Preliminary and incomplete

"Sexism and Women's Labor Market Outcomes"

by

Kerwin Kofi Charles
University of Chicago
Harris School and NBER

Jonathan Guryan
University of Chicago
Booth School of Business
and NBER

Jessica Pan
University of Chicago

October 7, 2009

Abstract

We examine the extent to which cross-market differences in women's relative labor market outcomes are determined by differences across markets in sexism – defined as views about the appropriate role women should play in society. Using data from the GSS to measure sexism, we show that selection-corrected gender wage gaps and relative employment rates are significantly related to the degree of sexist views held by the median male, but not with male sexism at the 10th or 90th percentile. Consistent with a standard labor supply model in which sexism lowers women's offered wage, we find lower relative employment of women in more sexist markets is concentrated among women who would have worked few hours in sexism's absence. Finally, we show that the patterns described for male sexism are not apparent for female responses to the GSS questions. The results are robust to a variety of extensions, including alternative strategies for correcting for gender skill differences, and selection. We argue that these results are consistent with a taste-based model of discrimination (Becker (1957)), and are especially striking in light of results from Charles and Guryan (2008) who find that racial wage differences are related to the left tail of the racial prejudice distribution, rather than the median or right tail - exactly as the prejudice model predicts for a group whose prevalence in the labor market much less than that for women. The results suggest that sexism has important implications for the workings of labor markets for men and women.

I. Introduction and Overview

A familiar result from the large literature on gender differences in the labor market is that, despite substantial convergence over the last half-century, women in the U.S. continue to receive lower wages and exhibit lower rates of employment than observationally equivalent men (Altonji and Blank, 1999, Blau and Kahn, 2000). Much less well known is that female-male gaps differ widely across labor markets. Using residuals from standard OLS individual regressions, Figure 1 shows geographic variation in state-level differences in male-female conditional employment and log wage gaps.[1] Table 1A and 1B summarize the wide cross-state variation in relative outcomes evident in the figure. The adjusted female-male employment gap ranges from a minimum of -0.24 percentage points in Louisiana, to -0.04 percentage points in D.C., with a standard deviation across states of 0.035 and mean of -0.170. The table shows that cross-state variation in the female-male log wage gaps are equally pronounced, whether at the mean or median. For example, state-level median wage gaps range from a low of -0.42 in Wyoming to -0.12 in D.C., with a cross-state standard deviation in the gap of 0.044. For all of these outcomes, the differences across various states exist for different racial groups and across education groups.

What accounts for these differences across labor markets? Unobserved physiological differences between men and women, or women's greater specialization in marriage, childbearing and home production might explain why, *within* a given labor market, women's labor market outcomes are inferior to men's,[2] but it is doubtful that these factors can account for the differences in male-female gaps *across* markets. This paper argues that cross-market differences in conditional gender gaps can be explained, in part, by differences across markets in the views held by men about the appropriate roles that women should play in the labor market and in society more generally.[3]

---

[1] The regressions are estimated on CPS MORG/May data from 1977-2002 and include controls for education, a quadratic in experience, an indicator for marital status, gender-specific year effects, and gender-specific state effects. The figure plots the state*female coefficients. The excluded female*year is 1990, so the levels of the gaps are for that year. Excluding later (earlier) years would produce smaller (larger) gaps, but the variation across states would be the same.

[2] Gender differences in differences in brawn may be one such possible physiological difference. A large literature (Becker, 1981) argues that because men and women usually form family units and consume the sum of their joint home and market product, family utility might be maximized by women concentrating on the production of the home-produced good, and men specializing in market work.

[3] One alternative reason for cross-market differences in female-male employment and wage gaps could be the spatial distribution of jobs at which women are especially able.

Suggestive evidence that men's negative view about women's appropriate roles matters for women's relative outcomes is provided by the two panels of Figure 2, each of which plots the relationship between the conditional white female-male employment and selection-corrected median wage gaps[4] to the mean number of men in a state responding "yes" to two questions in the General Social Survey (GSS): Whether a woman's working in the labor market hurts her children? And, whether a woman should help her husband's career first? For both of these questions, and for both labor market outcomes, there is a strongly negative association between how women fare relative to men in the labor market and men's views about women's appropriate roles.

In the analysis below we focus on gender differences among whites to avoid conflating the issues concerning gender differences that are our focus, with the potentially very different set of issues having to do with differential racial outcomes. We use men's individual-level responses to various questions, similar to the two above, to construct a comprehensive index of men's negative attitudes regarding women's roles. We label this measure of negative male views *sexism*. After accounting for various specification and other issues, our empirical analysis principally involves relating various measures of women's relative labor market performance in a state, to features of the distribution of sexism among men within a state. Our empirical work aims to establish not only whether sexism within a state is systematically related to relative female outcomes but also, and mainly, whether the form that relationship takes is consistent with the very specific predictions of a prejudice model of taste-based discrimination of the form introduced by Becker (1957) and subsequently analyzed by Arrow (1973) and more recently by Charles and Guryan (2008).

Using data on labor market outcomes from several years of the CPS, along with data from the General Social Survey about men's sexist beliefs, and after accounting carefully for selection and other concerns, we show that both conditional employment and median wage gaps across states vary negatively with differences across states in the 40th or 50th percentile of the male sexism distribution. By contrast, relative labor market outcomes do *not* depend on differences across states in other points of a state's male sexism distribution – whether in the far left (10th percentile) or far right (90th percentile) tails. Higher median sexism is especially strongly associated with reduced hours or work among women working relatively few hours, but is not

[4] Wage gaps are corrected for selection by imputing wages above or below the median based on educational attainment for women not working full-time and then estimating median regressions (see e.g. Neal (2004) and Olivetti and Petrongolo (2008)). This procedure is explained in detail below.

associated with cross-gender differences in hours for full time white workers. We argue that these findings are strongly consistent with the predictions forthcoming from a taste-based prejudice model.

In the various employer or co-worker versions of the standard taste-based model, it is assumed that prejudiced agents incur a psychic cost when interacting with the object of their prejudice and as a consequence demand some monetary premium to engage in any such interaction.[5] Since it is costly when the two types of agents interact, the market acts to separate the prejudiced persons from the objects of their bias. Moreover, because the cost of interaction rises as the relevant prejudice increases, sorting will be such, if there is to be any interaction in equilibrium, the *least* prejudiced persons will be those who interact with disadvantaged group members. This sorting mechanism guarantees that equilibrium prices will be determined not by the bias of the average, or "representative", prejudiced person but rather by a *marginal* discriminator whose prejudice is such that he is just indifferent between interacting with the disadvantaged group and not. The prejudice of this marginal discriminator - and thus equilibrium prices -- depends on the size of the disadvantaged group; the larger the disadvantaged group, the further up the frequency distribution of prejudice will be the marginal discriminator.

An especially useful prejudice model for illustrating the relationship between gender labor market differences and sexism is the model of employer-based prejudice. Following the standard formulation, suppose that sexism imposes a psychic cost on a man if he either interacts with directly or encounters a woman at work whose behavior is in contravention of his notion of women's appropriate roles. The cost to a sexist (male) employer of hiring a woman is thus the sum of the wage paid to women plus his sexism, whereas his cost of hiring an equally productive man is simply the male wage. When the number of women relative to men who supply labor is low, all women supplying labor are hired by un-prejudiced men. However, as the number of women relative to men increases, market clearance requires that women work with ever more sexist employers. Because these employers regard women as more expensive than they really are, they are only indifferent between hiring a woman and the man they could otherwise hire if the woman's money wage is less than the man's. As ever more women supply labor women's

[5] The insights from Becker's original model and those that extend his work are well established. For example, Charles and Guryan (2008) endogenize the roles of worker and employer, Becker (1957), Arrow (1972) and Welch (1967) discuss imperfect substitutability of labor, Arrow (1972) includes supervisors, Black (1995) incorporates search and monopsony. The main conclusions and insights from all of these models of taste-based discrimination are largely similar.

relative wages must fall even further before the last employer to hire women – the marginal employer – is indifferent between hiring men and women.

Since women constitute close to a half of all workers in most labor markets, the marginal employer would be drawn from somewhere close to the center of the sexism distribution. Gender wage differences across markets should be related to sexism close to median of the sexism distribution but not to other moments, as these persons should be infra-marginal. Nor should wages be the only labor market outcome affected by sexism; median sexism should also affect relative labor supply. Because any decrease in the market offer wage reduces the number of persons for whom the wage exceeds their reservation wage, it should unambiguously be the case that markets in which the median male is more sexist are places where a smaller share of women work. The effect of higher median sexism on hours worked among those who do work is theoretically ambiguous, and depends on the relative magnitudes of the income and substitution effects associated with the lower offer wage. For women who would otherwise work few hours, the low offer wage that comes from being in a market where the median man is very sexist is associated with a very large substitution effect and a small income effect. For these women, hours should decline and, for some should definitely fall to zero. For women working large hours, the effect of smaller offer wages on hours could be of either sign.

Our empirical findings support all of the predictions of a prejudice-based model of sexism. It is nonetheless useful to ask whether our relating of male sexism and women's labor market outcomes may actually be due to some other factors. For example, it could be that what we are calling male sexism merely reflects some unobserved feature of markets that also affects female wages and labor supply behavior. If women in a particular labor market, for whatever reason, do not wish to focus on work, they might be disinclined to invest in the sort of unobserved skills valued in the labor market. Men, living in those communities, when giving responses that we characterize as sexist may be merely reflecting that same community-wide sentiment rather than reporting something that would cause them psychic injury were they to encounter it.

Two things militate against this possibility in our study. The first is that virtually any such alternative account for what male sexism reflects would seem to imply that *all*, or at least several, moments of the ``sexism'' distribution should be systematically related to women's outcomes. Our finding that the median sexism -- and only the median – matters for women's market outcomes would seem to be inconsistent with any other explanation other than the standard prejudice argument. Moreover, in much of the analysis we control *directly* for women's own responses to very same sexism questions put to men. If male responses merely indicate some general community sentiment distinct from sexism, then these factors should load onto female responses to the questions in the various regressions. The fact that we find strong effects for median male sexism, even after controlling for women's own views about the tasks they should engage in, suggests that the male sexism measures capture something close to the prejudicial feeling we have described.

We have argued that our findings are consistent with sexism that operates mainly through costly interaction at work as implied by the prejudice model. But might our results be explained by some other consequence of sexism that does not rely on workplace interaction? One possibility is that male sexists might, through their control of the education or political systems, affect the distribution of skills possessed by men and women *before* they get to the labor market. However, as all of our results are conditioned on standard measures of skill, this argument is unlikely to account for our findings. Another possible non-workplace effect of sexism is that in a labor market in which men are strongly opposed to women working, a woman who chooses to work must overcome the objection or the bad feelings of the various men in her life – husband, friends, brothers, and sons.

The natural way to represent this possible effect of sexism is to say that a labor market with high sexism is one in which a woman faces a large (societal) fixed cost of labor supply. Large fixed labor costs should decrease the share of women who work, and may even affect, through selection, the observed wage distribution of women. There are nonetheless two reasons to doubt that this argument can explain our results. One is that there is no reason that fixed labor costs should affect women's *offer* wage distribution, as we find. In addition, the standard labor supply model would suggest that higher fixed labor costs should lead to an *increase* in hours of work for women working many hours – something that is absent in our data. Finally, if men's views impose a cost on market entry onto women, it is not at all clear why only the views of the median man should matter for these costs.[6]

Our results complement and extend recent work on prejudicial tastes and labor market discrimination in a variety of ways. One of these is to broaden the notion of prejudice beyond generalized *aversion*. Previous taste based papers on the racial wage gap, including Becker's seminal paper and Charles on Guryan (2008), conceptualizes prejudice as an antipathy towards all

[6] One hypothesis that seems reasonable is that the views of every individual in the area should matter, since each person can express his objection to behavior that breaks social norms, but that stronger views should matter more since those with strong views would be more likely to express those views publicly.

interaction with racial minorities. Casual observation and common sense suggests that, whatever their feelings about how women should spend their time, men are not averse to interacting with women in all settings. Yet, our findings suggests that *any* essentialist views about what particular agents can or should do – irrespective any other aversive sentiment is held towards them – might have important labor market consequences.

