

IZA DP No. 8797

**Why are There So Few Women in Executive Positions?
An Analysis of Gender Differences in the Life-Cycle of
Executive Employment**

Anders Frederiksen
Timothy Halliday

January 2015

Why Are There So Few Women in Executive Positions? An Analysis of Gender Differences in the Life-Cycle of Executive Employment

Anders Frederiksen

*Aarhus University
and IZA*

Timothy Halliday

*University of Hawaii at Manoa,
UHERO and IZA*

Discussion Paper No. 8797
January 2015

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Why Are There So Few Women in Executive Positions? An Analysis of Gender Differences in the Life-Cycle of Executive Employment

“Glass ceilings” and “sticky floors” are typical explanations for the low representation of women in top executive positions, but a focus on gender differences in promotions provides only a partial explanation. We consider the life-cycle of executive employment, which allows for a full characterization of the gender composition of executive management. We establish that there are few women in executive management because they have lower levels of human capital, are underrepresented in lower-level jobs, and are less likely to be perceived as high-productivity employees. We do not find that women have uniformly unfavorable promotion and demotion probabilities.

JEL Classification: J71, J62

Keywords: discrimination, dynamics, gender

Corresponding author:

Anders Frederiksen
Aarhus University
Birk Centerpark 15
7400 Herning
Denmark
E-mail: afr@asb.dk

1. Introduction

There are few women in top executive positions, and likewise, board memberships are dominated by men. “Glass ceilings” that prevent women from rising to top jobs and “sticky floors” that glue women to lower-level jobs are typical explanations for why women are rarely seen in top positions. For example, Olson and Becker (1983) show that the returns to promotions were similar for men and women, but women were held to higher promotion standards than men, with the result that fewer women are promoted. Blau and DeVaro (2007) survey the literature on gender differences in promotion rates and find that women generally face lower promotion probabilities, but they also stress that it may not be universally true (e.g., Hersch and Viscusi, 1996). In this analysis, we move beyond the literature’s focus on gender differences in promotion probabilities by broadening the perspective to include gender differences in the full life-cycle of executive employment. By considering both the flow into executive positions and the duration of executive employment we provide new and detailed evidence on why there are so few women in executive positions. In addition, we employ econometric techniques that allow us to condition on unobservable “productivity types,” which enables us to look more deeply into the relative dearth of women in top management.

The seminal work by Gary Becker, *The Economics of Discrimination* (1957), sparked a large literature on the topic. The main idea behind Becker’s theory is that agents may have “tastes for discrimination,” which can be used to rationalize discriminatory outcomes (at least in the short run). A different set of neoclassical models has explained discriminatory outcomes with a phenomenon known as statistical discrimination (see, e.g. Phelps (1972), McCall (1972), Arrow (1973), and Aigner and Cain (1977)). In these models, the outcome for minorities may be unfavorable because employers lack information about individual minority group members’ productive capacity; instead employers make decisions based on average minority group characteristics. A classic example is an employer who offers less training, lower wages, and fewer promotions to *all* women because *some* women leave the labor force when they marry (and men do not) and employee turnover is costly.

While early models of discrimination are static, contemporary models of discrimination are dynamic. The main advantage of dynamic models is that they allow for the study of promotions. Lazear and

Rosen (1990) study how men and women move within the hierarchy of a firm. They show that if men and women have similar ability distributions but women have non-market outside opportunities that are on average better than those for men, women will be subject to higher promotion standards than men. Fryer (2007) builds on this work by studying how initial discrimination through the application of higher standards in the hiring process for minority groups can be beneficial in the subsequent careers of those minority members who are actually hired. Booth, Francesconi, and Frank (2002) also build on the work of Lazear and Rosen (1990) by exploring two additional assumptions; the first is that women have worse market opportunities outside the firm, and the second is that firms are less likely to match women's outside offers. They establish that, in the first case, women are more likely to be promoted, and in the latter case, men and women have equal promotion probabilities. A somewhat different approach is taken by Bjerk (2008). He shows how skill differences or differences in signaling strength and possibilities between men and women can explain gender differences in promotion rates.

To understand why there are so few women in executive positions, the current empirical literature, with its focus on promotions, does not provide sufficient detail. Our analysis provides the additional details by considering gender differences in the life-cycle of executive employment, which involves estimating both promotion probabilities *and* the duration of executive employment. Because of the mixed evidence on gender differences in promotion rates provided in the literature (see Blau and DeVaro, 2007) and our lack of *a priori* expectations about gender differences in the duration of executive employment, we propose a generic statistical model of the life-cycle of executive employment. An important aspect of this model is that it is in line with existing theoretical models of discrimination, such as those developed by Bjerk (2008) and Freyer (2007), in that it can be used to predict career outcomes that are conditional on workers' (unobserved) productivity types.

We estimate the model using a register-based employer-employee dataset from Denmark that contains information on all workers and firms in Denmark. To avoid issues of education and retirement decisions and weak labor market attachment, we focus on a sample of employees who were continuously employed between 1992 and 2004 in private sector firms with 25 or more employees. In our sample, 6.5 percent of the employees are classified as executives, and the proportion of women in executive positions is 11.7 percent.

Our data provide some key insights into why there are so few women in executive positions. First, women make up only one-third of the employees in our sample of private sector employees. Second, education is an important driver for executive appointments, but in our sample only 7.7 percent of women are highly skilled (have obtained college/bachelor's degrees or above), whereas 21.9 percent of men are highly skilled. Given this gender and skill composition, even gender-neutral career paths would imply a low representation of women in executive employment.

