The Gender Wage Gap: A Product of Misogyny and Gender Norms

Elizabeth Maloney

October 14, 2021

1 Introduction

The wage gap between men and women in the United States is long established. While this topic is well studied, much work builds off of Becker (1957). Becker developed his model of wage discrimination in the context of race, allowing employer utility to depend on their level of distaste toward employing black workers. In the context of racial discrimination, where slavery and segregation are part of our not-so-distant past, the presence of true distaste along racial lines and it's relevance to the black-white wage gap is clear. And in fact, Charles and Guryan (2008) find evidence that the predictions of the Becker model are born out in the data. However, the story with regard to the female wage gap is perhaps murkier. While true distaste for women undoubtedly exists, we must also consider the role gender norms play in determining labor market outcomes for women. Even in the absence of explicit distaste for women, norms affect educational attainment, whether or not women exit the labor market after marriage or child birth, which jobs they apply to, and even salary expectations if they are expected to earn less than their husbands.

Due to data constraints, many attempts to explain the female wage gap muddle the role of discrimination and norms. Charles et al. (2018), for example, propose a theoretical framework which describes the role of both norms and discrimination in determining the wage gap, but model both as a function of only one kind of sexism, and estimate it using data from the General Social Survey (GSS). The problem with this approach is that the GSS asks questions about gender norms, not distaste for women. To address this, I construct a novel index of misogyny (or distaste for women) based on Google Trends data on searches for sexist terms and provide evidence that this measure does in fact capture sexist attitudes. Using this measure as well as a measure of normative sexism constructed from GSS data, I am able to tease out the role of each in determining the wage gap. Then using my measure of distaste for women, I can test the empirical predictions of the Becker model against alternative explanations.

1.1 The Wage Gap

Efforts to explain the gender wage gap have largely begun by exploring the difference between the observable characteristics of men and women, including education and work experience. Blau and Kahn (2017) note that the education gap between men and women has actually reversed. Since 2011, women have had higher levels of education on average than men. They also show that women have made substantial gains with regard to labor market experience. In 2011, the gap between the labor market experience of men and women was only 1.4 years despite having been just under seven years in 1981. Experience gaps may reflect differences in use of time outside the workplace (Becker, 1985; Mincer and Polachek, 1974). Time-use surveys show that women participate in non-market labor like household chores and child care responsibilities at higher rates than men (Bianchi et al., 2000). Faced with the burden of responsibilities at home, women may allocate effort away from the market and toward household labor.

Even accounting for differences in human capital accumulation and other observable differences between men and women, a wage gap persists. Of course, I cannot assume this residual wage gap entirely reflects discrimination. Unobserved differences between men and women that contribute to productivity will also factor into observed wages. If these unobserved differences, on average, reflect higher productivity among men relative to women, this residual wage gap will overstate the effect of discrimination on wages. Alternatively, as Blau and Kahn (2007) point out, discrimination could influence a woman's rate of human capital accumulation as well as her choice of occupation. If this is true, then examining the residual wage gap after accounting for these factors may understate the role of discrimination. For example, if women are pushed into lower paying fields due to discrimination, the residual wage gap after accounting for occupational choice will underestimate the portion of the wage gap caused by discrimination. A widening distribution of the returns to different fields exacerbates this issue as the returns to male-dominated fields advance relatively more quickly and women are left "swimming upstream" (Blau and Kahn, 2007). Despite these drawbacks, the use of residual wage gaps is a common method in the literature. The Becker model of discrimination provides the theoretical underpinning of a discrimination-based wage gap.

1.2 The Becker Model of Taste-Based Discrimination

Consider the taste-based discrimination model described by Becker (1957) with a perfectly competitive labor market and employers who are prejudiced against women. Employer utility V_i increases with profit π_i and decreases with the interaction between degree of prejudice d and the number of women he employs. That is,

$$V_i = \pi_i - d_i F_i.$$

Assuming men and women are perfect substitutes, $\pi_i = f(F + W) - w_M M - w_F F$ where w_M and w_F are the wages for men M and women F, respectively. The utility maximizing choice of men and women at an interior solution is given by

$$f' - w_M = 0$$
$$f' - w_F - d_i = 0.$$

The short-run equilibrium involves a completely segregated labor force, with women working for the least prejudiced employers and men workers for the most prejudiced employers. Assuming the distribution of prejudice is fairly smooth, there must exist an employer who is indifferent between hiring men and women. I refer to this employer as the "marginal discriminator" and his degree of prejudice as the "marginal" level of prejudice d^*. This employer hires at equilibrium wages

$$w_M = w_F + d^*.$$

More prejudiced employers hire only men at wage $w_M = f'$ and less prejudiced employers hire only women at wage $w_F = w_m - d^*$. As a result, the equilibrium wage is entirely determined by the marginal level of prejudice, while the average level of prejudice will not directly affect the wage gap. To see why this is the case, suppose average prejudice is very high because there is a lot of prejudice at the top of the prejudice distribution. If the employers who actually hire women are unprejudiced, there will still be no wage gap. Since we expect the average prejudice in an area to be greater than the marginal level of prejudice, the wage gap should be systematically related to the marginal, rather than the average, level of prejudice in a labor market.

Suppose prejudice increases at the top of the prejudice distribution. Any increase above the marginal prejudice level should have no effect on the wage gap. If prejudice increases at the bottom of the prejudice distribution on the other hand, the marginal employer will now be more prejudiced, and the age gap will widen.

Now consider the impact of increasing the number of women in the labor market. All else equal, as additional women enter the labor market, they must work for increasingly prejudiced employers, pushing up the marginal prejudice level, and increasing the wage gap. Thus, the wage gap will increase as the proportion of women in the labor force increases.

1.3 Modifying the Becker Model to Reflect a Segregated Workplace

We obviously do not observe perfect segregation between men and women in the labor force. Neumark (1988) shows that when we allow for employer utility V_i to depend on the relative share of women hired rather than the absolute number, we do not achieve perfect segregation. Consider employer utility

$$V_i = \pi_i - d_i \, \frac{F_i}{M_i}.$$

where F_i and M_i represent the number of men and women employed by i. Then, again assuming an interior solution, the first order conditions yield equilibrium wages

$$\begin{aligned} w_m &= f' + d_i \frac{F_i}{M_i^2} \\ w_w &= f' - d_i \cdot \frac{M_i - F_i}{M_i^2}. \end{aligned}$$

As before, increased discrimination yields a larger wage gap, though now this occurs through both a premium to men and a penalty to women relative to their mutual marginal product of labor f'. Notice also that both the premium/penalty are reduced as the share of men at firm i increases. This occurs through the denominator. As each additional man joins the firm, he reduces the proportion of women there by less and less.

These two opposing forces result in a labor market that is not perfectly segregated by gender. Low discrimination firms attract women only up to the point where the increased relative share of women outweighs the lower level of discrimination. Conversely, even for more prejudiced firms, the penalty for women will be tempered by the fact that fewer women work there in equilibrium. Ultimately, this produces a labor force that is not perfectly segregated. Rather, the proportion of women hired at a given firm is decreasing in the prejudice of that firm. This complicates the predictions of the Becker model.

Consider an increase in the overall proportion of women in the labor market. Because the market is not perfectly segregated, this increases the proportion of women at all firms in the labor market. The premium to men (and penalty to women) has now increased, because men are relatively scarcer. Though the mechanism is slightly different, this result matches that of the original Becker model: as the proportion of women in the labor market increases, the wage gap should increase.

The main distinction between this specification of employer utility and that of the original Becker model is that this model does not predict a marginal discriminator. Women are hired by firms with various levels of discrimination and the level of discrimination is tempered by the proportion of women drawn to the firm. Consequently, no one point of the prejudice distribution should have more impact on the wage gap than another.

1.4 Black's Search Model of Discrimination

Black (1995) proposes a model of discrimination that incorporates search frictions. Male and female workers who enter the labor market search sequentially for a suitable job. In each search period, applicants pay a fixed cost and are matched with a potential employer who chooses whether or not hire them. Sexist employers will not hire female applicants, which means that on average, women must search longer for jobs. Because search costs are higher for women, their reservation wage is lower and unprejudiced employers can offer them a lower wage relative to men. In this model, the key drivers of the wage gap are the share of employers that are prejudiced and the share of the workforce made up by women. As the share of sexist employers increases, female applicants face increasingly higher search costs and the power of unprejudiced firms also increases, lowering female wages further. As more women enter the work force, unprejudiced firms are able to hire more low-cost (female) workers and their profits increase relative to prejudiced firms. As unprejudiced firms become more profitable, some prejudiced firms are driven from the market, reducing the share of prejudiced firms. This reduces women's search costs and increases their reservation wage meaning unprejudiced firms must offer them higher wages to induce them to accept a job offer. In this model, employers are either prejudiced or not so intensity of discrimination is not a relevant factor, though we would still expect a relationship between average prejudice levels and the wage gap since they average prejudice increases as the proportion of prejudiced firms increases.

