# The Career Dynamics of High-Skilled Women and Men: Evidence from Sweden

James Albrecht* Mary Ann Bronson* Peter Skogman Thoursie†
Susan Vroman*

January 2017

**Abstract**

In this paper, we use matched worker-firm register data from Sweden to examine the career dynamics of high-skill women and men. Specifically, we track wages for 20 years among women and men born in the years 1960 - 70 who completed a university degree in business or economics. These women and men had essentially identical wages and earnings at the start of their careers, but their career paths diverge substantially as they age. We first look at whether firm effects account for the differences we observe between women's and men's wage profiles. The firms where men work differ in some important ways from the corresponding "female firms." However, a wage decomposition suggests that these firm effects do not play a major role in explaining the gender log wage gap among these workers. We then examine whether gender differences in firm-to-firm mobility help explain the patterns in wages that we see. Men are more likely than women are to switch from one firm to another, and this difference is especially pronounced in early career when the returns to mobility are highest. However, this gender difference in mobility rates does not explain – at least not directly – the glass ceiling effect that we see in the data. Instead, we find that the main driver of the gender difference in log wage profiles is that men experience higher wage gains than women do both as "switchers" and as "stayers."

*Department of Economics, Georgetown University.
†Department of Economics, Stockholm University.

# 1 Introduction

Gender differences in wages and earnings are particularly pronounced among top earners. This "glass ceiling" is seen across a range of countries with different institutions and policy environments. For example, in Sweden, a country with generous parental leave and an extensive public day care system, the gender gap in wages at the median is approximately 0.10 log points while the corresponding gap at the 90th percentile is approximately 0.30 log points (Albrecht et al. 2003, 2015). In the United States, a country with a very different policy environment, a similar though less extreme pattern can be seen (Blau and Kahn (2016)), and Arulampalam et al. (2007) document a glass ceiling pattern across several OECD countries. Increasing gender gaps in the higher percentiles of the wage and earnings distributions are observed even for narrowly defined skill groups, e.g., for lawyers (Wood et al. (1993)) and for MBA's working in the corporate and financial sectors (Bertrand et al. (2010)).

In this paper, we use matched worker-firm register data from Sweden to examine the career dynamics of high-skill men and women. Specifically, we track wages for 20 years among men and women born in the years 1960-70 who completed a university degree in business or economics. These men and women had essentially identical wages and earnings at the start of their careers, but their career paths diverge substantially as they age.

How can we understand these patterns? What accounts for different wage and earnings dynamics by gender? Because we have matched employer-employee data, we are able to focus on the role played by worker-firm interactions. The literature on wage dynamics suggests three potentially important channels. First, there is evidence for substantial firm-specific wage premiums (Abowd et al. (1999) and Card et al. (2016)), and wages may grow more rapidly in firms where men tend to work than they do in firms where women are more likely to be employed. That is, men may tend to work in "better" firms than women do. Second, wages increase with on-the-job search (Burdett and Mortensen (1998)). In particular, men and women may receive outside offers at different rates; the distributions from which offers are drawn may differ by gender; and offer acceptance choices may be different for men and women. In short, men may realize higher returns from firm-to-firm mobility than women do. Third, wage growth may differ by gender within similar firms (Gibbons and Waldman (1999)). This might arise for a host of reasons – gender differences in human capital accumulation, tracking, discrimination, etc.

These explanations are not mutually exclusive. For example, if men receive outside

offers at a higher rate than women do (or if the outside offers they receive tend to be better), then they are more likely to move to "better" firms and their wages are likely to increase more rapidly even if they don't move from one firm to another (e.g., via matching of outside offers as in Postel-Vinay and Robin (2002) or via wage-tenure contracts that are designed to reduce turnover as in Burdett and Coles (2003)). Nor are our explanations inconsistent with the common focus on motherhood as a driver of the gender log wage gap. We are able to follow the workers in our dataset year-by-year, and we observe when their first (and subsequent) children are born. We can thus see whether the probability of working in a high-paying firm changes as men and women become parents, whether parenthood affects firm-to-firm mobility and whether within-firm wage gains change after the first child is born.

The roadmap for the rest of our paper is as follows. In the next section, we describe the data that we use in our analysis. Then in Section III, we present some basic information about how wages, earnings and hours develop over the careers experienced by the men and women in our data. We look at how men's and women's wages evolve with age, and how these profiles are affected by becoming a parent, and we also present complementary information about earnings and hours profiles. In Section IV, we turn to possible explanations. First, we examine the question of whether firm effects are important for understanding the differences we observe between men's and women's wage profiles. The firms where men are more likely to work differ in some important ways from the corresponding "female firms." Men tend to work in smaller and higher-paying firms, but a series of wage decompositions suggests that these firm effects do not play a major role in explaining the gender log wage gap among these workers. We then examine whether gender differences in firm-to-firm mobility help explain the patterns in wages that we see. Men are more likely than women are to switch from one firm to another, and this difference is especially pronounced in early career when the returns to mobility are highest. However, this gender difference in mobility rates does not explain – at least not directly – the glass ceiling effect that we see in the data. Instead, we find that the main driver of the gender difference in log wage profiles is that men experience higher wage gains than women do both as "switchers" and as "stayers." Finally, in Section V, we conclude.

# 2 Data

In this paper, we use matched worker-firm data from Statistics Sweden to examine wage dynamics for a high-earner subset of the Swedish population. Our starting

point is the LOUISE database, which contains information from several registers for the Swedish population aged 16-75. This database includes annual labor earnings, age, education, region of residence, marital status, etc. From these data, we define our population of interest, namely, men and women born in the years 1960-70 who pursued a business or economics degree and completed at least 3 years of university education, which is equivalent to a US Bachelor's degree. This population comprises 30,735 individuals, almost equally split between men and women. We follow these individuals annually from age 25-45 over the period 1985-2013. For most of our analysis, we use the subpopulation of business and economics graduates who worked exclusively in the private sector (69.4% of the men versus 56.3% of the women).

The LOUISE data are combined with the multi-generational register, which links all children to their biological parents, including information on the birth month and year of individuals' children. This means that we observe all child births for our population.

Information on full-time equivalent monthly earnings (comparable to hourly wage rates) is obtained by matching the Wage Structure Statistics to our data. The Wage Structure Statistics dataset is based on individual earnings and contracted work hours during the survey month (typically September). These data are available only for individuals who had some working hours during the month of the survey. Full-time equivalent earnings are available for all public-sector employees with positive hours in the survey month. In the private sector, these data are available for all workers in firms with at least 500 employees. For firms with fewer than 500 employees, a stratified sample (based on industry and on firm size) is used each year. As a result, in any given year, approximately 50% of private-sector workers are included in our database. The analysis that we present below uses weights to adjust for the fact that full-time equivalent earnings are not observed for all private sector workers.

By using the Employment and Firm Registers we can match the Swedish population of workers to the population of firms. This means that we can construct measures of average individual characteristics at the firm level for our population of individuals. In addition, the Firm Register includes information on sector and industry classifications. By using the Firm Register we can also observe individual switches from one firm to another.

Finally, we also match the Parental Leave Register (compiled by the Social Insurance Agency) to our data. This register includes information on the number of paid parental leave days taken for the entire population of parents who use parental leave. As such, we can also create average parental leave take-out measures at the firm level

for our population.

# 3 Wage and Earnings Dynamics

We begin by examining patterns of wage and earnings growth for the men and women in our dataset. Figure 1 plots average log wages (monthly full-time equivalent earnings) for men and women by age, restricting our attention to those who worked exclusively in the private sector between the ages of 25 and 45. The starting wages of these men and women are close to the same (the male-female gap at ages 25 and 26 is less than .02 log points), but by age 45, there is a gender gap of approximately 0.25 log points.