An even more important extension of recent empirical work is that we focus on a disadvantaged group for which the prejudice model yields very different predictions than those forthcoming for race. Taste-based employer discrimination in both the case of race and that of gender leads to sorting, with the result that the relevant wage gap is determined by the marginal discriminator. But the *identity* of that marginal employer is very different for the two cases. In particular, unlike gender discrimination where the marginal discriminator should be close to the median of the relevant prejudice distribution as this paper finds, the fact that blacks everywhere are a small minority of employment means that the marginal discriminator in the race case should be drawn from the *left* tail of racial prejudice distribution – somewhere substantially below the median. Indeed, Charles and Guryan (2008) find that, for racial wage gaps among men, points close to the 10^{th} percentile of the racial prejudice distribution are systematically related to relative wage wags whereas other moments in racial prejudice distribution are not. The fact that both our study and the Charles and Guryan paper appear to confirm the sharply different predictions of the same underlying model buttresses our confidence that taste-based considerations do in fact play an important role in determining labor market outcomes for disadvantaged groups.

A final way in which this paper extends the previous literature is in the broad set of labor market outcomes we study. Most analyses of discrimination focus on wage differences and the recent empirical work relating direct measures of prejudice is no exception. For example, in the paper on the effect of racial prejudice among men, Charles and Guryan (2008) focus exclusively on how prejudice affects racial wage gaps. But precisely the same mechanism that drives wage differences should, by extension, also affect various dimensions of labor supply, as our results show. Understanding how prejudicial feeling affects outcomes other than market wages is important both because it is desirable to have as rich as possible a picture of the effects of various types of prejudice, and because, as with women, it is in the probability of working or the hours of work where some of the largest unexplained labor market outcome differences are to be found.

In the next section we describe the data used to measure labor market outcomes and sexism and document both cross-sectional patterns and trends in sexism over time. In section III, we present the main results showing the relationship between various labor market outcomes for women and male sexism, and in section IV we explore the relationship between female labor market outcomes and both male and female views on appropriate gender roles. In section V, we investigate various alternative explanations for the patterns we document in the data. We then conclude in section VI.

II. Data Description

This section describes the data on gender attitudes used in the empirical analysis to explore the relationship between male sexism and cross-state differences in female labor market outcomes. We summarize geographical differences in male and female views toward gender in the United States and how gender prejudice has evolved over the last three decades.

Data on Gender-Based Prejudice

To construct measures of gender-based prejudice, we use data from the 1972 to 2006 waves of the General Social Survey (GSS). The GSS is a nationally representative survey that asked respondents a variety of questions pertaining to their attitudes toward female gender roles. There are eight gender-related questions that were rotated through the years that touch on various aspects of sexism. These questions fall into three main groups. The first focuses on the appropriate role of women in society - for example, typical gender role questions ask respondents if they agree or disagree with statements like “women should take care of running their home and leave running the country up to men” or “it is much better for everyone involved if the man is the achiever outside the home and women takes care of the home and family”. The second set of questions touch on the perceived capabilities of women such as “would you vote a female for president” or “men are better suited emotionally for politics than are most women”. The final set of questions asks whether working mothers are able to juggle their dual roles effectively - for example, “A working mother can establish just as warm and secure a relationship with her children as a mother who does not work”.

The GSS variable abbreviation and summary for each of the eight gender attitude questions are listed in appendix table 1. In each survey year, different subsets of the full questions were asked, with the number of questions generally ranging from four to the full set of eight. Since each of these questions touch on some aspect of sexism, we construct our sexism index based on the full set of eight gender-related questions asked consistently across the survey years[7]. While using other subsets of questions may enable us to consider a wider time period, the trade-off is that fewer questions would be included in the sexism index. By using the full set of gender-related questions available in the GSS, we make no prior predictions on which attitude questions are more or less reflective of sexist sentiments. Throughout, we use responses for all races aged 18 and older, and responses are recoded so that higher values correspond to more gender-prejudiced answers.

To facilitate comparisons of the degree of sexism across individuals and across geographic areas, we combine the responses to the GSS questions into a unidimensional sexism index based on a procedure similar to Charles and Guryan (2008). To create an individual-level index for each GSS respondent, we need to ensure that the responses to each question is measured on the same scale and weighted equally in the index. This is done by creating a normalized measure that subtracts off from the individual responses to each question the mean of the response of the full population in 1977 and divides by the standard deviation in the first year the question was asked[8]. We normalize by the mean and standard deviation of both male and female responses to the gender-prejudice questions rather than the gender-specific mean and standard deviation to ensure that the sexism index for both men and women are on the same scale. The one-dimensional aggregate individual-level sexism index is computed by taking the average of the normalized responses in each survey year for each individual.

To explore the relationship between sexism and women's labor market outcomes across different geographic regions, we combine the individual-level sexism indices to obtain aggregate measures of sexism among males and females within a community. We do this by first regressing the individual-level sexism index on a full set of year dummies, separately for males and females. Next, we use the residuals from this regression to create a measure of "average" male (female) sexism, which is simply the mean across all years of the residual individual-level sexism index for males (females) in a community. Another set of measures uses the 10^{th}, 50^{th} and 90^{th} percentiles of the distribution of the residual individual-level sexism index for males and females within a state to capture different percentile points in the overall sexism distribution in a state. All

[7] This corresponds to the following years in which the GSS was conducted: 1977, 1985, 1986, 1988, 1989, 1990, 1991, 1993, 1994, 1996 and 1998.

[8] Following Charles and Guryan (2008), we normalize by the standard deviation in the first year the question was asked rather than, say, the overall standard deviation, because we want to avoid a mechanical relationship between trends in responses and the weight the question receives in the overall aggregate. We choose 1977 as the normalization year because it was the first year that the full set of eight gender-prejudice questions were asked.

the aggregate community sexism indices are constructed using the individual-level weights provided by the GSS.

B. Male and female views on gender

Cross-sectional variation

We begin by examining how men's and women's views about women's appropriate roles – what we call sexism – vary across both geographic areas and over time. Table 2A reports men's average responses to five of the GSS questions aimed at measuring views towards women. The table reveals a good deal of geographic variation in male views of appropriate gender roles. Male sexist views are strongest in the three Southern regions in all of the questions examined, and are least extreme in New England and the West. Though the ranking of regions by their levels of sexism is very similar to the cross-regions rankings for racial prejudice documented in Charles and Guryan (2008), cross-region differences in sexism are not as extreme. Thus, whereas the difference in the average racial prejudice index between the East South Central and New England divisions is about 0.5, the comparable difference for male sexism is 0.275. Although not as large as that for racial prejudice, cross sectional variation in male sexism is nonetheless quite substantial. For example, the median sexist East South Central male is as sexist as the 69th percentile sexist male in New England; and the median sexist New England male is as sexist as the 30th percentile sexist male in the East South Central division.

Relative to the variation in responses across individuals in the questions posed, there is about half as much cross-sectional variation across geographic areas in gender prejudice as racial prejudice. However, the variation is the case of sexism is more uniform across the distribution. Figure 3 shows the cumulative distribution function of male sexism, separately for each of the nine census divisions. There is somewhat more variation across divisions at the higher quantiles, e.g. the 80th v. the 20th, but the differences are not as significant as is the case for racial prejudice.

Table 2B documents the same patterns except for women's views about women's appropriate roles. Many of the geographic patterns for women's views about gender roles are similar to those found for men's views. Compared to women elsewhere, women in the South feel most strongly that there are specific roles that ought to be played by women in the labor market and in society more generally, while women in New England and the Pacific divisions have the least pronounced views about appropriate gender roles. Males and females views on female roles are most different in the Mountain and Middle Atlantic divisions. In the former, women hold more extreme views on gender roles; it is the views of men that are more extreme in the latter case. Interestingly, though the ranking of divisions is different by male and female gender prejudice, there is comparable geographic variation in tastes among men and women.

Intertemporal Variation

Figure 4 shows trends over time in the average index of male sexism separately for each of the nine census divisions. Several things are noteworthy about these patterns. First, there was a clear negative trend in men's measured sexism between 1977 and 1998. The decline in sexism over those 20 years is even larger than the difference on average over the period between the East South Central and New England census divisions. Second, the degree of cross-sectional variation across census divisions has remained fairly constant over time. There does not appear to have been a significant amount of convergence in gender attitudes over this period. And third, while there was clearly some change in relative rankings over time, the vast majority of the variation in sexism is cross-sectional or secular changes over time. A regression of division-by-year sexism on division and year fixed effects yields an R^2 of 0.78. Because of concerns that much of the remaining variation within divisions over time is noise and that the GSS has too few observations per state-year cell to reliably measure changes in individual quantiles of sexism, we follow Charles and Guryan (2008) and focus on pooled cross-sectional variation in measured prejudice in the regression analysis to follow.

Figure 5 shows trends in average responses by men to each of the GSS questions used to measure gender prejudice. For each question, men's views have become substantially less sexist over time. We have normalized the responses to be zero in the full population in 1977, so that variation in that year reveals differences in responses by men and women. In 1977 men had stronger views than women that a mother working was detrimental to her child; male and female views were more similar with regard to whether women were suited for politics. By the 1990's, male views had become less prejudicial on whether women were suited for politics and whether a wife should support her husband's career over her own.

III. Estimates of the effect of sexism on female labor market outcomes

A. The relationship between male sexism and gender wage gaps

Having examined the cross-sectional and intertemporal patterns in sexist attitudes, we now turn to the main empirical analysis of the paper. As described at the outset, we aim to investigate whether the relationships between various labor market outcomes, including wages, and particular moments of the distribution of sexism across states are consistent with the predictions of taste-based prejudice models, as originally outlined by Becker (1957) and explored by subsequent authors. We begin by examining the most basic prediction of the employer version of that model – the relationship between the distribution of male sexist preferences and the market female-male wage gap.

Recall from the earlier discussion that Becker-type prejudice models do not, in fact, make a prediction about the relationship between gender wage gaps and the *average* level of sexism in a labor market. Rather, the key prediction of such a model is that, because of the market tendency to sort women away from interaction with sexist market actors such as employers, gender wage gaps should be determined by the sexist sentiments of the marginal discriminator – the most sexist person with whom women are forced to interact after sorting has occurred. Since it is this interaction that determines wages, the model thus predicts a particular relationship between offer wage gaps and particular quantiles of the sexism distribution. It bears stressing that here that, that if racial prejudice and sexism play important roles in wage determination for blacks and women, respectively, the mechanism by which wages are determined will be the same but the predictions about how wage gaps relate to particular moments of the relevant distribution will be very different. In particular, because blacks are a numerical minority in most labor markets the marginal discriminator in the case of race is likely to be drawn from the left tail of the racial discrimination distribution. The empirical results in Charles and Guryan (2008), hereafter CG, confirm this prediction.

In contrast, women make up approximately 40-50 percent of most labor markets. Even as market forces push for segregation, significantly more sexist men will interact with women than racially prejudiced whites will interact with blacks in the labor market. As a result, the model predicts the marginal discriminator against women should be drawn from the middle (somewhere around the 40^{th}-50^{th} percentile) of the sexism distribution. To test this prediction, we estimate specifications similar to those shown in CG in which the selection-corrected female-male wage gap is regressed against the 10th, 50th, and 90th percentiles of male sexism. In contrast with some of the specifications reported in CG, the regressions do not include percent female as a regressor because the fraction of the workforce that is female is an outcome, which we examine in the next section.[9]

The predictions of the prejudice model concern the gender difference in *offered* wages, and not the difference in observed wages among those who work for pay. As it is well known, it is important to account for non-random selection when examining variation in female wages. We follow Neal (2004) and Olivetti and Petrongolo (2008) and calculate differences in median wages after imputing values above or below the median for those without valid wages based on educational attainment. Specifically, in our preferred specification we impute wages for a subset of workers not working full-time. Wages above the median are imputed for workers in the top quartile of the gender-specific education distribution, and wages below the median are imputed for workers in the bottom quartile of the gender-specific education distribution.

Following CG, we then estimate the regressions in two steps. First, we estimate a median regression of log hourly wages, with imputations, on controls for education, a quadratic in potential experience, marital status, gender specific year effects, state effects and state*female effects.[10] The coefficients on the state*gender dummies are then taken to be the dependent variable in the second step. These residual state gender wage gaps are then regressed against various measures of male sexism. The reported state-level regression has 45 observations because some small states are not separately identified or sampled in the GSS. These are the same 45 states included in the regressions reported in Charles and Guryan (2008).

Results for selection-corrected wage gap specifications are reported in table 3. The results in the table use wages imputed according to the rule described above for non-employed individuals. In appendix tables described below, we show that the results are not substantively sensitive to various alternative selection-correction imputations.

[9] Whereas variation across states in the fraction of the workforce that is black is mostly driven by racial residential patterns, variation in the fraction of the workforce that is female is wholly the result of gender differences in employment rates across states.
[10] Since it could be an outcome, controlling for marital status is arguably inappropriate. The results are insensitive to the inclusion or exclusion of a control for marital status in the individual level regressions.

