WORKING PAPER 2018:9

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

Jim Albrecht
Mary Ann Bronson
Peter Skogman Thoursie
Susan Vroman

IFAU
INSTITUTE FOR EVALUATION OF LABOUR MARKET AND EDUCATION POLICY

The Institute for Evaluation of Labour Market and Education Policy (IFAU) is a research institute under the Swedish Ministry of Employment, situated in Uppsala.

IFAU's objective is to promote, support and carry out scientific evaluations. The assignment includes: the effects of labour market and educational policies, studies of the functioning of the labour market and the labour market effects of social insurance policies. IFAU shall also disseminate its results so that they become accessible to different interested parties in Sweden and abroad.

Papers published in the Working Paper Series should, according to the IFAU policy, have been discussed at seminars held at IFAU and at least one other academic forum, and have been read by one external and one internal referee. They need not, however, have undergone the standard scrutiny for publication in a scientific journal. The purpose of the Working Paper Series is to provide a factual basis for public policy and the public policy discussion.

More information about IFAU and the institute's publications can be found on the website www.ifau.se

ISSN 1651-1166

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

by

Jim Albrecht[♣], Mary Ann Bronson[♣], Peter Skogman Thoursie[†] and Susan Vroman[♣]

April 25, 2018

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 up to 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 have similar wages and earnings at the start of their careers, but their career paths diverge substantially as they age. These men and women also have substantial differences in wage paths associated with becoming a parent. We look at whether firm effects account for the differences we observe between women's and men's wage profiles. We document differences between the firms where men work and those where women work. However, a wage decomposition suggests that these differences in firm characteristics play only a small 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 and women both exhibit greater mobility early in their careers, but there is little gender difference in this firm-to-firm mobility. We find that the main driver of the gender difference in log wage profiles are that men experience higher wage gains than women do both as "switchers" and as "stayers".

Keywords: Wages, Earnings, Gender gaps, Firms
JEL-codes: J16, J31

[♣] Department of Economics, Georgetown University.
[†] Department of Economics, Stockholm University.

Table of contents

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 10 log points while the corresponding gap at the 90th percentile is approximately 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 up to 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)). Perhaps wages grow more rapidly in firms where men work than they do in firms where women are 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, e.g., gender differences in human capital accumulation, tracking, and/or discrimination.

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 3, 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 4, 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 characteristics 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. Gender differences in mobility rates do 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".[1] Finally, in Section 5, 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 for up to 20 years starting at age 25 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 childbirths 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,

[1] A related paper is Reshid (2017) who uses the same data source as we do and study the role of job and upward occupational mobility in explaining gender differences early career wage growth of college graduates. Results show that job mobility and upward occupational mobility significantly contribute to the early career wage growth but that the returns to both types of mobility are significantly lower for females. Further analyses show that female penalty in returns to job is to a large extent related to the timing of childbirth.

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.

In our analysis, we use the sub-population of business and economics graduates who worked exclusively in the private sector (69.4% of the men versus 56.3% of the women). One motivation for this restriction is that we want to look at similar men and women and dropping the subset of business and economics majors who spend part or all of their careers in the public sector (arguably a subsample with somewhat different motivations) serves that purpose. Another motivation is that we want to consider "firm effects", and these are ill defined for public-sector employees. One drawback to this approach is that by focusing only on private-sector workers, we may miss potentially important margins of adjustment, such as switching between sectors especially around the time of first birth. In the Appendix, we replicate our key results for the entire population of business and economics majors born in the years 1960-70. Interestingly, movements to and from the public sector, including around the time of first birth, are less common for business and economics majors than they are for highly educated men and women in general. Overall, our findings are similar when we include workers who spend all or part of their careers in the public sector.

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 2 log points), but by age 45, there is a gender gap of approximately 25 log points.

Figure 1 Log Wage Profiles

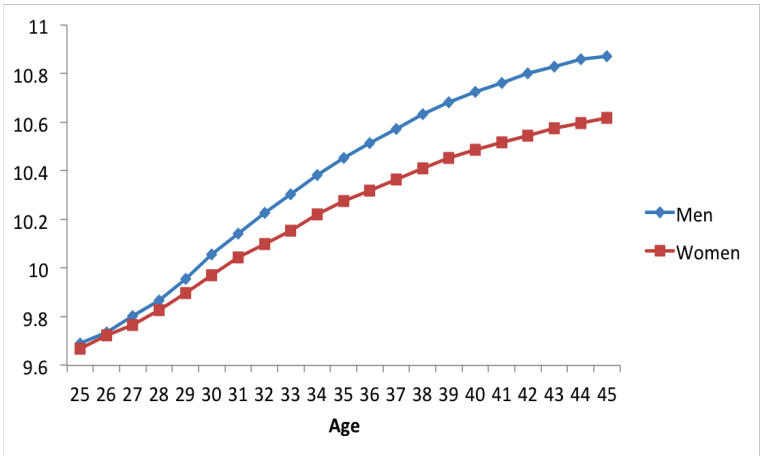

How do wages and earnings change as these individuals accumulate work experience, and how are they affected by the timing of the birth of a first child? To address these questions, we estimate regressions of the form

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

Here y_{it} is an outcome variable – either log wage or log earnings, for individual i in year t – which depends on a quartic in X_{it}, potential experience (age minus 25), and on a set of dummies corresponding to years relative to first birth, d_{it}^k, and a year fixed effect, λ_t.[2] Specifically, $d_{it}^k = 1$ if in year t individual i is k years away from the birth of his or her first child, and zero otherwise. We estimate the coefficients γ_k for $k = -5$ to $k = 10$; that is, we estimate the increase or decrease associated with being 5 years prior to the first

[2] We use potential experience (age-25) rather than actual experience because this makes the description of lifecycle profiles clearer. In practice, our results are the same whether we use potential or actual experience.

birth through the effect associated with being 10 years after first birth.[3] We take 1986 as our base year, so β_0 gives the expected earnings in 1986 net of age effects for an individual who remains childless throughout our sample period. We estimate the regressions separately for men and for women, and we follow individuals from age 25 up to a maximum age of 45 years. Note that the coefficients γ_k map out the dynamic effects associated with first becoming a parent, holding the age effects constant; the estimates of β_j, in turn, capture how labor market outcomes evolve as individuals age, holding constant the time to the birth of their first child.

The specification we use in equation (1) is closely related to those presented in Angelov et al. (2016) and in Kleven et al. (2017). As we do, Angelov et al. (2016) use Swedish administrative data. The unit of observation in their analysis is couples with children, and they estimate how the within-couple gender gap (their dependent variable) varies with time to first birth. Kleven et al. (2017) use Danish administrative data. They analyze a balanced panel of parents whom they observe every year between $k = -5$ (5 years before the first birth) and $k = 10$. They estimate separate regressions for men and for women with dummies for time to first birth, age, and calendar time (year). Other than the set of outcome variables we focus on, the main differences between our approach and the ones used in Angelov et al. (2016) and Kleven et al. (2017) are: (i) we use a high-earner subset of the population, namely, university graduates with business and economics majors, (ii) we limit our analysis to workers who spend their careers (up to age 45) in the private sector (as discussed above), and (iii) we include non-parents in our analysis. We choose to include non-parents because the variation in their labor market experience contributes to the estimation of the age effects. In the Appendix, we show that our results are robust with respect to the exclusion of childless men and women.

Throughout the analysis, we present not only the estimated time-to-first-birth ("parenthood") effects as in Angelov et al. (2016) and Kleven et al. (2017), but also estimated age effects, which capture all other gender differences in lifecycle labor market outcomes not directly associated with time relative to first birth. We incorporate an

[3] We observe some individuals – those who were relatively old when their first child was born – more than 5 years before first birth. We include the effect of being 6 years before first birth, 7 years before first birth, etc. in d_{it}^{-5} so that γ_{-5} for these individuals captures the effect of being 5 or more years before the birth of a first child. Similarly, γ_{10} captures our estimate of the effect associated with being 10 years or more past the time of first birth. The left-out category, i.e., those for whom $d_{it}^{k} = 0$ for all k and all t, is thus individuals who were childless over our observation period.

analysis of these age effects for two reasons. First, this allows us to assess the importance of dynamic factors associated with childbirth, relative to all other gender differences in lifecycle outcomes that are not directly related to parenthood. Second, a number of gender differences in labor market outcomes not related to the specific timing of the first birth are still of interest to researchers.[4]

We begin our analysis by documenting the main gender differences in wages, earnings, and hours worked, over the lifecycle and relative to the first birth. Panel A of Figure 2 graphs log wages over the lifecycle for male and female business and economics majors, controlling for time-to-first-birth effects and year fixed effects. Men's and women's starting wages are approximately the same, but as these workers age, even abstracting from the effect of becoming a parent, men have steeper wage profiles than women. Panel B graphs the difference in the effect of potential experience on men's and women's log wages at each age and shows the 95% confidence bands. As Panel B shows, the gender gap in wages, net of birth effects, increases monotonically with age, at least up to age 43.

