A GLASS ESCALATOR FOR FEMALE UVA GRADUATES?
GENDER GAPS ACROSS THE STARTING SALARY DISTRIBUTION

by

Hema Shah

Submitted to the Distinguished Majors Program
Department of Economics
University of Virginia
April 17, 2020
Advisor: Amalia R. Miller

Acknowledgments

Many thanks are in order:

First and foremost, to my advisor, Professor Amalia Miller. For her support of my research goals and ambitious ideas, for her expertise and thoughtful suggestions, and for her belief in my ability, I will always be grateful. This project would not have been possible without her. I also thank my DMP cohort - though our time together was cut short, I am thankful for their friendship and for their invaluable advice on every step of this project.

To Professor Sarah Turner, for her thoughtful and honest feedback on many stages of this project. Her guidance this year has surely made me a better economist. And to Professor Sandip Sukhtankar, my first research mentor, for inspiring my love of economics and teaching me to be a curious, compassionate, and rigorous researcher.

To Julia Lapan, Director of Engineering Career Development at the UVA Career Center. Her work to improve gender equity in career outcomes inspired this study, and her persistent advocacy allowed me to access the data I used. I also thank Everette Fortner, Kim Sauerwein, and Nathan Hunsaker from the UVA Career Center and Sarah Schultz Robinson from the UVA Office of Institutional Research and Analytics for their assistance with data access.

To Professor Ken Elzinga, Sarah Lapp, and the Marshall Jevons fund for providing resources for this project that I ultimately didn't use. And, to Debby Stanford and Ashley Watkins, for assistance and support when my research plan changed. I also thank the University of Virginia Economics Department as a whole for providing me the excellent training that brought me to this point.

And finally, to my loved ones. To my parents, for providing me with the world but encouraging me to be humble and work hard. To my grandparents, for pushing me to be the best I can be. To my twin brother Keshav, for being my partner in crime, and to my little sister Leela, for the title of this thesis. To Jacob Cunningham, for supporting me in all of my endeavors. Lastly, to Miss Julia, for inspiring my feminism and reminding me to be an advocate for women in all I do.

A Glass Escalator for Female UVA Graduates?
Gender Gaps Across the Starting Salary Distribution

Hema Shah

Abstract

Despite unprecedented increases in women's college attendance and labor force participation over the last century, large gender differences in pay persist even for graduates of elite universities. Can these gender differences in salary be explained by observed differences in graduates' skills and preparation, such as choice of major? Using unique data from the University of Virginia Career Center, I apply linear Oaxaca-Blinder decomposition models as well as unconditional quantile decomposition methods to measure the extent to which differences in early career compensation can be explained by observable differences in productive characteristics. I find that the explained share varies greatly across the pay distribution. In particular, at the low end of the salary distribution, the gender pay gap is larger and is only partially accounted for by gender differences in college major choice, industry choice, and internship experience. At the middle of the pay distribution, the gender pay gap can be almost entirely attributed to observable gender differences in these characteristics. Interestingly, at the high end of the salary distribution, gender differences in college major choice, industry choice, and internship experience suggest that the gender pay gap should be even larger than what is observed in the data. My results suggest that female graduates in high-paying majors and industries do not encounter a "glass ceiling" at the beginning of their careers. Rather, female UVA graduates either receive preferential labor market treatment, are more competent than male peers in the same majors and industries, or demand compensating wage differentials to account for differences in preferences for certain high-paying positions.

Contents

1 Introduction

Long-term trends in the United States indicate substantial reductions in the gender pay gap since the 1950s (Blau & Kahn, 2017). This convergence of male and female earnings is driven in part by changing trends in college attendance: The gender gap in college attendance has decreased continuously since the 1950s, with female college graduates now outnumbering males (Goldin et al., 2006). In addition to changes in college attendance trends, gender differences in collegiate schooling content have also narrowed. Goldin (2005) finds that the gender gap in college majors closed significantly between 1970 and 1985, a period she termed the "quiet revolution."

Gender gaps in the skills that students develop during college, however, have failed to converge since the 1980s (Shauman, 2016; Turner & Bowen, 1999). Furthermore, convergence in the mapping of major to occupation has been modest (Sloane et al., 2019). Progress in closing the gender pay gap has also slowed since the 1980s, suggesting an important link between gender differences in schooling content, industry choice, and earnings.

Previous literature demonstrates that a significant portion of the existing gender pay gap for college educated workers can be explained by gender differences in college major choice and industry selection (Blau & Kahn, 2017). It is unclear, however, whether this explanation holds for all college-educated workers. In particular, the extent to which gender differences in schooling content and industry selection can explain gender differences in earnings may vary across the earnings distribution.

This thesis examines the impact of gender differences in educational investments on gender differences in career outcomes for graduates of a public flagship university. I use new self-reported data from the University of Virginia Career Center's "First Destinations" Survey to analyze the gender gap in starting salary for UVA students who are employed full-time immediately after graduation. Using linear Oaxaca-Blinder decomposition models as well as unconditional quantile decomposition methods, I analyze the gender pay gap across the starting salary distribution.

I find that the size of the gender pay gap, as well as the extent to which it can be explained by observable gender differences in qualifications, varies greatly across the earnings distribution. The gender gap in 25th percentile salary is over 35 log points, narrowing to just over 9 log points at the 90th percentile of the salary distribution. Furthermore, only 75% of the 25th percentile pay gap can be explained by gender differences in major choice, industry selection, and internship experience. Interestingly, at the upper end of the salary distribution, gender difference in these characteristics "over-explain" the gender pay gap. My results suggest that female graduates in high-paying majors and career industries either receive preferential labor market treatment or are more qualified on dimensions that are unobservable in the data. Alternatively, male and female graduates' preferences for certain kinds of work may differ, creating a need for employers to compensate women more highly in order to attract qualified female graduates.

A small amount of existing literature analyzes the gender wage gap for graduates of a single prestigious university. Graham et al. (2000) use institutional data on graduates of a "highly prestigious" university employed by firms with at least 10 graduates of that university. Employing linear decomposition methods, the authors find that within firms, gender gaps in salary are mainly due to differences in field of study. Bertrand et al. (2010) analyze gender differences in outcomes for graduates of the University of Chicago Booth School of Business. They find that, while men and women's earnings diverge later in their careers, immediately after graduation their labor force participation rates and earnings are nearly identical.

I expand on previous work by studying a public flagship institution with a larger and more varied sample of employed graduates. I also apply novel estimation methods to study the wage gap beyond the mean, a technique not found in the literature on recent college graduates. Applying unconditional quantile decomposition methods in addition to standard linear Oaxaca-Blinder decomposition models allows me to identify potential "glass ceilings" experienced by female college graduates.

The term “glass ceiling” is used to describe explicit or implicit barriers that exclude women from the highest-paying jobs. Glass ceilings are often found even in developed Western nations and could be due to a scarcity of senior women in high-skilled, high-paying occupations; women’s higher propensity to exit the labor force temporarily and interrupt their acquisition of skills; women’s hesitation to negotiate salary offers (Babcock & Laschever, 2003); and the exclusionary aspects of “boys’ club networks that are prominent in certain industries (Xiu & Gunderson, 2014).

Glass ceilings are likely to be a particular concern for female graduates of elite universities, many of whom will graduate with extremely high earnings potential. As such, my work provides important insight into the pre-market human capital specialization and subsequent labor market treatment of highly qualified female college graduates. The remainder of the paper proceeds as follows: Section 2 presents an overview of relevant theoretical and empirical literature; Section 3 describes the empirical specifications used; Section 4 summarizes the data; Section 5 details results; and Section 6 concludes.

2 Literature Review

2.1 Insights from Human Capital Theory

Human capital theory provides insight into why gender differences in educational investments may persist. In a simple human capital model, an individual will invest in education up to the point at which the benefits from schooling are equal to the costs (Becker, 1964). The benefits and costs of schooling, and consequently the ideal schooling decision, may differ between men and women, causing gender gaps in decisions such as college major choice. Gender differences in the costs and benefits of college majors may arise for several reasons.

First, men and women may possess innate differences in academic ability; thus, the return to investing in academically rigorous majors with potentially high salary returns may be lower for women than for men. This notion of comparative advantage influencing occupational choice and subsequent earnings has been explored extensively in the theoretical literature, beginning with Roy (1951). The theory of comparative advantage amongst college educated workers has also been tested empirically. Willis and Rosen (1979) use data on male World War II veterans' IQ and test scores to show that comparative advantage partly determines the decision to invest in higher education. Using SAT and GRE test scores, Paglin and Rufolo (1990) find that mathematical ability is an important determinant of college major choice, and that differences in earnings across fields are largely explained as returns to the use of quantitative abilities that not all students possess.

The theory of comparative advantage may be particularly useful in the context of quantitative majors such as STEM (science, technology, engineering, and mathematics) fields, which continue to be male dominated despite changing gender attitudes towards women's work. If women have lower mathematical aptitude than men on average, then women will likely experience greater costs to majoring in STEM fields, such as time spent studying and stress imposed by examinations. Furthermore, women may also experience smaller salary returns to majoring in STEM fields if they have lower mathematical aptitude. Previous literature, however, finds that gender differences in mathematical ability, as measured by quantitative standardized test scores, are likely too small to explain existing gender gaps in college major choice (Ceci et al., 2014; Hyde et al., 2008; Riegle-Crumb et al., 2012).