The point estimate in column (1) indicates that the relationship between female-male offer wage gaps and average male sexism is negative and statistically significant. The association is slightly stronger with median male sexism, and the magnitude is economically significant. A one standard deviation increase in median male sexism in a state (0.134) is associated with a 0.031 log point decrease in female offer wages relative to males. This decline is 6.3 percent of the mean female-male offer wage gap in 1977 and 52 percent of the cross-state standard deviation of gender offer wage gaps.

The results comparing the different points of the male sexism distribution are striking. As columns (2) and (3) show, adding either the 10^{th} or 90^{th} percentiles separately to a regression with the median level of sexism increases the estimated negative effect of median male sexism on the gender offer wage gap. In all specifications, including those that control for all three points of the male sexism distribution at once, the female-male offer wage gap is significantly related to median male sexism. In no specification is either the 10^{th} or 90^{th} percentile of male sexism significantly related with the offer wage gap. The point estimates of the 10^{th} and 90^{th} percentiles are actually positive, though insignificant and smaller in magnitude than the estimate for the median. In our preferred specification, where we include controls for all three quantiles of the male sexism distribution, the coefficient on median male sexism is negative, significant and large, while the coefficients on the 10^{th} and 90^{th} percentiles are positive, relatively smaller in magnitude and statistically indistinguishable from zero. This result is strongly consistent with the prediction of a prejudice model that women's wages are determined by a marginal discriminator, drawn from close to the middle male sexism distribution as a result of women's representation in the labor force.

Not only are these results of independent interest to those interested in gender differences in labor markets; they are also very striking to those interested in the role of prejudice in determining labor market outcomes more generally. Using measures of racial prejudice from the GSS and computed in a very similar fashion to our sexism measures, Charles and Guryan (2008) find that black-white wage gaps are related to the 10^{th} percentile and *not* to the median and 90^{th} percentile of white prejudice – precisely as the prejudice model predicts for racial minorities. That the data appear to be consistent with the very different predictions about which quantile points in the relevant distribution are marginal and infra-marginal in the case of race and gender suggests that prejudice – both racial and gender – are important determinants of relative wages.

B. Sexism and Labor Supply: Employment and Hours Gaps

The focus of Becker's original prejudice discrimination model and of the subsequent work devoted to empirically assessing its importance has been on differences in prices, such as wages. But it should be clear that gender wage gap is not the only important labor market outcome for which the model has implications. If male sexism generates lower female offer wages, one might expect to observe a labor supply response as well. Male sexism, operating through a Becker-style mechanism, might thus also partly explain variation in female-male differences in employment rates.

A standard labor supply model, in which increases in a worker's offer wage generate income and substitution effects, implies that there might be a response in terms of the probability of employment and in hours worked arising from sexism-induced offer wage changes. For persons working a few hours, or not initially working at all, a reduction generates only a small income effect; a decrease in the offered wage thus weakly reduces the likelihood of being employed, particularly for people working relatively few hours initially. For workers already working positive hours, the income effect associated with a wage reduction increases with hours worked, so the overall effect on hours worked is theoretically ambiguous. Thus, if wages fall the standard labor supply model offers an ambiguous prediction about the share of people working, say, 35 hours of work.

In this section, we study both labor supply outcomes, focusing first on employment rates, and then on the distribution of hours worked. Table 4 presents results for male-female employment gaps. The table shows estimates from models corresponding to those estimated for wage gaps and shown in Table 3. The dependent variable is generated in the same way as wage estimates shown earlier: the initial regression is of an indicator for employment on education, a quadratic in experience, marital status, gender specific year effects, state effects and state*female effects. As with wages, the dependent variable in the second step is the state*female effects and can be interpreted as the state's conditional female-male employment gap.

Column (1) reports the univariate relationship between the residual female-male employment gap and average male sexism. The relationship is strongly negative and statistically significant. In states with more sexist men, female employment rates are lower relative to males. The magnitude of the correlation implies that a one standard deviation change in average male sexism is associated with a 1.7 percentage point decline in the female-male employment gap, which is 4.4 percent of the mean gap in 1977 and 47 percent of the standard deviation of gaps across states. The strong negative relationship between the unconditional female-male employment gap and average male sexism can be seen clearly in the scatter plot shown in figure 6. The District of Columbia (DC) is a clear outlier, and results from specifications dropping DC are reported in appendix tables; the results are unchanged substantively by the inclusion or exclusion of DC.

Column (2) shows the relationship between gender employment gaps and the median of male sexism, without controls for the 10^{th} and 90^{th} percentiles. Not surprisingly, the coefficient on median male sexism is very similar to that for average male sexism. In columns (3) and (4), the 10^{th} and 90^{th} percentile of male sexism, respectively, are added to the regression. Both the 10^{th} and 90^{th} percentile of male sexism are insignificantly related to female-male employment gaps. Column (5) presents the preferred specification, which jointly includes all three quantile points of the male sexism distribution. We find that while median sexism is strongly related to female-male employment gaps, these gaps are unrelated to the 10^{th} and 90^{th} percentiles of the male sexism distribution conditional on the median. This result, in light of the earlier results for wages, strongly suggests that wage gaps associated with Becker-style prejudice affect relative employment probabilities.

Were it the case that sexism only affected employment through its effect on wages, the estimates on wage and employment effects could be combined to calculate a female uncompensated labor supply elasticity. Under this assumption, median male sexism is essentially an instrument variable for female wages. Taking the ratio of the estimated effects of median male sexism on employment to that for wages, and adjusting by the female employment rate, the estimates from tables 3 and 4 imply female uncompensated labor supply elasticities that range from 0.8 to 0.9.

As noted, offer wage differences across markets arising from differences in male sexism should also lead to differences in the distribution of hours worked among working women.
In Table 5, we examine the relationship between the full-time gender employment gap and male sexism. Whereas the results presented in Table 4 defined employment based on working for pay at least one hour, the results in Table 5 define employment based on working at least 35 hours.

The results for the two dimensions of labor supply are strikingly different. When employment is defined based on working at least 35 hours, there is no relationship between the female-male employment gap and male sexism unlike the strong relationship in table 4 between median male sexism and female relative employment rates. The point estimates for both univariate regressions in the table are very close to zero. When the various quantiles of male sexism are included, the coefficient on median male sexism is negative but insignificant and about half the size as for part- and full-time employment. Furthermore, the relationship with the 10th percentile of male sexism is actually positive and significant when it is included. It does not appear that women in more sexist labor markets are any less likely to be employed full time than women in less sexist labor markets.

The results in tables 4 and 5 are based on definitions based on two particular points in the hours CDF (zero and 35 hours). By defining employment based on various cutoffs throughout the hours distribution, we can explore semi-parametrically how the female-male hours distribution changes as we move from less to more sexist labor markets. The 8 columns of table 6 present results from regressions of female-male employment gaps on median male sexism. In each column employment is defined based on working at least *X* hours, where *X* ranges from zero to 35. Each coefficient can be interpreted as the amount that the CDF of female-male hours worked declines as one compares more to less sexist labor markets. The results in columns 1 and 8 are drawn from tables 4 and 5, respectively.

The results show that moving up the CDF of female-male hours, the association between employment gaps and sexism grows weaker. In regressions that only include median male sexism, there is a strong negative relationship between sexism and gender employment gaps (i.e. working at all in the formal sector). The relationship gets progressively weaker as employment is defined based on more hours of work, ultimately reaching 0 for the relationship between male median sexism in a state and the gender full-time differential. It appears that non-employment due to male sexism is most common among women who would work less than 20 hours a week in a less sexist area. When the 10th, 50th and 90th percentiles of the male sexism distribution are all included, the coefficient on the median also decreases as employment is defined based on more hours of work. The point estimates on the median are larger in magnitude and remain negative though insignificant for full-time employment.

This pattern – in which male sexism appears to have a stronger influence on hours worked by women who work fewer hours – is consistent with the predictions of a labor supply model in which sexism reduces the offer wage. For women working a small number of hours, or who are on the margin between working and not working, income effects are negligible. With no income effect to offset the substitution effect, lower wages resulting from greater sexism in a labor market should have a relatively larger effect on labor supply. For women working more hours, the income effect may be substantial, and may completely offset the substitution effect. For these women, the effect of lower wages due to sexism on labor supply should be smaller and is ambiguous in sign.

Some of the facts we have presented thus far would seem consistent with an alternative mechanism by which sexism might affect women's labor supply behavior other than through its effect on wages. Whereas Becker, and most work in economics that follows, models taste-based discrimination as an aversion to interaction, other forms of discrimination are quite possible. Consider, for example the social penalties that sexists might impose on women in the community who violate what they consider to be appropriate female roles. Women who chose to work in such communities must overcome either explicit, or unstated but no less intensely felt, objection of men in her community. Their views about women's appropriate roles, in effect, impose a fixed cost of labor supply on women, with the cost increasing in the intensity of sexism.

Some of the patterns shown in table 6 are, in fact, consistent with the idea that men's sexism imposes a fixed cost of labor supply cost on women. For one thing, the standard labor supply model implies that higher fixed costs of labor supply lead unambiguously to lower reduced participation, exactly as we find for more sexist markets. But this notion of sexism operating as a kind of social penalty or fixed tax is not consistent with other things that we find. First, there is no reason to believe that how large the social penalty (or fixed cost) is should be related to the middle of the sexism distribution but *not* to either tail. Indeed, conditional on the sexism of the median sexist man in a community, one might well expect the sexism of the most sexist men matter importantly for any social penalty women incur by working. Because the social penalty form of sexism is not dependent on interaction at work, it is not as clear that the segregation mechanism would lessen the importance of the most sexist males' views. Furthermore, the most sexist males are probably those most likely to expend the energy making their negative views about what women ought to do known and felt.

A second consideration which lowers the plausibility of the social penalty/fixed cost explanation is that higher fixed costs have a prediction not borne out in the data at all. For a woman who would choose to work full-time, higher fixed costs of working represent a pure negative income effect. So long as leisure is a normal good, the model thus predicts we should observe an increase in hours worked in more sexist labor markets among those women do work – something we do not find in the data.

IV. The Roles of Male and Female Sexism

The empirical analysis thus far has focused on the effect of male sexism on female labor market outcomes. Economists have long wondered whether differences in labor market outcomes for women were driven by supply or demand side considerations. Inquiry in this vein has examined, for example, whether large-scale movement of women into the formal labor market was caused by supply side choices and skill upgrading by women or by reductions in discrimination in the labor market (Blau and Kahn (2006)). Recent work has documented a reversal of the trend among highly educated women and asked whether this behavior was a choice to "opt out" of the labor market (Stone (2007), Bertrand, Goldin and Katz (2008), Fortin (2009)).

The GSS data allow us to directly assess whether gender employment gaps are the result of male sexism, female choices, or some combination. As shown earlier, the GSS questions used to measure male sexism were asked of both men and women. The particular form of the questions allows us to compute indices of both male and female views of the appropriate role of women with regard to the labor market. Female responses can thus be interpreted as women's internalization of social norms. These sexist views by women may be personally developed beliefs, or they may be partly or fully influenced by male sexist preferences. The ability to consistently measure prejudicial, or stereotypical, beliefs for both the potential perpetrators and objects of discrimination is unique and fortunate.

Having a measure of female sexist views is also fortunate because it allows us to directly address a natural concern regarding the regressions presented thus far. The patterns in table 2 show that male and female sexism exhibit similar geographic patterns. One might reasonably worry that regressions that include only measures of male sexism suffer from an omitted variables bias, picking up the effect of unmeasured female preferences towards work.

Whereas the Becker model has a distinct prediction about which quantiles of the sexism distribution should be determinative on the demand side of the labor market, it has no such predictions about the role of tastes on the supply side of the market. The mechanism that leads to a "marginal discriminator" is one in which segregation and price differences are substitutes. Because individuals cannot escape themselves or their own tastes, there is no analog to segregation on the supply side of the labor market. Based on this logic, we hypothesize that the middle of the female sexism distribution should be no more strongly related to wage and employment gaps than other parts of the distribution. The regressions reported in table 7 address both this hypothesis and the questions raised in the prior two paragraphs.

