NBER WORKING PAPER SERIES

THE EFFECTS OF SEXISM ON AMERICAN WOMEN: THE ROLE OF NORMS VS. DISCRIMINATION

Kerwin Kofi Charles
Jonathan Guryan
Jessica Pan

Working Paper 24904
http://www.nber.org/papers/w24904

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2018

We thank seminar participants at the NBER, ASSA Meetings, University of Chicago, UC-Berkeley, Iowa State University, American University, Princeton University, Purdue University, the University of Michigan and many colleagues for helpful comments and discussions. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Effects of Sexism on American Women: The Role of Norms vs. Discrimination
Kerwin Kofi Charles, Jonathan Guryan, and Jessica Pan
NBER Working Paper No. 24904
August 2018
JEL No. J12,J13,J16,J22,J31,J7,Z10

ABSTRACT

We study how reported sexism in the population affects American women. Fixed-effects and TSLS estimates show that higher prevailing sexism where she was born (background sexism) and where she currently lives (residential sexism) both lower a woman's wages, labor force participation and ages of marriage and childbearing. We argue that background sexism affects outcomes through the influence of previously-encountered norms, and that estimated associations regarding specific percentiles and male versus female sexism suggest that residential sexism affects labor market outcomes through prejudice-based discrimination by men, and non-labor market outcomes through the influence of current norms of other women.

Kerwin Kofi Charles
Harris School of Public Policy
University of Chicago
1155 East 60th Street
Chicago, IL 60637
and NBER
kcharles@uchicago.edu

Jonathan Guryan
Northwestern University
Institute for Policy Research
2040 Sheridan Road
Evanston, IL 60208
and NBER
j-guryan@northwestern.edu

Jessica Pan
Department of Economics
National University of Singapore
1 Arts Link
Singapore
jesspan@nus.edu.sg

1 Introduction

The average American woman's socioeconomic outcomes have changed dramatically over the past fifty years. In the labor market, her wages and probability of employment have risen substantially compared to the average man's (Blau and Kahn, 2017); at every age, she is less likely to have ever been married (Isen and Stevenson, 2008); and she has fewer children over her lifetime and is older when she bears her first (Bailey et al., 2014). A large literature analyzes the evolution of these average changes, but virtually no work has studied heterogeneity in women's outcomes across markets in the U.S.[1] We document large cross-state differences in women's outcomes, even within the same geographic region of the country. These gaps have persisted for decades and we show that that they are not accounted for by standard individual controls or market-level factors.

This paper assesses how, and by what mechanisms, cross-market differences in prevailing sexism affect these persistent residual cross-market differences in women's outcomes. In principle, sexism might take many forms but we focus on negative or stereotypical beliefs concerning the ability or appropriateness of women engaging in market work rather than home production. Our analysis defines prevailing sexism in a market as the extent to which its residents believe that: (i) that women's capacities are inferior to men's; (ii) that the family unit is hurt when women focus on activities outside the home; or (iii) that men and women should occupy specific, distinct roles in society.

The prevalence of sexist beliefs in the population seems naturally relevant for two forces that the literature has long speculated partly determine women's outcomes. One of these potential forces is gender discrimination. If their sexist beliefs lead some persons to take actions that make it more difficult or less financially rewarding for women to engage in

[1]Some factors whose roles in driving average national changes have been studied include the contraceptive pill (Goldin and Katz, 2002; Bailey et al., 2014); technological changes in home and market production (Greenwood et al., 2005; Weinburg, 2000; Black and Spitz-Oener, 2010); shifts in occupational sorting (Hsieh et al., 2013); and women's educational gains (Charles and Luoh, 2003; Goldin et al., 2006; Blau and Kahn, 2006). Two of the small number of papers to have studied cross-market differences are Beaudry and Lewis (2014) who study cross-market wage differences; and Black et al. (2014), who study differences in married women's labor supply across cities.

a market activity compared to otherwise identical men, then women would experience a form of "taste-" or "prejudice-based" discrimination (Becker, 1957). Since the aspect of sexism we study concerns views about women working outside the home, we focus on gender discrimination in the labor market.[2] Prevailing sexism could also affect a woman's outcomes by partly shaping her preferences. Since at least Smith (1759), economists have understood that an individual's utility function partly depends on the social influences to which they are exposed, and this idea has been formally analyzed in more recent theoretical work (Akerlof and Kranton, 2000; Benabou and Tirole, 2006; Acemoglu and Jackson, 2017). As an element of the social norms that she partly internalizes or adheres to, the sexist beliefs of people around her might affect a woman's tastes, expectations, and beliefs, and thereby her choices and outcomes.