An alternative explanation for self-selection of men and women into different college majors may be that men and women simply have different preferences regarding college majors and the associated career outcomes (Zafar, 2012). For example, women may prefer majors with less quantitative coursework or majors associated with career fields that have less demanding hours. This could be the result of innate differences in preferences or of "sex role socialization" (Corcoran & Courant, 1985; Eccles & Hoffman, 1984), the conditioning of women to prefer certain occupations over others due to social factors. Women may also anticipate differential treatment in the labor market, due to discrimination in hiring or discrimination in promotions (Lazear & Rosen, 1990). Economic theory, however, predicts that discrimination is not sustainable in a competitive labor market (Arrow, 1972).

Uncertainty about success in schooling and in the labor market may also impact gender differences in college major choice (Altonji, 1993). Men and women may have different beliefs regarding the outcomes associated with certain majors, due to differences in information or other factors (Wiswall & Zafar, 2014; Zafar, 2012). This could lead them to systematically choose different fields of study.

2.2 Wage Decomposition Methods and Findings

A large amount of previous literature uses linear decomposition methods to analyze the contribution of gender differences in education to the gender pay gap. Much of this literature utilizes large, nationally representative survey data sets rather than institution-level data. Brown and Corcoran use data from the Survey of Income and Program Participation (SIPP) and the National Longitudinal Study class of 1972 (NLS-72). The authors find that, for college-educated workers, 50% of the portion of the gender wage gap that is unexplained by work experience and demographic factors can be explained by gender differences in college major choice (Brown & Corcoran, 1997). Loury (1997) finds similar results using the NLS-72 and the High School and Beyond Senior Cohort (Class of 1980).

Other literature uses similar linear decomposition methods on different nationally representative survey data, yet finds drastically different empirical results. Bobbitt-Zeher (2007) uses the 2000 National Educational Longitudinal Survey (NELS), restricting analysis to college-educated workers in their mid-20s. The author finds that, even for workers at the start of their careers, only 14% of the gender gap in annual salary can be explained by field of study. Joy (2003) uses data from the 1993-94 NCES Baccalaureate and Beyond Longitudinal Study and finds that college major accounts for just 1% of the gender salary gap.

More recent literature attempts to explain the gender wage gap for college-educated workers using nonlinear decomposition methods and novel data sources. Black et al. (2008) use a nonparametric matching procedure on data from the 1993 National Survey of College

Graduates (NSCG) and find that between 44 and 73% of the gender wage gap is accounted for by highest degree, major, and age (Black et al., 2008). A study by Glassdoor economists finds somewhat contradictory results using linear decomposition methods on anonymous, self-reported salary data from the Glassdoor platform. While 68% of the gender gap in base salary was accounted for by gender differences in observable characteristics, only 14% was accounted for by gender differences in education and experience. The remaining 54% was accounted for by sorting of men and women into different industries and occupations (Chamberlain, 2016).

Relevant to my analysis, additional literature explores college major choice and the gender wage gap specifically for starting salary offers. This eliminates the potential influence of gender differences in promotions, in parental leave time, and in years of work experience (although expected years of labor force participation may still differ). McDonald and Thornton use data from the annual surveys of the National Association of Colleges and Employers (NACE) and find that as much as 95% of the overall gender gap in starting salary offers can be attributed to differences in college major choice (McDonald & Thornton, 2007).

2.3 Institution-Level Analysis

The previously described analyses of nationally representative data may obscure differences between higher education institutions. These differences could include disparities in teaching quality, career services, alumni networks, and other factors potentially impacting labor market outcomes. Some existing literature controls for differences between institutions by evaluating college majors and career outcomes for graduates of individual colleges and universities.

In particular, Todd Stinebrickner and Ralph Stinebrickner designed and administered the Berea Panel Study (BPS), a multipurpose longitudinal survey conducted on students at Berea College in central Kentucky. Respondents in two cohorts were surveyed 12 times per year while in school and annually thereafter about expectations towards uncertain outcomes, the factors that might influence these outcomes, and eventually the outcomes themselves (T. Stinebrickner & Stinebrickner, 2011). This study has been the basis for research papers exploring college major choice (T. Stinebrickner & Stinebrickner, 2011); persistence within college major (R. Stinebrickner & Stinebrickner, 2014); and, relevant to my analysis, labor market outcomes such as the beauty wage premium (R. Stinebrickner et al., 2019) and the wage consequences to over-education (Agopsowicz et al., 2017). My work will contribute to this literature by examining career outcomes within a different institutional setting. Berea College is a small, private liberal arts college with a focus on providing an education to students from low income backgrounds. As such, Berea offers a full tuition subsidy to all students. UVA is a large, public flagship university with students coming from more diverse socioeconomic backgrounds and receiving a spectrum of financial aid packages. UVA is also slightly more selective, with an acceptance rate of 26% compared to Berea's 38% in 2018.[1]

Other literature specifically analyzes college major choice at individual institutions. (Zafar, 2012) and (Wiswall & Zafar, 2014) study gender differences in college major choice in the context of Northwestern University and New York University, respectively. Similarly to R. Stinebrickner and Stinebrickner (2014), Wiswall and Zafar collect data on students' major preferences and their subjective expectations regarding the academic, personal, and professional outcomes associated with various majors. The authors find that most of the gender gap in major choice is due to gender differences in preferences for major-specific outcomes, such as salary and work hours.

[1]Acceptance figures were taken from the Integrated Postsecondary Education Data system (IPEDS) at https://nces.ed.gov/ipeds/about-ipeds.

3 Empirical Specifications

3.1 Linear Oaxaca-Blinder Decomposition

I begin by estimating a standard linear Oaxaca-Blinder decomposition model (Blinder, 1973; Oaxaca, 1973).[2] First, I estimate three separate linear regression models — one for males, one for females, and one over all graduates as follows:

$$\overline{Y}_M = \hat{\beta}_M \overline{X}_M \tag{1}$$

$$\overline{Y}_F = \hat{\beta}_F \overline{X}_F \tag{2}$$

$$\overline{Y} = \hat{\beta} \overline{X} \tag{3}$$

From these estimates, I obtain the Oaxaca-Blinder decomposition model:

$$\overline{Y}_M - \overline{Y}_F = [(\overline{X}_M - \overline{X}_F) \times \hat{\beta}] + [(\hat{\beta}_M - \hat{\beta}_F) \times \overline{X}] \tag{4}$$

Here, $\hat{\beta}_M$ and $\hat{\beta}_F$ are the estimated coefficient vectors in the male and female group models, respectively. $\hat{\beta}$ is the coefficient vector of the pooled model, the "nondiscriminatory" coefficients vector representing a counterfactual situation in which male and female workers receive the same salary returns to observable skills (Jann, 2008). $\overline{Y}_M$ and $\overline{Y}_F$ represent average salary for male and female graduates working full-time; $\overline{Y}_M - \overline{Y}_F$ is therefore the raw gender gap in average starting salary. $\overline{X}_M$ and $\overline{X}_F$ are vectors representing average characteristics of male and female graduates, including major, ethnicity, internship experience, and career industry. $\overline{X}$ is a vector of characteristics averaged over all graduates.

The left hand side of the equation represents the gender gap in average salary. The first

[2] To estimate my linear Oaxaca-Blinder decomposition models in Stata I use Ben Jann's `oaxaca` command (Jann, 2008).

term on the right hand side of the equation, $[(\overline{X}_M - \overline{X}_F) \times \hat{\beta}]$, represents the portion of the gender gap in average salary that can be explained by observable outcomes, while the second term, $[(\hat{\beta}_M - \hat{\beta}_F) \times \overline{X}]$, represents the portion of the gender gap that cannot be explained using the variables in the data. This portion of the gap is due to gender differences in returns, rather than stock, of human capital. The unexplained portion of the gender gap is often attributed to discrimination; however, this portion of the gap may also be the result of variables omitted from the regression. Notably, I do not control for any measure of academic ability due to data constraints discussed in Section 4. I am primarily interested in the relative magnitude of $[(\overline{X}_M - \overline{X}_F) \times \hat{\beta}]$; that is, the portion of the gender wage gap that can be explained by gender differences in observable characteristics.

One advantage of using a pooled decomposition model is that standard errors are much lower for the coefficients of sex-atypical majors using this approach. This is because, for majors with very few women or very few men, the pooled regression estimates draw from a much larger sample size and therefore produce more precise estimates (Brown & Corcoran, 1997).

3.2 Unconditional Quantile Decomposition

There exists a significant amount of recent literature debating the merits of various quantile regression models and quantile decomposition methods. Albrecht et al. (2003) employ conditional quantile regression methods on 1998 data from Sweden containing a representative sample of workers aged 15-75. They find that the log gender wage gap increases throughout the wage distribution and sharply accelerates at the upper tail even after controlling for gender differences in age, education, industry, and occupation. This result is interpreted by the authors as evidence of a glass ceiling for female workers. Machado and Mata (2005) propose applying conditional quantile regression techniques to generalize linear decomposition models, a technique used in several subsequent papers to analyze wage gaps (Arulampalam et al., 2007; Lucifora & Meurs, 2006). Melly (2005) proposes a similar method of applying conditional quantile regression to decomposition models.

More recent literature has focused on the use of unconditional quantile regression models, such as the reduced influence function (RIF) regression model popularized by Firpo et al. (Firpo et al., 2009; Fortin et al., 2011). This procedure allows for the generalization of linear Oaxaca-Blinder models to distributional statistics other than the mean (Firpo et al., 2018), a technique used in recent literature on the gender wage gap in China (Chi & Li, 2008; Xiu & Gunderson, 2014), the United States (Kassenboehmer & Sinning, 2014), and various countries in Latin America (Carrillo et al., 2014). Unlike conditional quantile regression decomposition techniques, this method allows quantiles to be decomposed non-sequentially in the same way means can be decomposed using the conventional Oaxaca-Blinder methodology (Firpo et al., 2018).

