UNIVERSITY OF CALIFORNIA,
IRVINE

Three Essays in the Economics of Discrimination

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Economics

by

Ian Burn

Dissertation Committee:
Professor David Neumark, Chair
Professor John Duffy
Associate Professor Damon Clark

2017

DEDICATION

To my friends and family, who were always understanding of their place in my work-life balance. Your support made the last five years possible.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

Page

ACKNOWLEDGMENTS

I need to thank my dissertation committee for their help. First, I need to thank my committee chair, David Neumark. His advice was instrumental in shaping this dissertation and smoothing out its rough edges. His rapid-fire questions kept me on my toes, but always yielded valuable insights that made my research better. I am also indebted to John Duffy and Damon Clark. Without their classes and feedback, my researcher's toolkit would be sparse. To the Department of Economics and the University of California, thank you for the generous financial support over the years.

CURRICULUM VITAE

Ian Burn

EDUCATION

Doctor of Philosophy in Economics — **2017**
University of California, Irvine — *Irvine, CA*

Master of Arts in Applied Economics — **2012**
Northeastern University — *Boston, MA*

Bachelor of Science in Economics — **2011**
Northeastern University — *Boston, MA*

ABSTRACT OF THE DISSERTATION

Three Essays in the Economics of Discrimination

By

Ian Burn

Doctor of Philosophy in Economics

University of California, Irvine, 2017

Professor David Neumark, Chair

This dissertation explores the causes, consequences, and remedies of wage penalties in the labor market. The first chapter explores the relationship between prejudice and wages for gay men in the United States. I show that search models of taste-based discrimination can predict the empirical relationship between prejudice towards gay men and their wages. The second chapter explores how individuals use wage penalties when deciding which college major to select. Using a laboratory experiment, I show that higher female wage penalties in the labor market deter female students from selecting a major. Since female students expect discrimination to be worse in STEM fields, this preference for majors with lower wage penalties leads to a gender participation gap in STEM. My experiment showed that correcting misinformation about wage penalties in the labor market can increase female interest in STEM majors. The final chapter explores how effective public policy has been at reducing the wage penalty against gay men and lesbian women in the United States. I show that the heterogeneous nature of state-level employment non-discrimination laws in the United States has important consequences for their effectiveness. Stronger laws are more effective at reducing the wage penalty for gay men but may lead to lower levels of employment for lesbian women. Weak laws in the United States had no effect on the labor market outcomes of gay men and lesbian women.

Chapter 1

Pride, Prejudice, and Wages

1.1 Introduction

For the past 20 years, economists have documented differences in the labor market outcomes of gay men and heterosexual men in the United States. Beginning with Badgett (1995) studies have consistently found evidence of significant unexplained negative wage differentials for gay men after controlling for observable characteristics (Klawitter 2015). Data from the U.S. Census suggests that gay men experience a wage penalty of between 11% and 15% (Klawitter 2015).[1] In addition to the well-documented wage penalty, there is evidence of prejudice against gay men (Badgett, Lau, Sears and Ho 2007, Pew Research Center 2013). As a result of this prejudice, 27% of gay men reported experiencing harassment or discrimination in the workplace due to their sexual orientation (Sears and Mallory 2011).

The existence of both an unexplained wage penalty and sizable prejudice towards gay men would seem to imply prejudice leads to discrimination, resulting in the observed wage penalty.

[1] There is little evidence of wage discrimination against lesbians (Klawitter 2015). Research has shown the lesbian wage premium can be explained by non-discriminatory factors (Antecol and Steinberger 2013, Daneshvary, Wassoups and Wimmer 2009, Jepsen 2007, Jepsen and Jepsen 2015). Therefore, the focus of this paper is on the wage penalty between gay men and heterosexual men.

The Becker model of discrimination presents the most well-known mechanism by which prejudicial attitudes lead to differences in labor market outcomes that are interpreted as discriminatory (Becker 1971). In the Becker model, employers and workers can perfectly sort themselves, so the least prejudiced employers hire the minority workers. The wage penalty in the Becker model is determined by the prejudice of the "marginal employer," who is the most prejudiced employer to hire a minority worker. Charles and Guryan (2008) showed the Becker model was consistent with the relationship between racial prejudice and the black wage penalty in the United States. They argue that their estimates show a quarter of the racial wage gap is due to prejudice.

Given that the predictions of the Becker model are consistent with the empirical facts of the black wage penalty, one might expect similar results for the gay wage penalty. However, data from the General Social Survey suggests that the predictions of the Becker model are inconsistent with the empirics of the gay wage penalty. The lack of search frictions in the Becker model means there is no wage penalty as long as there are enough unprejudiced employers (i.e. employers who would pay their gay workers a wage equal to their marginal product of labor) to hire all the gay workers. In the United States, 3% of men are gay or bisexual (Gates and Newton 2013). So if more than 3% of employers are unprejudiced, the wage penalty should be zero. Data from the General Social Survey suggest that 16% of Americans are unprejudiced against homosexuals, which would imply no wage discrimination towards gay men in the Becker model.[2]

If prejudice is the cause of the wage penalty, the wages of gay men must be responding to the prejudice of employers the Becker model does not consider important. Search models of discrimination with sequential search (e.g. Black (1995) and Bowlus and Eckstein (2002)) allow for search frictions in the labor market and have workers search for jobs from randomly matched employers. Therefore, gay workers are unable to select only unprejudiced employers

[2]Unprejudiced means giving the most favorable answers to all questions about homosexuality in the General Social Survey.

to interview with. Because searching for a job is costly, gay workers will accept lower wages from unprejudiced employers to avoid more periods of matches with prejudiced employers. The result is that the wage penalty is not determined by the marginal employer, as occurs in the Becker model. Instead, it is determined by the share of employers who would never hire a gay man.

Motivated by contradictions between the gay wage penalty and the Becker model, I go beyond the framework developed by Charles and Guryan (2008) and test how well the search model of discrimination predicts the empirical relationship between prejudice against gay men and their wages. I provide some of the first evidence that search frictions play an important role in the development of a wage penalty. To date, there has been no published work empirically testing the ability of the search model to explain wage penalties using survey data on prejudicial attitudes.[3] Earlier work estimated the share of employers that are biased against minorities by calibrating theoretical search models and testing the relationships between labor market outcomes (Flabbi and Tejada 2015, Rosen 2003). While these papers estimated the parameters needed to approximate the dynamics of the labor market, they did not consider the prejudice of real individuals.

To test the relationship between observed prejudice and the wages of gay men, I follow the methodology of Charles and Guryan (2008). I construct a distribution of prejudice against gay men for each state in the United States using responses from the General Social Survey. If the search model of discrimination is accurate, states with more prejudiced populations should have larger wage penalties for gay men. I show that gay wage penalties are positively correlated with the prejudiced share of the population in a state. I find a one standard deviation increase in the percent of the population that is prejudiced towards gay men is associated with a 2.7 percentage point increase in the wage penalty for gay men.

[3] A recent working paper by Bond and Lehmann (2015) uses General Social Survey data to test a search model of discrimination for African Americans in the United States. The authors find evidence that the share of prejudiced employers is correlated with wage penalties and lower match quality.

1.2 Prejudice against Gay Men in the United States

The General Social Survey began tracking prejudice towards homosexuals in the 1970s. It is a nationally representative survey administered every two years.[4] Table 1.1 lists the questions asked in the General Social Survey about homosexuals. The first four questions in Table 1.1 are asked in every wave of the survey. The last two questions were recently added.

The first question, SEX, asks whether respondents think sexual relations between two adults of the same sex is wrong. The next two questions, BOOK and SPEAK, touch on support for speech in favor of homosexuality. BOOK asks respondents if they would support removing books in favor of homosexuality from their public library. SPEAK asks respondents if an admitted homosexual should be allowed to make a speech in public. COLLEGE asks respondents if homosexuals should be allowed to teach in colleges. The questions added to the General Social Survey reflect the changing societal concerns about homosexuality. MARRIAGE was added in 2006 to gauge respondents' views about same-sex marriage. CHILD was added in 2012 and asks respondents if a same-sex male couple can raise a child as well as a heterosexual couple.

[4] Between 1977 and 1994, the General Social Survey was administered every year with a few exceptions. In 1994, it switched to being every two years. See https://www.gssdataexplorer.norc.org for publicly available data and documentation from the General Social Survey.

Table 1.1: Questions from the General Social Survey

Question	Question Text
SEX	What about sexual relations between two adults of the same sex– do you think it is always wrong, almost always wrong, wrong only sometimes, or not wrong at all? Asked between 1990 and 2014 GSS Mnemonic: HOMOSEX
BOOK	If some people in your community suggested that a book he wrote in favor of homosexuality should be taken out of your public library, would you favor removing this book, or not? Asked between 1990 and 2014 GSS Mnemonic: LIBHOMO
SPEAK	Suppose this admitted homosexual wanted to make a speech in your community. Should he be allowed to speak, or not? Asked between 1990 and 2014 GSS Mnemonic: SPKHOMO
COLLEGE	And what about a man who admits that he is a homosexual? Should such a person be allowed to teach in a college or university, or not? Asked between 1990 and 2014 GSS Mnemonic: COLHOMO
MARRIAGE	Do you agree or disagree? Homosexual couples should have the right to marry one another. Asked between 2006 and 2014 GSS Mnemonic: MARHOMO
CHILD	To what extent do you agree or disagree with the following statements? A same-sex male couple can bring up a child as well as a male-female couple. GSS Mnemonic: SSMCHILD Asked in 2012

Note: Questions come from the pooled General Social Survey, 1990 to 2014.

There has been a steady decline in prejudiced responses since 1990 in each question. Figure 1.1 highlights how prejudiced responses have changed. Sexual relations and marriage between same-sex adults attract the highest disapproval in the General Social Survey in every survey. Between 1990 and 2014, the share of respondents that felt that sexual relations between two adults of the same sex was always wrong fell from 73% to 39%. In 2006, 51% of respondents disapproved of same-sex marriage. By 2014, only 31% of respondents disapproved. There has always been less support for banning books about homosexuals from public libraries, banning gay men from speaking in public, and banning gay men from teaching in colleges. In 1990, approximately a third of respondents supported these positions, but fewer than 20% of respondents supported them in 2014.

Figure 1.1: Share of Prejudiced Responses in the GSS by Year: 1990 to 2014

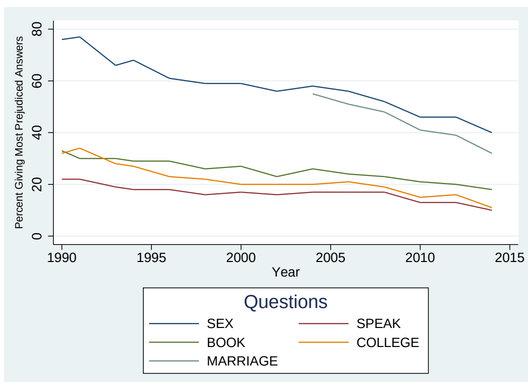

Note: Data on prejudice come from the pooled General Social Survey, 1990 to 2014. CHILD is excluded because it was only asked in 2012. See Table 3.6 for the text of each question.

1.3 Evidence of a Gay Wage Penalty

The most common sources of data used to calculate the gay wage penalty are the General Social Survey, the U.S. Decennial Census, and the Current Population Survey. The General Social Survey is the most detailed source of sexual orientation available because it reports sexual orientation based on both behavior and identity. The main drawback of the General Social Survey is the sample sizes are very small. Between 1990 and 2014, there were 368 gay or bisexual men identified in the data. The Decennial Census and the Current Population Survey provide larger sample sizes for gay men, but these sources do not ask about sexual orientation directly. In these sources, same-sex couples are identified if the householder indicates they are cohabiting with an unmarried partner of the same sex. Same-sex couples who indicate that they are married are re-coded as unmarried partners.[5] Data from the General Social Survey suggests that about 30% of gay men are cohabiting, meaning the Census definition will not capture 70% of the gay population.

Across all previous studies, the gay wage penalty averaged 11%, but individual estimates range from 0% to as high as 30% (Klawitter 2015).[6] In these studies, researchers control for education, age, race, hours worked, occupation and industry, and some control for the presence of children (Klawitter 2015). The estimated wage penalty has declined since 1990 (Cushing-Daniels and Yeung 2009, Elmslie and Tebaldi 2014, Klawitter 2015). Wage penalties in the General Social Survey have declined from 30% in the early 1990s to 11% in the mid-2000s (Badgett 1995, Cushing-Daniels and Yeung 2009). The estimated penalties in the Census data have fallen from 15% in the 1990 Census to approximately 6% in the American Community Survey. Wage penalties in the Current Population Survey have fallen from 8%

[5] One-third of gay respondents in the General Social Survey are cohabiting. Heterosexual men are more likely to cohabitate. Approximately 60% of heterosexual men in the General Social Survey are cohabiting. The large difference in cohabitation rates explains the difference in the gay share between sources that identify only cohabiting gay men and sources that identify all gay men.

[6] Note that the estimates of no wage penalty come from sources either using data from outside the United States (e.g. Frank (2006)) or from a single state (e.g. Carpenter (2005)).

in 1995 to 4% in 2011 (Elmslie and Tebaldi 2014).

Non-discriminatory explanations of the wage penalty are not consistent with the data. Many researchers, therefore, attribute the unexplained wage penalty to taste-based discrimination (Badgett 1995, Antecol, Jong and Steinberger 2008, Martell 2013a). Differences in preferences for time spent on leisure could explain the wage penalty, but this theory is not consistent with evidence that estimates of the wage penalty increase in magnitude after controlling for selection into the labor market (Berg and Lien 2002, Elmslie and Tebaldi 2014, Klawitter 2015). Research has also shown that occupational segregation is not the cause of the wage penalty (Antecol et al. 2008, Klawitter 2015). If gay men had the same distribution across occupations as heterosexual men, there would still be a significant wage penalty (Antecol et al. 2008). Unobserved differences between the personality of gay men and heterosexual men do not appear to drive the wage penalty. Controlling for personality characteristics (e.g. extroversion, neuroticism, agreeableness, conscientiousness, open to experience/intellect/imagination, optimism, and anger- hostility), physical appearance (e.g. BMI and self-rated attractiveness), or mental health (e.g. self-reported levels of stress, ADHD, and anxiety) also does not result in large changes in the estimated wage penalty (Sabia 2014).

Discriminatory explanations of the wage penalty are bolstered by the experimental evidence found in correspondence studies. Research has found evidence that gay men and lesbian women are discriminated against in hiring (Drydakis 2015, Mishel 2016, Tilcsik 2011). Tilcsik (2011) sent pairs of fictitious resumes, one gay and one heterosexual, to employers and measured the callback rates. The author found significantly lower callback rates for gay resumes than heterosexual resumes. The discrimination was concentrated in job ads where employers emphasized the importance of stereotypically male traits. Drydakis (2015) found that in addition to the lower callback rates, employers who called back gay resumes offered lower wages than employers who called back heterosexual resumes. Previous research has also shown that when states make it illegal for companies to discriminate based on sexual orientation, the gay wage penalty shrinks (Baumle and Poston Jr. 2011, Klawitter 2011, Martell 2013b).

1.4 Models of Taste-Based Discrimination

Using data on prejudice and wages, I can test whether there is a relationship between prejudice and wage penalties. There are two models in economics often used to explain why there might be an unexplained wage penalty for gay men observed in the data. First is the neoclassical model described by Becker (1971). The other is the search model of discrimination described by Black (1995).

1.4.1 Comparing the Predictions of the Models

I can differentiate between the Becker model and the search model by comparing how different parts of the prejudice distribution are correlated with the wages of gay men. Table 1.2 compares the predictions of each model. Both models predict a negative correlation between the wages of gay men and prejudice. In the Becker model, the smaller size of the minority group means the marginal employer comes from the low end of the prejudice distribution. However, search models predict the upper tail of the distribution, where employers would not hire a gay man, drives the wage penalty. In the search model, unprejudiced employers are reacting to the number of highly prejudiced employers when they determine the wages to pay gay men.

The models differ in the predicted relationship between the wage penalties and the size of the gay population. Without search frictions, gay workers sort towards the least prejudiced employers. Increasing the number of gay men in the labor market forces gay men to accept jobs with more prejudiced employers. As the disutility of the marginal employer increases, the wages of gay men fall. In the search model, increases in the number of gay men increase the reservation utility of gay men. This forces unprejudiced employers to pay gay workers more, increasing the wages of gay men. Therefore, if there is a negative correlation between the size of the gay population and the wages of gay men, this is evidence the Becker model better explains the wage penalty. If, however, there is a positive correlation between the size of the gay population and the wages of gay men, then this is evidence for the search model.

Table 1.2: Comparing Predictions of Taste-Based Discrimination Models

	Becker 1971	**Black 1995**
Prejudice and Wages of Gay Men		
Marginal Employer	Negative Correlation	
Average Employer	No Correlation	
Share Prejudiced		Negative Correlation
Share Unprejudiced	No Wage Penalty if Share Unprej>Minority Pop	
Population and Wages of Gay Men		
Minority Share	Negative Correlation	Positive Correlation

Note: Testable predictions are drawn from Charles and Guryan (2008) for Becker's model and from Black (1995) for the search model.

1.4.2 The Becker Model of Discrimination

In the Becker model of discrimination, firms operate in a perfectly competitive environment. There are two sets of agents in the model: employers and workers. Workers can either be gay or heterosexual. Employers know the sexual orientation of workers.

An employer's utility (u_e) depends on both their profits and the number of gay workers they hire.[7] Their utility depends positively on profits (π) and negatively on the number of

[7]Neumark (1988) shows the Becker model can be extended to the case where employers do not care about the absolute number of minorities they hire but care only about the relative share of minorities. When the disutility is a result of the relative share ($\frac{L_g}{L_g+L_h}$), there is no longer perfect segregation. The relationship

gay workers (L_g). There is a distribution of prejudice, with d_e representing the employer's specific level of disutility. Employers who are more prejudiced experience more disutility when hiring a gay worker.

$$u_e = \pi - d_e L_g \tag{1.1}$$

The profit function can be expressed as

$$\pi = f(L_h + L_g) - w_h L_h - w_g L_g \tag{1.2}$$

where w_h and w_g are the wages of heterosexual workers and gay workers, and f is a production function with constant returns to scale. Employers choose the number of heterosexual workers, L_h, and the number of gay workers, L_g, that maximizes Equation 1.1. These choices, L_h^* and L_g^*, satisfy the following first-order conditions:

$$\begin{aligned} f'(L_h^* + L_g^*) - w_h = 0 \quad if \quad L_h^* > 0 \\ f'(L_h^* + L_g^*) - w_g - d_e = 0 \quad if \quad L_g^* > 0 \end{aligned} \tag{1.3}$$

The first-order conditions state that each employer will hire a particular type of labor until the point where its marginal product is equal to its marginal cost. For heterosexual workers, the marginal cost is simply the wage, w_h. For gay workers, the marginal cost is the wage, w_g, and the disutility from hiring a gay worker, d_e.

By assumption, the marginal productivity of the two groups is identical.[8] The result is

between the prejudice of the marginal employer and the wage penalty is no longer as simple as in the Becker model. The lack of search frictions still results in gay workers sorting towards the least prejudiced employers first. So the prejudice in the lower tail of the prejudice distribution should matter more for the wage penalty than prejudice in the upper tail of the distribution. The relationship between wage penalties and the size of the gay population should still be negative. When there are no search frictions, minority workers will always take jobs from the least prejudiced employers first. So increasing the number of minority workers should increase the wage penalty as they are forced to take jobs from increasingly prejudiced employers.

[8] Any differences in productivity should be captured by the controls in the regression, so the wage penalties estimated in this paper are net of marginal productivity.

perfect segregation by worker type. Employers hire only heterosexual workers if $w_h < w_g + d_e$. Employers hire only gay workers if $w_h \geq w_g + d_e$. Gay workers will sort towards the least prejudiced employers, while heterosexual workers sort towards the more prejudiced employers.

In equilibrium, the market clears at wages w_g^* and w_h^*. If the distribution of prejudice is smooth enough, there will be an employer who is perfectly indifferent between hiring a heterosexual worker and a gay worker. The prejudice of the marginal employer, d_e^*, is equal to the gay wage penalty in equilibrium because their indifference between hiring heterosexual workers and gay workers implies that

$$w_h^* = w_g^* + d_e^* \tag{1.4}$$

Any employer with prejudice greater than d_e^* will hire only heterosexual workers and employers with prejudice less than d_e^* will hire only gay workers.

Charles and Guryan (2008) show the Becker model has two testable predictions. First, the more prejudiced the marginal employer is, the larger the wage penalties will be. The wage penalty for gay men increases as the prejudice of the marginal employer increases. Second, as the gay population grows, the marginal employer will change. Each additional gay worker will shift who is the marginal employer higher in the distribution of prejudice, where the employers have higher values of d_e. Therefore, the wage penalty will be larger when a larger share of the population is gay.

1.4.3 Search Model of Discrimination

The predictions of the Becker model depend on a lack of search frictions. The fact that searching for a job is costless allows gay workers to find jobs from the least prejudiced employers. If search frictions make searching for a job costly, then it may not be utility maximizing for gay men to continue searching until they find a less prejudiced employer. When only 16% of employers are unprejudiced towards gay men, gay men may be willing to accept lower wages to avoid repeated periods of search when they are matched with one of the 84% of employers who are prejudiced against them. The result is that the wage penalty in a search model can be influenced by infra-marginal employers and employers who would never hire a gay worker.

In this paper, I consider the model of discrimination using sequential search described in Black (1995). There are two types of employers in the search model: prejudiced and unprejudiced. For simplicity, assume that prejudice is binary. Unprejudiced employers have a disutility of hiring a gay man equal to zero. They are willing to hire both gay and heterosexual workers. Prejudiced employers have a level of disutility large enough that for any wage greater than zero they will not hire gay men.[9] Prejudiced employers only hire heterosexual workers and are θ percent of all employers.

1.4.3.1 Worker Behavior

First, consider the behavior of the workers. When a worker becomes unemployed, they search sequentially for a job from a set of potential employers looking to hire. Workers are either gay (g) or heterosexual (h). The type of worker is denoted with the subscript i. Employers are either unprejudiced (u) or prejudiced (p). The type of employer is denoted with the superscript j. Searching for a job costs a worker κ each period. After a worker is matched with an employer, they are offered a wage, w_i^j, and learn about their satisfaction for the

[9] This is the same environment as described in Black (1995). Black (1995) argues that the results hold even if one allows for prejudice to be continuous (as is the case in the Becker model).

match, α. The worker's utility for the job is given by:

$$u_i = w_i^j + \alpha \quad for \quad i = \{h, g\} \quad and \quad j = \{u, p\} \tag{1.5}$$

The distribution of α is given as $F(\alpha)$ and the density function as $f(\alpha)$. The distribution of α is strictly log-concave, implying that the inverse hazard function is strictly decreasing. This assumption ensures that employers are monopsonistic competitors and that they are facing an upward-sloping labor supply function (Black 1995).

Each employer chooses a wage offer that maximizes their utility. Given that heterosexual workers do not care if they work for a prejudiced or unprejudiced employer, the labor supply of heterosexual workers is independent of employer type. Assuming a constant returns to scale production function results in unprejudiced and prejudiced employers offering heterosexual workers the same wage ($w_h^p = w_h^u = w_h$).

We can write the expected value of search for heterosexual workers as

$$V_h = \theta Emax\{w_h + \alpha, V_h\} + (1 - \theta)Emax\{w_h + \alpha, V_h\} - \kappa \tag{1.6}$$

and the expected value of search for gay workers as

$$V_g = (1 - \theta)Emax\{w_g^u + \alpha, V_g\} - \kappa. \tag{1.7}$$

The expected value of search for gay workers depends only on unprejudiced employers, so the probability that they receive an offer is $1 - \theta$. Equations 1.6 and 1.7 can be rearranged using the distribution of α to obtain

$$\kappa = \int_{V_h - w_h}^{\infty} (w_h + \alpha - V_h) f(\alpha) d\alpha \tag{1.8}$$

for heterosexual workers and

$$\frac{\kappa}{1-\theta} = \int_{V_g - w_g}^{\infty} (w_g + \alpha - V_g) f(\alpha) d\alpha \tag{1.9}$$

for gay workers. The left-hand sides of these equations represent the expected cost of searching another period for a job, and the right-hand sides represent the expected benefit of waiting an extra period for a new job offer. A worker searches for a job until their expected cost of search is equal to their reservation utility (V). Because prejudiced employers never hire gay workers, the relative cost of a job search is higher for gay workers. The comparative statics for gay workers yield

$$\begin{aligned} &\frac{dV_g}{dw_g} = 1 \\ &\frac{dV_g}{d\theta} < 0. \end{aligned} \tag{1.10}$$

Note that when the wages for gay men increase the reservation utility of gay men increases at the same rate. An increase in the number of prejudiced employers decreases the reservation utility of gay workers because the increase in prejudiced employers increases the expected cost of searching for gay workers. If a gay worker does not accept the current offer, they have to wait more periods on average for the next job offer to arrive.

1.4.3.2 Employer Behavior

Unprejudiced employers receive no disutility from hiring gay men. Their utility is only determined by their profit, so their goal is to maximize their per applicant profit (π^u). If a worker accepts a job offer, the employer earns a profit of $MPL - w_i^u$, where MPL is the marginal product of the worker. The employer receives no profit if the worker rejects their wage offer and they fail to hire someone that period.

$$\pi_i^u = [1 - F(V_i - w_i^u)](MPL - w_i^u) \tag{1.11}$$

The necessary condition for profit maximization is

$$MPL - w_i^u - m(V_i - w_i^u) = 0 \tag{1.12}$$

where $m(V_i - w_i^u)$ is the inverse hazard function.