Figure 2 Log Wage: Estimated Age Effects

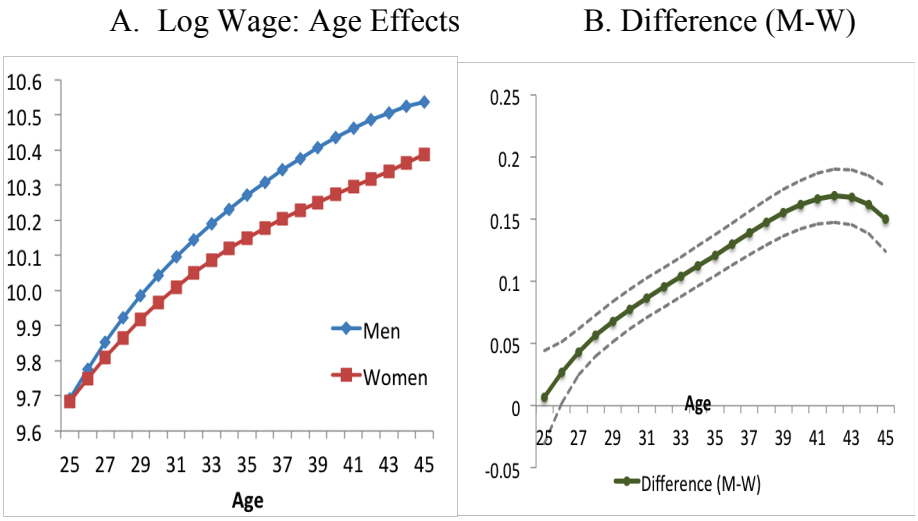

[4] For example, employers may discriminate against women of childbearing age, regardless of their fertility status or intention to have children; in addition, women may select into different types of firms than men from the start of their careers. Such differences would not be captured by parenthood effects since they are not necessarily related to the specific timing relative to first birth.

Figure 3 Log Wage: Estimated Time to First Birth Effects

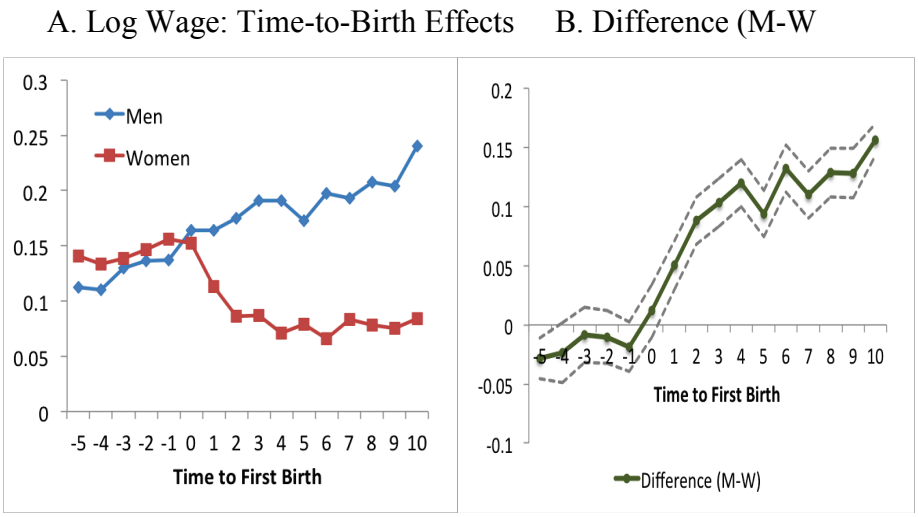

Figure 3, Panel A, shows the effects of the time to first birth on log wages, while controlling for age and year fixed effects. Specifically, Panel A graphs the estimated coefficients γ_k from equation (1). These time-to-first-birth effects illustrate two features in the data. First, they document how the wages of individuals who eventually become parents compare to those of individuals who remain childless over the entire sample period. Second, they capture the dynamic effects of parenthood on wages.

Consider, for example, a man whose first child is born when he is 30. At age 25, this man has five years before his first child is born. Panel A documents that, on average, his wage is approximately 11 log points higher than the wage of a 25-year-old man who never has children by age 45. Similarly, this man's wage is approximately 13 log points higher at age 27 – the effect of being three years prior to first birth. This pattern of wage gains associated with the impending birth continues up to age 30, when the child is born, and the man's wages continue to increase more or less monotonically throughout the years after the birth. 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, the wage gain associated with impending motherhood is slightly larger than the corresponding gain associated with impending fatherhood. The wage gain for women peaks a year before the first birth at approximately 16 log points. These positive time-to-first-birth coefficients prior to childbirth likely reflect a combination of two factors: positive selection into parenthood, as well as the incentives of the parental leave system.

Prior to the birth of their first child, women in particular have substantial incentive to accumulate human capital and to move to more highly paid jobs, as compensation during parental leave depends on income earned in the prior year.

Once the first child is born, the time-to-first-birth effects for men and women diverge significantly. Panel B graphs the gender difference in time-to-birth effects, sometimes referred to as the "motherhood penalty". As Panel B shows, the motherhood penalty increases from 0 in the year prior to first birth, to more than 15 log points ten years after first birth. This reflects the pattern in Panel A that for women, the time-to-first-birth effect drops sharply in the first four years after the first birth (from 15 to 7 log points) and stays at this lower level thereafter, while for men time-to-first-birth effects continue to increase.

Table 1 summarizes the relative importance of age versus time-to-first-birth effects. In Table 1 we consider a man and a woman, each having a first child at age 30, just below the median age at first birth in Sweden for this cohort. At age 25 the man's wage is on average 2.1 log points lower than the woman's, primarily reflecting the difference in male versus female time-to-first-birth effects at age 25 (approximately 2.8 log points in the woman's favor). However, by age 30 – the age when the first child is born – there is an approximately 9 log point gap in the man's favor, mostly due to the age effect (7.8 log wage points) and partly due to the time-to-first-birth effect (1.2 log wage points). By age 35 – 5 years after the birth of the first child – the gap widens to approximately 21.5 log points, and the gap continues to widen up to age 40. In short, there is a gap associated with age that opens up between men and women early in the lifecycle and increases until around age 40. Once the man and woman each have their first child, the effect of having had that child also becomes an important driver of the gender gap in wages. At age 45, the gender gap associated with parenthood accounts for 51% of the total gender wage differential. Note that the relationship between age and the gender wage gap may depend on several factors. For example, it may capture gender differences in firm-to-firm mobility and wage gains associated with such mobility. Gender differences in firm-to-firm mobility will be analysed in the next Section.

Table 1 Gender Log Wage Gap over the Lifecycle

Age	Lifecycle effect	Time-to-First Birth Effect†	Total Gender Gap† (M-W)
25	0.007	-0.028***	-0.021
	(0.020)	(0.009)	(0.018)
30	0.078***	0.012	0.090***
	(0.008)	(0.011)	(0.012)
35	0.121***	0.094***	0.215***
	(0.008)	(0.015)	(0.012)
40	0.162***	0.157***	0.319***
	(0.010)	(0.015)	(0.011)
45	0.151***	0.157***	0.308***
	(0.013)	(0.014)	(0.013)

† For individual whose first birth occurred at age 30. *** Significant at 1% level.

Table 2 carries out a similar exercise for the gender gap in log annual earnings, which captures the combined effect of gender differences in wages and hours worked.[5] Column 1 documents how the gender gap in earnings evolves with age, controlling for time-to-first-birth and year fixed effects. This gap is initially close to zero and insignificant, but rapidly becomes positive as individuals age, peaking in the mid-30s. Gender differences in time-to-first-birth effects (column 2) are in the woman's favor up to and including the year prior to first birth, reflecting both higher wages as well as women's higher hours worked prior to first birth. Thereafter, the time-to-first-birth effects favor men. In the year the first child is born, a large gender gap in earnings associated with parenthood emerges, of about 66 log points. This effect eventually decreases over time but remains substantial up through age 45. At age 30, when the first child is born, we observe a total difference in earnings of approximately 103 log points (=37.1 + 66.3). The gender gap in wages for these men and women is approximately 9 log points (= 7.8 + 1.2). 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.[6] As the men and women whose first children were 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

[5] Annual earnings are expressed in hundreds SEK. In order to be able to use log annual earnings we use log(earnings+0.1).

[6] Sweden has a relatively generous parental leave system allowing parents to take 480 days with parental leave benefits for each child. The replacement rate is approximately 80 percent up to ceiling for 390 of these days (60 days are reserved for each parent). The parental leave system provides a great degree of flexibility. During the first 18 months after childbirth, both parents have the right to be at home full time with job protection. The decrease in female earnings after childbirth is expected given the institutional setting and the fact that women take the lion part (around two-thirds) of all parental leave days.

wage gap. On the contrary, the gender wage gap increases 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.