We assume that any discrimination that people in the market where she lives and works direct towards a woman is a function of their own sexism and not of the sexism of persons elsewhere. By contrast, because norms might alter preferences in a long-lasting or even permanent way, they could affect a person's decision-making beyond the place and time of her exposure to them. Thus, wherever she lives as an adult, a woman's outcomes could be partly affected by the norms that prevailed where she spent her formative years. Building on this idea, the paper introduces a distinction between sexism where an adult woman spent her formative years ("background sexism"), and where she currently lives ("residential sexism"). Residential sexism, we argue, affects an adult woman through two possible channels: the norms she currently confronts and perhaps internalizes, and sexism-based discrimination in the labor market. Background sexism affects an adult woman's outcomes through the lasting influence of norms she previously internalized and skills she may have obtained when younger.

We measure sexism in different states using data from the General Social Survey (GSS), which for several years has asked respondents a series of questions concerning their beliefs

[2] Of course, women may be the victims of sexism-based discrimination in other settings that we do not study, such as in the housing market.

about women's capacities, and roles, and place in society. Using a sample of internal migrants (persons born in one state but currently living in another) drawn from several waves of data from the Census and American Community Survey (ACS), the first part of our analysis examines how outcomes for these adult migrants are affected by sexism in their states of birth and state of residence.

Estimating the causal effect of background sexism is relatively straightforward, under the assumption that whether people are born in one place rather than another (and thus into one level of background sexism rather than another) is as good as random. OLS or fixed effects estimates of the causal effect of background sexism would therefore not be biased because of correlation between people's innate productive traits and average sexism where they happened to be born. Of course, background sexism might affect productive skills that women *acquire* at some point in their lives, even if it is unrelated to traits like innate intelligence that they were endowed with at birth. We estimate versions of the models we use to measure the effect of background sexism with and without controls for own schooling. The education a person obtains is the most important of these acquired traits and likely strongly correlated with other acquired skills that we do not observe.[3]

Unlike background sexism, which one can argue is plausibly exogenously assigned at birth, an adult migrant undoubtedly selects the residential sexism she faces through her choice of where to live. To account for potential endogneity bias problems arising from any tendency of women with certain latent productive traits to sort to locations with particular levels of sexism, we estimate the causal effect of residential sexism using Two Stage Least Squares (TSLS). The instruments we use are based on two exogenous costs of migration suggested by previous work: the physical distance between markets; and the settlement patterns of previous generations of migrants. The variation that under-girds these instruments has been used to study domestic and international migration, but never for the questions analyzed in

[3] If prevailing sexism in their market affects people's schooling, adding controls for schooling would be '"over-controlling" since any difference in outcomes arising from the change in schooling caused by larger background sexism is part of the *effect* of higher background sexism.

this paper.[4] The instruments strongly predict variation in residential sexism, conditional on background sexism, even for internal migrants from within the same geographic region of the country.

We find that among male and female internal migrants, who currently live in a given state, gender gaps in wages and employment are larger the higher the level of sexism where they were born. Also, compared to other female migrants living in the same state, women who were born in more sexist places marry and have their first child at appreciably younger ages. All of these effects are economically large, ranging between one-quarter to one-third of a standard deviation for a standard deviation higher level of background sexism. These results indicate that background norms that she was previously exposed continue to affect American women's adult life outcomes, even after they have left the market where they were exposed to these beliefs.

To estimate the causal effect of residential sexism, the TSLS models compare persons who migrated from states with very similar levels of sexism but who were induced to locate into different levels of residential sexism because of exogenous migration costs. We find that, holding background sexism constant, adults induced to move to more sexist places by the instrumental variables marry and bear children younger, and have larger gender gaps in employment. Interestingly, our TSLS results show no larger gender gap in wages among workers in states with higher residential sexism. Except for wages, the estimated effects of higher residential sexism are of the same sign but are much bigger than the corresponding estimates for background sexism.

We argue that the measured effect of residential sexism is a blend of the influence of internalized residential norms plus externally-imposed discrimination. The final part of the

[4]See Card (2001) and Munshi (2003) for influential work linking pre-existing settlement patterns of people from ethnic or national groups to the location decisions of subsequent waves of international immigrants. Several subsequent papers use of the same idea to instrument for immigrant inflows into particular cities (Card, 2001; Cortes, 2008; Cortes and Tesseda, 2011). In terms of domestic migration, Boustan (2010) exploits established settlement patterns among black migrants to instrument for black in-migration to different cities. The idea that distance is a fundamental cost of migration an assumption in numerous models in urban economics and of migration. Among the papers that have used relative distance to instrument for where movers locate are Boustan et al. (2010) and Ortega and Peri (2014).

paper explores the importance of the two mechanisms. We first examine how the estimated effect of residential sexism varies across outcomes and by whether mean sexism is measured among men or among women. We reason that discrimination should account for a larger part of the overall effect of residential sexism for labor market outcomes than for non-labor market outcomes, which it affects only indirectly. In addition, we expect any effect attributable to discrimination to be at least as strongly related to the sexism of men as to that of women, since men's greater prevalence in positions of decision-making authority in the labor market gives them greater latitude to engage in discriminatory actions. For this part of the analysis, we measure labor market outcomes using the Current Population Survey (CPS), and we create selection-corrected wages to approximate the offer wage distribution.