Due to its popularity in recent literature as well as its analogies to the standard Oaxaca-Blinder model, I use recentered influence function (RIF) regression analysis to evaluate the gender pay gap at various points along the salary distribution.[3] My RIF regression model replaces the dependent variable in a standard linear regression model with the recentered influence function of the quantile of interest.[4] In this regression model, the coefficients correspond to the marginal effect on the unconditional quantile of shifts in the distribution of covariates, holding everything else constant.

For each quantile of interest τ, I estimate three RIF unconditional quantile regressions: one for the male earnings distribution, one for the female earnings distribution, and one counterfactual distribution in which females have the same characteristics as males. Analogous to OLS regressions, the RIF regression functions assume a linear specification as follows:

[3] To estimate my RIF decomposition models in Stata I use Fernando Rios-Avila's `oaxaca_rif` command (Rios-Avila, 2019).

[4] An influence function of a distributional statistic represents the influence of a single observation on the value of that distributional statistic. Adding back the distributional statistic to the influence function yields the recentered influence function (RIF). Conveniently, the expectation of the RIF is equal to the distributional statistic. In this case, the distributional statistics of interest are quantiles of the salary distribution (Firpo et al., 2009).

$$\nu_M = E[RIF(Y; q_{M,\tau}|X)] = \hat{\beta}_M \overline{X}_M \tag{5}$$

$$\nu_F = E[RIF(Y; q_{F,\tau}|X)] = \hat{\beta}_F \overline{X}_F \tag{6}$$

$$\nu_C = E[RIF(Y; q_{C,\tau}|X)] = \hat{\beta}_C \overline{X}_C \tag{7}$$

Similarly to the linear Oaxaca-Blinder model, I use these estimates to obtain the following decomposition:

$$\nu_M - \nu_F = [(\overline{X}_M - \overline{X}_F) \times \hat{\gamma}_C] + [(\hat{\beta}_M - \hat{\beta}_F) \times \overline{X}_C] \tag{8}$$

Here, $q_{M,\tau}$ and $q_{F,\tau}$ are the population τ - quantiles of the unconditional distribution of Y, total salary, for males and females, respectively. $q_{C,\tau}$ is the population τ - quantile of the counterfactual distribution. Thus, $\nu_M - \nu_F$ is the estimated raw gender gap in τ - quantile salary. $\hat{\beta}_M$ and $\hat{\beta}_F$ are the estimated coefficient vectors in the male and female group RIF models. $\hat{\beta}_C$ is the estimated coefficient vector in the counterfactual model. $\overline{X}_M$ and $\overline{X}_F$, as before, are vectors representing average characteristics of male and female graduates, and $\overline{X}_C$ represents average characteristics over all graduates in the counterfactual model.

The interpretation of this equation matches the standard Oaxaca-Blinder interpretation. The first term on the right hand side of the equation, $[(\overline{X}_M - \overline{X}_F) \times \hat{\gamma}_C]$, represents the portion of the gender gap in τ - quantile salary that can be explained by observable outcomes, while the second term, $[(\hat{\beta}_M - \hat{\beta}_F) \times \overline{X}_C]$, represents the portion of the gender gap that cannot be explained using the variables in the data.

4 Data Description

4.1 Overview

My data come from the University of Virginia Career Center's "First Destinations" Survey (FDS)[5] for the years 2016-2018. Due to data privacy restrictions, the years in my data set have been anonymized: I refer to them as Year A, Year B, and Year C. The FDS data include self-reported information from recent UVA graduates. The First Destinations Survey is made available to students beginning in December of their last year at the University of Virginia, and remains open for approximately one year. This means that students who graduate in May have seven months after graduation to report their post-graduation plans. The data includes the following information: primary school of enrollment, degree attained, major(s), minor(s), enrollment in higher education, undergraduate internship experience, post-graduation salary,[6] post-graduation job location, and post-graduation career industry, along with demographic information such as students' race and gender. I restrict my analysis to UVA bachelor's degree recipients, excluding graduates of UVA's professional schools and graduate programs.

Despite being self-reported, this data arguably provides the most accurate estimates of salary outcomes for recent University of Virginia graduates. The State Council of Higher Education for Virginia (SCHEV) reports wage outcomes using unemployment tax data from the Virginia Employment Commission (VEC)[7]. This administrative data includes information only on graduates who are employed in the state of Virginia and who meet the following

[5]Information about the survey can be found on the UVA Career Center's website at https://career.virginia.edu/uva-career-outcomes. The data is collected primarily for the purpose of constructing the Career Center's annual reports on student outcomes. Data for Year A and Year B was self-reported in the UVA Student Outcome Activity Report (SOAR) and data for Year C was self-reported through Handshake, an online recruiting platform used by the University of Virginia. Students from the McIntire School of Commerce reported their outcomes in the McIntire Portfolio Destination Survey (PDS). The Career Center website states that, in a limited number of cases, information was captured "through other sources, including faculty, employers, and social media (LinkedIn)."

[6]I measure earnings using log - base salary in all of my analysis.

[7]SCHEV salary information on University of Virginia graduates is available at https://research.schev.edu/iprofile.asp?UID=234076.

criteria:

1. *Graduates successfully matched to the Unemployment Insurance Wage records collected by the Virginia Employment Commission (VEC).*
2. *Graduates employed in Virginia by an entity that reports to the VEC. This excludes federal employees, including those within the Department of Defense.*

Additionally, the State Council of Higher Education for Virginia reports salary outcomes beginning 18 months after graduation, rather than immediately after. The FDS data therefore provides better information on students' earliest outcomes after graduation.

While the FDS data is likely subject to some amount of self-reporting bias[8] in comparison to the administrative data reported by SCHEV, the FDS data has the advantage of including salary information on graduates who are employed outside of the state of Virginia. As many UVA graduates who pursue high-paying urban jobs settle in large metropolitan areas outside of the state of Virginia, my salary estimates are higher on average than the estimates reported by SCHEV. For 2017 graduates (the most recent cohort for which salary data is available from SCHEV), the State Council of Higher Education reported average salary 18 months after graduation to be $47,880. In my data, the average salary over all graduates employed full-time immediately after graduation is $56,548. Assuming that graduates' salaries do not decrease substantially in the 18 months following graduation, it is therefore likely that my sample includes more high-earning graduates. Thus, my analysis likely provides a better view of gender gaps in outcomes at the upper tail of the salary distribution.[9]

My sample has 7,918 observations: 2,120 from Year A; 2,762 from Year B; and 3,036 from Year C. It is important to note that the First Destinations Survey was not identical over

[8]Self-reporting bias may be of particular concern if we expect that bias differs between male and female students. For example, if male students are more likely than female students to report higher compensation than they are actually receiving, my results will over-estimate the size of the gender pay gap.

[9]By merging college records with both in-state and national earnings data from the Longitudinal Employer-Household Dynamics (LEHD), Foote and Stange (2019) demonstrate that the effect of graduating from a public flagship university is underestimated by 26% when exclusively using in-state earnings. Additionally, earnings differences across majors are underestimated when using in-state earnings alone.

the three years for which I have data. During the sample period, the University of Virginia switched survey distribution platforms, removed several questions, and simplified response options for many questions.[10] Despite this change, the samples do not differ significantly across years, as demonstrated in Table 1 below. Therefore, I argue that survey differences across years do not pose a significant problem for my subsequent analysis. I do, however, include time fixed effects in my primary specification to account for changes in data collection practices as well as changes in labor market conditions over the sample period.

Table 1: Means by Year

Variable	**Year A**	**Year B**	**Year C**
Annual Salary, Full-Time Workers	$54,374.99	$56,837.10	$57,914.75
Log Annual Salary, Full-Time Workers	10.83	10.87	10.80
Number of Internships Completed	2.49	2.36	3.13
Ethnicity Dummy Variables			
American Indian or Alaska Native	0.0010	0.0004	0.0007
Asian	0.1220	0.1255	0.1526
Black	0.0560	0.0581	0.0666
Hispanic	0.0585	0.0624	0.0828
Multi-Race	0.0446	0.0478	0.0004
Non-Resident Alien[11]	0.0407	0.0516	N/A
White	0.6772	0.6542	0.6933
Outcome Type Dummy Variables			
Working	0.5481	0.4692	0.5178
Continuing Education	0.1774	0.1463	0.1792
Other[12]	0.2745	0.3845	0.3030

[10] Notably, I do not have information on Year C graduates' post-graduation job locations; therefore, I do not control for job location in my analysis. Year C graduates were also provided with a different set of industry categorizations to choose from in the survey. To account for this discrepancy, I recoded my industry variable to make the industry categories broader, fewer in number, and consistent across years.

My sample is fairly representative of the racial composition of the University of Virginia undergraduate student body, although white students are slightly overrepresented and black students are slightly underrepresented.[13] As of Fall 2016, 6.45% of all undergraduates at the University of Virginia were Black; 12.82% were Asian; 6.28% were Hispanic; 4.40% were Multi-Race; 4.75% were Non-Resident Aliens; and 59.47% were White. The remainder reported their race as "Unknown" or "Other," with the "Other" category including Native American and Alaskan Native students.

As demonstrated in Table 2 below, one variable that does differ across years is my set of industry dummies. In particular, a higher percentage of students identify their industry as "Other" in Year C than in Year A or Year B. This is likely due to changes in industry categorization options in the First Destination Survey.

[11] Year C graduates did not have the option to identify their ethnicity as "Non-Resident Alien" in the First Destinations Survey.

[12] Since this data is collected less than a year after graduation, the "Other" category includes many students who are taking time off after graduation to apply to graduate and professional schools; participate in volunteer work or unpaid fellowship programs; travel; or seek employment.