**Figure 1:** Log Wage Profiles

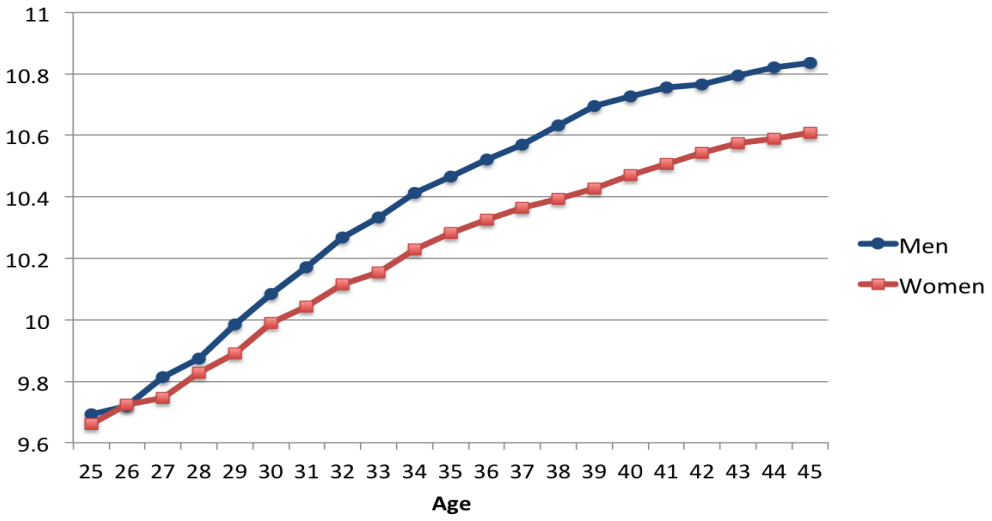


How do wages and earnings change as these individuals accumulate work experience, and how are they affected by the timing of the birth of a first child? To address these questions, we estimate regressions of the form

$$y_{it} = \beta_0 + \sum_{j=1}^{4} \beta_j X_{it}^j + \sum_k \gamma_k d_{it}^k + \varepsilon_{it} \tag{1}$$

Here $y_{it}$ is an outcome variable – either log wage or log earnings, which depends on a quartic in $X_{it}$, age (actually, age - 25), and on a set of dummies corresponding to years relative to first birth, $d_{it}^k$.[1] Specifically, $d_{it}^k = 1$ if in year $t$ individual $i$ is $k$ years

[1] We use potential experience (age - 25) rather than actual experience because this makes the

**Figure 2:** Log Wage: Age Effects

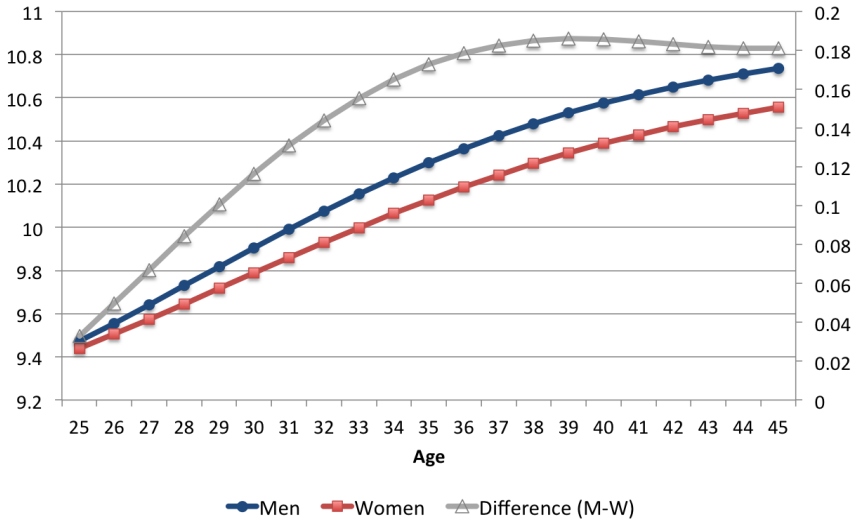


away from the birth of his or her first child, and zero otherwise. We estimate the coefficients $\gamma_k$ for $k = -8$ to $k = 15$; that is, we estimate the increase or decrease associated with being 8 years prior to the first birth through the effect associated with being 15 years after first birth. The left-out category, i.e., those for whom $d_{it}^k = 0$ for all $k$ and all $t$, is thus the individuals who were childless over our observation period.[2] We estimate these regressions separately for men and for women, and we follow these individuals from age 25 up to a maximum age of 45 years. In contrast to some of the literature on gender wage gaps that focuses solely on the effect of becoming a parent, we present estimates for both age effects and time-to-first-birth effects. This allows us to see how the gender gap evolves as individuals age holding constant the time to the birth of their first child and, in addition, to hold the age effects constant and see the dynamic effects of first becoming a parent.

Figure 2 shows the effects of age on log wages for men and for women, controlling for time-to-first-birth effects. Men's and women's starting wages are approximately the same, but by age 30, even abstracting from the effect of becoming a parent, there is a substantial gender gap in wages. Reading off the vertical axis on the left-hand side of the figure, men's wages have increased by approximately 0.43 log points by

description of lifecycle profiles clearer. In practice, our results are the same whether we use potential or actual experience.

[2]We observe some individuals – those who were relatively old when their first child was born – more than 8 years before first birth. We include the effect of being 9 years before first birth, 10 years before first birth, etc. in $d_{it}^{-8}$ so that $\hat{1}^3\ \gamma_{-8}$ for these individuals captures the effect of being 8 or more years before the birth of a first child. Similarly, $\gamma_{15}$ captures our estimate of the effect associated with being 15 years or more past the time of first birth.

**Figure 3:** Log Wage: Time to First Birth Effects

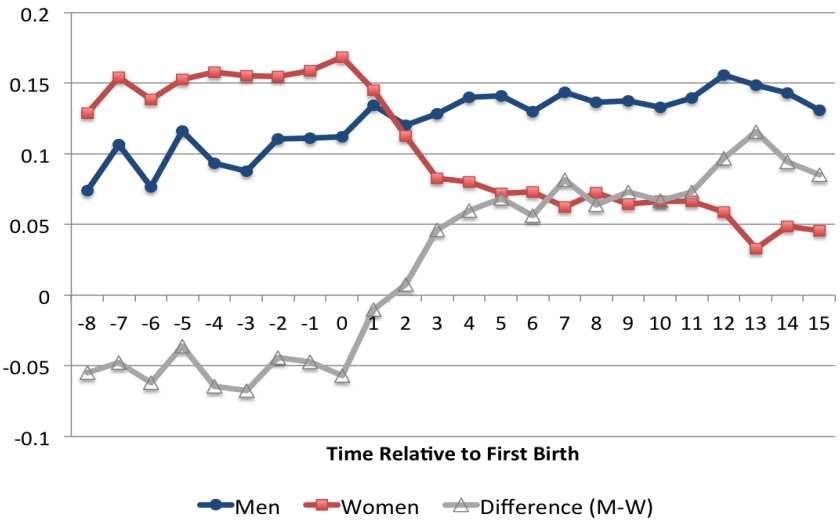


age 30 whereas women's wages have increased by approximately 0.35 log points, both relative to their starting wages. The gender difference in wages ascribable to age up to age 30 (0.117 log points), is read off the right-hand side axis. The gender gap in wages associated with age peaks at age 39 at 0.186 log points, declining somewhat thereafter to a difference of 0.181 log points by age 45.