The statistical analysis reveals a set of more subtle points. As discussed, theoretical models of discrimination such as Bjerk's (2008) and Freyer's (2007) build on employers' perceptions about employees' productivity types. Our statistical model uses insights from Heckman and Singer (1984), which allow us to identify the proportion of high- and low-productivity types by gender and education group. The regression results reveal that, for all skill groups, a higher proportion of men are considered high-productivity types. The results also show that the promotion probabilities for high-productivity types are relatively high and increasing over their careers, whereas the demotion probabilities are moderate and declining as their careers progress. This implies that, among the highly skilled, the proportion of high-productivity types in executive positions increases in age, and by the end of their careers almost 80 percent of high-productivity type women and 70 percent of high-productivity type men work in executive positions. In contrast, when low-productivity types are promoted the "mistake" is quickly corrected, which implies that the proportion of low-productivity types in executive positions remains low throughout their careers. This result accords with Lange (2007), who shows that firms quickly learn about workers' unobservable productivity. Thus, the strong sorting based on type and the observation that fewer women are perceived to be of the high-productivity type are additional explanations for why there are so few women in executive positions. In the discussion section of this paper, we provide some explanations for why private sector employment, educational achievement, and the proportion of high-productivity types are so low among women and to what extent this is likely to change in the future.

In the next section we present the data. In section 3 our empirical model is outlined in detail. The empirical results are presented in Section 4, which is followed by a discussion in Section 5. Our conclusions are presented in Section 6.

2. Data description

We use register-based employer-employee data from Denmark for the period 1992 to 2004. We impose four restrictions on the data. First, we focus our attention on the private sector. Second, to be able to clearly distinguish between executives and non-executives we sample only individuals working in firms with 25 or more employees. Third, we focus on employees with stable labor market careers (no registered spells of unemployment in the sample period). Finally, to avoid issues of education and retirement decisions and the fact that almost no one is an executive before age 30, we restrict the sample to employees aged 30 to 58. That is, our sample consists of prime aged employees in the private sector who are continuously employed in firms with 25 or more employees between 1992 and 2004.

Our definition of a top executive is based on the International Standard Classification of Occupations (ISCO) from the International Labor Organization (ILO). We define an executive as a person who has an ISCO code from Major Group 1 comprising corporate executives and general executives. All other employees are defined as non-executives. According to this definition, 6.5 percent of our sample is classified as executives.

Descriptive statistics are presented in Table 1 for women and men, respectively. A first observation is that the male sample is twice as big as the female sample. In addition, the probability of promotion to an executive position is almost three times higher for men than for women. Because the clear majority of employees are men and they are more likely to be promoted, many more men are promoted. This outcome, however, should be seen in the context that transitions out of executive positions also are three times higher for men, leaving the net inflow undetermined.

Splitting the samples by education (Low-skilled, Skilled, and High-skilled)¹ reveals the importance of human capital for appointment to a top executive position. High-skilled women are almost five times more likely to become executives than low-skilled women. For men the ratio is three to one. Further, low-skilled men are 2.6 times more likely than low-skilled women to work as executives. This ratio increases to 3.6 for skilled workers and drops to 2 for the highly skilled. These transitions into executive positions are matched with almost similar transitions out of executive positions for each gender and skill level.

While this data description sheds some light on why there are so few women in executive positions, it does not fully address the life-cycle of executive employment. To fully address this issue we turn to estimation of a dynamic panel data model with unobserved heterogeneity.

3. Statistical model

In this section, we lay out a simple statistical model for the life-cycle of executive employment. We view this model as a statistical model as opposed to a structural model of economic behavior, but in the presentation we point out structural interpretations of different aspects of the model. The model serves two purposes. First, the model provides a framework for estimating the transitions into executive positions and the duration of executive employment by gender and conditional on a worker's productivity type. Second, the model can be used to predict the employees' career outcomes conditional on their types.

3.1 The model

We consider a simple model of employment dynamics within the firm. There are a total of T periods which are indexed $t \in \{0, 1, \dots, T\}$. In each period, an employer must decide how to assign workers to one of two layers of the firm's hierarchy: executive positions or non-executive positions. We denote worker i 's position within the hierarchy at time t by $m_{it} \in \{0, 1\}$. If $m_{it} = 0$ ($m_{it} = 1$), then worker i is in a non-executive position (executive position) at time t . An employer's decision about where to assign workers is based on his or her beliefs about worker productivity. We consider a simple scenario

¹ "Low-skilled" employees have completed compulsory primary school education. "Skilled" employees are individuals with high-school and vocational degrees, and "High-skilled" employees have college degrees (the equivalent of a bachelor's degree in the Danish context) and above.

in which there are heterogeneous beliefs about productivity and transaction costs. For $t > 0$, productivity beliefs about workers in non-executive and executive positions are given by

$$\beta_{0it} = \alpha_{0i} + \gamma_{0i}(1 - m_{i(t-1)}) + f_0(t) + u_{0it} \quad (1)$$

and

$$\beta_{1it} = \alpha_{1i} + \gamma_{1i}m_{i(t-1)} + f_1(t) + u_{1it} \quad (2)$$

where $\gamma_{0i} > 0$ and $\gamma_{1i} > 0$. One interpretation of the first terms in (1) and (2) is that they model beliefs about baseline productivity in either the non-executive or executive state. Importantly, this may reflect either actual productivity or perceptions about productivity. We remain agnostic about how these beliefs are formed. These are assumed to be heterogeneous to account for productivity differences within the firm. The second terms are adjustment costs. If the worker has only been in a given state for one period, then productivity beliefs are given by α_{0i} in non-executive positions and α_{1i} in executive positions. However, once the worker has been in that state for more than one period, she learns the processes and procedures that go along with her position, and productivity (or its perception) increases to $\alpha_{0i} + \gamma_{0i}$ in non-executive positions and to $\alpha_{1i} + \gamma_{1i}$ in executive positions. We assume that if the worker leaves an executive position or a non-executive position and then re-enters it, she will have to re-learn the position's processes and procedures. The third term is a homogeneous function of age and can be interpreted as employee learning. The final term is an idiosyncratic component of productivity beliefs. For $t > 0$, the employer will assign a worker to an executive positions if he or she believes that the employee will be more productive in an executive positions than in a non-executive positions—*i.e.*,

$$m_{it} = 1(\beta_{1it} > \beta_{0it}) = 1(\alpha_i + \gamma_i m_{i(t-1)} + f(t) + u_{it} > 0)$$

where $\alpha_i \equiv \alpha_{1i} - \alpha_{0i} - \gamma_{0i}$, $\gamma_i \equiv \gamma_{1i} - \gamma_{0i}$, $f(t) \equiv f_1(t) - f_0(t)$ and $u_{it} \equiv u_{1it} - u_{0it}$.