[13] Data on the racial composition of the UVA student body comes from Institutional Research and Analytics (IRA) at the University of Virginia. Data was retrieved online at https://ira.virginia.edu/university-stats-facts/enrollment.

Table 2: Proportion of Students Working in Each Industry, by Year

Industry	**Year A**	**Year B**	**Year C**
Accounting	0.0029	0.0023	0.0081
Arts, Media, and Entertainment	0.0385	0.0153	0.0305
Communications	0.0584	0.0397	0.0285
Construction	0.0292	0.0351	0.0332
Consulting	0.1803	0.2313	0.1220
Consumer Products/Retail	0.0207	0.0260	0.0169
Education	0.0520	0.0412	0.0800
Financial Services	0.1326	0.1733	0.0807
Government	0.0371	0.0443	0.0671
Healthcare	0.0919	0.0710	0.1512
IT and Engineering	0.2060	0.1702	0.1810
Natural Resources[14]	0.0100	0.0084	0.0176
Nonprofit/NGO	0.0378	0.0313	0.0393
Real Estate	0.0007	0.0084	0.0176
Services	0.0656	0.0626	0.0231
Other	0.0364	0.0397	0.1031

4,419 (55.81%) of the observations in my sample are female and 3.499 (44.19%) are male. This is representative of overall undergraduate enrollment at the University of Virginia: In Fall 2016, 54.71% of undergraduates were female and 45.29% were male.[15] Summary statistics for male and female graduates are presented in Table 3.

[14]Full industry title: Natural Resources, Agriculture, and Environmental Science

[15]This information comes from Institutional Research and Analytics (IRA) at the University of Virginia and was retrieved online at https://ira.virginia.edu/university-stats-facts/enrollment.

Table 3: Means by Gender

Variable	Male	Female	Gender Gap
Annual Salary, Full-Time Workers	\$62,781.37	\$51,204.27	\$11,577.10***
Log Annual Salary, Full-Time Workers	10.97	10.72	0.25***
Number of Internships Completed	2.59	2.71	-0.12***
Ethnicity Dummy Variables			
American Indian or Alaska Native	0.0012	0.0002	0.0010
Asian	0.1317	0.1368	-0.0051
Black	0.0511	0.0683	-0.0172
Hispanic	0.0738	0.0649	0.0089
Multi-Race	0.0317	0.0300	0.0017
Non-Resident Alien	0.0258	0.0324	-0.0066*
White	0.6846	0.6673	0.0173
Outcome Type Dummy Variables			
Working	0.5316	0.4911	0.0405***
Continuing Education	0.1569	0.1754	-0.0185**
Other	0.3115	0.3335	-0.2220**

Results of a two-sample t-test are indicated as follows: *** p<0.01, ** p<0.05, * p<0.1

Table 3 demonstrates that the gender gap in average salary is over \$11,000, significant at the 1% confidence level. It is unlikely that this gap is due to gender differences in work experience, as the average number of internships completed by female graduates is slightly higher than the average for males. Racial differences are also unlikely to explain the gap, as the racial composition of the male and female samples is not significantly different.

4.2 Gender Differences in Earnings

In addition to the gender gap in mean earnings, I am interested in the gender gap at other points in the salary distribution. Figure 1 plots kernel density estimates of the log-earnings distributions for male and female graduates. Figure 2 plots the raw gender earnings differential at different quantiles. It is evident from these figures that the raw gender gap in earnings varies greatly across the pay distribution. Specifically, the pay gap is much larger at the bottom of the salary distribution and smaller at the top of the distribution. This suggests the absence of a "glass ceiling" for female graduates directly after graduation.

Figure 1: Kernel Density Estimates of the Log Earnings Distribution

Figure 2: Gender Pay Gap by Percentile

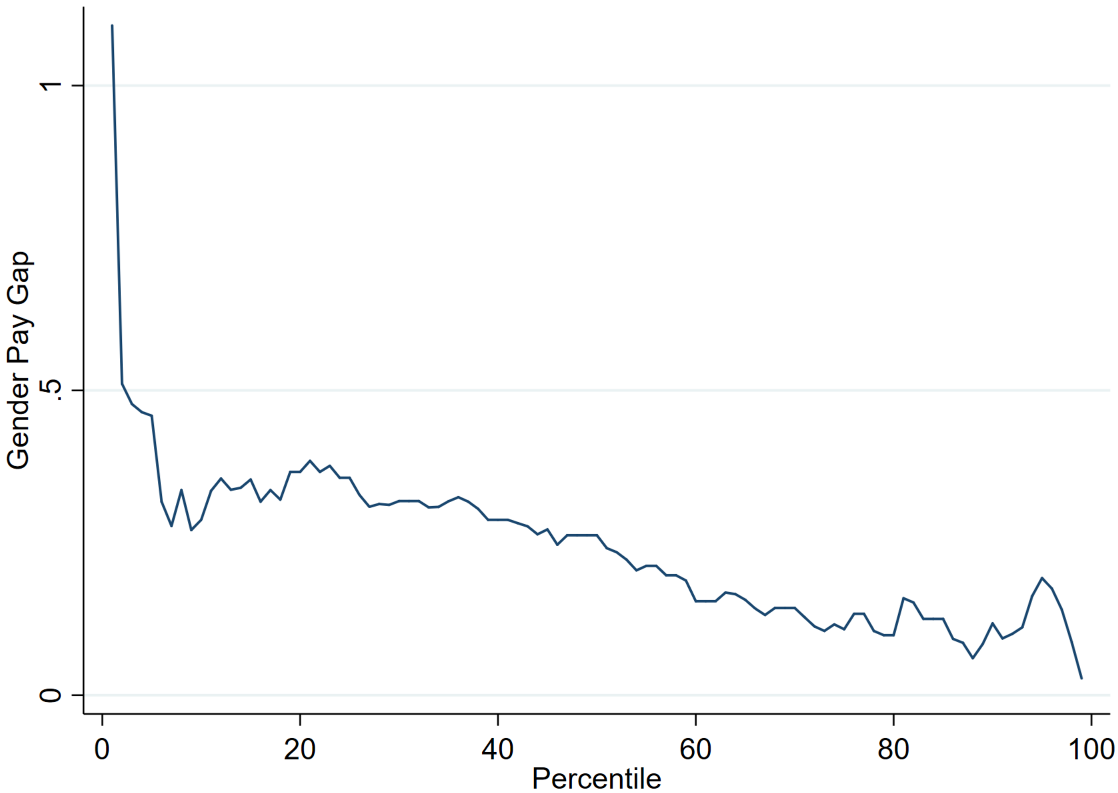

4.3 Gender Segregation in Majors and Industries

Tables 4 and 5 examine gender segregation within majors and within industries. Gender differences in participation are significant for a large majority of majors and career industries.

Table 4: Proportion of Male and Female Students in the 20 Largest Majors[16]

Major	Male	Female	Gender Gap
Economics	0.0921	0.0436	0.0485***
Biology	0.0475	0.0766	-0.0291***
Foreign Affairs	0.0455	0.0550	-0.0095*
Psychology	0.0203	0.0631	-0.0429***
Computer Science (B.S.)[17]	0.0669	0.0186	0.0483***
Media Studies	0.0160	0.0561	-0.0401***
Systems Engineering	0.0515	0.0200	0.0315***
English	0.0157	0.0421	-0.0263***
Biomedical Engineering	0.0323	0.0271	0.0053
History	0.0360	0.0221	0.0140***
Mechanical Engineering	0.0481	0.0118	0.0362***
Computer Science (B.A)	0.0386	0.0184	0.0202***
Public Policy and Leadership	0.0243	0.0289	-0.0046
Government	0.0229	0.0250	-0.0021
Chemistry	0.0254	0.0211	0.0043
Cognitive Science	0.0117	0.0286	-0.0169***
Global Studies	0.0086	0.0309	0.0223***
Civil Engineering	0.0215	0.0186	0.0028
Finance	0.0300	0.0098	0.0203***
Environmental Science	0.0146	0.0207	-0.0061**

Results of a two-sample t-test are indicated as follows: *** $p<0.01$, ** $p<0.05$, * $p<0.1$

Major is defined as primary major, excluding second majors.

Table 5: Proportion of Male and Female Students Working in Each Industry

Industry	**Male**	**Female**	**Gender Gap**
Accounting	0.0038	0.0051	-0.0013
Arts, Media, and Entertainment	0.0147	0.0392	-0.0246***
Communications	0.0271	0.0537	-0.0266***
Construction	0.0456	0.0222	0.0234***
Consulting	0.1758	0.1757	0.0001
Consumer Products/Retail	0.0222	0.0200	0.0022
Education	0.0282	0.0823	-0.0541***
Financial Services	0.1742	0.0900	0.0842***
Government	0.0564	0.0448	0.0117*
Healthcare	0.0646	0.1390	-0.0745***
IT and Engineering	0.2572	0.1301	0.1271***
Natural Resources	0.0130	0.0115	0.0015
Nonprofit/NGO	0.0179	0.0507	-0.0328***
Real Estate	0.0109	0.0077	0.0032
Services	0.0369	0.0597	-0.0228***
Other	0.0515	0.0682	-0.0167**

Results of a two-sample t-test are indicated as follows: *** p<0.01, ** p<0.05, * p<0.1

[16]Notably, engineering majors are overrepresented in my sample due to the University of Virginia Center for Engineering Career Development's intensive advertising of the First Destinations Survey. Thus, the 20 largest majors by number of graduates at the University of Virginia differ slightly from the 20 largest majors in my sample. In 2018, the 20 majors with the largest numbers of graduates were (in no particular order): Economics, Biology, Foreign Affairs, Psychology, Media Studies, English, Computer Science (B.A.), History, Government, Computer Science (B.S.), Global Studies, Cognitive Science, Spanish, Mathematics, Chemistry, Finance, Accounting, Statistics, Environmental Science, and Systems Engineering. This information was retrieved from Institutional Research and Analytics (IRA) at the University of Virginia: https://ira.virginia.edu/university-stats-facts/degrees-awarded.