Table 2 Gender Log Annual Earnings Gap over the Lifecycle

Age	Lifecycle effect	Time-to-First Birth Effect†	Total Gender Gap† (M-W)
25	-0.084	-0.243***	-0.327***
	(0.080)	(0.040)	(0.075)
30	0.371***	0.663***	1.03***
	(0.035)	(0.050)	(0.051)
35	0.396***	0.647***	1.04***
	(0.033)	(0.050)	(0.052)
40	0.291***	0.475***	0.766***
	(0.041)	(0.034)	(0.045)
45	0.317***	0.475***	0.793***
	(0.061)	(0.034)	(0.059)

† For individual whose first birth occurred at age 30. *** Significant at 1% level.

The effects of age and time to first birth on hours worked[7] can be seen in Figure 4 and 5. Holding time to first birth constant, Figure 4 shows that 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. In other words, even among the men and women who remain childless, men work longer hours. However, this difference is partly compensated for by the hours worked of individuals who eventually have children. Figure 5 shows that men who eventually have children work almost identical hours as those who remain childless. However, women who eventually have children work about 10% more hours prior to childbirth than their childless counterparts. Once their first child is born, however, women experience a large drop in hours worked. For men, the effect of parenthood on hours worked is relatively small, and observed primarily in the year after first birth, when men in Sweden typically take the two months of parental leave allocated specifically to fathers. About 10 years after the first birth, women's hours return to levels similar to men's.

[7] Our measure of hours worked is a proxy constructed by dividing earnings by wages. Specifically, we divide annual labor earnings by our wage measure, full-time equivalent monthly earnings (times twelve). Thus, for an individual who works full-time (40 hours per week), we would record hours worked equal to 1. For an individual working 20 hour per week, we would record hours worked equal to 0.5, etc. Note that annual earnings also include bonuses. To the extent that men are more likely to receive bonuses, this proxy may overstate men's hours relative to women's.

Figure 4 Hours: Estimated Age Effects

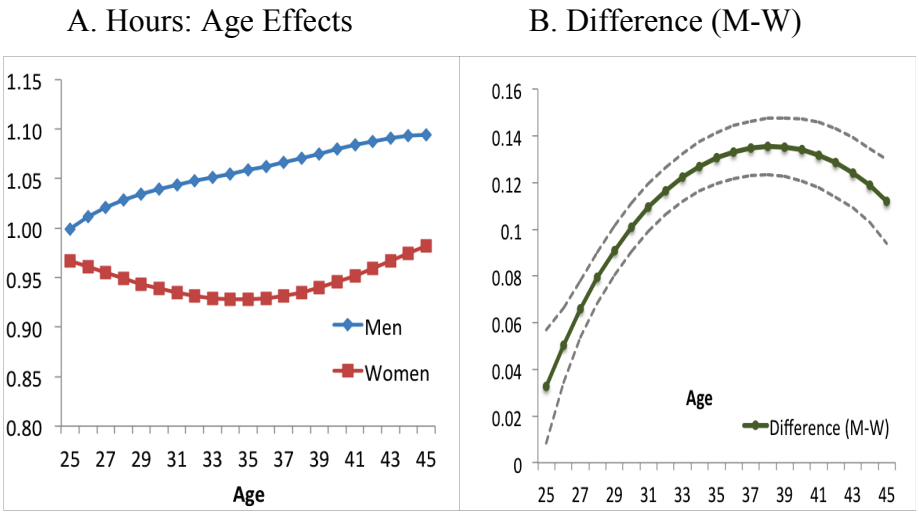

Figure 5 Hours: Estimated Time to First Birth Effects

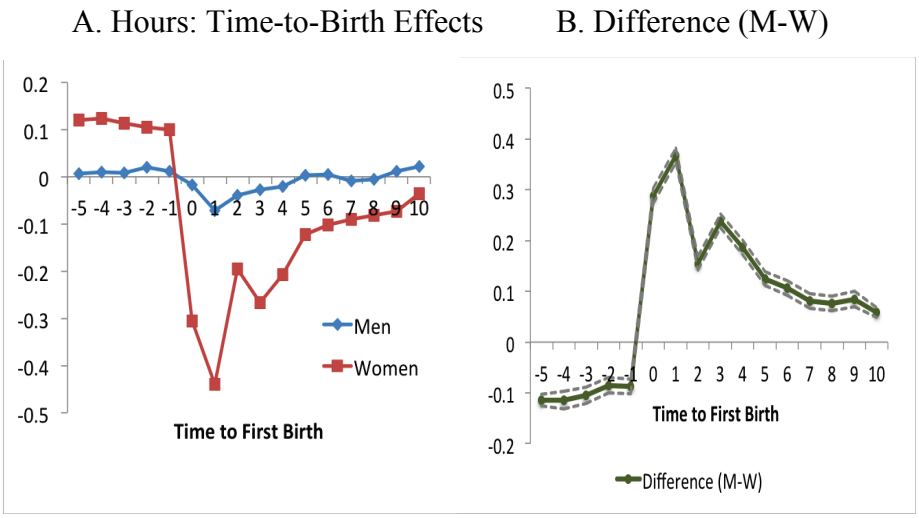

Clearly, differences in hours worked affect the gender gap in earnings directly. These gender differences in hours may reflect different choices by gender related to parenthood. The gender gap in earnings is also the result of differences in wages, which, as shown above, evolve differently for men and women, especially after the first birth. In the next section, we explore these gender wage differences over workers' careers in more detail.

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. 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, do men experience faster wage growth than women do within firms?

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 in larger firms than men do), (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), and (iii) "family friendliness" (e.g., women are more likely to work in firms in which part-time work is more common).

The firm size effect is striking. Figure 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 considerably larger than the firms where men work. The time-to-first-birth effect on firm size for women is non-monotonic. The tendency for a woman to work in a larger firm accelerates around the time she becomes a parent, reaching a peak two years after the first birth. After that, the time-to-first-birth effect for women falls, reaching a level close to zero ten years after becoming a parent. For men, the time-to-first-birth effect on firm size is small.

Figure 6 Firm Size: Estimated Age Effects

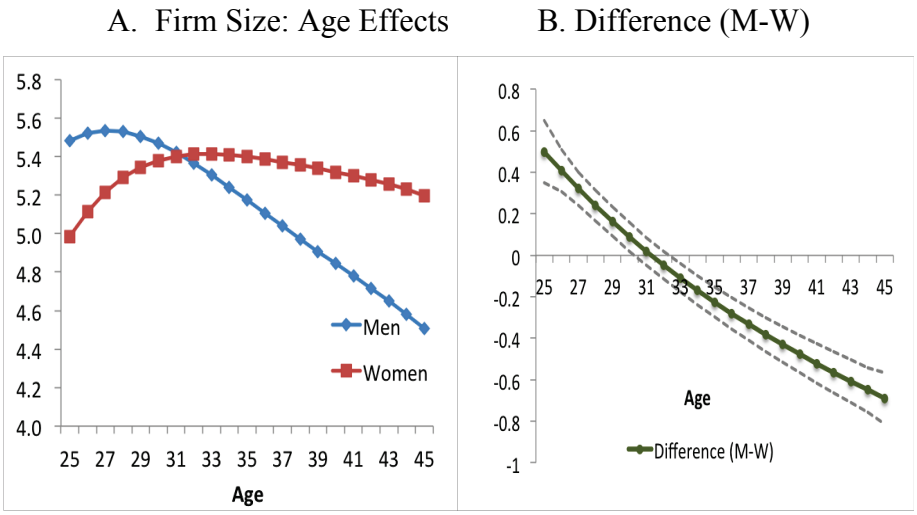

Figure 7 Firm Size: Estimated Time to First Birth Effects

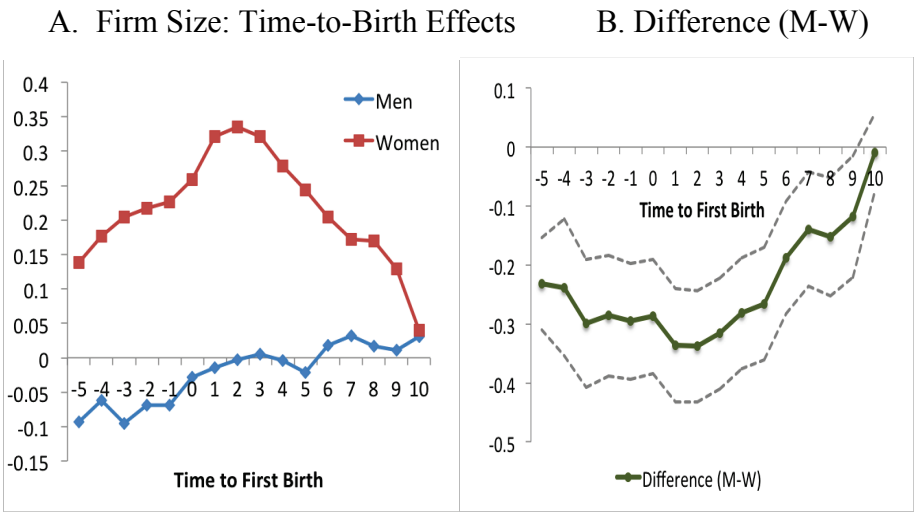

The tendency for men to sort into smaller firms while women tend towards larger employers is consistent with the argument that women sort into (or are tracked into) work environments that allow for more flexible work arrangements, including lower hours worked. For example, Goldin (2014), Bronson (2015) and Cortes and Pan (2017) 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 accommodate parental leave, i.e., to fill a position that is left vacant when an important worker is absent.