3.2 Maximum Likelihood Estimation

We assume that u_{it} follows a logistic distribution, is serially uncorrelated, and is independent of (θ_i, m_{i0}) for all t where $\theta_i \equiv (\alpha_i, \gamma_i)$. This implies that

$$P(m_{it}|m_{i(t-1)}, \theta_i) = \Lambda\left(\left(\alpha_i + \gamma_i m_{i(t-1)} + f(t)\right) (2m_{it} - 1)\right). \quad (3)$$

Note that the econometrician only observes employment outcomes for $t \in \{\tau_i, \dots, T_i\}$ where $\tau_i \geq 0$. Next, we model the unobserved heterogeneity as a series of discrete mass points in the spirit of Heckman and Singer (1984), so that $\theta_a \in \{\theta_L, \theta_H\}$ where we assign probability π_a to the a th mass point. Throughout we refer to θ_L as the “low-productivity type” to reflect that these individuals have a lower probability of being appointed to an executive position; similarly, θ_H is the “high-productivity type,” reflecting a greater likelihood of being appointed to an executive position. As discussed above, the pair $\{\theta_L, \theta_H\}$ can be thought of as reflecting actual (or beliefs about) worker productivity.

We obtain the likelihood of a sequence of outcomes from $t = \tau_i$ to $t = T_i$ as

$$P(m_{iT_i}, \dots, m_{i\tau_i}) = \sum_{a \in \{L, H\}} \prod_{t=\tau_i+1}^{T_i} P(m_{it}|m_{i(t-1)}, \theta_a) P(m_{i\tau_i}|\theta_a) \pi_a. \quad (4)$$

Now, because the process has been running prior to $t = \tau_i$, we must compute $P(m_{i\tau_i}|\theta_a)$, as in Heckman (1981) and Halliday (2008). This is the famous initial-conditions problem. To address this, we note that any probability of being in an executive position conditional on θ_i can be written as

$$P(m_{it}|\theta_i) = \sum_{d=0}^1 P(m_{it}|m_{i(t-1)} = d, \theta_i) P(m_{i(t-1)} = d|\theta_i). \quad (5)$$

We then evaluate this expression at $t = \tau_i$ and substitute back into this equation recursively until $t = 1$. Because we model the unobserved heterogeneity as discrete mass points, the probabilities of the initial conditions, $P(m_{i0}|\theta_i)$, can be estimated directly from the data without any additional assumptions. In this sense, this MLE procedure provides a solution to the initial-conditions problem that requires no

auxiliary distributional assumptions. Based on this, the likelihood function can be obtained by taking the log of the probabilities in equation (4) and summing across individuals in the data.

4. Results

We estimated the model separately for women and men and for three skill groups: High-skilled, Skilled, and Low-skilled, each of which corresponds to a separate educational category as described above. We report parameter estimates and their standard errors in Tables A1 and A2 in the Appendix. Recall that within each skill category there are two types: high and low, reflecting employer beliefs about the employees' productivity types. In this section, we present our estimation and then offer some thoughts about their economic implications and underpinnings in the next discussion section.

In Figure 1, we report the raw probability of being in an executive position by gender and skill group. These probabilities are *not* conditional on productivity type. Instead, they reflect a weighted average across types. These profiles are depicted by the solid lines in each panel of the figure, corresponding to the left axis. In addition, we report the proportion of managers who are women. This profile is depicted by the dashed line in each panel and corresponds to the right axis.

The figure shows four facts. First, within skill groups, men are more likely than women to be executives. Second, the proportion of executives increases with education for both men and women. Third, except for an initial drop for highly educated women, the probability of being an executive increases monotonically with age. Fourth, the figure shows that, for all skill groups, between 10% and 20% of managers are women. These profiles increase for the high-skilled and the skilled but decrease for the low-skilled.

In Figure 2, we report promotion probabilities conditional on productivity type by gender and skill group. These figures reveal a level of detail that is new to the literature. For highly skilled employees, women of the high-productivity type (light dotted lines) initially have lower promotion probabilities than men (dark dotted lines), but they overtake men during their early forties, and these higher promotion probabilities prevail for the remainder of their careers. For skilled employees, high-productivity type women have higher promotion probabilities than men at all ages. For the lowest skill

group, however, women have higher promotion probabilities only when they are in their early thirties. Employees of the low-productivity type (solid lines) have very low promotion probabilities regardless of gender and skills.

While the promotion probabilities reveal interesting new patterns, they should be seen in the light that relatively few women are considered to be of the “high-productivity” type (the quotation marks here are deliberate to indicate that this may be a mere perception). Our estimation results reveal that in the high-skilled group only 11.4 percent of women are perceived to be of the high-productivity type, whereas 21 percent of men are considered to be in this category. In the skilled group high-productivity types constitute 2.7 percent of women and 10.8 percent of men, and among the low-skilled only 3.2 percent of women are considered to be of the high-productivity type; for low-skilled men the figure is 8.3 percent.