[17]The University of Virginia offers two majors in Computer Science: A Bachelor of Science (B.S.) degree in the School of Engineering and Applied Science and a Bachelor of Arts (B.A.) degree in the College of Arts and Sciences.

It is evident from Table 4 that nearly every major has a significant gender gap in participation. Similarly, Table 5 demonstrates that most industries also demonstrate significant gender gaps in participation. This suggests that gender segregation in college majors and career industries may explain a large portion of the gender pay gap.

Figure 3 demonstrates that many of the highest paying majors have high percentages of male graduates and, conversely, many majors with lower average salaries have high percentages of female graduates. This further supports the hypothesis that gender segregation in college majors is a large determinant of the gender pay gap for recent UVA graduates. Figure 4 shows a similar result for career industries: Many of the highest paying industries employ high percentages of male graduates and many industries with lower average salaries employ high percentages of female graduates.

Figure 3: Gender Segregation of Majors

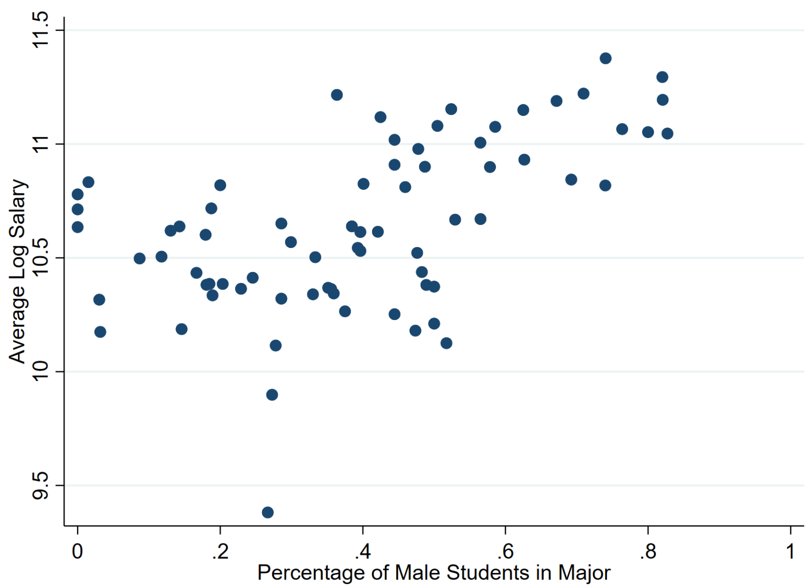

Figure 4: Gender Segregation of Industries

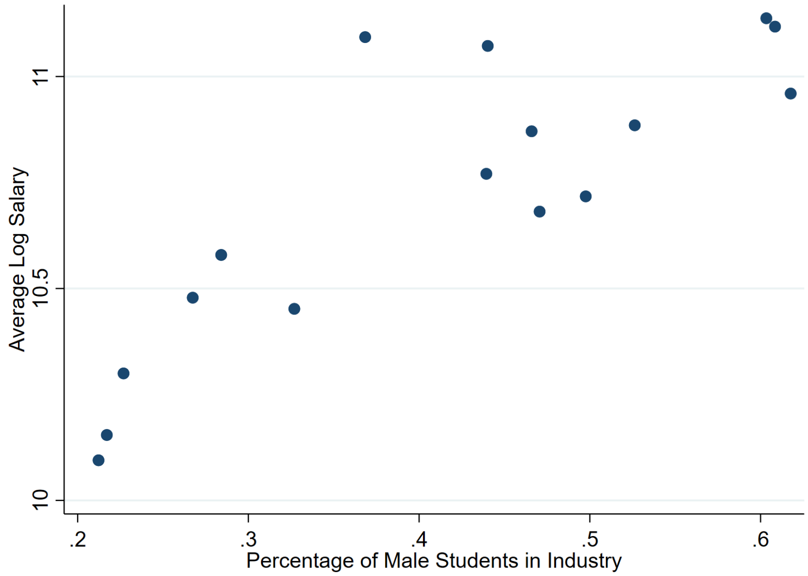

One limitation to my data is the absence of any individual-level measure of academic ability. In particular, I do not have access to student-level standardized test score or grade point average (GPA) data. However, within-major comparisons of average GPA between male and female graduates, provided in Table 6,[18] suggest that gender differences in academic performance are negligible. In fact, out of the largest 20 majors in my sample, women have higher average GPAs than men in all but two majors. Thus, all subsequent analysis assumes that academic performance is similar, on average, between male and female graduates. This assumption may not hold if male and female students within the same major tend to take different types of coursework. In particular, it is possible that, within majors, male students take courses that are more highly valued by employers. Since I do not have information on students' course records, including detailed information on the courses taken by male and female students, I do not examine this possibility. This is an area for future work to expand upon, should data linking academic performance and career outcomes become available.

[18]Summary statistics for this table were provided by UVA Institutional Research and Analytics, as I do not have permission to access raw data on student GPA. I do not have information on the standard deviation of GPA within each major, so I am unable to conduct two-sample t-tests to determine whether gender differences are statistically significant. Data is on 2018 graduates.

Table 6: Within-Major Gender Differences in Mean GPA

Major	Mean GPA Male	Female	Gender Gap[19]
Economics	3.4614	3.3905	-0.0709
Biology	3.3308	3.2884	-0.0424
Foreign Affairs	3.3196	3.1959	-0.1236
Psychology	3.3665	3.2778	-0.0887
Computer Science (B.S.)	3.3866	3.3666	-0.0200
Media Studies	3.4185	3.2623	-0.1562
Systems Engineering	3.4023	3.3709	-0.0315
English	3.4500	3.3176	-0.1324
Biomedical Engineering	3.4018	3.4329	0.0311
History	3.3664	3.1786	-0.1878
Mechanical Engineering	3.3826	3.3691	-0.0135
Computer Science (B.A)	3.4325	3.4498	0.0173
Public Policy and Leadership	3.6186	3.4586	-0.1600
Government	3.3956	3.2556	-0.1401
Chemistry	3.3255	3.3099	-0.0155
Cognitive Science	3.3431	3.2764	-0.0667
Global Studies	3.5612	3.5124	-0.0488
Civil Engineering	3.3123	3.1217	-0.1906
Commerce[20]	3.6046	3.6147	0.0101
Environmental Science	3.3397	3.2380	-0.1017

Two-sample t-tests are not conducted due to data limitations.

Included are the 20 largest majors, as in Table 4

[19] Gender Gap is defined as Mean Female GPA - Mean Male GPA

[20] In this data provided by UVA Institutional Research and Analytics, all majors within the McIntire

A final limitation to my data is the nontrivial proportion of missing values in certain fields of the First Destination Survey. Specifically, many graduates report academic and salary information but leave the questions regarding internship experience and career industry blank. To increase the statistical power of my analysis, I impute missing values for number of internships completed and career industry. I assume that graduates who leave the internship question blank have completed 0 internships. For graduates who leave the career industry question blank but indicate that they are working full-time, I classify their industry as "Other." I include in my regressions dummy variables indicating whether number of internships and career industry were imputed. The estimated coefficients of these dummy variables are not statistically significant, indicating that my imputation procedure does not introduce significant bias in my results. To further confirm this, I conduct a robustness exercise in which I exclude observations with missing values from my analysis. Results can be found in Appendix A.

5 Results

5.1 Decomposition Results

Table 7 presents results from pooled Oaxaca-Blinder and RIF decomposition models, corresponding to specifications (4) and (8) from Section 3. The gender gap in log salary is larger at the lower end of the pay distribution and cannot be entirely explained by observable gender differences in major, industry, and experience. In particular, the gap is largest around the 25th percentile of the salary distribution: At the 25th percentile, the log salary gap is approximately 0.356. Only 0.266, or 74.73%, of this gap can be explained by gender differences in observable endowments. This suggests that, within low-paying majors and industries, female graduates still enter lower-paying jobs immediately after graduation. This could be due in gender differences in preferences for very low-paying occupations, such as

School of Commerce are grouped together.

teaching. It could also be due to anticipated or realized labor market discrimination.

At the very top of the wage distribution, the gender gap in log salary narrows. Furthermore, the gap at the 90th percentile is "over-explained" by gender differences in endowments. One explanation for this result is that female graduates at the top of the pay distribution are more qualified than male graduates on dimensions not measured in the data. This is certainly plausible: As demonstrated in Table 6, women have higher average grade point averages than men in nearly all of the largest majors in my sample. Thus, they could be preferred by high-paying employers on the basis of demonstrated academic achievement. Alternatively, female students within high-paying majors could take coursework that is more highly valued by employers than the coursework taken by male students. Female students could also be more qualified on non-academic dimensions, such as interview skills, resume quality, and other "soft" skills.

An alternate interpretation for the "over-explanation" of the gender pay gap at the top of the salary distribution is that female graduates receive preferential labor market treatment. This is also somewhat likely, as high-paying industries such as finance and technology are known be extremely male-dominated. As companies attempt to increase gender diversity, some employers may give preferential labor market treatment to female graduates.

Finally, it is possible that the theory of compensating differentials, as discussed by Rosen (1986), explains my result at the top of the salary distribution. In particular, male and female graduates may have different preferences for high-paying positions in industries such as finance. If female graduates find the work involved in such industries to be unpleasant, then employers in these industries who seek to hire women will be forced to offer a salary premium in order to attract qualified female graduates.

