NBER WORKING PAPER SERIES

DOES THE GENDER WAGE GAP ACTUALLY REFLECT TASTE DISCRIMINATION AGAINST WOMEN?

Molly Maloney
David Neumark

Working Paper 33405
http://www.nber.org/papers/w33405

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2025

This paper was originally sole-authored work in Maloney's dissertation. Neumark became a co-author subsequently to help finalize the paper and move it towards publication. We thank seminar participants at University of California, Irvine for their helpful suggestions. All errors are our own. We use GSS Sensitive Data Files for some of the analyses in this paper. These were obtained under contractual arrangements with the GSS in order to protect the anonymity of respondents. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Molly Maloney and David Neumark. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Does the Gender Wage Gap Actually Reflect Taste Discrimination Against Women?
Molly Maloney and David Neumark
NBER Working Paper No. 33405
January 2025
JEL No. J16, J29, J7

ABSTRACT

One explanation of the gender wage gap is taste discrimination, as in Becker (1957). We test for taste discrimination by constructing a novel measure of misogyny using Google Trends data on searches that include derogatory terms for women. We find—surprisingly, in our view—that misogyny is an economically meaningful and statistically significant predictor of the wage gap. We also test more explicit implications of taste discrimination. The data are inconsistent with the Becker taste discrimination model, based on the tests used in Charles and Guryan (2008). But the data are consistent with the effects of taste discrimination against women in search models (Black, 1995), in which discrimination on the part of even a small group of misogynists can result in a wage gap.

Molly Maloney
University of California, Irvine
molly.maloney@brattle.com

David Neumark
Department of Economics
University of California, Irvine
3151 Social Science Plaza
Irvine, CA 92697
and NBER
dneumark@uci.edu

1 Introduction

The wage gap between men and women in the United States is long established, and cannot be fully explained by observable wage determinants such as education and experience (Blau & Kahn, 2017). As a result, it is common to attribute all or part of the unexplained portion of the wage gap to discrimination. The workhorse model for explaining a discriminatory gap in wages is the taste discrimination model of Becker (1957). Others argue that statistical discrimination is a more plausible explanation, in part because under competitive market conditions equilibrium wage differences owing to taste discrimination may not persist (Arrow, 1973; Phelps, 1972).

In addition, there is a legitimate question as to whether the taste discrimination model of Becker fits the gender context. Becker developed his model of wage discrimination in the context of race, allowing employers' 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 its relevance to the black-white wage gap is clear and has been empirically validated. Charles and Guryan (2008) lay out several testable implications of Becker's model and evaluate how well these fit the data on the wage gap between black and white workers, finding evidence supporting Becker's model.

However, the role of distaste for women with regard to the gender wage gap is less obvious. First, most men are romantically entwined with women over a large portion of their lives, and men and women have always been joined by reproductive necessity. Accordingly, men and women have never been truly segregated from one another and, in fact, have had to coordinate with one another at least within families. Moreover, employers interact with their mothers, wives, and sisters, so it is not obvious why workplace interactions—and the role of employer tastes—would be substantively different. This type of interaction and cooperation makes it hard to believe that the gender wage gap is driven by disutility from interacting with women, as in the taste discrimination model.

On the other hand, there is apparently strong distaste for or dislike of women among some, as evidenced, for example, by the emergence of "incels" or "involuntary celibates," an internet subculture organized around misogynistic ideas (Tranchese & Sugiura, 2021; S. Zimmerman, 2024). It seems likely that these are fringe misogynystic tastes held by a small minority. In the standard Becker model, discriminatory tastes held by a small group do not affect equilibrium wages, and thus

do not lead to a discriminatory wage gap. But in a search model, as shown by Black (1995), a small group with discriminatory tastes shifts the offer distribution of wages in the labor market, making it perhaps more plausible that taste discrimination against women, even if held by only a small minority, can contribute to the gender wage gap. Hence, in studying the evidence on misogyny and the gender gap in wages, we consider the implications of both the Becker model and the Black model.

In an unpublished paper, Charles et al. (2018) extend their tests of the Becker taste discrimination model from race to gender, asking whether gender prejudice can explain the gender gap in wages in a manner consistent with the Becker model. But they measure gender prejudice using only General Social Survey data, which—as we argue—capture norms more than distaste. While gender norms may influence the gender gap in wages—for example by influencing specialization in market vs. home production (Mincer & Polachek, 1974)—it is not clear that norms play the same role as distaste for women as characterized by the Becker model.

To measure the role of distaste for women (or misogyny) in determining the wage gap, we construct a novel index of misogyny based on Google Trends data on search terms that are derogatory toward women. We find that the gender wage gap is increasing in misogyny and that this relationship is statistically significant, even after accounting for education and experience, and also controlling for the strength of gender norms. We then use this measure to test the empirical implications of two different models of taste discrimination—the Becker model of discrimination and Black's search model.[1] We find that the data are consistent with the implications of taste discrimination in a search framework.

We also find that misogyny is strongly related to other labor market (and related) outcomes including the proportion of women who remain unmarried between the ages of 20 and 40 and the average age at which women have their first child, as well as full-time employment and college completion. Because we estimate the gender wage gap conditional on employment and observables like education, the fact that misogyny is an important predictor of these variables that themselves affect earnings means that the role of misogyny on the residual gender wage gap may understate the role of misogyny in generating overall gender differences in earnings.

[1] In doing this, we draw on the work of Charles and Guryan (2008) and Burn (2017).

2 Background

Blau and Kahn (2017) document the evolution of the raw gender wage gap over time. Prior to 1980, they show that the ratio of women's to men's wages hovers around .6. During the 1980's and 1990's the ratio improved until it was just under .8. In more recent years, however, progress has stagnated. Efforts to explain changes in the gender wage gap to a large extent focus on the role of differences in observable characteristics between women and men, especially education and work experience.[2] Blau and Kahn (2017) note that the education gap between men and women has actually reversed in recent decades, with women's average education exceeding men's since 2011. Women have also 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, compared to a gap of almost 7 years in 1981.[3]

Even accounting for differences in human capital accumulation and other observable differences between women and men, a wage gap persists. This residual wage gap is the portion we consider to be potentially due to discrimination. Of course, this residual wage gap need not solely reflect discrimination. Unobserved differences between men and women that contribute to productivity can also influence the residual wage gap (Hellerstein et al., 1999). If these unobserved differences, on average, reflect higher productivity among men, the residual wage gap will overstate the effect of discrimination on wages. Conversely, as Blau and Kahn (2007) point out, discrimination could lower women's return to human capital investment or push them into lower-paying occupations, in which case the residual wage gap (depending on what controls are included) could understate discrimination. Other non-discriminatory explanations that may be partly responsible for the gender wage gap include bargaining, competitiveness, turnover, and compensating differentials and endogenous job design (Goldin, 2014).

Despite these potential limitations, estimation of residual wage gaps is common and these gaps are frequently interpreted as informative about discrimination. We follow this literature in interpreting the residual wage gap as potentially reflecting discrimination—possibly taste discrimi-

[2] See work referenced in Blau and Kahn (2017).

[3] Experience gaps may reflect differences in the use of time outside the workplace (Becker, 1985; Mincer & Polachek, 1974). Time-use surveys have shown that women participate in non-market labor like household chores and child care responsibilities at higher rates than men (Bianchi et al., 2000), and this gender gap in household productivity continues to persist although it has shrunk (see, e.g., Sayer, 2016). Thus, the decline in the gap in labor market experience could also be reflected in a declining gap in time spent in non-market production.

nation. Our core approach and question, however, is to try to directly measure distaste for women (misogyny), to see if it can account for variation in the residual gender wage gap.

3 Models of Discrimination

Our empirical analysis is framed in terms of testing whether taste discrimination can help explain the residual wage gap between women and men. Statistical discrimination is an alternative model. In the models of Arrow (1973) and Phelps (1972), employers try to infer the productivity of applicants based on a noisy signal of productivity (e.g., education), as well as beliefs about the average productivity of men and women. Because of the weight employers place on the group average, any female shortfall in expected productivity (even if inaccurate) will lead to lower wages for women than men.[4]

Our sense is that, for the reasons outlined in the Introduction—in particular, the view that taste discrimination against women is less tenable than, say, taste discrimination against blacks—labor economists view statistical discrimination as more likely to account for unexplained gender differences in labor market outcomes. However, our goal in this paper is to provide evidence on this question, by implementing an explicit test of the role of taste discrimination against women (misogyny) in the residual gender wage gap. Our core approach is to construct a measure of misogyny at the state level to see whether it helps explain variation in the residual wage gap by state. We then go on to test specific implications of two alternative models of taste discrimination to provide further evidence on whether taste discrimination helps explain the residual wage gap.