To fully characterize the life-cycle of executive employment, we present demotion probabilities conditional on type by gender and skill group in Figure 3. A clear pattern emerges from the graphs: employees considered to be of the low-productivity type have relatively high demotion probabilities throughout their careers regardless of gender and skill, and these demotion probabilities decline over the course of their careers. Note that we allow the coefficient on $m_{i(t-1)}$ denoted earlier by γ_i to be a heterogeneous parameter, which implies that the higher demotion rates are not merely a mechanical aspect of the model as they would have been had γ_i been a homogeneous parameter. So, low-productivity types not only have low promotion probabilities; once promoted, they do not last long in executive positions. Hence, sorting of low-productivity type employees seems to be immediate: when firms mistakenly promote a low-productivity type employee to an executive position this is quickly realized and corrected.

Employees of the high-productivity type have relatively low demotion probabilities, yet these probabilities are around 20-30 percent early in their careers and then decline monotonically with age to a level around 10 percent when nearing retirement. Only low-skilled women deviate from this pattern: their demotion probability is virtually flat at a level of 20 percent throughout their career. Interestingly, high-skilled women have demotion probabilities that exceed those for men early in their career, but

around the age of 40 the demotion probabilities drop below those for men and stay lower for the remainder of their career. For skilled employees there are virtually no gender differences in demotion probabilities for high-productivity types. But for the low-skilled the demotion probabilities start out at 20 percent for both men and women; although the probability does not change for women as their careers progress, it declines steadily for men.

The promotion and demotion probabilities combined with the stocks of employees in, respectively, executive and non-executive positions determine the probability that a given type of employee is working in an executive position. In Figure 4 we present these probabilities by gender and skill level. For all groups, relatively few employees work in executive positions early in their careers. We also find that employees considered to be of the low-productivity type have a probability close to zero of working in an executive position throughout their careers because of the relatively low likelihood of promotion and the relatively high likelihood of demotion. It is only the highly educated employees in this group who have a discernable increase in the probability of working in an executive position as their careers progress. Yet even at its peak close to retirement, the probability is less than 10 percent.

For employees considered to be of the high-productivity type, the dynamics are more interesting. For high-skilled individuals, we have established that men have relatively high promotion and relatively low demotion probabilities early in their careers. This implies that the proportion of these men working in an executive position increases relatively quickly early in their careers. Over time, however, the growth rate stabilizes at a somewhat lower level, and by the end of their careers 70 percent are working in executive positions. For women the dynamics are different. Women have relatively low promotion probabilities and relatively high demotion probabilities early in their careers, which implies a convex growth path. Of particular interest is that by the end of their careers more than 80 percent of the women who are considered to be of the high-productivity type are working in executive positions, a proportion that exceeds the proportion for men.

For the high-productivity types who are skilled or low-skilled, the probability of working in an executive position follows a concave growth path. Skilled women have a higher probability of working in the executive ranks than comparable men throughout their careers, and by the age of 60 more than

70 percent of these women are executives, whereas only 60 percent of men are executives. For the low-skilled the curves are almost identical for men and women up to age 37, at which time the probability of working in an executive position starts to decline for women and continues to increase for men. This implies that by the end of their careers 60 percent of the low-skilled men who are considered to be of the high-productivity type are working in executive jobs; for women it is only 35 percent.

Finally, in Figure 5, we plot the proportion of high-productivity types in management by gender and skill group.² For all three skill groups there is a sharp increase in the proportion of high-productivity types in management from age 30 to age 35. After that, there is a slow decline in this proportion for high-skilled people and a slower decline for skilled people. However, for the low skilled, the proportion of high-productivity types is steady from age 35 to age 60 at just under 90 percent. Another notable feature of this figure is that the percentage of high-productivity types in management declines with education. These findings indicate that firms quickly learn about latent productivity (as established in Lange (2006)), but also that employers are continuously trying to hit a moving target (Kahn and Lange, forthcoming).

5. Discussion

Our results show that the gender gap in executive employment is caused primarily by women’s relatively low representation in larger private sector firms, women’s lower education levels, and the fact that relatively few women are perceived to be of the high-productivity type. Gender differences in promotion and demotion probabilities (conditional on type) do not provide a uniform explanation for why there are so few women in executive positions. It is the case that highly educated women experience “late promotions” and men do not, but later in their careers high-skilled women perceived to be of the high-productivity type have a higher likelihood of promotion than comparable men. Furthermore, throughout their careers, skilled women of the high-productivity type have higher promotion rates than men.

² Specifically, we plot

$$P(\theta_H | m_{it} = 1) = \frac{P(m_{it} = 1 | \theta_H) \pi_H}{P(m_{it} = 1)}$$

A key driver for executive appointments is education. This is a variable that can be affected by public policy or simply through structural shifts in society. In fact, recent works indicate that the education gap between men and women has narrowed substantially and that it has been reversed in many OECD countries, including Denmark (Goldin et al. (2006)). Clearly, as women acquire more education relative to men their representation in executive positions will increase. However, and importantly, our results indicate that, within education groups, important differences in career prospects are likely to ameliorate the under-representation of women in executive positions only to some extent. As has been long understood, picking the (relatively) low-hanging fruit of equalizing educational attainment across genders will not deliver equal representation of women in executive employment. Hence, a better understanding of the distribution of productivity types is in order. Unfortunately, understanding why so few women are perceived to be high-productivity types is a very thorny issue.

An equally important issue is that a relatively small share of women work in large private sector firms (one-third in our data). The direct consequence of this is that, absent strong affirmative action favoring women or strong positive selection of underlying productivity types into the private sector, women's representation in executive positions will be less than 50 percent. A more interesting question is, however, why so few women are working in the private sector. Altonji and Blank (1999) survey theories of occupational segregation, but it is also likely that some of the mounting evidence from, e.g. laboratory experiments on gender differences in risk attitudes, behavior in competitions and negotiations, preferences, and attitudes toward incentive pay may have explanatory power. Bertrand (2011) discusses the importance of these gender differences for labor market and career outcomes. We stress that such a discussion also should address differences between men and women in sector choice, as sector choice appears to be an important factor affecting the representation of women in executive employment.