Male and female sexism and wage gaps

We begin by exploring the unconditional relationship between selection-corrected gender wage gaps and the measured gender role views of women in a state. These results are presented in columns (3)-(5) of table 7. Column (3) reports the results from a regression of female-male employment gaps on average female sexism in the state. The correlation is negative and significant, and almost identical to the relationship with male sexism reproduced in column (1). The relationship with median female sexism, shown in column (4) is also negative and significant and of a similar magnitude. Not surprisingly, states in which women have stronger sex-specific views about women's appropriate role regarding work are also those in which women are relatively less likely to work in the formal labor market. In column (5), we report the results from a regression that includes the same three percentiles of the female sexism distribution that were included for men in table 3. In contrast with the results for male sexism, there is no clear relationship with any particular portion of the female sexism distribution. Also unlike the results for male sexism, the coefficient on the median is the smallest in magnitude and none of the estimates are statistically distinguishable from zero.

In columns (6)-(8) we report results from regressions that include measures of both male and female sexism. In the regression reported in column (6), the female-male wage gap is regressed on the median male and average female sexism in the state. When both are included together, the coefficient on male median sexism is reduced by 16 percent (from -0.270 to -0.226), but remains statistically significant and negative. In contrast, the coefficient on female sexism is cut by more than half and is indistinguishable from zero when we control for male median sexism. In column (7), the specification includes the 10^{th}, median and 90^{th} percentile of male sexism along with a control for median female sexism. The patterns for male sexism, shown in table 3, are not affected by the inclusion of a control for female sexism. Even conditional on how strongly women themselves feel about appropriate gender roles, female-male wage gaps are significantly negatively related with the median male sexism in the state, but not with the 10^{th} or 90^{th} percentile of male sexism. Also, conditioning on female views on gender roles reduces the point estimate on median male sexism only slightly. It appears that the results shown in table 3 are not spurious correlations driven by the unmeasured effect of female views on gender roles.

Finally in column (8) we present a specification that includes the 10^{th}, 50^{th} and 90^{th} percentiles of both male and female sexism all at once. Including the female quantiles does not affect the pattern or significance of results for male sexism. Furthermore, the female results do not match the pattern shown for male sexism. This result is interesting because the Becker model does not predict that this pattern – in which the middle of the distribution is significantly related to employment gaps while the right and left tails are not – would be present for female gender preferences. If this pattern were present one might have been concerned that some other phenomenon was behind the male sexism patterns. That it is not for female sexism lends further credence to the results for male sexism and the Becker model.

Male and female sexism and employment gaps

While the results in table 7 suggest strongly that the association between wage gaps and male sexism are not the result of a correlation with an omitted measure of female notions of gender roles, there is perhaps even more reason to be concerned about the interpretation of the relationship between male sexism and gender employment gaps. Because many of the GSS questions inquire about the appropriateness of women working, one might reasonably expect a strong relationship between women's responses to this question and their own choices about whether to work in the formal labor market.

We explore the relationship between male and female sexism and employment gaps in table 8. Columns (3) and (4) confirm that the association between female views about gender norms are indeed strongly related to female-male employment gaps. In states where women have more sexist views on average, women are less likely to work for pay relative to men. The relationship is strong – the t-statistic is 3.44; the association of employment gaps with median female sexism is similar. The employment gap's association with female sexism is no stronger, however, than with male sexism.

In column (5), we present results from a regression that includes all three quantiles of the female sexism distrubution. In contrast with the results shown in table 4 for male sexism, there is no apparent pattern across the quantiles. All three coefficients are small in magnitude and negative, and the coefficient on the median is closest to zero. It is remarkable that the pattern seen for male views about women's appropriate roles – that employment and wage gaps are strongly related to the median but not to either tail of the sexism distribution – is not there for women's views.

In column (6), we relate the female-male employment gap to both median male and female sexism. The coefficient on male median sexism is reduced by about one-quarter but remains significant; the coefficient on female median sexism is cut in half and is statistically insignificant. In columns (7) and (8) we include all three quantiles of male sexism along with the median of female sexism, and with all three quantiles of female sexism, respectively. Even conditional on how women themselves feel about appropriate gender roles, female-male employment gaps are strongly related to the median of male sexism but not with the 10^{th} or 90^{th} percentile of male sexism. Remarkably, the coefficient on male median sexism is the same in the full model in column (8) as it is when median male sexism is included alone. In the full model, none of the female sexism quantiles are significantly related to gender employment gaps. These results suggest that: (a) the patterns seen for male sexism in tables 3 and 4 were not the result of female sexism being an omitted variable, and (b) low female employment rates observed in some areas of the US may be more the result of male sexism than of female choices to "opt out".

V. Extensions: Accounting for alternative explanations

Having shown the strong basic relationship between female labor market outcomes and both male and female sexism, we now address the possibility that these patterns are merely spurious correlations and explore whether the patterns shown so far are sensitive to alternative specifications.

Accounting for unobserved differences in skills

We begin by addressing whether the relationship between state-level differences in labor market outcomes are the result not of male sexism but rather unobserved differences in female skills that happen to be correlated with the level of male sexism across states. Before examining other empirical results, we pause to discuss how the results shown thus far speak to this concern. The results shown so far control for female-male differences in education and experience, so any unobserved skill would have to be conditional on these standard measures of labor market skill. Furthermore, for such unobserved differences in female skill to explain the patterns in employment and wages we have shown, it would have to be the case that these skill differences were correlated with the median but not the 10th or 90th percentiles of male sexism. And, it would have to be the case that these unobserved female skill differences were not correlated strongly with female views on gender roles because the main results are unaffected by including these controls.

To further address this concern about potential unmeasured skill, the regressions in table 9 add controls for the female-male difference in National Assessment of Educational Progress Long Term Trends (NAEP-LTT) math and reading scores in each state. The NAEP-LTT is a set of standardized tests given to a random sample of U.S. 9-, 13- and 17-year old students enrolled in public schools. Tests in math or reading have been given almost every other year since 1971. We calculate the difference in normalized scores (i.e. z-scores) between females and males in each state for math and reading and include both as controls. The results for wage gaps, shown in table 9, are essentially unchanged by the inclusion of these controls for differences in academic knowledge. The estimated effects of NAEP-LTT skills are insignificant and many are of the opposite sign one would predict. The results for employment gaps, not shown, are also essentially unchanged by the inclusion of these controls.

Alternative imputation rules for to account for selection into the workforce

As discussed above, it is particularly important to correct for selection when analyzing variation in female wages. The results in table 3 correct for selection by estimating median regressions after imputing wages either above or below the median based on whether the individual is in the top or bottom quartile of the gender-specific education distribution. In table 10, we show results for the base wage gap specification with alternative imputation rules. The results show that the pattern presented in table 3 is not sensitive to these changes in how wages are imputed to account for selection.

In columns (1) to (3) wages are imputed for non-employed individuals, while in columns (4) to (6) wages are imputed for non full-time employed individuals. Columns (1) and (4) impute log wages below the median if the number of years of education is less than the 10th percentile of the education distribution for each gender in each year and impute log wages above the median if education is greater than the 90th percentile. Columns (2) and (5) are similar except that the 25th

and 75th percentiles are used as the cut-offs. The results in column (5) correspond to the results in column (5) of table 3. Columns (3) and (6) restrict the sample to individuals who are household heads or spouses of household heads. For males and females with no recorded spouse wage, the imputation rule follows columns (2) and (5). For females with spouse wage information, we impute log wages below the median if the number of years of education is in the bottom quartile of the female education distribution in that year and log spouse wage is in the bottom quartile of the spouse wage distribution in that year. Log wages are imputed above the median if number of years of education is in the top quartile of the female education distribution and log spouse wage is in the top quartile of the spouse wage distribution. In each of the six regressions, the female-male wage gap is strongly and significantly negatively related to median male sexism, but not with the 10^{th} or 90^{th} percentile of male sexism

Sensitivity to the inclusion of outlier states

The scatter plot shown in figure 6 shows a strong negative relationship between average male prejudice and the conditional female-male employment gap. The other obvious feature of the scatter plot is that the District of Columbia (DC) is a clear outlier. So motivated, we have run all of the results shown after dropping DC. The results are largely insensitive to dropping these observations. We show the results for the base wage gap specifications dropping in table 11. As can be seen in the table, all of the basic patterns are the same as those reported in table 3.

Taking the middle of the sexism distribution to be the 40^{th} percentile

The Becker discrimination model predicts that for gender wage gaps, the marginal discriminator should be drawn from the middle of the sexism distribution. For the purposes of the empirical work and to be comparable with CG, we have defined the middle to be the median thus far. While women are now close to making up half of the labor force, they on average accounted for about 40 percent of labor markets during the period under study, suggesting that some percentile point smaller than the median is where the marginal discriminator will be. In table 12, we present results from models that replace the median with the 40^{th} percentile of male sexism. Maybe not surprisingly, the results are quite similar to those reported for the median in table 3.

Measuring sexism only among the high-skilled

Throughout our analysis thus far, we have measured sexism among all men, rather than among employers. This decision is consistent with arguments made in Charles and Guryan (2008) that, since the decision to become an employer is endogenous, the relevant prejudice is that found in the population overall. Nonetheless, one could argue that the views of those whose skills would prevent them from being employers should not determine the extent of market discrimination. In table 13, we present results from models in which sexism is measured only among those with at least a college education. We take education to be exogenous and treat college-educated persons as those with the skills necessary to be an employer. Because sample sizes get small once we restrict to college-educated individuals, we combine male and female college graduate to measure the quantiles of sexism. The results show similar patterns to those reported in table 3. For the wage gap, the coefficient on median sexism is negative and the largest in magnitude, though it is smaller than in the base specification and insignificant. For the employment gap, the coefficient on median sexism is negative and significant, while the coefficients on the 10th and 90th percentiles are close to zero and insignificant.

VI. Conclusion

We examine the relationship between male and female views on the appropriate roles women should play, which we call sexism, and the labor market outcomes of women. We motivate this examination with an observation that there is a great deal of variation across areas of the US in both wages and employment rates of women, which we document. The main goals of the paper have been to document the relationship between measured sexism and female labor market outcomes, and to test to what extent Becker's model of taste-based discrimination can account for this relationship. To do this, we collect data on male and female responses to relevant questions from the GSS and compute measures of the distribution of sexism in US states.

The empirical patterns we find are remarkably consistent with the predictions of the Becker model. We show that female-male wage gaps are significantly larger, i.e. female wages are relatively lower, in states with more sexist males. Gender wage gaps are not, however, simply related to the average amount of sexism in a labor market. Rather, they are significantly related to the degree of sexist views held by the median male, and not to the degree of sexism held by the most or least sexist males in the state. This pattern is just as predicted by the Becker discrimination model, and by the many variations that have followed in the literature. Whereas the marginal discriminator against blacks should be drawn from the left, or least prejudiced, tail of the racial prejudice distribution, the prevalence of women in the labor market makes it more difficult for market induced segregation to separate sexists from the female objects of their discriminatory tastes. As a result, Becker discrimination models predict that the marginal discriminator against women should be drawn from the middle of the sexism distribution.

We extend the usual focus on wage gaps to other labor market outcomes for women. We show that female-male employment gaps, outcomes of particular interest when it comes to females in the labor market, are also negatively related to the median amount of male sexism, but not with the 10th or 90th percentile. We describe how this result would be predicted as a direct labor supply response to reduced female wages resulting from sexism. We then examine hours worked and show that the reduced labor supply by women in sexist labor markets is concentrated among those who would have worked few hours in sexism's absence. We show that this pattern is just as predicted by income and substitution effects in a standard labor supply model applied to sexism-reduced wages.

Finally, we show that the patterns described for male sexism are not apparent for female responses to the GSS questions. We take this as supportive evidence for the model, which predicts the patterns for the quantiles of sexism only among males. We also show that adding controls for gender-skill differences, alternative strategies for wage selection corrections, and various other sensitivity analyses do not appreciably affect the results.

The paper has focused squarely on the question of whether taste-based discrimination against women, or sexism, plays an important role in the labor market success of women in the U.S. The results suggest strongly that indeed sexism has important implications for the workings of labor markets for men and women. The significant differences across areas of the country in the employment and pay of women relative to men appear to be driven in part by differences in sexist preferences.

The results of the paper are also notable when compared with the results in Charles and Guryan (2008), the paper to which it is closest in spirit. In many ways, the analyses of prejudice and sexism, black-white and female-male gaps, are similar. In both racial prejudice and sexism there is a great deal of geographic variation. In both cases, labor market gaps are significantly related with the degree of animus by discriminators. There is, however, a key difference that is just as predicted by models of taste-based discrimination. Charles and Guryan (2008) emphasize that a key feature of taste-based discrimination models is that wage gaps and segregation are substitutes. Markets generate incentives to segregate, or separate, and the more this happens, the smaller are wage gaps in equilibrium. A key difference in labor markets between blacks and women is that women are not a significant minority. It is thus harder for the market to separate women from those who hold animus towards them. Taste-based discrimination models are clear in their prediction that, all else equal, this prevalence should make the marginal discriminator towards women someone relatively more discriminatory than the marginal discriminator against blacks. That we find results that are different here than in Charles and Guryan (2008), but in precisely the way predicted by the model, speaks to the power of the model and to the possible widespread importance of taste-based discrimination in labor markets.