Regressing adults' outcomes simultaneously on average sexism among men and among women in the state of residence, we find that selection-corrected wage and employment gaps are strongly related to mean male sexism among men but not to average sexism of women. In striking contrast, when a woman marries or has her first child systematically vary only with average female residential sexism but not at all with men's mean sexism. We take these results to show that the effect of residential sexism on a woman's labor market outcomes is due almost entirely to discrimination from men, whereas its effect on her non-labor market outcomes derives from her internalization of (or adherence to) the sexist norms of other women where she currently lives.

To test whether residential sexism's effect on labor market outcomes is truly due to discrimination, we next relate gender gaps in these outcomes to different percentiles of the male and female sexism distributions. Consistent with what a model of taste-based labor market discrimination against women by sexist men would predict, we find that the gender gaps in selection-corrected wages and in employment probability in a market: (a) are not related to any percentile of female sexism; and (b) are negatively related to the median of male sexism, but *not* to any other percentile of the distribution of male sexist beliefs.

The results for women's non-labor market outcomes are strikingly different. These outcomes are not related to *any* percentile of male sexism. They also are not related to the median of female sexism, varying systematically instead with other percentiles of female sexism in the market. The unimportance of the median compared to other percentiles of female residential sexism contrasted with the findings for median male residential sexism and labor market outcomes suggests that the latter results arise not because outcomes always load on the "middle" of a sexism distribution but rather from the causal mechanism of discrimination. It is also striking that non-labor market outcomes do not load onto any specific part of the distribution of female residential sexism. This pattern is perfectly consistent with the operation of norms which, unlike the sharp associations predicted by prejudice-based discrimination models, are not predicted to operate chiefly through any particular percentiles.

One body of work the paper extends is the previous literature studying gender-related beliefs and norms. A large literature in sociology has stressed the importance of norms for women's outcomes (Kiecolt and Acock, 1988; Burt and Scott, 2002). There is an active literature in economics on gender norms as well. Some of this work has examined how beliefs about gender roles affects women's outcomes across countries (Fortin, 2005) and over time within the U.S. (Fortin, 2010). Other work has examined how female employment in immigrant families' origin countries persist to later generations in the U.S. (Fernandez and Fogli, 2009; Blau et al., 2011). Alesina et al. (2013) examine the cultural persistence of gender norms over the very long term. How norms affect decision-making within a family has been explored by Fernandez et al. (2004), Bertrand et al. (2015) and Bursztyn et al. (2017). Drawing the formal distinction between background and residential sexism and the mechanisms by which they operate; accounting for the endogeneity of norms one is currently exposed to; separating the effect of men's versus women's beliefs; and evaluating how sexism's influence differs across outcomes and over time are all ways in which our paper extends the previous literature.

We also extend the massive literature on discrimination - particularly work studying taste- (or prejudice-) based discrimination. Charles and Guryan (2008) show that blacks' labor market outcomes are related to different percentiles of white racial animus exactly as taste-based model of discrimination of Becker (1957) would predict. These specific predictions are very different from those we confirm regarding percentiles of sexism and women's outcomes. Our results help to bolster the argument that taste-based discrimination may be an important determinant of outcomes for disadvantaged groups in the economy.[5]

The remainder of the paper proceeds as follows. Section 2 describes the data and presents new descriptive facts that motivate the rest of the analysis. In Section 3, we describe background and residential sexism and outline the empirical framework for estimating their causal effects. Section 4 discusses the two instrumental variables and presents first-stage results. Section 5 presents estimates of causal effects. Section 6 evaluates the relative importance of discrimination versus norms for the causal effect of residential sexism. Section 7 concludes.

2 Data and Summary Statistics

2.1 Cross-State Differences in Outcomes

We study four socioeconomic outcomes over the period from 1970 to 2017: women's hourly wages and labor force participation relative to those of observationally identical men, and the ages at which women aged 20-40 married and had their first child. Information on marriage and childbearing comes from the 1980-2000 Decennial Censuses and from the 2012 three-year aggregate American Community Survey (ACS) (2010-2012), which we combine to create a "2012" sample. Because the Census/ACS does not track children across households, we infer how old a woman was when her first child was born from the reported age of her eldest child living in the same household. This part of the analysis is restricted to women aged 20 to

[5]For an alternative view, see Flabbi (2010), who estimates a search model of the labor market with matching, bargaining, employers prejudice and workers labor force participation decisions and argues that prejudice is not a relevant factor in explaining the slow-down in gender wage convergence in the U.S. in the 1990s.

40 so her the oldest child is likely to be still residing in the mother's household. Whenever possible, we follow the convention in labor economics of using the Current Population Survey (CPS) to measure wages and labor force participation.[6] However, as we describe below, much of our analysis can only be done using Census/ACS data. The two data sources yield qualitatively similar results for various descriptive exercises, which is reassuring about the quality of the Census/ACS labor market data. A final note about the data is that in order to avoid conflating issues concerning gender with the potentially different set of considerations having to do with race, the paper only studies adult whites.