Figure 8 and 9 shows that men are more likely to work in higher-paying firms than women are, primarily among those who eventually have children. In Figure 8 and 9, we show the associations of age and time to first birth with average pay at the firm. We define average pay as the mean log earnings of high-skill male co-workers (i.e., those with three or more years of university education) in the firms that employ the men and women in our data. Figure 8 shows that men on average work at higher-paying firms when they first enter the labor market, although this gender difference disappears over time. Figure 9 shows that men and women who eventually have children work at higher-paying firms than their childless counterparts, in line with the findings in Figure 3 that these individuals are higher earners than the childless. Around the time of first birth, women switch to somewhat lower-paying firms, although the size of this change is moderate. The gender difference in firm pay increases from about 3% in the year before first birth to a maximum of about 5% six to nine years after first birth.

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

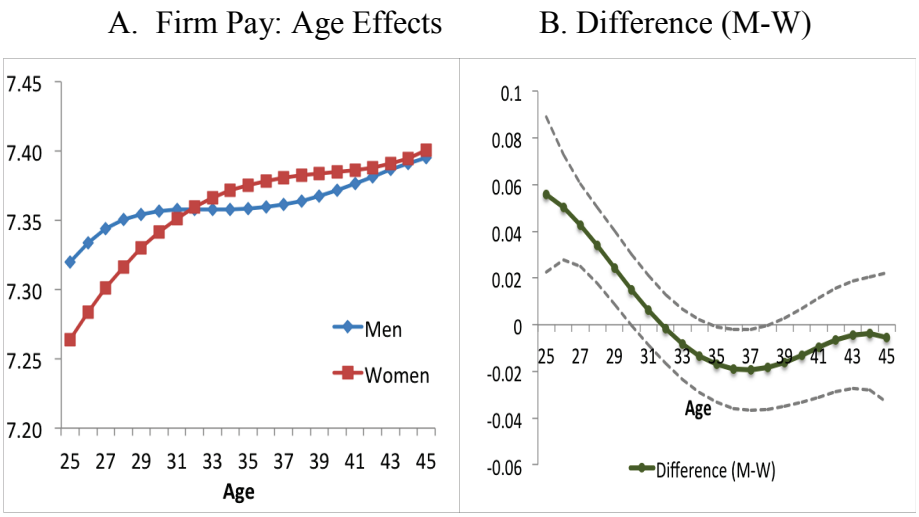

Figure 9 Average Log Earnings of High-Skill Male Coworkers: Estimated Time to First Birth Effects

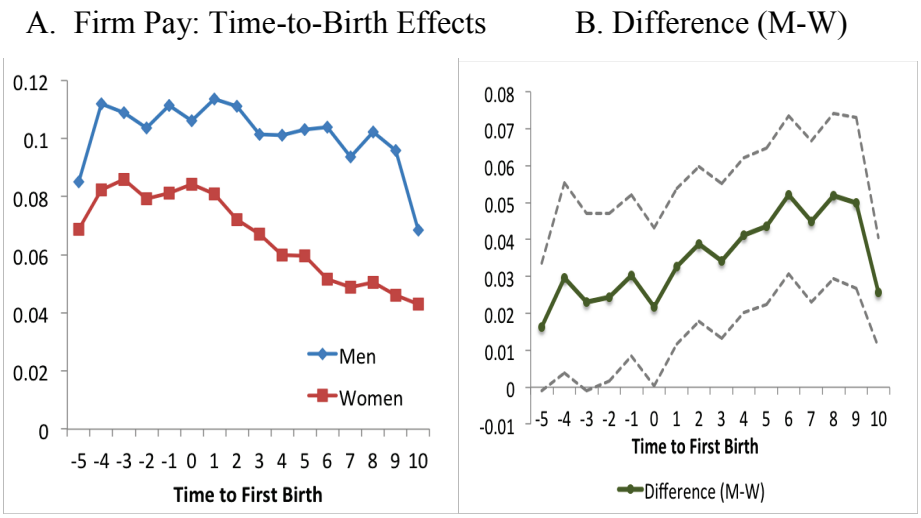

The association between gender and firm size is potentially important for accounting for gender differences in wages. Figure 10 overlays the distributions of firm size for high-paying firms and for firms that pay somewhat lower wages, i.e. for firms in the top and middle third of the distribution of firm pay. As illustrated in Figure 10, on average, pay is higher in relatively small firms. Specifically, relatively small firms are the ones in which high-skill male co-workers tend to have the highest earnings and men are more likely than women to work in firms where other high-skill men earn high wages.

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

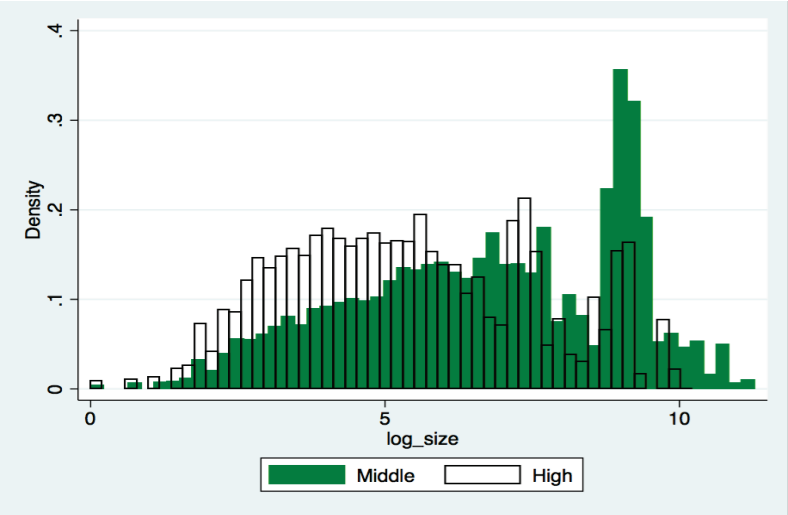

Note: Firm rank is determined by ordering all firms according to the mean earnings of their high-skill male employees. “High”-ranked firms are those that fall in the top third of this ranking; “middle”-ranked firms are those that fall between the 33rd and 66th percentiles.

Figure 11 Share of High-Skill Co-Workers Working Part Time: Estimated Age Effects

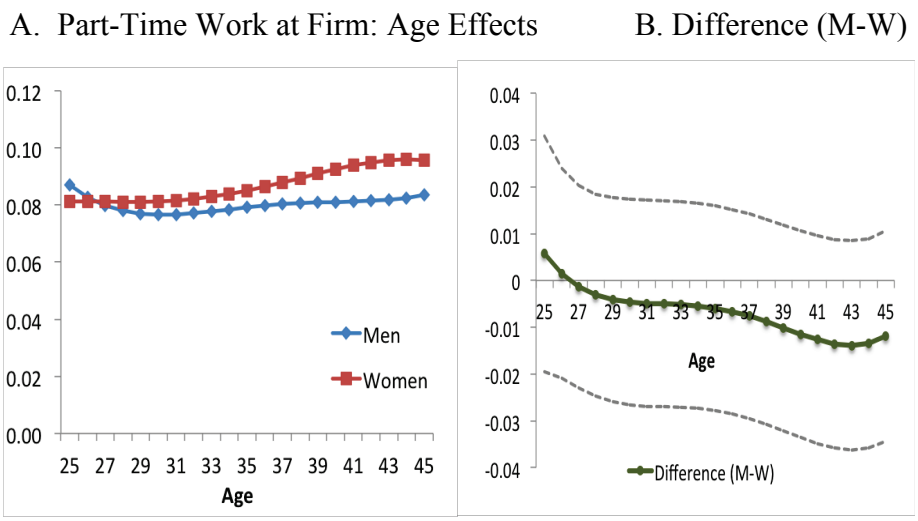

Figure 12 Share of High-Skill Co-Workers Working Part Time: Estimated Time to First Birth Effects