1.5 Testing the Becker Model Empirically in the Context of Race

Charles and Guryan (2008) test the unmodified Becker model's predictions about the racial wage gap using responses about racial attitudes from the GSS. Respondents to the survey are asked questions like, "In general, how warm or cool do you feel toward blacks?," "How strongly would you object if a family member brought a black friend home for dinner?," and "Would you object to sending your kids to a school that had few/half/mostly black students?" Note that these types of questions reflect distaste for or a desire to segregate from black individuals.

Charles and Guryan recode responses so that the most prejudiced responses have highest numeric value and normalize each question by its mean and standard deviation in 1977. For each question k, individual i, and year t, the normalized individual prejudice questions are given by

$$\tilde{d}_{it}^{k} = \frac{d_{it}^{k} - E[d_{i,77}]}{\sqrt{Var(d_{i,t*}^{k})}}.$$

After normalizing each question, they take the average prejudice score among all of the relevant prejudice questions asked in that year. That is, they compute

$$D_{it} = \sum_{k} \tilde{d}_{it}^{k} / K_t$$

where K_t is the number of racial prejudice questions asked in year t. Finally, they regress D_{it} on a full set of year dummies to get $\tilde{D}_{it}$ which they aggregate by geographic region to construct the distributional prejudice measures. They take the average prejudice in a state to be the mean of $\tilde{D}_{it}$ in each state. They also compute the 10th, 50th, and 90th percentiles on $\tilde{D}_{it}$ for each state. This yields a measure prejudice at the lower, middle, and upper tail of the prejudice distribution.

They approximate the marginal prejudice level by computing the p^{th} percentile of the prejudice distribution where p is the percent of black people in a given area's workforce.

They find that, as predicted by the original Becker model, the wage gap is increasing in the marginal level of prejudice, the marginal level of prejudice is a stronger predictor of the wage gap than the average level of prejudice, and the bottom of the prejudice distribution matters more in predicting the wage gap than does the top of prejudice distribution. They also find that the the wage gap is increasing in the fraction of the labor market that is black, as predicted by the Becker model.

It is interesting that they are able to find evidence of the marginal prejudice level given that we do not observe a perfectly segregated labor market with regard to race. Recall that the modified Becker model which allows for such integration, does not predict a marginal discriminator. When they incorporate degree of workplace integration into their wage gap regressions, they find that wage gaps are larger in states with more integrated workplaces and that the relationship is statistically significant. However, this relationship decreases and becomes statistically insignificant when the marginal level of significance is accounted for, which they argue indicates that wage gaps are being driven by the ability of the market to segregate workers.

We can apply similar tests to the female wage gap considering the relationship between the wage gap and various points of the prejudice distribution, the proportion of women in the work force, and the proportion of employers who are prejudiced. To do this, however, we need a measure that captures distaste toward women.

2 Prejudice, Norms, and the Labor Market

An initial approach to measuring distaste might be to use questions about women from the GSS just as Charles and Guryan (2008) do for race. And in fact Charles et al. (2018) do just that. But there is an important distinction in how the GSS asks about race and how it asks about women. As I already laid out, questions on the GSS about racial attitudes get at distaste for interaction with black individuals. But the questions the GSS asks with regard to women, get at the role of men and women in society, rather than at disdain for women. For example, the survey asks whether respondents agree that, "It is much better for everyone involved if the man is the achiever outside the home and the woman takes care of the home and family" or that, "A working mother can establish just as warm and secure a relationship with her child as a mother who does not work." Table 1 provides the full list of GSS questions about women. These questions are really about gender norms rather than disdain for women.

This begs the question about the role of norms versus the role of discrimination in determining the labor market outcomes of women. While a more traditional view of gender roles in society can certainly be considered sexist, it may not result directly in the kind of labor market discrimination the Becker and Black models describe. Consider someone raised to think that children are better off when their mother stays home with them and expect their own wife to fulfill this role. This same person may not get disutility from working with other women and may be happy to hire women at his firm. This type of normative sexism may not result in a wage gap the way distaste for women would– at least not through the mechanisms proposed by the discrimination models described above. This is not to say that we wouldn't expect norms to impact the wages of women. We would expect gender norms to influence how much education a woman obtains, which fields she opts into, and even how much money she negotiates for. But the implications of a wage gap due to norms are very different than those of a wage gap due to discrimination. And an estimate of the impact of discrimination in the labor market will likely be overstated if norms are not accounted for.

2.1 Measuring Normative Sexism

In order to measure normative sexism, I use restricted-access GSS data on survey responses to questions about gender roles and perform the same transformations used by Charles and Guryan (2008) which I describe above. These GSS questions are listed in Table 1. This procedure is similar to the one used by Charles et al. (2018). Figure 2 shows the trends in average prejudice against women for each census division over time. There is a general decrease in prejudice over time in all regions. Figure 3 depicts this index and the questions used to construct it, across the US over time. Respondents give less prejudiced responses in more recent years. Note that sexist responses to the question about how warm and secure a relationship working mothers can have with their children have decreased much more slowly than responses about whether or not a preschooler suffers when his or her mother works. This difference between answers to very similar questions may indicate that responses are very sensitive to exactly how a question is phrased, and highlights one concern with survey-based measures. Interestingly, sexist responses to the question about voting for a female president have declined much more slowly than those of other questions.

To get a sense of how responses to these questions vary by demographic characteristics, Table 3 gives regressions showing how age, education, and gender relate to the aggregate level of individual prejudice as well as to responses to the prejudice questions used to construct the aggregate index. These regressions include year and state fixed effects and are clustered at the state level. Recall that responses were coded so that more prejudiced responses were given higher values, meaning positive coefficients indicate an increase in sexist responses. Across almost all questions, prejudice against women increases with age, decreases with education, and is generally lower for women relative to men. The one exception is that women are actually more likely than men to say they would not vote for a female president if their party nominated a qualified candidate, though this difference is not statistically significant.

Table 2 gives the average responses to the GSS questions about women as well as the average prejudice index by census division, after regressing on years to account for changes over time. Prejudice of this kind is highest in the south east of the country and lowest in New England. Most survey questions seem to follow the same geographic trends as the aggregate index of normative sexism, though there is some variation across regions especially in questions about a working mother's relationship with her children, whether or not married women should work, and whether a wife should put her husband's career first. Figure 1 gives the variation in the aggregate index of normative sexism across states. As reflected in the census level tables, normative sexism is highest in the south east of the country, but there are a few states out west with high levels of normative sexism including Utah and Nebraska.

Note that in contrast with the questions used by Charles and Guryan (2008) to measure racial prejudice, GSS questions about women do not identify a desire for segregation between men and women, but rather opinions about the appropriate domain of men and women. To see this more clearly, consider Figure 4, which gives the relationship between the normative sexism index when we only include responses by men compared to the same index constructed using responses by women. If the GSS measure reflected distaste for women, we would expect women's answers to these questions to differ strongly from those of men. Yet, the two are strongly correlated (p-value less than 0.001). This supports the notion that this measure gets at societal norms, which are more likely to be aligned between men and women in a given community as compared to distaste for women.

2.2 Measuring Distaste for Women Using a Search-based Sexism Index

There are two main concerns about using GSS survey responses to construct an index of sexism. As already discussed, I believe these survey questions actually get at norms rather than discrimination. Additionally, survey-based estimates of prejudice receive criticism because of social censoring that may bias estimates of prejudice downward, especially in areas where it is less socially acceptable to express prejudiced views (Berinsky, 1999). In his paper on the role of racial prejudice on vote shares in either of Obama's presidential runs, Stephens-Davidowitz (2014) uses Google Trends data on searches for "the n-word" as a proxy for racial prejudice. He finds that his proxy strongly correlates with GSS measures of racial prejudice, but argues that it improves upon survey-based measures because it is not tempered by social attitudes toward prejudice. He points out that the vote-share for John Kerry (a democrat) is negatively correlated with more prejudiced responses to GSS questions, but that it does not have a statistically significant correlation with his google search proxy. He argues that in more democratic areas, it is socially unacceptable to admit racial prejudice so measures like the GSS underestimate the degree of racial prejudice in such areas. Compared to other surveys, the GSS is especially vulnerable to this critique since it is conducted via face-to-face interviews. Google searches on the other hand, are conducted anonymously. Beyond addressing the concern of social-censoring, this method allows me to construct a measure that gets at distaste toward women rather than societal norms.

Google Trends provides data on the rate of Google searches for a given phrase at the US state or direct market area levels over various segments of time based on a random sample of all searches. They report these relative to the maximum rate of searches for the term experienced in any region during a specified time period, scaled by 100. For example, if Ohio has the highest rate of searches for the term "weather" across the US, it would receive a score of 100. Then a score of 25 for Indiana would reflect that the search rate for "weather" in Indiana is one quarter of Ohio's during the time period. I can obtain these for various periods of time, but since each is normalized relative to the maximum in that time period, these cannot be readily compared across time.