3.1 The Becker Model of Taste-Based Discrimination

Consider first the taste-based discrimination model described by Becker (1957), with a perfectly competitive labor market and employers who are prejudiced against women, denoted F (men are denoted M). Employer utility V_i increases with profit π_i and decreases with the interaction between

[4] Aigner and Cain (1977) point out that if employers correctly identify the average productivity of men and women (which we would expect them to do in the long run), this practice does not constitute economic discrimination against a group in the typical sense since expected productivity equals true productivity on average. But this does not mean equally productive men and women will be paid the same, if the reliability of the signal of productivity differs.

the employer's degree of prejudice d_i and the number of women he employs:

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

Assuming women and men 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 and women, respectively. The utility maximizing choices of F and M at an interior solution are given by

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

The short-run equilibrium involves an almost completely segregated labor force, with women working for the less discriminatory employers and men working for the more discriminatory employers. Assuming the distribution of prejudice is fairly smooth, there is an employer who is indifferent between hiring women and men at the prevailing wages. We refer to this employer as the "marginal discriminator" and his degree of prejudice as the "marginal" level of prejudice d^*. More prejudiced employers ($d_i > d*$) hire only men at wage $w_M = f'$, and less prejudiced employers ($d_i < d*$) 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.

In contrast, the average level of prejudice will not directly affect the wage gap. To see why, suppose, for example, that average prejudice is very high because there is a lot of prejudice at the top of the prejudice distribution. Since the highly-prejudiced employers do not hire women, their prejudice is irrelevant to the wage gap. Conversely, there could be some unprejudiced employers who would hire women even if $w_f = w_m$. But if there is not a sufficient number to employ all women, the marginal discriminator still generates a wage gap. Thus, the Becker model predicts a relationship between the marginal level of prejudice and the wage gap, but not one between the average level of prejudice and the wage gap.[5]

The model also implies that prejudice increases will affect the wage gap differently depending on where they come from in the prejudice distribution. For example, suppose prejudice increases at the top of the prejudice distribution. Any increase above the marginal prejudice level should

[5]Note that if the marginal level of prejudice is very close to the average, they may be difficult to disentangle empirically.

have no effect on the wage gap. If, on the other hand, prejudice increases at the bottom of the prejudice distribution, this is likely to push up the prejudice level of the marginal discriminator, and the wage gap will widen.

Finally, consider the impact of increasing the number of women in the labor market. All else equal, as additional women enter, 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. These are the ideas that underlie the Charles and Guryan (2008) test of the applicability of the Becker model to race differences in wages (with substitution of references to blacks and whites rather than women and men).

Note that under this framework, a small contingent of very prejudiced employers should not contribute to the wage gap as long as the number of women is not so high that some will be employed by these very prejudiced employers. As we might expect the number of misogynysts to be low, misogyny is unlikely to be able to account for the residual wage gap based on the Becker model.

3.2 Black's Search Model of Taste-Based Discrimination

Black (1995) offers an alternative model of how taste discrimination matters, by embedding discriminatory tastes in a search model with frictions. In this model, men (who are not discriminated against) and women (who are) enter the labor market and search sequentially for a job. In each period of search, applicants pay a fixed cost and are matched with a potential employer who chooses whether or not to hire him or her. Prejudiced 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. Note that in contrast to the Becker model, in the search framework a small contingent of very prejudiced employers does contribute to wage gaps, as they generate longer expected search times for women.

In this model, the key drivers of the wage gap are the share of employers that are discriminatory and the share of women in the workforce. As the share of discriminatory employers increases, women face longer searches, on average, thus increasing the bargaining power of non-discriminatory firms and resulting in lower wages for women.

As more women enter the work force, non-discriminatory firms are able to hire more low-cost (female) employees and the profits of these firms increase relative to discriminatory firms. As non-discriminatory firms become more profitable, some discriminatory firms are driven from the market, reducing the share of such firms. This reduces search costs for women and increases their reservation wage. Now, non-discriminatory firms must offer them higher wages to induce them to accept a job offer, and the wage gap shrinks. Note that this prediction about the relationship between the wage gap and the proportion of women in the labor market is the reverse of Becker's, allowing us to sharply distinguish between the two models.

In this search framework, employers are either prejudiced or not and the intensity of discrimination (measured by d in the Becker model) is not a relevant factor. Despite this, we can still expect a relationship between the average level of prejudice and the wage gap since average prejudice increases systematically with the proportion of prejudiced employers.

4 Empirical Approach and Measurement

Our empirical analysis tests whether taste discrimination helps explain the residual gender gap in wages. Our novel contribution is the construction of a measure of misogyny. Because this is available at the state level, our empirical analysis focuses on cross-state variation in the residual gender wage gap. To test the different versions of the taste discrimination model, we construct measures of the proportion of women in the labor market as well as the proportion of prejudiced employers in the labor market and other moments of the prejudice distribution. We also contrast results for the role of misogyny vs. gender norms in explaining the residual gender wage gap, although this analysis does not directly relate to testing the taste discrimination model.

4.1 Residual gender gap and basic approach

We first estimate the residual gender wage gap by state. We regress log wages on education, a quadratic in potential experience, and gender-specific year and state dummy variables, using data from the CPS May Outgoing Rotation Group between 1979 and 2019 and weighting by the Census earnings weights.[6] The coefficients on the state-by-female indicators give the residual wage gap in

[6]Specifically, we use the cleaned and standardized version of this data provided by the Center for Economic and Policy Research.

each state. We also use these data to estimate the proportion of women in the workforce, which we compute as the average ratio of women workers to all workers in the state across years.

Following Charles and Guryan (2008), we estimate the distribution of employer prejudice under the assumption that it mirrors that of the general population. For now, suppose we have an index of individual prejudice that varies within each state (we discuss the data we use below). We can take, e.g., the 10th and 90th percentiles of the prejudice distribution as measures of its lower and upper tails, respectively. We then approximate the marginal prejudice level by computing the p^{th} percentile of the prejudice distribution where p is the percent of women in a state's workforce. This reflects Becker's definition of the marginal discriminator as the least prejudiced employer who hires women. We use the same index of prejudice to compute the average level of prejudice in the state.

Then, in a second stage, we use estimates of the residual wage gap as the dependent variable in a regression on the proportion of women in the labor market and various points of the prejudice distribution (mean, marginal, and each tail), weighted by the precision with which we estimate the residual wage gap in the first stage. We then evaluate the results to see if they are consistent with the predictions of either the Becker or Black model.

To reiterate, under the Becker model we expect the wage gap to be increasing in the proportion of women in the labor force. We can also test for evidence that the marginal discriminator strongly influences the wage gap by examining which moments of the prejudice distribution influence the wage gap more strongly. If the marginal discriminator determines the wage gap, as predicted by the Becker model, we expect the wage gap to be systematically related to the marginal, rather than the average level of prejudice. We would also expect prejudice at the bottom of the prejudice distribution to be a better predictor of the wage gap than prejudice at the top of the distribution.

Under Black's model (2005), on the other hand, we expect the wage gap to be decreasing, rather than increasing, in the proportion of the labor force made up by women. We also expect the wage gap to be increasing in the proportion of prejudiced employers, as the Becker model predicts; so it is the evidence on the effect of the proportion female that is dispositive between these two models.[7]

[7]To flesh out this argument even more fully, recall that the Becker model with the form of discriminatory tastes given above implies that there is near-perfect segregation in the labor market, depending on whether an employer has weaker or stronger discriminatory tastes than the marginal discriminator. Neumark (1988) suggests an alternative model where an employer's disutility depends on the ratio of, in our case, female to male workers. In this formulation of the model, there is not such strong segregation, but less discriminatory employers employ relatively more women, and vice versa. As a result, the Charles and Guryan test for the influence of a marginal discriminator is no longer informative, but the test for the negative influence of the proportion of women in the workforce still holds, and hence

The Black model does not yield predictions about various points of the prejudice distribution, as the wage gap does not rely on a marginal discriminator. But we still expect the wage gap to be increasing in average prejudice levels. To be clear, though, our core goal is not to distinguish between the two models, but to see whether the evidence indicates that taste discrimination—whether as modeled by Becker or by Black—helps explain the residual gender wage gap.