The comparative statics show that unprejudiced employers set wages in response to the reservation utility of workers. This leads to two testable predictions. First, the wages of gay men are negatively correlated with the share of prejudiced employers (θ). Second, the wages of gay men are positively correlated with the share of the population that is gay (γ).

$$\begin{aligned} 0 &< \frac{dw_g}{dV_g} < 1 \\ 0 &< \frac{dw_h}{dV_h} < 1 \end{aligned} \tag{1.13}$$

Using the comparative statics, it is easy to show that $\frac{dw_g}{d\theta} < 0$. When the number of prejudiced employers in the market increases, the reservation utility of gay workers decreases ($\frac{dV_g}{d\theta} < 0$). The decrease in the reservation of utility lowers the wages unprejudiced employers pay gay workers ($0 < \frac{dw_g}{dV_g} < 1$). This decline in the wages paid to gay workers increases the wage penalty.

$$\frac{dw_g}{d\theta} = \frac{dw_g}{dV_g}\frac{V_g}{d\theta} < 0 \tag{1.14}$$

Using the comparative statics in Equation 16, it is easy to show that $\frac{dw_g}{d\gamma} > 0$.[10] To obtain this result, Black (1995) assumes that employers have a distribution of entrepreneurial ability

[10] See Appendix of Black (1995) for the formal derivation of this result.

that results in a distribution of fixed costs. When the proportion of gay men increases, prejudiced employers are matched with gay workers at a higher rate. Since each employer is only matched with one worker each period, this results in more periods of zero profit for prejudiced employers, which reduces their expected profit. For employers with sufficiently high fixed costs, the decrease in expected profit due to fewer heterosexual workers drives them from the market. The result is fewer prejudiced employers as the proportion of gay men increases ($\frac{d\theta}{d\gamma} < 0$). As the share of prejudiced employers falls, the reservation utilities of gay workers increase ($\frac{dV_g}{d\theta} < 0$). The increase in reservation utilities forces unprejudiced employers to pay higher wages ($0 < \frac{dw_g}{dV_g} < 1$). The net effect is that increases in the gay population increase the wages for gay men.

$$\frac{dw_g}{d\gamma} = \frac{dw_g}{dV_g}\frac{dV_g}{d\theta}\frac{d\theta}{d\gamma} > 0 \tag{1.15}$$

1.5 Data and Methodology

Testing the relationships predicted by the theory requires data on wages for gay men and heterosexual men and data on prejudice against gay men in the United States. The data on social attitudes come from the General Social Survey. I obtain restricted access data that allows me to match each respondent to their state of residence because the General Social Survey only reports the Census division of residence in the publicly available data.[11] The data on wages of gay men and heterosexual men come from the Census Bureau. The Census began collecting data on cohabiting homosexuals in 1990. To match the availability of data on homosexuals in the two data sources, I pool state-level data from the 1990 to 2014 General Social Surveys and merge it with the wage data from the 1990 Census, the 2000 Census, and the 2008 through 2014 American Community Surveys.

[11]The restricted nature of the state of residence means that I cannot show data from the General Social Survey at the state level.

1.5.1 Data on Prejudice

The General Social Survey combines data across individuals and states. In the pooled sample of General Social Surveys between 1990 and 2014, there are 34,706 respondents. In Table 1.3, I present the demographics of General Social Survey respondents. Respondents in the General Social Survey are 45.57 years old on average. They have obtained 13.33 years of schooling, with 26% obtaining a degree higher than a high school diploma. One thing to note is that there are more women in the sample than men. This may lead the General Social Survey to underestimate the prejudice against homosexuals since female respondents in the General Social Survey are less prejudiced against gay men than male respondents. This over-representation of women in the sample is consistent across Census divisions, so the relative differences in prejudice across states will not be biased, even if the average level is too low.

Table 1.3: Demographics of GSS Respondents: 1990 to 2014

White	78%
Black	14%
Other	8%
Years of Schooling	13.33
High School Diploma	52%
Bachelor's Degree	17%
Graduate Degree	9%
Age	45.57
Male	44%
Female	56%
New England	5%
Middle Atlantic	14%
E. N. Central	17%
W. N. Central	7%
South Atlantic	20%
E. S. Central	7%
W. S. Central	10%
Mountain	7%
Pacific	14%
N	34,706

Note: Data on GSS respondents come from the pooled General Social Survey, 1990 to 2014.

The goal of using the responses from the General Social Survey is to create a distribution of prejudice. There is not enough data in the General Social Survey to calculate a distribution of prejudice for each state in each year. Because the Becker model requires prejudice to be sufficiently smooth for there to be a marginal employer, I do not have enough data to calculate the marginal employer for each state-year cell. Therefore, I follow the methodology used in Charles and Guryan (2008) and calculate the distribution of prejudice for each state.

I begin by aggregating each respondent's responses into a single index. If there are differences in the number of questions used to construct the index, then the tails of the distribution will be drawn disproportionately from observations with fewer questions due to the higher variance of the indexes. Therefore, I limit the data to four questions asked in every survey (i.e. SEX, BOOK, SPEAK, and COLLEGE). This excludes MARRIAGE and CHILD.

When constructing the distribution of prejudice, I use the same formula as Charles and Guryan (2008), which treats all questions as having the same weight in the index. Because the questions are coded on different scales, I normalize the responses for question k to have a mean of zero and a standard deviation of one in 1990. These normalized responses can be written as

$$\tilde{d}_{i,t,k} = \frac{d_{i,t,k} - E[d_{1990,k}]}{\sqrt{Var\ d_{1990,k}}} \tag{1.16}$$

Individual responses are aggregated to create a single index for individual i. These indexes are weighted by the total number of questions a respondent answered, in this case, four.

$$D_{i,t} = \frac{\sum_K \tilde{d}_{i,t,k}}{4} \tag{1.17}$$

These individual responses are regressed on a full set of year dummies to capture the average prejudice ($\tilde{D}_{i,t}$). The prejudice is aggregated to the state level to create a distribution of prejudice for each state. From these distributions, I determine the average prejudice, percentiles of the prejudice distribution, the share of respondents who gave prejudiced responses to all questions, and the prejudice of the marginal individual. The marginal individual can be approximated for by using the percentile of the prejudice distribution equal to the share of the gay population in the sample (Charles and Guryan 2008).

1.5.2 Data on Wages

The data used to determine the gay wage penalty come from the 2008 through 2014 American Community Survey (ACS) 1-Year Samples, the 1990 U.S. Census 5% Sample, and the 2000 U.S. Census 5% Sample (Ruggles, Alexander, Gendadek, Goeken, Schroeder and Sobek 2010). I restrict the sample to adults over the age of 18 and younger than 65. To identify gay men in the United States, the Census collects information on householders and the relationships of everyone in the household to the householder. A same-sex couple is identified when the gender of the householder and the gender of the unmarried partner (or spouse) of the householder are the same. There is no information on single gay men in the Census data or ACS, only cohabiting gay men. Also missing are gay men in a household where one of the partners is not the household head (such as living with one's parents). Therefore, the sample is restricted to comparisons between cohabiting individuals.

By combining the Census data with the American Community Survey data, I construct a sample of 72,239 cohabiting gay men and 10,635,623 cohabiting heterosexual men (Ruggles et al. 2010). Table 1.4 highlights the differences between cohabiting heterosexual men and gay men in the Census data. Gay men are younger than heterosexual men but have obtained more years of schooling. They are also less likely to have children (10% of gay men have children, while 54% of cohabiting heterosexuals do). Regarding where they live, gay men are more likely to be living in an urban area (56% vs. 48%). Gay men are also more likely to live on the coasts, most notably the Pacific coast, where 22% of gay men live compared to 16% of heterosexuals.

I restrict the sample used in the analysis to consist only of gay men and married cohabiting men. I further impose the sample restriction that the men must be in the labor force and earning a wage. Observations are dropped from three states because the states do not have enough observations in the General Social Survey to calculate the marginal employer. Each of the states dropped has fewer than 50 observations in the General Socials when the 1990 through 2014 General Social Surveys are combined. After imposing these restrictions, I am left with a total sample size of 6,268,265 individuals.

Table 1.4: Demographics of ACS/Census Respondents: 1990 to 2014

	Heterosexual Men	Gay Men
Annual Income (1999 dollars)	$42,245	$41,885
White	80%	83%
Black	10%	7%
Other	10%	10%
Years of Schooling	14.5	15.4
Age	44.3	42.7
Urban	48%	56%
Kids	54%	10%
New England	5%	6%
Middle Atlantic	13%	15%
E. N. Central	16%	13%
W. N. Central	7%	5%
South Atlantic	19%	21%
E. S. Central	6%	4%
W. S. Central	12%	9%
Mountain	7%	7%
Pacific	16%	22%
N	10,635,623	72,239

Note: Data on wages come from the 1990 Decennial Census 5% PUMS, the 2000 Decennial Census 5% PUMS, and the 2008 through 2014 American Community Surveys. All respondents are cohabiting or married.

1.5.3 Testing the Relationship Between Prejudice and Wages

As noted in Charles and Guryan (2008), there are two ways to estimate the relationship between prejudice and wages. The first method they propose is to aggregate the prejudice data to the state level and use it as a control in the individual-level wage equation. The base regression would estimate an OLS regression of log wages on education, potential experience, marginalized group-specific year dummies, a dummy for the marginalized group, and the interaction between the prejudice measure and the dummy for the marginalized group, with state and year fixed effects. The coefficient of interest in the regression is the estimated effect on the interaction of the marginalized group dummy with the prejudice measure. A negative coefficient would indicate the marginalized group experiences a larger wage penalty in states with more prejudiced populations. Charles and Guryan (2008) argue against this methodology since it could result in standard errors that are too small because they do not take into account that prejudice is calculated at the state level.

Instead, Charles and Guryan (2008) propose using a two-step process to estimate the relationship between prejudice and wages. First, they estimate the residual wage gap in each state by interacting a dummy for being black with a state indicator. The estimated effects on each of these interactions become the dependent variable in the second step. In this second regression, the measure of prejudice is then used as the independent variable. A negative coefficient on the prejudice measure would indicate that states with more prejudiced populations have larger wage gaps.

In this paper, I use the first methodology to estimate the baseline results.[12] Using the estimation sample described in Section 1.5.2, I regress log hourly wages on a dummy for being gay (G), the interaction between the prejudice measure (P) and the gay dummy, and the interaction between the gay share of the cohabiting men and the gay dummy. I use the same controls as Charles and Guryan (2008). I control for the quadratic of potential experience (Exp and Exp^2), for years of schooling (S), and for race (either being black

[12] Charles and Guryan (2008) argue that the two-step method produces standard errors that are more conservative than those produced by the first method. I have estimated the results using their method and the standard errors produced are smaller in the two-step procedure than when estimated using the individual level data and clustering the standard errors at the state level, as shown in Table 1.5. The results of this exercise are available upon request.

($Black$) or of other race ($Other$). I include state fixed effects (I_s), year fixed effects (I_t), and state-by-year fixed effects ($I_s \times I_t$). The standard errors (ϵ) are clustered at the state level to take into account that the prejudice and gay population are calculated at the state level.

$$\begin{aligned} LnY_{i,s,t} = \alpha_0 + \delta_1 G_{i,s,t} + \delta_2 (G_{i,s,t} \times P_s) + \delta_3 (G_{i,s,t} \times Share_s) \\ + \beta_1 Schooling_{i,s,t} + \beta_2 Exp_{i,s,t} + \beta_3 Exp^2_{i,s,t} + \beta_4 Black_{i,s,t} + \beta_5 Other_{i,s,t} \\ + \theta_s I_s + \theta_t I_t + \theta_{s,t} (I_s \times I_t) + \epsilon_{i,s,t} \end{aligned} \tag{1.18}$$

For the measure of prejudice (P), I use the measures of prejudice that each model predicts to be significant. I include the prejudice of the marginal employer ($Marginal_s$) or the share of individuals in a state that gave prejudiced answers to all the questions in the General Social Survey ($Prejudiced_s$). I also test the effect of the share of cohabiting men in each state in the sample that are gay ($Share_s$).

Note here that this wage penalty is calculated under the assumption that sexual orientation is binary and known by the employer. This is a strong assumption, but not unreasonable given that 65% of gay men are out to more than half of their coworkers in the General Social Survey. An additional 20% of gay men were out at work but to less than half of their coworkers.

I test the robustness of my estimation strategy in Section 1.6.2. First, I test how robust the results are to changes in how the prejudice measure are calculated. Specifically, I focus on how the questions are weighted in the index and using the prejudice of managers. Second, I test how robust the results are after controlling for time-varying factors that impact gay men and heterosexual men differently. The coefficients of interest (δ_1, δ_2, and δ_3) assume that the time-varying factors in each state influenced the wages of gay men and heterosexual men equally. I test this assumption by controlling for changes in employment non-discrimination acts and legal recognition of same-sex marriages in a state.

1.6 Results

Similar to the previous literature, I find there is a significant wage penalty for gay men in the United States. Between 1990 and 2014, the average wage penalty for gay men was 10.4%. Figure 1.2 shows the wage penalty by states. There is a large variation in the wage penalty, with the penalty being largest in Wyoming (-27.4%) and smallest in the District of Columbia (+2.5%). The wage penalty is largest in the Great Plains states and the Midwest.

Figure 1.2: Gay Wage Penalty By State

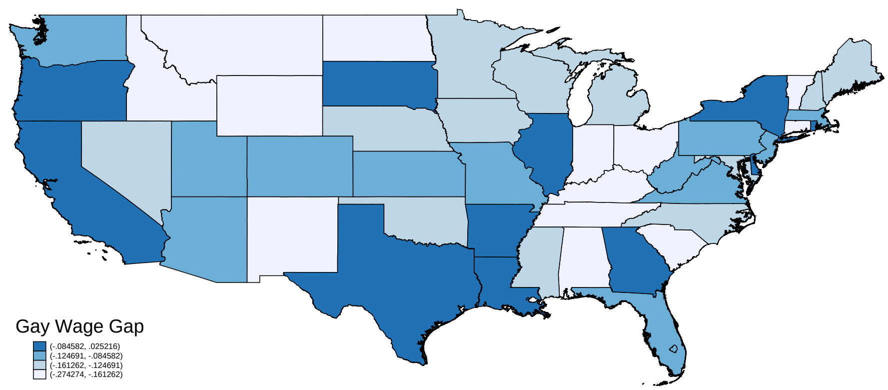

Note: Data on wages come from the 1990 Decennial Census 5% PUMS, the 2000 Decennial Census 5% PUMS, and the 2008 through 2014 American Community Surveys. Wage penalties are the time invariant wage penalties found by modifying Equation 1.18 to identify the average wage penalty in a state using an interaction term between a dummy for being gay and a state indicator. No controls for prejudice or the gay population are included.

1.6.1 Baseline Results

I begin by testing the relationship between prejudice and the wage penalty in Table 1.5 using Equation 1.18.[13] In columns 1 and 2, I test each model independently of the other. By ignoring the prejudice measure from the other model, I am potentially biasing the results with an omitted variable. To account for this, I control for both the prejudice of the marginal employer and the share of individuals in a state who are prejudiced against gay men in column 3.

Table 1.5: Testing the Predictions of Models of Taste-Based Discrimination

	(1) Becker	(2) Search	(3) Both
Marginal Prejudice	-0.197		-0.245
	(0.167)		(0.173)
Share Prejudiced		-0.438**	-0.451**
		(0.178)	(0.192)
Share Gay	3.138**	2.358*	2.347**
	(1.491)	(1.192)	(1.167)
adj. R^2	0.21	0.21	0.21
States	48	48	48
Obs	6,268,265	6,268,265	6,268,265

*** p<0.01, ** p<0.05, * p<0.1

Note: The dependent variable is log hourly wages in constant 1999 dollars. The sample has been restricted to only include gay or married cohabiting men in the labor force. On average gay men experience a wage penalty of 10.4% relative to married heterosexual men. Data on wages come from the 1990 Decennial Census 5% PUMS, the 2000 Decennial Census 5% PUMS, and the 2008 through 2014 American Community Surveys. Data on prejudice come from the 1990 through 2014 waves of the GSS. Standard errors are clustered at the state level and are reported in parentheses. Census sample weights are used to weight the observations. Three states have been dropped from the sample because they have too few respondents in the General Social Survey.

[13] Charles and Guryan (2008) argue that testing the relationship between various prejudice percentiles and the wage penalty was a sharper test of the Becker model. I conduct a similar exercise for the search model, which predicts prejudice in the right-hand tail of the distribution matters more. Table A.2 shows that the only significant correlations between percentiles of the prejudice distribution and the penalty for gay men are found above the marginal employers.

Table 1.5 shows the coefficient on the prejudice of the marginal employers is always negative, but never statistically significant. This result suggests that the lower tail of the prejudice distribution does not play a large role in determining the gay wage penalty. The fact that there is no significant relationship between the marginal employer and the wage penalty is not unexpected given the gay share is less than the unprejudiced share in the population. If prejudice is driving the wage penalty, the prejudice responsible must come from a part of the prejudice distribution the Becker model predicts does should not matter.

I find strong evidence that the wage penalty is correlated with the share of prejudiced individuals, regardless of whether the marginal employer is prejudiced towards gay men. In each of the specifications, there is a negative relationship between the wage penalty and the share of individuals in a state who gave prejudiced answers to all the questions in the General Social Survey. I find that a 1% increase in the share of individuals who are prejudiced increases the wage penalty for gay men by 0.45 percentage points. A one standard deviation increase in the percent of the population that is prejudiced towards gay men would increase the wage penalty for gay men in that state by 2.7 percentage points.

If the Becker model was correct, there should be a negative relationship between the wages of gay men and the share of cohabiting men who are gay. Only in the search model of discrimination would we expect to find the wages of gay men positively correlated with the share of cohabiting men who are gay. This is strong evidence that search frictions are playing an important role in the gay wage penalty. Across all three specifications, I find that a 1% increase in the share of cohabiting men who are gay shrinks the wage penalty between 2 and 3 percentage points. Overall, the results show that if the wage penalty is discriminatory, the Becker model is not the model that explains the wage penalty.

1.6.2 Robustness to Alternative Prejudice Measures

So far, the search model has been shown to be consistent with the evidence for the wage penalty for cohabiting gay men. In this section, I test how robust the results are to alternative measures of prejudice.

The prejudice used in the baseline estimation could be biased because the measures were calculated using all of the General Social Survey respondents. These responses may not reflect the prejudice of those making the hiring decisions. I can refine the method used to calculate the prejudice measures by restricting the General Social Survey sample to include only respondents employed as managers. When I calculate the share of prejudiced managers in each state, I have a better measure of the prejudice of those responsible for setting the wages of gay men. The problem is that managers only make up 8.6% of General Social Survey respondents. Shrinking the sample size potentially increases the measurement error because there is likely to be some errors in who claims to be a manager. The measurement error may result in attenuation bias, which would bias the results towards not finding a relationship between prejudice and the wage penalties.

To correct for the measurement error, I use an instrumental variable approach. The share of managers who are prejudiced will be correlated with the share of non-managers who are prejudiced. States with more prejudiced managers are likely to have more prejudiced non-managers. For the share of prejudiced non-managers to be a valid instrument, the measurement error in who claims to be a manager must be uncorrelated with the share of non-managers who are prejudiced which is likely since the propensity to misstate one's occupation is unlikely to be correlated with the prejudice of other individuals in the state. To test if the measurement error in the share of managers who are prejudiced is uncorrelated with the "signal" from the share of non-managers who are prejudiced, I regress the difference between the share of prejudiced managers and the share of prejudiced non-managers on the share of prejudiced non-mangers. I find that there is no significant correlation between them. The coefficient on the difference is -.129 and the t-statistic is -0.91.

The first column of Table 1.6 reports the results when I estimate the prejudice of managers using OLS. The magnitudes are smaller and less significant when estimated using OLS than the results found using the IV approach in column 2. A 1% increase in the share of managers who are prejudiced is correlated with an increase in the wage penalty of 0.17 percentage points. This relationship is significant only at the 10% level. The effect of the gay share of cohabiting men becomes much more significant, with the magnitude of the coefficient tripling in size. The second column of Table 1.6 reports the results when I instrument for the share of prejudiced managers using the share of prejudiced non-managers in a state. Here, I find that the effect of prejudice on the wage penalty has increased. A 1% increase in the number of prejudiced managers increases the wage penalty by 0.52 percentage points. The effect of the gay share returns to the sizes shown in the baseline estimation.

Table 1.6: Robustness of Results Using Prejudice of Managers

		First Stage
Share of Prejudiced Non-Managers		0.782***
		(0.117)
F-stat		43.59
	OLS	2SLS
Share of Prejudiced Managers	-0.173*	-0.517**
	(0.102)	(0.229)
Share Gay	6.820**	2.443**
	(3.299)	(1.062)

*** p<0.01, ** p<0.05, * p<0.1

Note: The dependent variable is log hourly wages in constant 1999 dollars. The sample has been restricted to only include gay or married cohabiting men in the labor force. On average gay men experience a wage penalty of 10.4% relative to married heterosexual men. Data on wages come from the 1990 Decennial Census 5% PUMS, the 2000 Decennial Census 5% PUMS, and the 2008 through 2014 American Community Surveys. Data on prejudice come from the 1990 through 2014 waves of the GSS. Standard errors are clustered at the state level and reported in parentheses. Census sample weights are used to weight the observations. Three states have been dropped from the sample because they have too few respondents in the General Social Survey. In the 2SLS, the share of prejudiced managers in a state has been instrumented for using the share of prejudiced non-managers in a state as the instrument.

Even if the full sample of individuals is a good sample to use, treating all four of the prejudice questions as having equal weight when calculating the prejudice measure is a strong assumption. There was a noticeable pattern in Figure 1.1 that questions attracted different levels of prejudice. Sexual relations between two adults attracted more opposition than homosexuals speaking in public. The fact that three of the questions touched on free speech for homosexuals may be biasing the results if this does not truly reflect prejudice, but instead, reflect concerns about free speech.

To account for this, I use a principal component analysis to determine the weights of each question when calculating the prejudice indexes. The principal component analysis identifies the eigenvalues of the matrix of prejudice questions. The first two components show a clear division between sexual relations and the free speech questions. Figure 1.3 shows how the questions vary over components 1 and 2. Component 1 can be viewed as a distaste towards gay men in public and being in favor of curtailing free speech for homosexual content. The second component is a distaste for the act of homosexuality itself. By counting SPKHOMO, LIBHOMO, and COLHOMO equally when calculating the index, the original estimation strategy gave opposition to gay men in public three times the weight as it gave distaste for sexual relations.

Figure 1.3: Principal Component Analysis Loading Plot

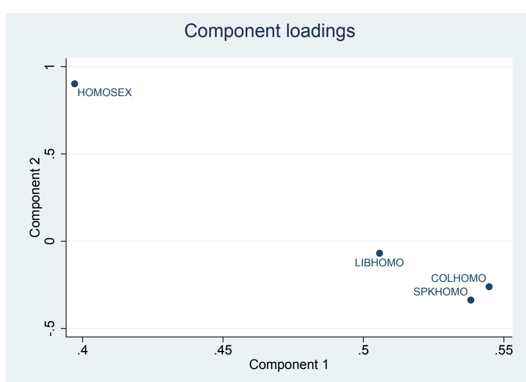

Note: Data on prejudice come from the pooled General Social Survey, 1990 to 2014. See Table 3.6 for the text of each question.

As a robustness check, I use the weights proposed by the principal component analysis determine the weights used to aggregate the individual's questions into the single prejudice index. For each respondent to the GSS, I predict their PCA component scores for components 1 and 2. I then average these two scores together to calculate the individual prejudice index. This index is then regressed on the year fixed effects to calculate the average prejudice. The indexes are then aggregated to the state level, the same as in the baseline analysis. Table 1.7 repeats the baseline analysis but instead uses the prejudice of the marginal individual from the distribution created using the PCA method. I find that using the weighted prejudice index reduces the effect of the prejudice of the marginal individual. This decline is most noticeable in the nested model shown in Column 3. Note that re-weighting the questions in the prejudice index did not affect the share of the population that was prejudiced, so the results in Column 2 remain the same.

Table 1.7: Testing the Effect of Re-weighting the Questions in the Prejudice Index

	(1) Becker	(2) Search	(3) Both
Marginal Prejudice	-0.088		-0.042
	(0.092)		(0.097)
Share Prejudiced		-0.438**	-0.415*
		(0.178)	(0.210)
Share Gay	3.296**	2.358*	2.479**
	(1.477)	(1.192)	(1.234)
adj. R^2	0.21	0.21	0.21
States	48	48	48
Obs	6,268,265	6,268,265	6,268,265

*** p<0.01, ** p<0.05, * p<0.1

Note: The dependent variable is log hourly wages in constant 1999 dollars. The sample has been restricted to only include gay or married cohabiting men in the labor force. On average gay men experience a wage penalty of 10.4% relative to married heterosexual men. Data on wages come from the 1990 Decennial Census 5% PUMS, the 2000 Decennial Census 5% PUMS, and the 2008 through 2014 American Community Surveys. Data on prejudice come from the 1990 through 2014 waves of the GSS. Standard errors are clustered at the state level and are reported in parentheses. Census sample weights are used to weight the observations. Three states have been dropped from the sample because they have too few respondents in the General Social Survey. The questions in the prejudice index have been reweighted based on the principal component analysis component loadings.

1.6.3 Gay Rights Movement

The baseline estimation does not account for changes in a state that impact gay men and heterosexuals differently. To test how this may bias the results, consider the changes in laws due to the gay rights movement. The passage of gay rights laws could potentially increase the wages for gay men without changing the wages for heterosexual men. If the passage of the laws is correlated with the share of prejudiced employers in a state, then the relationship between prejudice and wages in the baseline estimates will suffer from omitted variable bias. The coefficients would contain both the effect of the laws and the effect of the prejudice.

Table 1.8: Gay Rights Laws

State	ENDA	Same-Sex Marriage
California	1992	
Colorado	2007	
Connecticut	1991	2008
Delaware	2009	2013
District of Columbia	1977	2010
Hawaii	1991	
Illinois	2006	
Iowa	2007	
Maine	2005	
Maryland	2001	2013
Massachusetts	1989	2004
Minnesota	1993	2013
Nevada	1999	
New Hampshire	1998	2010
New Jersey	1992	2013
New Mexico	2003	
New York	2003	2011
Oregon	2008	
Rhode Island	1995	2013
Vermont	1991	2010
Washington	2006	2013
Wisconsin	1982	

Note: See Human Rights Campaign (2013), Human Rights Campaign (2012), and Sears, Hunter and Mallory (2009) for more details on these laws.