Equalizing observable gender differences in education and sector choice will contribute to raising the representation of women in executive positions, but it is unlikely that it will equalize the gender composition of top management. The reason is that underlying differences in the proportions of women and men who are perceived to be high-productivity types will continue to influence the gender composition of executive employment. Thus, only in the special case, where an equalization of

observable differences implies equalization of perceptions about underlying productivity types, will the gender gap vanish. If differences in underlying productivity types between men and women persist, we will continue to observe a lower representation of women in executive management. This is an important insight because it points toward an explanation for the gender gap that is unrelated to education and sector choice.

There are several potential explanations for the gender difference in perceived productivity types. One is simply that there are fewer high-productivity women than men in private sector jobs. A related but more subtle explanation is a story of statistical discrimination in which employers perceive that women have lower lifetime productivity; as a consequence, women rationally make fewer investments in human capital (other than education) and are therefore less productive (Altonji and Blank, 1999). A final explanation is that the difference in the productivity distribution is an employer error, although as has long been understood, it is not clear that such errors can persist in competitive markets.

6. Conclusion

The low representation of women in executive employment is universal. While many attempts have been made in the literature to explain the gender gap, its focus on gender differences in promotion rates provides insufficient detail. In this analysis we have broadened the perspective to include gender differences in the life-cycle of executive employment. The estimation of both promotion probabilities and the duration of executive employment allows us to characterize the flows into and out of the top ranks, which determine the gender composition in top management.

Our results show that observable gender differences in the employment composition at lower ranks and education levels explain part of the gender gap in executive employment. But our results also reveal that women do not have uniformly unfavorable promotion and demotion probabilities (conditional on productivity types). Instead, a large part of the gender gap in executive employment is caused by the fact that relatively few women are perceived to be of the high-productivity type. This implies that the gender gap is caused by factors that are unrelated to education and sector choice, which makes it particularly difficult to influence through policy.

The main unanswered question that arises from our work is why employers perceive that there are relatively fewer high-productivity women than men. Some straightforward field work could survey employees in firms where productivity is easily measured and compare measures of actual productivity to employer surveys about perceived productivity. Although, as discussed above, less direct methods seem to indicate that employers have fairly accurate knowledge about employee productivity. Thus the productivity types we have identified here would seem to reflect real as opposed to perceived productivity. If this is indeed the case, then future work should devise novel surveys and field experiments that can test models of statistical discrimination to see if there are more fundamental explanations for the productivity differences.

References

- Aigner, D. J. and G. G. Cain, 1977, Statistical Theories of Discrimination in Labor Markets, *Industrial and Labor Relations Review*, vol. 30, pp. 175-187.
- Altonji, J. G. and R. M. Blank, 1999, Race and Gender in the Labor Market, *Handbook of Labor Economics*, Volume 3, Edited by O. Ashenfelter and D. Card, Elsevier science B. V.
- Arrow, K., 1973, The Theory of Discrimination, in O. Ashenfelter and A. Rees, eds., *Discrimination in the Labor Market*, Princeton University Press, Princeton.
- Becker, G. S., 1957, *The Economics of Discrimination*, University of Chicago Press, Chicago.
- Bertrand, M., 2011, New Perspectives on Gender, *Handbook of Labor Economics*, Volume 4b, Edited by D. Card and O. Ashenfelter, Elsevier science B. V.
- Bjerk, D., 2008, Glass Ceilings or Sticky Floors? Statistical Discrimination in a Dynamic Model of Hiring and Promotion, *Economic Journal*, vol. 118, pp. 961-982.
- Blau, F. D. and J. DeVaro, 2007, New Evidence on Gender Differences in Promotion Rates: An Empirical Analysis of a Sample of New Hires, *Industrial Relations*, vol. 46, pp. 511-550.
- Booth, A. L., M. Francesconi and J. Frank, 2003, A Sticky Floors Model of Promotion, Pay, and Gender, *European Economic Review*, vol. 47, pp. 295-322.
- Fryer, R., 2007, Belief Flipping in a Dynamic Model of Statistical Discrimination, *Journal of Public Economics*, vol. 91, pp. 1151-1166.
- Goldin, C., L. F. Katz, and I. Kuziemko, 2006, The Homecoming of American College Women: The Reversal of the College Gender Gap, *Journal of Economic Perspectives*, vol. 20, pp. 133-156.
- Hersch, J. and W. K. Viscusi, 1996, Gender Differences in Promotions and Wages, *Industrial Relations*, vol. 35, pp. 461-472.
- Halliday, T. J., 2008, Heterogeneity, State Dependence and Health, *The Econometrics Journal*, vol. 11, pp. 499-516.

Heckman, J. J., 1981, The Incidental Parameters Problem and the Problem of Initial Conditions in Estimating a Discrete Time-Discrete Data Stochastic Process, in C. Manski and D. McFadden (Eds.), *Structural Analysis of Discrete Data*, Cambridge, MA: MIT Press.

Heckman, J. J. and B. Singer, 1984, A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data, *Econometrica*, vol. 52, pp. 271–320.

Kahn, L. B. and F. Lange, *forthcoming*, Employer Learning, Productivity and the Earnings Distribution: Evidence from Performance Measures, *Review of Economic Studies*.

Lange, F., 2007, The Speed of Employer Learning, *Journal of Labor Economics*, vol. 25, pp. 1-35.

Lazear, E. P. and S. Rosen, 1990, Male-Female Wage Differentials in Job Ladders, *Journal of Labor Economics*, vol. 8, pp. S106-S123.