To construct a measure of misogyny or distaste for women, I take a series of sexist search terms and collect the relative interest by region for each term over the period of a year. I do this for both the state and direct market area (DMA) levels. Because the google search index is a sample, I collect 10 samples for each year and search term and average over all samples. I normalize the interest in each word by its mean and standard deviation in 2016. For each observation, I take the average over all the words I want to include in my index, and regress each of these values on year dummies, taking the residuals as my index of misogyny or distaste for women. Figure 9 gives the average levels of misogyny by state. This differs quite a bit from the average levels of normative sexism by state. For example, it is much lower in the south east of the United States.

To construct such an index, I need to identify which which search terms to include in my prejudice index. Following Stephens-Davidowitz (2014), I initially consider slurs against women. This is not without basis in the literature. In a paper aimed at describing misogynistic language used on the Reddit "Manosphere" (a portion of the internet devoted to promoting misogyny), Farrell et al. (2019) identify words such as "bitch," "cunt," and "whore" as part of group they call "Hostility" which "includes violent verbs, and slurs that are not immediately racist or homophobic." They find that this group of words is rampant across the manosphere, and is actually the most prominent of all categories they examine. I take these three words as well as synonyms for them as my main search terms of interest. I call this group of terms "Derogatory." I also collect data on search trends for three other categories of words which may also be informative if not as easily interpreted. The categories are "Violent," "Manosphere," and "Reactionary." The "Violent" search terms describe violent sexual acts that are often (though not always) perpetrated against women. The "Manosphere" group includes terms often used by members of the portion of the internet devoted to men's rights. For example, I include "misandry" and "men's rights." Finally, the "Reactionary" terms are those one might search in response to a negative incident and include the words "sexual harassment" and "workplace harassment." Table 4 lists all words in each category.

Tables 5, 6, 7 and 8 list the top 20 related search queries for each of the terms given in Table 4. Note that these tables contain a lot of offensive language, but I include them in order to give a sense of the intent behind some of these searches. While some of the related queries are obviously in search of anti-woman content, others pertain to song lyrics or searches for definitions of the terms. Similarly noisy search terms have been able to explain regional behavior differences well in other work. For example, Stephens-Davidowitz (2014) shows that variation in search interest for the word “God” strongly predicts regional variation in the the percent of people who believe in God (R^2 of 0.65), despite the fact that the top query related to searches of the term “god” during the time period he examine is “god of war,” a popular video game released during in the middle search period. Interest in “god of war” actually had almost double the search interest of ”church of god” during this time period. “Greek god”, ”god of war walkthrough”, “oh god”, “egyptian god”, their eyes were watching god”, “sun god”, and “god bless america,” are also found in the top search queries related to searches for the word “god.”

Google Trends also provides data on the rate of Google searches for a given topic over time for a given region. These are provided by week and are normalized by the maximum rate of searches on the topic in a given week, scaled by 100. Figures 5, 6, 7, and 8 give the interest over time for the words listed in Table 4 across the US between 2006 and 2020.

Interest in the “Derogatory”, “Manosphere”, and “Violent” groups of words have been steadily trending downward in recent years. By contrast, search interest in the terms sexual harassment and workplace harassment, seen in Figure 5 has been much more stable, without an obvious up- or downward trend. We do see two notable spikes in reactionary words. The first occurs in search interest for “workplace harassment” in the Fall of 2009. This may be due to David Letterman’s October 2009 revelation that he had slept with female members of his staff and had been extorted over the affairs. The second big spike in both “workplace harassment” and “sexual harassment” occur in late 2017. This is likely due to the “me too” movement, which exploded in late 2017 after sexual assault allegations against Harvey Weinstein were brought to the public’s attention. These spikes likely reflects that interest in these words may be driven by external media buzz rather than by individual experiences.

Turning to the group of “Manosphere” words shown in Figure 7, note the spike in the searches for “SJW” during late 2016. “SJW” is an abbreviation of the term “social justice warrior” and was a common critique of the Hillary Clinton coalition during the 2016 presidential election. Search interest in the term “misandry” spikes in late 2014. Interestingly, this is likely due to the increasingly frequent ironic use of the word in feminist portions of the internet, which was highlighted in articles published by Slate, Time Magazine, and The Guardian in August and December of 2014.

It would be reasonable to worry that the words included in my sexism index reflect porn interest rather than sexism. Figure 10 gives the correlation between my misogyny index, the average search interest in each each group of words individually, the normative sexism index constructed using GSS survey data, and google search interest in the word "porn". Note that search interest in the the word porn include searches for sites like "Porn Hub" or other sites with porn in the name, so it should be correlated with porn usage even if it's an imperfect measure. While there is a weak correlation between searches for porn and my index of misogyny this correlation is statistically insignificant and much weaker than the correlation between porn and the index of normative sexism.

Figure 10 also includes the correlation between the various sexism measures and the partisan voting index (PVI), aggregated to state level and re-coded so that positive values indicate a propensity to vote for republicans and negative values indicate a propensity to vote for democrats. Note the inverse correlation of the "Reactionary" word index with both the normative sexism index. This could indicate that issues of sexual and workplace harassment are more salient in areas with less normative sexism. Note also that the "Manosphere" word index is also negatively correlated with the normative sexism, but positively associated with misogyny. This could be consistent with men in areas where it is less socially acceptable to express sexist ideologies turning to the internet as an outlet for expressing their sexism. If these searches are representative of a large swath of the population, then these searches could account for some unspoken, but pervasive sexist ideologies. If, however, these are limited to a fringe group of men unlikely to be in employment or management positions, these would be unlikely to impact female labor force outcomes like the wage gap. Anecdotal evidence about these groups supports the notion that these men tend to exist on the outskirts of society, feeling left out by the mainstream and emasculated by women. If this is the case, than interest in these specific "Manosphere" words may be less relevant for labor market outcomes.

Figure 11 plots the index of gender norms against the index of misogyny. The two are correlated, but weakly so. Notice states like Alabama which have very strong gender norms, but below average levels of misogyny. On the flip side states like Montana have very high levels of misogyny despite weaker gender norms.

2.3 Testing the Predictions of Various Discrimination Models

To test whether or not I see evidence that the wage gap is consistent with the predictions of either formulation of the Becker model, I follow Charles and Guryan (2008). Specifically, I use a multi-stage regression model. In the first stage, I regress log wages on education, a quadratic in potential experience, and gender-specific year and state dummies using data from the May CPS Outgoing Rotation Groups which has been cleaned and standardized by the Centre for Economic Policy Research. The coefficients on the state by female indicators give the residual wage gap in each given state, and I use these as the dependent variable in a second regression on various parts of the prejudice distribution.

To the distribution of misogyny, I use DMA level data on search interest for the derogatory search terms after residualizing out year fixed effects. This gives me within-state variation in misogyny. I can then compute the median, tenth, and ninetieth percentiles of the misogyny distribution within states, weighting each DMA by its population. I also compute the marginal level of misogyny to be the level of misogyny at the p^{th} percentile of the misogyny distribution where p is the percent of women in the labor force, again weighting by DMA population. This reflects the model's definition of the marginal discriminator as the least prejudiced employer who hires women. To compute various points in the distribution of normative sexism, I use the aggregated response score to GSS survey questions after normalizing and residualizing out year effects and compute the various percentiles of the prejudice distribution. Finally, I estimate the proportion of the population that is sexist by taking the percent of GSS individual aggregated prejudice levels that are greater than the normalized average.

If our data is consistent with the original Becker model (Becker, 1957), I expect the wage gap to be increasing in the marginal level of sexism and that this relationship is stronger than that of the average prejudice level and the wage gap. The wage gap should also be more sensitive to increases in prejudice at the lower tail of the prejudice distribution relative to its upper tail, which is likely to be above the marginal level of prejudice. As more women enter the labor force, women must work for increasingly prejudiced employers which drives up the wage gap, so the wage gap should be increasing in the proportion of women in the labor market. Neumark's 1988 modified Becker model, which allows for integration between men and women in the labor market, does not predict a marginal discriminator, but still predicts an increase in the wage gap as more women enter the labor market. Finally, Black's 1995 search model of discrimination predicts a segregated labor market that is driven by the search costs women face as prejudiced employers refuse to hire them in a sequential search setting. This model does not consider intensity of prejudice and the wage gap is principally determined by the share of prejudiced employers in the labor market. As the share of prejudiced firms increase, women face higher search costs since they must wait longer on average to match with an unprejudiced employer and they are thus willing to accept a lower wage when they finally match with an employer who will hire them. Unlike either formulation of the Becker model, Black predicts that the wage gap is decreasing in the the share of female labor force participation because increased female labor force participation increases the amount of low-cost labor available to unprejudiced firms giving them a market advantage over prejudiced farms who only hire the more expensive male labor. As prejudiced firms are driven out of the labor force by the more cost-effective unprejudiced firms, the search costs faced by women decrease and they can demand higher wages.