In Table 1.9, I test the robustness of the results to the passage of employment non-discrimination laws and legal recognition of same-sex marriages. Table 1.8 details the states with each law and the years they passed it. Earlier work has found that non-discrimination laws and legal recognition of same-sex marriage can increase the wage of gay men (Baumle and Poston Jr. 2011, Burn and Jackson 2014, Klawitter 2011, Martell 2013b). To control for these laws, I modify the estimation process used earlier to control for changes in each state's laws. I include dummies for whether a state had an employment non-discrimination act and whether a state had legal recognition of same-sex marriage.

$$\begin{aligned} LnY_{i,s,t} = \alpha_0 + \delta_1 G_{i,s,t} + \delta_2(G_{i,s,t} \times P_s) + \delta_3(G_{i,s,t} \times Share_s) \\ + \mu_1(G_{i,s,t} \times ENDA_{s,t}) + \mu_2(G_{i,s,t} \times Marriage_{s,t}) \\ + \beta_1 Schooling_{i,s,t} + \beta_2 Exp_{i,s,t} + \beta_3 Exp^2_{i,s,t} + \beta_4 Black_{i,s,t} + \beta_5 Other_{i,s,t} \\ + \theta_s I_s + \theta_t I_t + \theta_{s,t}(I_s \times I_t) + \epsilon_{i,s,t} \end{aligned} \tag{1.19}$$

Table 1.9 reports the results of this analysis. The results are still the same, with only small changes in the magnitude occurring. Columns 1 and 2 test the models separately and confirm the results of the baseline estimation results in Table 1.5. Column 3 provides the strongest evidence that the changing of same-sex marriage laws is not driving differences in the average wage penalty across states. In column 3, the effect of the prejudice of the marginal employer on the wages penalty for gay men is still not significant. The coefficient on the share of the population that is prejudiced has fallen. A 1% increase in the share of the population that is prejudiced increases the wage penalty by 0.32 percentage points. The point estimates for the effect of the size of the gay population has not changed.

Table 1.9: Robustness of Results Controlling for Changes in Laws

	(1) Becker	(2) Search	(3) Both
Marginal Prejudice	-0.173		-0.239
	(0.169)		(0.171)
Share Prejudiced		-0.291*	-0.318**
		(0.147)	(0.155)
Share Gay	2.709**	2.459**	2.463**
	(1.149)	(1.077)	(1.059)
adj. R^2	0.21	0.21	0.21
States	48	48	48
Obs	6,268,265	6,268,265	6,268,265

*** p<0.01, ** p<0.05, * p<0.1

Note: The dependent variable is log hourly wages in constant 1999 dollars. The sample has been restricted to only include gay or married cohabiting men in the labor force. On average gay men experience a wage penalty of 10.4% relative to married heterosexual men. Data on wages come from the 1990 Decennial Census 5% PUMS, the 2000 Decennial Census 5% PUMS, and the 2008 through 2014 American Community Surveys. Data on prejudice come from the 1990 through 2014 waves of the GSS. Standard errors are clustered at the state level and are reported in parentheses. Census sample weights are used to weight the observations. Three states have been dropped from the sample because they have too few respondents in the General Social Survey.

1.7 Alternative Estimation Strategy

A major drawback of the General Social Survey data is that it does not allow the reliable estimation of prejudice by state and year. The state-level sample sizes are small, resulting in very noisy estimates when calculating the prejudice at the state-year level. This is why my estimation strategy and Charles and Guryan's (2008) used average state-level prejudice measures. As a robustness exercise, I use the prejudice found in the General Social Survey to predict the prejudice at the state by year level using Census data. This allows me to calculate the predicted distribution of prejudice in each state in each year and use this to derive the measures of prejudice.

To do this, I begin by estimating the relationship between prejudice and demographic characteristics. I recode the General Social Survey variables and the Census variables to be identical. This entails combining some categories into broader categories since the Census data has more granular categorizations than the General Social Survey. I regress the individual prejudice index of General Social Survey respondents on their observable characteristics. In the regression, I include interactions between demographic characteristics and state and year fixed effects. This allows the prejudice of different groups to vary across states and over time.

Equation 1.20 shows the full estimation equation. The independent variable in this regression is $P_{i,s,t}$ which is the prejudice index of individual i who lives in state s in year t. I control for observable characteristics that are common across the two data sets. I control for gender, age, race, education, marital status, occupation, and industry. Because the General Social Survey is only asked in even years, linear time trends are used rather than time fixed effects.

These linear trends allow me to interpolate the odd years in the Census data.

$$\begin{aligned}
P_{i,s,t} = \alpha &+ \beta_1 Female_{i,s,t} + \beta_2(Female_{i,s,t} \times t) + \beta_3(Female_{i,s,t} \times I_s) \\
&+\beta_4 Age_{i,s,t} + \beta_5(Age_{i,s,t} \times t) + \beta_6(Age_{i,s,t} \times I_s) \\
&+\beta_7 Black_{i,s,t} + \beta_8(Black_{i,s,t} \times t) + \beta_9(Black_{i,s,t} \times I_s) \\
&+\beta_{10} Other_{i,s,t} + \beta_{11}(Other_{i,s,t} \times t) + \beta_{12}(Other_{i,s,t} \times I_s) \\
&+\beta_{13} Schooling_{i,s,t} + \beta_{14}(Schooling_{i,s,t} \times t) + \beta_{15}(Schooling_{i,s,t} \times I_s) \\
&+\beta_{16} Married_{i,s,t} + \beta_{17}(Married_{i,s,t} \times t) + \beta_{18}(Married_{i,s,t} \times I_s) \\
&+\beta_{19} Divorced_{i,s,t} + \beta_{20}(Divorced_{i,s,t} \times t) + \beta_{21}(Divorced_{i,s,t} \times I_s) \\
&+\beta_{22} Widowed_{i,s,t} + \beta_{23}(Widowed_{i,s,t} \times t) + \beta_{24}(Widowed_{i,s,t} \times I_s) \\
&+\beta_{25} Occupation_{i,s,t} + \beta_{26}(Occupation_{i,s,t} \times t) + \beta_{27}(Occupation_{i,s,t} \times I_s) \\
&+\beta_{28} Industry_{i,s,t} + \beta_{29}(Industry_{i,s,t} \times t) + \beta_{30}(Industry_{i,s,t} \times I_s) \\
&+\gamma_s I_s + \omega t + \delta_{s,t}(I_s \times t) + \epsilon_{i,s,t}
\end{aligned} \tag{1.20}$$

Using these covariates, I can explain a significant portion of the variation in prejudice indexes across individuals. The R^2 of this regression is 0.327, suggesting that these covariates can explain 33% of the prejudice observed in the General Social Survey.[14]

The next step is to use the estimated coefficients from this regression and predict the prejudice score of individuals in the Census data. I construct the distribution of predicted prejudice in each state for each year in the Census data using the predicted prejudice.[15] I calculate the prejudice of the marginal individual the same as before. To calculate the share of prejudiced individuals in a state, I calculate the maximum prejudice index for an individual who did not give prejudiced answers to all questions in the General Social Survey.

[14]Including controls for religion or political ideology increase the R^2 but are not available in the Census data, so I do not include them.

[15]Because this data has been constructed using regression estimates and not directly from the restricted General Social Survey samples, they can be shown and distributed. The estimated coefficients from the General Social Survey data and the predicted distributions of prejudice at the state level from the Census data are available from the author.

Any individual in the Census data with a predicted prejudice score greater than that value was coded as being prejudiced. I combine these measures of prejudice with the share of cohabiting men who are gay in the Census data to re-estimate Equation 1.18 directly.

$$
\begin{aligned}
LnY_{i,s,t} = \alpha_0 + \delta_1 G_{i,s,t} + \delta_2 (G_{i,s,t} \times \hat{P}_{s,t}) + \delta_3 (G_{i,s,t} \times Share_{s,t}) \\
+\beta_1 Schooling_{i,s,t} + \beta_2 Exp_{i,s,t} + \beta_3 Exp^2_{i,s,t} + \beta_4 Black_{i,s,t} + \beta_5 Other_{i,s,t} \\
+\theta_s I_s + \theta_t I_t + \theta_{s,t}(I_s \times I_t) + \epsilon_{i,s,t}
\end{aligned}
\tag{1.21}
$$

Because the General Social Survey data is used to predict the prejudice in the Census data, the standard error of δ_2 will be biased if I use OLS. To estimate the standard error of δ_2 in Equation 1.21, I bootstrap the process of calculating the prejudiced share in a state. I randomly sample 10,000 General Social Survey respondents with replacement. From these 10,000 respondents, I estimate the coefficients in Equation 1.20. I predict the prejudice respondents in the Census data with these estimated coefficients and calculate the share of individuals who are prejudiced. I then estimate Equation 1.21 using the share prejudiced calculated from the bootstrapped sample. I repeat this 100 times and collect 100 $\hat{\delta_2}$. The bootstrapped standard error of δ_2 is then the standard deviation of these estimates.

The results of this exercise are reported in Table 1.10. Using the predicted share of the population that is prejudiced and share of cohabiting men who are gay at the state-by-year level results in similar effects to those observed in all the previous tables. There is still no evidence in favor of Becker's model of discrimination, while I find evidence in favor of the search model of discrimination.

Table 1.10: Effect of Estimating Prejudice in the Census Data Using Data from the GSS

	(1)
Share Prejudiced	-0.568***
	(0.071)
Share Gay	4.756***
	(1.498)
adj. R^2	0.21
Obs.	6,268,265

*** p<0.01, ** p<0.05, * p<0.1

Note: The dependent variable is log hourly wages in constant 1999 dollars. The sample has been restricted to only include gay or married cohabiting men in the labor force. On average gay men experience a wage penalty of 10.4% relative to married heterosexual men. Data on wages come from the 1990 Decennial Census 5% PUMS, the 2000 Decennial Census 5% PUMS, and the 2008 through 2014 American Community Surveys. Data on prejudice come from the 1990 through 2014 waves of the GSS. Standard errors are clustered at the state-level and are reported in parentheses. Census sample weights are used to weight the observations. Three states have been dropped from the sample because they have too few respondents in the General Social Survey.. The share of prejudiced individuals in a state has been estimated using the predicted prejudice. The standard error of the share of prejudiced individuals has been estimated using bootstrapping. The standard error reported is the standard deviation the mean coefficient found using 100 repetitions. See the text for more detail on the bootstrapping process.

1.8 Discussion

Using the estimates from Table 1.10, I can decompose how much of the change in the wage penalty is due to changes in prejudice. A 1 percentage point decline in the share of the population that is prejudiced reduces the wage penalty by 0.57 percentage points. The share of the population that was prejudiced fell 6.5 percentage points between 1990 and 2014 (Table 1.11). This change in prejudice would reduce the wage penalty by 3.7 percentage points. The total decline in the wage penalty over between 1990 and 2014 was 10.9 percentage points. A back of the envelope calculation suggests that declining prejudice was responsible for 34% of this (3.7%/10.9%). Using the cross-sectional results from the baseline estimation would result in a similar size effect.[16]

Table 1.11: Changes in Factors Related to Taste-Based Discrimination: 1990 to 2014

Year	Prejudiced Share	Gay Share	Wage Penalty
1990	10.5%	0.4%	19.7%
2000	6.3%	1.2%	13.9%
2008	6.4%	1.0%	8.4%
2009	5.4%	1.1%	10.1%
2010	4.5%	1.1%	8.7%
2011	4.5%	1.1%	10.3%
2012	4.6%	1.2%	8.4%
2013	4.3%	0.9%	10.5%
2014	4.0%	0.9%	8.8%

Note: Data on wages come from the 1990 Decennial Census 5% PUMS, the 2000 Decennial Census 5% PUMS, and the 2008 through 2014 American Community Surveys. Data on prejudice come from the 1990 through 2014 waves of the GSS.

[16] The measure of prejudice used here is only a proxy for the true prejudice. Measurement error may bias the results toward zero, raising the possibility that prejudice matters even more.

While Table 1.11 shows that prejudice in the upper tail of the distribution does change and this can lead to declines in the wage penalty, it also shows that this process can be slow. This suggests there is a role for active enforcement of non-discrimination laws. A better understanding of the model of discrimination underlying the wage penalty helps policymakers craft more effective non-discrimination laws. Burn (2017) showed that employment non-discrimination laws with stronger provisions for damages were more effective at reducing the wage penalty for gay men. Weaker provisions may work when the wage penalty is being driven by the lower tail of the prejudice distribution, but under the search model, weaker provisions may not have enough bite to change the behavior of employers in the upper tail of the prejudice distribution.

In addition to crafting the non-discrimination laws, the search model also provides important guidance for enforcing them. If the mechanism behind the wage penalty is similar to the Becker model, enforcement should focus on pay discrimination. However, in the search model, enforcement needs to focus on hiring discrimination. The enforcement of non-discrimination laws relies in large part on potential damages awarded to the plaintiffs, which are used to pay the plaintiffs' attorneys (Bloch 1994). The enforcement of the law is then skewed towards cases with large damages (Bloch 1994). This results in more focus on discrimination in pay and termination, due to the damages awarded for not being hired can be low (Neumark and Button 2014).

1.9 Conclusion

The results of this paper provide evidence that taste-based discrimination plays a role in the wage penalty for gay men. The evidence suggests the Becker model of taste-based discrimination does not explain the wage penalty. A search model, such as that described by Black (1995), correctly predicts the relationships between prejudice and wage penalties found in the United States. It also correctly predicts the relationship between the size of the gay population and the wage penalty. The results suggest that declining prejudice towards gay men can explain up to 34% of the decline in the gay wage penalty since 1990.

Using restricted access data from the General Social Survey, I constructed state-level distributions of prejudice towards homosexuals. I estimated the relationship between measures of prejudice and the wages of gay men observed in the United States. I found no significant relationship between the prejudice of the marginal employer in a state and the wages of gay men. There was a negative relationship between the wages of gay men and prejudice in the top half of the prejudice distribution. As the share of prejudiced individuals in a state increases, so do the observed gay wage penalties. A 1% increase in the share of prejudiced individuals is correlated with a 0.45 percentage point increase in the wage penalty gay men. This effect is even larger when I compare the wage penalties to the share of prejudiced managers in a state.

The relationship between prejudice and wage penalties is not being driven by the passage of gay rights laws. I showed the results held when I controlled for the presence of employment non-discrimination acts and the legal recognition of same-sex marriages. After controlling for the passage of these laws, there was still a negative and significant relationship between the prejudiced and the wage penalty.

The results described in this paper suggest that there is still a lot that researchers do not understand about the economics of discrimination. Future work needs to reconcile the fact that the Becker model of discrimination appears to explain the black wage penalty and a search model appears to explain the gay wage penalty. This difference suggests that search frictions may have heterogeneous effects for different minority groups. African-Americans may be better able to infer the prejudice of an employer from the race of their supervisor or by the number of other African-Americans they observe at the firm (Bond and Lehmann 2015).

While it is possible that both models are wrong and the results are spurious, a more promising research agenda should explore the role of labor market networks. Previous work has shown that labor market networks play a large role in how minority individuals find a job (Hellerstein, McInerney and Neumark 2011). This leads to the question of how an invisible minority trait impacts the utilization of labor market networks. The invisible minority trait potentially impacts the formation of labor market networks in two ways. First, it may raise the search frictions associated with finding an unprejudiced employer, making it harder for gay men to identify unprejudiced employers since they cannot use the identity of the manager or other employees to gauge the prejudice of the firm. Second, the fact that gay men do not often have parents who are also gay may inhibit the formation of gay labor market networks. Labor market networks based on sexual orientation may be generation specific since there is no need for heterosexual parents to cultivate a network to pass down to their children or for homosexual parents to pass down the network to their heterosexual children.

Chapter 2

Why Aren't Women Majoring in STEM Majors?

2.1 Introduction

Higher levels of participation in Science, Technology, Engineering, and Math (STEM) are needed to help meet the evolving needs of the U.S. economy. The Council of Economic Advisers argued the United States needs an additional one million graduates with degrees in STEM to satisfy future demand for STEM workers (Council of Economic Advisers 2012). The federal government spends $3 billion a year funding over 200 different federal programs designed to increase the number of STEM graduates (Altonji, Arcidiacono and Maurel 2016, Scott 2013). The goal of many of these programs is to increase the retention of STEM majors because 48% of students who start with a STEM major eventually switch to a different major.

Since female students switch out of STEM majors at a higher rate, a large gap between male and female participation in STEM has developed (Chen and Soldner 2013, Gemici and Wiswall 2014, Turner and Bowen 1999). Data from the 2013 American Community Survey show that women hold just 18% of Engineering and Computer Science degrees and 37% of degrees in Physical Sciences and Math (Ruggles et al. 2010). Increasing female persistence in STEM is an important factor in providing the STEM graduates needed in the future (Chen and Soldner 2013).

Avoiding STEM majors has important implications for the economic well-being of women. Because STEM majors pay more than many other majors, women often select lower-paying majors. Furthermore, these lower-paying majors have lower wage ratios between women and men (Figure 2.1). This has led policymakers to argue that increasing STEM participation would help reduce the gender wage gap (Council of Economic Advisors 2015).

Figure 2.1: Female Degrees and Ratio of Female to Male Earnings by Major

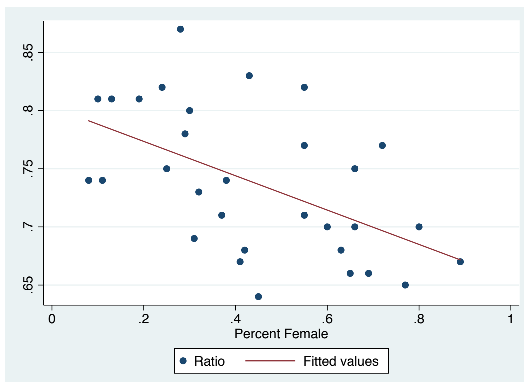

Note: Data comes from the 2013 American Community Survey (Ruggles et al. 2010). Percent female is the percent of individuals with a degree who are female. Wage ratio is the ratio of average female earnings to average male earnings for all full-time individuals with a degree in that major. The majors used here are the degree majors reported by the ACS.

In this paper, I explore the role of labor market discrimination in the decision to select a college major and what this means for female participation in STEM. Students potentially use wage ratios when deciding their major because they believe wage ratios provide information about the discrimination women face in these majors. While economists are careful to describe what is meant by discrimination and the wage ratio, media sources often conflate the two (American Association of University Women 2015, O'Brien 2015, North 2010). Therefore, majors with lower perceived wage ratios may be viewed as worse long-term investments for female students interested in pursuing a career in STEM.

Using a laboratory experiment, I collect data on how students view wage ratios in different majors, how close their expectations are to the truth, and the implications of their misinformation. If expectations of gender wage ratios influence the choice of a college major, overestimating wage ratios in STEM majors may explain the gap in female participation. In the experiment, I compare the self-reported probability of selecting a major with the expected wage ratio in that major. I test whether correcting the errors in their expectations of wage ratios increases female interest in STEM majors.

In my experiment, I find students are misinformed about the labor market. The misinformation about average earnings is worse in STEM majors than in other majors. Subjects overestimate the average annual earnings of a graduate with a bachelor's degree in Biology and Life Sciences by \$43,204. Subjects underestimate the average annual earnings of an Arts and Humanities graduate by \$87. Subjects are similarly misinformed about the gender wage ratio in the United States. They expect the gender wage ratio is highest in the Social Sciences and lowest in Business and Economics. Subjects were closest to guessing the gender wage ratio in Social Sciences, where they overestimated the wage ratio by 3.9%. Subjects were the most off when guessing the gender wage ratio in the Physical Sciences and Math, where they underestimated the ratio by 23.8%. I find the effect of correcting the misinformation was a 5.7% increase in female interest in STEM majors. The underlying cause of this change was female students underestimating the gender wage ratios in STEM majors by 0.18. By providing subjects in the treatment group with information about the actual wage ratios, female subjects increased their interest in STEM majors.

This project provides an important contribution to the economics of discrimination literature by studying how discrimination influences human capital investment along the intensive margin. Previously, Lang and Manove (2011) studied how statistical discrimination led to African-Americans obtain more years of schooling. In this paper, I compliment the previous work by asking how discrimination influences the choices that students make while in school.

As part of the experiment, I extend the experimental methodology devised by Wiswall and Zafar (2015a) and Wiswall and Zafar (2015b) to study how wage ratios impact the choice of a college major. I provide an important methodological improvement to their previous study by randomly assigning subjects to treatment and control groups. This allows me to obtain better causal estimates of the effect of information.

2.2 Origins of the STEM Participation Gap

While the existence and persistence of the gender gap in STEM careers has been well documented, there is no clear consensus about the underlying causes (American Association of University Women 2010, Ceci, Ginther, Kahn and Williams 2014, Leslie, Cimpian, Meyer and Freeland 2015, Moss-Racusin, Dovidio, Brescoll, Graham and Handelsman 2012, Reuben, Wiswall and Zafar 2017, Singh, Allen, Scheckler and Darlington 2007). The debate centers on the relative importance of society and individual factors in shaping preferences for careers in STEM. The economics literature has argued that differences in preferences for work-life balances explain why men and women choose to pursue different careers (Ceci, Williams and Banett 2009, Ceci et al. 2014, Goldin 2014). The sociology literature has argued that tastes and preferences are malleable by societal forces and women have been taught to prefer other fields to STEM (Singh et al. 2007, Gunderson, Ramirez, Levine and Beilock 2012, Nosek, Banaji and Greenwald 2002).

In this paper, it is important to consider why students select the major they do and why women may select different majors than men.[1] Potential earnings are a factor when deciding a major, but are overshadowed by tastes and preferences for a major (Arcidiacono 2004, Long, Goldhaber and Huntington-Klein 2015, Montmarquette, Cannings and Mahseredjian 2002, Reuben et al. 2017, Wiswall and Zafar 2015a, Zafar 2013). Enjoying their coursework is the

[1] In this paper, I focus on providing an overview of how college major choices vary by gender. See Altonji et al. (2016) for an in-depth review of previous studies on how students select their college major.

number one reason students select a college major (Zafar 2013). Students' enjoyment of a major may be influenced by performing well in their classes and receiving good grades. While students do not appear to sort into a major based on their actual ability for a major, there is evidence students respond to the grades they receive when they select a major (Arcidiacono 2004, Butcher, McEwan and Weerapana 2014, Ost 2010). Therefore, grade inflation in some majors may induce students to select into less rigorous majors, even if they are more adept at other majors (Butcher et al. 2014, Ost 2010, Sjoquist and Winters 2015, Stinebrickner and Stinebrickner 2014). Equalizing the grading across majors has been shown to increase female participation in STEM, where grades are curved, at the expense of the Arts and Humanities, where there are no curves (Butcher et al. 2014).

The gender gap in STEM majors converged during the 1960s and 1970s but grew larger in the 1980s and 1990s (Gemici and Wiswall 2014, Turner and Bowen 1999). The cause of this gender gap in majors appears to be differences in tastes and preferences between men and women. Turner and Bowen (1999) found that academic ability, as measured by SAT scores, could only explain between 32% and 45% of the gender gap in STEM. While research has shown that confidence in one's own academic ability influences the choice of a college major (Reuben et al. 2017), Zafar (2013) found that female under-confidence did not explain the gender gap in college majors. Zafar (2013) found that tastes and preferences for a major explained 86% of the choice of a college major for women, but only 54% of the choice for men. The most important factors for students were enjoying course work and parental approval, which explained 47% of a woman's choice of a college major and 46% of a man's. Reconciling work and family and enjoying work explained 21% of the choice for women and only 5% for men.

The formation of preferences for STEM fields begins early in a student's education (Eccles, Jacobs and Harold 1990, Gunderson et al. 2012). Throughout their time in school, a student's perception of STEM classes is shaped by the views of teachers (Gunderson et al. 2012).

Teachers that view math and science as male fields discourage female students from exploring these fields further (Gunderson et al. 2012, Steele 1997). The academic performance of female students declines when they are stereotyped by their teachers because they pay less attention in class and are less engaged (Adams, Garcia, Purdie-Vaughns and Steele 2006, Spencer, Steele, and Quinn 1999). By college, female students have developed strong implicit associations between men and math (Nosek et al. 2002). The result is that women avoid pursuing careers in fields where innate ability is seen as the driver of success (Leslie et al. 2015). Because women do not view themselves as having the same level of quantitative skills as their male peers, they pursue careers in STEM at lower rates (Leslie et al. 2015).

2.3 Methodology

The ideal experiment to test whether discrimination influences the choice of a college major would randomly assign students to experience discrimination. Based on these negative experiences, subjects would form expectations of future discrimination that were uncorrelated with unobservable characteristics. In reality, it is very difficult to study experiences of past discrimination in a way that is not confounded by unobservables.

While I cannot ask how past discrimination determines the choice of a college major, I can ask how the expectation of future discrimination influence the choice. Using an experiment similar to Wiswall and Zafar (2015a) and Wiswall and Zafar (2015b), I utilize within-subject variation to identify the causal effect of changing expectations of future discrimination on expected future outcomes.

2.3.1 Theoretical Model

It is useful to model a student's choice of a college major to understand the identification strategy behind the experiment. I construct a model similar to the one used in Wiswall and Zafar (2015a) to help illustrate how the experiment removes the endogeneity of individual preferences.

During their time in college, students pick one of K majors. Students have not selected a major at the beginning of the first period (i.e. all students enter college as undeclared majors). During period $t = 0$, students select their college major. At the end of period $t = 0$, students graduate from college. From period $t = 1$ onward, assume their choice of human capital is fixed. Individuals earn wages based on their human capital investment until they retire in period $t = T$. At the beginning of period $t = 0$, the utility for each major (V_k) is given by:

$$V_{0,k} = \gamma_k + \eta_{0,k} + EV_{1,k} \tag{2.1}$$

where γ_k represents the tastes and preferences for each major k and $\eta_{0,k}$ represents shocks to the utility of a major during period 0. $EV_{1,k}$ is the expected value of all future utility.