Figure 3 shows the effects of the timing of first birth while controlling for age. Consider, for example, a man whose first child is born when he is aged 30. At age 25, this man has five years before his first child is born. This is associated with a wage gain of approximately 0.12 log points relative to a 25-year-old man who remains childless over the sample period. Similarly, this man has approximately 0.09 additional log points at age 26 (the effect of being four years prior to first birth), etc. This pattern of wage gains associated with the impending birth continues up to age 30, when the child is born, and then continues to increase up to a peak of about 0.16 log points by age 42. The pattern of time-to-first-birth effects for a woman is approximately the same as the pattern for the corresponding man up to the time the first child is born. Indeed, if anything, the wage gain associated with impending motherhood is clearly larger than the corresponding gain associated with impending fatherhood. The wage gain for women peaks at the time of birth at 0.17 log points. The pattern of positive time-to-first-birth effects prior to childbirth likely reflects the incentives of the parental leave system. Prior to the birth of their first child, women have substantial incentive to accumulate human capital and move to higher paid jobs, as

**Table 1:** Log Wage Gap

| Age | Age Effect | Time-to-First Birth Effect First birth at Age 30 | Time-to-First Birth Effect First birth at Age 35 |
|---|---|---|---|
| 25 | 0.032* | -0.036** | -0.055*** |
| | (0.018) | (0.015) | (0.014) |
| 30 | 0.117*** | -0.057*** | -0.036** |
| | (0.008) | (0.010) | (0.015) |
| 35 | 0.173*** | 0.069*** | -0.057*** |
| | (0.006) | (0.011) | (0.010) |
| 40 | 0.186*** | 0.067*** | 0.069*** |
| | (0.005) | (0.011) | (0.011) |
| 45 | 0.181*** | 0.085*** | 0.067*** |
| | (0.005) | (0.011) | (0.011) |

*Significant at 1% level. **Significant at 5% level. ***Significant at 1% level.

compensation during parental leave depends on income earned in the prior year.[3]

Once the first child is born, the time-to-first-birth effects for men and women diverge significantly. There is a slightly increasing time-to-first-birth effect for men up to 12 years after the first birth. For women, the time-to-first-birth effect drops sharply in the first three years (from 0.17 to 0.08 log points) and then decreases at a slower rate. Controlling for age, the wage of a man whose first child was born 12 years earlier is approximately 0.16 log points higher than the wage of a childless man. In contrast, the wage of a woman whose first child was born 12 years earlier is approximately 0.05 log points higher than the wage of a childless woman.

To get a clearer sense of the relative importance of age versus time-to-first-birth effects, consider a man and a woman, each having a first child at age 30. At age 25 we expect the man's wage to be approximately the same as the woman's. This reflects the sum of the difference in male versus female[4] age effects at age 25 (approximately .032 log points in the man's favor) and the difference in male versus female time-to-first-birth effects at age 25 (approximately .036 log points in the woman's favor). This is summarized in Table 1. However, by age 30 – the age when the first child is born – there is an approximately 0.06 log point gap in the man's favor, entirely due to the age effect, which is stronger than the time-to-first-birth effect. By age 35 – 5 years after the birth of the first child – the gap widens to approximately 0.24 log points,

[3]Of course, an alternative possibility is that, in this population, women who have children are positively selected.

[4]The age effects shown in Table 1 include the difference between the estimated intercept in the male log wage regression (9.473) and the corresponding estimated intercept in the female log wage regression (9.440). Thus, the "age effect" at age 25 is simply the difference in estimated intercepts.

**Table 2:** Log Earnings Gap

| Age | Age Effect | Time-to-First Birth Effect First birth at Age 30 | Time-to-First Birth Effect First birth at Age 35 |
|---|---|---|---|
| 25 | 0.048* | -0.255*** | -0.248** |
| | (0.025) | (0.022) | (0.011) |
| 30 | 0.330*** | 0.510*** | -0.255*** |
| | (0.010) | (0.015) | (0.014) |
| 35 | 0.336*** | 0.471*** | 0.510*** |
| | (0.007) | (0.014) | (0.014) |
| 40 | 0.243*** | 0.227*** | 0.471*** |
| | (0.007) | (0.015) | (0.014) |
| 45 | 0.176*** | 0.171*** | 0.227*** |
| | (0.007) | (0.018) | (0.014) |

*Significant at 1% level. **Significant at 5% level. ***Significant at 1% level.

and the gap continues to widen up to and past age 40, peaking at age 43, when the gap is approximately 0.30 log points. In short, there is a substantial gap associated with age that opens up between men and women up to the age of first birth. Once the man and woman each have their first child, the effect of having that child adds to the size of the gender gap. The pattern is similar for men and women who have their first child at age 35. Again, the age effect accounts for most of the gender gap up to the age at first birth. After age 35, the increase in the gender gap is driven by both the age and the time-to-first-birth effects.

Carrying out a similar exercise for the gender gap in log earnings reveals some interesting differences. Table 2 shows the gender gap in log earnings for (i) men and women who are childless over our observation period (this is captured by the age effect 1 ), (ii) men and women who have a first child at age 30, and (iii) men and women who have a first child at age 35. Holding time to first birth constant, age has a strong positive effect on earnings, which increases up to the mid 30's and declines somewhat thereafter. Regarding time-to- first-birth effects, these are in the woman's favor up to the year of first birth, but once the first child is born, a very large gender gap in earnings emerges and then persists through age 45. Consider, for example, men and women who have a first child at age 30. At age 30, when the first child is born, we observe difference in earnings of approximately 0.84 log points (= 0.330 + 0.510). The gender gap in wages for these men and women is approximately 0.06 log points (= 0.117 - 0.057). That is, as we would expect given that mothers take more parental leave than fathers do, most of the gender gap in earnings at age 30 can be ascribed to the gender difference in hours worked. As the men and women whose first

**Figure 4:** Hours: Age Effects

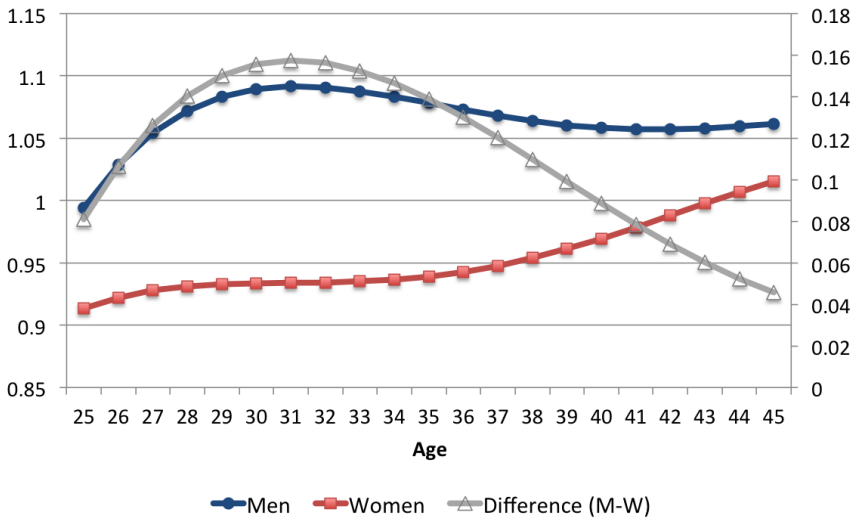


**Figure 5:** Hours: Time to First Birth Effects

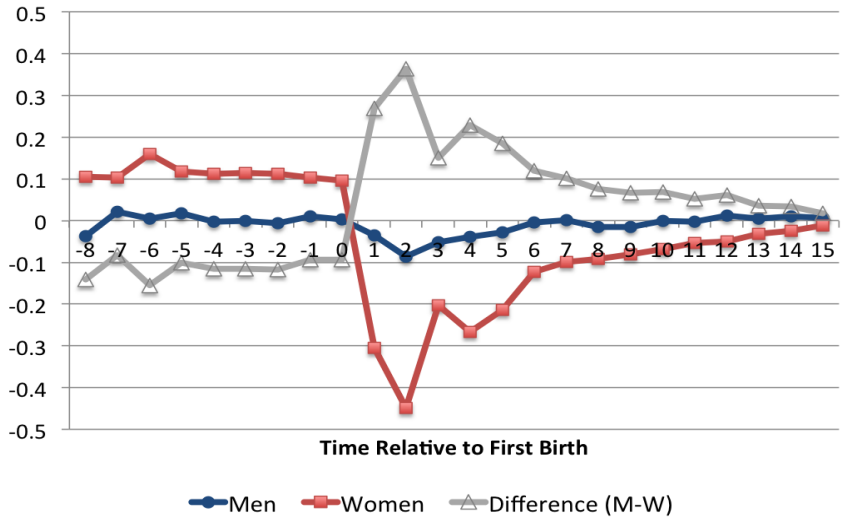


child was born when they were 30 get older, a substantial earnings gap remains, but the difference gets smaller. This reduction in the gender earnings gap after the year of first birth is not due to a reduction in the gender wage gap. On the contrary, the gender wage gap increases slightly for these individuals. Rather, the reduction in the gender earnings gap is due to a convergence in the hours worked by these men and women. The pattern for men and women who have their first child at 35 is similar with the exception that the large age effect on the earnings gap at age 30 means that men already have higher earnings even before the birth of the first child.