4.2 Other contexts vs. gender discrimination and misogyny

Other work has demonstrated the relevance of both models to residual wage gaps between different groups. Charles and Guryan (2008) use a similar procedure to empirically validate the Becker model's predictions in the context of the wage gap between black and white workers. They find that the black-white wage gap is increasing in the proportion of black individuals in the labor force, that the marginal level of prejudice is more strongly related to the wage gap than the average prejudice level, and that the lower tail of the prejudice distribution is a strong predictor of the wage gap whereas the upper tail is not. They therefore conclude that the Becker model describes the racial wage gap well. Burn (2019), on the other hand, performs some of these tests on the wage gap between gay and straight men and finds evidence that Black's search model is more consistent with data on the gay-straight wage gap.

In our view, the hypothesis that the Becker model of taste discrimination helps explain race discrimination is a priori more plausible than it is for gender discrimination. With regard to gender discrimination, it seems reasonable to believe that prejudicial views are held by a small part of the population, and hence that discriminatory employers are not marginal, in which case the search model may better fit the data. And we suspect that similar considerations may explain the evidence for discrimination against gays—although in this case there may be more prejudice, but the share gay in the workforce is low, again making it less likely that a wage gap could be driven by a marginal discriminatory employer.

---

again distinguishes the Becker model (reformulated) from the Black model.

5 Measuring Distaste for Women

We do not think that the GSS questions on attitudes towards women can be used to construct an index of misogyny against women. These questions are listed in Table 1. Perusing these questions suggests that they pertain more to gender norms than to misogyny (where misogyny is the equivalent to racism, but applied to women). In contrast, for example, the questions about race used by Charles and Guryan (2008) are more related to distaste, such as: "How strongly would you object if a family member brought a black friend home for dinner?" and "Do you think blacks should have as good a chance as anyone to get any kind of job, or do you think white people should have the first chance at any kind of job?"

We therefore instead construct our own measure of distaste for women, using Google Trends data on search interest for misogynystic terms. Using Google Trends data has some important advantages over survey-based measures of misogyny (or sexism more generally). 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). For example, in his paper on the role of racial prejudice on vote shares in 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 interprets this as suggesting 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.

To construct a misogyny measure, we need to identify appropriate search terms to include. Following Stephens-Davidowitz (2014), we initially consider slurs against women. There is some research on which to base this. In a paper aimed at describing misogynistic language used on Reddit's "Manosphere" (a portion of the internet devoted to promoting misogyny), Farrell et al.

(2019) identify a small set of words (three) that are derogatory towards women as part of a group including "violent verbs, and slurs that are not immediately racist or homophobic."[8] They find that this group of words is rampant across the manosphere, and is actually the most prominent of all categories they examine. We use these three words as well as a few synonyms for them to construct our index of misogyny. We refer to this group of terms as "Derogatory" throughout this section. We also collect data on search interest for three other categories of words that 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 (e.g., rape). The "Manosphere" group includes terms often used within the manosphere (e.g., 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 2 lists all the words in the Reactionary and Manosphere categories.[9] Note that Google censors search interest when it falls below some threshold of minimum interest. We exclude search terms that are censored too frequently. For example, "femoid" was considered for inclusion in the group of "Manosphere" words, but was excluded because it is censored by Google in all but a handful of states.

Tables 3 and 4 list the top 20 related search queries for each of the terms related to "Reactionary" and the "Manosphere," as examples.[10] These tables 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 (including more offensive terms not shown here) pertain to song lyrics or searches for definitions of the terms. This is not overly concerning, as 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

[8] In this version of our paper, we elide the actual derogatory terms from the text, as well as tables containing them, as they may be offensive to some. For readers who want to see the uncensored paper that provides all of the details involved in constructing our misogyny index, please see: https://sites.socsci.uci.edu/~dneumark/research.html (with the forewarning that some highly offensive language is presented in that version).

[9] We say more about some of these words below. The corresponding table in the full paper available at the url listed above also includes words in the "Derogatory" and "Violent" categories. Note some terms in the latter category appear to be from searches related mainly to pornography; these are omitted from the construction of our misogyny index. In contrast, most of the "Derogatory" search terms clearly reflect anti-women sentiment, and we use interest in these terms to construct our index of misogyny.

[10] The tables for searches related to the "Derogatory" and "Violent" terms are reported in the full paper available at the url listed above, because of particularly offensive language.

to searches for "god" during the time period he examines is "god of war," a popular video game released during the 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."

Figures 1 and 2 show interest over time in the words related to "Reactionary" and the "Manosphere" listed in Table 2 across the US between 2006 and 2020 for each category.[11] These are normalized by the maximum rate of searches on the topic in a given period and scaled by 100. Interest in the "Derogatory," "Manosphere," and "Violent" groups of words have trended downward in recent years.

In contrast, search interest in the terms sexual harassment and workplace harassment, seen in Figure 1, has been much more stable without an obvious upward or downward trend, but with two notable spikes. 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. The second big spike in both "workplace harassment" and "sexual harassment" occurs toward the end of 2017. This is likely due to the "#MeToo" movement, which exploded in late 2017 after sexual assault allegations against Harvey Weinstein were brought to the public's attention. These spikes likely reflect the effect of external media buzz rather than personal experience. Moreover, Table 3 indicates that much of the search interest in these words reflects searches for laws and policies about the topic. Due to concerns that search interest in these terms may reflect media reports or legal interest rather than personal experiences, we do not include them in our measure of misogyny.

Turning to the group of "Manosphere" words shown in Figure 2, note the spike in 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. While this term is commonly used in the manosphere, it is also used by the political right, and may reflect political ideology rather than activity in the manosphere. The use of the word misandry may also

[11] Note that we use CPS data from 1979 onward. This is intended to align with the availability of the GSS data, which we also incorporate. The state wage gaps are very similar whether we compute them using the CPS data since 1979 or restrict to 2006 onward. The correlation coefficient between the state wage gaps computed using these two subsets of data is 0.8698 (p-value < 0.0001). The corresponding figures for searches related to "Derogatory" and "Violent" terms are reported in the full paper provided at the url listed above.

be hard to interpret because it was adopted and used ironically by members of feminist groups. This usage was highlighted in articles published by Slate, Time Magazine, and The Guardian in August and December of 2014 (Begley, 2014; Hess, 2014; J. Zimmerman, 2014). And we do see a corresponding spike in search interest for "misandry" in late 2014. Such uses of these words make it hard to interpret them as reflective of misogyny, hence we exclude these words out of concern that they may reflect political ideology instead of misogyny.

In addition to interest over time, Google Trends also provides data on the rate of Google search interest for a given phrase at the US state or designated market area (DMA) level within a specified period of time, based on a random sample of all searches.[12] This regional variation is central to our empirical strategy. Google Trends reports search interest in a given region relative to the maximum rate of searches for the term experienced in any region during that time period, scaled by 100. For example, if Ohio has the highest rate of searches for the term "weather" across the U.S., 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. We can obtain these for various periods of time, but since each is normalized relative to the maximum in that time period, they cannot be readily compared across time.

To construct a measure of misogyny or distaste for women, we take each term in the group of "Derogatory" words and collect relative search interest in that term for each region (state or DMA) over the period of a year. We do this for both the state and DMA levels. Because the Google search index is a sample, we collect 10 samples for each year and search term and average over all samples. We normalize the interest in each word by its mean and standard deviation in 2016. For each observation, we take the average over all five "Derogatory" search terms, regress it on year dummies, and use the residuals as our index of misogyny.

To get the average level of misogyny in a state, we simply compute the mean of this index, constructed on state-level data. Figure 3 displays information on the average level of misogyny by state. To obtain other parts of the misogyny distribution, we need additional within-state variation, so we compute this index of misogyny using DMA level data. We then compute the 10th and 90th percentiles of the DMA-level index for each state, weighting each DMA by its population.[13] We

[12] DMAs are geographic areas used by Nielsen for for television ratings. See https://www.nielsen.com/dma-regions.

[13] We compute DMA populations by summing the populations of all counties within a given DMA, using county populations by year provided by the U.S. Census Bureau, and a mapping between Google trend DMAs and U.S.

also estimate 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 full-time employed, again weighting by DMA population.