Bibliography

Altonji, Joseph G. and Rebecca M. Blank (1999). "Race and Gender in the Labor Market," in O.A. Ashenfelter and D. Card, eds. *Handbook of Labor Economics,* Volume 3, pp. 3143-3259.

Arrow, Kenneth (1973). "The Theory of Discrimination," in O.A. Ashenfelter and A. Rees, eds. *Discrimination in Labor Markets*. Princeton, NJ: Princeton University Press, pp. 3-33.

Becker, Gary S. (1957, 1971). *The Economics of Discrimination*. Chicago: The University of Chicago Press, 2nd edition.

Becker, Gary S. (1981). *A Treatise on the Family*. Cambridge, Massachusetts: Harvard University Press.

Bertrand, Marianne, Claudia Goldin and Lawrence F. Katz (2008). "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors," Working Paper.

Black, Dan A. 1995. "Discrimination in an Equilibrium Search model," *Journal of Labor Economics*, 13(2):309-34.

Blau, Francine D. and Lawrence Kahn (2000). "Gender Differences in Pay," *Journal of Economic Perspectives,* 14(4): 75-99.

Blau, Francine D. and Lawrence Kahn (2006). "The US Gender Pay Gap in the 1990s: Slowing Convergence", *Industrial Relations and Labor Relations Review*, 60(1): 45-66.

Charles, Kerwin Kofi and Jonathan Guryan (2008). "Prejudice and Wages: An Empirical Assessment of Becker's *The Economics of Discrimination*," *Journal of Political Economy*, 116(5): 773-809.

Fortin, Nicole M. (2009). "Gender Role Attitudes and Women's Labor Market Participation: Opting-Out, AIDS, and the Persistent Appeal of Housewifery," Working Paper.

Neal, Derek (2004). "The Measured Black-White Wage Gap Among Women is Too Small," *Journal of Political Economy*, 112(2):S1-S28.

Olivetti, Claudia and Petrongolo, Barbara (2008). "Unequal Pay or Unequal Employment? A Cross-Country Analysis of Gender Gaps," *Journal of Labor Economics*, 26(4): 621-654.

Stone, Pamela (2007). *Opting Out? Why Women Really Quit Careers and Head Home.* Berkeley: University of California Press.

Welch, Finis (1967). "Labor Market Discrimination: An Interpretation of Income Differences in the Rural South," *Journal of Political Economy*, 75(3):225-240.

Figure 1: State-level Plots of Residual Female-Male Employment Gaps and Residual Female-Male Wage Gaps

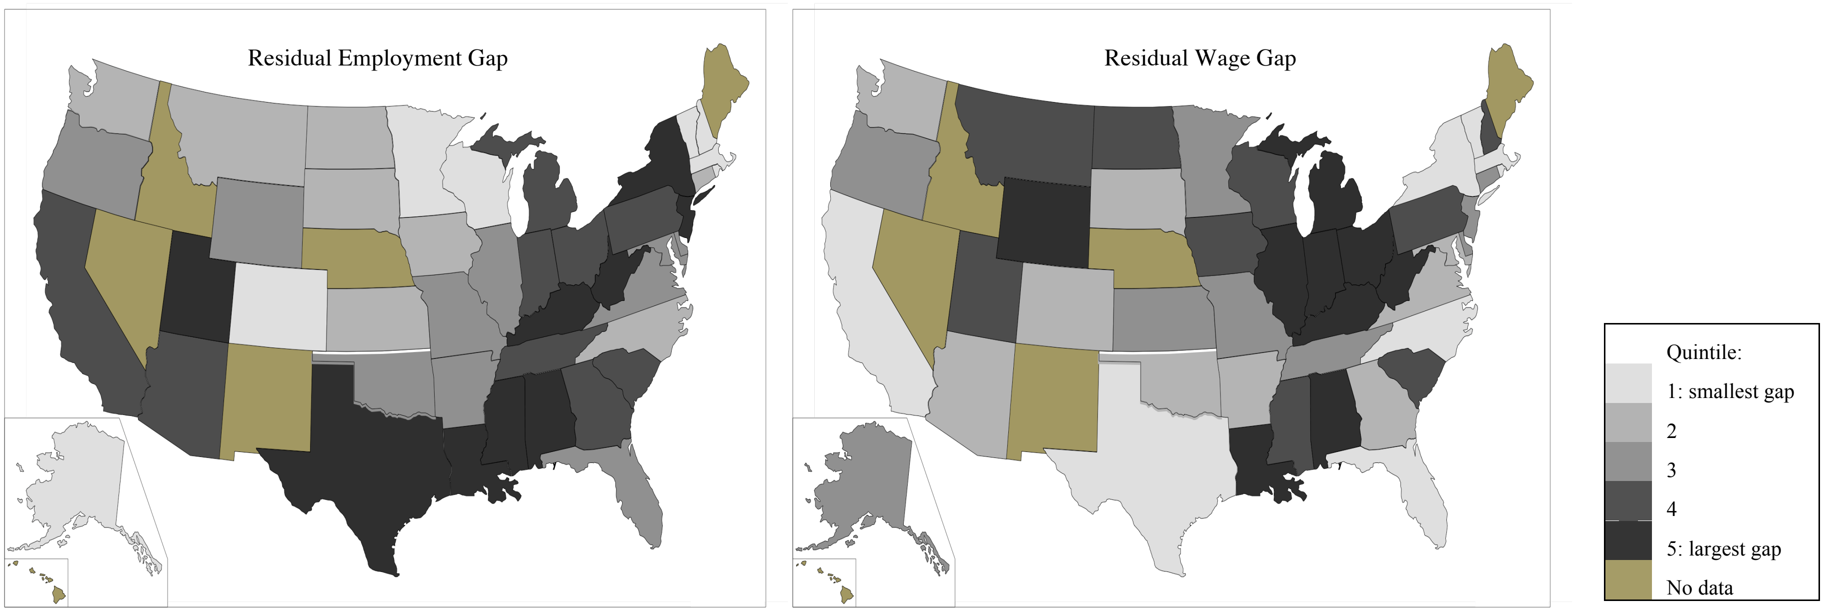

Note: The figure documents variation in residual female-male gaps in wages and gaps in employment at the state level. The gaps are from regressions that are estimated on CPS MORG/May data from 1977-2002 and include controls for education, a quadratic in experience, an indicator for marital status, gender-specific year effects, and gender-specific state effects. The estimated state*female coefficients are grouped into 5 separate quintiles and the colors on the state map correspond to the quintile of the corresponding employment or wage gap for each state. Lighter colors indicate states with smaller female-male gaps (less negative state*female coefficients), while darker colors indicate states with larger gaps female-male gaps (more negative state*female coefficients). The tan color represents states in which there is no data available in the GSS.

Figure 2A: Relationship between conditional white female-male employment gaps and average male response to two questions in the GSS

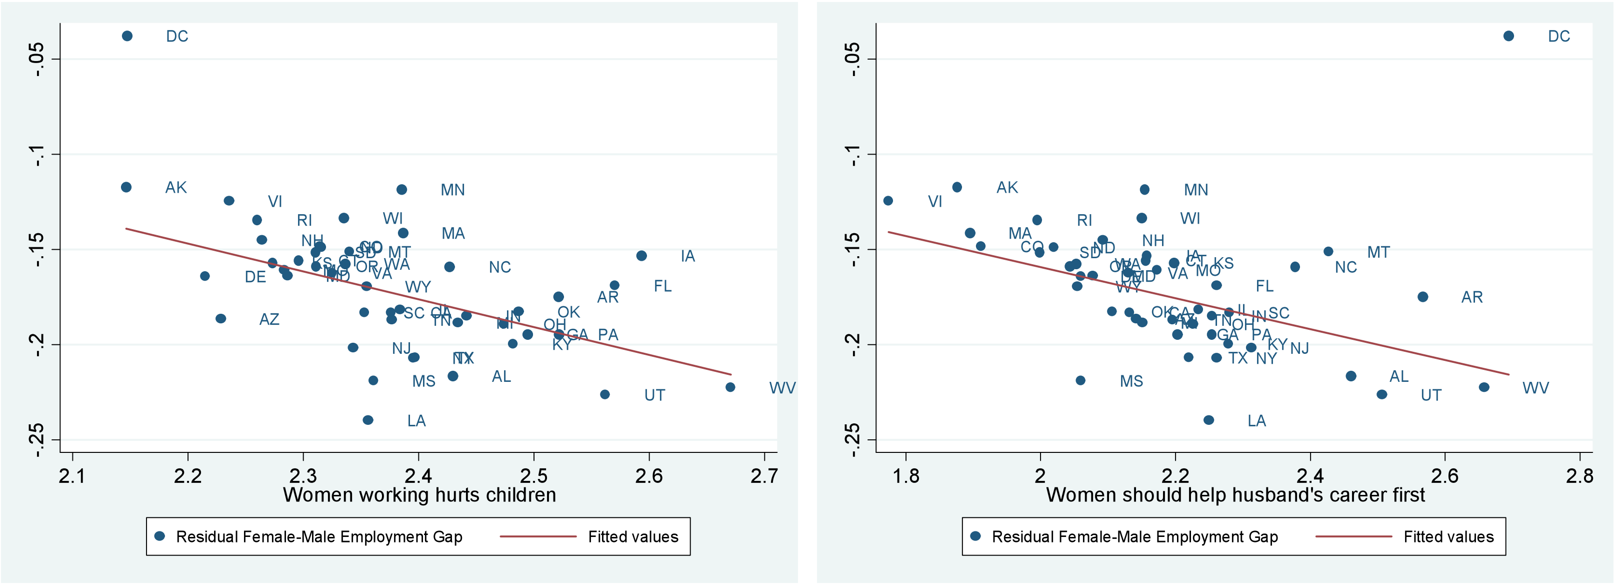

Notes: The figure plots residual female-male employment gaps at the state level against average responses by white males to two questions in the GSS. The questions are listed in Appendix table 1. The gaps are from regressions that are estimated on CPS MORG/May data from 1977-2002 and include controls for education, a quadratic in experience, an indicator for marital status, gender-specific year effects, and gender-specific state effects. The figure plots the state*female coefficients. The excluded female*year is 1990, so the levels of the gaps are for that year. Excluding later years would produce smaller gaps, but the variation across states would be the same.

Figure 2B: Relationship between conditional white female-male selection-corrected wage gaps and average male response to two questions in the GSS

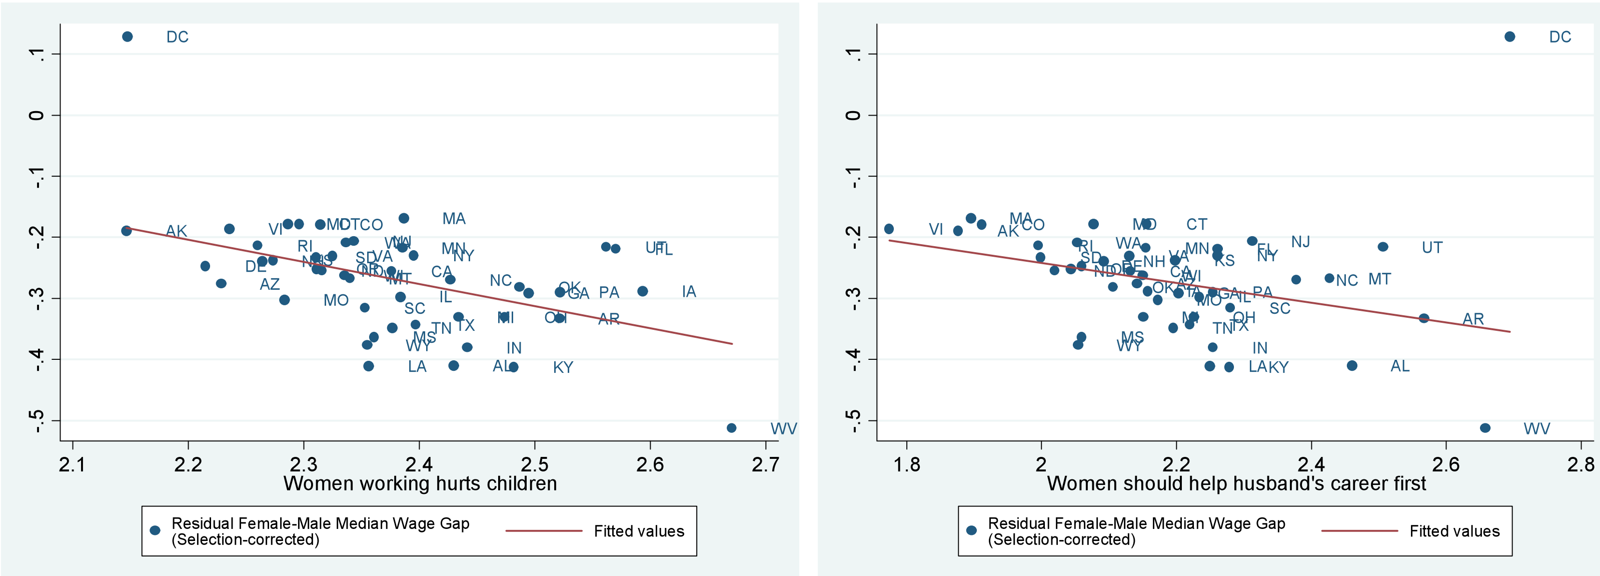

Notes: The figure plots residual female-male selection-corrected wage gaps at the state level against average responses by white males to two questions in the GSS. The questions are listed in Appendix table 1. The gaps are from regressions that are estimated on CPS MORG/May data from 1977-2002 and include controls for education, a quadratic in experience, an indicator for marital status, gender-specific year effects, and gender-specific state effects. The figure plots the state*female coefficients. The excluded female*year is 1990, so the levels of the gaps are for that year. Excluding later years would produce smaller gaps, but the variation across states would be the same.