The effects of age and time to first birth on hours[5] worked can be seen in Figures 4 and 5. Holding time to first birth constant, men work longer hours than women do in their late 20's to middle 30's, on the order of 10-15% more each year. Then as the men and women enter their 40's - again, holding the effect of time to first birth constant - there is a convergence in hours worked. The difference is then on the order of 5-10% each year. In other words, even among the men and women who remain childless, men work longer hours when they are young, and hours are more similar only once the women are reaching the end of their potential childbearing years. The effect of time to first birth on the gender difference in hours worked is substantial, as expected. Holding age constant, prior to the first birth, women tend to work a bit more than men. Once their first child is born, however, women experience a precipitous drop in hours worked while the effect on men's hours is negligible. Only 15 years after the first birth do women's hours return to levels similar to men's.

Clearly, these differences in hours worked affect the gender gap in earnings. To some extent, these gender differences in hours reflect different choices by gender related to parenthood. The gender gap in earnings is also the result of differences in wages over careers that may not reflect differences in individual choices but may reflect different opportunities offered by firms. We turn to the gender gap in career wage profiles in the next section.

# 4 Accounting for Gender Differences in Career Wage Dynamics

The men and women we observe start their careers with essentially the same wages, but their career paths quickly diverge. This gender difference is present even before having children complicates the picture. Then, once men and women become parents, the difference in career paths between men and women becomes even more pronounced. We focus on three factors that can account for these patterns. First, do men tend to work in "better" firms than women do? Second do men increase their wages via firm-to-firm mobility at a faster rate than women do? Third, within similar firms, do men experience faster wage growth than women do?

[5]Our measure of hours worked is a proxy constructed by dividing earnings by wages. To the extent that men are more likely to receive bonuses, this proxy may overstate men's hours relative to women's.

**Figure 6:** Firm Size: Age Effects

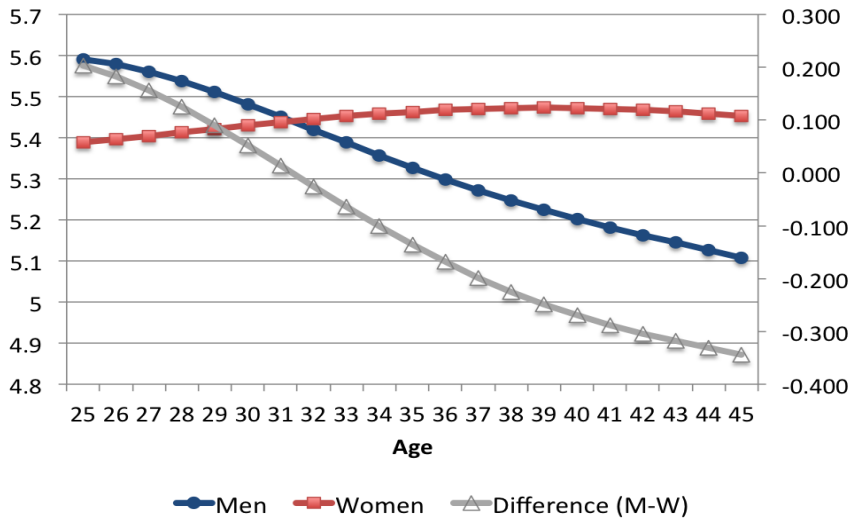


## 4.1 Firm Characteristics

A potentially important factor in accounting for gender differences in career dynamics is that men and women sort into different types of firms. The firm characteristics that seem to differ the most between the men and women in our data are (i) firm size (women tend to work for larger firms than men do) and (ii) the earnings of high-skill male co-workers (men are more likely to work in high-paying firms, that is, firms in which other high-skill men have high earnings). The firm size effect is particularly striking. Figures 6 and 7 show the effects of age and time to first birth on the log number of employees at the individual's firm. Holding time to first birth constant, women tend to start at smaller firms than men do, but as they age, they tend to move toward somewhat larger firms while at the same time, men are switching to smaller firms. By the time they reach their 40's – again, holding time to first birth constant – women are on average working in firms that are on the order of 0.3 log points larger. The time-to-first-birth effect on firm size for women is non-monotonic. Impending parenthood is associated with working in a larger firm, and once a woman has her first child, the tendency to work in a larger firm is even stronger. This tendency reverses 3 years after the birth of the first child when women start moving toward smaller firms. For men, the time to first birth effect on firm size is negligible. Holding age constant, the peak difference in time to birth effect on firm size occurs 4 years after the first birth when women work at firms that are on average 30 percent larger.

The tendency for men to sort into smaller firms while women tend towards larger employers is consistent with the argument that women sort into (or are tracked into)

**Figure 7:** Firm Size: Time to First Birth Effects

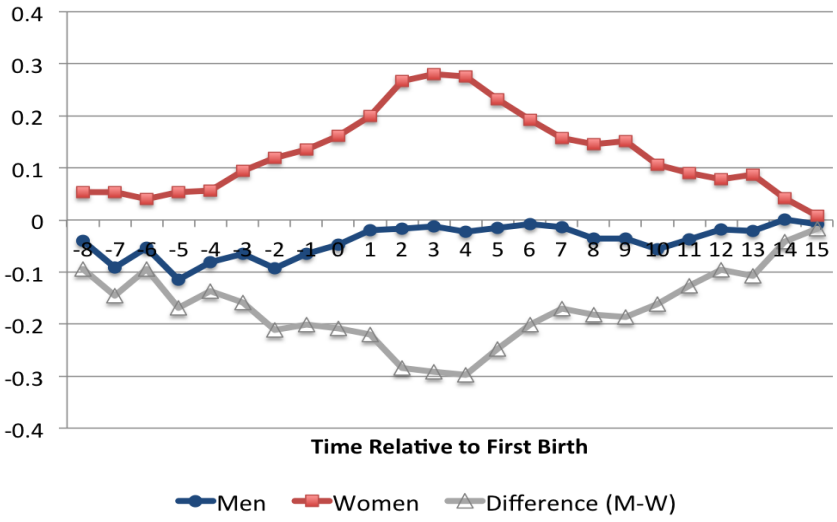


**Table 3:** Firm Size and Average Hours Worked, Relative to Full-Time (Full-Time = 1)

| Number of employees | Men | Women |
|---|---|---|
| Fewer than 20 | 1.04 | 0.91 |
| 21-50 | 1.08 | 0.94 |
| 51-100 | 1.05 | 0.92 |
| 101-250 | 1.07 | 0.92 |
| 251-500 | 1.08 | 0.92 |
| 501-1,000 | 1.03 | 0.90 |
| 1,001-2,500 | 1.04 | 0.89 |
| 2,501-5,000 | 1.03 | 0.89 |
| 5,001-1,0000 | 1.00 | 0.88 |
| More than 10,000 | 1.01 | 0.90 |

work environments that do not require longer and less regular hours of work. For example, Goldin (2014), Bronson (2015) and Cortes and Pan (2016) make this argument in the context of the US labor market. Large firms have relatively many highly qualified workers to choose among when someone needs to work late to meet a deadline, and large firms have more ways to fill a position that is left vacant when an important worker goes on parental leave. The evidence on firm size and hours worked in Table 3 suggests that both men and women in our population work somewhat higher hours at small and medium-sized firms, particularly firms with more than 20 and fewer than 500 employees.