Our misogyny measure, which captures the use of vulgar and disparaging terms for women, is designed to reflect true misogyny, not unlike the kind of animus we associate with explicit racism.[14] But as discussed in the Introduction, it is worth considering whether or not we really believe this kind of distaste makes sense in the context of gender where men and women have always been linked within (and outside of) families. The majority of men fall in love with, marry, and reproduce with women and most were raised by at least one female guardian. While the presence of animus toward women certainly exists among fringe individuals, it is hard to imagine that enough men exhibit it to induce a wage gap under Becker's (1957) model of discrimination, because we would not expect men at the top of the prejudice distribution to influence the gender wage gap in this framework. However, a small group of misogynistic men can influcence the wage gap in the Black (2005) model.

Consider how this dynamic differs from that of racial integration between whites and blacks. Though racial segregation has declined over time, integration began less than 100 years ago in many parts of the country, and De la Roca et al. (2014) show that segregation between between whites and blacks remains quite high even in recent years. As such, it is much easier to imagine racial animus playing a role in reducing wages of black workers.

Given this difference, it may be more reasonable to expect women's wages (and other labor market outcomes) to be influenced by gender norms or more subtle forms of prejudice rather than the overt misogyny captured by our measure. To explore this question, we also study the influence of normative sexism on women's labor market outcomes.

counties obtained via Kaggle (Pastor, 2021). In some cases, DMAs span multiple states, and we assign the DMA to the state with fewer uncensored words in it (due to inadequate search interest). For example, if our search terms were censored more often in AL than FL, the "Mobile AL-Pensacola (Ft. Walton Beach) FL" DMA would be assigned to FL.

[14]Some of the words in the "Derogatory" group are slang for female genitalia, suggesting it might be reasonable to worry that the index reflects porn interest rather than sexism. While there is some correlation between searches for porn and our index of misogyny, this correlation is weak and statistically insignificant. Appendix Table A.1 provides the correlation between average state misogyny and other measures including Google search interest in the word "porn." Search interest in the the word porn includes searches for sites like "Porn Hub" or other sites with porn in the name, so it should be correlated with traffic to porn sites even if it is an imperfect measure.

6 Normative Sexism

Now consider the role of gender norms or what we call normative sexism in the labor market. Gender norms may directly influence a woman's decisions about human capital accumulation and whether to continue working after having children. Norms about who is expected to provide for a family could influence wage negotiations both from the employer's perspective as well as from that of a wife conscientious about out-earning her husband. In fact, Fortin (2005, 2008) has shown the relevance of norms for female labor market outcomes including labor force participation in the U.S. and the gender pay gap across countries.

Thus, to isolate the effect of misogyny—as well as to provide evidence of the impact of gender norms—we supplement our index of misogyny with a measure of normative sexism using restricted-access GSS responses to questions about the role of men and women in society. We use the questions listed in Table 1. They cover opinions about whether a woman should work outside the home or focus on caring for her family, whether children of working mothers are worse off, and whether men are better leaders than women. As argued above, rather than identifying misogyny or hatred toward women, these questions identify opinions about the appropriate domains of men and women.

Following Charles and Guryan (2008), we recode responses to these questions so that answers reflecting more traditional gender norms have higher numeric value and then normalize responses to each question by the mean and standard deviation in 1977.[15] For each question k, individual i, and year t, the normalized individual scores are given by

$$\tilde{d}_{it}^{k} = \frac{d_{it}^{k} - E[d_{i,77}]}{\sqrt{Var(d_{i,1977}^{k})}}.$$

We then average over all questions asked in a given year, computing

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

where K_t is the number of these questions asked in year t. Finally, we regress D_{it} on a full set of year dummies to get $\tilde{D}_{it}$ which we aggregate by state to construct a normative sexism measure

[15] We normalize by mean and standard deviation in 1977 because this is one of the few years in which all questions are asked.

comparable to the misogyny measure we described earlier.

Figure 4 shows the trends in average normative sexism for each census division over time. There is a general decrease in the rigidity of gender norms over time across all regions. Figure 5 depicts this index and the questions used to construct it, across the U.S. over time. Respondents give less sexist 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 sensitive to exactly how a question is phrased, and highlights one concern with survey-based measures like those in the GSS. Interestingly, sexist responses to the question about voting for a female president have declined much more slowly than those of other questions.

Table 5 gives the average responses to the GSS questions about women as well as the average normative sexism index by census division, after regressing on years to account for changes over time. Sexism of this kind is highest in the southeast of the country and lowest in New England. Most survey questions seem to exhibit similar geographic patterns as the aggregate index of normative sexism, though there are some differences from question to question. Figure 6 gives information on the variation in the aggregate index of normative sexism across states. As reflected in Table 5, normative sexism is highest in the southeast of the country, but there are a few states in the West with high levels of normative sexism including Utah and Nebraska. This map differs quite a bit from that of misogyny (Figure 3), where we see low misogyny levels across the Southeast.

6.1 Misogyny and Normative Sexism

Regional differences in the distributions of the measures of misogyny and normative sexism suggest that they are not simply two manifestations of the same type of sexism. Figure 7, which plots the index of gender norms against the index of misogyny, emphasizes this point. The two measures are positively correlated, but weakly so. Notice states like Alabama that have very strong gender norms, but below average levels of misogyny. Conversely, states like Montana have very high levels of misogyny despite weaker gender norms.

Figure 8 gives plots the average normative sexism by state constructed using only responses from men against responses of women. Note the strong positive correlation (p-value less than 0.001)

between normative sexism levels of men and women. This also highlights the difference between this kind of sexism and misogyny toward women. Normative sexism reflects societal expectations about the role of men and women that are shared across genders, rather than distaste for or animus toward women, which we would expect to be higher among men, on average (although we cannot tell from Google Trends data).

In our view, it is easy to see why misogyny and normative sexism (which one might call two different kinds of sexism) can be quite distinct. Consider a man who has grown up in a community in which women nearly always stay home and care for children. He may think this is the best way for a child to grow up and may want this for his children, but may not have any animosity toward women. In fact, he may value the role women play in raising children and caring for the home. He may have been taught to treat women with respect and the emphasis on gender roles may make him less likely to disparage women in the way that would be picked up by our misogyny measure. And note that women in such a community may share these values, as we observe empirically. On the other hand, imagine a man who has been unsuccessful with women romantically and outperformed by them in the workplace. Such a man may have developed a distaste for all women regardless of which roles they occupy since they have spurned or surpassed him in both public and private domains. While both men exhibit "sexism" broadly defined, the two sets of attitudes are quite different, may vary across people as well as regions, and may have different implications for wages and other labor market (and related) outcomes.

7 Analysis of the Effects of Misogyny and Normative Sexism

7.1 Analyses

The Becker and Black models have predictions for misogyny—i.e., a distaste for interacting with women—that likely pertain to wages only, since both models capture how discriminatory tastes affect equilibrium prices of female and male labor. There are no clear implications of these models for normative sexism, although we certainly might expect norms to affect the degree to which women work, invest in human capital, specialize in market vs. home production, etc.

Thus, in addition to empirically testing the predictions of the Black and Becker models with regard to misogyny and the gender wage gap, we also estimate the relationships between our

measures of misogyny and of normative sexism with other labor market and related outcomes (and wages), including the proportion of women aged 20 to 40 who remain unmarried, the average age at which women give birth to their first child, full-time employment, and college degree attainment, after conditioning on relevant observables.[16] We regress each outcome on our measures of misogyny and normative sexism, both individually and together, to compare the effects of the two kinds of tastes or attitudes.

These analyses are interesting for a number of reasons. First and foremost, our regressions with only the misogyny measure test for the role of taste discrimination in the gender wage gap, including testing whether the evidence is consistent with the Becker or the Black models. Second, our analysis of the relationship between the misogyny index and the other outcomes is interesting because if misogyny is reducing women's educational attainment and labor market participation (which could lead to more childbearing), then estimates of its impact on the residual wage gap will be biased towards zero. Third, when we include measures of both misogyny and normative sexism, the regressions provide evidence about the role of both misogyny and normative sexism in determining labor market outcomes for women. Examining the relative strengths of misogyny and gender norms is relevant to policy. If misogyny places an important role—especially in the gender wage gap—then traditional anti-discrimination tools that raise the cost to employers of discriminating can help mitigate gender differences. These anti-discrimination tools are also likely more effective at mitigating the effects of taste discrimination than statistical discrimination. In contrast, if normative sexism holds back women, but misogyny does not play a role, then policies implemented to prevent discrimination might have very little impact.