By picking a major, an individual is making a choice not to select another major. Therefore, a shock to one major potentially impacts the ranking of preferences for all other majors. To account for this, I construct the choice of a college major to be relative to a reference major $\tilde{k}$.

$$r_{0,k,i} = EV_{1,ki} - EV_{1,\tilde{k},i} + \psi_{k,i} \tag{2.2}$$

where $\psi_{k,i} = \gamma_{k,i} - \gamma_{\tilde{k},i} + \eta_{0,k,i} - \eta_{0,\tilde{k},i} + \epsilon_{0,k,i}$ is the combined unobservable that reflects individual specific tastes and additional sources of error.

The experiment works by exposing subjects to information and observing the changes in r. By shocking beliefs using an information experiment, I can form a panel of subjects' beliefs. For any r, let r be the relative pre-treatment beliefs and r' be the relative post-treatment beliefs. The change in the relative odds of selecting a major is then

$$r'_{0,k,i} - r_{0,k,i} = (EV'_{1,k,i} - EV'_{1,\tilde{k},i}) - (EV_{1,k,i} - EV_{1,\tilde{k},i}) + \epsilon'_{k,i} - \epsilon_{k,i}$$

Since the time between the beginning and the end of the experiment is less than an hour, $\eta_{0,k,i}$ is assumed to be 0. Because $\gamma_{k,i}$ does not vary over time, $\psi'_{k,i} - \psi_{k,i}$ simplifies to $\epsilon'_{k,i} - \epsilon_{k,i}$. Therefore, individual-specific unobservables have been removed, allowing for a causal identification of the randomly assigned treatment.[2] Any differences in $r'_{0,k,i} - r_{0,k,i}$ between the treated group and the control group will be due to the treatment shifting expectations of the future. If the information has a significant impact on $r'_{0,k,i} - r_{0,k,i}$, it means the information is relevant for deciding which major to pursue.

It is important to note the updating occurs at the individual level. Subjects may have unique processes for deciding human capital investments. For that reason, men and women may respond to information in ways that are not symmetric. Indeed, the literature on stereotyping threats suggests that female students respond to stereotyping, while male students do not (Adams et al. 2006).

2.3.2 Experimental Subjects

Subjects were randomly drawn from a pool of students who volunteered to be part of experiments at the UCI Experimental Social Sciences Lab (ESSL). A total of 66 subjects were

[2] Only one assumption is needed to obtain causal estimates. For the effect of the information to be causal, the information presented must be new to the subjects. Wiswall and Zafar (2015a) argue that the information must also be relevant to obtain a causal estimate. Since the question in this paper is whether or not this information is relevant, a causal estimate of 0 is still informative. This would indicate that subjects do not use this information when updating their beliefs.

recruited.[3] Three sessions of the experiment were conducted. Subjects were assigned to the treatment or control group with probability equal to 50%. Across all three sessions, 38 subjects were randomly assigned to the control group, and 28 were randomly assigned to the treatment group.

The experiment took place in a computer lab under the supervision of the researcher. Subjects were given 90 minutes to complete the experiment. On average, subjects took 32 minutes to complete the experiment. The fastest subject finished the experiment in 20 minutes, and the slowest subject took 56 minutes. Subjects were compensated using a combination of a show-up fee, scoring of responses, and risk elicitation.

Table 2.1 reports the descriptive statistics of the sample. The subjects were 20 years old on average. Subjects were evenly split between men and women. Asian individuals made up 70% of the sample, and only 45% of subjects grew up speaking only English at home. The average college GPA of a subject was a 3.0. On average, they had taken 4.6 AP classes in high school. STEM majors accounted for 50% of the sample, while only 5% of subjects were majoring in Arts and Humanities. Biology and Life Sciences was the largest major, followed by Social Sciences.

[3] Because subjects were not recruited based on their year of school, each subject had a different level of commitment to a major.

Table 2.1: Summary Statistics

	Control Group	Treatment Group	All Subjects
Age	20.07	20.32	20.18
Sophomore	34%	18%	28%
Junior	32%	43%	36%
Senior	34%	39%	36%
Male	55%	43%	50%
Female	45%	57%	50%
White	18%	21%	20%
Asian	71%	68%	70%
Black	3%	7%	5%
Native English Speaker	42%	50%	45%
English Second Language	58%	50%	55%
LGBT	5%	18%	11%
Heterosexual	95%	82%	89%
College GPA	2.94	3.07	3.00
AP Classes	4.42	4.82	4.59
Arts and Humanities Major	5%	4%	5%
Biology and Life Sciences Major	32%	29%	30%
Business and Economics Major	19%	21%	20%
Engineering and Computer Science	13%	14%	14%
Physical Sciences and Math	5%	7%	6%
Social Sciences	26%	25%	26%
Risk Adverse	39%	43%	41%
Risk Neutral	53%	39%	47%
Risk Loving	8%	18%	12%
Cognitive Reflection Score	1.18	1.64	1.38
Observations	38	28	66

2.3.3 Experimental Design

I utilized a survey design software to conduct the information experiment. The experiment consisted of three stages. In *Stage 1*, subjects were asked their expectations of labor market outcomes for the average worker and their own labor market outcomes. In *Stage 2*, subjects were randomly assigned to the treatment or control group. Subjects selected for the treatment learned about the true values of the labor market outcomes of the average worker. The control group received irrelevant and uninformative information. In *Stage 3*, subjects were asked to restate their expectations of their own labor market outcomes from *Stage 1*.

Stage 1 focused on subjects' expectations at age 30.[4] Subjects first answered questions about their expected education. UCI offers more than 80 undergraduate degrees. Asking subjects to rank their preferences for each of them would be unrealistic. Majors were aggregated up to six different categories to simplify the decisions. The six categories were Arts and Humanities, Biology and Life Sciences, Business and Economics, Engineering and Computer Science, Physical Sciences and Mathematics, Social Sciences.[5] A category called No Degree/Dropped Out was included to assess the likelihood a student did not graduate college. Subjects stated the probability they would earn a degree in each of the major categories. A subject's responses must add to 100. A response of 100 meant a subject had no interest in any other major and 0 meant they had no interest in that major. Subjects were then asked to rate their ability in each major relative to all other students in that major. Responses could vary from 1 to 100, with 1 being the lowest relative ability and 100 being the highest relative ability.

Subjects were asked a series of questions that was repeated for each of the major categories. The order of the questions was randomized to avoid an order effect, but the order of the

[4]Please see Figure A.1 and Figure A.2 for an example of the questions asked as it appeared to the subjects in Qualtrics.

[5] Due to the focus on STEM majors, the STEM groups are more granular than those used in in Wiswall and Zafar (2015a).

possible responses was not. Subjects were asked what they expected to earn at age 30 conditional on selecting that major and working full-time. They were also asked their expected labor supply. Conditional on majoring in that subject, subjects were asked the highest degree they expected to obtain. The degrees students were asked about were limited to Bachelor's Degree, Master's Degree, Ph.D./M.D., or Professional Degree (M.B.A., J.D., etc.).

Because an individual's job is correlated with their major and not determined by the major, subjects were asked about characteristics of the job they expected to hold. These questions included whether they would use the knowledge from that major in the job they expect to hold, and how family-friendly did they expect their job would be. Additionally, subjects reported how important they expected grades would be in determining their ability to get a job in each major.[6]

The third part of *Stage 1* asked subjects to guess, to the best of knowledge, average earnings, the average wage ratio, and the proportion of individuals who think women should tend the home while their husband works for each of the seven categories of majors. Subjects were instructed that their responses should be for workers age 30, working full-time with a Bachelor's degree.[7]

During *Stage 2*, subjects were randomly assigned to either the treatment or the control group. The control group was given information unrelated to the labor market or their educational choices. In this experiment, the students learned about monthly high and low temperatures in California. The treatment revealed average earnings, wage ratio, and measures of gender bias for the majors.

[6] The answers to these questions will be used as controls in the analysis.

[7] Their responses were scored using a scoring rule for a random subset of their responses. There are three questions with parts for each of the seven majors. One question was randomly selected (either average earnings, average wage ratio, or the proportion of individuals who think women should tend the home), and they received a payout based on how correct their responses were to the question. Scoring results should reduce the number of subjects simply guessing. Students' responses were scored using a linear scoring rule. As a student's response gets further from the true value, they earn less money. In this experiment, every 10% away the student was from the true value cost them 5¢.

Figure 2.2: Experimental Treatment: True Values of Population Averages by Major

These are the answers to the questions you were just asked previously. Please take a moment to read through them. You will then be asked a series of questions to test your comprehension of the table.

	Average Annual Earnings	Female to Male Earnings Ratio	Percent in occupation who believe that it is better for a man to work and a woman to stay at home
Arts and Humanities	$48,354	96%	26%
Biology and Life Sciences	$52,262	96%	26%
Business and Economics	$60,332	81%	28%
Engineering and Computer Science	$69,569	81%	26%
Physical Sciences and Mathematics	$51,700	94%	28%
Social Sciences	$49,932	80%	28%
No Degree	$37,121	77%	32%

Table 2.2 shows the information subjects in the treatment group were given. The average earnings and wage ratio were calculated using data from the 2013 American Community Survey (Ruggles et al. 2010).[8] Subjects were given the average earnings from wages for individuals at age 30 who are employed full-time that have graduated from college with a Bachelor's degree in the given major. The wage ratio was the ratio of average wage earnings between men and women (women divided by men). The ratio was calculated for individuals at age 30 who are employed full-time that have graduated from college with a Bachelor's degree in the given major. Drawing on the methodology of Charles and Guryan (2008) to measure bias, data from the 2012 General Social Survey is used to calculate the percent of individuals in an occupation who believe that women should tend to the home and their husbands should work. For each major, the gender bias is averaged over all the occupations where the graduates work. This measure provides a proxy for the implicit discrimination or micro-aggression that many argue keep women from pursuing STEM degrees or occupations (American Association of University Women 2010, Moss-Racusin et al. 2012, Williams, Phillips and Hall 2014).

In *Stage 3*, subjects were asked to re-state their beliefs that were elicited in *Stage 1*.[9] These

[8]The wage ratio has remained unchanged between 2000 and 2010 (Goldin 2014). It fell from 0.74 in 2000 to 0.72 in 2010. Therefore, in this experiment, I assumed that the next decade will see similar stability in wage ratio.)

[9]See Figure A.3 shows an example of how this was presented to subjects.

questions were identical to those previously. Within each major, the order of the questions was randomized. Due to the length of the experiment, subjects may have difficulty remembering all of the information presented to them. The treatment group was shown the true value for that major and then proceeded to answer questions about that major to minimize the probability of forgetting the information. The control group received no additional information.[10]

After completing the experiment, subjects were given a survey to collect demographic information and other potential controls. Subjects were asked about their parents' backgrounds, their academic performance, their optimism about the future labor market in each major, standard demographics, and their usage of career resources on campus. I elicited subjects' risk aversion using an incentivized risk elicitation task. Subjects could earn between $2.80 to $7 for this task. The average expected earnings for this task was $5. Figure A.4 illustrates these risk elicitation questions. Subjects were also given a cognitive reflection test that was not incentivized.

2.4 Hypotheses and Estimation Strategy

The goal of the experiment was to identify the effect of expectations of labor market discrimination on plans for human capital investment. I do this by comparing the error in the student's expectations to how they updated their human capital choices.

For the experiment to identify a causal effect, the information presented to students must be new. If students are wrong when they guess the population average, then giving them the true value of the population averages would constitute new information (Wiswall and Zafar 2015a). Based on the results from Wiswall and Zafar (2015a), I expect that students

[10]The control group was also administered a survey because the act of taking a survey may shift perceptions.

will be off by a significant amount when they guess the population averages. In Wiswall and Zafar (2015a), subjects overestimated earnings by \$32,620 for college-educated workers.

Hypothesis 1a: On average, subjects are misinformed about the labor market.

Hypothesis 1b: Students' misinformation is worse in STEM majors than in the Arts and Humanities or the Social Science.

To test this first hypothesis, I calculate the average error of each subject's guesses. I first test whether the errors are significantly different from zero. Then I regress the average error on a dummy for each major (with Arts and Humanities being the omitted group) and individual fixed effects to measure if subjects have more misinformation about STEM majors.

If Hypothesis 1a is true, then I can test the effect the experiment had on the choices of the subjects. This leads to two hypothesis about the effect of the experiment.

Hypothesis 2a: Correcting misinformation about the labor market can increase interest in STEM majors.

Hypothesis 2b: Women are more sensitive to the treatment than men.

To test the second hypothesis, I use a difference-in-difference estimation. To obtain a causal estimate of the impact of discrimination on the choice of a major, I subtract the *Stage 1* results from the *Stage 3* results. This removes individual fixed effects and any unobservables that do not vary by major or period. Because there may be unobservables that occur when an individual compares their choices across majors, each observation is made relative to a baseline major. In this paper, I select Arts and Humanities as the baseline major (denoted as m^*). This removes any unobservables that occur at the major-pair level. Regressing the change in the relative outcome on a dummy for receiving the information treatment will then identify the causal effect of the experiment on the expectations of educational obtainment.

$$\begin{aligned}(Y_{i,m,t+1} - Y_{i,m^*,t+1}) - (Y_{i,m,t} - Y_{i,m^*,t}) &= \alpha + \beta_1 T_i + \beta_2 (T_i \times F_i) \\ &+ [(X_{i,m,t+1} - X_{i,m,t}) - (X_{i,m^*,t+1} - X_{i,m^*,t})]\eta + \epsilon_{i,m,t}\end{aligned} \quad (2.3)$$

The outcome variable, $Y_{i,m,t}$, is a given belief in major m for individual i. In this paper, $Y_{i,m}$ is the probability that individual i will obtain a major in major m or $Y_{i,m}$ is the number of years of schooling that individual i plans to obtain in major m. Years of schooling is defined relative to a Bachelor's degree (i.e. a B.A. is worth 0, and a Ph.D. is worth five years).

The dummy for being in the treated group is T_i. The experiment asked a number of questions in both *Stage 1* and *Stage 3*.[11] The relative changes in these variables are included as controls in the regression. These controls were the change in expectations about the relative family-friendliness of the major, the change in the relative probability of using the knowledge gained in the major for their job, the change in relative perceived ability, and the change in relative future income. If the treatment affects the outcome, β_1 should be statistically significant. If β_2 is statistically significant, that means the effect of the treatment was different for men and women.

If the experiment shifts the subjects' choices, I can decompose the effect of the experiment to identify how errors in each of the measures shown contributed to the change in human capital choices.

Hypothesis 3: Women who overestimate the wage ratio or gender bias will adjust their choices to account for their error.

[11] These questions were whether they would use the knowledge from that major in the job they expect to hold, how family-friendly did they expect their job would be, their expectations for growth of the occupations, and how important they expected grades would be in determining their ability to get a job in each major.

To test the third hypothesis, I use a more flexible model than Equation 2.3. In this case, the errors in a subject's expectations are included. To test the effect of errors in expectations on outcomes, I estimate the following equation:

$$
\begin{aligned}
(Y_{i,m,t+1} - Y_{i,m^*,t+1}) - (Y_{i,m,t} - Y_{i,m^*,t}) = \alpha + \beta_1 T_i + \beta_2 (T_i \times F_i) \\
+\delta_1[(e^*_m - e_{i,m}) - (e^*_{m^*} - e_{i,m^*})] + \delta_2[(e^*_{i,m} - e_{i,m}) - (e^*_{i,m^*} - e_{i,m^*})] \times T_i \\
+\delta_3[(e^*_{i,m} - e_{i,m}) - (e^*_{i,m^*} - e_{i,m^*})] \times T_i \times F_i + \theta_1[(d^*_m - d_{i,m}) - (d^*_{m*} - d_{i,m*})] \\
+\theta_2[(d^*_m - d_{i,m}) - (d^*_{m*} - d_{i,m*})] \times T_i + \theta_3[(d^*_m - d_{i,m}) - (d^*_{m*} - d_{i,m*})] \times T_i \times F_i \\
+\gamma_1[(b^*_m - b_{i,m}) - (b^*_{m*} - b_{i,m^*})] + \gamma_2[(b^*_m - b_{i,m}) - (b^*_{m*} - b_{i,m^*})] \times T_i \\
+\gamma_3[(b^*_m - b_{i,m}) - (b^*_{m*} - b_{i,m^*})] \times T_i \times F_i \\
+[(X_{i,m,t+1} - X_{i,m,t}) - (X_{i,m^*,t+1} - X_{i,m^*,t})]\eta + \epsilon_{i,m,t}
\end{aligned}
\tag{2.4}
$$

I include interactions between the information treatment and the error in beliefs to test which pieces of information are driving the effects. I define $e_{i,m}$ as a subject's guess of the average earnings at age 30 if someone majored in major m. The true value of average earnings is given as e^*_m. The error in their beliefs is then $e^*_m - e_{i,m}$. Their expectations of the wage ratio $(d_{i,m})$ and their expectations of the gender bias $(b_{i,m})$ are also included. Errors are similarly defined as the true value minus the subject's expectation. Similar to before, m^* is the Arts and Humanities. All errors are then relative to the error in Arts and Humanities.

In this specification, δ_2 captures the effect of learning about the error in expected average earnings (learning about errors is the level of income). The effect of learning about the error in expected wage ratios is captured by θ_2. The effect of learning about the error in expected gender bias is captured by γ_2.

I interact a dummy for being female with the experimental treatment to isolate the effect of expectations of labor market discrimination on the human capital investment of female subjects. Therefore, θ_3 can be interpreted as the effect of learning about relative errors in expected wage ratios. If θ_3 is positive, it means female subjects who underestimate the relative wage ratio increase their probability of selecting that major. Assuming the assumptions of the model hold (e.g. the information presented is new and salient to students), I can interpret θ_3 and γ_3 as causal estimates of the effect of expectations of labor market discrimination on human capital investment for female students.

2.5 Results

Among the experimental subjects, there was a STEM interest gap between male and female subjects. In Stage 1, there was a 56% chance a male subject would graduate with major in STEM (Biology and Life Sciences, Engineering and Computer Science, and Physical Science and Mathematics). There was a 43% chance female subjects would graduate with a STEM major. Therefore, the STEM interest gap was 13 percentage points. The gap was larger in the treatment group than the control group (3% in the control group and 27% in the treatment group).

Similar to earlier experiments, I find subjects significantly misinformed about the labor market (Wiswall and Zafar 2015b, Wiswall and Zafar 2015a). Table 2.2 shows the average error in subjects' guesses for each major category. The results in Table 2.2 lend strong support for Hypothesis 1. In the majority of cases in Table 2.2, subjects errors are significantly different from zero. For all three measures, subjects were more incorrect in their guesses for STEM majors than in non-STEM majors.

Table 2.2: Errors in Beliefs

	Average Guess	Average Error
	Annual Income	
Arts and Humanities	$48266.50	87.50
	(18320.69)	(18320.69)
Biology and Life Sciences	$95466.26	-43204.26***
	(37153.83)	(37153.83)
Business and Economics	$81754.39	-21422.39***
	(29738.12)	(29738.12)
Engineering and Computer Science	$101656.00	-32087.03***
	(36196.74)	(36196.74)
Physical Sciences and Mathematics	$81146.89	-29446.89***
	(31749.78)	(31749.78)
Social Sciences	$61134.91	-11202.91***
	(27359.28)	(27359.28)
No Degree	$31974.86	5146.14
	(19444.14)	(19444.14)
	Female-Male Wage Ratio	
Arts and Humanities	.8200	.1400***
	(.2605)	(.2605)
Biology and Life Sciences	.7756	.1844***
	(.2450)	(.2450)
Business and Economics	.6685	.1415***
	(.2441)	(.2441)
Engineering and Computer Science	.6973	.1027***
	(.3133)	(.3133)
Physical Sciences and Mathematics	.7123	.2377***
	(.2741)	(.2741)
Social Sciences	.8385	-.0385
	(.2830)	(.2829)
No Degree	.62561	.1444***
	(.3051)	(.3051)
	Percent of Gender Biased Co-Workers	
Arts and Humanities	26.94%	-.94%
	(19.81)	(19.81)
Biology and Life Sciences	40.42%	-14.42%***
	(23.70)	(23.69)
Business and Economics	53.11%	-25.11%***
	(24.26)	(24.26)
Engineering and Computer Science	54.83%	-28.83%***
	(23.75)	(23.75)
Physical Sciences and Mathematics	48.55%	-20.55%***
	(24.25)	(24.25)
Social Sciences	29.85%	-1.85%
	(20.25)	(20.25)
No Degree	51.14%	-19.14%***
	(32.57)	(32.57)
Observations	66	66

Note: In both columns, the first number is the average and standard deviations are reported in parentheses below. In the second column, the error is measured as the true value minus the guess. The stars indicate the significance level of the difference from zero.
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

The first panel of Table 2.2 shows the average expected annual earnings and the average error in expectations. Subjects expect individuals who major in the Arts and Humanities to earn \$48,266 and Engineering and Computer Science majors to earn \$101,656. Subjects were furthest from the truth when they overestimated the average annual earnings of a graduate with a bachelor's degree in Biology and Life Sciences by \$43,204. Subjects were closest to the truth when they underestimated the average annual earnings of an Arts and Humanities graduate by \$87. The average errors are higher in STEM majors than in the other majors. On average, students overestimate what STEM majors will earn by \$34,912.73. The average overestimation in Arts, Humanities, and the Social Sciences is only \$5,645.21. This difference is statistically significant at the 1% level.

The second panel of Table 2.2 focuses on the gender wage ratio. Subjects expect the wage ratio between females and males will be highest in the Social Sciences or the Arts and the Humanities (0.84 and 0.82 respectively). STEM majors are expected to have a wage ratio between 0.70 and 0.78.[12] Subjects were closest to guessing the gender wage ratio in Social Sciences, where they overestimated the wage ratio by 3.9%. Subjects were the most off when guessing the gender wage ratio in the Physical Sciences and Math, where they underestimated the ratio by 23.8%. Subjects underestimate the wage ratio in STEM majors by 0.18. They underestimated the wage ratio in Arts, Humanities, and the Social Sciences by 0.05. This difference is statistically significant at the 1% level. While not reported in the table, female subjects had larger errors when guessing the wage ratio than male subjects. In STEM, female errors in the wage ratio were larger by 0.07. This suggests that subjects believe that the slope of Figure 2.1 is positive, and female subjects believe it to be steeper than male subjects.

[12]Higher expectations of gender bias were correlated with lower expectations for the wage ratio between female and male workers. The correlation coefficient between the two expectations was -0.14 and was highly significant. This suggests that subjects view gender bias is correlated with lower annual earnings for female workers. This relationship provides additional evidence that subjects are viewing the wage ratios as informing them about the discrimination in the labor market.

The third panel of Table 2.2 shows that expectations about gender bias are similarly misinformed. In Figure 2.2, I show the differences in gender bias by major are quite small (a 2% difference from the most biased to the least biased major). This is because college major and occupation are only loosely correlated. On average, 27% of coworkers for a college graduate will believe that it is better for a wife to tend to the home and a husband to work. Subjects were most accurate when guessing the gender bias of coworkers of individuals in the Arts and Humanities, where they were only off by 0.9%. Subjects were furthest from the truth in Engineering and Computer Science where they were off by 28.8%. Female subjects overestimate the percent of a STEM majors coworkers who will be biased against women in the workforce by 24 percentage points. Male subjects also overestimate gender bias against women in STEM, but by a smaller margin (11 percentage points). Subjects are more accurate in their estimates for the Arts, Humanities, and Social Sciences. In these majors, they are only off by two percentage points. There is no difference between male and female estimates of gender bias in these majors.

These errors in expectations partially explain the pattern seen in Figure 2.1. The reason why women concentrate in the lower-paying majors with lower wage ratios is that women expect the wage ratio to be higher in those majors. While the average earnings are less, the expected pay equality is attractive to female students. These majors also have lower expected levels of gender bias. This suggests that women are willing to trade higher incomes for work where they expect the pay to be more equal and there to be less gender bias.

To test if the information experiment shifted the preferences of students, I begin by looking at the updating of beliefs about the choice of a major and years of schooling. There were 193 observations of subjects shifting their preference for a major. The average absolute shift in those observations is 8.1% change in selecting a major. There were 184 observations of subjects changing their expected years of graduate education. The average change in years of graduate education is -0.12 years.

Table 2.3 reports the results for Hypothesis 2. The first three columns of Table 2.3 focus on the relative probability of selecting a major. The last three columns compare the effect of the treatment across majors on the expected years of schooling. The coefficient on $Treatment$ is the effect of the experiment on men in the treated group. In every case, I find the experiment did not result in men in the treatment group updating their preferences for a major or the number of years of graduate school they expected to obtain. The coefficient on $Treatment \times Female$ identifies whether the effect of the treatment for female subjects was different than for men. As predicted, female subjects in the treatment group are responding to the treatment by increasing their preference for non-Arts and Humanities majors. Female subjects in the treated group increased the probability of picking a non-Arts and Humanities major by 4.8%. The effect was larger in STEM majors than in non-STEM majors. Female subjects increased their interest in STEM majors by 5.7% and saw no significant change in their preferences for non-STEM majors. The fact that the magnitude of the change in non-STEM majors is also smaller is further evidence that the effect of the experiment was stronger in STEM majors. In the last three columns, I find there was no effect of the experiment on expected years of graduate schooling. Regardless of the comparison group that one looks at, there does not appear to be any difference in how subjects are behaving.

Table 2.3: Human Capital Investment Updating in Response to the Experiment

	Major Probability			Years of School		
	All Majors	STEM	Non-STEM	All Majors	STEM Only	Non-STEM
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	-0.0079	-0.0030	-0.0234	-0.6464	-0.5986	-0.7908
	(0.0157)	(0.0160)	(0.0347)	(0.5317)	(0.5794)	(0.6095)
Female	-0.0247	-0.0025	-0.0640*	-0.5189	-0.5243	-0.5483
	(0.0181)	(0.0172)	(0.0335)	(0.3767)	(0.4038)	(0.4470)
Treatment × Female	0.0481*	0.0568**	0.0471	0.7301	0.6034	1.0793
	(0.0264)	(0.0265)	(0.0517)	(0.6767)	(0.7470)	(0.7396)
Observations	330	198	132	330	198	132
R-squared	0.0221	0.0849	0.0524	0.0588	0.0755	0.0598

Note: Outcomes in this regression were relative to the Arts and Humanities major. Not shown, but included in the regression, are controls for the change in expectations about the relative family-friendliness of the major, the change in the relative probability of using the knowledge gained in the major, relative changes in perceived ability, and relative changes in future income. No demographic controls are included because they do not vary within an individual. Dropping Out is excluded from this analysis due to it being so unlikely to occur (mean probability less than 1%). Robust standard errors are reported in parentheses. Observations are clustered at the individual level.
*** p<0.01, ** p<0.05, * p<0.1

The baseline results showed the experiment shifted interest in STEM majors for female subjects. The net effect of the experiment was to decrease the STEM interest gap for women. In the treatment group, male subjects were initially 26.6% more likely to select a STEM major. The experiment reduced this STEM interest gap by 21% (5.7%/26.6%). The next step is to test Hypothesis 3 and determine what pieces of information drove this change. The experiment exposed students to three different pieces of information. Using Equation 2.4, I determine the impact of each piece of information. Table 2.4 reports the effect of errors in expectations on the probability of selecting a major relative to the Arts and Humanities and the relative years of graduate education a student plans to receive.