McCall, J. J., 1972, *Income Mobility, Racial Discrimination, and Economic Growth*, Lexington Books, D. C. Heath and Co., Lexington.

Olson, C. A. and B. E. Becker, 1983, Sex Discrimination in the Promotion Process, *Industrial and Labor Relations Review*, vol. 36, pp. 624-641.

Phelps, E. S., 1972, The Statistical Theory of Racism and Sexism, *American Economic Review*, vol. 62, pp. 659-661.

Table 1: Descriptive statistics: Women and Men in Executive Positions in Denmark, 1992-2004

	Women			
	Low-skilled	Skilled	High-skilled	All
Age	44.88 (5.31)	44.50 (5.76)	43.00 (5.60)	44.47 (5.66)
Percent employed in executive positions	1.829	1.886	7.982	2.344
Yearly transition into executive positions	0.004	0.005	0.019	0.006
Yearly transition out of executive positions	0.004	0.004	0.015	0.005
Observations	64,788	207,240	22,764	294,792

	Men			
	Low-skilled	Skilled	High-skilled	All
Age	44.73 (5.65)	44.30 (5.72)	43.88 (5.78)	44.28 (5.73)
Percent employed in executive positions	4.825	6.737	16.083	8.455
Yearly transition into executive positions	0.011	0.014	0.033	0.018
Yearly transition out of executive positions	0.009	0.012	0.027	0.015
Observations	105,468	376,044	134,892	616,404