7.2 Results for Tests of Taste Discrimination

Table 6 reports our explicit tests of the taste discrimination model. Here, we regress the residual wage gap on various parts of the misogyny distribution as well as on the percent of employed women and the percent of the population that is sexist. We also include our normative sexism measure

[16] We use CDC data on natality between 2016 and 2019 to obtain the average age of women when they give birth to their first child. We take the average in each state after partialling out year effects. For the remaining outcomes, we use the same Census Outgoing Rotation Group data that we used to compute the residual wage gaps. We estimate college degree attainment controlling for age and year dummies. For the proportion of women never married between the ages of 20 and 40 and full-time employment rates, we control for education, age, and year indicators. In all cases, we weight estimates by the census earnings weight.

in some regressions and show that the effect of misogyny on the wage gap remains statistically significant, even after controlling for normative sexism. Note that a negative coefficient indicates an increase in the wage gap between women and men.

The estimate in Column (1) indicates that the residual wage gap is increasing in the misogyny index, and that the relationship is statistically significant. Column (2) shows that the normative sexism index is also negatively related to the gender gap in wages. Column (3) indicates that these indices are quite independent statistically, as the estimates change very little when both are included in the model.

The strength of the relationship between misogyny and the residual wage gap is striking. A one standard deviation increase in misogyny is associated with with a reduction in women's relative wages by about .0167 log points or approximately 4.3 percent of the average residual wage gap.[17] This effect is on the same order of magnitude (though slightly larger) as that of normative sexism. A one standard deviation increase in normative sexism is associated with a reduction in women's relative wages by about 0.0144 log points or approximately 3.7 percent of the mean residual wage gap. Despite the high levels of interaction and coordination between men and women in private domains, overt misogyny appears to be an important factor in determining the wage gap. So far, then, the results point to an important role for misogyny in generating the gender wage gap.

We next turn to more explicit tests of taste discrimination models. Based on the estimates in Columns (4)-(7), we do not find evidence that the wage gap is determined by a marginal discriminator as predicted by the Becker model, since marginal misogyny is a weaker predictor of the wage gap than is mean misogyny (the variable in the first row); in particular, it does not enter the model with a significant coefficient when the mean measure is included (Column (5)). Moreover, the wage gap is not more strongly related to the bottom of the misogyny distribution than to the top (Column (7)). Per the earlier discussion about the absence of the stark labor market segregation predicted by the Becker model, it is perhaps not surprising that the marginal discriminator plays little role in determining the gender wage gap (despite Charles and Guryan finding such evidence for race). What is more telling in rejecting the Becker model, however, is the evidence that the wage gap is decreasing, rather than increasing, in the percent of women employed full-time (Columns

[17]The standard deviation of the state-average misogyny level is equal to 0.154. The average residual wage gap is -0.389.

(6) and (8)).[18] This is inconsistent with Becker's predictions, even when we allow for labor market integration as Neumark (1988) does. On the other hand, the negative relationship between the percent of employed women and the gender wage gap is consistent with Black's search model of discrimination (as is the absence of an effect of the marginal discriminator). Given that Black's model allows for the influence of a small contingent of very misogynistic employers, it follows that this model better fits the data on the gender wage gap than does the Becker model.

Table 7 reports the relationships between our measures of sexism (both misogyny and normative sexism) and various labor market and related outcomes for women—education, marriage, fertility, and full-time employment. Panel A of Table 7 shows that normative sexism is associated with an increase in the gender gap in college degree attainment and that this relationship is statistically significant at the 10 percent level. Misogyny is actually associated with a smaller gap in degree attainment though this relationship is statistically insignificant. Panel B shows that though the gap between men and women is smaller in states with higher misogyny, women in these states are still getting less education (conditional on age and year dummies) than in states with lower misogyny, all else equal. This is consistent with lower education among both men and women in high-misogyny areas. Panel B makes clear that strongly gendered norms and misogyny are both related to reduced educational attainment among women.

Panel C of Table 7 shows that both misogyny and normative sexism are negatively associated with the proportion of women between ages 20 and 40 who have never married, after controlling for level of education, age, and year. This suggests that women in areas with stronger gender norms and more misogyny are more likely to marry young, all else equal. Similarly, Panel D shows that both measures are associated with lower average ages at first birth, meaning women tend to start having children earlier in areas with more traditional gender norms and greater misogyny, all else equal. While it is not surprising that gender norms influence the probability of marriage and timing of child bearing, the role of misogyny in these decisions is less obvious. Given the relationship between misogyny and the wage gap, however, it is possible that women respond to labor market discrimination by exiting the labor force in favor of marriage and motherhood.

Panel F of Table 7 shows that full-time employment among women is lower where normative sexism is higher and where misogyny is higher, with the misogyny coefficient just shy of statistical

[18] A positive coefficient on Percent women implies that a higher percentage lowers the gender wage gap.

significance at the 5 percent level. When we consider the gap between the full-time employment rate of women relative to men, given in Panel E of Table 7, we see that the gap in participation is increasing in both types of sexism and the relationships are statistically significant, suggesting that both are important determinants of women's employment relative to men's.

Overall, it is clear that both misogyny and gender norms play an important role in determining the residual wage gap as well as in other labor market outcomes. We cannot simply attribute differences in outcomes to differences in choices driven by gender norms, though these do play an important role. Moreover, the fact that misogyny strongly correlates with the choices women make regarding human capital accumulation and employment suggests that estimates linking taste discrimination to the residual wage gap understate the impact of discrimination.

8 Conclusion

In this paper, we construct an entirely novel measure of misogyny based on Google search interest in misogynistic terms used to describe women, and present evidence that this measure captures true distaste for women. In addition to our measure of misogyny, we also construct a measure of gender norms based on survey responses regarding the role of men and women in society. We test for taste-based discrimination by estimating whether misogyny helps explain the gender wage gap. Given that we find such evidence, we also test the predictions of both the Becker and Black models of taste discrimination in the context of the gender wage gap, based on the misogyny index. Finally, we examine the relationships between this index, as well as our measure of normative sexism, and various labor market outcomes that can influence wages.

Quite surprisingly, we find that misogyny is strongly related to the residual wage gap, and this is true even accounting for gender norms. It is also an important determinant of other labor market indicators including the probability that women aged 20 to 40 remain unmarried, the age at which women have their first child, college degree attainment, and the gender gap in full-time employment. While subtle discrimination against or underestimation of women may be unsurprising in the workplace, this misogyny measure is intended to capture outright disdain for or dislike of women. Its relationship with the wage gap implies that a much more overt kind of discrimination is acting against women in the labor market. Moreover, the impact of misogynistic discrimination

against women on factors that affect wages, such as educational attainment and labor market participation, implies that estimates of the impact of discrimination on the residual wage gap understate the problem.

We also use the measure of misogyny to test the predictions of two variants of models of employer taste discrimination: the original Becker model of discrimination (1957) and Black's (1995) search model of discrimination. We find evidence consistent with Black's search model, and show that the data is inconsistent with Becker-type models.

This paper highlights that even accounting for gender norms, distaste for women is an important predictor of the gender wage gap. This type of sexism reflects more than subtle or implicit bias against women, as the measure we construct and analyze reflects language that is related to anger and contempt toward women. In our view, then, this paper provides evidence that the gender wage gap cannot be waved away as merely reflecting women's choices, gender norms, or subtle discrimination against women. Nor, to the extent it reflects discrimination, does it reflect only statistical discrimination. Instead, misogyny appears to play an important and previously unmeasured role in determining the wage gap between women and men.

Table 1: GSS Questions on Attitudes Toward Women

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.

This table describes questions about gender norms from the GSS. The left-hand column gives the variable name used by the GSS. The italicized text in the right-hand column is the summary of the question topic. The non-italicized text in the right-hand column is the question as phrased in the GSS.

Table 2: Sexist Search Terms

Manosphere
misandry
men's rights
sjw
Reactionary
sexual harassment
workplace harassment

This table gives the search terms, by category, used to explore potential measures of misogyny based on Google Trends data.

Table 3: Top 20 Queries and Relative Interest - Reactionary 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

This table provides the top 20 related searches that occur in the same search session as each "Reactionary" search term. Interest (Int.) reflects relative search volumes for related queries. Scores of 100 imply that the query is most common and interest in other queries is given relative to it.

Table 4: Top 20 Queries and Relative Interest - Manosphere Words

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

This table provides the top 20 related searches that occur in the same search session as each "Manosphere" search term. Interest (Int.) reflects relative search volumes for related queries. Scores of 100 imply that the query is most common and interest in other queries is given relative to it.

Table 5: GSS Responses by Census Division

	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

This table gives average GSS responses by census division after partialling out year indicators.