Men tend to work in higher-paying firms than women do. Of course, there are many ways to define "higher-paying." In Figures 8 and 9, we show the associations of

**Figure 8:** Detrended Average Log Earnings of High-Skill Male Coworkers: Age Effects

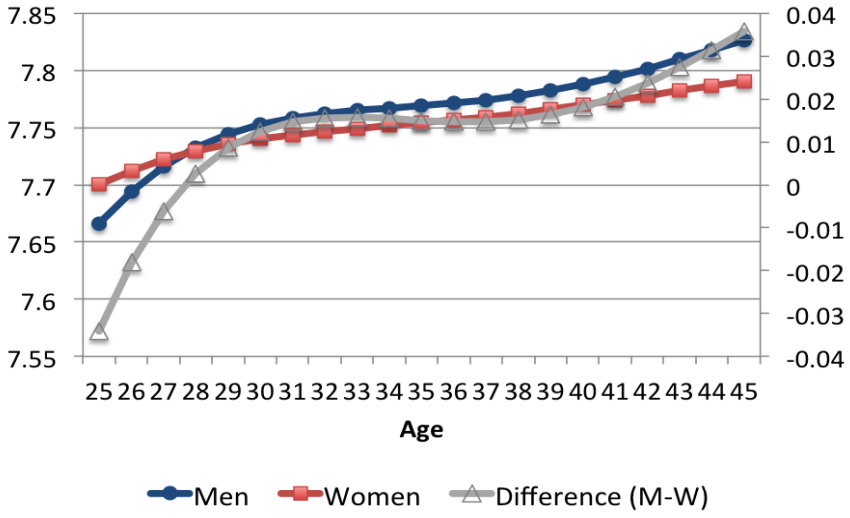


**Figure 9:** Detrended Average Log Earnings of High-Skill Male Coworkers: Time to First Birth Effects

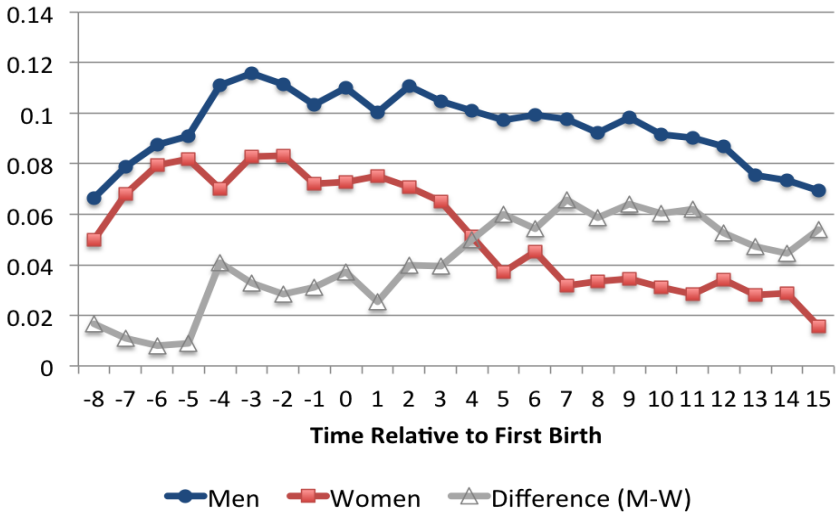


age and time to first birth with detrended average log earnings of other high-skill men (i.e., male coworkers with any university education) in the firms that employ the men and women in our data. We remove the time trend to control for the fact that the average log earnings of a worker's male high-skill coworkers would tend to increase over time even if the worker's age and time relative to first birth were held fixed. These figures show a tendency for both men and women to move towards higher-paying firms as they age, more so for men than for women, as well as a positive effect associated with becoming a parent.

The association between gender and firm size is potentially important for accounting for gender differences in wages. As illustrated in Figure 10, there is a tendency towards higher pay in relatively small firms. Specifically, relatively small firms are the ones in which high-skill male co-workers tend to have the highest earnings. Men are more likely than women to work in firms where other high-skill men earn high wages and work long hours. The high earnings among the high-skill males in these firms is driven to some extent by relatively long hours of work, as indicated in Table 3.

**Figure 10:** Firm Size Distributions by Rank of Mean Earnings of Male Workers

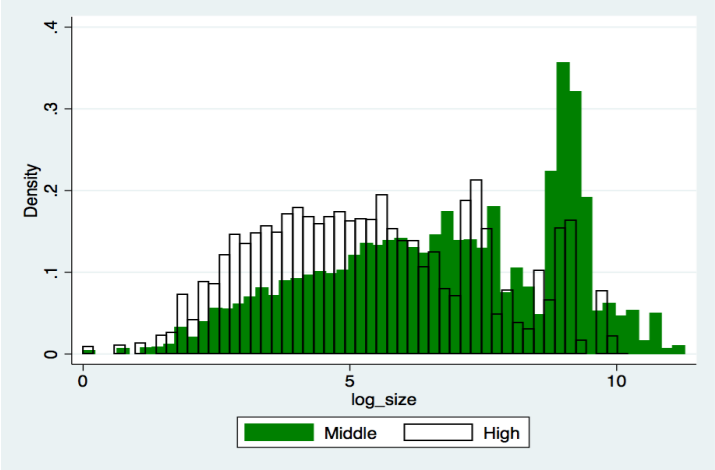


Note: Firm rank is determined by ordering all firms according to the mean earnings of their high-skill male employees. "High"-ranked firms are those that fall in the top third of this ranking; "middle"-ranked firms are those that fall between the 33rd and 66th percentiles.

To carry out a more systematic analysis of the effect of firm characteristics on the gender gap in log wages we use the following wage decompositions. Specifically, we regress male and female log wages separately on a set of individual and firm characteristics. We run a separate regression at each age, from age 25 to 45, to allow for the most flexible possible specification. Then we compare the observed gender gap in log wages at each age to (i) the gap we would expect to observe if the firms where women work had the same characteristics as do the firms where men work, (ii) the gap we would expect to observe if women had the same individual characteristics as men do as well as the same firm characteristics, and (iii) the gap we would expect to observe if women's characteristics – both their own characteristics and the characteristics of the firms where they work – were rewarded with men's coefficients.