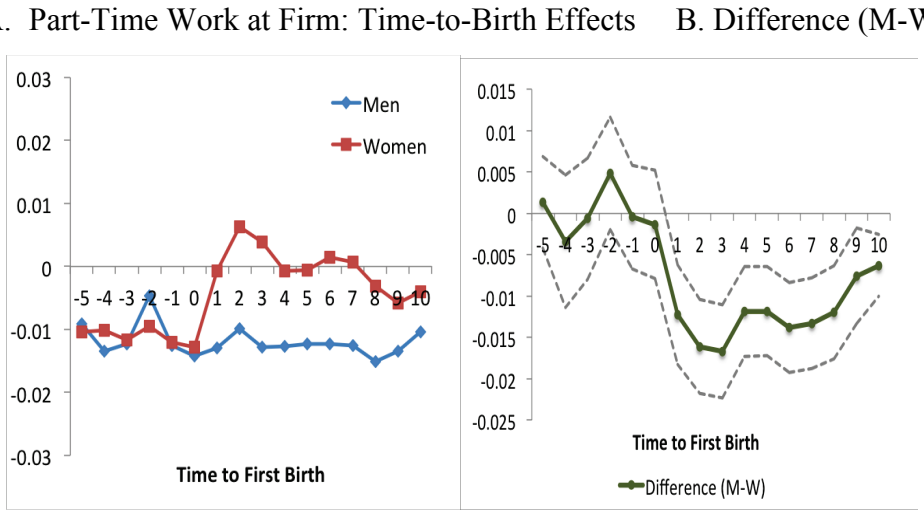

Finally, women are more likely than men are to work in "family-friendly" firms. Of course, family friendliness can be characterized in many ways. We focus on the fraction of high-skill workers in the firm who work part time (35 hours or less per week). Holding time to first birth constant, Figure 11 shows that women are slightly more likely to work at firms in which the share of part-time employment among high-skill workers is relatively high, although this difference is not statistically significant. The time-to-first-birth effects indicate that relative to childless individuals, men and women who eventually have children are more likely to work at firms with a lower share of part-time workers. After first birth, men continue to work at such firms, while women are more likely to switch to firms with more part-time workers, with a gender difference of about

1 to 1.5 percentage points. This a moderate to sizable effect, as the average part-time rate for high-skill workers at the firms employing individuals in our sample is about 11 percent.[8]

To carry out a more systematic analysis of the effect of firm characteristics on the gender gap in log wages we use the following Oaxaca-Blinder style 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 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 actual experience, (ii) dummies for years of completed university education, (iii) the fraction of full time hours worked in the survey month, 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 female among high-skill employees at 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 13. The log wage profile labeled "Baseline" shows the observed difference between the average log wage paid to men and the corresponding average log wage paid to women at each age. The profile labeled "Women's Coefficients with Male Firm Characteristics" gives the log wage gap

[8] Using Swedish administrative data, Karimi et al. (2016) document that women tend to move to more family-friendly firms around the time of first birth. Although this effect can be seen in our Figure 12, it is quite small. The difference between the pattern seen in our data and the one described in Karimi et al. (2016) is due to the fact that we focus on highly skilled men and women.

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 mid-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 40s, 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.

Figure 13 Wage Gap Decompositions

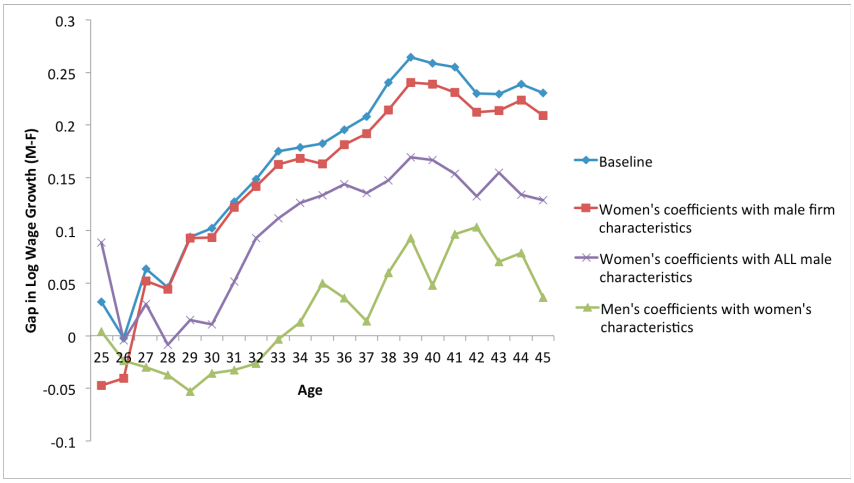

In sum, these decompositions indicate that differences between the firms where men and women work explain little of the gender log wage gap. Differences in labor market characteristics such as having worked part time explain somewhat more of the gap, but the main driver of the gender gap appears to be differences in returns to labor market characteristics. In the next section, we look more closely at this issue to see whether differences in firm-to-firm mobility can account for some of the gender log wage gap. That is, if women are switching firms at different rates (less opportunity to get a large wage increase) or bargaining less effectively when they do switch firms, this may account for some of the gender difference in the returns to labor market characteristics. Alternatively, women may simply experience lower wage growth within firms, even after controlling for part-time work and part-time history.

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? We begin our analysis of this question by first considering gender differences in the propensity to switch firms over the lifecycle, summarized in Figure 14 and 15. As shown in Figure 14, the fraction of men and women who move from one firm to another is relatively high in this population, and strongly declines with age, from more than 35% annually at age 25 to around 18% at age 45. Holding time-to-first-birth constant, men are slightly 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 39.[9] Figure 15 shows time-to-first-birth

[9] 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

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 than men. Then, from one year prior to three years after the first birth, women are substantially less likely to switch firms than men are, by as much as 10 percentage points in the year of first birth. For men, parenthood effects on the propensity to switch firms are negligible after the first birth.

Figure 14 Probability of Switching Firms: Estimated Age Effects

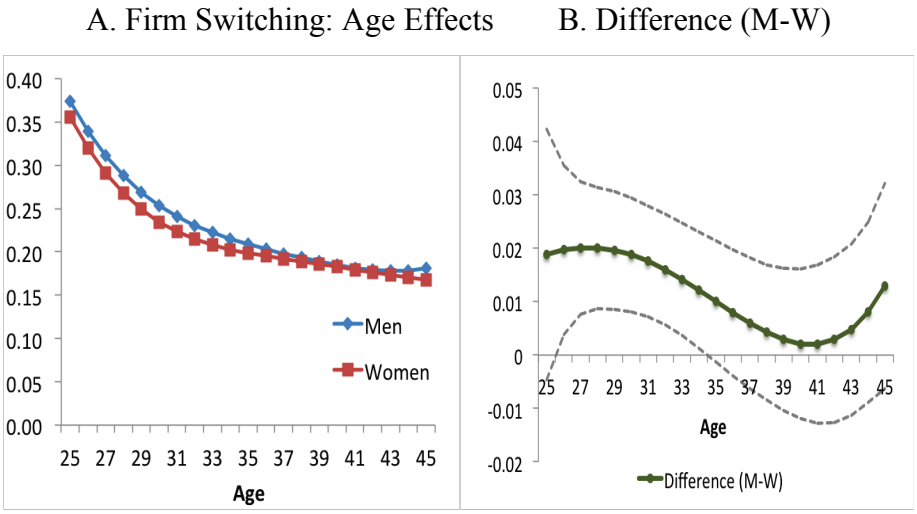

into the public sector after childbirth. In our population of high-skilled men and women, we also observe such movements, though at a lower rate, and not entirely concentrated around the time of first birth. At age 25, women are about 6.2 percentage points more likely to be in the public sector (30.6% 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. See Appendix.

Figure 15 Probability of Switching Firms: Estimated Time to First Birth Effects

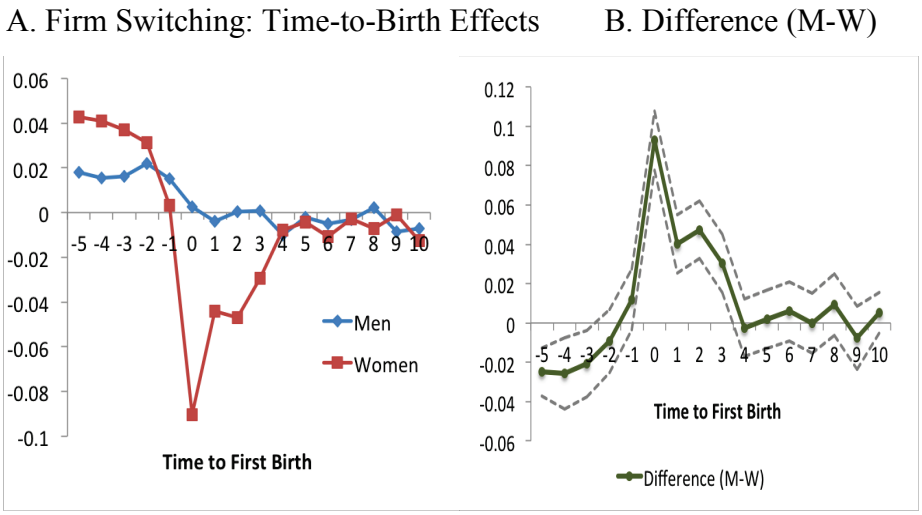

Men and women are both more likely to move from one firm to another early in their careers and it is this early career firm-to-firm mobility that is most associated with substantial wage gains. Figure 16 and 17 show the age and time-to-first-birth effects on log wage gains among those who switch firms ("switchers") at a particular age. Holding time-to-first-birth effects constant, the wage gains associated with movement from one firm to another for men are highest early on in the lifecycle, and similarly for women, with a gender difference that expands over time. Mobility gains fall with age for both men and women. For example, a man who switches firms at age 28 can expect a wage gain of around 11 log wage points while a man switching firms at age 45 can expect a wage gain of about 4 log wage points. For women, the corresponding expected wage gains are on average about two percentage points lower at all ages, although these differences are not statistically significant. The time-to-first-birth effects in Figure 17 on the wage gains for switchers are noisy and relatively small, but it is noteworthy that women who switch firms between the start of their careers and the first birth tend to receive wage gains while the effects of a switch in the years immediately after becoming a parent tend to be negative.