Figure 3: Cumulative distribution function of the male sexism index across nine census divisions

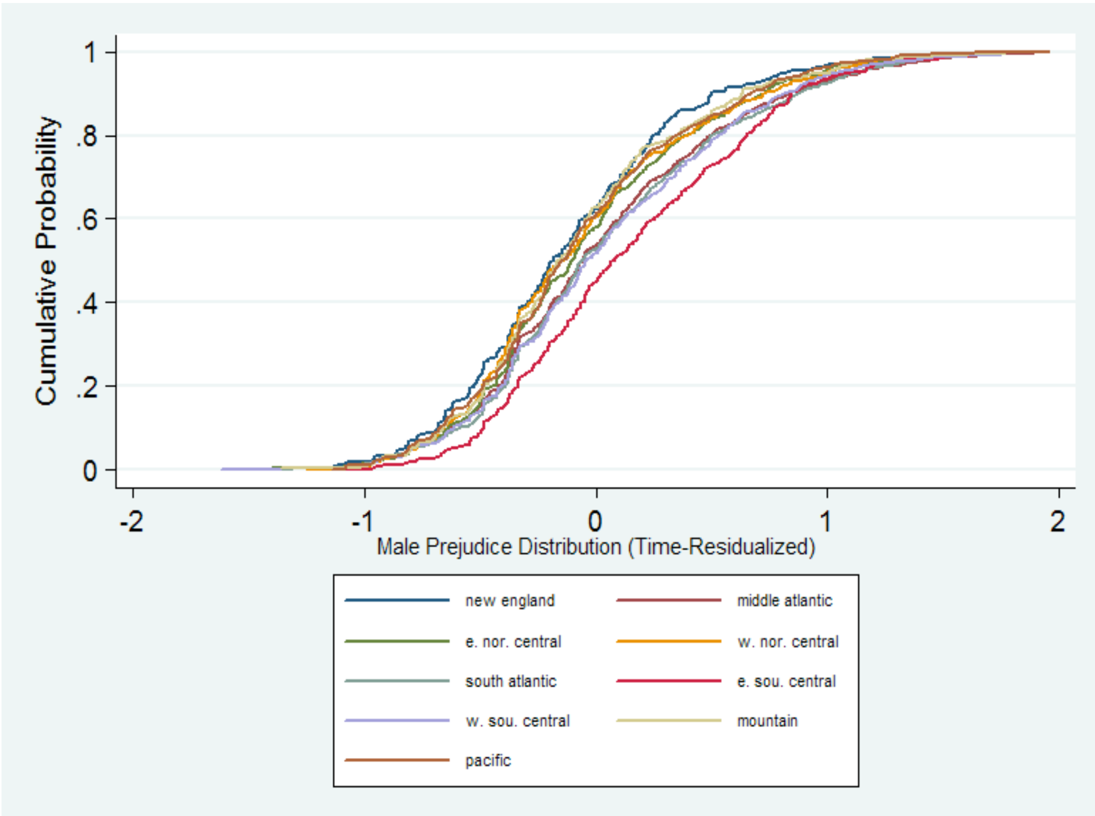

Notes: The figure plots the cumulative distribution function of male sexism in each of the nine U.S. Census divisions. Male sexism is measured with an index constructed as described in the text from the eight GSS questions listed in Appendix table 1.

Figure 4: Trends over time in the average index of male sexism by census divisions

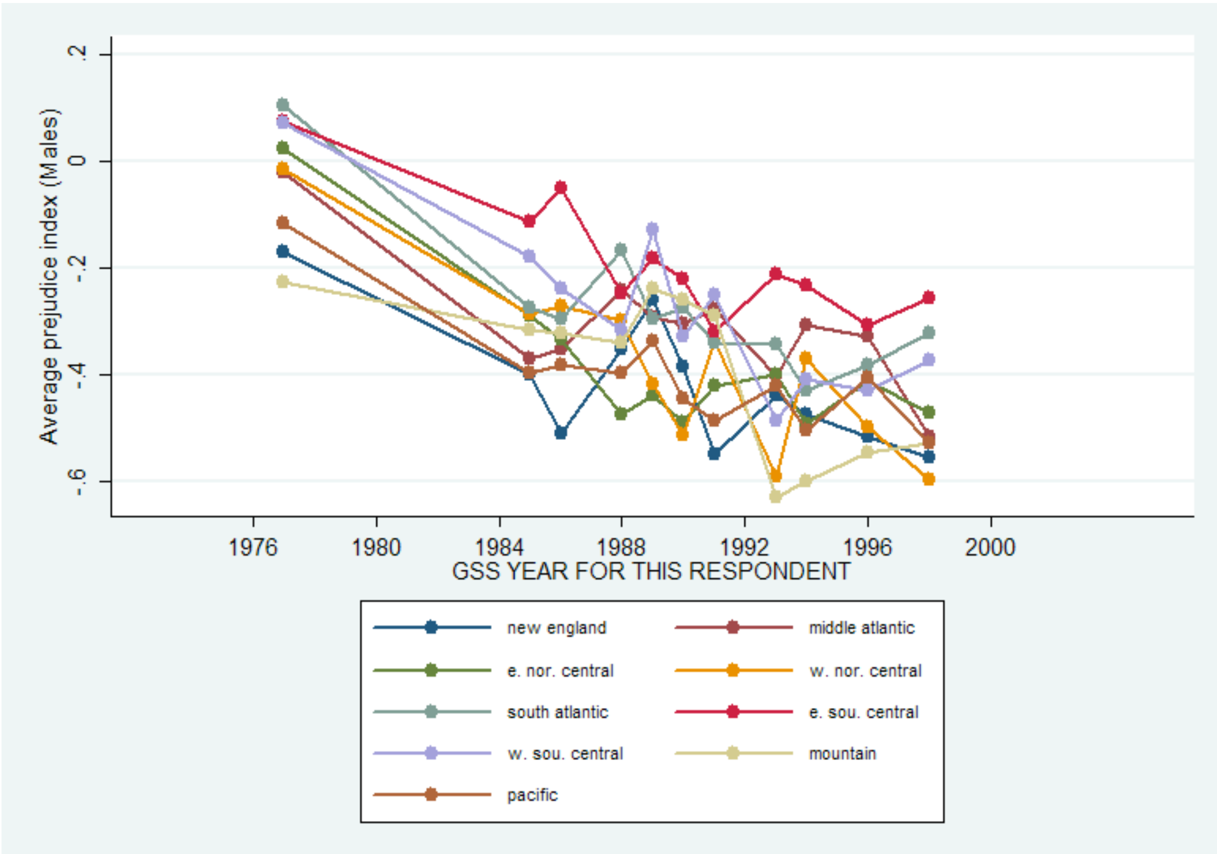

Notes: The figure plots average male sexism over time for each of the nine U.S. Census divisions. Averages are plotted for each of the years in which the eight GSS sexism questions were asked.

Figure 5: Trends in male average responses to each of the GSS questions in the sexism index

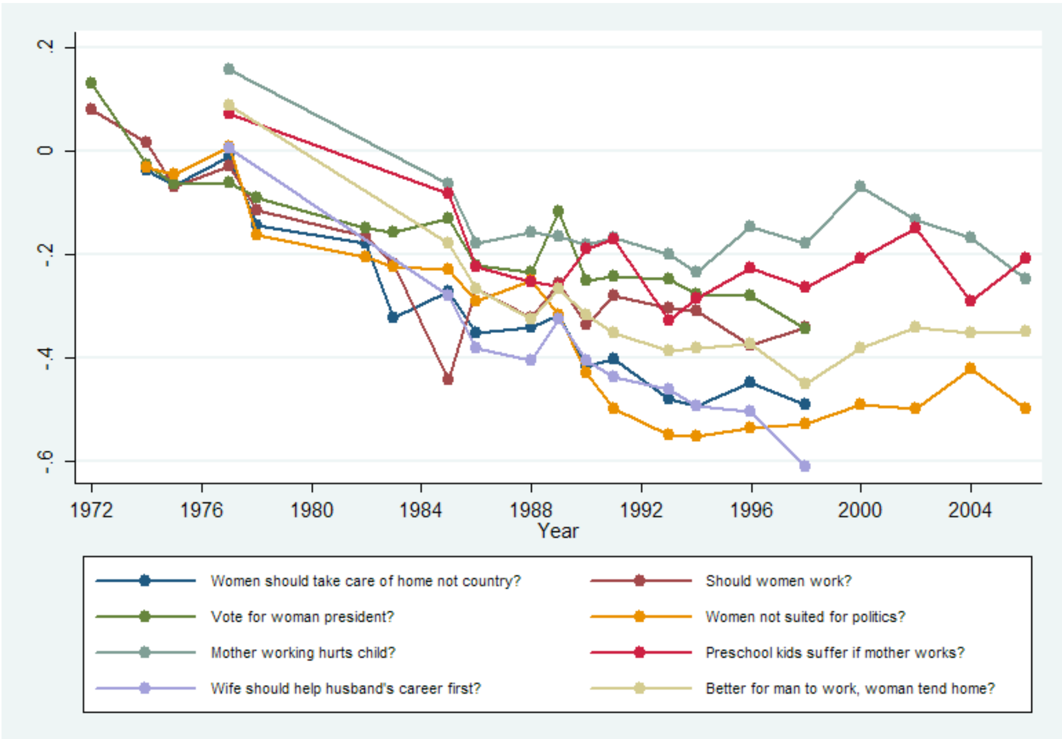

Notes: The figure plots average responses by white males to each of the eight GSS sexism questions over time. Responses are plotted for years that are not used in the regression analysis because that analysis is restricted to years in which all eight questions are asked together. Responses are normalized in 1977 so that the average response in that year by men and women together is zero.