The paper focuses on persistent differences across states in women's outcomes since 1970. Table 1 presents cross-state summary statistics of the four outcomes studied over the entire study period, and separately for the early and later halves of that time. The left hand side of the table shows unadjusted versions of the statistics, and the right hand side presents conditional measures that have been regression-adjusted for age and education.

Over the earlier half of our study period, gender wage and labor force participation gaps in the different states averaged 26.6% and 39.3 percentage points, respectively. Although these gaps are still disturbingly large, they have declined sharply since the mid-1990s, falling to 15.2% and 21.8 percentage points today, respectively. Women's non-labor market outcomes have also changed substantially. For example, whereas the share of women aged 20-40 in a state who have never married averaged 24% across states between the early 1970s and mid-1990s, during the latter half of the study period, this mean rose to 41.8%. Similarly, the cross-state mean age of first birth among women aged 20-40 who have ever given birth grew by 0.7 years between the late and early halves of the time frame we study.

While these mean trends may be familiar from previous work, two things in the table are likely not as well-known. First, the close similarity between the regression-adjusted and

[6] As Autor et al. (2008) note the point-in-time wage and employment information provided by the May/ORG make it a superior source of labor market information, and especially information related to the distribution of wages, compared to the March CPS and Census/ACS, which provide only retrospective annual earnings information. See the Data Appendix for more details on the construction of the hourly wage measure.

Table 1: Cross-State Summary Statistics of Women's Labor Market and Non-Labor Market Outcomes

	Panel A. Labor Market Outcomes					
	Unconditional			Residual		
	1977-2017	1977-1997	1998-2017	1977-2017	1977-1997	1998-2017
	(1)	(2)	(3)	(4)	(5)	(6)
	Female-Male LFP Gap					
Mean	-0.162	-0.241	-0.147	-0.159	-0.232	-0.148
SD	0.031	0.029	0.038	0.030	0.028	0.037
Max-Min	0.122	0.121	0.141	0.118	0.123	0.135
	Female-Male Wage Gap, Conditional on Working					
Mean	-0.254	-0.360	-0.206	-0.269	-0.361	-0.251
SD	0.036	0.039	0.038	0.033	0.036	0.033
Max-Min	0.191	0.189	0.189	0.172	0.169	0.174
	Panel B. Non-Labor Market Outcomes					
	1980-2012	1980-1990	2000-2012	1980-2012	1980-1990	2000-2012
	Female Share Nevermarried (Age 20 to 40)					
Mean	0.246	0.243	0.418	0.168	0.161	0.316
SD	0.053	0.052	0.054	0.053	0.049	0.059
Max-Min	0.227	0.219	0.234	0.217	0.194	0.243
	Female Age at First Birth (Age 20 to 40)					
Mean	23.58	23.58	24.28	23.67	23.58	24.37
SD	0.655	0.554	0.789	0.428	0.381	0.484
Max-Min	2.66	2.30	3.19	1.75	1.55	1.99
No. of states	44	44	44	44	44	44

Note: LFP and wage gaps (log hourly wages) are for whites aged 25-64 in MORG/ORG CPS. Residualized results control for gender-specific year effects and state fixed effects, years of schooling and age dummies. The non-labor market outcomes are estimated using Census/ACS data on white women aged 20 to 40 at time t. Residualized versions control for year fixed effects, number of years of schooling and age dummies. Data are from the sample of 44 states used in the main analysis (ie, states with necessary information in GSS on sexist beliefs).

raw statistics indicate that individual-level controls like age and education explain very little of average cross-state differences in womens outcomes. The other result in table that has received very little previous attention is the remarkable stability of cross-state differences in all four outcomes. Whether measured in terms of the standard deviation or the max-min gap, neither raw nor regression-adjusted average state-level outcomes have converged over time. Instead, cross-state differences either remained essentially constant or widened slightly.

Not only have the size of mean cross-state gaps changed little over time, but each state's ranking in the cross-state distribution of outcomes has also been relatively constant. States where women's outcomes ranked in 1980s at the top, middle or bottom for any outcomes relative to other states held essentially the relative positions throughout our period of study.[7] Formally, state$\times$year effects add no further explanatory power in state-level regressions with state and time fixed effects.[8] Figure 1, which plots the outcomes over time after collapsing the states into 9 Census divisions, illustrates the lack of both cardinal and ordinal convergence in women's outcomes across states.

2.2 Sexism Data From GSS

Our information on sexism comes from the General Social Survey (GSS). The GSS is a nationally representative survey that, for several years, has asked respondents various questions concerning their attitudes or beliefs about women's place in society. Our analysis uses responses to the eight questions, which were asked most consistently in the 1977-1998 waves of the GSS.