Table 6: Testing the Predictions of Black and Becker Models Using the Index of Misogyny

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Misogyny	-0.108***		-0.098**		-0.215**	-0.052		
	(0.038)		(0.038)		(0.082)	(0.059)		
Normative Sexism		-0.134**	-0.114**					
		(0.059)	(0.057)					
Marginal - Derogatory				-0.095**	0.092	0.037		
				(0.043)	(0.082)	(0.055)		
Percent Women						2.031***		2.074***
						(0.277)		(0.261)
Lower Tail - Derogatory							-0.008	-0.004
							(0.043)	(0.027)
Upper Tail - Derogatory							-0.025	-0.003
							(0.021)	(0.014)
N	51	51	51	46	46	46	46	46
R^2	0.14	0.09	0.21	0.10	0.22	0.66	0.14	0.66

This table gives the relationships between the residual wage gap (controlling for education, a quadratic in potential experience, and year) and the percent of women employed full-time, the average percent of the population that responds with sexism on the GSS, and various parts of the misogyny distribution (the mean, the marginal, and the upper and lower tails). All regressions are weighted by the precision with which the residual wage gap is estimated. Standard errors are given in parentheses.

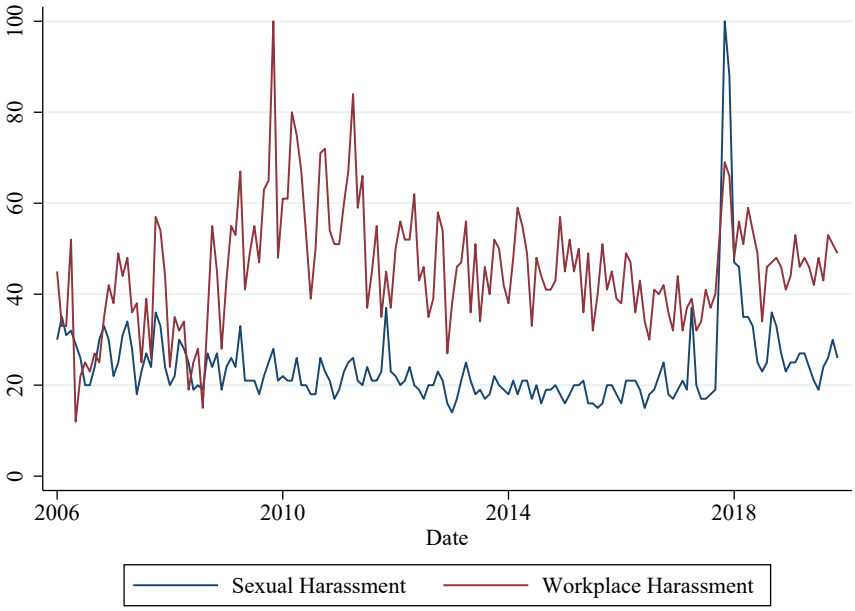
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table 7: Relationship between Sexism Indices and Labor Market Outcomes

	(1)	(2)	(3)
Panel A. Gender Gap in College Degree Attainment			
Misogyny	0.011		0.015
	(0.014)		(0.014)
Normative Sexism		-0.039*	-0.043*
		(0.021)	(0.021)
N	51	51	51
R^2	0.01	0.07	0.09
Panel B. Female College Degree Attainment			
Misogyny	-0.366***		-0.322***
	(0.100)		(0.092)
Normative Sexism		-0.535***	-0.465***
		(0.153)	(0.140)
N	51	51	51
R^2	0.22	0.20	0.36
Panel C. Proportion of Women Aged 20-40 Never Married			
Misogyny	-0.216***		-0.195***
	(0.065)		(0.056)
Normative Sexism		-0.388***	-0.362***
		(0.095)	(0.086)
N	49	49	49
R^2	0.19	0.26	0.41
Panel D. Average Age at First Birth			
Misogyny	-4.660***		-3.925***
	(1.200)		(0.982)
Normative Sexism		-8.108***	-7.293***
		(1.594)	(1.410)
N	51	51	51
R^2	0.24	0.35	0.51
Panel E. Full-Time Employment Gap			
Misogyny	-0.111**		-0.100**
	(0.034)		(0.033)
Normative Sexism		-0.139*	-0.117*
		(0.054)	(0.050)
N	51	51	51
R^2	0.18	0.12	0.26
Panel F. Full-Time Female Employment Rate			
Misogyny	-0.060		-0.054
	(0.033)		(0.033)
Normative Sexism		-0.075	-0.063
		(0.051)	(0.051)
N	51	51	51
R^2	0.06	0.04	0.09

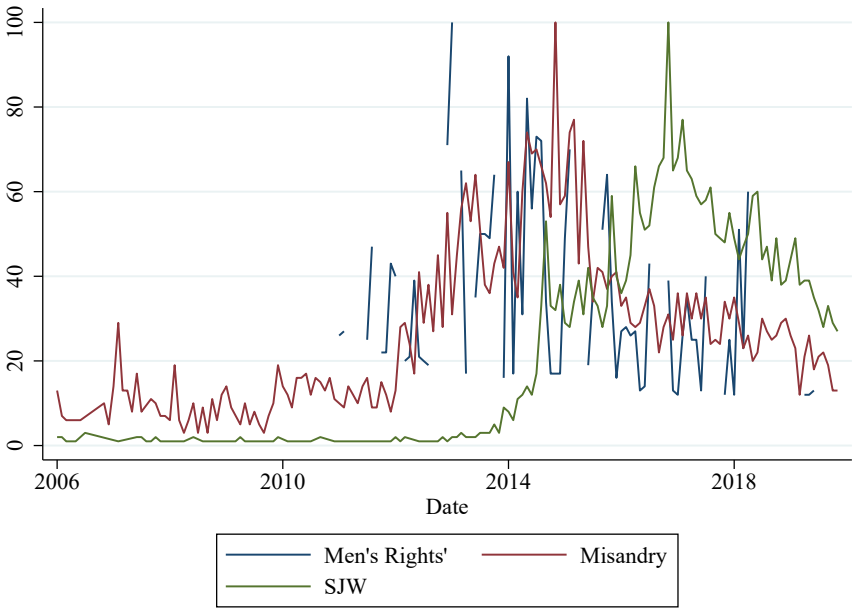
This table gives the relationships between various labor market outcomes and both misogyny and normative sexism. Estimates for college completion gaps and rates (Panels A and B) control for age and year indicators and are weighted by the census earnings weight. Estimates for the proportion of women aged 20-40 never married (Panel C) control for age, education level, and year indicators and are weighted by the census earnings weight. Estimates for the average age at first birth (Panel D) control for year indicators. Estimates for full-time employment gaps and rates (Panels E and F) control for age, education, and year indicators, and are weighted by the census earnings weight. All regressions are weighted by the precision at which the outcome is estimated. Standard errors are given in parentheses.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Figure 1: Reactionary Search Terms Over Time



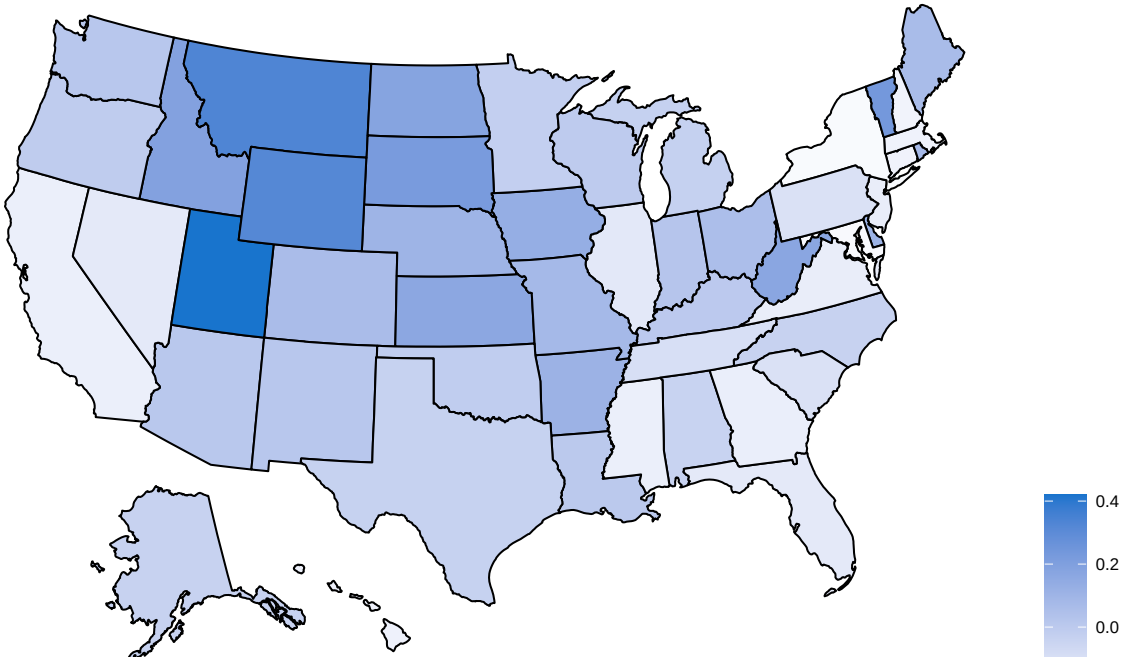
This figure gives Google search interest between 2006 and 2020 for the "Reactionary" search terms. These are normalized by the maximum rate of searches in the time period and scaled by 100.