Figure 6: Scatter plot of female-male employment gaps against average male sexism

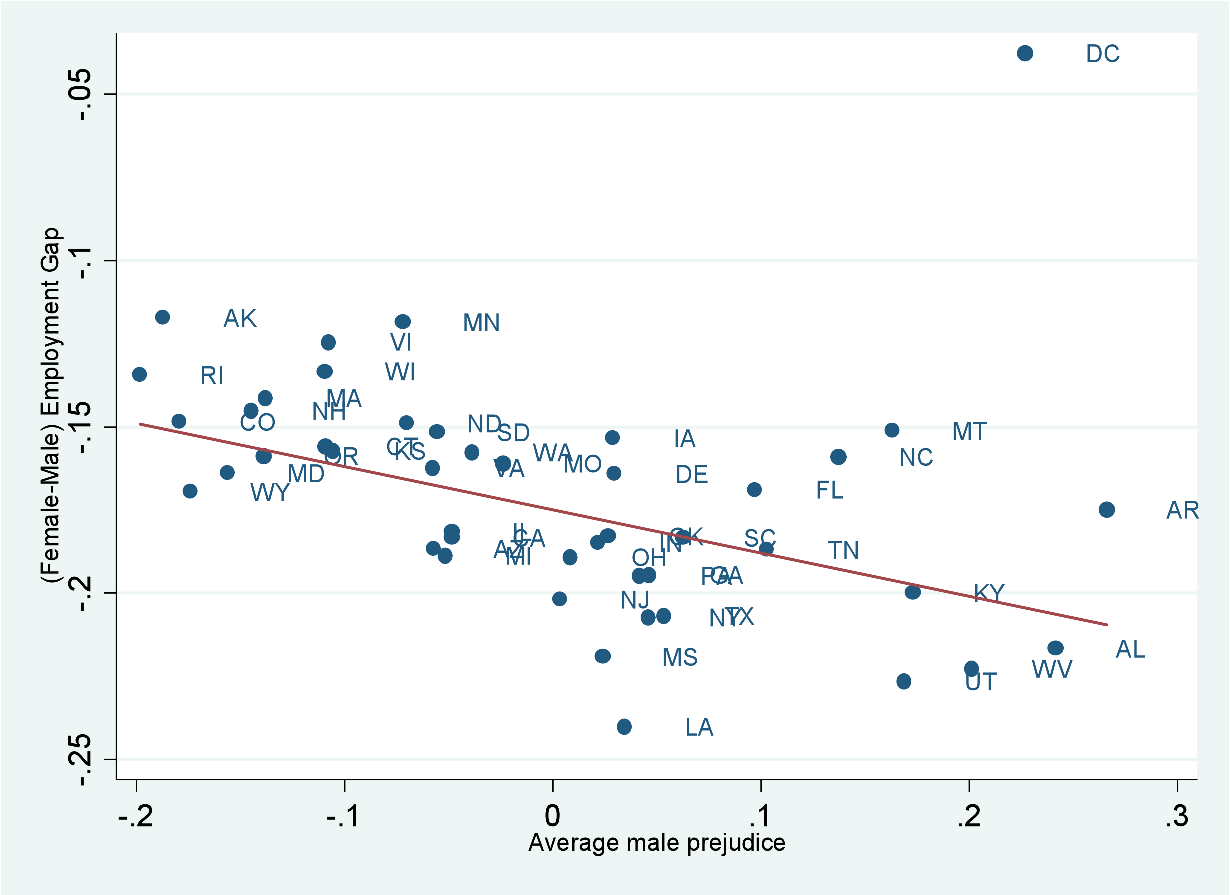

Notes: The figure plots residual female-male employment gaps at the state level against average male sexism. The gaps are from regressions that are estimated on CPS MORG/May data from 1977-2002 and include controls for education, a quadratic in experience, an indicator for marital status, gender-specific year effects, and gender-specific state effects. The figure plots the state*female coefficients. The excluded female*year is 1990, so the levels of the gaps are for that year. Excluding later (earlier) years would produce smaller (larger) gaps, but the variation across states would be the same.

Table 1A: Cross-state variation in residual female-male employment gaps

	Residual female-male employment gap			
	Whites			
	All	HS or Less	Some college or more	Blacks
Full period: 1977 to 2002 (relative to 1990)				
Mean	-0.170	-0.207	-0.137	-0.072
SD	0.035	0.040	0.028	0.049
Max-Min	0.202	0.210	0.193	0.260
Earlier period: 1977 to 1989 (relative to 1983)				
Mean	-0.257	-0.292	-0.214	-0.122
SD	0.039	0.043	0.030	0.081
Max-Min	0.250	0.267	0.220	0.436
Later period: 1990 to 2002 (relative to 1996)				
Mean	-0.152	-0.200	-0.123	-0.078
SD	0.037	0.048	0.029	0.037
Max-Min	0.163	0.193	0.172	0.153
Number of states	45	45	45	45

Notes: Residual female-male employment gaps are estimated using the 1977-2002 May/ORG CPS data and control for education, a quadratic in potential experience, marital status, gender-specific year effect and state effects. Data from 1973-1976 are dropped because CPS reports states in groups in those years. States are dropped if they are not sampled in the GSS in the years necessary to measure male prejudice.

Table 1B: Cross-state variation in residual female-male wage gaps

	Residual female-male wage gap			
	Whites			
	All	HS or Less	Some college or more	Blacks
Full period: 1977 to 2002 (relative to 1990)				
Mean	-0.315	-0.355	-0.262	-0.183
SD	0.044	0.057	0.035	0.059
Max-Min	0.298	0.345	0.243	0.289
Earlier period: 1977 to 1989 (relative to 1983)				
Mean	-0.388	-0.432	-0.322	-0.227
SD	0.046	0.063	0.037	0.076
Max-Min	0.286	0.376	0.212	0.405
Later period: 1990 to 2002 (relative to 1996)				
Mean	-0.265	-0.308	-0.226	-0.151
SD	0.044	0.053	0.037	0.079
Max-Min	0.302	0.306	0.255	0.456
Number of states	45	45	45	45

Notes: Residual female-male wage gaps are estimated for all employed workers using the 1977-2002 May/ORG CPS data and control for education, a quadratic in potential experience, marital status, gender-specific year effect and state effects. Data from 1973-1976 are dropped because CPS reports states in groups in those years. States are dropped if they are not sampled in the GSS in the years necessary to measure male prejudice.

Table 2A: Male Sexism across Census Regions

	Should women work?	Vote for women president?	Mother working hurts child?	Wife should help husband's career first?	Better for man to work, woman to tend home?	Average Sexism Index	
						Overall	High Skilled
East South Central	0.276	0.197	2.418	2.267	2.593	0.144	-0.079
West South Central	0.259	0.185	2.438	2.238	2.505	0.065	-0.077
South Atlantic	0.252	0.153	2.430	2.259	2.484	0.050	-0.099
Middle Atlantic	0.232	0.140	2.381	2.269	2.481	0.036	-0.176
East North Central	0.236	0.127	2.408	2.200	2.354	-0.035	-0.200
West North Central	0.254	0.111	2.336	2.142	2.305	-0.053	-0.178
Pacific	0.195	0.116	2.374	2.109	2.331	-0.058	-0.157
Mountain	0.204	0.102	2.291	2.088	2.283	-0.085	-0.122
New England	0.184	0.086	2.313	1.996	2.241	-0.131	-0.249

Note: The first five columns of the table show the average responses by white males for each of the nine Census divisions to five of the GSS questions measuring sexist attitudes among men. The first two questions have binary answers, whereas the latter three have answers that range from 1 to 5. Responses are coded so that higher numbers denote more sexist attitudes. The table also shows for each of the nine Census divisions the average sexism index for all white men and for college-educated white men.

Table 2B: Female Sexism across Census Regions

	Should women work?	Vote for women president?	Mother working hurts child?	Wife should help husband's career first?	Better for man to work, woman to tend home?	Average Prejudice Index	
						Overall	High Skilled
East South Central	0.261	0.221	2.189	2.234	2.484	0.186	-0.061
West South Central	0.249	0.163	2.054	2.181	2.371	0.043	-0.158
South Atlantic	0.228	0.171	2.093	2.191	2.341	0.029	-0.214
West North Central	0.239	0.148	2.074	2.071	2.231	-0.010	-0.232
East North Central	0.226	0.127	2.084	2.106	2.261	-0.035	-0.297
Mountain	0.236	0.112	2.036	1.998	2.235	-0.036	-0.231
Middle Atlantic	0.200	0.111	2.096	2.062	2.215	-0.056	-0.249
Pacific	0.185	0.102	2.104	2.008	2.224	-0.079	-0.224
New England	0.157	0.087	2.008	1.992	2.140	-0.135	-0.380

Note: The first five columns of the table show the average responses by white females for each of the nine Census divisions to five of the GSS questions measuring sexist attitudes among women. The first two questions have binary answers, whereas the latter three have answers that range from 1 to 5. Responses are coded so that higher numbers denote more sexist attitudes. The table also shows for each of the nine Census divisions the average sexism index for all white women and for college-educated white women.

Table 3: Male Sexism and Selection-corrected Gender Wage Gaps

	Impute wages for non fulltime employed persons				
Male Sexism	(1)	(2)	(3)	(4)	(5)
Average	-0.234				
	[0.095]*				
10th Percentile			0.101		0.070
			[0.132]		[0.132]
50th Percentile		-0.27	-0.326	-0.385	-0.418
		[0.082]**	[0.110]**	[0.109]**	[0.126]**
90th Percentile				0.112	0.106
				[0.072]	[0.074]
Observations	45	45	45	45	45
R-squared	0.12	0.2	0.21	0.25	0.25

Notes: Table reports coefficients (standard errors) from OLS regressions of residual state-level female-male wage gaps on various measures of sexism among all males. Residual female-male wage gaps are estimated by median regression after imputing log wages for men and women who are non-employed [Columns (1)-(5)] or non full-time employed [Columns (6)-(10)]. Log wages are imputed based on the following rule: Impute log wages equal to zero if number of years of education is less than the bottom quartile of the education distribution for each gender in each year. Impute log wages equal to 20 if education is greater than the top quartile of the education distribution for each gender in each year. Just as in the base specifications, female-male wage gaps are estimated using 1977-2002 May/ORG CPS data and control for education, a quadratic in experience, marital status, gender-specific year effects and state effects. Data from 1973-1976 are dropped because CPS reports states in groups in those years. States are dropped if they are not sampled in the GSS in the years necessary to measure male prejudice. *-5%, **-1% significance.

Table 4: Male Sexism and Gender Employment Gaps

	Residual Female-Male Employment Gap				
Male Sexism	(1)	(2)	(3)	(4)	(5)
Average	-0.124				
	[0.037]**				
10th Percentile			0.042		0.036
			[0.052]		[0.053]
50th Percentile		-0.126	-0.149	-0.152	-0.169
		[0.032]**	[0.043]**	[0.044]**	[0.050]**
90th Percentile				0.026	0.023
				[0.029]	[0.030]
Observations	45	45	45	45	45
R-squared	0.21	0.26	0.27	0.28	0.29

Notes: Table reports coefficients (standard errors) from OLS regressions of residual state-level female-male employment gaps on various measures of sexism among all males. Residual female-male employment gaps are estimated using the 1977-2002 May/ORG CPS data and control for education, a quadratic in potential experience, marital status, gender-specific year effect and state effects. Data from 1973-1976 are dropped because CPS reports states in groups in those years. States are dropped if they are not sampled in the GSS in the years necessary to measure male prejudice. *-5%, **-1% significance.

Table 5: Male Sexism and Full-time Gender Employment Gaps

	Residual Female-Male Full-time Employment Gap				
Male Sexism	(1)	(2)	(3)	(4)	(5)
Average	0.020				
	[0.042]				
10th Percentile			0.131		0.128
			[0.058]*		[0.059]*
50th Percentile		0.000	-0.072	-0.020	-0.080
		[0.038]	[0.048]	[0.051]	[0.056]
90th Percentile				0.020	0.009
				[0.034]	[0.033]
Observations	45	45	45	45	45
R-squared	0.01	0	0.11	0.01	0.11

Notes: Table reports coefficients (standard errors) from OLS regressions of residual state-level female-male full-time employment gaps on various measures of sexism among all males. Residual female-male employment gaps are estimated using the 1977-2002 May/ORG CPS data and control for education, a quadratic in potential experience, marital status, gender-specific year effect and state effects. Data from 1973-1976 are dropped because CPS reports states in groups in those years. States are dropped if they are not sampled in the GSS in the years necessary to measure male prejudice. *-5%, **-1% significance.

Table 6: The Effect of Male Sexism on Different Points of the Hours Worked Distribution

	employed (≥ 0 hours)	≥ 5 hours	≥ 10 hours	≥ 15 hours	≥ 20 hours	≥ 25 hours	≥ 30 hours	≥ 35 hours
Panel I. Median male sexism								
50th percentile	-0.126	-0.112	-0.105	-0.092	-0.074	-0.037	-0.020	0.000
	[0.032]**	[0.032]**	[0.032]**	[0.033]**	[0.033]*	[0.036]	[0.037]	[0.038]
Panel II. Moments of male sexism distribution								
10th percentile	0.036	0.048	0.057	0.066	0.084	0.104	0.122	0.128
	[0.053]	[0.053]	[0.053]	[0.053]	[0.054]	[0.057]	[0.059]*	[0.059]*
50th percentile	-0.169	-0.149	-0.145	-0.14	-0.133	-0.111	-0.101	-0.08
	[0.050]**	[0.051]**	[0.050]**	[0.051]**	[0.051]*	[0.054]*	[0.056]	[0.056]
90th percentile	0.023	0.011	0.009	0.011	0.013	0.016	0.014	0.009
	[0.030]	[0.030]	[0.030]	[0.030]	[0.030]	[0.032]	[0.033]	[0.033]

Note: The table shows estimates from regressions in which the dependent variables are female-male gaps in the fraction that work at least X hours where X is equal to 0, 5, 10, 15, 20, 25, 30, or 35. Panel I shows results from univariate regressions with the median of male sexism as the regressor, while panel II shows results from regressions that include the 10^{th}, 50^{th}, and 90^{th} percentiles of male sexism. The regressions shown in the first column correspond to the regressions shown in columns 2 and 5 of Table 4. Together, the regressions show how median male sexism relates to the difference in CDFs of hours worked by females and males in the state. *-5%, **-1% significance.