Two of these questions ("Women should take care of running their home and leave running the country up to men" and "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") elicit beliefs about women's and men's appropriate roles inside and outside the home. Two other questions touch on

[7] The correlation coefficient between states' ranking across all states between the first and last year of our study period for residualized labor force gaps, wage gaps, share ever married by age 40 and mean age of first child were, respectively, 0.72, 0.6, 0.93, and 0.84.

[8] See Appendix Table A1 for these results.

Figure 1: Mean Residual Labor Market and Non-Labor Market Outcomes by Census Divisions

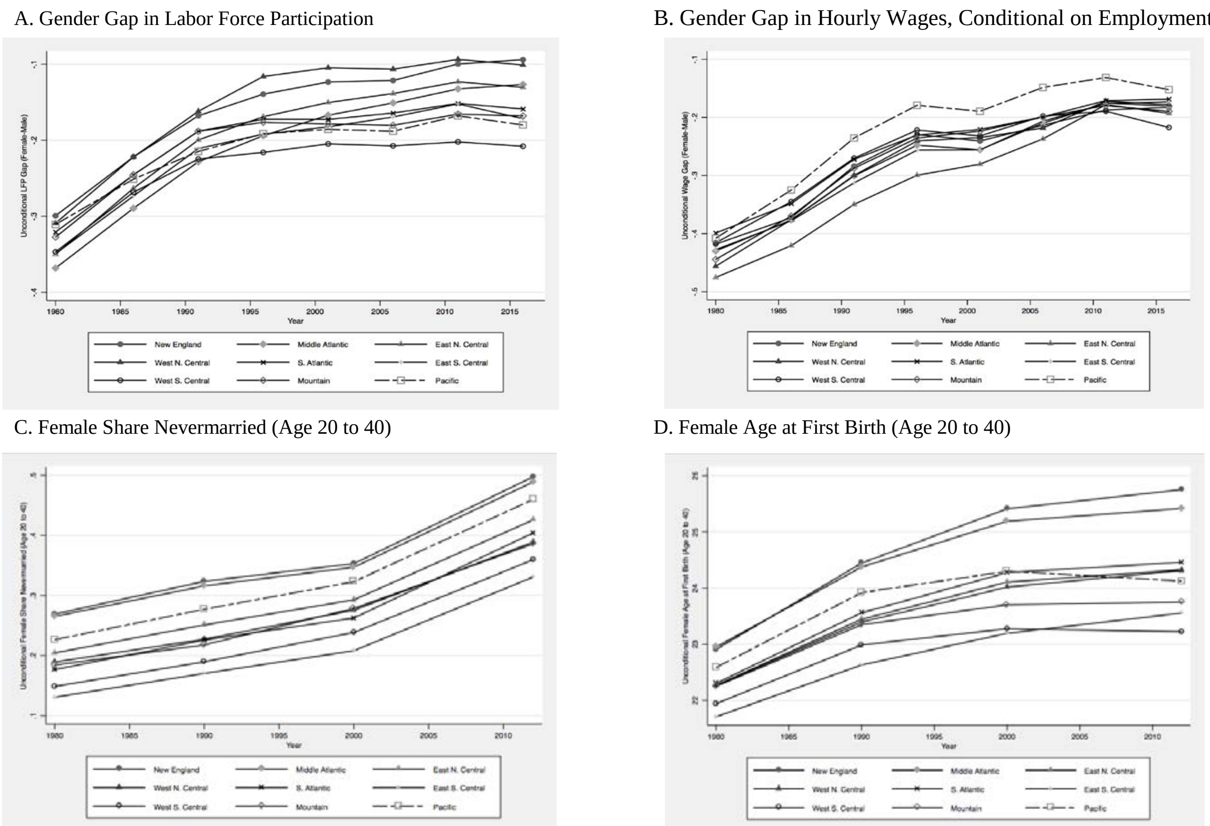

Note: LFP and wage gaps are estimated using MORG/ORG CPS data for sample of whites aged 25 to 64. The non-labor market outcomes are estimated using Census/ACS data for women between 20 and 40. Outcomes are residualized of gender-specific year effects, state fixed effects, individual years of schooling, and age dummies. Means are for 9 Census Divisions over time.

beliefs about women's capacities ("Would you vote a female for President?", "Are men better suited emotionally for politics than are most women?"). Finally, respondents' stated agreement or disagreement with statements such as "A working mother can establish just as warm and secure a relationship with her children as a mother who does not work," reveals beliefs about whether working mothers can juggle their dual roles effectively."[9]

We drop GSS respondents who are younger than 18, and recode the data so that larger values reflect higher sexism. Using their responses to all of the eight questions, we create a unidimensional index of a person's overall sexism. The sexism index subtracts off from the individual's response to each question the average response from the entire population to the same question in 1977, then divides by the standard deviation in the first year that the question is asked.[10] We next create an aggregated state-level version of the standardized index by averaging these individual-level measure at each quantile point across survey years. Each state's sexism at a given percentile is the average over all years for which we have data for the state. There happens to be substantial stability in cross-state differences in sexism over time, even as sexism has declined everywhere. As a result, we lose little useful variation by taking means over all years that the state is observed.[11]

Figure 2, which identifies states by their average overall sexism, shows the geographic distribution of sexist beliefs. Sexism is highest in the Southeast and least extreme in New England and the West. The figure shows that there is substantial variation in mean sexism across states *within* each geographic region of the country - a fact which we will exploit in the analysis to follow.