Figure 16 Wage Gains for Switchers: Estimated Age Effects

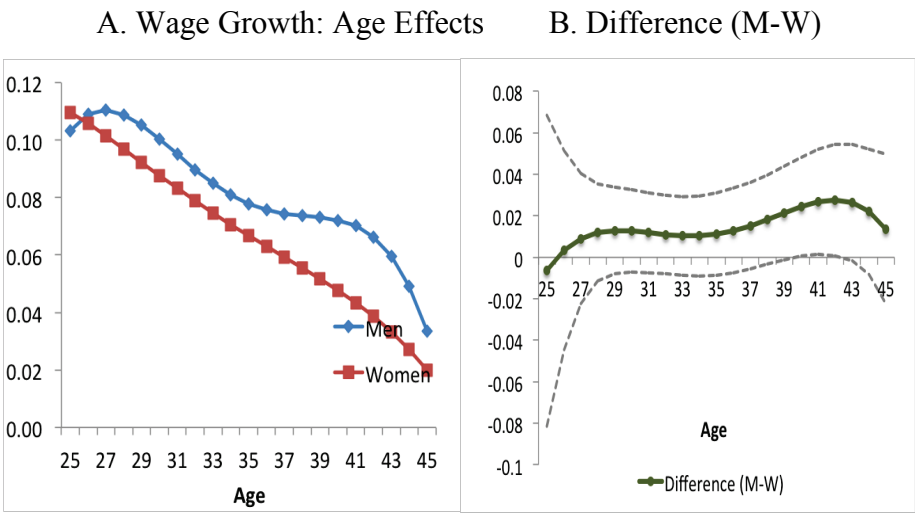

Figure 17 Wage Gains for Switchers: Estimated Time to First Birth Effects

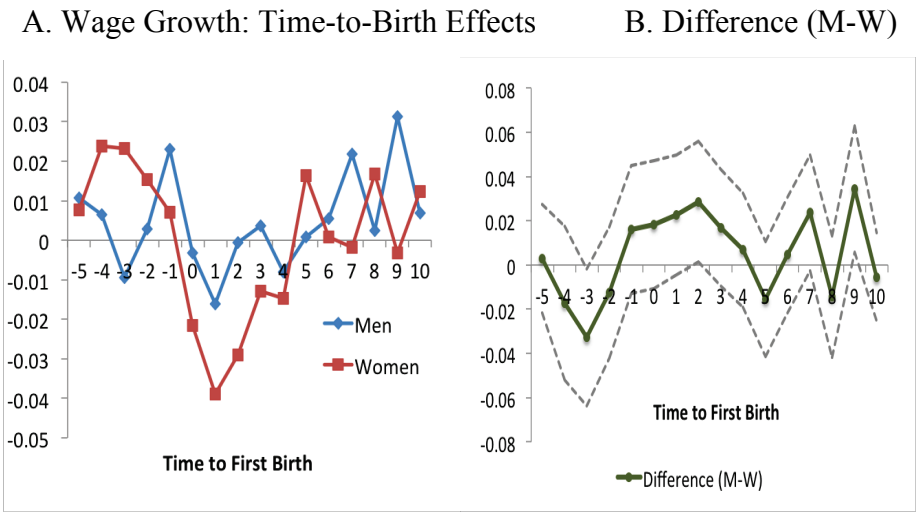

To recap, men exhibit marginally higher firm-to-firm mobility rates than women, most of this mobility occurs early in their careers, and the immediate payoff associated with a change in firms is about one to two percentage points higher for men than it is for women.

The men and women who don't move from one firm to another ("stayers") also experience different rates of wage growth. Figure 18 and 19 show the age and time-to-first-birth effects on log wage gains for workers who stay with the same employer from one year to the next. Holding time to first birth constant, the log wage gains for workers who continue to work for the same firm are highest at the start of both men's and women's careers, as 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 until age 35, with a statistically significant difference of around 1 to 2.5 percentage points. The pattern of age effects is qualitatively similar to the one shown in Figure 16 for switchers, though the gender difference in within-firm growth is most pronounced around ages 28 to 31, even controlling for time-to-birth effects, while the difference in wage gains for switchers expands somewhat over the lifecycle. The gender difference in time-to-first-birth effects on log wage gains is substantially more pronounced for stayers than it is for switchers. On average, women's wage growth in the year of the first birth and the immediately after is as much as 4 to 5 percentage points lower than men's. In the following four years, i.e., two to five years after the first birth, women's wage growth remains at least 1 percentage point lower than men's. About 6 years after first birth, this gender difference disappears. In the years prior to first birth, both men and women experience higher within-firm wage growth, relative to individuals who remain childless.

Figure 18 Wage Gains for Stayers: Estimated Age Effects

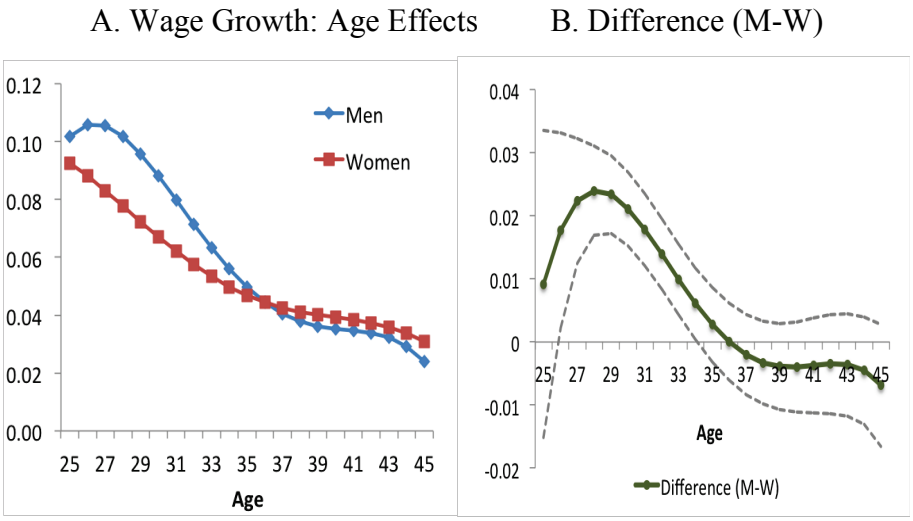

Figure 19 Wage Gains for Stayers: Estimated Time to First Birth Effects

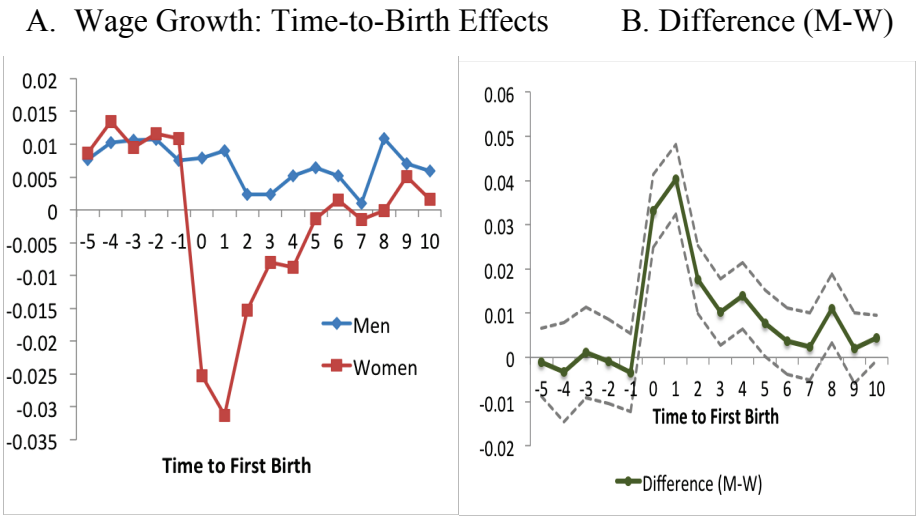

A simple decomposition can illustrate how differences in mobility and in the log wage gains associated with switching and with staying affect the gender difference in log wage profiles. For each gender, the average log wage gain from one year to the next can be written as

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

Recall that at each age, the probability of switching from one firm to another ($P[Switch]$) is very similar for men and for women. To understand the role of wage gains for switchers versus the corresponding gains for stayers, we ask what fraction of the observed gender gap in log wage gains at each age is accounted for by the gender difference in gains for (i) switchers, (ii) stayers and (iii) both switchers and stayers.