Table 7: Pooled Decomposition of Male-Female Salary Differentials

		Gap Explained by					
Statistic	Raw Log-Salary Gap	Year	Ethnicity	Major	Industry	Internships	Total Explained
Mean	0.2492	0.0005	0.0024	0.1340	0.0997	0.0006	0.2421
10th Percentile	0.2654	-0.0011	0.0030	0.0770	0.1407	0.0010	0.2030
25th Percentile	0.3563	-0.0015	0.0037	0.1322	0.1400	0.0009	0.2663
Median	0.2314	0.0014	0.0003	0.1487	0.1400	0.0007	0.2688
75th Percentile	0.1078	0.0025	0.0001	0.1304	0.1141	0.0004	0.1786
90th Percentile	0.0932	0.0015	0.0003	0.1067	0.0410	0.0003	0.1514

n = 3649

5.2 STEM Major Choice and the Gender Wage Gap

One notable area in which female college students continue to lag behind males is STEM (science, technology, engineering, and mathematics). Within these fields, women are particularly underrepresented in mathematically intensive majors (Ceci et al., 2014). The gender gap in STEM degrees and in STEM careers has been studied extensively by policymakers and researchers. Due to the "STEM earnings premium" earned by STEM workers relative to their non-STEM counterparts (Beede, McKittrick, et al., 2011), it has been hypothesized that the gender gap in STEM majors could be a significant determinant of the gender pay gap among college-educated workers (Beede, Julian, et al., 2011).

I explore the impact of STEM major choice on the gender pay gap for UVA graduates by replacing the set of major dummy variables in my primary specification with a single STEM indicator variable. This variable is equal to 1 for students who majored in a STEM field and equal to 0 for those who did not.[21] Results from this specification are presented in Table 8.

I find that STEM major choice explains a greater proportion of the gender pay gap at the upper tail of the salary distribution: While STEM major choice explains only 0.9% of the gap at the 25th percentile, it explains 15.2% of the gap at the 90th percentile. Across the earnings distribution, the explanatory power of the STEM indicator is quite small in comparison to the full set of major variables. At the 25th percentile, just 2.3% of the salary variation attributable to major choice can be explained by the STEM major indicator; even at the 90th percentile, STEM major choice explains only 13.3% of the salary variation attributable to major choice. My results suggest that broadly encouraging female college students to pursue STEM fields obscures important gender differences in major choice within STEM. Thus, attempts to narrow the gender gap in wages through encouraging women to pursue STEM fields must be focused on particular majors within the STEM categorization.

[21] I define the following majors as STEM majors: Aerospace Engineering; Astronomy; Astronomy-Physics; Biology; Biomedical Engineering; Chemical Engineering; Chemistry; Civil Engineering; Cognitive Science; Computer Engineering; Computer Science (B.A.); Computer Science (B.S.); Electrical Engineering; Engineering Science; Environmental Sciences; Materials Science and Engineering; Mathematics; Mechanical Engineering; Neuroscience; Nursing; Physics; Statistics; and Systems Engineering.

Table 8: Pooled Decomposition of Male-Female Salary Differentials using STEM Major Indicator

		Gap Explained by					
Statistic	Raw Log-Salary Gap	Year	Ethnicity	STEM Major	Industry	Internships	Total Explained
Mean	0.2492	0.0021	0.0010	0.0182	0.1402	-0.0002	0.1802
10th Percentile	0.2654	-0.0014	0.0019	0.0082	0.1690	-0.0001	0.1742
25th Percentile	0.3563	-0.0031	0.0031	0.0023	0.1901	-0.0002	0.1990
Median	0.2314	0.0022	-0.0008	0.0187	0.1601	-0.0001	0.1995
75th Percentile	0.1078	0.0042	-0.0007	0.0241	0.0731	-0.0001	0.1159
90th Percentile	0.0932	0.0033	-0.0010	0.0142	0.0687	-0.0000	0.0945

n = 3649

5.3 Gender Differences in Labor Force Participation

A concern in the gender wage gap literature is that male and female graduates' propensity to join the labor force may differ due to gender differences in unobservable factors. In particular, men and women may receive different distributions of salary offers (Xiu & Gunderson, 2014). If this is the case, and students have a salary threshold below which they will not accept a job offer, then male and female graduates will differ in their propensity to join the labor force immediately after graduation. Since I can only observe salary offers that students accept, I am unable to control for gender differences in the distribution of salary offers received.

Alternatively, male and female graduates may receive similar distributions of salary offers but have different salary thresholds that they apply when accepting or rejecting job offers. For example, if more female graduates than male graduates expect to gain access to spousal income soon after graduation, then female graduates may set higher wage thresholds to accept a job offer (Ge et al., 2018). This would imply that, even if men and women received similar distributions of salary offers, women would enter the labor force at a lower rate based on spousal income expectations that I cannot observe. (Hersch, 2013) demonstrates that this phenomenon of "opting out" is more common for female graduates of elite universities, such as the University of Virginia. (Bertrand et al., 2010), however, find that opting out of the labor market occurs later in life, and that labor market participation rates are similar between male and female MBA graduates immediately after graduation.

If men and women differ in their propensity to join the labor force based on gender differences in factors that I am unable to observe in the data, then the samples of male and female graduates who have opted in to full-time employment may have had systematically different salary offers than the general populations of male and female University of Virginia students. This would necessitate some sort of correction procedure, such as those proposed by Fang and Sakellariou (2011) and Gunewardena et al. (2008), to mitigate selection bias in my decomposition results. I argue that such a correction procedure is unnecessary. A recent report by the Economic Policy Institute (Gould et al., 2019) finds that for United States college graduates in the class of 2019, gender differences in labor market participation immediately after graduation are extremely small.

Specifically in my sample, gender differences in labor force participation are quite small. While the raw gender gap in labor market participation is significant at the 1% confidence level, I show using a simple linear probability model that the gap can be accounted for entirely by college major. The model I use is as follows, where $working_i$ is an indicator for whether individual i is employed full-time after graduation; $male_i$ is an indicator for whether individual i is male; and γ_i are major fixed effects, a set of indicators for each major offered at the University of Virginia. ϵ_i represents the error term.

$$working_i = \beta_0 + \beta_1 male_i + \beta_2 \gamma_i + \epsilon_i \tag{9}$$

Estimating this model yields a coefficient of $\hat{\beta}_1 = -.046$, significant at the 1% confidence level. This result indicates that within majors, female graduates are more likely than male graduates to be employed full-time immediately after graduation.

Next, I replace $working_i$ in equation (9) with $education_i$, an indicator for whether individual i pursues higher education immediately after graduation.

$$education_i = \beta_0 + \beta_1 male_i + \beta_2 \gamma_i + \epsilon_i \tag{10}$$

Estimating this model yields insignificant results for $\hat{\beta}_1$, suggesting that, within majors, male and female graduates are equally likely to pursue higher education. Taken together, the results of these two models along with the significant raw gender gap in labor force participation indicate that female students disproportionately choose majors that are less likely to lead to full-time employment immediately after graduation and more likely to lead to higher education.

6 Discussion and Conclusion

In this thesis, I document a novel result contradicting prior literature on the gender pay gap: Female graduates of the University of Virginia do not face a glass ceiling immediately after graduation. Rather, gender differences in college major and career industry choice suggest that the gender gap at the top of the salary distribution should be even larger than what is observed in the data. This suggests the possible existence of a "glass escalator" for female graduates, wherein women receive preferential treatment when entering traditionally male-dominated professions.[22]

Previous studies conducted in developed western nations tend to confirm the existence of a glass ceiling, as evidenced by large unexplained wage gaps at the upper end of the salary distribution or slowly falling wage gaps at the top of the salary distribution (Albrecht et al., 2003; Albrecht et al., 2009; Arulampalam et al., 2007; Blau & Kahn, 2006, 2017; Christofides et al., 2010; Chzhen & Mumford, 2009; de la Rica et al., 2008; del Rio et al., 2011; Jellal et al., 2008; Kassenboehmer & Sinning, 2014; Kee, 2006). Studies finding evidence of glass ceilings have used data on workers across the life cycle, typically restricting analysis to workers aged 18-65, or until retirement age. I restrict my analysis to recent bachelor's degree recipients from a single university, meaning my sample is much younger on average than in previous literature on the glass ceiling. It is well established in the literature that the gender wage gap increases with age. Using data on United States workers from 1978-1998, Blau and Kahn (2001) find that the pay gap increased from ages 25-34 to ages 35-44 and later stabilized or even decreased. Goldin (2015) has documented a similar life cycle pattern using data on college-educated workers born in the United States in the late 1940s and early 1950s.

The widening of the gender wage gap roughly between the ages of 25 and 45 suggests that responsibilities associated with marriage and family as well as barriers to promotion

[22]The term "glass escalator," coined by sociologist Christine Williams in 1992, is typically used in situations where men are promoted more quickly than women in traditionally female-dominated professions (Williams, 1992). My usage, in addition to reversing the gender roles typically associated with the term, does not imply anything about differences in promotions between men and women after their initial salary offers.

may contribute to the glass ceiling for female workers. Both phenomena have been well documented empirically. Early work by Waldfogel (1998) on pay differences between mothers and childless women — often called the "family gap" — finds that differences in the impacts of marriage of parenthood on earnings account for a large portion of the gender gap in earnings. More recent work has found evidence that the family gap in wages persists (Budig & England, 2001; Simonsen & Skipper, 2006) and that delaying motherhood is associated with career benefits (Hotz et al., 2005; Klepinger et al., 1999; A. Miller, 2011; Taniguchi, 1999). Other literature finds that expectations of long hours and continuous work force attachment contribute significantly to the underrepresentation of women in high-paying corporate leadership positions (Bertrand et al., 2010; Bertrand & Hallock, 2001).