Table 2.4: Updating of Human Capital Investment Plans In Response to Errors

	Major Probability	Years of School
	(1)	(2)
Treatment	-0.0195	-0.9632
	(0.0175)	(0.6292)
Female	-0.0190	-0.4043
	(0.0178)	(0.3289)
Treatment × Female	0.0483	1.4513*
	(0.0435)	(0.8485)
Error in Earnings	0.0002	0.0051
	(0.0002)	(0.0044)
Error in Earnings × Treatment	-0.0000	-0.0249*
	(0.0004)	(0.0128)
Error in Earnings × Treatment× Female	-0.0011	0.0440***
	(0.0007)	(0.0158)
Error in Wage Ratio	-0.0100	-0.6031
	(0.0144)	(0.5760)
Error in Wage Ratio×Treatment	-0.0274	-0.2310
	(0.0399)	(1.2806)
Error in Wage Ratio×Treatment×Female	0.1407**	1.5296
	(0.0622)	(1.2987)
Error in Bias	0.0304	0.1910
	(0.0240)	(0.6297)
Error in Bias×Treatment	-0.0727	1.0924
	(0.0491)	(1.7453)
Error in Bias×Treatment×Female	0.2359*	-2.6285
	(0.1334)	(1.8766)
Observations	330	330
R-squared	0.0709	0.1292

Note: Outcomes in this regression were relative to the Arts and Humanities major. Not shown, but included in the regression, are the change in relative family-friendliness of the major, the change in the relative probability of using the knowledge gained in the major, relative changes in perceived ability, and relative changes in future income. No demographic controls are included because they do not vary within an individual. Dropping Out is excluded from this analysis due to it being so unlikely to occur (mean probability less than 1%). Robust standard errors are reported in parentheses. Observations are clustered at the individual level.
*** p<0.01, ** p<0.05, * p<0.1

Table 2.4 use the specification in Equation 2.4 to test if there was a differential effect of the treatment on women. The first column of Table 2.4 report the results for the probability of selecting a major and the last column reports the results for graduate years of schooling. In column 1, I find errors cause women to update the probability they will select a major. Women who underestimate the wage ratio in a major by 10% are 1.4% more likely to select that major after being shown the true value. This effect is statistically significant at the 5% level. Because women expect lower wage ratios in STEM majors (their average error was 0.18), correcting their misconceptions leads to higher participation. Updating their expectations of STEM wage ratios increases the self-rated probability they would graduate in a STEM field by 2.5%. The effect of correcting errors in expected wage ratios was 43.8% of the net effect of the experiment (2.5%/5.7%).

In column 2, I find subjects respond to errors in expected earnings when deciding whether or not to obtain graduate schooling. The experimental treatment increased the planned years of schooling for female subjects by 1.45 years. Male and female subjects responded to errors in opposite ways. Male subjects who overestimate earnings expect to obtain more schooling, likely to make up for the decline in their own expected earnings. Overestimating expected annual earnings by $1,000 increased the years of schooling in graduate school by 0.025 years for male subjects. For female subjects, overestimating the expected annual earnings in a major caused a student to decrease their expected years of schooling by 0.019. The overall effect of the experiment was a net increase of 0.86 years of schooling for the average female subject when taking into account the average relative error in earnings and the effect of being in the treated group. Expected years of schooling increased 0.59 years for the average male subject.

Based on the results in Table 2.4, I find evidence female subjects are responding in ways that are consistent with labor market discrimination driving them away from a major. Where low wage ratios push female students away from a major, there is a small, but not significant, attraction of men to those majors. These results suggest that the Hypothesis 2a is correct for women, but not for men. Wage ratios influence the human capital investment of women and underestimating the wage ratios causes women to avoid those majors. The fact that the male responses are of the correct sign, but not significant is in line with the stereotyping threat literature which predicts weaker responses of men to information about gender bias.

The previous analysis pooled all college majors and all subjects together. There may be differences in the ways that students react to information that varies by major. Errors in expectations about your current major may matter more than errors in other majors, or it could be the case that new information shifts your beliefs more in majors you are less familiar with. To understand these patterns, I explore the effect of the experiment in a student's own major and outside their major.

Table 2.5: Human Capital Investment Plans Within Own-Major and Outside Major

	Major Probability		Years of School	
	Own Major	All Others	Own Major	All Others
	(1)	(2)	(3)	(4)
Error in Earnings × Treatment× Female	-0.0027	-0.0012**	0.0233	0.0322**
	(0.0025)	(0.0006)	(0.0209)	(0.0150)
Error in Wage Ratio×Treatment×Female	0.3607	0.1405*	0.0365	1.1303
	(0.2515)	(0.0753)	(2.1633)	(1.7445)
Error in Bias×Treatment×Female	0.6904*	0.1200	-6.5638**	-3.5843**
	(0.4082)	(0.0844)	(2.9236)	(1.6089)
Observations	63	267	63	267
R-squared	0.2359	0.1316	0.2328	0.1187

Note: Outcomes in this regression were relative to the Arts and Humanities major. These regressions used the specification laid out Equation 2.4. Controls included in the regression are the change in relative family-friendliness of the major, the change in the relative probability of finding full-time work, the change in the relative probability of using the knowledge gained in the major, relative changes in perceived ability, and relative changes in future income. No demographic controls are included because they do not vary within an individual. Dropping Out is excluded from this analysis due to it being so unlikely to occur (mean probability less than 1%). Robust standard errors are reported in parentheses. Observations are clustered at the individual level.
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 2.5 looks at the effect of the treatment on the probability of majoring in your intended major versus the probability of majoring in another major. Subjects may react differently to information about their major than they react to information about alternative majors. Subjects were more likely to increase their preference for their major while decreasing their preference for other majors. The average change in the probability of majoring in your own major was 1.35%, and the average change in the probability of majoring in a major that is not your current major was -0.41%. The probability of selecting other majors is driven by the expected income and the expected wage ratio.

For the choice of schooling, the bias affects the terminal degree you expect to get in your major. The effect of errors in expected bias is twice as strong in your own major as it is in other majors. This suggests that the information on the relative bias is more salient when it tells you something about your current major. Earnings data informs students about their outside options. Overestimating the average annual income decreases the odds a subject selected a major outside their own. This pattern suggests that these information experiments mainly cause students to update their information about majors outside their own. This is not unexpected given they have more information about their own major and less about other majors. The larger the amount of imperfect information, the more new information causes students to update their expectations.

2.6 Conclusion

In this paper, I showed that subjects have incorrect expectations of average annual earnings, average wage ratios, and gender bias in the labor market. The errors in STEM majors were larger than the errors in other majors. I found that female subjects have higher errors in their expectations of STEM majors than male subjects. I showed that incorrect perceptions of wage ratios in STEM majors are causing female to prefer other majors to STEM majors.

Relative to the Arts and Humanities, female subjects assume that the wage ratios are smaller in STEM majors and the levels of gender bias are higher. In reality, the differences between majors are smaller than subjects expect. I was able to increase the interest of female subjects in STEM majors by correcting their misconceptions of STEM majors. The experimental treatment was able to increase the interest in STEM majors for the female subjects in the treated group by 5.7 percentage points. Approximately 44% of this change was due to correcting errors in expected wage ratios.

It appears lower expected wage ratios push women away from a major by making other majors more attractive. This mechanism is similar to the way that grades influenced female choices of college major Butcher et al. (2014). Much like higher grades led women to switch from STEM to the liberal arts, higher expected wage ratios are leading women to choose non-STEM majors. The misinformation about wage ratios is responsible for 9% of the STEM interest gap (2.5%/26.6%) in the treatment group of the experiment.

I find the experiment's effects were larger in majors outside the subject's current major. The larger treatment effect corresponded with larger errors in annual earnings, wage ratios, and gender bias. This pattern suggests the experiment worked by giving students more accurate information. This channel would explain why the presence of female role models in STEM majors may increase interest in STEM majors (Bettinger and Long 2005, Carrell, Page and West 2010). Having a female role model gives students more accurate information about the state of the labor market and improves students' expectation of future outcomes.

In contrast to the choice of a major, the years of graduate schooling is driven by expectations of future income. Men are better at guessing the average earnings in the labor market, but when they are wrong, they react by increasing their years of schooling. This suggests that money plays a dominant role for men when thinking about human capital investment along the extensive margin. They want to earn as much as possible and will increase their human capital accumulation to compensate for lower expected earnings. Female subjects responded to the experiment by increasing their expected years of schooling by a much smaller amount. The updating of female students was not driven entirely by their responses to their errors.

Taking the results together, the evidence from this experiment suggests that there is a role for public policy and interventions to close the STEM participation gap. Because of the large amount of imperfect information that students have, they respond to new information by updating their beliefs and expectations. For those seeking to increase female participation in STEM majors, low-cost information interventions about the STEM labor market paired with female role models appears to be a promising combination. Both work by providing female students with more information about the true state of the labor market, which may be more favorable than they initially assume.

Chapter 3

Not All Laws are Created Equal

3.1 Introduction

Beginning with Badgett (1995), researchers have accumulated evidence of disparities in the labor market outcomes between homosexuals and heterosexuals. Research has found gay men are paid less than heterosexual men (Klawitter 2015). The evidence of a gay wage gap exists across different datasets and is robust to various methods for identifying who is gay (Klawitter 2015).[1] There is inconclusive evidence of wage differentials for lesbian women, with differences in fertility and selection into the labor market potentially explaining the differences (Klawitter 2015).[2] In addition to the evidence of disparities in pay, there is consistent evidence from resume correspondence studies that heterosexual are preferred by hiring managers to homosexuals (Bailey, Wallace and Wright 2013, Mishel 2016, Tilcsik

[1]Klawitter (2015) is a meta-study of the wage differentials for gay men and lesbian women. The meta-study shows that despite the large variance in the estimates, there is consistent evidence of a wage penalty for gay men. See Allegretto and Arthur (2001), Antecol et al. (2008), Black, Makar, Sanders and Taylor (2003), Blandford (2003), Carpenter (2004), Cushing-Daniels and Yeung (2009), Elmslie and Tebaldi (2007), and Sabia (2014) for more detailed discussions of the gay wage penalty.

[2]See Klawitter (2015) for the results of the meta-study for the lesbian wage differential. Antecol and Steinberger (2013) and Jepsen (2007) contain more specifics of the challenges in estimating the lesbian wage differential.

2011). These correspondence studies provide the best causal evidence of discrimination against gay men and lesbian women in the labor market.[3]

Historically, the policy implemented in the United States to reduce these disparities has been to make it illegal for an employer to discriminate against individuals based on their membership in a protected group. When employment non-discrimination acts work as intended, the relative labor market outcomes of the protected group gradually improve, as appears to have happened for black men (Collins 2003, Donohue and Heckman 1991, Landes 1968, Neumark and Stock 2006). In other cases, however, the increased protections may make the protected group relatively more expensive to hire and terminate. Therefore, employers may reduce the number of employees they hire from the protected group (Bloch 1994), as may have happened for women and older workers (Beegle and Stock 2003, Lahey 2008, Neumark and Stock 2006).[4]

Employment non-discrimination acts for gay men and lesbian women are the best way to study the impact of employment protections because the existence of a federal law limits the differences that can exist between states. If there is a federal law banning discrimination, it creates a lower bound on the state laws. Any state law that provides weaker protections than the federal law would be superseded by the federal law. Despite being limited to studying the effect of a state law being stronger than the federal law, previous work on non-discrimination laws has shown that heterogeneity in the laws can impact the effect. Jolls and Prescott (2004) used state-level variation in disability discrimination laws to show that the negative employment effects of the Americans with Disabilities Act (ADA) were primarily due to "reasonable accommodations" requirements and not firing costs. Neumark

[3]Correspondence studies are the gold standard by which economists can measure discrimination, but there is the potential that the estimates obtained from these studies are not properly identifying discrimination. See Heckman (1998) and Neumark (2012) for a discussion of how the variance of unobservables may bias the results from correspondence studies.

[4]Though it should be noted, the presence of negative employment effects and whether they are short-term or long-term effects is a heavily debated topic in the literature (Acemoglu and Angrist 2001, Adams 2004, Beegle and Stock 2003, DeLeire 2000, Jolls and Prescott 2004, Kruse and Schur 2003, Lahey 2008, Neumark and Stock 1999).

and Button (2014) showed that stronger state-level laws against age discrimination might have reduced the hiring of older workers during the Great Recession.

In the case of ENDAs for gay men and lesbian women, all state laws are binding. Therefore, I can study the effect of weak laws as well as strong laws. The ability to disaggregate the laws has been ignored by the previous research, which treated state laws as identical (Baumle and Poston Jr. 2011, Klawitter and Flatt 1998, Klawitter 2011, Martell 2013b). In this paper, I provide an important contribution to the literature by showing how differences in damages, employer size minimums, and the statute of limitations of complaints lead to differences in the outcomes for cohabiting gay men and lesbian women.

Using a difference-in-differences-in-differences methodology, I show that ENDAs were effective at reducing wage gaps between cohabiting gay men and married heterosexual men. After an ENDA had been passed, hourly wages for cohabiting gay men rose 2.7%. There was no significant effect on annual income from wages, employment, or hours worked for gay men. For lesbian women, the results are less positive. The passage of an ENDA had no significant effect on the wage differentials for lesbian women, but was associated with a 1.7% decline in the employment and a 0.733-hour decline in hours worked.

The results of this paper highlight how looking at the average effect of ENDAs ignores the heterogeneous impact the laws have had. The results clearly show the benefits of the law were dependent on the structure of the law. When comparing the effects across legal regimes, the most important determinant of the size of the wage increase was the type of damages available to plaintiffs. For gay men in states where successful plaintiffs cannot be awarded damages, there was no effect of an ENDA. Gay men in states that only allow compensatory damages experienced an increase in annual wages of 12.3% and an increase in hourly wages of 15.5%. Gay men in states that allow both compensatory damages and punitive damages experienced an increase in annual wages of 7.9% and an increase in hourly wages of 7.4%. In states with longer complaint periods, I found small increases in the employment of gay men.

A 1-month increase in the length of the complaint period above the average length increased the employment of gay men by 0.2%. After controlling for the differences in provisions, the average ENDA decreased the annual income of lesbian women by 11.4%. I find that stronger damages further decreased the employment of lesbian women. Allowing for punitive damages decreased the hours worked for lesbian women by an additional 0.850 hours.

3.2 Overview of Employment Non-Discrimination Legalization

Wisconsin passed employment protections for gay men and lesbian women in 1982. In 1990, only Wisconsin, Massachusetts, and D.C. had employment protections for gay men and lesbian women. By 2000, the number of states had grown to twelve. In 2014, twenty-one states and Washington, D.C. had employment non-discrimination acts protecting gay men and lesbian women. As shown in Figure 3.1, the Northeast, Midwest, and the West feature the highest concentration of these laws.[5]

[5]The regional pattern of ENDAs suggests that the states that pass these laws may not be selected at random. In Section 3.6, I show how prejudice changes in a region before and after the passage of a law to test whether the timing of these laws is exogenous to changes in prejudice.

Figure 3.1: Map of ENDA Laws in 2011

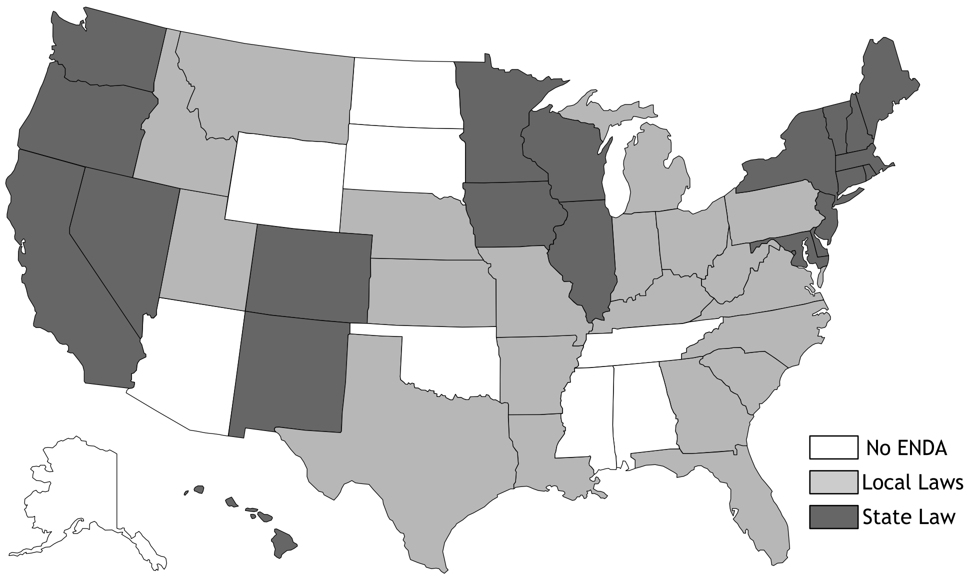

Note: Data on state and local laws comes from (Government Accountability Office 1997), (Government Accountability Office 2013), (Human Rights Campaign 2012), and (Sears et al. 2009), and by reading the state laws themselves.

Despite an increasing number of states passing laws to protect gay men and lesbian women, there is little research on how effective these policies have been at reducing disparities in the workplace for gay men and lesbian women. The research on LGBT non-discrimination acts has found mixed evidence for the effectiveness of employment protections at the state level. Klawitter and Flatt (1998) found there was no effect of employment protections on the wage or employment differences between cohabiting homosexuals and heterosexuals using the 1990 Census. Klawitter (2011) and Baumle and Poston Jr. (2011) revisited the question using 2000 Census data. In both papers, the authors found that ENDAs had no impact on the labor market outcomes of lesbian women. For gay men, Baumle and Poston Jr. (2011) found that ENDAs increased the annual earnings of gay men by 2.6%, but Klawitter (2011) showed that this increase was mainly due to ENDAs increasing weeks worked. Using the General Social Surveys, Martell (2013b) found that employment non-discrimination laws decrease the wage gap between gay men and heterosexual men between 2% and 15% each year that the law has been in place.

In each of these papers, the authors assumed that all laws were identical. I provide evidence of the wide variation in these state laws. When comparing state laws, the differences appear on three issues: who is protected under the law, how a complaint is resolved, and what damages and remedies are available for plaintiffs. Within these three groups of differences, there are thirteen provisions over which states differ. Table 3.1 details the areas where state laws diverge. Information about provisions comes from state laws and reports compiled by the Williams Institute and the Government Accountability Office (Sears et al. 2009, Government Accountability Office 2013). Tables A.3, A.4, A.5, A.6, and A.7 detail the differences between the state laws. In this paper, I focus on damage availability, employer size minimums, attorney's fees, and the statute of limitations.

Table 3.1: States with Each Legal Provision in its ENDA

	2000	2012
ENDA Law	12	22
Damage Awards	2000	2012
Equitable Relief	12	22
Compensatory Damages	11	20
Punitive Damages	9	14
Attorney's Fees	10	19
Statute of Limitations	2000	2012
120 Days	0	1
180 Days	7	12
300 Days	3	4
365 Days	3	5
Employer Size Minimums	2000	2012
1 Employee	6	9
3 Employees	1	1
4 Employees	1	4
5 Employees	1	1
6 Employees	2	2
8 Employees	0	1
15 Employees	1	4

Note: Information on state laws comes from (Sears et al. 2009), the Government Accountability Office (Government Accountability Office 2013), and information from state laws collected by the author.

There are three categories of damages: equitable relief, compensatory damages, and punitive damages. All states allow for equitable relief, which consists of remedies such as backed pay or being reinstated to your job. Compensatory damages are used to replace lost earnings and compensate for pain and suffering. Eighteen states provide for compensatory damages in their laws. Punitive damages are designed to punish egregious violations of the employment non-discrimination laws and are determined by the seriousness of the violation, not the damage done to the plaintiff. Thirteen states provide for punitive damage in their statutes. Damage awards may be capped, with the amount that damages are capped at varying by state. Some states cap the damages based on the size of the employer, while others cap the awards at a set amount. In eighteen states, it is possible for a successful plaintiff to recoup attorney's fees as part of the damage awards.

The statute of limitations for complaints determines how long employees have to file their complaints (e.g. the complaint period). The average statute of limitations in states with an ENDA is 241 days (approximately eight months). States range from 120 days to 365 days. There are 14 states with statutes of limitations of six months or shorter. The employer size minimums determine how large a firm must be before they have to comply with the law. The minimums range from one employee to fifteen employees. In states with a size minimum of one, all employers are covered. This occurs in nine states. There are four states with a size minimum of 15 employees, which is equal to the federal employer size minimum for discrimination laws.

These legal differences are important because they determine the expected cost of discriminating for employers. The expected cost of discriminating can be increased either by increasing the probability that an employer is sued or by increasing the cost of being sued. The probability that an employer is sued potentially increases when employees are given longer to file a complaint and they can recoup their attorney's fees in a successful lawsuit (which lowers the cost of a lawsuit). Stronger damage provisions increase the expected cost of discriminating since employers who do discriminate face stiffer penalties if caught.

3.3 Data

The data used in this paper comes from the 2008 through 2014 American Community Survey (ACS) 1-Year Samples, the 1990 U.S. Census 5% Sample, and the 2000 U.S. Census 5% Sample (Ruggles et al. 2010). A longitudinal database of all state laws was created using information from The Williams Institute at UCLA, the Government Accountability Office, and state laws (Government Accountability Office 2013, Sears et al. 2009, Sears and Mallory 2011).

To identify gay men and lesbian women in the United States, the Census collects information on householders and the relationships of everyone in the household to the householder. A same-sex couple is identified when the gender of the householder and the gender of the unmarried partner (or spouse) of the householder are the same. There is no information on single gay men and lesbian women in the Census data or ACS, only cohabiting gay men and lesbian women. Also missing from the sample are gay men and lesbian women in a household where one of the partners is not the household head (such as living with one's parents). Therefore, the sample in the analysis is restricted to comparisons between cohabiting individuals.

The sample used in the analysis begins with all adults older than 18 who claim to be the householder, spouse, or unmarried partner. In this paper, I focus on comparisons between cohabiting same-sex couples and married heterosexual couples. Cohabiting gay men and lesbian women are identified in the sample if they are cohabiting with an individual of the same gender.[6] Once I identify cohabiting same-sex couples, I restrict the sample to

[6] Cohabiting is defined as either being married or in an unmarried partnership. Unmarried partnerships are defined as relationships where the unmarried partner shares a close personal relationship with the reference

individuals over the age of 22 and under the age of 65. The age of 22 is selected to avoid school-aged individuals.

In the data, heterosexual couples are miscoded as same-sex cohabiting couples if the sex of one of the individuals is miscoded. Even though miscoding of sex is one of the least common errors made on Census forms, due to the small size of the gay and lesbian population, any miscoding in the heterosexual sample has the potential for increasing the number of gay men and lesbian women in a significant way (O'Connell and Golding 2006). When looking at the Census forms, researchers have that the miscoding is concentrated in the group of individuals who claimed to be married on their Census form and found smaller amounts of miscoding in the sample of same-sex couples who claim to be unmarried partners (O'Connell and Golding 2006, O'Connell and Loftquist 2009). Work on the wording of relationship questions in the ACS resulted in a large decline in the number of miscodings in the 2008 and subsequent surveys. Because before 2004 it was not possible for same-sex couples to be married, I exclude any gay or lesbian that had their marital status recoded in 2000 from married to cohabiting. This is a conservative portioning of the cohabiting gay and lesbian population, but it reduces the measurement error.

The other concern is that the passage of ENDAs laws may induce gay men and lesbian women to move to these states. Previous research has found mixed evidence that gay men and lesbian women migrate in response to the passage of pro-gay laws. Ueno, Vaghela and Ritter (2014) found that gay men were no more likely than heterosexual men to migrate to a different state, but lesbian women are more likely than heterosexual women to move. When gay men and lesbian women do move, there is inconclusive evidence that gay rights laws influence their choices. Colvin and Riccucci (2002) found no evidence that when they migrated that cohabiting gay men and lesbian women were more likely to move to a state that had passed an ENDA. Beaudin (2017) found evidence that cohabiting gay men and lesbian

person.

women were more likely than cohabiting heterosexuals to move to states with marriage equality. Because Colvin and Riccucci (2002) was estimated using a five-year window (while Beaudin (2017) used a one-year window), the results may not be as useful in evaluating the risk of current migration biasing the results. To avoid the migration problem completely, I remove any individual who has migrated within the past year from the ACS data. This means that whenever possible (i.e. the latter half of the sample), the effect of an ENDA will only be identified off of the wages of individuals who have been in the state for more than a year.