Before turning to our formal analysis, Table 2 presents simple correlations between the average level of sexism in a state and state-level means of the outcomes we study. Panel

[9]See the Appendix Table A2 for GSS variable abbreviation and summary for each of the eight questions about gender beliefs.

[10]This procedure is similar to that used by Charles and Guryan (2008) in their study of racial prejudice.

[11]We see this from regressions that relate state-level sexism to state, year and state×year fixed effects. Analysis of the R^2 indicates that state×year fixed effects account for 13-20% of the variation over times in state level outcomes, indicating that there is relatively little *differential* variation in state means over time and thus little convergence. See Appendix Table A3 for regression results.

Figure 2: Mean Overall Sexism in U.S.

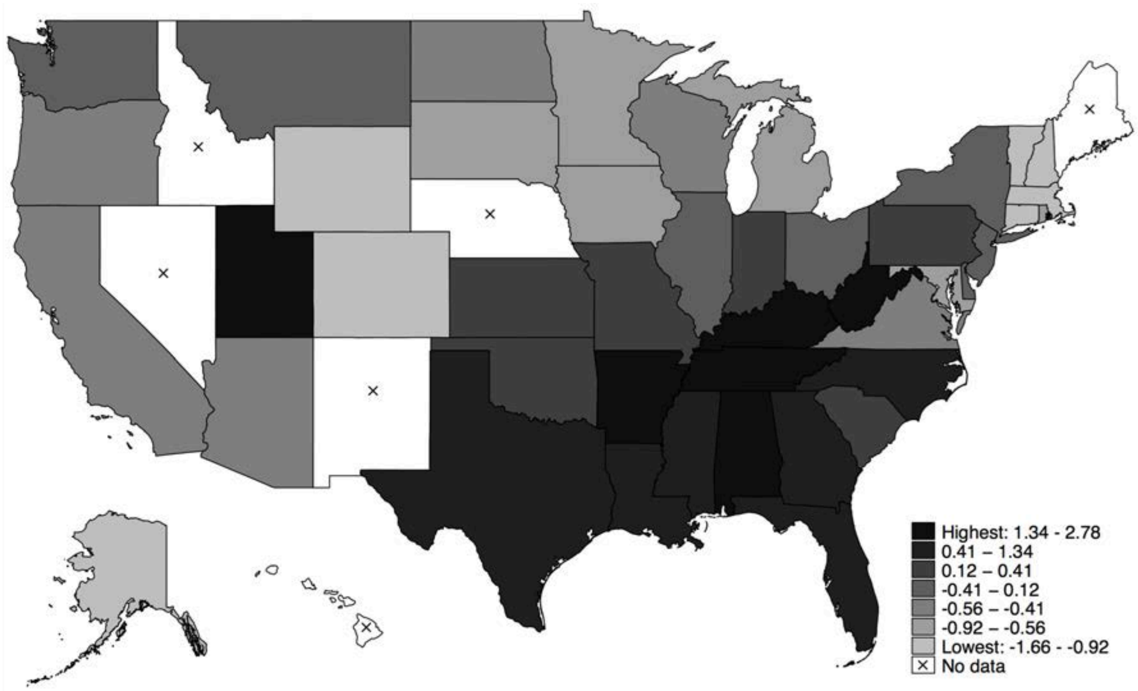

Note: Data are from several years of the General Social Survey (GSS). See text for further details.

(A) shows simple means of the various outcomes, which have been adjusted only for gender-specific year effects. Panel (B) shows results after the underlying individual-level data has been residualized of age and education and year effects. The results show that higher mean overall sexism in a state is correlated with a larger gender employment rate gap, and with earlier marriage and childbearing among women in the state. Interestingly, the simple gender wage gap among workers in a state is not correlated with mean sexism in the state. The same pattern is evident in Panel (B) of the table. For reasons made clear below, the results in Table 2 are merely suggestive of a causal link between average sexism in a state and women's outcomes. Moreover, it is not clear how these simple associations relate to the mechanisms through which we believe sexism operates. In the next section, we define these constructs more precisely and outline the framework and methods we use to estimate and interpret the causal effect of prevailing sexism on women's different socio-economic outcomes.