3 Results

3.1 Measures of Sexism and Labor Market Outcomes

Tables 9 and 10 give the average impact of both indices on outcomes related to the labor market. Panel A of table 9 shows that both measures are increasing in the residual wage gap (controlling for gender-specific year effects, education, experience, and experience squared). This relationship is statistically significant. Here, a negative coefficient indicates an increase in the wage gap. We can also see that the relationship between normative sexism and the wage gap is driven by normative sexism among women, rather than men. When we include each separately, normative sexism among women is a statistically significant predictor of the wage gap whether or not we include misogyny, but normative sexism among men is not. This bolsters the argument that the GSS survey picks up on gender norms, which is likely to be shared by women and men rather as opposed to discrimination against women.

Panel B of Table 9 shows the relationship of each measure with the proportion of women between ages 20 and 40 who have never married after controlling for level of education, age, and year. Both measures are associated with smaller shares of such women, implying that women in areas with stronger gender norms and greater misogyny are more likely to marry young, all else equal. Similarly, Panel C shows that both measures are also associated with lower average ages at first birth (after controlling for year fixed effects), meaning women tend to start having children earlier in areas with more traditional gender norms and greater misogyny, all else equal.

Panel A of Table 10 shows the relationship between each measure of sexism and the residual labor force participation gap after controlling for education level, age categories, and gender-specific year fixed effects. The gap between male and female participation rates is increasing in both types of sexism and these relationships are statistically significant. Panel B sheds additional light on this by showing that both types of sexism are associated with less female labor force participation after controlling for age, education, and year fixed effects, but only normative sexism among women is a statistically significant predictor at the 5% level. That is to say, when we look at women, irrespective of men, norms internalized by women are the strongest predictors of their labor force participation rates. But when we consider labor market participation relative to men, it is clear that both misogyny and norms are important determinants of participation.

Panel C of Table 10 shows that neither measure of sexism is statistically significantly related to the gap between men and women in obtaining a college degree, conditional on gender-specific year dummies and age categories. Note though that misogyny toward women is actually higher in areas where the gap is smaller. Panel D shows that though the gap between men and women is smaller in areas with higher misogyny, women there are still getting less education (conditional on age and year dummeis) than in areas with lower misogyny, all else equal. This is consistent with lower education among both men and women in high misogyny areas. Panel D makes clear that strongly gendered norms and misogyny are both related to reduced educational attainment among women.

3.1.1 Models of Discrimination

Table 11 gives the regressions of the residual wage gap on various parts of the misogyny distribution. A negative coefficient indicates an increase in the wage gap for women. The marginal level of prejudice is not a stronger predictor of the wage gap than the mean level of prejudice. Moreover, when I include various points of the prejudice distribution in the model, the wage gap is not more strongly related to the bottom of the prejudice distribution than it is to the top. Furthermore, the wage gap is decreasing in the percent of women in the labor force. Thus, data on the female wage gap is not consistent with the predictions of either formulation of the Becker model. Since the wage gap is increasing in the proportion of women in the labor market and decreasing in the proportion of sexist individuals, it is consistent with Black's model of labor market discrimination, though I would have expected the relationship between the wage gap and the proportion of sexist individuals to be stronger as this is the driving force of the Black model. It should be noted that we would expect the percent of women in the labor force to be strongly influenced by misogyny. In fact, 10 shows the strong negative correlation between the two.

Table 12 gives the same results using the measure of gender norms. The results are generally the same although the proportion of sexist individuals is actually decreasing in the wage gap rather than increasing in it. This is likely due to multi-colinearity between average normative sexism and the percent of sexist individuals, which can bee seen in Figure 10.

4 Conclusion

In this paper, I construct an entirely novel measure of gender prejudice based on Google search trends for misogynistic words. I present evidence that this measure captures distaste for women whereas previous measures have only been able to capture normative sexism or beliefs about gender norms. I show that both measures are strongly correlated with women's labor market outcomes including the wage gap, the probability of remaining unmarried for women between the ages of 20-40, the age at which women have their first child, labor force participation, and attainment of a college degree. These measures are not picking up the same variation as they are weakly correlated and are individually significant even when both are included in models of these labor market outcomes. This suggests that appropriate models of the wage differential between men and women will consider the separate impacts of misogyny and norms.

In addition, I test the predictions of several models of labor market discrimination using both measures of sexism. I find some evidence consistent with Black's search model of discrimination, and show that the data is inconsistent with Becker-type models of discrimination. Given the fact that these models only consider the role of distaste for women and not the separate influence of gender norms, perhaps it is not surprising that the data does not readily conform to any of them. More work is needed to properly model the interaction of norms and misogyny in determining the wage gap.

Much of the work on gender wage discrimination muddles normative sexism and misogyny or distaste for women, but this paper highlights the importance of capturing each individually in order to understand wage differentials between men and women. The policy implications of these different channels of discrimination also diverge. Policies aimed at narrowing the gender wage gap will need to address both sources of discrimination. For example, attempts to prevent employers from discriminating will do little to improve the normative factors that prevent women from entering the labor market, investing in human capital, and negotiating for higher salaries. On the other hand, I show that we cannot attribute the wage gap entirely to gender norms. True distaste for women plays an important and previously unmeasured role in determining the wage gap.

5 Tables

fework:	*Married Women Shouldn't Work* Do you approve or disapprove of a married woman earning money in business or industry if she has a husband capable of supporting her?
fehome:	*Women Should Run Homes* Do you agree or disagree with this statement? Women should take care of running their homes and leave running the country up to men.
fepol:	*Men Better for Politics* Tell me if you agree or disagree with this statement: Most men are better suited emotionally for politics than are most women.
fepres:	*Would Not Vote for Woman President* If your party nominated a woman for President, would you vote for her if she were qualified for the job?
fechld:	*Working Mothers Worse Relationship* A working mother can establish just as warm and secure a relationship with her children as a mother who does not work.
fepresch:	*Preschoolers Suffer When Mom Works* A preschool child is likely to suffer if his or her mother works.
fehelp:	*Husband's Career More Important than Wife's* It is more important for a wife to help her husband's career than to have one herself.
fefam:	*Wife Takes Care of Home and Family* It is much better for everyone involved if the man is the achiever outside the home and the women takes care of the home and family.

Table 1: GSS Survey Questions on Attitudes Toward Women

	Normative Sexism	Married Women Shouldn't Work	Women Should Run Home	Men Better for Politics	Would Not Vote for Woman President
Panel A.					
East South Central	0.147	0.102	0.297	0.234	0.154
West South Central	0.071	0.049	0.065	0.087	0.077
South Atlantic	0.060	0.023	0.126	0.078	0.060
East North Central	-0.027	-0.001	-0.050	-0.034	-0.027
West North Central	-0.035	0.058	-0.049	-0.043	-0.006
Mountain	-0.039	0.033	-0.085	-0.093	-0.044
Pacific	-0.042	-0.076	-0.138	-0.072	-0.058
Middle Atlantic	-0.043	-0.040	-0.040	-0.049	-0.050
New England	-0.128	-0.109	-0.145	-0.138	-0.099
	Working Mothers Worse Relationship	Preschoolers Suffer When Mom Works	Husband's Career More Important Than Wife's	Wife Takes Care of Home and Family	
Panel B.					
East South Central	0.073	0.046	0.145	0.227	
West South Central	0.012	0.033	0.080	0.141	
South Atlantic	0.026	0.007	0.083	0.068	
East North Central	-0.001	-0.024	-0.016	-0.048	
West North Central	-0.040	-0.065	-0.029	-0.090	
Mountain	-0.025	0.033	-0.095	-0.071	
Pacific	0.018	0.040	-0.118	-0.060	
Middle Atlantic	-0.030	-0.021	0.008	-0.041	
New England	-0.096	-0.077	-0.138	-0.176	

Table 2: GSS Survey Responses by Census Division

Derogatory	Manosphere
bitch	misandry
cunt	men's rights
whore	sjw
thot	
twat	
Violent	**Reactionary**
rape	sexual harassment
gangbang	workplace harassment