3.4 Estimation Strategy

The goal of this paper is to estimate the impact of employment non-discrimination acts on the labor market outcomes of cohabiting gay men. I focus on the relationship between the strength of the laws and the impact of the laws. Because the laws in question vary by state over time, I use the state and year variation to estimate a differences-in-differences-in-difference model. Since the passage of an ENDA does not impact heterosexuals, I can use the comparison of gay men and lesbian women and their married heterosexual counterparts to isolate the effect these laws have on the labor market outcomes of cohabiting gay men and lesbian women. I use a flexible model, saturated with year, state, and state-by-year fixed effects interacted with being homosexual:

$$\begin{aligned} Y_{ist} &= \beta_0 + \beta_1 Homo_{ist} + \beta_2 (Homo_{ist} \times ENDA_{st}) + I_s\gamma_s + I_t\gamma_t \\ &+ (I_s \times I_t)\gamma_{st} + (Homo_{ist} \times I_s)\theta_{H,s} + (Homo_{ist} \times I_t)\theta_{H,t} + \mathbf{X}_{ist}\delta + \epsilon_{ist} \end{aligned} \tag{3.1}$$

Y_{ist} is the dependent variable of interest. In this paper, I look at log annual income from wages, log hourly wages, the probability of being employed, and hours worked. $Homo_{ist}$ is a dummy for being a cohabiting gay man or lesbian woman. $ENDA_{st}$ is a dummy for state

s having an employment non-discrimination act that protected gay men and lesbian women in year t. The vector $\mathbf{X}$ contains controls for demographic, occupation, and geographic variables that may affect wages or employment. See Appendix Table A.1 for the full list of control variables.

Also included in the regression are state, year, and state-by-year fixed effects (γ_s, γ_t, and γ_{st}). These fixed effects will capture unobserved differences common to all observation in a state, year, and state-by-year cell. To account for the differences between cohabiting homosexuals and married heterosexuals that exist across states, state fixed effects are interacted with the dummy for cohabiting same-sex couples ($\theta_{H,s}$). To account for common trends that affect cohabiting same-sex couples differently than heterosexual couples (e.g. declines in prejudice, business cycle fluctuations, etc.), year fixed effects are interacted with the dummy for cohabiting same-sex couples ($\theta_{H,t}$). Due to individual preferences potentially being correlated within a state and treatment occurring at the state-level, the standard errors (ϵ) are clustered at the state-level.

I estimate the models for men and women separately. In the analysis of wages (both annual and hourly), the sample was restricted to individuals in the labor force.[7] For the analysis of employment and hours worked per week, the sample of all adults is used.

In Equation 3.1, the differential effect of being a cohabiting same-sex couple in the United States across all years is captured by β_1. How this differential changes over time is captured by $\theta_{H,t}$. The state fixed effects for cohabiting same-sex couples ($\theta_{H,s}$) capture how the wage differential varies across states. The parameter of interest is β_2, which captures how the differential between homosexuals and heterosexuals in states with ENDAs changed after the passage of a law.

To interpret β_2 as being the causal effect of enacting employment protections for sexual

[7]To test the effect of ENDAs across the labor force, I split the sample by full-time and part-time work status as a robustness check.

orientation, it must be the case that there are no other factors related to changes in the gay-heterosexual wage differential in states that do and do not have LGBT employment protections. I test these assumptions in Section 3.6. In addition to testing the assumptions of the DDD model, I test how robust the results are to confounding factors, specifically same-sex marriage and selection into the labor market. Results of the robustness checks are also reported in Section 3.6.

The main contribution of this paper is to test how the heterogeneity of the law affects the impact of these laws. To do this, I add a series of controls in Equation 3.1 to capture policy differences in the state law. I estimate the following equation:

$$\begin{aligned} Y_{ist} = \beta_0 + \beta_1 G_{ist} + \beta_2 (G_{ist} \times ENDA_{st}) + \beta_p (G_{ist} \times ENDA_{st} \times P_{s,t}) + \mathbf{X}_{ist}\delta + I_s \gamma_s \\ + I_t \gamma_t + (I_s \times I_t)\gamma_{st} + (G_{ist} \times I_s)\theta_{G,s} + (G_{ist} \times I_t)\theta_{G,t} + \epsilon_{ist} \end{aligned} \tag{3.2}$$

where $P_{s,t}$ is a vector of policy variables. States rarely change their non-discrimination laws, but there have been amendments or court cases that resulted in changes to the laws.[8] The policy variables used here are the availability of damages (compensatory and punitive), the statute of limitations for complaints, the employer size minimum, and the ability to recoup attorney's fees. The availability of damages and the ability to recoup attorney's fees enter into Equation 3.2 as dummy variables. The statute of limitations for complaints and employer size minimums have been normalized to the average value to make interpreting the coefficients easier. The average complaint period in a state with an ENDA is 245 days. State complaint periods are coded as the number of months relative to the average. The average size of firm size minimum for being covered by a law is five employees. For states with an ENDA, size minimums are coded as the employer size minimum minus five.

[8]See the tables about laws in the Appendix for details on how these provisions have changed within states over time. A case of this happening was Connecticut added compensatory damages to the list of remedies available after an appeals court ruled in 2000 that the statues prescription of such legal and equitable relief which the court deems appropriate and attorneys fees and costs included compensatory damages (Sears et al. 2009).

3.5 Main Results

In this section, I begin by discussing the effect of ENDAs on gay men and then move to discussing the results for lesbian women. For both groups, I first look at the average effect of the law before comparing the effects of different legal provisions. I focus on four labor market outcomes: log annual income from wages, log hourly wages, employment, and hours worked per week. ENDAs may increase the wages of gay men and lesbian women if the laws force employers to pay them the same wages they pay their heterosexual employees. These laws can also increase the employment of gay men and lesbian women if they make it easier for them to find a job. They may increase the hours worked if they allow gay men and lesbian women to obtain full-time employment or find a second job.

As a starting point, Table 3.2 shows how the hourly wage differential has evolved over time for gay men and lesbian women.[9] In Table 3.2, gay men earned 12.6% less than a comparable heterosexual man in 1990, but this penalty declined by 4.7 percentage points between 1990 and 2014. The decline has been most pronounced in states with an ENDA. States with an ENDA saw a decline of 10.2 percentage points and states with no ENDA saw an increase of 0.2 percentage points. Lesbian women earned 2.2% more than a comparable heterosexual woman, and this has not changed much since 1990. States with ENDAs and without ENDAs did not see as stark differences in the wage differentials as was observed for gay men.

[9] There is an extensive literature that discusses the nature of these wage differentials and whether they are driven by discrimination or other unobserved differences between the groups (Klawitter 2015), but from a legal perspective, the type of discrimination that leads to the observed wage penalty for gay men does not matter. The passage of an ENDA makes both statistical discrimination and taste-based discrimination illegal.

Table 3.2: Gay and Lesbian Wage Differentials by Year

	Men		
	All states	ENDA	No ENDA
	(1)	(2)	(3)
1990	-0.126***	-0.123***	-0.130***
	(0.008)	(0.012)	(0.010)
2000	-0.089***	-0.068***	-0.114***
	(0.013)	(0.015)	(0.012)
2010	-0.061***	-0.042	-0.081***
	(0.016)	(0.024)	(0.014)
2014	-0.079***	-0.021	-0.132***
	(0.018)	(0.015)	(0.016)

	Women		
	All states	ENDA	No ENDA
	(1)	(2)	(3)
1990	0.022	0.032	0.003
	(0.013)	(0.017)	(0.015)
2000	0.031***	0.030*	0.032**
	(0.008)	(0.011)	(0.011)
2010	0.056***	0.045**	0.068***
	(0.009)	(0.014)	(0.011)
2014	0.025*	0.042**	0.013
	(0.010)	(0.012)	(0.015)

*** $p<0.001$,** $p<0.01$, * $p<0.05$

Note: Author's calculations based on data from the 2008 through 2014 American Community Surveys 1% PUMS, the 1990 Decennial Census 5% PUMS, and the 2000 Decennial Census 5% PUMS. The outcome variable in this table is log hourly wages. Hourly wages are in constant 1999 dollars. The coefficients correspond to the interaction of being gay or lesbian and the year dummy. All regressions are estimated using OLS and include demographic and occupation controls, state fixed effects, year fixed effects, and state-by-year fixed effects. For a full list of control variables, see Table A.1. Standard errors clustered at the state level are in parentheses.

I begin by looking at the effect of ENDAs on cohabiting gay men in Table 3.3. In Panel A of Table 3.3, I replicate what the previous literature has done and use a single dummy variable for having an ENDA. The first column of Table 3.3 reports the results for log annual income, the second column reports the results for log hourly wages, the third column reports the results for employment, and the fourth column reports the results for usual hours worked per week. I find no effect of passing an ENDA on the log annual income of gay men. There was, however, a positive effect of passing an ENDA on log hourly wages. Between 1990 and 2014, the passage of an ENDA increased the hourly wages of cohabiting gay men by 2.7%. This increase in hourly wages was statistically significant at the 5% level. Similar to results found in the previous literature, the passage of an ENDA had no significant effect on the employment of cohabiting gay men.[10] I also find no statistically significant effect of the passage of an ENDA on hours worked per week.[11]

[10]The results reported in the paper are those from a linear probability model. The results using a probit are similar in magnitude and significance.

[11]While Klawitter (2011) looked at weeks worked per year in the 2000 Census, the ACS reports the usual weeks worked per year in discrete intervals, so I do not use this as an outcome.

Table 3.3: Effects of ENDAs on Labor Market Outcomes of Cohabiting Gay Men

	Baseline Effect			
Panel A.	Ln Annual Income	Ln Hourly Wages	Employment	Hours Worked
	(1)	(2)	(3)	(4)
Gay × ENDA	0.017	0.027*	0.005	0.196
	(0.013)	(0.012)	(0.005)	(0.194)
Observations	6,172,273	6,172,273	7,660,401	7,660,401

	Effect of Provisions			
Panel B.	Ln Annual Income	Ln Hourly Wages	Employment	Hours Worked
	(1)	(2)	(3)	(4)
Gay × ENDA	0.040	0.028	-0.006	0.097
	(0.027)	(0.021)	(0.009)	(0.361)
Gay × ENDA × Employer Size Minimum	-0.003	-0.005	0.001	-0.016
	(0.003)	(0.003)	(0.001)	(0.024)
Gay × ENDA × Complaint Period	-0.001	-0.006	0.002*	0.078
	(0.004)	(0.003)	(0.001)	(0.044)
Gay × ENDA × Compensatory Damages	0.083***	0.127***	0.003	-0.650
	(0.020)	(0.020)	(0.012)	(0.697)
Gay × ENDA × Punitive Damages	-0.044*	-0.081***	-0.001	0.089
	(0.020)	(0.018)	(0.012)	(0.390)
Gay × ENDA × Attorney's Fees	-0.066	-0.027	-0.017	0.114
	(0.034)	(0.030)	(0.012)	(0.845)
Observations	6,172,273	6,172,273	7,660,401	7,660,401

*** $p<0.001$, ** $p<0.01$, * $p<0.05$

Note: Author's calculations based on data from the 2008 through 2014 American Community Surveys 1% PUMS, the 1990 Decennial Census 5% PUMS, and the 2000 Decennial Census 5% PUMS. The sample is restricted to men, older than 22 and younger than 65. Men are either married heterosexual men or cohabiting gay men. Men who have moved between states in the past year are not included in the analysis. Annual income and hourly wages are in constant 1999 dollars. All regressions are estimated using OLS and include demographic and occupation controls, state fixed effects, year fixed effects, and state-by-year fixed effects. For a full list of control variables, see Table A.1. Standard errors clustered at the state level are in parentheses.

The results in Panel A suggest the declines in the wage penalty for gay men observed in Table 3.2 in states that passed an ENDA are due in part to the passage of these laws. In states with an ENDA, the wage penalty for gay men declined 10.2 percentage points between 1990 and 2014. The passage of ENDAs can explain 26% of this decline (2.7% divided by 10.2%).

Panel B in Table 3.3 reports the effect each policy had on outcomes. Among states with ENDAs, the average size of firm size minimum for being covered by the law is five employees. Decreasing the firm size minimum by one employee relative to the average minimum will increase the number of firms that are covered by the law. For all the outcomes studied here, I find no significant effect of smaller firm size minimum. I find similar results for the complaint period. When I compare the effects of the complaint period across the different outcomes, I find that a longer statute of limitations only has a significant effect on employment. Increasing the complaint period by 1-month above the average increases the employment of gay men by 0.2%.

Where I find the strongest effects is in damages. Damages have a significant effect on both annual income and hourly wages. In states that allow compensatory damages, the annual income of gay men increase 9.4%, and the hourly wages of gay men increase by 12.7%. When a state allows for punitive damages, the annual income of gay men falls 4.4%, and the hourly wages of gay men fall 8.1%. When looking at the damages provisions together, there is a distinct pattern of decreasing returns to strength. Providing for damages is important to the increase in earnings of gay men since it is a credible message to discriminating firms that there is a cost to engaging in discriminatory behavior. But, firms increase the wage of gay men less as the cost of discriminating increases. In states that allow only compensatory damages, there is a 12.8% net increase in annual wages. In states that allow both compensatory damages and punitive damages, there is an 8.8% net increase in annual wages. A similar pattern exists for hourly wages.

Table 3.4 shows the average effect of ENDAs for lesbian women appears similar to the effects found in earlier research on non-discrimination laws for women (Neumark and Stock 2006). In Panel A of Table 3.4, I show there was no increase in annual income or wages after the passage of the ENDA, consistent with what earlier work on ENDAs found for lesbian women (Klawitter and Flatt 1998, Klawitter 2011). My results for the employment effects for lesbian women show that ENDAs were detrimental to the employment of lesbian women, where the Klawitter and Flatt (1998) and Klawitter (2011) found no effect. There was a 1.7% decline in employment after an ENDA is passed. I also find a 0.733-hour decline in hours worked for lesbian women.

Table 3.4: Effects of ENDAs on Labor Market Outcomes of Cohabiting Lesbian Women

	Baseline Effect			
Panel A.	Ln Annual Income	Ln Hourly Wages	Employment	Hours Worked
	(1)	(2)	(3)	(4)
Lesbian × ENDA	-0.019	-0.003	-0.017*	-0.733**
	(0.017)	(0.014)	(0.007)	(0.255)
Observations	5,061,060	5,061,060	8,060,889	8,060,889

	Effect of Provisions			
Panel B.	Ln Annual Income	Ln Hourly Wages	Employment	Hours Worked
	(1)	(2)	(3)	(4)
Lesbian × ENDA	-0.114*	-0.027	-0.0156	-1.329***
	(0.049)	(0.037)	(0.012)	(0.337)
Lesbian × ENDA × Employer Size Minimum	0.005	0.004	0.002	0.017
	(0.004)	(0.003)	(0.001)	(0.037)
Lesbian × ENDA × Complaint Period	0.007	0.006	0.001	-0.022
	(0.006)	(0.005)	(0.002)	(0.060)
Lesbian × ENDA × Compensatory Damages	-0.025	-0.089	-0.016	0.295
	(0.056)	(0.046)	(0.022)	(0.414)
Lesbian × ENDA × Punitive Damages	-0.033	-0.003	0.005	-0.850*
	(0.028)	(0.024)	(0.016)	(0.322)
Lesbian × ENDA × Attorney's Fees	0.041	0.036	-0.008	0.672
	(0.064)	(0.051)	(0.021)	(0.518)
Observations	5,061,060	5,061,060	8,060,889	8,060,889

*** $p<0.001$,** $p<0.01$, * $p<0.05$

Note: Author's calculations based on data from the 2008 through 2014 American Community Surveys 1% PUMS, the 1990 Decennial Census 5% PUMS, and the 2000 Decennial Census 5% PUMS. The sample is restricted to women, older than 22 and younger than 65. Women are either married heterosexual women or cohabiting lesbian women. Women who have moved between states in the past year are not included in the analysis. Annual incomes and hourly wages are in constant 1999 dollars. All regressions are estimated using OLS and include demographic and occupation controls, state fixed effects, year fixed effects, and state-by-year fixed effects. For a full list of control variables, see Table A.1. Standard errors clustered at the state level are in parentheses.

The second panel of Table 3.4 explores how the legal provisions influence the effect of an ENDA for lesbian women. By breaking out the effect of an ENDA by provision, I show that for the most part there is no effect of the provisions on the wages or employment of lesbian women. There are two exceptions. First, after controlling for the differences in provisions, the average ENDA now has the effect of decreasing the annual income of lesbian women by 11.4%. The specific provisions do not appear to mitigate this effect, suggesting this result is driven by correlations amongst the provisions. Second, I find that stronger damages may further decrease the employment of lesbian women. Allowing for punitive damages decreases the hours worked for lesbian women by an additional 0.850 hours.

These costs and benefits of ENDAs for gay men and lesbian women are not spread evenly across the labor force. In Table 3.5, I compare the effect of ENDAs on wages for those working full-time and part-time.[12] I find significant differences in how ENDAs impacted full-time and part-time workers. For gay men who were working more than 30 hours a week, ENDAs increased their wages by 3.0%-3.2%. Gay men who were working part-time (less than 30 hours a week) did not experience a wage increase after an ENDA was passed. For lesbian women, the pattern is flipped. There is weak evidence that lesbian women working part-time may have received wage increases. I find a 12.0% increase in hourly wages of lesbian women working part-time, but no significant effect on annual income from wages. The effects for full-time lesbian women are negative for both groups, but not significant at the 5% level.

[12]See Appendix Table A.8 for the results broken out by provision.

Table 3.5: Effect of ENDA on Earnings by Full-Time and Part-Time Employment Status

	Men			
	Full-Time		Part-Time	
Panel A.	Ln Annual Income	Ln Hourly Wages	Ln Annual Income	Ln Hourly Wage
	(1)	(2)	(3)	(4)
Gay × ENDA	0.032*	0.030*	-0.154	-0.006
	(0.016)	(0.014)	(0.101)	(0.104)
Observations	5,985,370	5,985,370	186,903	186,903

	Women			
	Full-Time		Part-Time	
Panel B.	Log Annual Income	Log Hourly Wages	Log Annual Income	Log Hourly Wage
	(1)	(2)	(3)	(4)
Lesbian × ENDA	-0.007	-0.013	0.082	0.120*
	(0.015)	(0.014)	(0.093)	(0.055)
Observations	4,213,547	4,213,547	847,513	847,513

*** p<0.001,** p<0.01, * p<0.05

Note: Author's calculations based on data from the 2008 through 2014 American Community Surveys 1% PUMS, the 1990 Decennial Census 5% PUMS, and the 2000 Decennial Census 5% PUMS. The sample is restricted to adults, older than 22 and younger than 65. Individuals are either married heterosexuals or cohabiting gay men or lesbian women. Individuals who have moved between states in the past year are not included in the analysis. Annual earnings and hourly wages are in constant 1999 dollars. Full-time employment is defined as working 30 hours per week or more. Part-time employment is defined as working less than 30 hours per week. All regressions are estimated using OLS and include demographic and occupation controls, state fixed effects, year fixed effects, and state-by-year fixed effects. For a full list of control variables, see Table A.1. Standard errors clustered at the state level are in parentheses.

3.6 Robustness Checks and Threats to Validity

So far the results have shown increases in the wages of gay men and declines in employment and hours worked for lesbian women as a result of passing an ENDA. This next section addresses potential threats to the validity of these results.

The differences-in-differences methodology rests on the assumption that the unobservables are uncorrelated with the treatment. Error terms of the wage equation may not be parallel if the level of discrimination is changing faster in states with employment protections than in states without employment protections. Using responses from the General Social Survey, it is possible to calculate the percent of individuals that express prejudiced sentiments to questions about homosexuality. Research has shown that the wages of gay men are correlated with the share of individuals in a state who give prejudiced answers to questions about homosexuality in the General Social Survey (Burn 2017).[13]

There is no publicly available data on prejudice at the state level, so the publicly available Census division level data is used. In the General Social Survey, there are four questions about homosexuality asked in every wave. Table 3.6 details the text of each question and the possible answers. The questions in the GSS ask a respondent's feelings about sexual relations between adults of the same gender, whether they support homosexuals teaching in colleges, whether they support books promoting homosexuality to be housed in public libraries, and whether homosexuals should be able to give speeches in favor of homosexuality in public. For every Census division, I calculate the share of individuals giving the most prejudiced answer to all the questions. Since the General Social Survey is asked every two years, I impute the odd years using the mean of the preceding and succeeding shares of prejudiced individuals. I then estimate the rate at which states that pass an ENDA are growing less

[13]Research into wage penalties for black men has also found significant correlations between wage penalties and prejudice in the General Social Survey (Bond and Lehmann 2015, Charles and Guryan 2008). The calculation of the share prejudiced used here is similar to the definition of prejudiced used in Bond and Lehmann (2015).

Table 3.6: Questions from the General Social Survey

Question	Question Text
SEX	What about sexual relations between two adults of the same sex–do you think it is always wrong, almost always wrong, wrong only sometimes, or not wrong at all? Asked between 1990 and 2014 GSS Mnemonic: HOMOSEX
BOOK	If some people in your community suggested that a book he wrote in favor of homosexuality should be taken out of your public library, would you favor removing this book, or not? Asked between 1990 and 2014 GSS Mnemonic: LIBHOMO
SPEAK	Suppose this admitted homosexual wanted to make a speech in your community. Should he be allowed to speak, or not? Asked between 1990 and 2014 GSS Mnemonic: SPKHOMO
COLLEGE	And what about a man who admits that he is a homosexual? Should such a person be allowed to teach in a college or university, or not? Asked between 1990 and 2014 GSS Mnemonic: COLHOMO

Note: Questions come from the pooled General Social Survey, 1990 to 2014.

prejudiced relative to other states.

To test the pre-trends, I regress the share of GSS respondents giving all prejudiced answers ($Prejudice_{s,t}$) on dummies for the year relative to the passage of an ENDA (τ_t), where (τ_0) is equal to the year of passage and is the omitted category in the regression. I also include state and year fixed effects (δ_s and δ_t).

$$Prejudice_{s,t} = \beta + \sum_{t=-5}^{-1} \tau_t + \sum_{t=1}^{5} \tau_t + \delta_t I_t + \delta_s I_s + \epsilon_{s,t} \tag{3.3}$$

Figure 3.2 shows how the pre-trends in prejudice are evolving relative to the passage of an ENDA. I find no evidence states that pass an ENDA experience faster declines in prejudice. The share of individuals expressing prejudice against the LGBT community rises slightly four and five years after the passage of an ENDA, suggesting the passage of ENDAs may result in a small backlash against the LGBT community.

Figure 3.2: Change in Prejudice Relative to Year of Passage for ENDAs

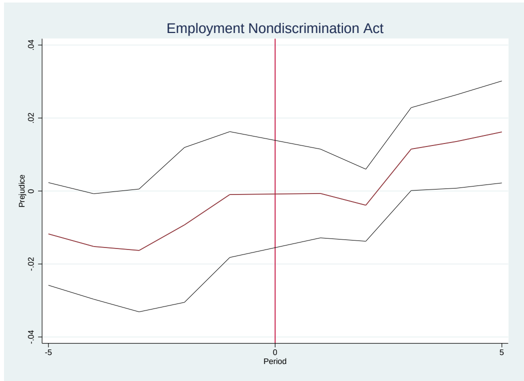

Note: Data comes from the 1990 through 2014 General Social Surveys. Prejudice is calculated as the percent of individuals in a Census division that give the most prejudice answers to questions regarding homosexuality in the General Social Survey. See Table 3.6 for the questions asked and the possible answers.

The second issue potentially biasing the results is that the estimates for log hourly wages are conditional on being in the labor market. It is possible that selection into the labor market is not random. To account for selection, I use a semi-parametric estimation strategy. I regress the indicator for being employed (E_{ist}) on all of the controls used in the baseline estimation. I also include additional controls: the number of children an individual has ($Kids_{i,s,t}$) and a dummy for whether any of those children are under the age of two ($Young_{i,s,t}$).

$$\begin{aligned} E_{ist} &= \beta_0 + \beta_1 Homo_{ist} + \beta_2(Homo_{ist} \times ENDA_{st}) + I_s\gamma_s + I_t\gamma_t \\ &+(I_s \times I_t)\gamma_{st} + (Homo_{ist} \times I_s)\theta_{H,s} + (Homo_{ist} \times I_t)\theta_{H,t} + \mathbf{X}_{ist}\delta \\ &+\alpha_1 Kids_{i,s,t} + \alpha_2 Young_{i,s,t} + \epsilon_{ist} \end{aligned} \tag{3.4}$$

Using the estimated coefficients from Equation 3.4, I estimate the predicted probability that an individual would be employed (ρ). I use a fifth order polynomial of this predicted probability as a control in the wage regression to control for selection into the labor market (Equation 3.5).

$$\begin{aligned} Y_{ist} &= \beta_0 + \beta_1 Homo_{ist} + \beta_2(Homo_{ist} \times ENDA_{st}) + I_s\gamma_s + I_t\gamma_t \\ &+(I_s \times I_t)\gamma_{st} + (Homo_{ist} \times I_s)\theta_{H,s} + (Homo_{ist} \times I_t)\theta_{H,t} + \mathbf{X}_{ist}\delta \\ &+\rho_{i,s,t} + \rho^2_{i,s,t} + \rho^3_{i,s,t} + \rho^4_{i,s,t} + \rho^5_{i,s,t} + \epsilon_{ist} \end{aligned} \tag{3.5}$$

I show in Table 3.7 that selection did not have a significant effect on the results. In column 1, I report the results for the average effect of an ENDA for gay men. The results do not appear to have been driven by selection into the labor force. The effect of an ENDA conditional on selection into employment remains similar to what it was in the baseline estimation. The effect on log annual income does not change if one controls for selection. The effect of an ENDA on log hourly wages increases slightly to 2.9%.[14] The results for lesbian women remain the same.

[14] Appendix Table A.10 reports the effect of controlling for selection on the results for legal provisions.

Table 3.7: Robustness of Results on Wages to Controlling for Selection into Employment

	Men		Women	
	Ln Annual Income	Ln Hourly Wages	Ln Annual Income	Ln Hourly Wage
	(1)	(2)	(3)	(4)
Homosexual × ENDA	0.018	0.029*	-0.019	-0.003
	(0.013)	(0.012)	(0.017)	(0.013)
Observations	6,172,273	6,172,273	5,061,060	5,061,060

*** p<0.001,** p<0.01, * p<0.05

Note: Author's calculations based on data from the 2008 through 2014 American Community Surveys 1% PUMS, the 1990 Decennial Census 5% PUMS, and the 2000 Decennial Census 5% PUMS. The sample is restricted to adults, older than 22 and younger than 65. Individuals are either married heterosexuals or cohabiting gay men or lesbian women. Individuals who have moved between states in the past year are not included in the analysis. Annual income and hourly wages are in constant 1999 dollars. All regressions are estimated using OLS and include demographic and occupation controls, state fixed effects, year fixed effects, and state-by-year fixed effects. For a full list of control variables, see Table A.1. Standard errors clustered at the state level are in parentheses.