3 Framework and Empirical Approach

3.1 Setup

Let br indicate the set of adults who were born in a particular state $b \in [b_1, b_2 \ldots, b_{50}]$ and who currently reside in a given state $r \in [r_1, r_2, \ldots, r_{50}]$.[12] Denote women in a br group by the indicator variable f_{br}. Assume that men's outcomes in the economy are determined exclusively by two (possibly unobserved) sets of factors: θ_b, which are individual productive traits like experience or ability; and φ_r which are characteristics of the market of residences, such as the quality of public transportation or availability of childcare there. By contrast, let women's outcomes depend not only on productive individual and market characteristics, but also on prevailing sexism, which affects them through two possible mechanisms: market discrimination and norms.

We assume that sexism among the people where she lives or has lived affects an adult

[12]Our empirical work will use data from only 44 states because of data availability in the GSS.

Table 2: Mean Overall Sexism in State and Women's Outcomes

	LFP Gap (Female - Male)	Log Wage Gaps, conditional on working (Female-Male)	Share of Females Nevermarried	Average Female Age at First Child
	(1)	(2)	(3)	(4)
	A. Unconditional Outcomes			
Sexist Beliefs in State of Residence	-0.018***	0.000	-0.046***	-0.520***
	[0.004]	[0.007]	[0.006]	[0.080]
R-squared	44	44	44	44
Observations	0.292	0.000	0.469	0.406
	B. Residualized Outcomes			
Sexist Beliefs in State of Residence	-0.017***	-0.005	-0.040***	-0.279***
	[0.004]	[0.006]	[0.005]	[0.046]
R-squared	44	44	44	44
Observations	0.292	0.017	0.458	0.351
Sample	CPS 1977-2017, Age 25 to 64		Census/ACS 1980-2012, Age 20 to 40	

Note: The table reports coefficients (standard errors) from OLS regressions of various outcomes on mean sexism in state. The labor market outcomes are estimated from white men and women aged 25-64 in 1977-2013 May/ORG CPS. Non-labor market outcomes in columns (3) and (4) use data on white women aged 20-40 from 1980 to 2000 Census and 2012 ACS. Unconditional regressions include state fixed effects and gender-specific year effects. Regressions for residualized results include controls for years of schooling and age dummies. Robust standard errors weighted by the inverse of the variance of the outcome reported in parentheses. ***significant at 1%, **5%, *10%.

woman through two types of mechanisms: norms and discrimination. We use the term "norms" to refer to the set of social conventions, mores and influences with which a person comes into contact. One type of norms, which we call *background norms*, N_b, are the social influences and mores in the place where she grew up and which affect the traits and preferences that a woman brings into her adulthood. Exposure to these norms during her formative years potentially determine a woman's preferences through internalization or unconscious emulation. In addition, through pre-market discrimination, norms where she was born and raised might have led people and institutions there to provide her with different training or education during her youth relative to what was provided to boys, or caused her to make different human capital investment relative to boys.

Apart from the skills and elements of her preferences that were determined earlier in her life partly from exposure to background norms, two other forces determine her adult outcomes. One of these is the influence of what we call *residential norms.* These are social mores and practices in the market where she currently lives. The second force operating on an adult woman of a given level of skill is *market discrimination*, D_r. Following the standard formulation in economics, we define market discrimination as the amount by which others' actions reduce the economic payoff an adult woman gets from market activity relative to that of an otherwise identical man doing the same thing.[13]

We assume that these three forces - the things that determine skills and preferences a woman brings into adulthood, the influence of social mores she encounters as an adult, and any market discrimination she faces - are functions of prevailing sexist beliefs in the relevant markets. Specifically, we assume that

$$
\begin{aligned}
N_b &= \delta^b \, \overline{S}_b + u_1 \\
N_r &= \delta^r \, \overline{S}_r + u_2 \\
D_r &= \beta \, \overline{S}_r + u_3
\end{aligned}
\tag{1}
$$

[13]The paper studies discrimination against adult women in the labor market and not in other settings, such as the housing market.

where $\overline{S}_b$ is mean overall sexism where the woman was born and $\overline{S}_r$ is average overall sexism where lives. Residential sexism is an element of the norms that she's currently exposed to and is the basis upon which some persons discriminate against her when she engages in labor market activity. These effects are captured by the terms β and δ^r. Although she no longer lives there, background sexism potentially affects her current outcomes because the norms where she grew up may have affected enduring aspects of her latent preferences and skills. This effect is captured by the parameter δ^b. The mean-zero terms u_1, u_2, and u_3 represent all other determinants of discrimination and the two types of norms.

Mean outcomes for br persons, Y_{br}, are given by the sum of the "neutral" components θ_b and φ_r, which affect both men and women, plus a portion attributable to discrimination and norms, which affects only women. That is,

$$Y_{br} = f_{br}(D_r + N_r + N_b) + \theta_b + \varphi_r. \tag{2}$$

To avoid clutter, we suppress observable controls in (2) and throughout the paper. Substituting from (1) yields the main regression equation in the paper

$$Y_{br} = f_{br}[(\beta + \delta^r)\overline{S}_r + \delta^b\overline{S}_b] + \theta_b + \varphi_r + v_{br} \tag{3}$$

where v_{br} is a random, mean-zero statistical error. We seek to estimate the causal parameters β, δ^r and δ^b.