Table 7: Male and Female Sexism and Selection-corrected Gender Wage Gaps

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Male Sexism								
Average	-0.234							
	[0.095]*							
10th Percentile							0.090	0.139
							[0.134]	[0.145]
50th Percentile		-0.270				-0.226	-0.371	-0.349
		[0.082]**				[0.107]*	[0.137]**	[0.142]*
90th Percentile							0.109	0.104
							[0.074]	[0.079]
Female Sexism								
Average			-0.236					
			[0.086]**					
10th Percentile					-0.080			0.004
					[0.175]			[0.182]
50th Percentile				-0.200	-0.043	-0.065	-0.088	-0.044
				[0.080]*	[0.175]	[0.100]	[0.100]	[0.181]
90th Percentile					-0.109			-0.085
					[0.070]			[0.076]
Observations	45	45	45	45	45	45	45	45
R-squared	0.12	0.2	0.15	0.13	0.18	0.21	0.27	0.29

Notes: Table reports coefficients (standard errors) from OLS regressions of residual state-level female-male wage gaps on various measures of sexism among males and females. Residual female-male wage gaps are estimated by median regression after imputing log wages for men and women who are non full-time employed. Log wages are imputed based on the following rule: Impute log wages equal to zero if number of years of education is less than the bottom quartile of the education distribution for each gender in each year. Impute log wages equal to 20 if education is greater than the top quartile of the education distribution for each gender in each year. Just as in the base specifications, female-male wage gaps are estimated using 1977-2002 May/ORG CPS data and control for education, a quadratic in experience, marital status, gender-specific year effects and state effects. Data from 1973-1976 are dropped because CPS reports states in groups in those years. States are dropped if they are not sampled in the GSS in the years necessary to measure male prejudice. *-5%, **-1% significance.

Table 8: Male and Female Sexism and Gender Employment Gaps

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Male sexism								
Average	-0.124							
	[0.037]**							
10th Percentile							0.050	0.068
							[0.053]	[0.057]
50th Percentile		-0.126				-0.092	-0.138	-0.126
		[0.032]**				[0.042]*	[0.054]*	[0.056]*
90th Percentile							0.025	0.016
							[0.029]	[0.031]
Female sexism								
Average			-0.117					
			[0.034]**					
10th Percentile					-0.068			-0.053
					[0.069]			[0.072]
50th Percentile				-0.105	-0.028	-0.050	-0.059	-0.008
				[0.031]**	[0.069]	[0.039]	[0.040]	[0.072]
90th Percentile					-0.03			-0.023
					[0.027]			[0.030]
Observations	45	45	45	45	45	45	45	45
R-squared	0.21	0.26	0.22	0.21	0.24	0.29	0.32	0.34

Notes: Table reports coefficients (standard errors) from OLS regressions of residual state-level female-male employment gaps on various measures of sexism among males and females. Residual female-male employment gaps are estimated using the 1977-2002 May/ORG CPS data and control for education, a quadratic in potential experience, marital status, gender-specific year effect and state effects. Data from 1973-1976 are dropped because CPS reports states in groups in those years. States are dropped if they are not sampled in the GSS in the years necessary to measure male prejudice. *-5%, **-1% significance.

Table 9: Male Sexism and Selection-Corrected Gender Wage Gaps Controlling for NAEP scores

	Imputed wages for non fulltime employed persons					
Male sexism	(1)	(2)	(3)	(4)	(5)	(6)
Average	-0.258					
	[0.103]*					
10th Percentile			0.127		0.094	0.169
			[0.134]		[0.131]	[0.154]
50th Percentile		-0.296	-0.368	-0.431	-0.476	-0.428
		[0.085]**	[0.114]**	[0.110]**	[0.127]**	[0.145]**
90th Percentile				0.143	0.135	0.135
				[0.076]	[0.078]	[0.086]
Female sexism						
10th Percentile						0.085
						[0.194]
50th Percentile						-0.114
						[0.187]
90th Percentile						-0.076
						[0.088]
M-F diff. in NAEP math	-0.068	-0.015	0.014	0.029	0.048	0.140
	[0.175]	[0.165]	[0.168]	[0.161]	[0.165]	[0.181]
M-F diff. in NAEP reading	0.493	0.457	0.439	0.388	0.378	0.307
	[0.293]	[0.272]	[0.273]	[0.266]	[0.268]	[0.284]
Observations	42	42	42	42	42	42
R-squared	0.18	0.28	0.29	0.34	0.35	0.39

Notes: Table reports coefficients (standard errors) from OLS regressions of residual state-level female-male wage gaps on various measures of sexism among males and females. Residual female-male wage gaps are estimated by median regression after imputing log wages for men and women who are non full-time employed. Log wages are imputed based on the following rule: Impute log wages equal to zero if number of years of education is less than the bottom quartile of the education distribution for each gender in each year. Impute log wages equal to 20 if education is greater than the top quartile of the education distribution for each gender in each year. Just as in the base specifications, female-male wage gaps are estimated using 1977-2002 May/ORG CPS data and control for education, a quadratic in experience, marital status, gender-specific year effects and state effects. Data from 1973-1976 are dropped because CPS reports states in groups in those years. States are dropped if they are not sampled in the GSS in the years necessary to measure male prejudice. *-5%, **-1% significance.

Table 10: Sexism and gender wage gaps with alternative selection imputation procedures

	Median Regression, imputing log wages for:					
	Nonemployed Females + Males			*Non full-time employed Females + Males*		
	Dependent variable: Residual Female-Male Wage Gap in Market					
Male sexism	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: All States						
10th Percentile	0.078	0.067	0.062	0.096	0.07	0.052
	[0.098]	[0.106]	[0.095]	[0.109]	[0.132]	[0.103]
50th Percentile	-0.271	-0.319	-0.259	-0.311	-0.418	-0.29
	[0.093]**	[0.101]**	[0.090]**	[0.104]**	[0.126]**	[0.098]**
90th Percentile	0.087	0.098	0.082	0.089	0.106	0.080
	[0.055]	[0.059]	[0.053]	[0.061]	[0.074]	[0.057]
Observations	45	45	45	45	45	45
R-squared	0.18	0.22	0.18	0.19	0.25	0.2

Notes: Tables reports coefficients (standard errors) from OLS regressions of residual state-level female-male wage gaps on various measures of sexism among all males. Residual female-male wage gaps are estimated by median regression using different imputation rules. Columns (1) to (3) imputes wages for nonemployed individuals while Columns (4) to (6) imputes wages for non full-time employed individuals. Columns (1) and (4) impute log wages equal to zero if the number of years of education is less than the education in the bottom 10th percentile of the education distribution for each gender in each year and impute log wages equal to 20 if education is greater than the top 10th percentile. Columns (2) and (5) is similar except that the bottom and top 25th percentile is used as the cut-off. Columns (3) and (6) restricts the sample to individuals who are household heads or spouses of household heads. For males and females with no recorded spouse wage, the imputation rule follows Columns (2) and (5). For females with spouse wage information, we impute log wages equal to zero if the number of years of education is less than the bottom quartile of the female education distribution in a specific year and log spouse wage is less than the bottom quartile of the spouse wage distribution in a specific year. Log wages are set equal to 20 if number of years of education is greater than the top quartile of the female education distribution and log spouse wage is greater than the top quartile of the spouse wage distribution. *-5%, **-1% significance.

Table 11: Sexism and Selection-Corrected Gender Wage Gaps dropping DC

	Impute wages for non fulltime employed persons				
	Excluding D.C.				
Male sexism	(1)	(2)	(3)	(4)	(5)
Average	-0.315				
	[0.078]**				
10th Percentile			-0.055		-0.054
			[0.114]		[0.116]
50th Percentile		-0.290	-0.260	-0.284	-0.256
		[0.067]**	[0.092]**	[0.095]**	[0.114]*
90th Percentile				-0.006	-0.005
				[0.067]	[0.068]
Observations	44	44	44	44	44
R-squared	0.28	0.31	0.31	0.31	0.31

Notes: Table reports coefficients (standard errors) from OLS regressions of residual state-level female-male wage gaps on various measures of sexism among males and females. Residual female-male wage gaps are estimated by median regression after imputing log wages for men and women who are non full-time employed. Log wages are imputed based on the following rule: Impute log wages equal to zero if number of years of education is less than the bottom quartile of the education distribution for each gender in each year. Impute log wages equal to 20 if education is greater than the top quartile of the education distribution for each gender in each year. Just as in the base specifications, female-male wage gaps are estimated using 1977-2002 May/ORG CPS data and control for education, a quadratic in experience, marital status, gender-specific year effects and state effects. Data from 1973-1976 are dropped because CPS reports states in groups in those years. States are dropped if they are not sampled in the GSS in the years necessary to measure male prejudice. *-5%, **-1% significance.

Table 12: Male Sexism and Women's Labor Market Outcomes - Using the 40th Percentile

	Selection-Corrected Female-Male Wage Gaps					Residual Female-Male Employment Gap				
Male sexism	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Average	-0.234					-0.124				
	[0.095]*					[0.037]**				
10th Percentile			0.195		0.183			0.087		0.085
			[0.147]		[0.145]			[0.058]		[0.058]
40th Percentile		-0.305	-0.443	-0.442	-0.566		-0.144	-0.205	-0.177	-0.234
		[0.092]**	[0.139]**	[0.124]**	[0.157]**		[0.036]**	[0.054]**	[0.050]**	[0.063]**
90th Percentile				0.117	0.112				0.028	0.026
				[0.072]	[0.072]				[0.029]	[0.029]
Observations	45	45	45	45	45	45	45	45	45	45
R-squared	0.12	0.2	0.24	0.25	0.28	0.21	0.27	0.31	0.28	0.32

Notes: Table reports coefficients (standard errors) from OLS regressions of residual state-level female-male wage gaps on various measures of sexism among males and females. Residual female-male wage gaps are estimated by median regression after imputing log wages for men and women who are non full-time employed. Log wages are imputed based on the following rule: Impute log wages equal to zero if number of years of education is less than the bottom quartile of the education distribution for each gender in each year. Impute log wages equal to 20 if education is greater than the top quartile of the education distribution for each gender in each year. Just as in the base specifications, female-male wage gaps are estimated using 1977-2002 May/ORG CPS data and control for education, a quadratic in experience, marital status, gender-specific year effects and state effects. Data from 1973-1976 are dropped because CPS reports states in groups in those years. States are dropped if they are not sampled in the GSS in the years necessary to measure male prejudice. *-5%, **-1% significance.

Table 13: Sexism among college-educated and employment and wage gaps

Skilled Prejudice Distribution (Both Males and Females)	Employment Gap (1)	Selection-Corrected Wage Gap (2)
Panel A: All States		
10th Percentile	0.033	0.023
	[0.032]	[0.084]
50th Percentile	-0.107	-0.132
	[0.035]**	[0.092]
90th Percentile	-0.023	-0.082
	[0.018]	[0.047]
Observations	45	45
R-squared	0.3	0.18

Note: The table reports results from specifications similar to those reported in column (5) of tables 3 and 4. Sexism is measured among both men and women, but only among those with at least a college education. There is no education restriction on the sample used to estimate wage and employment gaps. *-5%, **-1% significance.

Appendix Table 1: Gender-related question in the GSS

FEWORK	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	Do you agree or disagree with this statement? Women should take care of running their home and leave running the country up to men.
FEPRES	If your party nominated a woman for president, would you vote for her if she were qualified for the job?
FEPOL	Tell me if you agree or disagree with this statement: Most men are better suited emotionally for politics than are most women.
FECHILD	A working mother can establish just as warm and secure a relationship with her children as a mother who does not work.
FEPRESCH	A preschool child is likely to suffer if his or her mother works.
FEHELP	It is more important for a wife to help her husband's career than to have one herself.
FEFAM	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.
	Years where all questions overlap: (1977, 1985, 1986, 1988, 1989, 1990, 1991, 1993, 1994, 1996, 1998)