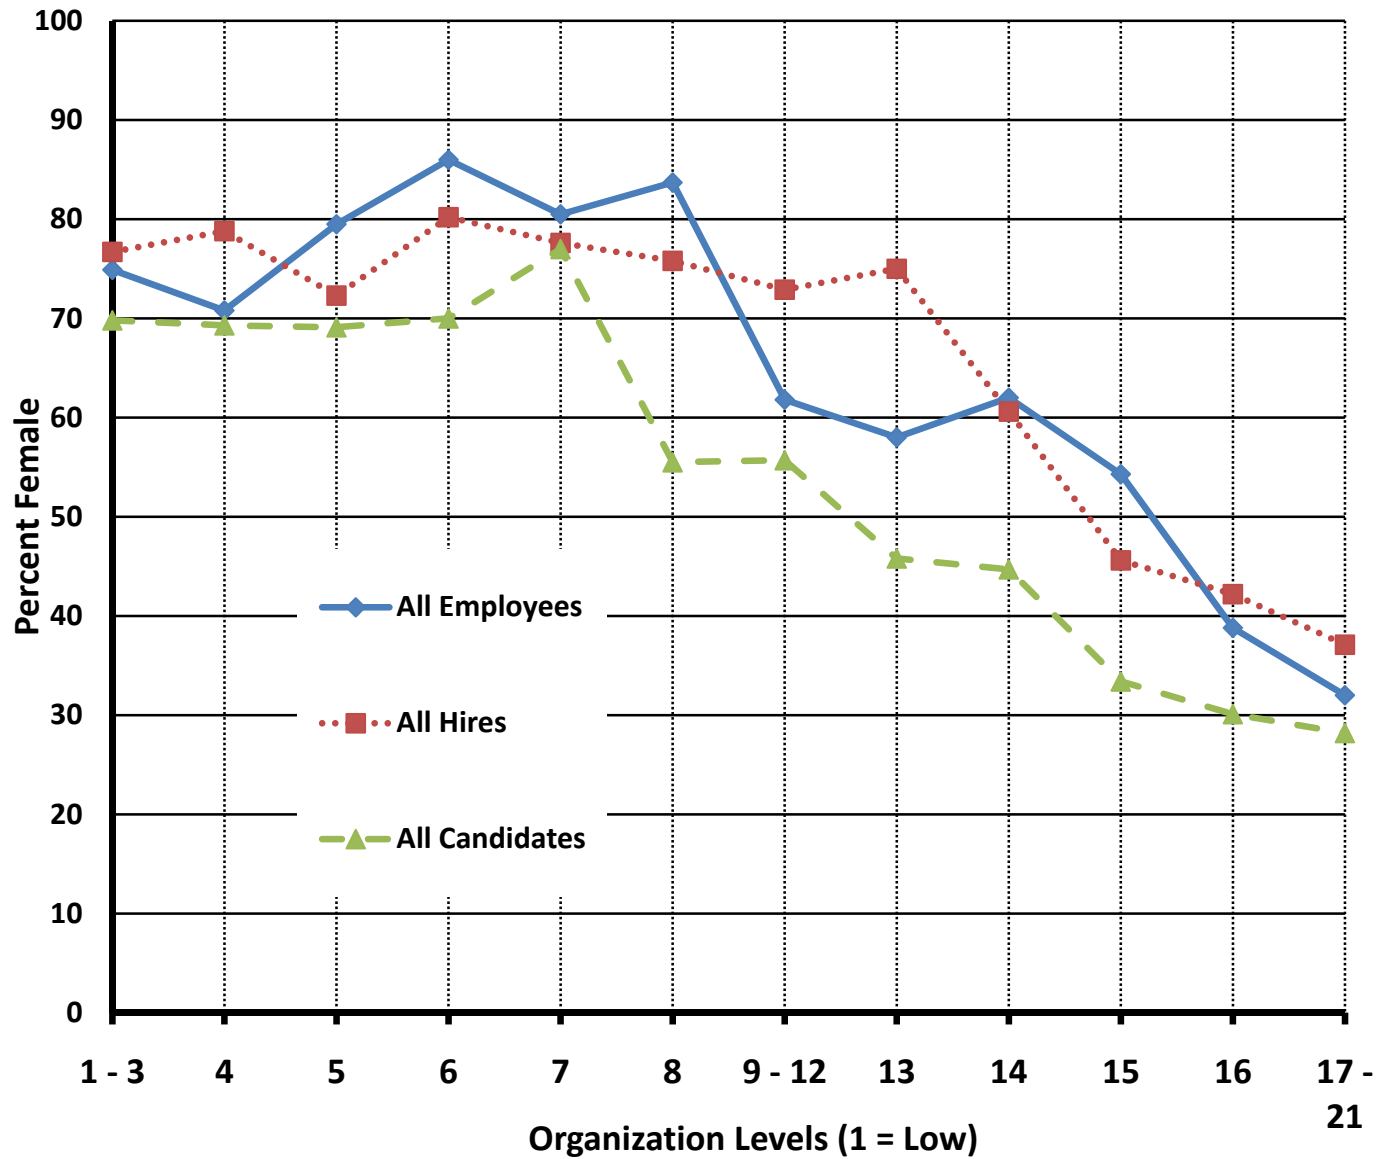


Figure 2. Hires Per Candidate for Female and Male