As a converse to the glass ceiling, another phenomenon that could emerge later in women's life cycles is the "sticky floor." Sticky floors, where female workers are stuck at the bottom of the salary distribution, could occur due to an excess supply of female workers who are excluded from high-paying jobs and consequently crowded into low-wage jobs. Many of the same family-related factors contributing to the glass ceiling in high-paying corporate jobs may also contribute to the sticky floor effect. This seems to be the case in many Asian countries, but not necessarily in developed nations (Xiu & Gunderson, 2014). There is, however, some evidence that certain types of workers in western countries experience the sticky floor effect (Christofides et al., 2010; de la Rica et al., 2008; del Rio et al., 2011; P. Miller, 2009).

Given this extensive literature on women's life cycle earnings, my results are perhaps less surprising. While the gender pay gap between UVA graduates at the start of their careers can largely be explained by gender differences in major and industry self-selection, it is likely that the unexplained portion of the gap widens over time. Future work should analyze graduates' career outcomes further after graduation to obtain a better understanding of gender differences in UVA graduates' life cycle earnings. Such analysis, however, will be greatly limited by available data sources, as the University of Virginia Career Center does not currently track graduates' career outcomes further than one year after graduation.

This work demonstrates the significant contribution of pre-market specialization to gender differences in salary outcomes. As such, my results have important policy implications for universities striving to increase gender equity in outcomes for their graduates. In particular, faculty advisers, career counselors, and other professionals who advise students on major selection and career choice can play an important role in informing students about the connection between gender segregation in majors and labor market outcomes (Lapan, 2018). In the future, the collection of more comprehensive salary data by the University of Virginia can aid researchers in better determining if, when, and why glass ceilings and sticky floors emerge for UVA graduates.

Appendices

A Robustness Checks

In my main specification, I include one set of major fixed effects, with one indicator variable for each major offered at the University of Virginia. For example, the indicator variable $Mathematics_i$ is equal to one if student i reports Mathematics as his or her primary or secondary major. It is possible, however, that students' self-reported major ranking matters. That is, students who report a major as their secondary major may have systematically different career outcomes than those who indicate that major as their primary field of study. This could be due to the procedures employers use to screen applicants by major, differences in chosen coursework rigor between primary and secondary majors, or differences between the types of students students who indicate certain majors as their primary versus secondary field. To see if the unranked major definitions used in my main specification bias my results, I estimate a secondary specification with two separate sets of major fixed effects, one for primary major and one for secondary major. Appendix Table A1 reports the results of this robustness exercise. I find that the results are largely unchanged, although slightly more of the gap at the lower end of the salary distribution is explained by the inclusion of ranked majors.

I also conduct a robustness check on my imputation procedure for observations missing industry and internship values, as described in Section 4. Results are presented in Appendix Table A2. Excluding observations with missing values reduces my sample size from 3,649 to 2,563 observations but leaves the gender pay gap similar in magnitude across the salary distribution. A larger portion of the gap is explained in this sample, but estimates are less precise.

Table A1: Pooled Decomposition of Male-Female Salary Differentials with Ranked Majors

		Gap Explained by					
Statistic	Raw Log-Salary Gap	Ethnicity	Major 1	Major 2	Industry	Internships	Total Explained
Mean	0.2485	0.0023	0.1237	0.0230	0.0951	0.0003	0.2444
10th Percentile	0.2635	0.0044	0.0511	0.0276	0.1410	0.0006	0.2247
25th Percentile	0.3559	0.0052	0.1074	0.0185	0.1414	0.0012	0.2736
Median	0.2315	0.0007	0.1295	0.0188	0.1115	0.0008	0.2613
75th Percentile	0.1765	0.0002	0.1236	0.0092	0.0385	0.0003	0.1718
90th Percentile	0.0937	0.0002	0.0871	0.0262	0.0412	0.0002	0.1549

n = 3712

Table A2: Pooled Decomposition of Male-Female Salary Differentials, Excluding Missing Values

		Gap Explained by					
Statistic	Raw Log-Salary Gap	Age	Ethnicity	Major	Industry	Internships	Total Explained
Mean	0.2579	0.0023	0.0020	0.1292	0.1344	-0.0004	0.2674
10th Percentile	0.2634	-0.0010	0.0044	0.0716	0.1967	-0.0004	0.2714
25th Percentile	0.3481	-0.0007	0.0044	0.1150	0.1582	-0.0005	0.2763
Median	0.2494	0.0026	0.0008	0.1546	0.1265	-0.0005	0.2840
75th Percentile	0.1480	0.0027	-0.0003	0.1660	0.0574	-0.0003	0.2256
90th Percentile	0.1240	0.0008	0.0010	0.1264	0.0542	-0.0002	0.1822

n = 2563

References

Agopsowicz, A., Robinson, C., Stinebrickner, R., & Stinebrickner, T. (2017). *Careers and mismatch for college graduates: College and non-college jobs* (tech. rep. No. 4). Centre for Human Capital and Productivity (CHCP).

Albrecht, J., Bjorklund, A., & Vroman, S. (2003). Is There a Glass Ceiling in Sweden? *Journal of Labor Economics*, *21*(1), 145–177.

Albrecht, J., Vuuren, A. V., & Vroman, S. (2009). Counterfactual distributions with sample selection adjustments: econometric theory and an application to the Netherlands. *Labor Economics*, *16*(4), 383–396.

Altonji, J. (1993). The Demand for and Return to Education when Education Outcomes are Uncertain. *Journal of Labor Economics*, *11*(1), 48–83.

Arrow, K. (1972). Some Models of Racial Discrimination in the Labor Market (A. Pascal, Ed.). In A. Pascal (Ed.), *Racial discrimination in economic life*. D.C. Heath.

Arulampalam, W., Booth, A., & Bryan, M. (2007). Is There a Glass Ceiling over Europe? Exploring the Gender Pay Gap Across the Wage Distribution. *Industrial and Labor Relations Review*, *60*(2), 163–186.

Babcock, L., & Laschever, S. (2003). *Women don't ask: Negotiation and the gender divide*. Princeton University Press.

Becker, G. (1964). *Human capital*. Columbia University Press.

Beede, D., Julian, T., Langdon, D., McKittrick, G., Khan, B., & Doms, M. (2011). *Women in STEM: A Gender Gap to Innovation* (Issue Brief No. 04-11). Economics and Statistics Administration, U.S. Department of Commerce.

Beede, D., McKittrick, G., Langdon, D., Khan, B., & Doms, M. (2011). *STEM: Good Jobs Now and for the Future* (Issue Brief No. 03-11). Economics and Statistics Administration, U.S. Department of Commerce.

Bertrand, M., Goldin, C., & Katz, L. F. (2010). Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors. *American Economic Journal: Applied Economics*, *2*(3), 228–255.

Bertrand, M., & Hallock, K. (2001). The Gender Gap in Corporate Jobs. *Industrial and Labor Relations Review*, *55*(1), 3–21.

Black, D., Haviland, A., Sanders, S., & Taylor, L. (2008). Gender Wage Disparities among the Highly Educated. *Journal of Human Resources*, *43*(3), 630–659.

Blau, F., & Kahn, L. (2001). Gender Differences in Pay. *Journal of Economic Perspectives*, *14*(4), 75–99.

Blau, F., & Kahn, L. (2006). The U.S. Gender Pay Gap in the 1990s: Slowing Convergence. *Industrial and Labor Relations Review*, *60*(1), 45–66.

Blau, F., & Kahn, L. (2017). The Gender Wage Gap: Extent, Trends, and Explanations. *Journal of Economic Literature*, *55*(3), 789–865.

Blinder, A. (1973). Wage Discrimination: Reduced Form and Structural Estimates. *The Journal of Human Resources*, *8*(4), 436–455.

Bobbitt-Zeher, D. (2007). The Gender Income Gap and the Role of Education. *Sociology of Education*, *80*(1), 1–22.

Brown, C., & Corcoran, M. (1997). Sex-based differences in school content and the male-female wage gap. *Journal of Labor Economics*, *15*(3), 431–465.

Budig, M., & England, P. (2001). The wage penalty for motherhood. *American Sociological Review*, *66*(3), 204–225.

Carrillo, P., Gandelman, N., & Robano, V. (2014). Sticky floors and glass ceilings in Latin America. *Journal of Economic Inequality*, *12*, 339–361.

Ceci, S., Ginther, D., Kahn, S., & Williams, W. (2014). Women in Academic Science: A Changing Landscape. *Psychological Science*, *15*(3), 75–141.

Chamberlain, A. (2016). *Demystifying the Gender Pay Gap: Evidence From Glassdoor Salary Data* (Research Report). Glassdoor.

Chi, W., & Li, B. (2008). Glass ceiling or sticky floor? Examining the gender earnings differential across the earnings distribution in urban China, 1987–2004. *Journal of Comparative Economics*, *36*(2), 243–263.

Christofides, L., Polycarpou, A., & Vrachimis, K. (2010). *The Gender Wage Gaps, 'Sticky Floors' and 'Glass Ceilings' of the European Union* (Discussion Paper No. 5044). Institute of Labor Economics (IZA).

Chzhen, Y., & Mumford, K. (2009). *Gender Gaps across the Earnings Distribution in Britain: Are Women Bossy Enough?* (Discussion Paper No. 4331). Institute of Labor Economics (IZA).