Table 4: Sexist Search Terms

	Normative Sexism				Married Women Shouldn't Work				Women Should Run Homes			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A.												
Age	0.104***			0.088***	0.102***			0.078***	0.128***			0.094***
	(0.004)			(0.004)	(0.007)			(0.006)	(0.005)			(0.004)
Education		-0.059***		-0.050***		-0.063***		-0.051***		-0.088***		-0.073***
		(0.001)		(0.002)		(0.002)		(0.003)		(0.003)		(0.002)
Female			-0.120***	-0.142***			-0.005	-0.032*			-0.004	-0.039*
			(0.008)	(0.007)			(0.013)	(0.012)			(0.019)	(0.017)
N	38275	38527	38630	38182	22646	22762	22824	22592	22363	22477	22538	22309
R^2	0.10	0.10	0.03	0.17	0.07	0.07	0.03	0.10	0.12	0.15	0.06	0.18
	Men Better for Politics				Would Not Vote for Woman President				Working Mothers Worse Relationship			
Panel B.												
Age	0.069***			0.059***	0.077***			0.065***	0.081***			0.074***
	(0.005)			(0.005)	(0.005)			(0.005)	(0.004)			(0.004)
Education		-0.042***		-0.036***		-0.039***		-0.030***		-0.043***		-0.039***
		(0.001)		(0.001)		(0.002)		(0.001)		(0.002)		(0.002)
Female			-0.023	-0.039**			0.021	0.004			-0.308***	-0.328***
			(0.013)	(0.012)			(0.013)	(0.012)			(0.011)	(0.011)
N	34252	34471	34555	34174	25052	25188	25259	24989	30094	30293	30372	30020
R^2	0.08	0.08	0.06	0.10	0.07	0.06	0.04	0.08	0.05	0.05	0.06	0.10
	Husband's Career More Important than Wife's				Preschoolers Suffer When Mom Works				Wife Takes Care of Home and Family			
Panel C.												
Age	0.188***			0.165***	0.111***			0.104***	0.162***			0.143***
	(0.006)			(0.005)	(0.005)			(0.005)	(0.005)			(0.004)
Education		-0.084***		-0.062***		-0.044***		-0.037***		-0.089***		-0.077***
		(0.003)		(0.003)		(0.003)		(0.004)		(0.003)		(0.004)
Female			-0.095***	-0.151***			-0.292***	-0.313***			-0.157***	-0.193***
			(0.023)	(0.020)			(0.010)	(0.010)			(0.013)	(0.012)
N	15194	15267	15311	15154	29773	29977	30050	29704	29799	30004	30075	29732
R^2	0.18	0.14	0.08	0.22	0.08	0.06	0.07	0.12	0.12	0.11	0.05	0.17

Table 3: GSS Survey Responses and Demographic Traits

Bitch		*Cunt*		*Whore*	
Int.	Query	Int.	Query	Int.	Query
100	you bitch	100	fuck	100	slut
49	that bitch	62	definition	51	whore house
45	bitch lyrics	61	cunt definition	44	whore definition
39	bad bitch	57	what is cunt	29	man whore
35	bitch ass	54	cunt mean	27	hoe
26	bitch song	54	what does cunt	26	whore lyrics
22	bitch meme	52	what a cunt	26	whore meaning
22	this bitch	47	what does cunt mean	26	whore meme
21	im a bitch	44	cunt meaning	25	bitch
15	fuck you bitch	44	ass	22	what is a whore
15	bitch nigga	41	what is a cunt	21	whore of babylon
14	lil bitch	33	word cunt	21	boo you whore
13	bitch in spanish	33	anal cunt	18	whore gif
12	bitch face	29	bitch	17	definition of whore
12	son of a bitch	22	definition of cunt	15	whore movie
11	bitch quotes	21	slut	15	attention whore
11	little bitch	19	cunt wars	14	prostitute
10	crazy bitch	18	cunt meme	14	whore mouth
10	bitch gif	18	define cunt	14	whore song
10	fat bitch	16	whore	13	whore in this moment

Thot		*Twat*	
Int.	Query	Int.	Query
100	what is thot	100	twat waffle
74	thot meaning	64	definition twat
74	what is a thot	61	definition
69	definition thot	53	twat meaning
68	thot urban	51	what is twat
67	definition	45	what is a twat
67	urban dictionary thot	38	what does twat
65	she a thot thot	35	what does twat mean
64	urban dictionary	33	cunt
57	thot thot lyrics	23	twat urban
56	thot mean	23	twat define
56	what does thot	21	twat dictionary
47	what does thot mean	19	urban dictionary twat
39	thot meme	16	urban dictionary
36	begone thot	15	definition of twat
36	white thot	14	twat meme
31	thot spot	13	twit
30	hoe	11	ass
27	duckie thot	11	twat slang
26	thots	11	twatt

Table 5: Top 20 Queries and Relative Interest - Derogatory Search Terms

Rape		*Gangbang*	
Int.	Query	Int.	Query
100	trump rape	100	gangbang porn
92	gay rape	55	gangbang creampie
89	rape victim	51	black gangbang
81	rape video	49	anal
79	rape videos	49	gangbang anal
78	what is rape	48	wife gangbang
76	statutory	43	teen gangbang
74	gang rape	34	girl gangbang
73	statutory rape	33	bbc gangbang
71	date rape	30	big gangbang
64	rape movie	28	gay gangbang
61	teen rape	27	gangbang sex
56	rape me	25	cum gangbang
46	rape definition	21	interracial gangbang
46	rape statistics	20	gangbang dp
45	ear rape	19	forced gangbang
43	child rape	18	ebony gangbang
42	rape stories	18	gangbang videos
39	real rape	17	free gangbang
37	prison rape	17	gangbang milf

Table 6: Top 20 Queries and Relative Interest - Violent Search Terms

Workplace Harassment		*Sexual Harassment*	
Int.	Query	Int.	Query
100	harassment in workplace	100	sexual harassment definition
89	harassment in the workplace	96	harassment definition
74	workplace sexual harassment	89	sexual harassment workplace
74	sexual harassment	80	sexual harassment training
57	sexual harassment in workplace	78	what is sexual harassment
54	sexual harassment in the workplace	67	sexual harassment in workplace
24	work harassment	60	sexual harassment in the workplace
19	harassment at workplace	53	sexual assault
18	what is harassment	53	sexual harassment law
17	what is workplace harassment	39	sexual harrassment
14	discrimination	39	sexual harassment policy
14	workplace harassment laws	35	sexual harassment news
14	harassment laws	34	sexual harassment laws
14	workplace harassment definition	33	sexual harassment at work
13	harassment definition	33	accused of sexual harassment
13	workplace discrimination	32	definition of sexual harassment
11	harassment at the workplace	31	discrimination
11	what is harassment in the workplace	30	sexual harassment california
10	harassment at work	30	quid pro quo harassment
10	workplace bullying	29	quid pro quo sexual harassment

Table 7: Top 20 Queries and Relative Interest - Reactionary Search Terms

Men's Rights		*Misandry*		*SJW*	
Int.	Query	Int.	Query	Int.	Query
100	mens rights reddit	100	misogyny	100	sjw meaning
64	mens rights movement	60	feminism	60	what is sjw
36	mens rights activists	48	define misandry	42	sjw meme
36	mens rights activist	34	what is misandry	39	what does sjw
26	mens rea	23	misogynist	31	sjw reddit
25	mra	23	feminist	27	what does sjw mean
		23	misandrist	14	marvel sjw
		20	misogyny definition	14	tumblr sjw
		20	misandry meaning	13	anti sjw
		18	definition of misandry	12	youtube sjw
		13	misogynistic	11	sjw star wars
		13	misandry today	10	what is an sjw
		12	misandry game of thrones	9	sjw hate
		11	misandry pronunciation	9	what is a sjw
		11	misandry meme	9	sjw cringe
		10	feminism definition	8	sjw feminist
		10	misandry gif	8	sjw memes
		10	misandry bubble	8	whats sjw
		10	misandry def	8	sjws
		7	misogyny meaning	7	define sjw

Table 8: Top 20 Queries and Relative Interest - Manosphere Words

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. Dependent Variable - Wage Gap							
Misogyny	-0.108**				-0.098*	-0.100*	-0.101**
	(0.038)				(0.038)	(0.038)	(0.037)
Normative Sexism		-0.134*			-0.114*		
		(0.059)			(0.057)		
Normative Sexism - Men			-0.094			-0.075	
			(0.054)			(0.052)	
Normative Sexism - Women				-0.131*			-0.118*
				(0.058)			(0.055)
N	51	51	51	51	51	51	51
R^2	0.14	0.09	0.06	0.09	0.21	0.17	0.21
Panel B. Dependent Variable - Proportion of Women Aged 20-40 Never Married							
Misogyny	-0.173***				-0.150***	-0.151**	-0.159***
	(0.047)				(0.042)	(0.043)	(0.042)
Normative Sexism		-0.290***			-0.258***		
		(0.070)			(0.063)		
Normative Sexism - Men			-0.230***			-0.199**	
			(0.065)			(0.059)	
Normative Sexism - Women				-0.264***			-0.243***
				(0.070)			(0.063)
N	51	51	51	51	51	51	51
R^2	0.22	0.26	0.20	0.22	0.42	0.37	0.40
Panel C. Dependent Variable - Average Age at First Birth							
Misogyny	-4.660***				-3.925***	-4.024***	-4.154***
	(1.200)				(0.982)	(1.031)	(1.009)
Normative Sexism		-8.108***			-7.293***		
		(1.594)			(1.410)		
Normative Sexism - Men			-6.577***			-5.862***	
			(1.485)			(1.320)	
Normative Sexism - Women				-7.273***			-6.654***
				(1.617)			(1.412)
N	51	51	51	51	51	51	51
R^2	0.24	0.35	0.29	0.29	0.51	0.46	0.48