Figure 2: Manosphere Search Terms Over Time



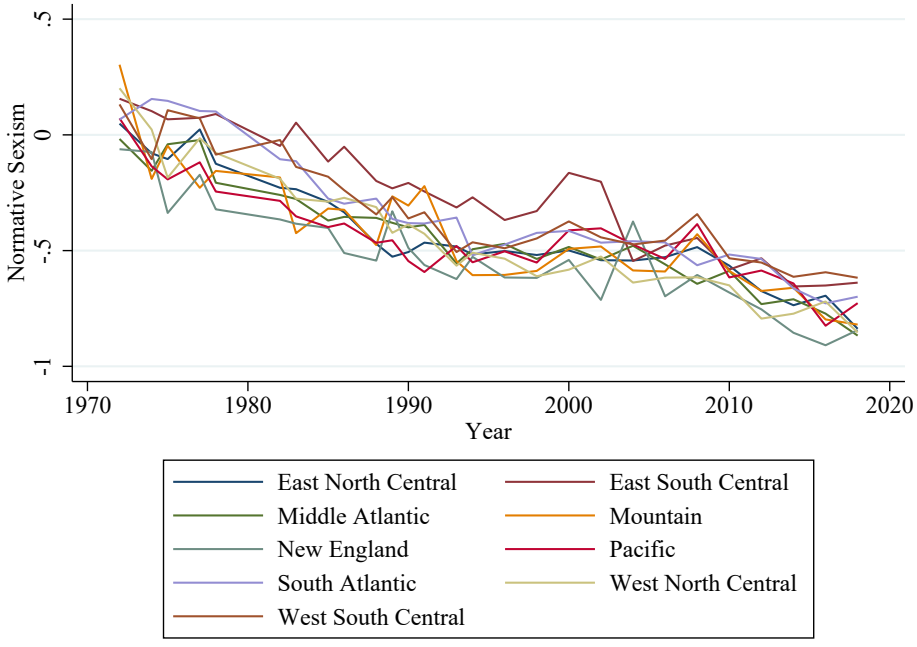
This figure gives Google search interest between 2006 and 2020 for the "Manosphere" search terms. These are normalized by the maximum rate of searches in the time period and scaled by 100.

Figure 3: Index of Misogyny by State



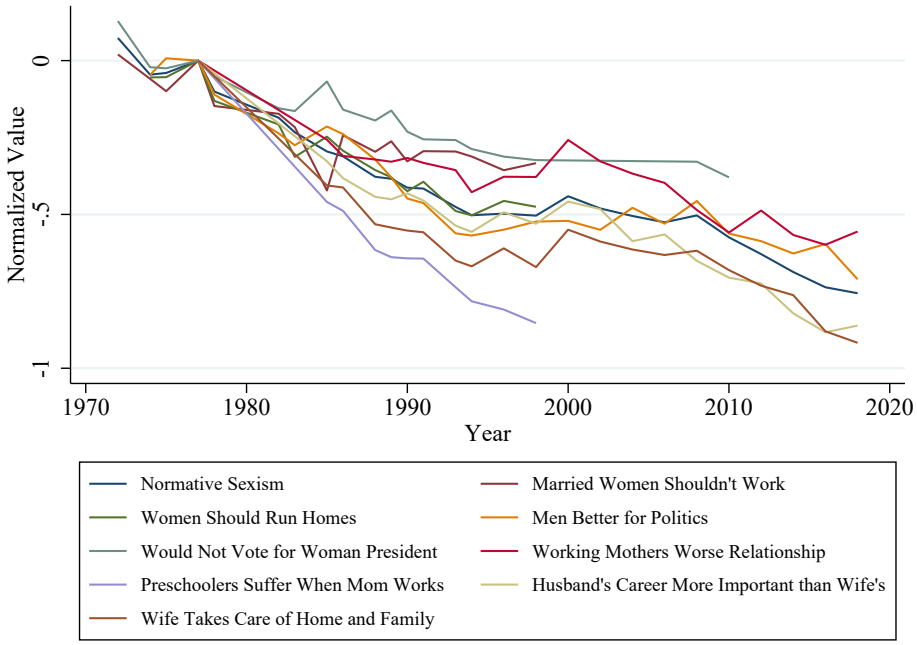
This figure gives average misogyny levels by state after partialling out year effects.

Figure 4: Average Normative Sexism Index Over Time by Census Division



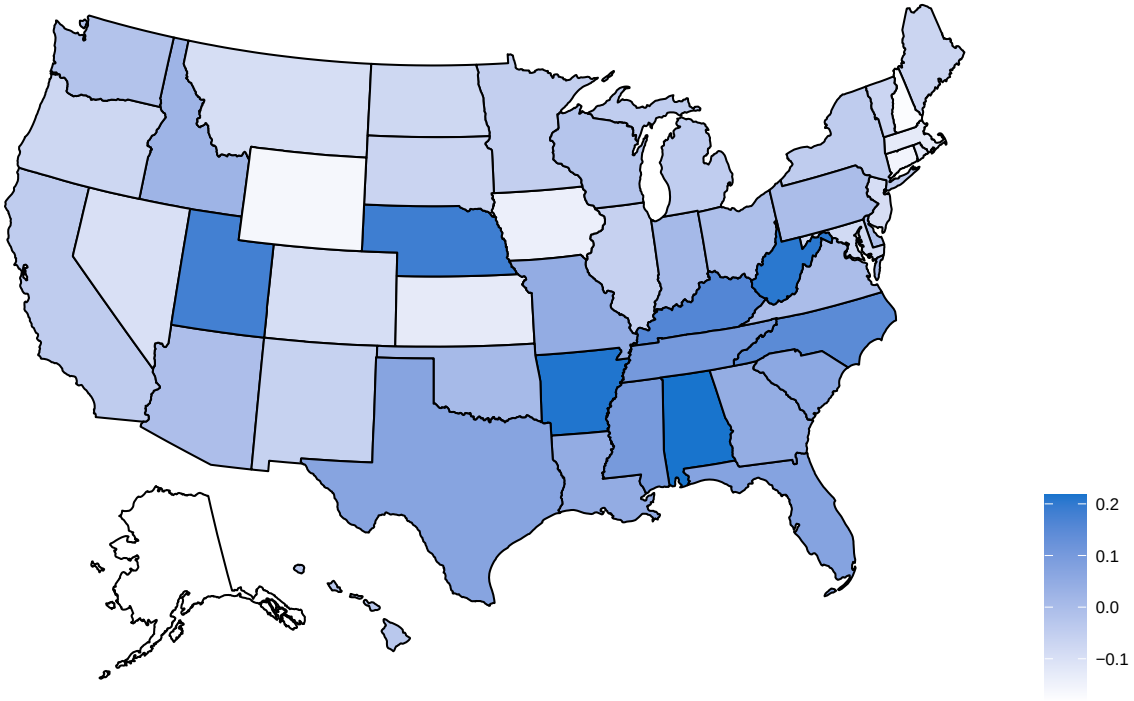
This figure gives average levels of normative sexism over time for each census division.