Figure 1: Proportion of employees in executive positions in Denmark, 1992–2004

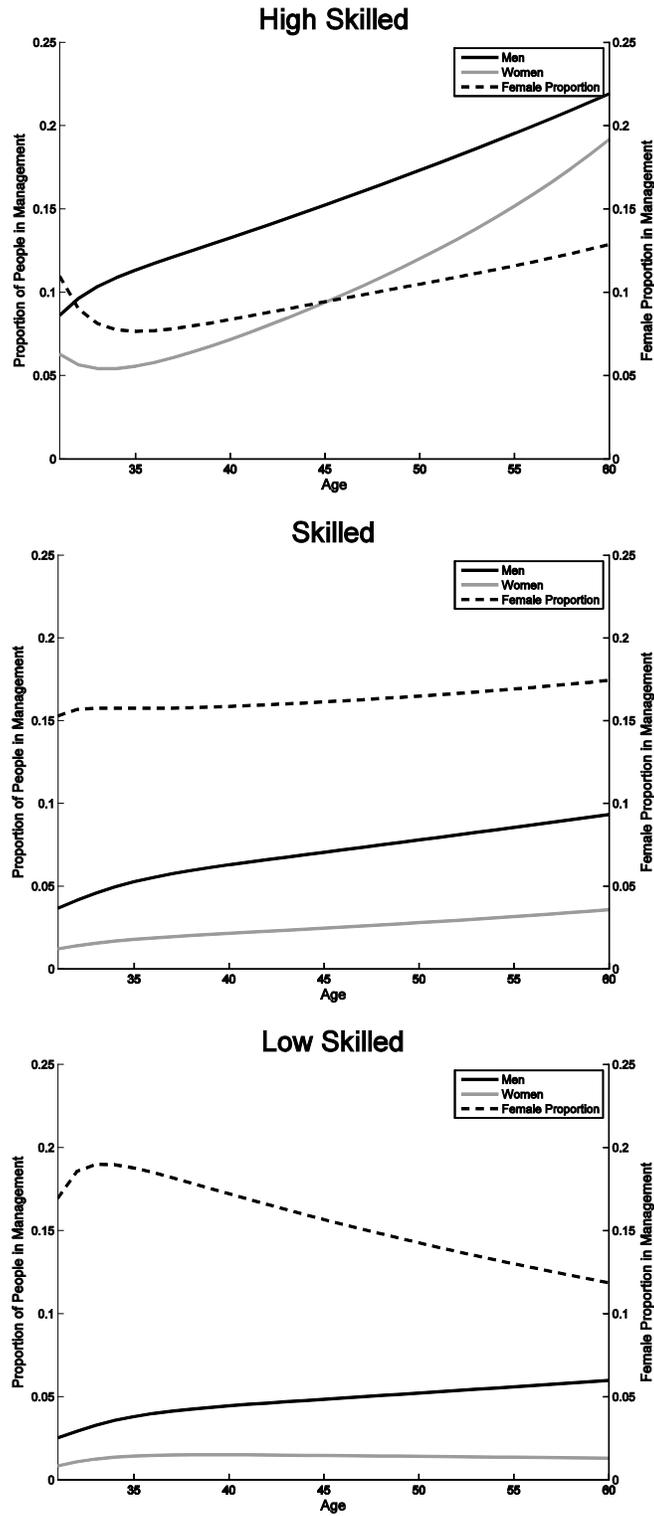


Figure 2: Promotion probabilities

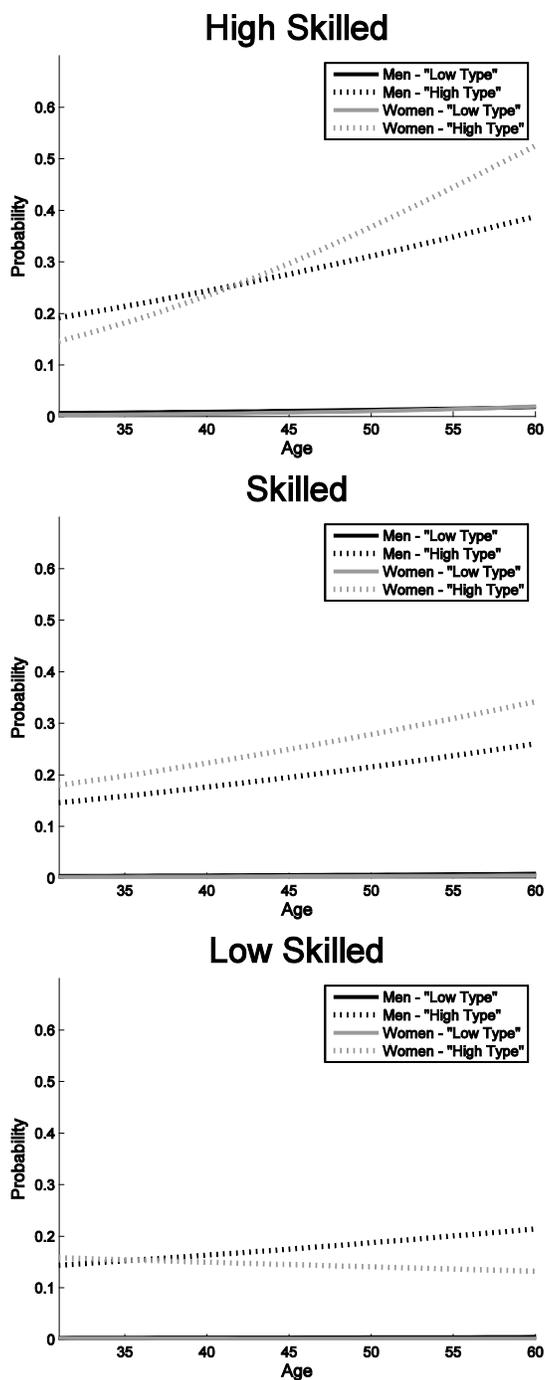


Figure 3: Demotion probabilities

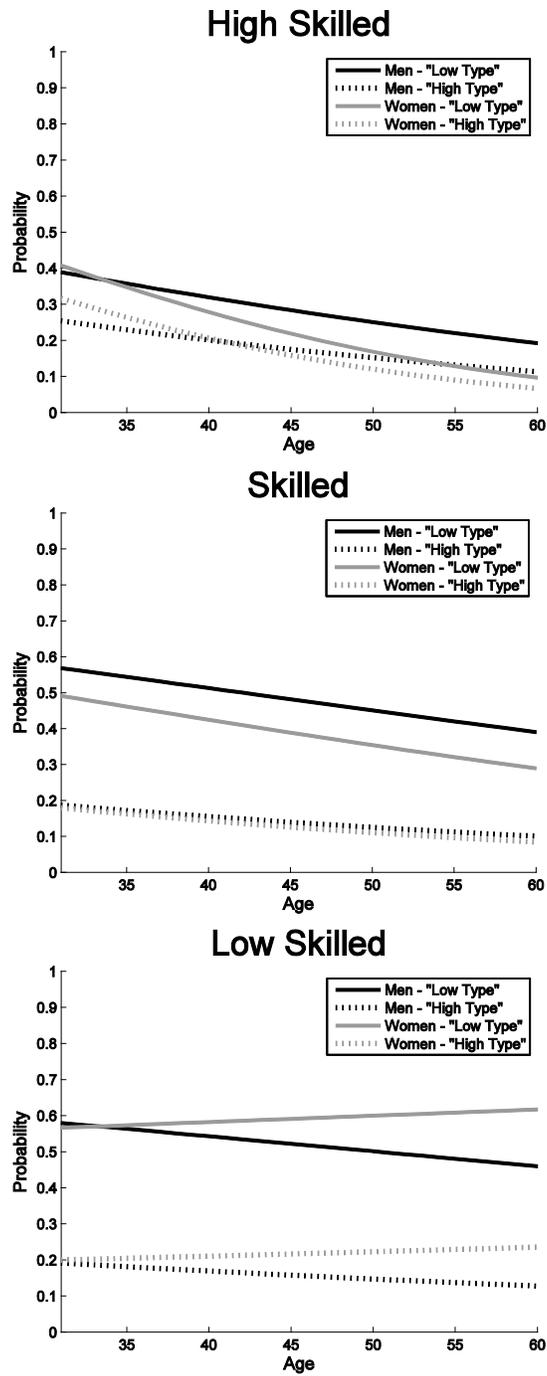


Figure 4: Proportion in executive employment

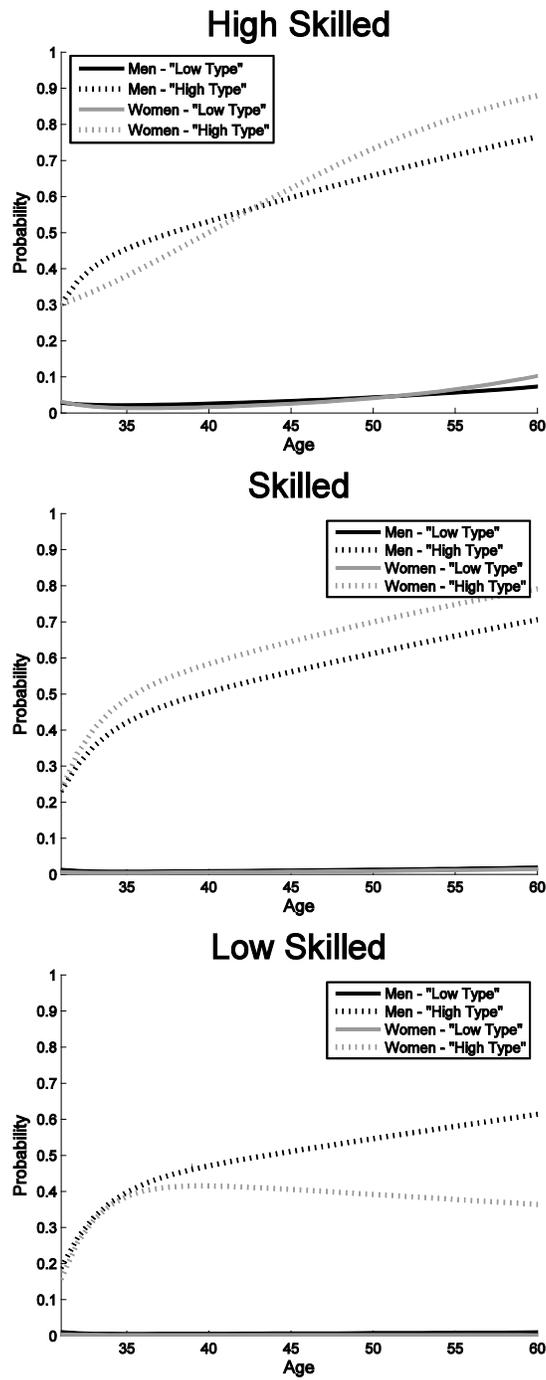
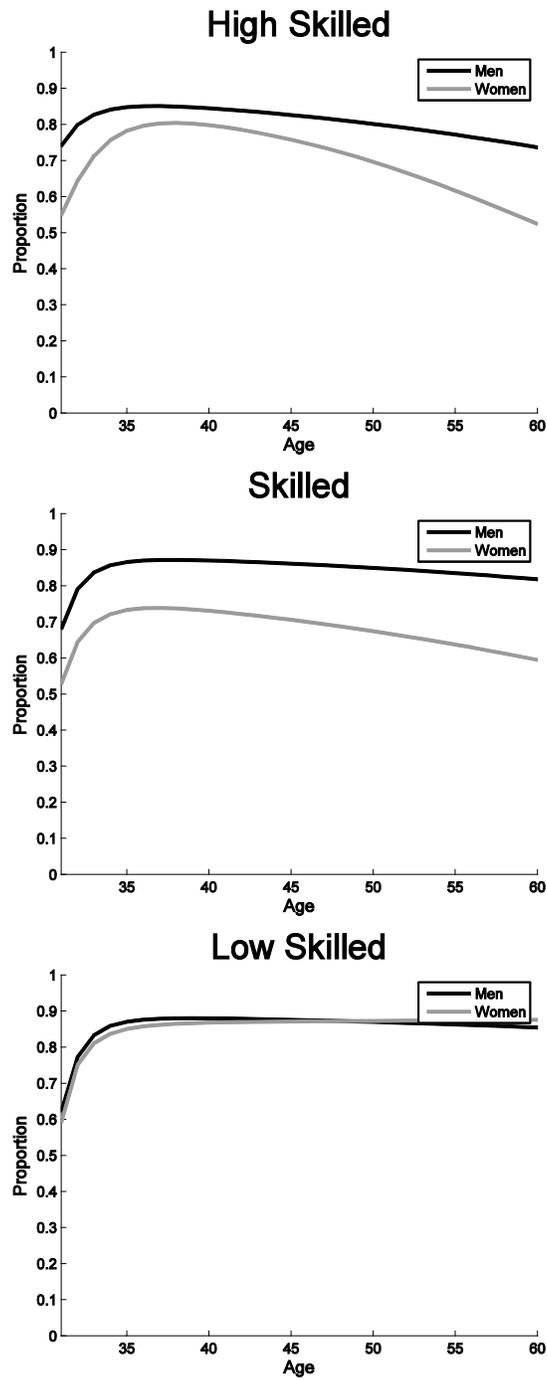


Figure 5: Proportion of high-productivity types in executive employment



Appendix

Table A1: Parameter estimates, women

	Women		
	Low-skilled	Skilled	High-skilled
γ_1 (lagged management – low type)	6.399 (0.379)	6.219 (0.173)	6.147 (0.573)
γ_2 (lagged management – high type)	3.060 (0.167)	3.037 (0.123)	2.532 (0.229)
ρ_1 (age)	-0.074 (0.165)	0.290 (0.088)	0.643 (0.284)
ρ_2 (age ²)	0.013 (0.011)	-0.015 (0.006)	-0.006 (0.014)
α_1 (constant – low type)	-6.438 (0.567)	-7.109 (0.304)	-7.767 (1.043)
α_2 (constant – high type)	-1.446 (0.595)	-2.449 (0.309)	-3.755 (1.023)
Probability Management - 30 (Low Type)	0.0052	0.0074	0.0491