Corcoran, M., & Courant, P. (1985). Sex Role Socialization and Labor Market Outcomes. *The American Economic Review*, *75*(2), 275–278.

de la Rica, S., Dolado, J., & Llorens, V. (2008). Ceilings or floors? Gender wage gaps by education in Spain. *Journal of Population Economics*, *21*(3), 751–776.

del Rio, C., Gradin, C., & Canto, O. (2011). The measurement of gender wage discrimination: the distributional approach revisited. *Journal of Economic Inequality*, *9*(1), 57–86.

Eccles, J., & Hoffman, L. (1984). Sex Differences in Preparation for Occupational Roles (H. Stevenson & A. Siegel, Eds.). In H. Stevenson & A. Siegel (Eds.), *Child development and social policy*. University of Chicago Press.

Fang, Z., & Sakellariou, C. (2011). A case of sticky floors: gender wage differentials in Thailand. *Asian Economic Journal*, *25*(1), 35–54.

Firpo, S., Fortin, N., & Lemieux, T. (2009). Unconditional Quantile Regressions. *Econometrica*, *77*(3), 953–973.

Firpo, S., Fortin, N., & Lemieux, T. (2018). Decomposing Wage Distributions Using Recentered Influence Function Regressions. *Econometrics*, *6*(2), 1–40.

Foote, A., & Stange, K. (2019). Attrition from administrative data: Problems and solutions with an application to higher education. NBER Education Program.

Fortin, N., Lemieux, T., & Firpo, S. (2011). Decomposition methods in economics (O. Ashenfelter & D. Card, Eds.; 1st ed.). In O. Ashenfelter & D. Card (Eds.), *Handbook of labor economics* (1st ed.). Elsevier.

Ge, S., Isaac, E., & Miller, A. (2018). *Elite Schools and Opting In: Effects of College Selectivity on Career and Family Outcomes* (Working Paper No. 25315). NBER.

Goldin, C. (2005). The Quiet Revolution That Transformed Women's Employment, Education, and Family. *American Economic Journal*, *96*(2), 1–21.

Goldin, C. (2015). A Grand Gender Convergence: Its Last Chapter. *American Economic Review*, *104*(4), 1091–1119.

Goldin, C., Katz, L. F., & Kuziemko, I. (2006). The Homecoming of American College Women: The Reversal of the College Gender Gap. *Journal of Economic Perspectives*, *20*(4), 133–156.

Gould, E., Mokhiber, Z., & Wolfe, J. (2019). *Class of 2019: College edition* (Report). Economic Policy Institute.

Graham, M., Hotchkins, J., & Gerhart, B. (2000). Discrimination by Parts: A Fixed-Effects Analysis of Starting Pay Differences across Gender. *Eastern Economic Journal*, *26*(1), 9–27.

Gunewardena, D., Abeyrathna, D., Ellagala, A., Rajakaruna, K., & Rajendran, S. (2008). *Glass ceiling, sticky floors or sticky doors? A quantile regression approach to exploring gender wage gaps in Sri Lanka* (Working Paper). Poverty and Economic Policy Research Network.

Hersch, J. (2013). Opting out among women with elite education. *Review of Economics of the Household*, *11*(4), 469–506.

Hotz, V., McElroy, S. W., & Sanders, S. (2005). Teenage childbearing and its life cycle consequences: exploiting a natural experiment. *Journal of Human Resources*, *60*(3), 683–715.

Hyde, J., Lindberg, S., Linn, M., Ellis, A., & Williams, C. (2008). Gender Similarities Characterize Math Performance. *Science*, *321*(5888), 494–495.

Jann, B. (2008). The Blinder-Oaxaca Decomposition for linear regression models. *The Stata Journal*, *8*(4), 453–479.

Jellal, M., Nordman, C., & Wolff, F.-C. (2008). Evidence on the glass ceiling effect in France using matched worker-firm data. *Applied Economics*, *40*(24), 3233–3250.

Joy, L. (2003). Salaries of Recent Male and Female Graduates: Educational and Labor Market Effects. *Industrial and Labor Relations Review*, *56*(4), 606–621.

Kassenboehmer, S., & Sinning, M. (2014). Distributional Changes in the Gender Wage Gap. *Industrial and Labor Relations Review*, *67*(2), 335–361.

Kee, H. J. (2006). Glass ceiling or sticky floor? Exploring the Australian gender pay gap. *The Economic Record*, *82*(259), 408–427.

Klepinger, D., Lundberg, S., & Plotnick, R. (1999). How does adolescent fertility affect the human capital and wages of young women? *Journal of Human Resources*, *34*(3), 421–448.

Lapan, J. (2018). Equal Opportunity, Unequal Outcomes: Exploring Gender Inequality in Post-College Career Outcomes. *National Association of Colleges and Employers Journal*.

Lazear, E., & Rosen, S. (1990). Male-Female Wage Differentials in Job Ladders. *Journal of Labor Economics*, *8*(1), S106–S123.

Loury, L. (1997). The Gender Earnings Gap Among College-Educated Workers. *Industrial and Labor Relations Review*, *50*(4), 580–593.

Lucifora, C., & Meurs, D. (2006). The Public Sector Pay Gap in France, Great Britain, and Italy. *Review of Income and Wealth*, *52*(1), 43–59.

Machado, J., & Mata, J. (2005). Counterfactual Decomposition of Changes in Wage Distributions Using Quantile Regression. *Journal of Applied Econometrics*, *20*(4), 445–465.

McDonald, J., & Thornton, R. (2007). Do New Male and Female College Graduates Receive Unequal Pay? *The Journal of Human Resources*, *42*(1), 32–48.

Melly, B. (2005). Decomposition of differences in distribution using quantile regression. *Labor Economics*, *12*(4), 577–590.

Miller, A. (2011). The effects of motherhood timing on career path. *Journal of Population Economics*, *24*, 1071–1100.

Miller, P. (2009). The gender pay gap in the US: does sector make a difference? *Journal of Labor Research*, *30*(1), 52–74.

Oaxaca, R. (1973). Male-female wage differentials in urban labor markets. *International Economic Review*, *14*(3), 693–709.

Paglin, M., & Rufolo, A. M. (1990). Heterogeneous Human Capital, Occupational Choice, and Male-Female Earnings Differences. *Journal of Labor Economics*, *8*(1), 123–144.

Riegle-Crumb, C., King, B., Grodsky, E., & Muller, C. (2012). The More Things Change, the More They Stay the Same? Prior Achievement Fails to Explain Gender Inequality in Entry Into STEM College Majors Over Time. *American Educational Research Journal*, *49*(6), 1048–1073.

Rios-Avila, F. (2019). *Recentered Influence Functions in Stata: Methods for Analyzing the Determinants of Poverty and Inequality* (Working Paper No. 927). Levy Economics Institute of Bard College.

Rosen, S. (1986). The theory of equalizing differences (O. Ashenfelter & R. Layard, Eds.). In O. Ashenfelter & R. Layard (Eds.), *Handbook of labor economics*. North Holland Publishing Company.

Roy, A. (1951). Some Thoughts on the Distribution of Earnings. *Oxford Economic Papers*, *3*(2), 135–146.

Shauman, K. (2016). Gender Differences in the Early Career Outcomes of College Graduates: The Influence of Sex-Type of Degree Field Across Four Cohorts. *RSF: The Russell Sage Foundation Journal of the Social Sciences*, *2*(4), 152–193.

Simonsen, M., & Skipper, L. (2006). The costs of motherhood: an analysis using matching estimators. *Journal of Applied Economics*, *21*(7), 919–934.

Sloane, C., Hurst, E., & Black, D. (2019). *A Cross-Cohort Analysis of Human Capital Specialization and the College Gender Wage Gap* (Working Paper No. 2019-121). Becker Friedman Institute for Economics.

Stinebrickner, R., & Stinebrickner, T. (2014). Initial Beliefs and Final Outcomes for College Major and Dropout. *The Review of Economic Studies*, *81*(1), 426–472.

Stinebrickner, R., Stinebrickner, T., & Sullivan, P. (2019). Beauty, Job Tasks, and Wages: A New Conclusion about Employer Taste-Based Discrimination. *The Review of Economics and Statistics*, *101*(4), 602–615.

Stinebrickner, T., & Stinebrickner, R. (2011). *Math or Science? Using Longitudinal Expectations Data to Examine the Process of Choosing a College Major* (Working Paper No. 16869). NBER.

Taniguchi, H. (1999). Work and family-the timing of childbearing and women's wages. *Journal of Marriage and Family*, *61*(4), 1008–1020.

Turner, S., & Bowen, W. (1999). Choice of Major: The Changing (Unchanging) Gender Gap. *Industrial and Labor Relations Review*, *52*(2), 289–313.

Waldfogel, J. (1998). Understanding the "Family Gap" in Pay for Women with Children. *Journal of Economic Perspectives*, *12*(1), 137–156.

Williams, C. (1992). The Glass Escalator: Hidden Advantages for Men in the "Female" Professions. *Social Problems*, *39*(3), 253–267.

Willis, R., & Rosen, S. (1979). Education and Self-Selection. *Journal of Political Economy*, *87*(5), S7–S36.

Wiswall, M., & Zafar, B. (2014). Determinants of College Major Choice: Identification Using an Information Experiment. *The Review of Economic Studies*, *82*(2), 791–824.

Xiu, L., & Gunderson, M. (2014). Glass ceiling or sticky floor? Quantile regression decomposition of the gender pay gap in China. *International Journal of Manpower*, *35*(3), 306–326.

Zafar, B. (2012). College Major Choice and the Gender Gap. *The Journal of Human Resources*, *48*(3), 545–595.