We use the following individual and firm characteristics. At the individual level, we use (i) a quadratic in years of experience, (ii) dummies for years of completed university education, (iii) the fraction of full time hours worked in the survey month,

**Figure 11:** Wage Gap Decompositions

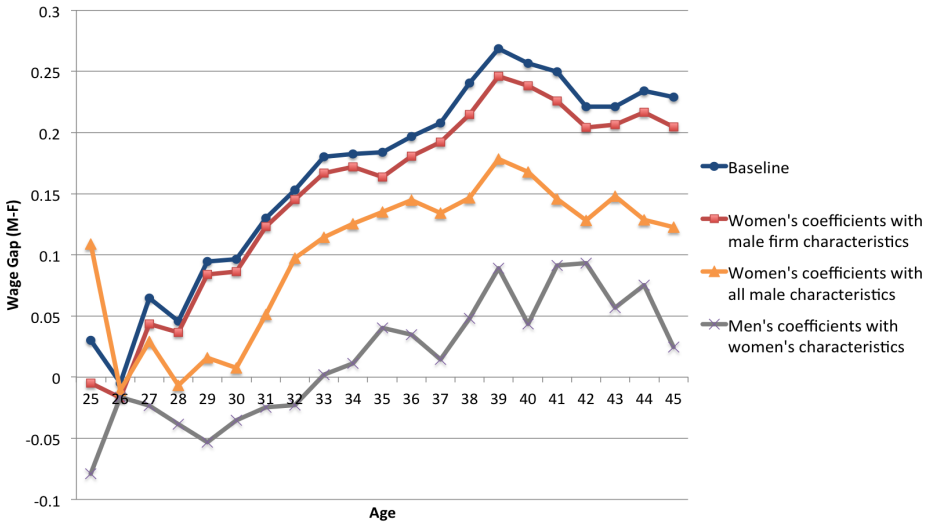


and (iv) the share of previous years worked part time. At the firm level, we use (i) industry, (ii) average earnings of (other) high-skill males employed in the firm, (iii) the ratio of earnings at the 90th percentile to earnings at the 50th percentile among high-skill workers in the firm, (iv) firm size dummies (one dummy for fewer than 100 employees; a second dummy for between 100 and 500 employees), (v) log firm size, (vi) the share of females in high-skill employment in the firm, and (vii) the average number of parental leave days taken by men employed in the firm in the first two years after their first child was born.

The results from our decompositions are shown in Figure 11. The log wage profile labeled "Baseline" shows the observed difference between the average log wage paid to men and the corresponding average paid to women at each age. The profile labeled "Women's Coefficients with Male Firm Characteristics" gives the log wage gap by age that we would expect to observe if women were employed in the same types of firms as men are; that is, if women worked in smaller firms, in higher-paying firms, etc. There is a gap between the baseline profile and this second profile, but the difference is not particularly large. The differences in firm characteristics account for at most 2.0-2.5 percentage points, or about 9% of the wage gap between men and women. That is, men do tend to work in "better" firms than women do, but the fraction of the gender log wage gap accounted for by firm differences is relatively small. Gender differences in individual characteristics account for a larger fraction of the gap. The profile labeled "Women's Coefficients with All Male Characteristics" shows the log wage gap by age that we would expect to observe if women had the same individual

characteristics as men do, e.g., the same share of years worked part time, as well as the same firm characteristics as men. Gender differences in individual characteristics account for more of the gender gap in log wages, especially after men and women reach their early 30's. The increasing importance of individual characteristics starting in this age range reflects the impact of having worked part time in the past on wages, and as men and women enter their late 30's and early 40's, women are more likely to have previously worked part time. Finally, the profile labeled "Men's Coefficients with Women's Characteristics" gives the gender log wage gap that we would expect to observe if women's characteristics – both their individual characteristics and their firm characteristics – were rewarded in the same way as men's characteristics are. This counterfactual profile accounts for almost all of the observed gender log wage gap through the early 30's; after the early 30's, gender differences in characteristics – especially a history of part-time work – also plays a significant role in accounting for the gender difference in log wages.

## 4.2 The Gender Gap in Wage Gains for Switches and for Stayers

What role does firm-to-firm mobility play in accounting for gender differences in career paths? As shown in Figure 12, the fraction of men and women who move from one firm to another strongly declines with age. Holding time-to-first-birth constant, men are somewhat more likely to switch firms early in their careers (up to age 35 or so), and about equally likely to switch firms as women after age 40.[6] Figure 13 shows time-to-first-birth effects on firm-to-firm mobility. Holding age constant, up to one year prior to the birth of the first child, women are more likely to switch from one firm to another. Then, in the first four years after the first birth, women are substantially less likely to switch firms than men are, reflecting, presumably, the anchoring effects of parental leave benefits. The relationship between time to first birth and the propensity to switch firms is smaller for men, although we similarly observe that the propensity to switch firms after the first birth declines for men.

[6]As in our previous analysis, we restrict our attention to men and women who spend their careers (at least the 20 years of their careers that we can observe) entirely in the private sector, and therefore exclude individuals who switch between sectors from the analysis. Recent work by Kleven et al. (2016) using Danish data has emphasized movement of women into the public sector after childbirth. In our population of high-skilled men and women, we also observe such movements, though at a lower rate, and not entirely concentrated around the time of first birth. At age 25, women are about 6.2 percentage points more likely to be in the public sector (30.6% of women, compared to 24.3% of men). By age 45, this gender difference is around 11.1 percentage points, and is mostly driven by men's higher propensity to switch out of the public sector in their late twenties.

**Figure 12:** Probability of Switching Firms: Age Effects

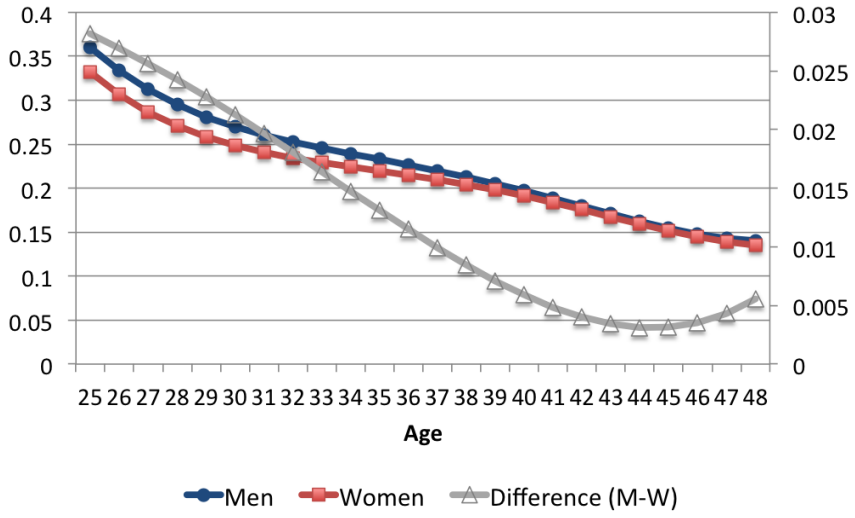


**Figure 13:** Probability of Switching Firms: Time to First Birth Effects

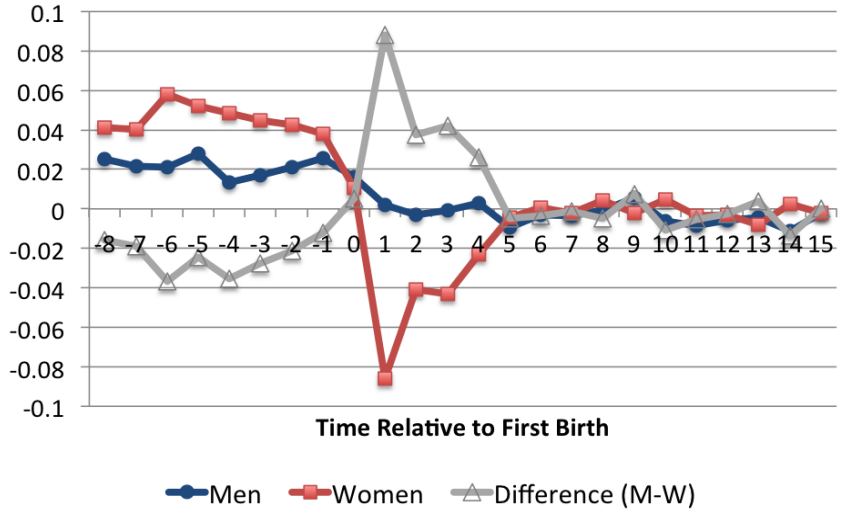


The overall effect of age and time to first birth on firm-to-firm mobility is shown in Figure 14. The propensity to switch firms declines strongly with age, and up to approximately age 40, men are somewhat more likely to move from one firm to another than are women.