Figure 20 shows the annual gender difference, which is substantial when cumulated, in log wage gains between ages 25 and 45. The observed gender gap in annual log wage gains is shown in the Baseline profile. As expected, both the gender difference in log wage gains for switchers and the corresponding difference for stayers contribute to the overall gender gap in log wage gains, with the gender gap for stayers having a larger effect. The combination of the two effects accounts almost completely for the gender gap in annual log wage gains; that is, gender differences in mobility account for essentially none of the observed gap.

Figure 20 Decomposing Log Wage Gains

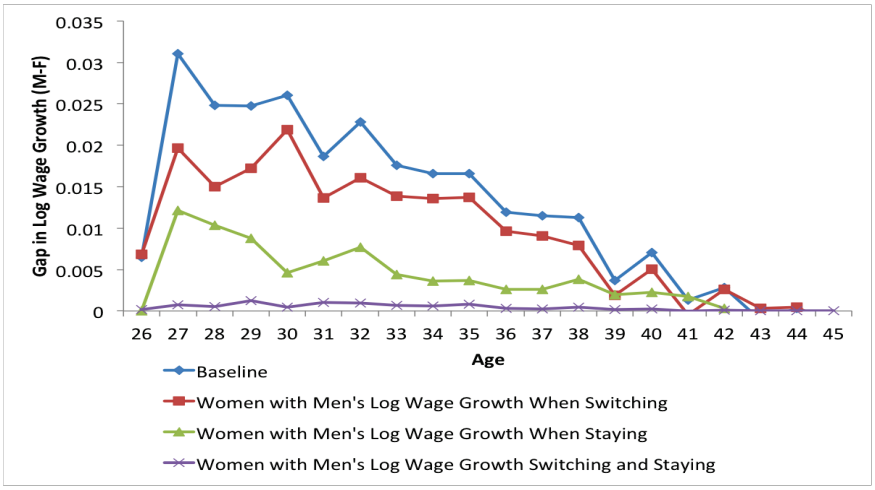

We emphasize that the patterns observed in Figure 20 – specifically, the absence of a direct effect of gender differences in the share of switchers in the decomposition – do not mean that gender differences in firm-to-firm mobility are unimportant for understanding the different log wage profiles that we observe for men and women. 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 individuals born between 1960 and 1970, and we follow these men and women for up to 20 years.

We emphasize three findings. First, much of the related literature emphasizes the effect of parenthood on the gender log wage gap. That effect is certainly present in our data, but we also find that gender differences in parenthood effects account for only about half of total gender wage differentials by age 45, and substantially less than half early in

the lifecycle. Thus, even though 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.

Second, the firms where highly skilled men tend to work differ from the firms where highly skilled women tend to work. Most strikingly, the men tend to work in substantially smaller firms than the women do. Men also tend to work in higher-paying (measured by the mean log wage of highly skilled male co-workers) and less family friendly (measured by the fraction of highly skilled workers who work less than 35 hours per week) firms than women do. One would expect that these firm differences would account for much of the gender log wage gap. However, when we carry out a Oaxaca-Blinder style wage decomposition, most of the gender gap is instead explained by gender differences in individual characteristics and by gender differences in the rewards to individual and firm characteristics.

Finally, as is emphasized in the equilibrium search literature, firm-to-firm mobility is very important for wage growth. One might expect, therefore, that highly skilled men are substantially more likely to switch firms than are highly skilled women. We do not see this difference in our data. Rather, what we see is that the wage gains among men who move from one firm to another tend to be larger than the corresponding gains for women who switch employers, and, similarly, the wage gains among stayers are higher for men than for women. In short, gender differences in annual log wage gains are accounted for by gender differences in the returns to switching and in the returns to staying, not by gender differences in mobility rates.

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.

Angelov, Nikolay,Per Johansson and Erica Lindahl (2016), "Parenthood and the Gender Gap in Pay," *Journal of Labor Economics*, 34(3), 545-77.

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.

Cortés, Patricia and Jessica Pan (2017), "When Time Binds: Substitutes for Household Production, Returns to Working Long Hours, and the Gender Wage Gap Among the Highly Skilled," *Journal of Labor Economics*, forthcoming.

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.

Karimi, Arizo, V. Joseph Hotz, and Per Johansson (2016), "Family-Friendly Firms? Worker Mobility, Firm Attributes and Wage Trajectories of Women and Men," working paper.

Kleven, Henrik Jacobsen, Camille Landais, and Jakob Egholt Søgaard (2017), "Children and Gender Inequality: 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.

Reshid, Abdulaziz Abrar (2017), "The gender gap in early career wage growth: the role of children, job mobility and occupational mobility," IFAU Working Paper 2017:15

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.

Appendix

Including Public-Sector Workers

Although we restrict our attention to men and women who spend their careers (at least the part of their careers that we can observe) entirely in the private sector, we note that among the workers we observe with business and economics degrees, women are substantially more likely to be working in the public sector. Figure A1 shows the fraction of the men and women in our dataset who are working in the public sector at each age. Women are more likely than are men to start their careers in the public sector, and they are more likely to switch from private-sector jobs to the public sector as they age and as they become parents. Specifically, approximately 30.6% of the women versus 24.3% of the men start their careers in the public sector. As men and women age, there is net mobility out of the public sector, at about equal rates for men and women, holding time-to-birth effects constant. However, women with children are somewhat more likely to move into public sector once the first child is born than men are. This mobility into the public sector is gradual and takes place over the ten years after first birth. It accounts for the small divergence in public sector participation observed in Figure A1 toward the end of the lifecycle.

Figure A1 Share Working in the Public Sector

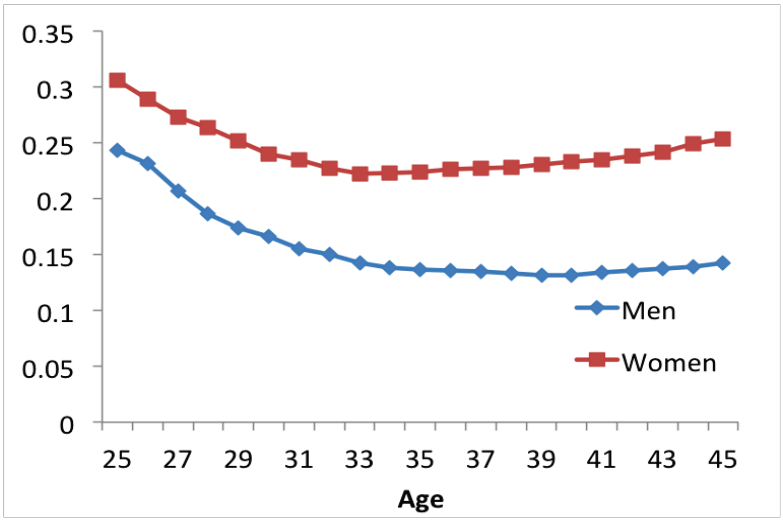

As discussed in the text, we prefer to work with the subpopulation of university graduates in business and economics who spend the part of their careers that we observe entirely in

the private sector. This restriction allows us to focus on a relatively homogeneous group of individuals. An interesting question is whether our key results change substantially when we include all business and economics majors in our analysis. Figure A2 replicates Figure 2 in the text for a sample including those who work in the public sector at some time between ages 25 and 45. That is, Panel A gives the age effects from our basic specification using the larger sample. The pattern of age effects is virtually the same. The initial age effects at 25 are identical, but now at age 45, the level is lower for both men and for women. The gap between the age effects (Panel B) is essentially the same. Figure A3 presents the time-to-first-birth effects including the public sector. Again, these are essentially the same as the effects shown in text Figure 3, i.e., our results do not change much when we add those who work part or all their careers in the public sector.

Figure A2 Log Wage: Estimated Age Effects, Including Public Sector

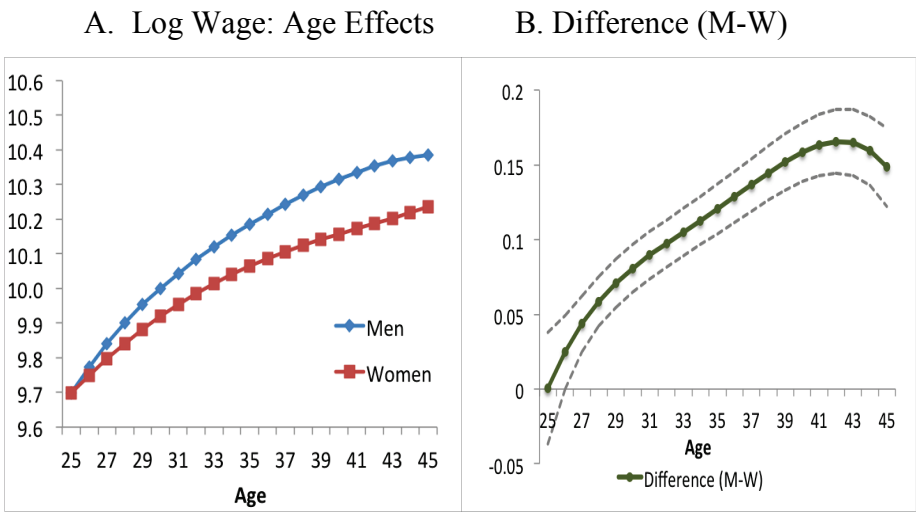

Figure A3 Log Wage: Estimated Time to First Birth Effects, Including Public Sector