Table 9: Relationship between Sexism Indices and Labor Market Outcomes (Part 1)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. Dependent Variable - Labor Force Participation Gap							
Misogyny	-0.112**				-0.101**	-0.103**	-0.105**
	(0.034)				(0.033)	(0.034)	(0.033)
Normative Sexism		-0.140*			-0.118*		
		(0.054)			(0.050)		
Normative Sexism - Men			-0.100*			-0.078	
			(0.049)			(0.046)	
Normative Sexism - Women				-0.136*			-0.122*
				(0.053)			(0.049)
N	51	51	51	51	51	51	51
R^2	0.18	0.12	0.08	0.12	0.26	0.23	0.28
Panel B. Dependent Variable - Female Labor Force Participation Rate							
Misogyny	-0.061				-0.055	-0.059	-0.055
	(0.033)				(0.033)	(0.034)	(0.032)
Normative Sexism		-0.076			-0.063		
		(0.051)			(0.051)		
Normative Sexism - Men			-0.031			-0.019	
			(0.047)			(0.046)	
Normative Sexism - Women				-0.106*			-0.099*
				(0.049)			(0.048)
N	51	51	51	51	51	51	51
R^2	0.06	0.04	0.01	0.09	0.09	0.07	0.14
Panel C. Dependent Variable - Gender Gap in College Degree Attainment							
Misogyny	0.011				0.015	0.015	0.014
	(0.014)				(0.014)	(0.014)	(0.014)
Normative Sexism		-0.039			-0.043		
		(0.021)			(0.021)		
Normative Sexism - Men			-0.029			-0.032	
			(0.019)			(0.019)	
Normative Sexism - Women				-0.039			-0.041
				(0.021)			(0.021)
N	51	51	51	51	51	51	51
R^2	0.01	0.07	0.04	0.07	0.09	0.06	0.09
Panel D. Dependent Variable - Female College Degree Attainment							
Misogyny	-0.366***				-0.322***	-0.322**	-0.342***
	(0.100)				(0.092)	(0.093)	(0.094)
Normative Sexism		-0.535**			-0.465**		
		(0.153)			(0.140)		
Normative Sexism - Men			-0.466**			-0.399**	
			(0.138)			(0.126)	
Normative Sexism - Women				-0.433**			-0.386**
				(0.156)			(0.140)
N	51	51	51	51	51	51	51
R^2	0.22	0.20	0.19	0.14	0.36	0.35	0.32

Table 10: Relationship between Sexism Indices and Labor Market Outcomes (Part 2)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Mean	-0.108**		-0.215*	-0.053			-0.012
	(0.038)		(0.082)	(0.059)			(0.027)
Marginal		-0.095*	0.092	0.039			
		(0.043)	(0.082)	(0.055)			
Percent Women				2.022***		2.087***	2.013***
				(0.276)		(0.271)	(0.247)
Lower Tail					0.004	-0.006	
					(0.044)	(0.029)	
Median					-0.066	0.013	
					(0.063)	(0.042)	
Upper Tail					-0.015	-0.004	
					(0.023)	(0.015)	
Percent Sexist							-0.047
							(0.052)
N	51	46	46	46	46	46	51
R^2	0.14	0.10	0.22	0.66	0.16	0.66	0.66

Standard errors in parentheses

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

Table 11: Relationship between Gender Wage Gap, Fraction of Women in the Workforce, and Misogyny

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Mean	-0.134*		-0.425**	-0.035			-0.207
	(0.059)		(0.125)	(0.095)			(0.117)
Marginal		-0.050	0.301*	-0.021			
		(0.057)	(0.116)	(0.086)			
Percent Women				2.055***		1.953***	2.024***
				(0.255)		(0.247)	(0.222)
Lower Tail					-0.230**	-0.050	
					(0.080)	(0.057)	
Median					-0.057	-0.028	
					(0.075)	(0.049)	
Upper Tail					0.028	-0.000	
					(0.057)	(0.038)	
Percent Sexist							0.215
							(0.157)
N	51	51	51	51	51	51	51
R^2	0.09	0.02	0.21	0.67	0.22	0.67	0.68

Standard errors in parentheses

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

Table 12: Relationship between Gender Wage Gap, Fraction of Women in the Workforce, and Normative Sexism