Men and women are both most likely to move from one firm to another early in their careers and it is this early career firm-to-firm mobility that is most associated with substantial wage gains. This is true for both men and women, but the early-career wage gains associated with firm-to-firm mobility are substantially higher for

**Figure 14:** Share Switching Firms

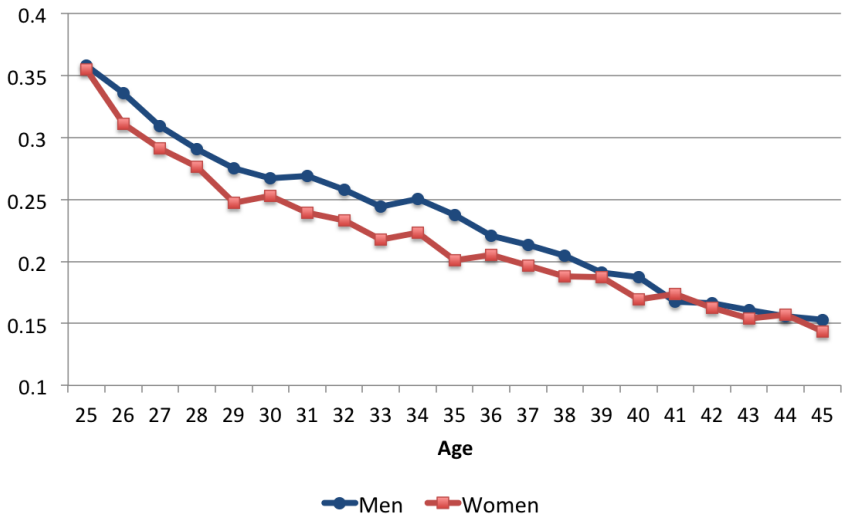


men than they are for women. Figures 15 and 16 show the age and time-to-first-birth effects on wage gains among those who switch firms ("switchers") at a particular age. So, for example, holding time-to-first-birth effects constant, women who switch from one firm to another at age 30 average a wage gain of approximately 0.12 log points while men who switch at that age average a gain of approximately 0.16 log points, around 33% more. In other words, men are more likely to switch firms when they are young, and, when they do switch, they tend to receive a higher payoff from doing so. Once men and women reach their late-30's, the wage gains associated with mobility are much smaller, and, holding time-to-first-birth effects constant, it is only after age 40 that the wage gains for women who switch firms become roughly the same as the corresponding wage gains for male switchers. These patterns we observe are consistent with the tendency of men and women to move to higher-paying firms as they age. There are also time-to-first-birth effects on the wage gains for switchers (Figure 16), but these are relatively small and similar for men and women, though women appear to be slightly more likely than men to switch to a lower-wage job in the first 4 years after childbirth.

To recap, there is a substantial gender difference in movement between firms. Most firm-to-firm mobility occurs early in workers' careers, and in early career the probability with which men move from one firm to another exceeds the corresponding probability for women. Further, the payoff associated with firm-to-firm mobility is higher for men than for women, especially early in their careers when most of the movement is taking place.

**Figure 15:** Wage Gains for Switchers: Age Effects

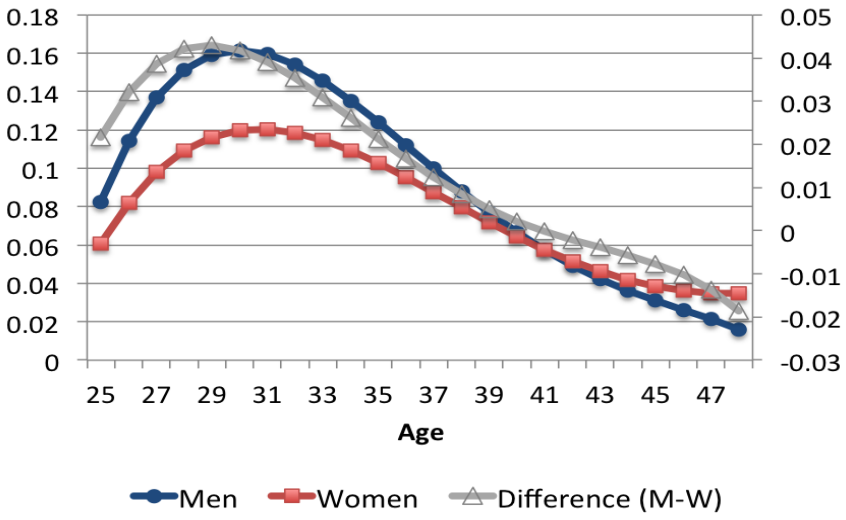


**Figure 16:** Wage Gains for Switchers: Time to First Birth Effects

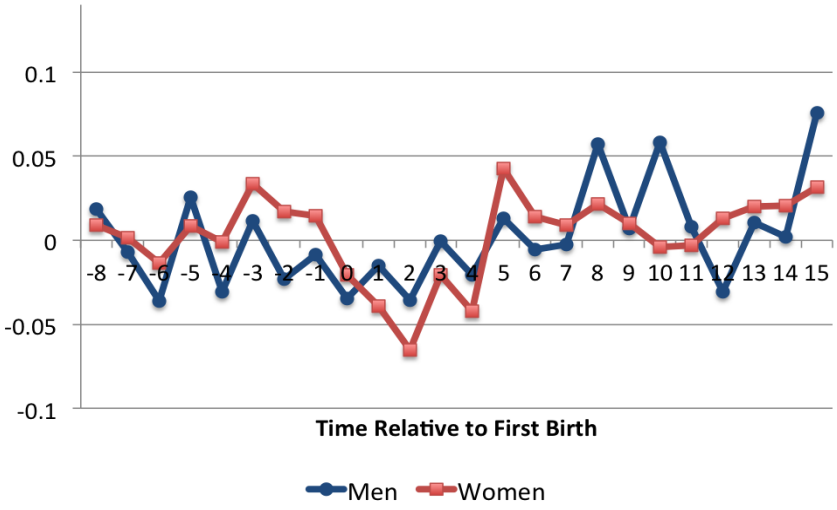


Of course, the men and women who don't move from one firm to another ("stayers") also experience different rates of wage growth. Table 4 compares, by age group, the average wage gains for workers who stay with the same employer from one year to the next with those who switch employers, which include both the age and time-to-first-birth effects. As is the case with workers who move from one firm to another, the log wage gains realized by men who stay with the same employer are on average larger than the corresponding gains for women, and more pronounced early on in the lifecycle. Overall, the differences in wage gains between young men and women are

**Table 4:** Annual Log Wage Gains, by Gender and Mobility Status

| | Switched Employers in Past Year | | | Did Not Switch Employers | | |
|---|---|---|---|---|---|---|
| | Men | Women | Diff. (M-W) | Men | Women | Diff. (M-W) |
| Age 25-30 | 0.135 | 0.092 | 0.042*** | 0.088 | 0.069 | 0.019*** |
| | (0.010) | (0.006) | | (0.002) | (0.002) | |
| Age 30-35 | 0.122 | 0.105 | 0.016** | 0.079 | 0.066 | 0.014*** |
| | (0.005) | (0.004) | | (0.002) | (0.001) | |
| Age 35-40 | 0.086 | 0.080 | 0.006 | 0.054 | 0.046 | 0.008*** |
| | (0.006) | (0.005) | | (0.001) | (0.001) | |
| Age 40-45 | 0.051 | 0.050 | 0.002 | 0.032 | 0.031 | 0.001 |
| | (0.006) | (0.003) | | (0.001) | (0.001) | |

**Significant at 5% level. ***Significant at 1% level.

higher among those who switched firms in the last year, than among those who did not.