The third potential concern is that at the same time that states were passing ENDAs, some states were also granting the right of same-sex marriage to gay men and lesbian women. It is possible that the passage of same-sex marriage increases the wages of gay men and lesbian women, though the previous literature has not found there is a wage premium associated with cohabiting for gay men (Zavodny 2007). Table 3.8 reports the results controlling for the passage of same-sex marriage. The top panel of Table 3.8 shows the results for men. After controlling for the passage of same-sex marriage, the average effect of an ENDA remains relatively unchanged at 2.8%. The results for employment and hours worked are still not statistically significant. The bottom panel of Table 3.8 reports the results for women. The passage of same-sex marriage similarly does not change the results in a significant way. There is still a negative and significant effect of ENDAs on employment and hours worked for lesbian women.[15]

[15] Appendix Table A.9 reports the effect of controlling for same-sex marriage on the results for legal provisions.

Table 3.8: Robustness of Results to Controlling for Same-Sex Marriage

Panel A.	Men Ln Annual Income	Ln Hourly Wages	Employment	Hours Worked
	(1)	(2)	(3)	(4)
Gay × ENDA	0.017	0.028*	0.005	0.190
	(0.013)	(0.012)	(0.005)	(0.194)
Observations	6,172,273	6,172,273	7,660,401	7,660,401

Panel B.	Women Ln Annual Income	Ln Hourly Wages	Employment	Hours Worked
	(1)	(2)	(3)	(4)
Lesbian × ENDA	-0.018	-0.003	-0.017*	-0.738**
	(0.016)	(0.013)	(0.008)	(0.251)
Observations	5,061,060	5,061,060	8,060,889	8,060,889

*** $p<0.001$,** $p<0.01$, * $p<0.05$

Note: Author's calculations based on data from the 2008 through 2014 American Community Surveys 1% PUMS, the 1990 Decennial Census 5% PUMS, and the 2000 Decennial Census 5% PUMS. The sample is restricted to adults, older than 22 and younger than 65. Individuals are either married heterosexuals or cohabiting gay men or lesbian women. Individuals who have moved between states in the past year are not included in the analysis. Annual income and hourly wages are in constant 1999 dollars. All regressions are estimated using OLS and include demographic and occupation controls, state fixed effects, year fixed effects, and state-by-year fixed effects. For a full list of control variables, see Table A.1. Standard errors clustered at the state level are in parentheses.

3.7 Conclusion

In this paper, I explored the effect that the passage of ENDAs at the state-level had on the labor market outcomes of gay men and lesbian women. The results showed that ENDAs led to a decline in the wage gap between gay men and married heterosexual men, but also reduced the employment of lesbian women. The fact that gay men would benefit from the passage of an ENDA and lesbian women would suffer from the passage of an ENDA can be interpreted as employers reducing the importance of sexual orientation in employment decisions. Because lesbian women may be favored over heterosexual women, the passage of an ENDA reduces the benefits they experience and increases the employment prospects of married heterosexual women. The passage of an ENDA increases the employment prospects of gay men relative to heterosexual men.

When I treat all ENDAs as identical as was done in the previous literature, I find ENDAs increased the wages of gay men by 2.7% and had no effect on employment and hours worked per week. This is similar to the effects in Baumle and Poston Jr. (2011), which found ENDAs increased wages by 2.6%. The results are also similar to the magnitudes found in the General Social Survey by Martell (2013b).

The key contribution of this paper is that I show the effects of ENDAs for gay men are concentrated in a handful states. The evidence suggests that strong and weak laws had different effects that reduced the aggregate effect observed in previous research (Baumle and Poston Jr. 2011, Klawitter and Flatt 1998, Klawitter 2011, Martell 2013b). By looking at the state-level variation in these non-discrimination laws, I can estimate the effect of strong versus weak laws. For gay men, the effects of the law depend heavily on the structure of the law. If I look at what would be considered a strong law and a weak law separately, I find that strong laws had larger wage increases and larger employment increases. Based on the estimates from Panel B of Table 3.3, a strong law (e.g. a law with compensatory damages, punitive damages, attorney's fees, a size minimum of 1, and a statute of limitations of 1 year) saw an average increase in annual wages of 12.3% and an increase in employment of 0.8%. A weak law (e.g. a law with no compensatory damages and no punitive damages, a size minimum of 15 employees, and a statute of limitations of 180 days) saw no increase in wages and a decrease in employment of 0.4%.

For lesbian women, the specific provisions of the laws have little effect on their labor market outcomes (Table 3.4). The average law had no effect on wages, but decreased employment by 1.7% and decreased hours by 0.733 hours. States with punitive damages saw larger declines in hours worked than states without punitive damages.

These results highlight the care that must be taken when crafting a law to protect a marginalized group. Simply using the same template as existing laws that protect employees against racial and gender discrimination may be detrimental. The most current version of the federal Employment Non-Discrimination Act to pass the Senate in 2013 had a statute of limitations of 180 days. This is lower than the average of state statutes of limitations by two months. The employer size minimum in the proposed ENDA is 15 employees. This is ten more employees than the average state law. The federal ENDA allows for compensatory and punitive damages and attorney's fees to be awarded. Using the results from Table 3.3 and 3.4, it is possible to calculate the effect of the federal law. For gay men, this law would have an estimated wage effect of 7.4%-7.9% (depending on whether one uses annual income from wages or hourly wages) and an estimated employment effect of -0.4%. For lesbian women, this law would decrease employment by 1.7% and decrease hours worked by 2.179 hours. These estimates suggest that policymakers must think carefully about the strength of the provisions when crafting the bill.

The effect of these employment protection laws can be very nuanced. Future research into similar laws where state laws are not constrained by federal protections, such as pay secrecy bans or transgender protections, need to take the legal differences into account when estimating the effect of the laws. Failure to do so may result in inconclusive results that are driven by the differences between weak and strong laws. My results also suggest that other policies that are implemented differently across states may experience similar differences depending on the structure of the implementation.

Bibliography

Acemoglu, Daron and Joshua D. Angrist, "Consequences of Employment Protection? The Case of the Americans with Disabilities Act," Journal of Political Economy, 2001, 109 (5), 915–957.

Adams, Glenn, Donna M. Garcia, Valerie Purdie-Vaughns, and Claude M. Steele, "The Detrimental Effects of a Suggestion of Sexism in an Instruction Situation," Journal of Experimental Social Psychology, 2006, 42 (5), 602 – 615.

Adams, Scott J., "Age Discrimination Legislation and the Employment of Older Workers," Labour Economics, 2004, 11 (2), 219 – 241.

Allegretto, Sylvia A. and Michelle M. Arthur, "An Empirical Analysis of Homosexual/Heterosexual Male Earnings Differentials: Unmarried and Unequal?," Industrial and Labor Relations Review, April 2001, 54 (3), 631.

Altonji, J.G., P. Arcidiacono, and A. Maurel, "Chapter 7 - The Analysis of Field Choice in College and Graduate School: Determinants and Wage Effects," in Stephen Machin Eric A. Hanushek and Ludger Woessmann, eds., , Vol. 5 of Handbook of the Economics of Education, Elsevier, 2016, pp. 305 – 396.

American Association of University Women, "Why So Few? Women in Science, Technology, and Mathematics," 2010. http://www.aauw.org/files/2013/02/Why-So-Few-Women-in-Science-Technology-Engineering-and-Mathematics.pdf.

____, "The Simple Truth about the Gender Pay Gap," 2015. http://www.aauw.org/files/2015/02/The-Simple-Truth-Spring-2015.

Antecol, Heather and Michael D. Steinberger, "Labor Supply Differences Between Married Heterosexual Women and Partnered Lesbians: A Semi-Parametric Decomposition Approach," Economic Inquiry, 2013, 51 (1), 783–805.

____**, Anneke Jong, and Michael Steinberger**, "The Sexual Orientation Wage Gap: The Role of Occupational Sorting and Human Capital," Industrial and Labor Relations Review, 2008, 61 (4), 518–543.

Arcidiacono, Peter, "Ability Sorting and the Returns to College Major," Journal of Econometrics, 2004, 121 (12), 343 – 375.

Badgett, M V Lee, "The Wage Effects of Sexual Orientation Discrimination," Industrial and Labor Relations Review, June 1995, 48 (4), 726–739.

___, **Holning Lau, Brad Sears, and Deborah Ho**, "Bias in the Workplace : Consistent Evidence of Sexual Orientation and Gender Identity Discrimination," UC Los Angeles: The Williams Institute, 2007, (June), 1–8.

Bailey, John, Michael Wallace, and Bradley Wright, "Are Gay Men and Lesbians Discriminated Against When Applying for Jobs? A Four-City, Internet-Based Field Experiment," Journal of Homosexuality, 2013, 60 (6), 873–894.

Baumle, Amanda K. and Dudley L. Poston Jr., "The Economic Cost of Homosexuality: Multilevel Analyses," Social Forces, 2011, 89 (3), 1005–1031.

Beaudin, Laura, "Marriage Equality and Interstate Migration," Applied Economics, 2017, 49 (30), 2956–2973.

Becker, Gary, The Economics of Discrimination, Vol. 2, The University of Chicago Press, 1971.

Beegle, Kathleen and Wendy A. Stock, "The Labor Market Effects of Disability Discrimination Laws," Journal of Human Resources, 2003, 38 (4), 806–859.

Berg, Nathan and Donald Lien, "Measureing the Effect of Sexual Orientation on Income: Evidence of Discrimination?," Contemporary Economic Policy, 2002, 20 (4), 394–414.

Bettinger, Eric P. and Bridget Terry Long, "Do Faculty Serve as Role Models? The Impact of Instructor Gender on Female Students," The American Economic Review, 2005, 95 (2), pp. 152–157.

Black, Dan, "Discrimination in an Equilibrium Search Model," Journal of Labor Economics, 1995, 13 (2), 309–334.

Black, Dan A., Hoda R. Makar, Seth G. Sanders, and Lowell J. Taylor, "The Earnings Effects of Sexual Orientation," Industrial and Labor Relations Review, April 2003, 56 (3), 449–469.

Blandford, John M., "The Nexus of Sexual Orientation and Gender in the Determination of Earnings," Industrial and Labor Relations Review, 2003, 56 (4), 622–643.

Bloch, Farrel, Antidiscrimination Law and Minority Employment: Recruitment Practices and Regulatory Constraints, Chicago: University of Chicago Press, 1994.

Bond, Timothy N. and Jee-Yeon K. Lehmann, "Prejudice and Racial Matches in Employment," Working Paper, 2015.

Bowlus, Audra J. and Zvi Eckstein, "Discrimination and Skill Differences in an Equilibrium Search Model," International Economic Review, 2002, 43 (4), 1309–1345.

Burn, Ian, "Pride, Prejudice, and Wages: An Empirical Assessment of Models of Taste-Based Discrimination for Gay Men," Working Paper, February 2017.

___ **and Osborne Jackson**, "Valuable Vows: An Examination of the Marriage Premium Using Same-Sex Marriage Legalization," Working Paper, 2014.

Butcher, Kristin F., Patrick J. McEwan, and Akila Weerapana, "The Effects of an Anti-grade-Inflation Policy at Wellesley College," Journal of Economic Perspectives, 2014, 28 (3), 189–204.

Carpenter, Christopher S., "New Evidence on Gay and Lesbian Household Incomes," Contemporary Economic Policy, 2004, 22 (1), 78–94.

___ , "Self-Reported Sexual Orientation and Earnings: Evidence from California," Industrial and Labor Relations Review, 2005, 58 (2), 258–273.

Carrell, Scott E., Marianne E. Page, and James E. West, "Sex and Science: How Professor Gender Perpetuates the Gender Gap," The Quarterly Journal of Economics, 2010, 125 (3), 1101–1144.

Ceci, Stephen J., Donna K. Ginther, Shulamit Kahn, and Wendy M. Williams, "Women in Academic Science: A Changing Landscape," Psychological Science in the Public Interest, December 2014, 15, 75–141.

___ **, Wendy M. Williams, and Susan M. Banett**, "Womens Underrepresentation in Science: Sociocultural and Biological Considerations," Psychological Bulletin, 2009, 135 (2), 218–261.

Charles, Kerwin and Jonathan Guryan, "Prejudice and Wages: An Empirical Assessment of Beckers The Economics of Discrimination," Journal of Political Economy, 2008, 116 (5), 773–808.

Chen, Xianglei and Matthew Soldner, "STEM Attrition: College Students' Paths Into and Out of STEM Fields ," November 2013, NCES 2014-001. http://files.eric.ed.gov/fulltext/ED541511.pdf.

Collins, William J., "The Labor Market Impact of State-Level Anti-Discrimination Laws, 1940-1960," Industrial and Labor Relations Review, January 2003, 56 (2), 244–272.

Colvin, Roddrick A. and Norma M. Riccucci, "Employment Nondiscrimation Policies: Assessing Implementation and Measuring Effectiveness," International Journal of Public Administration, 2002, 25 (1), 95–108.

Council of Economic Advisers, "Engage to Excel: Producing One Million Additional College Graduates with Degrees in Science, Technology, Engineering, and Mathematics," February 2012. http://files.eric.ed.gov/fulltext/ED541511.pdf.

Council of Economic Advisors, "Gender Pay Gap: Recent Trends and Explanations," April 2015. https://www.whitehouse.gov/sites/default/files/docs/equal-pay-issue-brief-final.

Cushing-Daniels, Brendan and Tsz-Ying Yeung, "Wage Penalties and Sexual Orientation: an Update Using the General Social Survey," Contemporary Economic Policy, April 2009, 27 (2), 164–175.

Daneshvary, Nasser, C. Jeffrey Wassoups, and Bradley S. Wimmer, "Previous Marriage and the Lesbian Wage Premium," Industrial Relations: A Journal of Economy and Society, 2009, 48 (3), 432–453.

DeLeire, Thomas, "The Wage and Employment Effects of the Americans with Disabilities Act," The Journal of Human Resources, 2000, 35 (4), 693–715.

Donohue, John J. and James Heckman, "Continuous Versus Episodic Change: The Impact of Civil Rights Policy on the Economic Status of Blacks," Journal of Economic Literature, December 1991, 29 (4), 1603–1643.

Drydakis, Nick, "Sexual Orientation Discrimination in the United Kingdoms Labour Market: A Field Experiment," Human Relations, 2015, 68 (11), 1769–1796.

Eccles, Jacquelynne S., Janis E. Jacobs, and Rena D. Harold, "Gender Role Stereotypes, Expectancy Effects, and Parents' Socialization of Gender Differences," Journal of Social Issues, 1990, 46 (2), 183–201.

Elmslie, Bruce and Edinaldo Tebaldi, "Sexual Orientation and Labor Market Discrimination," Journal of Labor Research, July 2007, 28 (3), 436–453.

___ **and** ___ , "The Wage Gap against Gay Men: The Leveling of the Playing Field," Kyklos, 2014, 67 (3), 330–345.

Flabbi, Luca and Mauricio Tejada, Gender Gaps in Wages and Employment: The Role of Employers' Prejudice, Oxford University Press, 5 2015.

Frank, Jeff, "Gay Glass Ceilings," Economica, 2006, 73 (291), 485–508.

Gates, Gary J. and Frank Newton, "LGBT Percentage Highest in D.C., Lowest in North Dakota," February 2013. http://www.gallup.com/poll/160517/lgbt-percentage-highest-lowest-north-dakota.

Gemici, Ahu and Matthew Wiswall, "Evolution of Gender Differences in Post-Secondary Human Capital Investments: College Majors," International Economic Review, 2014, 55 (1), 23–56.

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

Government Accountability Office, "Sexual-Orientation-Based Employment Discrimination: States' Experience With Statutory Prohibitions," October 1997. GAO-98-7R.

___ , "Update on State Statutes and Administrative Complaint Data on Employment Discrimination Based on Sexual Orientation and Gender Identity," July 2013. GAO-13-700R.

Gunderson, ElizabethA., Gerardo Ramirez, SusanC. Levine, and SianL. Beilock, "The Role of Parents and Teachers in the Development of Gender-Related Math Attitudes," Sex Roles, 2012, 66 (3-4), 153–166.

Heckman, James J., "Detecting Discrimination," Journal of Economic Perspectives, 1998, 12 (2), 101–116.

Hellerstein, Judith K., Melissa McInerney, and David Neumark, "Neighbors and Coworkers: The Importance of Residential Labor Market Networks," Journal of Labor Economics, 2011, 29 (4), 659–695.

Human Rights Campaign, "Municipality Equality Index 2012," 2012. www.hrc.org/mei.

——, "Marriage Equality and Other Relationship Recognition Laws," December 2013. www.hrc.org/state_maps.

Jepsen, Christopher and Lisa K. Jepsen, "Labor-Market Specialization within Same-Sex and Difference-Sex Couples," Industrial Relations: A Journal of Economy and Society, 2015, 54 (1), 109–130.

Jepsen, Lisa K., "Comparing the Earnings of Cohabiting Lesbians, Cohabiting Heterosexual Women, and Married Women: Evidence from the 2000 Census," Industrial Relations: A Journal of Economy and Society, 2007, 46 (4), 699–727.

Jolls, Christine and J.J. Prescott, "Disaggregating Employment Protection: The Case of Disability Discrimination," September 2004, (10740).

Klawitter, Marieka, "Multilevel Analysis of the Effects of Antidiscrimination Policies on Earnings by Sexual Orientation," Journal of Policy Analysis and Management, 2011, 30 (2), 334–358.

Klawitter, Marieka M., "Meta-Analysis of the Effects of Sexual Orientation on Earnings," Industrial Relations, 2015, 54 (1), 4–32.

—— **and Victor Flatt**, "The Effects of State and Local Antidiscrimination Policies on Earnings for Gays and Lesbians," Journal of Policy Analysis and Management, 1998, 17 (4), 658–686.

Kruse, Douglas and Lisa Schur, "Employment of People with Disabilities Following the ADA," Industrial Relations: A Journal of Economy and Society, 2003, 42 (1), 31–66.

Lahey, Joanna, "State Age Protection Laws and the Age Discrimination in Employment Act," Journal of Law and Economics, 2008, 51 (3), 433–460.

Landes, William M., "The Economics of Fair Employment Laws," Journal of Political Economy, July 1968, 76 (4), 507–552.

Lang, Kevin and Michael Manove, "Education and Labor Market Discrimination," American Economic Review, 2011, 101 (4), 1467–1496.

Leslie, Sarah-Jane, Andrei Cimpian, Meredith Meyer, and Edward Freeland, "Expectations of Brilliance Underlie Gender Distributions Across Academic Disciplines," Science, 2015, (347), 262–265.

Long, Mark C., Dan Goldhaber, and Nick Huntington-Klein, "Do Completed College Majors Respond to Changes in Wages?," Economics of Education Review, 2015, 49, 1 – 14.

Martell, E. Michael, "Differences Do Not Matter: Exploring the Wage Gap for Same-Sex Behaving Men," Eastern Economic Journal, 2013, 39 (1), 45–71.

Martell, Michael E., "Do ENDAs End Discrimination for Behaviorally Gay Men?," Journal of Labor Research, November 2013, 34, 147–169.

Mishel, Emma, "Discrimination against Queer Women in the U.S. Workforce: A Resume Audit Study," Socius: Sociological Research for a Dynamic World, 2016, 2.

Montmarquette, Claude, Kathy Cannings, and Sophie Mahseredjian, "How Do Young People Choose College Majors?," Economics of Education Review, 2002, 21 (6), 543 – 556.

Moss-Racusin, Corinne A., John F. Dovidio, Victoria L. Brescoll, Mark J. Graham, and Jo Handelsman, "Science Facultys Subtle Gender Biases Favor Male Students," Proceedings of the National Academy of Sciences, 2012, 109 (41), 16474–16479.

Neumark, David, "Employers' Discriminatory Behavior and the Estimation of Wage Discrimination," Journal of Human Resources, 1988, 23 (3), 279–295.

———, "Detecting Discrimination in Audit and Correspondence Studies," Journal of Human Resources, 2012, 47 (4), 1128–1157.

——— **and Patrick Button**, "Did Age Discrimination Protections Help Older Workers Weather the Great Recession?," Journal of Policy Analysis and Management, 2014, 33 (3), 566–601.

——— **and Wendy A. Stock**, "Age Discrimination Laws and Labor Market Efficiency," Journal of Political Economy, 1999, 107 (5), 1081–1125.

——— **and** ———, "The Labor Market Effects of Sex and Race Discrimination Laws," Economic Inquiry, July 2006, 44 (4), 385–419.

North, Anna, "Does The Wage Gap Really Exist?," Jezebel, April 21, 2010. http://www.jezebel.com/5521000/does-the-wage-gap-really-exist.

Nosek, B.A., M.R. Banaji, and G Greenwald A, "Math = Male, Me = Female, Therefore Math Me," Journal of Personality and Social Psychology, 2002, 83 (1), 101–115.

O'Brien, Sara Ashley, "78 Cents on the Dollar: The Facts About the Gender Wage Gap," CNN Money, April 14, 2015. http://www.money.cnn.com/2015/04/13/news/economy/equal-pay-day-2015.

O'Connell, Martin and Daphne Loftquist, "Counting Same-sex Couples: Official Estimates and Unofficial Guesses ," 2009. Paper presented at the Population Association of America Meeting, Detroit.

___ **and Gretchen Golding**, "The Use of First Names to Evaluate Reports of Gender and Its Efffect on the Distribution of Married and Unmarried Couple Households," 2006. Paper presented at the Population Association of America Meeting, Los Angeles.

Ost, Ben, "The Role of Peers and Grades in Determining Major Persistence in the Sciences," Economics of Education Review, 2010, 29 (6), 923–934.

Pew Research Center, "A Survey of LGBT Americans: Attitudes, Experiences and Values in Changing Times," June 2013.

Reuben, Ernesto, Matthew Wiswall, and Basit Zafar, "Preferences and Biases in Educational Choices and Labour Market Expectations: Shrinking the Black Box of Gender," The Economic Journal, 2017.

Rosen, Asa, "Search, Bargaining, and Employer Discrimination," Journal of Labor Economics, 2003, 21 (4), 807–829.

Ruggles, Steven, J. Trent Alexander, Katie Gendadek, Ronald Goeken, Matthew B. Schroeder, and Matthew Sobek, "Integrated Public Use Microdata Series: Version 5.0 [Machine-readable database]," 2010. Minneapolis: University of Minnesota.

Sabia, Joseph J., "Sexual Orientation and Wages in Young Adulthood: New Evidence from Add Health," Industrial and Labor Relations Review, 2014, 67 (1), 239–267.

Scott, George, "Testimony before the subcommittee on early childhood, elementary, and secondary education, committee on education and the workforce, house of representatives," April 2013. http://www.gao.gov/assets/660/653661.pdf.

Sears, Brad and Christy Mallory, "Documented Evidence of Employment Discrimination and Its Effects on LGBT People ," The Williams Institute, 2011.

___ **, Nan Hunter, and Christy Mallory**, "Analysis of Scope and Enforcement of State Laws and Executive Orders Prohibiting Employment Discrimination Against LGBT People," The Williams Institute, September 2009.

Singh, Kusum, Katherine R. Allen, Rebecca Scheckler, and Lisa Darlington, "Women in Computer-Related Majors:A Critical Synthesis of Research and Theory From 1994 to 2005," Review of Educational Research, 2007, 77 (4), 500–533.

Sjoquist, David L. and John V. Winters, “The Effect of Georgia’s HOPE Scholarship on College Major: A Focus on STEM,” IZA Journal of Labor Economics, 2015, 4 (1), 15.

Spencer, Steven J, Claude M. Steele, , and Diane M Quinn, “Stereotype Threat and Womens Math Performance,” Journal of Experimental Social Psychology, 1999, 35, 4 – 28.

Steele, Claude M., “A Threat in the Air: How Stereotypes Shape Intellectual Identity and Performance,” American Psychologist, 1997, 52 (6), 613 – 629.

Stinebrickner, Ralph and Todd R. Stinebrickner, “A Major in Science? Initial Beliefs and Final Outcomes for College Major and Dropout,” The Review of Economic Studies, 2014, 81 (1), 426.

Tilcsik, Andras, “Pride and Prejudice: Employment Discrimination against Openly Gay Men in the United States,” American Journal of Sociology, September 2011, 117 (2), 586–626.

Turner, Sarah E. and William G. Bowen, “Choice of Major: The Changing (Unchanging) Gender Gap,” Industrial and Labor Relations Review, 1999, 52 (2), 289–313.

Ueno, Koji, Preeti Vaghela, and Lacey J. Ritter, “Sexual Orientation, Internal Migration, and Mental Health during the Transition to Adulthood,” Journal of Health and Social Behavior, 2014, 55 (4), 461–481.

Williams, Joan C., Katherine W. Phillips, and Erika V. Hall, “Double Jeopardy: Gender Bias Against Women of Color in Sciences,” UC Hastings College of Law WorkLifeLaw Center, 2014. http://www.uchastings.edu/news/articles/2015/01/double-jeopardy-report.pdf.

Wiswall, Matthew and Basit Zafar, “Determinants of College Major Choice: Identification using an Information Experiment,” The Review of Economic Studies, 2015, 82 (2), 791.

___ and ___ , “How Do College Students Respond to Public Information about Earnings?,” Journal of Human Capital, 2015, 9 (2), 117–169.

Zafar, Basit, “College Major Choice and the Gender Gap,” Journal of Human Resources, 2013, 48 (3), 545–595.

Zavodny, Madeline, “Is There a “Marriage Premium” for Gay Men?,” Review of Economics of the Household, October 2007, 6 (4), 369–389.

Appendix A

Appendix Tables and Figures

A.1 Chapter 1

Table A.1: List of Variables

Variable Definition
Dependent Variable:
Annual earnings
Natural logarithm of hourly earnings (= total annual salary earnings divided by total number of hours worked per year) in previous year, in constant 1999 USD
Control Variables:
Sexual Orientation (=1 if Homosexual, =0 if Heterosexual)
Experience (Potential, =Age - Schooling - 5)
Experience Squared
Black (=1 if True, =0 if False)
Other Race (=1 if True, =0 if False)
Years of Schooling
State
Year
Gay Rights Movement Variables
Employment Non-Discrimination Act Protections for Homosexuals in State (=1 if True, =0 if False)
Legal Recognition of Same-Sex Marriages in State (=1 if True, =0 if False)

Note: Sources of all variables are the 2008 through 2014 American Community Survey 1-Year Sample, the 1990 to 2012 General Social Surveys, and Sears et al. (2009).