6 Figures

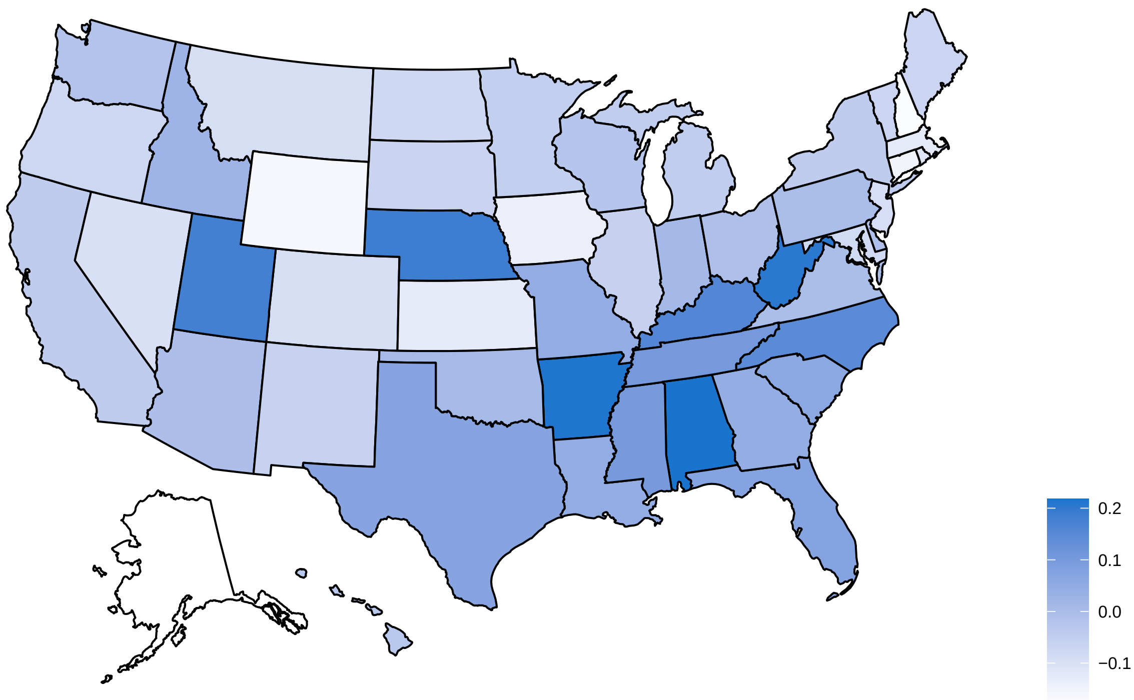

Figure 1: Normative Sexism by State

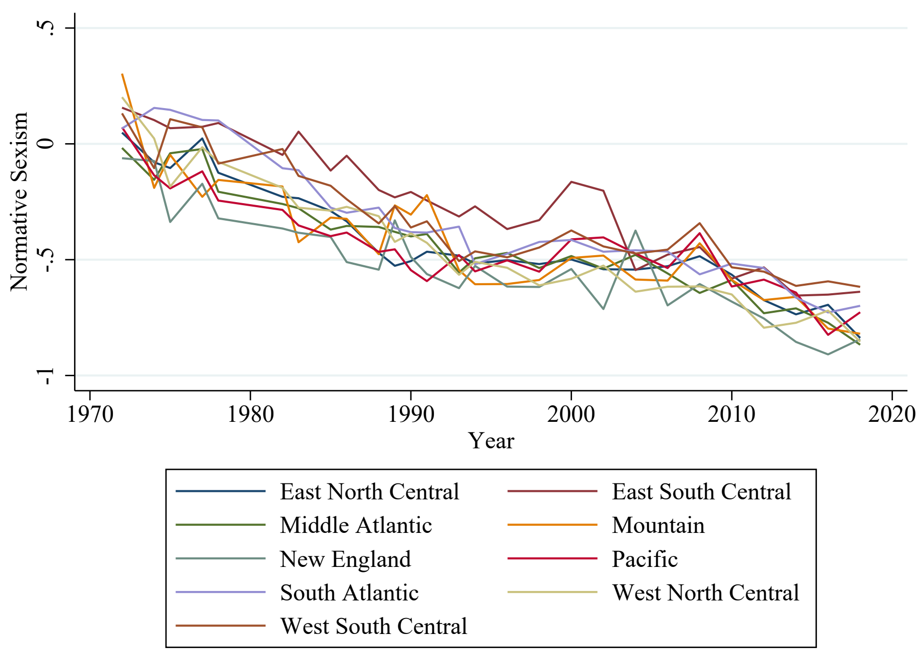

Figure 2: Average Prejudice Index Over Time by Census Division

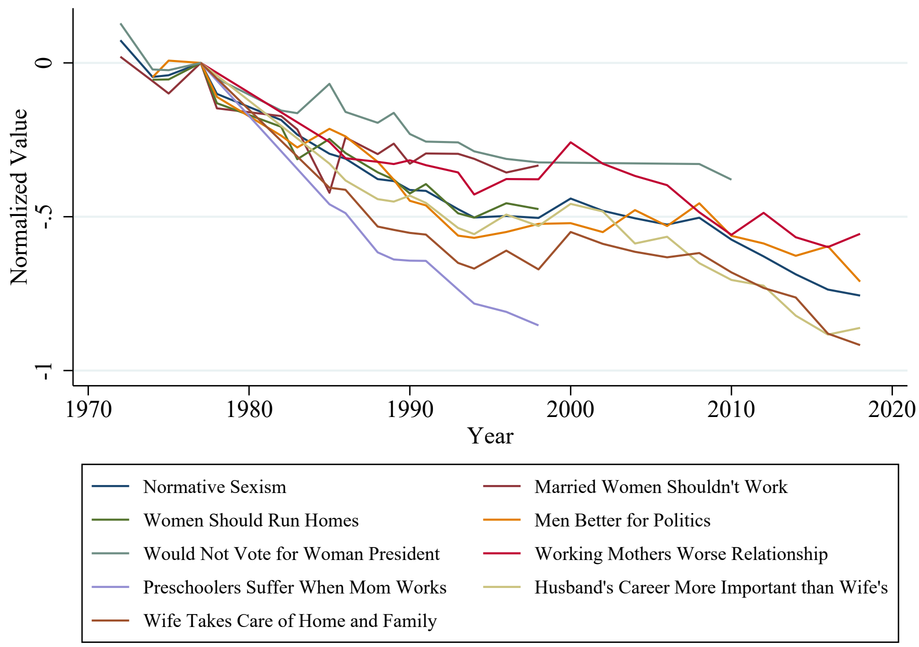

Figure 3: Responses to GSS Prejudice Questions Over Time

Figure 4: Normative Sexism - Men vs. Women

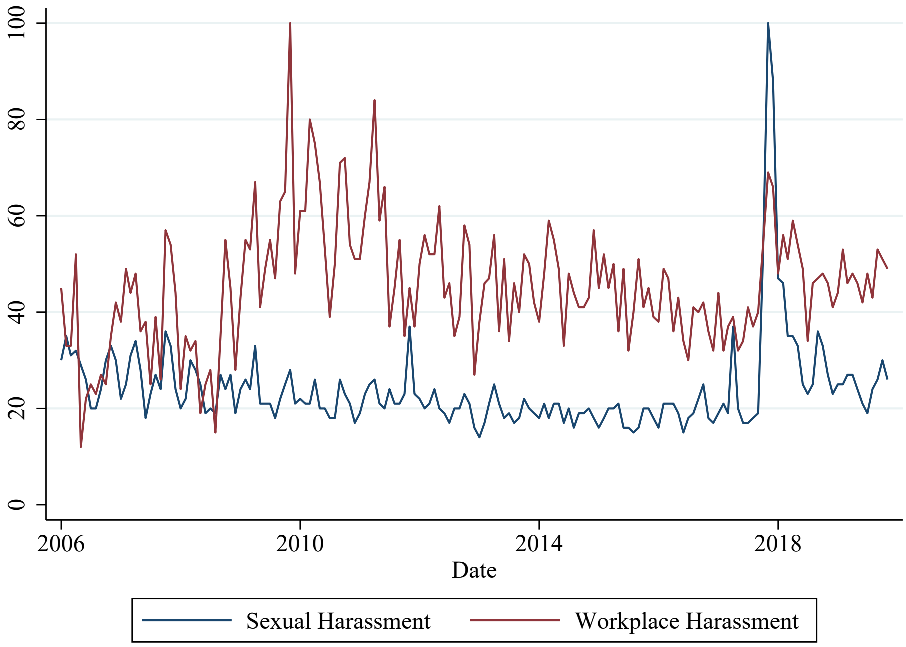

Figure 5: Reactionary Search Terms Over Time

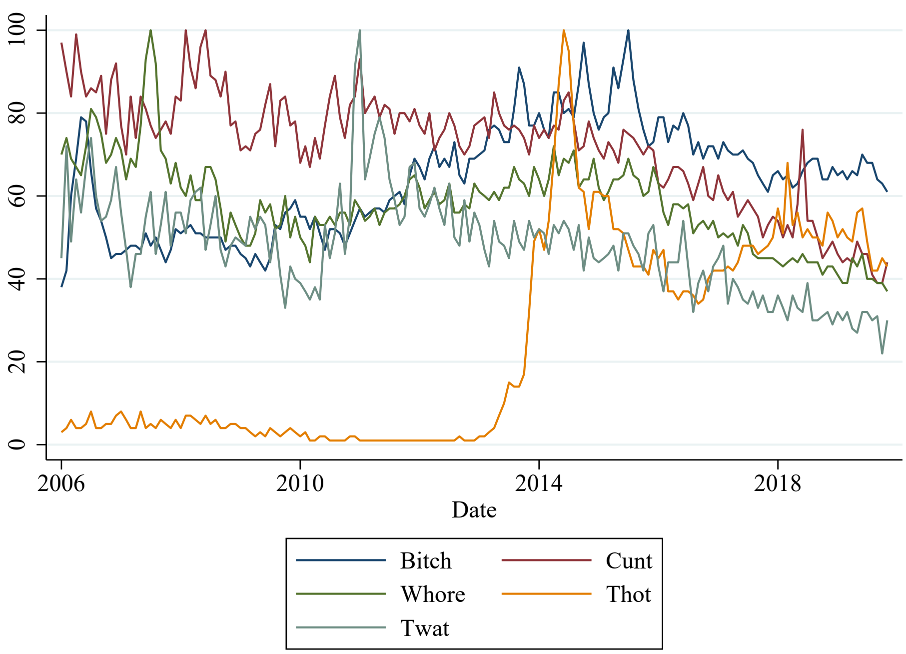

Figure 6: Derogatory Search Terms Over Time

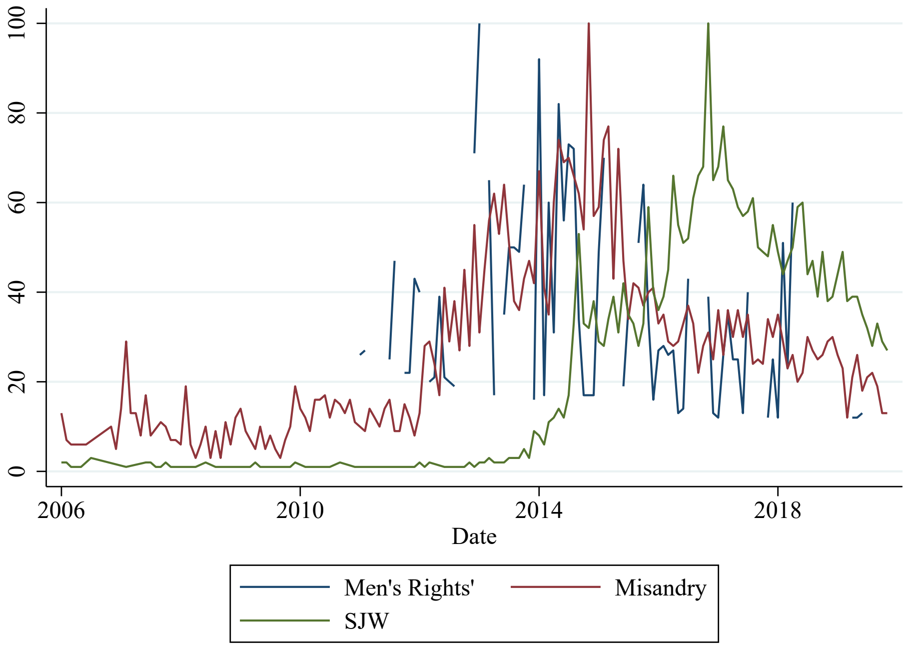

Figure 7: Manosphere Search Terms Over Time

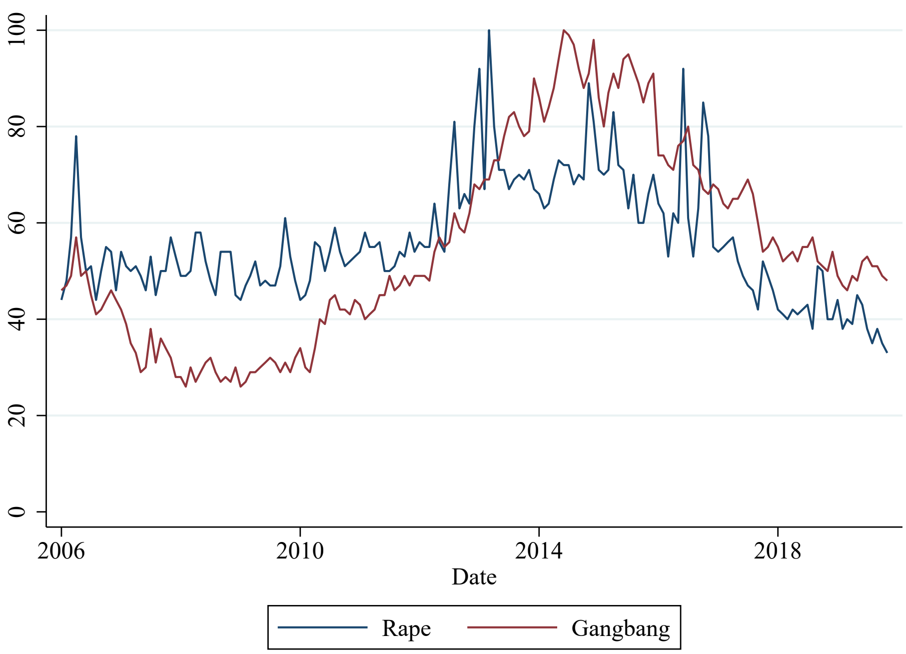

Figure 8: Violent Search Terms Over Time

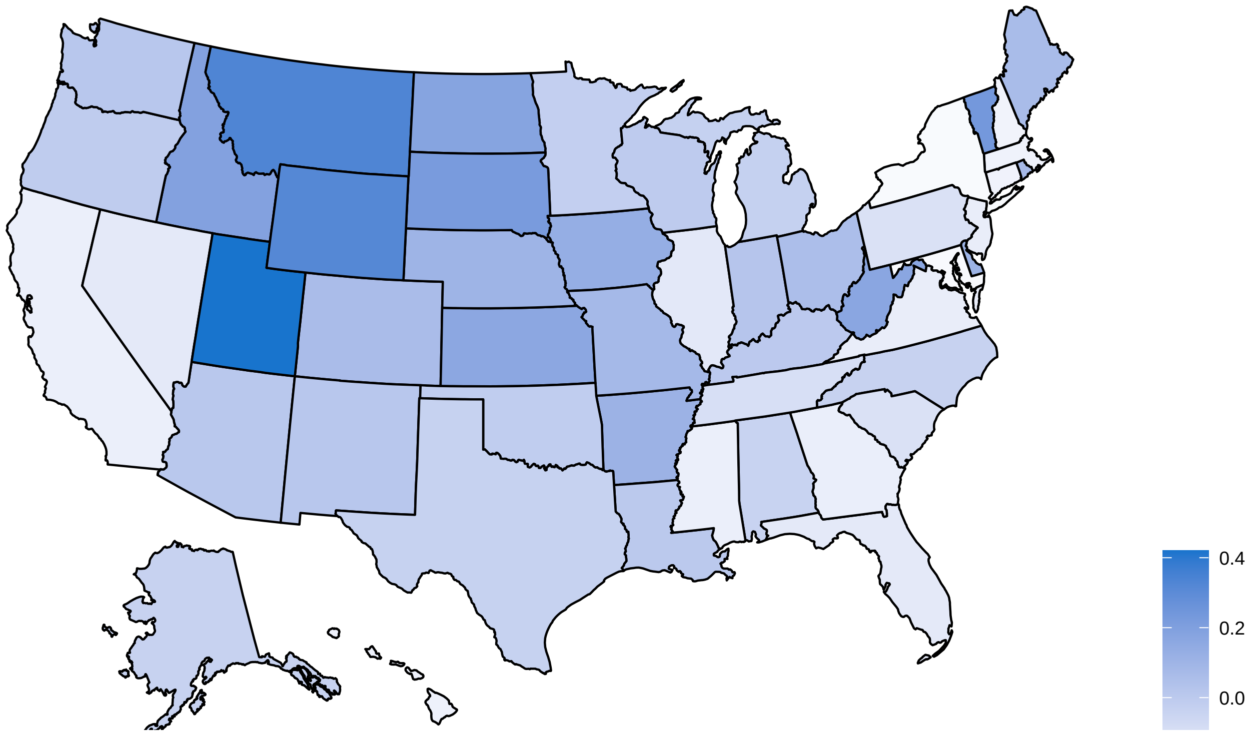

Figure 9: Index of Misogyny by State

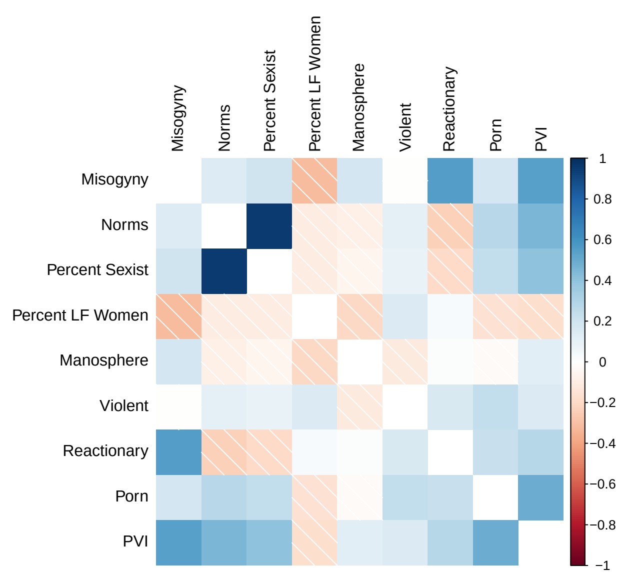

Figure 10: Correlation Between Various Measures

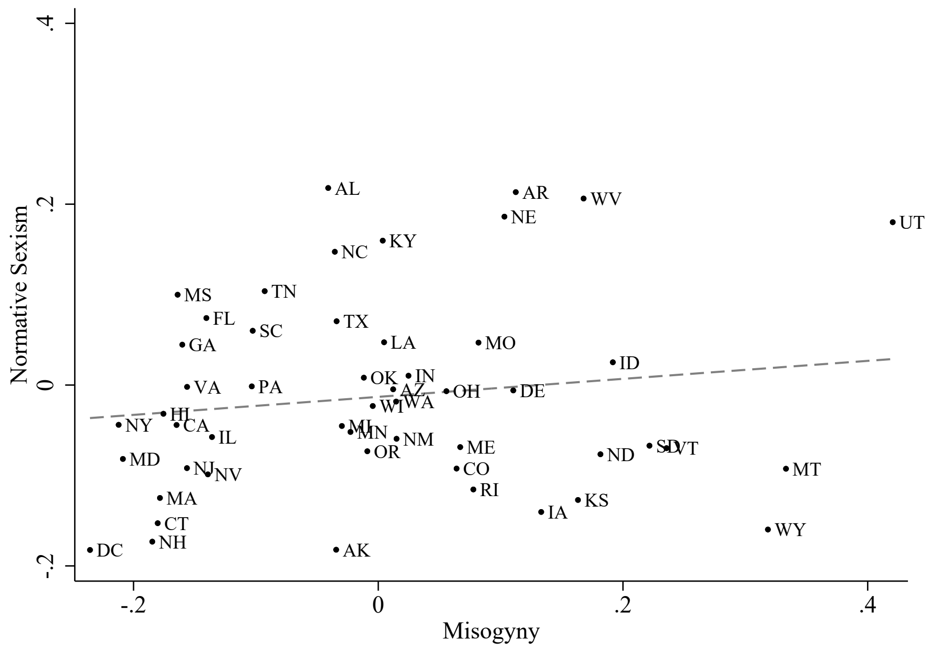

Figure 11: Misogyny vs. Normative Sexism Indices

References

Becker, G. S. (1957). *The economics of discrimination*. University of Chicago press.

Becker, G. S. (1985). Human capital, effort, and the sexual division of labor. *Journal of labor economics*, 3(1, Part 2):S33–S58.

Berinsky, A. J. (1999). The two faces of public opinion. *American Journal of Political Science*, pages 1209–1230.

Bianchi, S. M., Milkie, M. A., Sayer, L. C., and Robinson, J. P. (2000). Is anyone doing the housework? trends in the gender division of household labor. *Social forces*, 79(1):191–228.