Figure 5: Responses to GSS Normative Sexism Questions Over Time



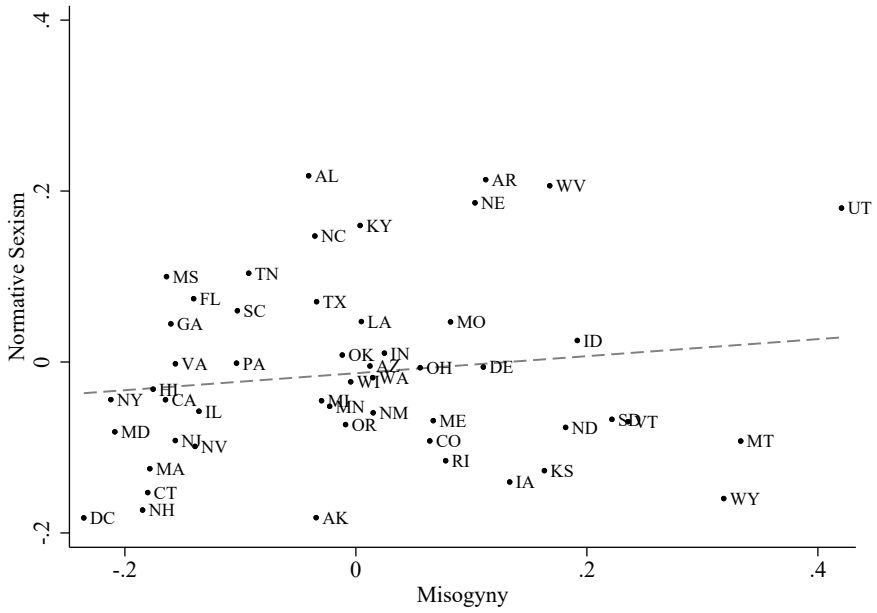
This figure gives the time trends in responses to GSS questions about women as well as the normative sexism index.

Figure 6: Normative Sexism by State



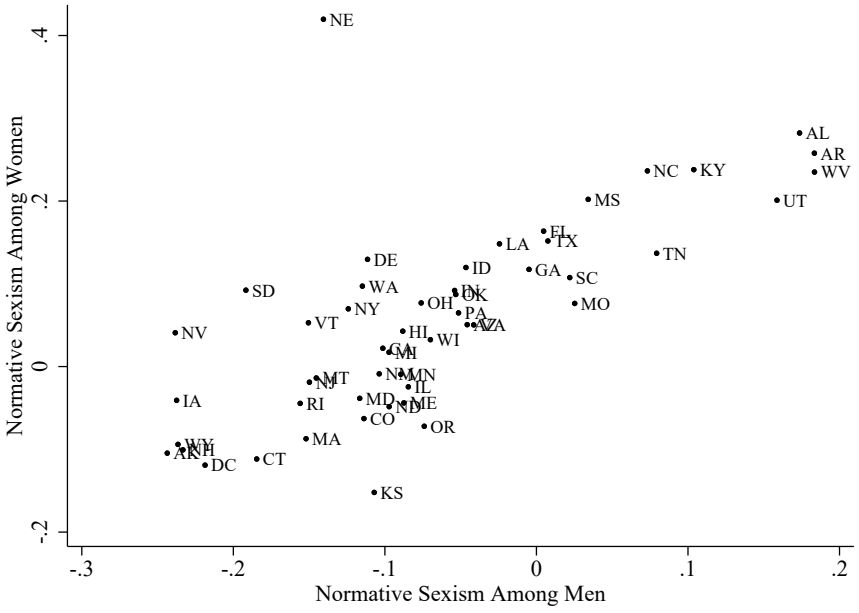
This figure is based on average levels of normative sexism in each state.

Figure 7: Misogyny vs. Normative Sexism Indices



This figure plots average misogyny against average normative sexism for each state. The dashed grey line gives the linear relationship between the two. The correlation coefficient between the two measures is 0.1430 (p-value = 0.3168).

Figure 8: Normative Sexism - Men vs. Women



This figure gives the average level of normative sexism constructed from responses by men against the average level of normative sexism constructed from responses by women. The correlation coefficient between the two measures is 0.7076 (p-value < 0.0001).

Table A.1: Correlation Between Various Measures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Misogyny (1)	1.00							
Normative Sexism (2)	0.14	1.00						
Percent Sexist (Normative) (3)	0.12	0.89***	1.00					
Percent Women LF (4)	-0.31**	-0.11	-0.11	1.00				
Manosphere (5)	0.19	-0.09	-0.03	-0.21	1.00			
Violent (6)	-0.00	0.11	0.06	0.15	-0.11	1.00		
Reactionary (7)	0.55***	-0.23	-0.25*	0.03	0.01	0.16	1.00	
Porn Interest (8)	0.19	0.28*	0.16	-0.16	-0.02	0.24*	0.22	1.00

This figure shows the correlation between various measures including misogyny, normative sexism, the percent of the population whose average GSS responses reveal normative sexism, the percent of women employed full-time, average normalized interest in the manosphere, violent, and reactionary search terms, and porn search interest.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

References

Aigner, D. J., & Cain, G. G. (1977). Statistical theories of discrimination in labor markets. *ILR Review*, *30*(2), 175–187.

Arrow, K. J. (1973). *The theory of discrimination*. Princeton University Press.

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.

Begley, S. (2014). Ironic misandry: Why feminists pretending to hate men isn't funny. *Time Magazine*.

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

Bianchi, S. M., Milkie, M. A., Sayer, L. C., & 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., & 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., & Kahn, L. M. (2017). The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature*, *55*(3), 789–865.

Burn, I. (2017). *Pride, prejudice, and wages: An empirical assessment of models of taste-based discrimination for gay men* (tech. rep.). Working Paper, February 2017.

Burn, I. (2019). The relationship between prejudice and wage penalties for gay men in the united states. *ILR Review*, *73*(3), 650–675.

Charles, K. K., & 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., & Pan, J. (2018). *The effects of sexism on american women: The role of norms vs. discrimination* (tech. rep.). National Bureau of Economic Research.

Cps org uniform extracts, version 2.5. (2020). *Center for Economic and Policy Research*.

De la Roca, J., Ellen, I. G., & O'Regan, K. M. (2014). Race and neighborhoods in the 21st century: What does segregation mean today? *Regional Science and Urban Economics*, *47*, 138–151.

Farrell, T., Fernandez, M., Novotny, J., & Alani, H. (2019). Exploring misogyny across the manosphere in reddit. *Proceedings of the 10th ACM Conference on Web Science*, 87–96.

Fortin, N. (2005). Gender role attitudes and the labour-market outcomes of women across oecd countries. *Oxford Review of Economic Policy*, *21*(3), 416–438.

Fortin, N. (2008). Gender role attitudes and women's labor market participation: Opting-out and the persistent appeal of housewifery. *Unpublished manuscript, University of British Columbia*.

Goldin, C. (2014). A grand gender convergence: Its last chapter. *American Economic Review*, *104*(4), 1091–1119.

Google trends [https://www.google.com/trends.]. (2021).

Hellerstein, J. K., Neumark, D., & Troske, K. R. (1999). Wages, productivity, and worker characteristics: Evidence from plant-level production functions and wage equations. *Journal of Labor Economics*, *17*(3), 409–446.

Hess, A. (2014). The rise of the ironic man-hater. *Slate Magazine*.

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

National vital statistics system, natality on cdc wonder online database [Data are from the Natality Records 2016-2020, as compiled from data provided by the 57 vital statistics jurisdictions through the Vital Statistics Cooperative Program]. (2020).

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

Pastor, K. (2021). Google trends county-dma mapping. https://www.kaggle.com/kapastor/google-trends-countydma-mapping.

Phelps, E. S. (1972). The statistical theory of racism and sexism. *American Economic Review*, *62*(4), 659–661.

Population estimates 2010-2019. (2019). https://census.hawaii.gov/wp-content/uploads/2021/05/co-est2020-alldata.xlsx.

Sayer, L. C. (2016). Trends in women's and men's time use, 1965–2012: Back to the future? *Gender and Couple Relationships*, 43–77.

Smith, T. W., Marsden, P., Hout, M., & Kim, J. (2020). General social surveys, 1972-2018. *NORC at the University of Chicago*.

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.

Tranchese, A., & Sugiura, L. (2021). "i don't hate all women, just those stuck-up bitches": How incels and mainstream pornography speak the same extreme language of misogyny. *Violence Against Women*, *27*(14), 2709–2734.

Zimmerman, J. (2014). No, we don't literally want to ban men. but 2014 was the year women got even. *The Guardian*.

Zimmerman, S. (2024). The ideology of incels: Misogyny and victimhood as justification for political violence. *Terrorism and political violence*, *36*(2), 166–179.