Table A.2: Testing Higher Percentiles of the Prejudice Distribution

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1^{st} Percentile	-0.163							
	(0.148)							
5^{th} Percentile		-0.209						
		(0.178)						
25^{th} Percentile			-0.061					
			(0.046)					
50^{th} Percentile				-0.067				
				(0.047)				
75^{th} Percentile					-0.063*			
					(0.034)			
95^{th} Percentile						-0.058		
						(0.063)		
99^{th} Percentile							0.069	
							(0.046)	
Mean								-0.092*
								(0.050)
Share Gay	3.0093**	2.987**	2.924**	2.671*	2.539**	2.451**	2.705*	2.924**
	(1.488)	(1.458)	(1.432)	(1.356)	(1.251)	(1.304)	(1.491)	(1.431)
States	48	48	48	48	48	48	48	48
Obs.	6,268,265	6,268,265	6,268,265	6,268,265	6,268,265	6,268,265	6,268,265	6,268,265

*** p<0.01, ** p<0.05, * p<0.1

Note: The dependent variable is log hourly wages in constant 1999 dollars. The sample has been restricted to only include gay or married cohabiting men in the labor force. On average gay men experience a wage penalty of 10.4% relative to married heterosexual men. Data on wages come from the 1990 Decennial Census 5% PUMS, the 2000 Decennial Census 5% PUMS, and the 2008 through 2014 American Community Surveys. Data on prejudice come from the 1990 through 2014 waves of the GSS. Standard errors are clustered at the state level and are reported in parentheses. Census sample weights are used to weight the observations. Three states have been dropped from the sample because they have too few respondents in the General Social Survey.

A.2 Chapter 2

Figure A.1: Example: Probability of Selecting Major

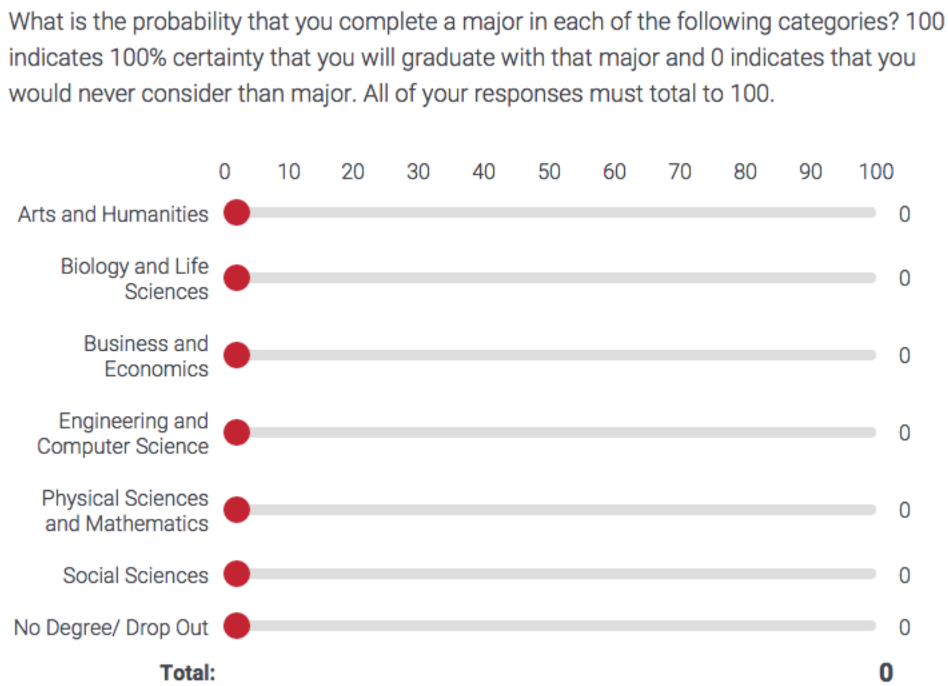

Figure A.2: Example: Stage 1 Questions

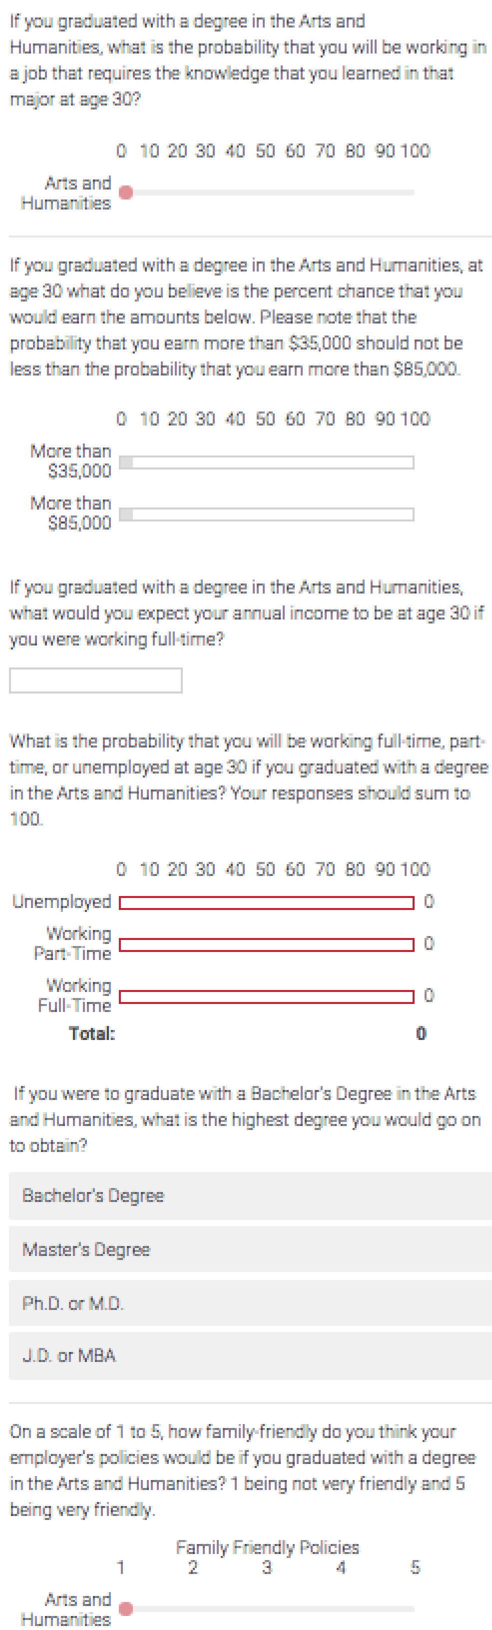
If you graduated with a degree in the Arts and Humanities, what is the probability that you will be working in a job that requires the knowledge that you learned in that major at age 30?

0 10 20 30 40 50 60 70 80 90 100

Arts and Humanities

If you graduated with a degree in the Arts and Humanities, at age 30 what do you believe is the percent chance that you would earn the amounts below. Please note that the probability that you earn more than $35,000 should not be less than the probability that you earn more than $85,000.

0 10 20 30 40 50 60 70 80 90 100

More than $35,000

More than $85,000

If you graduated with a degree in the Arts and Humanities, what would you expect your annual income to be at age 30 if you were working full-time?

What is the probability that you will be working full-time, part-time, or unemployed at age 30 if you graduated with a degree in the Arts and Humanities? Your responses should sum to 100.

0 10 20 30 40 50 60 70 80 90 100

Unemployed 0

Working Part-Time 0

Working Full-Time 0

Total: **0**

If you were to graduate with a Bachelor's Degree in the Arts and Humanities, what is the highest degree you would go on to obtain?

Bachelor's Degree

Master's Degree

Ph.D. or M.D.

J.D. or MBA

On a scale of 1 to 5, how family-friendly do you think your employer's policies would be if you graduated with a degree in the Arts and Humanities? 1 being not very friendly and 5 being very friendly.

Family Friendly Policies

1 2 3 4 5

Arts and Humanities

Figure A.3: Example: Stage 3 Directions

The following questions will ask you about your expectations and beliefs for when you are 30 years old. For this next section, consider the situation where you graduate with a Bachelor's degree in the Arts or the Humanities.

Below are the answers to the questions that were just asked.

At age 30, people with a Bachelor's Degree in the Arts and Humanities earn $48,354 a year when they work full time.

At age 30, a woman with a Bachelor's Degree in the Arts and Humanities earns 96 cents for every dollar a man with a Bachelor's Degree in the Arts and Humanities earns.

In the occupations that individuals with a Bachelor's Degree in the Arts and Humanities work, 26% of people believe it is better for a man to work and a woman to stay at home.

	Average Annual Earnings	Female to Male Earnings Ratio	Percent in occupation who believe that it is better for a man to work and a woman to stay at home
Arts and Humanities	$48,354	96%	26%

Figure A.4: Example: Risk Elicitation

	Outcome 1	Outcome 2
Lottery A	\$5.00	\$5.00
Lottery B	\$4.00	\$6.05
Lottery C	\$2.80	\$7.00

Consider the three lotteries presented above. In each of the three lotteries, you have a 50% chance of winning Outcome 1 or Outcome 2 (i.e. each outcome is equally likely). Please pick which of the lotteries that you prefer (Lottery A, Lottery B, or Lottery C). At the end of the experiment, you will play which every lottery you selected for that line (either A, B, or C). You will be awarded the amount of money that results from the lottery.

Lottery A

Lottery B

Lottery C

A.3 Chapter 3

Table A.3: Employment Non-Discrimination Acts

State	Effective Date	Law	Enforcement Agency
California	1992	CAL. GOV. CODE 12940	California Department of Fair Employment and Housing
Colorado	2007	COLO. REV. STAT. 24-34-401(7.5)	Colorado Civil Rights Commission
Connecticut	1991	CONN. GEN. STAT. 46a-81a	Connecticut Commission on Human Rights and Opportunities
Delaware	2009	DEL. CODE ANN. tit. 19 710	Delaware Department of Labor
District of Columbia	1997	D.C. Code 1-2512	Washington, D.C. Commission on Human Rights
Hawaii	1991	HAW. REV. STAT. 378-1	Hawaii Civil Rights Commission
Illinois	2006	775 ILCS 5/1-102(O-1)	Illinois Department of Human Rights
Iowa	2007	IOWA CODE 216.2(14)	Iowa Civil Rights Commission
Maine	2005	ME. REV. STAT. ANN. tit. 5 4553(9-C)	Maine Human Rights Commission
Maryland	2001	MD Code, State Government, 20-606	Maryland Human Rights Commission
Massachusetts	1989	MASS. GEN. LAWS ch. 151B, 3(6)	Massachusetts Commission Against Discrimination
Minnesota	1993	MINN. STAT. 363A.03 subd. 44	Minnesota Department of Human Rights
Nevada	1999	NEV. REV. STAT. 613.310(6)	Nevada Equal Rights Commission
New Hampshire	1998	N.H. REV. STAT. Ann. 354-A:2(XIV-c)	The New Hampshire State Commission on Human Rights
New Jersey	1992	N.J. STAT. 10:5-5(hh)	The New Jersey Division of Civil Rights
New Mexico	2003	N.M. STAT. 28-1-2(P)	The New Mexico Human Rights Division
New York	2003	N.Y. EXEC. LAW 292(27)	New York Division of Human Rights
Oregon	2008	OR. REV. STAT. 174.100(6)	Oregon Bureau of Labor and Industry
Rhode Island	1995	R.I. GEN. LAWS 28-5-6(7)	The Rhode Island Commission for Human Rights
Vermont	1991	1 VT. STAT. ANN. 143	The Vermont Human Rights Commission
Washington	2006	WASH. REV. CODE 49.60.040(15)	The Washington State Human Rights Commission
Wisconsin	1982	WIS. STAT. 111.32(13m)	The Wisconsin Department of Workforce Development

Table A.4: Coverage Provisions

State	Effective Date	Domestic Workers	Family Members	Perceived Orientation	Employer Size	Complaint Period
California	2001	No	No	Yes	5	300
California	1992	No	No	No	5	300
Colorado	2007	No	No	Yes	1	180
Connecticut	1991	No	Yes	Yes	3	180
Delaware	2009	Yes	Yes	No	4	120
District of Columbia	1977	No	No	Yes	1	365
Hawaii	1991	Yes	Yes	Yes	1	180
Illinois	2006	No	Yes	Yes	15	180
Iowa	2007	No	Yes	Yes	4	180
Maine	2005	Yes	No	Yes	1	180
Maryland	2001	Yes	Yes	Yes	15	180
Massachusetts	2002	No	No	Yes	6	300
Massachusetts	1989	No	No	Yes	6	180
Minnesota	1993	No	No	Yes	1	180
Nevada	1999	Yes	Yes	Yes	15	180
New Hampshire	1998	No	No	Yes	6	180
New Jersey	1992	No	Yes	Yes	1	180
New Mexico	2003	Yes	Yes	Yes	15	300
New York	2003	No	No	Yes	4	365
Oregon	2008	No	No	Yes	1	365
Rhode Island	1995	No	No	Yes	4	365
Vermont	1991	Yes	Yes	No	1	365
Washington	2006	No	No	No	8	180
Wisconsin	1982	Yes	No	Yes	1	300

Table A.5: Compensatory Damages

State	Effective Date	Back Pay	Compensatory Damages	Administrative Compensatory Damages	Damages Cap	Cap Amount
California	1992	Yes	Yes	No	Yes	150,000
Colorado	2007	Yes	No	No		
Connecticut	2002	Yes	Yes	No	No	
Connecticut	1991	Yes	No	No		
Delaware	2009	Yes	Yes	Yes	Yes	Graduated Scale
District of Columbia	1977	Yes	Yes	Yes	No	
Hawaii	1991	Yes	Yes	Yes	No	
Illinois	2006	Yes	Yes	Yes	No	
Iowa	2007	Yes	Yes	Yes	No	
Maine	2005	Yes	Yes	No	Yes	Graduated Scale
Maryland	2001	Yes	Yes	Yes	Yes	Graduated Scale
Massachusetts	1989	Yes	Yes	No	Yes	3x the damage
Minnesota	2006	Yes	Yes	Yes	Yes	3x the damage
Nevada	1999	Yes	No	No		
New Hampshire	1998	Yes	Yes	Yes	No	
New Jersey	1992	Yes	Yes	Yes	No	
New Mexico	2003	Yes	Yes	Yes	No	
New York	2003	Yes	Yes	Yes	No	
Oregon	2008	Yes	Yes	Yes	No	
Rhode Island	1995	Yes	Yes	Yes	No	
Vermont	1991	Yes	No	No		
Washington	2006	Yes	Yes	Yes	Yes	20,000
Wisconsin	1982	Yes	Yes	No	Yes	Graduated Scale

Table A.6: Punitive Damages

State	Effective Date	Punitive Damages	Administrative Punitive Damages	Damages Cap	Cap Amount
California	1992	No	No		
Colorado	2007	No	No		
Connecticut	2002	Yes	No	No	
Connecticut	1991	No	No		
Delaware	2009	Yes	Yes	Yes	Graduated Scale
District of Columbia	1977	Yes	Yes	Yes	Graduated Scale
Hawaii	1991	Yes	Yes	No	
Illinois	2006	No	No		
Iowa	2007	No	No	No	
Maine	2005	Yes	Yes	Yes	Graduated Scale
Maryland	2001	Yes	No	Yes	Graduated Scale
Massachusetts	1989	Yes	No	Yes	3x the damage
Minnesota	2006	Yes	Yes	Yes	25,000
Minnesota	1993	Yes	Yes	Yes	8500
Nevada	1999	No	No		
New Hampshire	1998	Yes	Yes	Yes	Graduated Scale
New Jersey	1992	Yes	No	No	
New Mexico	2003	No	No		
New York	2003	No	No		
Oregon	2008	Yes	No	No	
Rhode Island	1995	Yes	No	No	
Vermont	1991	Yes	No	No	
Washington	2006	No	No		
Wisconsin	1982	Yes	No	Yes	Graduated Scale

Table A.7: Process Provisions

State	Year	Requirement to Exhaust Administrative Remedies Before Civil Suit	Recoup Attorney's Fees	Recoup Attorney's Fees for Administrative Processes	Enforcement Agency Can Act On Its Own
California	1992	Yes	Yes	No	Yes
Colorado	2007	Yes	No	No	No
Connecticut	1991	Yes	Yes	No	Yes
Delaware	2009	Yes	Yes	Yes	Yes
District of Columbia	1977	Yes	Yes	Yes	Yes
Hawaii	1991	Yes	Yes	Yes	Yes
Illinois	2006	Yes	Yes	Yes	Yes
Iowa	2007	Yes	Yes	Yes	No
Maine	2005	Yes	Yes	Yes	Yes
Maryland	2001	Yes	Yes	Yes	Yes
Massachusetts	1989	Yes	Yes	Yes	No
Minnesota	1993	No	Yes	Yes	Yes
Nevada	1999	Yes	No	No	No
New Hampshire	1998	Yes	No	No	No
New Jersey	1992	No	Yes	Yes	Yes
New Mexico	2003	Yes	Yes	Yes	Yes
New York	2003	No	No	No	Yes
Oregon	2008	No	Yes	No	Yes
Rhode Island	1995	Yes	Yes	Yes	Yes
Vermont	1991	No	Yes	No	Yes
Washington	2006	No	Yes	No	Yes
Wisconsin	1982	Yes	Yes	No	Yes

Table A.8: Effect of ENDA Provisions on Earnings by Full-Time and Part-Time Employment Status

	Men			
	Full-Time		Part-Time	
Panel A.	Ln Annual Income	Ln Hourly Wages	Log Annual Income	Log Hourly Wage
	(1)	(2)	(3)	(4)
Gay × ENDA	0.046	0.025	-0.037	0.129
	(0.029)	(0.027)	(0.190)	(0.235)
Gay × ENDA × Employer Size Minimum	-0.005	-0.005	-0.008	0.007
	(0.003)	(0.003)	(0.013)	(0.016)
Gay × ENDA × Complaint Period	-0.004	-0.004	-0.025	-0.041
	(0.003)	(0.003)	(0.021)	(0.020)
Gay × ENDA × Compensatory Damages	0.141***	0.158***	-0.097	-0.572*
	(0.021)	(0.024)	(0.356)	(0.226)
Gay × ENDA × Punitive Damages	-0.100***	-0.104***	0.259	0.163
	(0.015)	(0.014)	(0.148)	(0.157)
Gay × ENDA × Attorney's Fees	-0.064	-0.054	0.204	0.729*
	(0.035)	(0.033)	(0.44)	(0.282)
Observations	5,985,370	5,985,370	186,903	186,903
	Women			
Panel B.	Ln Annual Income	Ln Hourly Wages	Employment	Hours Worked
	(1)	(2)	(3)	(4)
Lesbian × ENDA	-0.030	-0.0030	-0.418**	-0.102
	(0.044)	(0.033)	(0.129)	(0.116)
Lesbian × ENDA × Employer Size Minimum	0.000	0.001	0.045***	0.040***
	(0.004)	(0.003)	(0.006)	(0.006)
Lesbian × ENDA × Complaint Period	0.003	0.002	0.056***	0.044***
	(0.005)	(0.004)	(0.011)	(0.009)
Lesbian × ENDA × Compensatory Damages	-0.052	-0.075*	-0.238	-0.597
	(0.040)	(0.035)	(0.362)	(0.401)
Lesbian × ENDA × Punitive Damages	-0.013	-0.013	-0.114	0.015
	(0.028)	(0.023)	(0.063)	(0.065)
Lesbian × ENDA × Attorney's Fees	0.049	0.047	0.011	0.179
	(0.034)	(0.033)	(0.414)	(0.475)
Observations	4,213,547	4,213,547	847,513	847,513

*** $p<0.001$,** $p<0.01$, * $p<0.05$

Note: Author's calculations based on data from the 2008 through 2014 American Community Surveys 1% PUMS, the 1990 Decennial Census 5% PUMS, and the 2000 Decennial Census 5% PUMS. The sample is restricted to adults, older than 22 and younger than 65. Individuals are either married heterosexuals or cohabitating gay men or lesbian women. Individuals who have moved between states in the past year are not included in the analysis. Annual earnings and hourly wages are in constant 1999 dollars. Full-time employment is defined as working 30 hours per week or more. Part-time employment is defined as working less than 30 hours per week. All regressions are estimated using OLS and include demographic and occupation controls, state fixed effects, year fixed effects, and state-by-year fixed effects. For a full list of control variables, see Table A.1. Standard errors clustered at the state level are in parentheses.

Table A.9: Robustness of Results to Controlling for Same-Sex Marriage

			Men	
Panel A.	Ln Annual Income	Ln Hourly Wages	Employment	Hours Worked
	(1)	(2)	(3)	(4)
Gay × ENDA	0.038	0.032	-0.008	0.013
	(0.027)	(0.022)	(0.010)	(0.328)
Gay × ENDA × Employer Size Minimum	-0.003	-0.005	0.001	-0.011
	(0.003)	(0.003)	(0.001)	(0.024)
Gay × ENDA × Complaint Period	-0.001	-0.006	0.003	0.091
	(0.004)	(0.003)	(0.001)	(0.047)
Gay × ENDA × Compensatory Damages	0.086***	0.122***	0.006	-0.518
	(0.022)	(0.023)	(0.010)	(0.713)
Gay × ENDA × Punitive Damages	-0.042*	-0.085***	0.001	0.174
	(0.019)	(0.017)	(0.012)	(0.436)
Gay × ENDA × Attorney's Fees	-0.071	-0.018	-0.023*	-0.109
	(0.036)	(0.031)	(0.011)	(0.895)
Observations	6,172,273	6,172,273	7,660,401	7,660,401
			Women	
Panel B.	Ln Annual Income	Ln Hourly Wages	Employment	Hours Worked
	(1)	(2)	(3)	(4)
Lesbian × ENDA	-0.109*	-0.022	-0.013	-1.381***
	(0.046)	(0.034)	(0.012)	(0.317)
Lesbian × ENDA × Employer Size Minimum	0.006	0.003	0.002	0.020
	(0.004)	(0.003)	(0.001)	(0.035)
Lesbian × ENDA × Complaint Period	0.006	0.005	0.000	-0.013
	(0.005)	(0.004)	(0.002)	(0.051)
Lesbian × ENDA × Compensatory Damages	-0.030	-0.095	-0.019	0.366
	(0.059)	(0.049)	(0.022)	(0.432)
Lesbian × ENDA × Punitive Damages	-0.036	-0.007	0.004	-0.813*
	(0.026)	(0.022)	(0.015)	(0.312)
Lesbian × ENDA × Attorney's Fees	0.050	0.047	-0.003	0.546
	(0.064)	(0.054)	(0.023)	(0.513)
Observations	5,061,060	5,061,060	8,060,889	8,060,889

*** $p<0.001$,** $p<0.01$, * $p<0.05$

Note: Author's calculations based on data from the 2008 through 2014 American Community Surveys 1% PUMS, the 1990 Decennial Census 5% PUMS, and the 2000 Decennial Census 5% PUMS. The sample is restricted to adults, older than 22 and younger than 65. Individuals are either married heterosexuals or cohabitating gay men or lesbian women. Individuals who have moved between states in the past year are not included in the analysis. Annual earnings and hourly wages are in constant 1999 dollars. Full-time employment is defined as working 30 hours per week or more. Part-time employment is defined as working less than 30 hours per week. All regressions are estimated using OLS and include demographic and occupation controls, state fixed effects, year fixed effects, and state-by-year fixed effects. For a full list of control variables, see Table A.1. Standard errors clustered at the state level are in parentheses.

Table A.10: Robustness of Results on Wages to Controlling for Selection into Labor Market

Panel A.	Men: Ln Annual Income	Men: Ln Hourly Wages
	(1)	(2)
Gay × ENDA	0.040	0.029
	(0.027)	(0.021)
Gay × ENDA × Employer Size Minimum	-0.003	-0.005
	(0.003)	(0.003)
Gay × ENDA × Complaint Period	-0.001	-0.006
	(0.004)	(0.003)
Gay × ENDA × Compensatory Damages	0.084***	0.128***
	(0.021)	(0.021)
Gay × ENDA × Punitive Damages	-0.042*	-0.080***
	(0.020)	(0.018)
Gay × ENDA × Attorney's Fees	-0.069*	-0.029
	(0.033)	(0.031)
Observations	6,172,273	6,172,273
Panel B.	**Women: Ln Annual Income**	**Women: Ln Hourly Wages**
	(1)	(2)
Lesbian × ENDA	-0.115*	-0.027
	(0.049)	(0.038)
Lesbian × ENDA × Employer Size Minimum	0.005	0.004
	(0.004)	(0.003)
Lesbian × ENDA × Complaint Period	0.007	0.006
	(0.006)	(0.005)
Lesbian × ENDA × Compensatory Damages	-0.017	-0.088
	(0.054)	(0.046)
Lesbian × ENDA × Punitive Damages	-0.041	-0.007
	(0.027)	(0.023)
Lesbian × ENDA × Attorney's Fees	0.034	0.033
	(0.064)	(0.052)
Observations	5,061,060	5,061,060

*** $p<0.001$,** $p<0.01$, * $p<0.05$

Note: Author's calculations based on data from the 2008 through 2014 American Community Surveys 1% PUMS, the 1990 Decennial Census 5% PUMS, and the 2000 Decennial Census 5% PUMS. The sample is restricted to adults, older than 22 and younger than 65. Individuals are either married heterosexuals or cohabitating gay men or lesbian women. Individuals who have moved between states in the past year are not included in the analysis. Annual earnings and hourly wages are in constant 1999 dollars. Full-time employment is defined as working 30 hours per week or more. Part-time employment is defined as working less than 30 hours per week. All regressions are estimated using OLS and include demographic and occupation controls, state fixed effects, year fixed effects, and state-by-year fixed effects. For a full list of control variables, see Table A.1. Standard errors clustered at the state level are in parentheses.