Black, D. A. (1995). Discrimination in an equilibrium search model. *Journal of labor Economics*, 13(2):309–334.

Blau, F. D. and Kahn, L. M. (2007). The gender pay gap: Have women gone as far as they can? *Academy of Management Perspectives*, 21(1):7–23.

Blau, F. D. and Kahn, L. M. (2017). The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature*, 55(3):789–865.

Charles, K. K. and Guryan, J. (2008). Prejudice and wages: an empirical assessment of becker's the economics of discrimination. *Journal of political economy*, 116(5):773–809.

Charles, K. K., Guryan, J., and Pan, J. (2018). The effects of sexism on american women: The role of norms vs. discrimination. Technical report, National Bureau of Economic Research.

Farrell, T., Fernandez, M., Novotny, J., and Alani, H. (2019). Exploring misogyny across the manosphere in reddit.

Mincer, J. and Polachek, S. (1974). Family investments in human capital: Earnings of women. *Journal of political Economy*, 82(2, Part 2):S76–S108.

Neumark, D. (1988). Employers' discriminatory behavior and the estimation of wage discrimination. *Journal of Human resources*, pages 279–295.

Stephens-Davidowitz, S. (2014). The cost of racial animus on a black candidate: Evidence using google search data. *Journal of Public Economics*, 118:26–40.