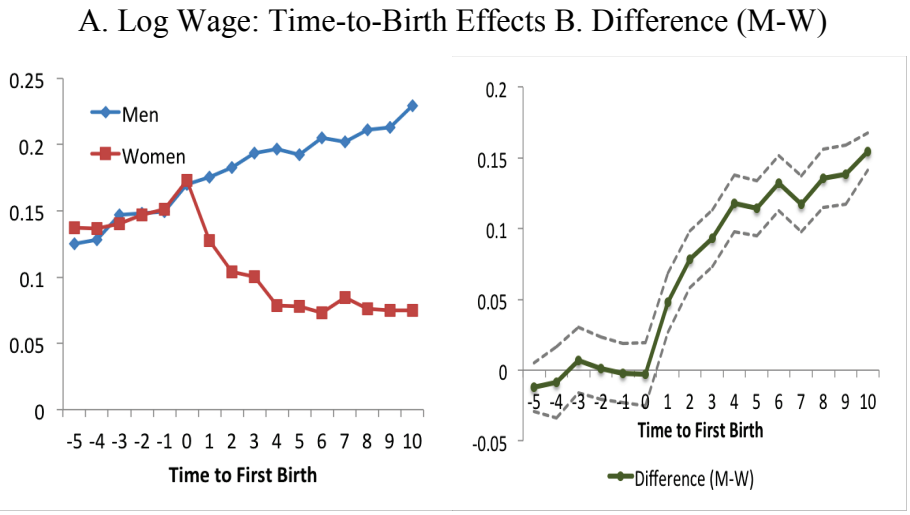

In Figure A4 and A5 we consider mean earnings of high-skill male co-workers at the individual's firm or public-sector establishment, corresponding to Figure 8 and 9 in the text. For consistency, we use the same employer identifier variable to construct this measure for public sector workers as we do for private sector workers. Note that the definition of employer in the public sector can be broad, e.g., an employer could be a municipality. Thus, some employer (firm) characteristics we study in the main text, such as number of employees, are potentially less informative for public sector workers. Similarly, the mean co-worker's earnings measure we use is potentially noisier for public sector workers. Nevertheless, comparing mean co-worker's earnings for exclusively private sector workers and for the full population of workers is a useful exercise, since the population we focus on in the text may potentially be less likely to switch to lower-paying, more flexible workplaces, even within the private sector. Interestingly, Figure A4 and A5 document that the patterns are in fact almost identical in the full population to those shown in Figure 8 and 9 for exclusively private sector workers.

Figure A4 Average Log Earnings of High-Skill Male Coworkers: Estimated Age

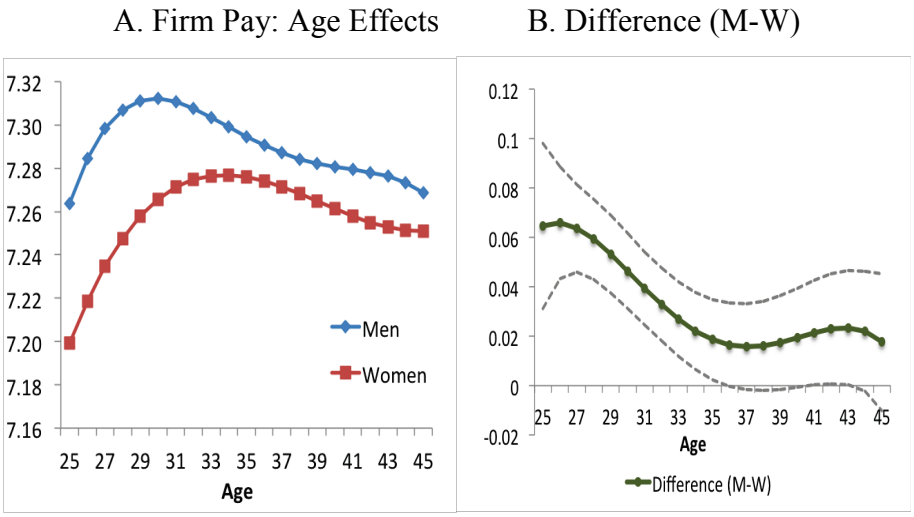

Figure A5 Average Log Earnings of High-Skill Male Coworkers: Estimated Time to First Birth Effects

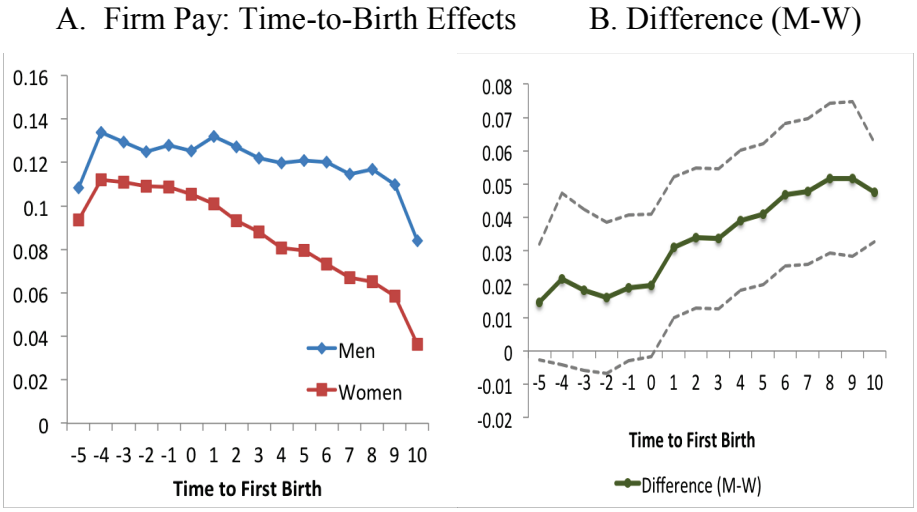

Excluding the Childless

We also consider the subpopulation consisting of only those who are parents and work in the private sector, about 76% of the population we follow in our main analysis. This gives us a subpopulation of high-skill workers analogous to the one that Angelov et al. (2016) use for the entire labor force. Similarly, Kleven et al. (2017) focus only on the population of individuals who ever have children. Figure 6 shows the age effects for this group. As is clear, these are essentially the same as the effects when the childless are included. The age effects for both men and for women start slightly higher and end slightly higher. The basic shape of the curves is the same and the gender gap in these coefficients rises as

before, but here it rises monotonically while when the childless were included, the gap fell slightly in the 40's.

Figure A6 Log Wage: Estimated Age Effects, Parents Only

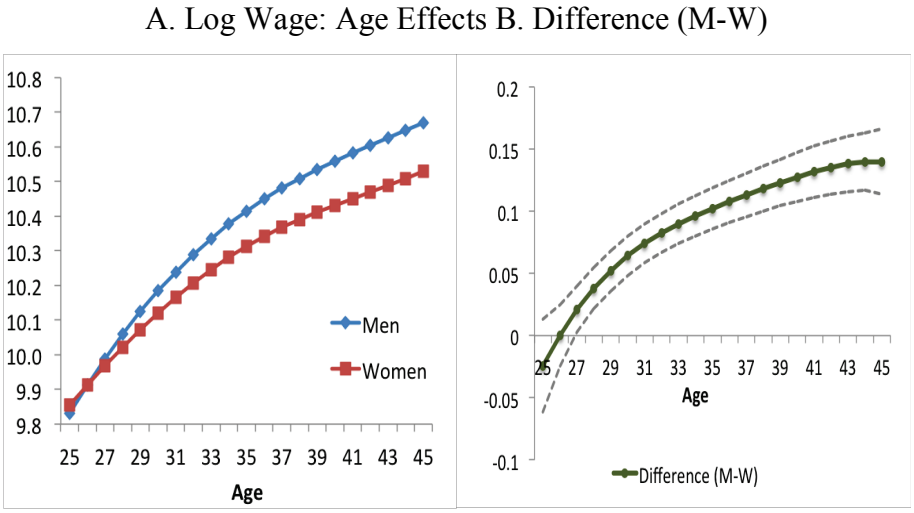

Finally, Figure A7 gives the time-to-first-birth effects including only parents. These of course can no longer be measured relative to the childless as was done before. Following the convention in prior studies, we select as the omitted category individuals one year prior to the birth of the first child. The curves for the men and women in Panel A of Figure A7 are therefore identical to those in the text but are simply normalized to zero for each gender in the year prior to first birth. Specifically, they are shifted down by about 15 log points, i.e., the difference between wages of childless individuals and parents in the year prior to first birth, controlling for age and time fixed effects. Correspondingly, Panel B graphs the difference in time-to-birth effects for men and women, normalizing the difference to zero in the year prior to first birth. In sum, our results do not change significantly when we exclude the childless.

Figure A7 Log Wage: Estimated Time to First Birth Effects, Parents Only