A simple decomposition can illustrate how differences in mobility and in the log wage gains associated with switching and with staying affect the gender difference in log wage profiles. For each gender, the average log wage gain from one year to the next can be written as

$$\Delta \ln w = P[Switch] \cdot (\Delta \ln w | Switch) + (1 - P[Switch]) \cdot (\Delta \ln w | Stay)$$

As a first step, suppose that at each age, the probability of firm-to-firm mobility were the same for men and women. Specifically, suppose that at each age, women switched from one firm to another with the observed men's frequency. Since the male switching probabilities are uniformly higher than women's up to age 40 (Figure 14), one might expect that adjusting the share of women switching firms at each age would affect the gender difference in the age profile of log wage gains. In fact, this has essentially no effect whatsoever. This is because (1) the difference in switching rates is relatively small and (2) the wage growth that women realize as switchers differs relatively little from the wage growth they realize as stayers. As a result, the difference between the log wage gains we observe in the data and the log wage gains we would expect to observe if women switch employers at the rate men do is essentially zero. Instead, what drives the gender difference in log wage gains is simply the fact that men realize higher wage gains whether they switch firms or not.

Figure 17 shows a gender difference, which is substantial when cumulated, in log wage gains between ages 26 and 45. The observed gender gap in annual log wage gains is shown in the Baseline profile. The majority of the difference is ascribable to

**Figure 17:** Decomposing Log Wage Gains

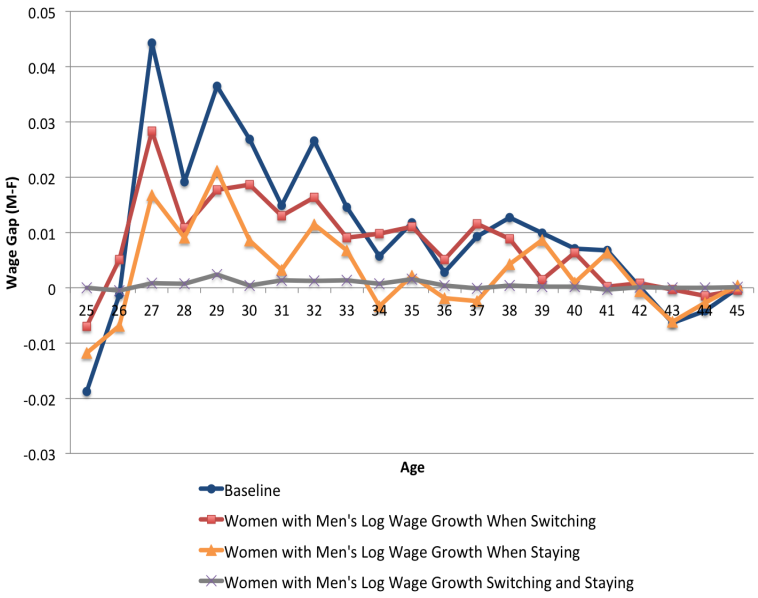


the fact that when men and women switch firms, the returns to mobility are higher for men, but it is also the case that the log wage gains realized by male stayers exceed the corresponding log wage gains for women. This can be seen in the fact that virtually all of the gender gap is accounted for if we leave the share of women switching firms constant but give women the wage gains that men receive switching or staying. We emphasize that the patterns observed in Figure 17 – and the absence of a direct effect of gender differences in the share of switchers in the decomposition – do not mean that gender differences in firm-to-firm mobility are unimportant for understanding the different log wage profiles that we observe for men and women, as there may be indirect effects. For example, if it is easier for men to generate outside offers or if men are less constrained in the offers they are able to consider, then men can be pickier about which offers they accept. That is, the ability to generate and accept outside offers affects the expected log wage gains attached to accepted offers. Similarly, having better outside options affects a worker's bargaining position even if he or she doesn't exercise the option to switch firms. Teasing out these equilibrium effects is a subject for future research.

# 5 Conclusion

There is a significant gender gap in log wages among highly skilled Swedish men and women. To investigate this phenomenon, we use matched employee-employer data on men and women who have completed university education in business and economics. Our data are for men and women born between 1960 and 1970 and we follow these individuals for 20 years. While men and women start with virtually the same wages, a gender gap is apparent before these individuals have children. Once these men and women become parents, the gender gap widens further.

Since we have matched data, we are able to examine whether firm characteristics and/or mobility between firms can account for the different wage paths of men and women. We find that men tend to work in smaller firms and in firms that are higher paying (in the sense that the other highly skilled male workers there are highly paid) than those where women work. Nonetheless when we carry out a series of wage decompositions, we find that these firm characteristics account only for a small part of the gender difference in wage profiles. We then add an adjustment for gender differences in individual characteristics including part-time work and can account for a bit more of the difference in wage profiles. When we have women receive men's returns for the women's characteristics (both individual and firm), we are able to account for most of the difference in wage profiles. Finally, we look at differences in mobility. Firm-to-firm mobility is higher for men than it is for women, especially in their careers, and the log wage gains realized by workers who switch from one firm to another are greater than those realized by stayers. However, when we carry out a simple decomposition in which we assign men's switching probabilities to women (and vice versa), the gender difference in mobility explains virtually none of the observed difference in annual log wage gains. Instead, what matters is that men's wages tend to grow faster than women's from one year to the next whether or not they switch firms.

# 6 References


Abowd, John, Francis Kramarz, and David Margolis (1999), "High Wage Workers and High Wage Firms," *Econometrica*, 67(2), 251-333.

Albrecht, James, Anders Björklund and Susan Vroman (2003), "Is There a Glass Ceiling over Sweden?" *Journal of Labor Economics*, 21(1), 145-77.

Albrecht, James, Peter Thoursie and Susan Vroman (2015), "Parental Leave and the Glass Ceiling in Sweden," *Research in Labor Economics*, 41, 89-114.

Arulampalam, Wiji, Booth, Alison, and Mark Bryan (2007), "Is There a Glass Ceiling over Europe? Exploring the Gender Pay Gap across the Wage Distribution," *Industrial and Labor Relations Review,* 60(2), 163-86.

Bertrand, Marianne, Claudia Goldin and Lawrence Katz (2010), "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors," *American Economics Journal: Applied Economics*, 2, 228-55.

Blau, Francine and Lawrence Kahn (2016), "The Gender Wage Gap: Extent, Trends and Explanations," *Journal of Economic Literature*, forthcoming.

Bronson, Mary Ann (2015), "Degrees are Forever: Marriage, Educational Investment, and Lifecycle Labor Decisions of Men and Women," working paper.

Burdett, Ken and Melvyn Coles (2003), "Equilibrium Wage-Tenure Contracts," *Econometrica*, 71(5), 1377-1404.

Burdett, Kenneth and Dale Mortensen (1998), "Wage Differentials, Employer Size and Unemployment," *International Economic Review*, 39, 257-73.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline (2016), "Firms and Labor Market Inequality: Evidence and Some Theory," working paper.

Cortes, Patricia and Jessica Pan (2016), "When Time Binds: Returns to Working Long Hours and the Gender Wage Gap Among the Highly Skilled," working paper.

Gibbons, Robert and Michael Waldman (1999), "Careers in Organizations: Theory and Evidence," Chapter 36 in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*, Volume 3, Elsevier Science, 2374-2437.

Goldin, Claudia (2014), "A Gender Convergence: Its Last Chapter," *American Economic Review*, 104(4), 1091-1119.

Kleven, Henrik Jacobsen, Camille Landais, and Jakob Egholt Søgaard (2016), "Parenthood and the Gender Gap: Evidence from Denmark," working paper.

Postel-Vinay, Fabien and Jean-Marc Robin (2002), "Equilibrium Wage Dispersion with Worker and Firm Heterogeneity," *Econometrica*, 70(6), 2295-2350.

Wood, Robert, Mary Corcoran and Paul Courant (1993), "Pay Differences Among the Highly Paid: The Male-Female Earnings Gap in Lawyers Salaries," *Journal of Labor Economics*, 11(3), 417-41.