Probability Management - 30 (High Type)	0.0000	0.0872	0.2888
Probability High Type	0.032	0.027	0.114
Log likelihood	-2,564	-9,075	-3,110
Observations	64,788	207,240	22,764

Note: Standard errors in parentheses. Age is divided by 10 and age square is divided by 100.

Table A2: Parameter estimates, men

	Men		
	Low-skilled	Skilled	High-skilled
γ_1 (lagged management – low type)	5.555 (0.278)	5.262 (0.129)	5.416 (0.120)
γ_2 (lagged management – high type)	3.224 (0.093)	3.236 (0.046)	2.514 (0.058)
ρ_1 (age)	0.167 (0.090)	0.249 (0.045)	0.338 (0.050)
ρ_2 (age ²)	0.011 (0.005)	-0.002 (0.003)	0.006 (0.003)
α_1 (constant – low type)	-6.392 (0.331)	-6.310 (0.160)	-6.012 (0.173)
α_2 (constant – high type)	-2.300 (0.333)	-2.538 (0.167)	-2.486 (0.175)
Probability Management - 30 (Low Type)	0.0183	0.0216	0.0355
Probability Management - 30 (High Type)	0.0676	0.1274	0.1996
Probability High Type	0.083	0.108	0.211
Log likelihood	-9,343	-4,295	-3,017
Observations	105,468	376,044	134,892

Note: Standard errors in parenthesis. Age is divided by 10 and age square is divided by 100.