3.2 OLS: Specification and Bias Concerns

Inherited, pre-determined ability such as innate "intelligence" or "motivation" are typically assumed to be randomly distributed across people at birth. OLS or fixed-effects estimates of the causal effect of background sexism, δ^b, should therefore not be biased because of correlation between these latent traits people are born with and the level of sexism that prevailed where they happen to have been born.[14]

Unlike pre-determined abilities, skills that people invest in or are provided during their formative years, most notably the amount and quality of their schooling, might be affected by the amount of sexism where they were born and raised. To account for the possibility that the causal effect of background sexism on adult outcomes operates in part by changing agents' education, our preferred OLS models for estimating δ^b do not control for education. The resulting estimated coefficient on the background sexism captures the part of its effect on outcomes that operates through schooling. By contrast, our preferred specifications for estimating the effect of residential sexism always control for schooling, which is in general determined *before* the person's exposure to residential sexism. This approach follows standard procedure in the discrimination literature, which focuses on differences in outcomes between adults, holding constant their level of observable skill (education). In the Appendix, we present results showing that for both background and residential sexism, our main results are generally unaffected by whether we control for education or not.

Because of possible systematic sorting, OLS performed on (3) is unlikely to yield unbiased estimates of the two causal effects of residential sexism, β and δ^r. Women might choose where to live based on locations' levels of sexism directly or on the basis of other market characteristics that are correlated with market sexism. Residential sexism could therefore be endogenous in a naive OLS regression like (3), so that the observed relationship between higher sexism in a market and the outcomes of women living there reflects both the causal effect of sexism plus the fact that women with particular latent traits are more likely to live in such places in the first place. We use Two Stage Least Squares (TSLS) methods to deal with the possible endogeneity of residential sexism. Using TSLS also addresses any measurement error bias arising from the fact that our analysis treats a person's labor market as their state of residence which almost certainly mis-measures their true labor market - the place where they live, work and socially interact.

[14]This assumption is implicit in previous papers studying the effect on immigrants' outcomes of cultural practice in countries from which they moved to the U.S.

It is worth emphasizing that even after endogeneity bias has been accounted for, the estimated effect on residential market sexism from the TSLS analysis is the sum of the causal effect of discrimination plus the causal effect of internalized residential norms - $(\beta+\delta^r)$. In the final part of our analysis we attempt to shed light on the importance of the two mechanisms of discrimination and internalized residential norms for residential sexism. We focus first on the more fundamental problem of obtaining an unbiased estimate of the causal effect of $\overline{S}_r$ in equation (3). Below, we outline our approach for isolating exogenous variation in residential sexism for use in the TSLS analysis.

4 Exogenous Variation in Residential Sexism

4.1 Determinants of Migration Decision

The potential endogeneity problem in (3) for which TSLS is intended to correct is that the level of residential sexism observed for migrants born into a given level of background sexism, is a function of the latent productive factors θ_b and φ_r. We assume that, in general, the residential sexism chosen by migrants born in state b, who therefore had background sexism $\overline{S}_b$, may be written

$$\overline{S}_r(b) = \gamma\overline{S}_b + \rho_1 Z_1(\Delta_{bj}) + \rho_2 Z_2(\kappa_{bj}) + \theta_b + \varphi_r + \xi_{br}, \tag{4}$$

where ξ_{br} is a random error term. The two terms Z_1 and Z_2, are functions that depend on fixed costs of migration that must be incurred by anyone (whatever their individual traits) who moves between locations. They are the source of exogenous variation in the TSLS analysis.

The variable Δ_{bj} is a vector measuring the distance between the state b from which a migrant moves to every other state j (measured as the distance between the states' population-weighted centroids). Relative distance is a fixed migration cost, meant to reflect the fact that it is simply logistically more difficult for people to move farther from their initial locations, irrespective of their individual traits or features of the two markets. (See, for example, Boustan et al. (2010) and Ortega and Peri (2014)).

Figure 3 assesses the plausibility of this assumption for internal migrants in the U.S. whom we study. Focusing on all adults across multiple Census years living in states different from where they were born, the figure is a histogram of the distances between the migrants' states of residence and states of birth. Since other factors besides relative distance determine migrants destinations the plotted relationship does not decline perfectly smoothly. Nonetheless, the figure clearly shows that internal migrants in the U.S. are systematically less likely to move to states that are farther away, irrespective of the state from which they are moving. If the probability that a migrant leaving b ended up in j, in fact, depended systematically *only* on how relatively far j is from b compared to other other destinations, plus some random error, the expected level of residential sexism observed among migrants from b would be given by

$$Z_1(\Delta_{bj}) = \sum_{b \neq j} \overline{S}_j / \Delta_{bj}. \tag{5}$$

An equivalent way to think of $Z_1(\Delta bj)$, then, is as the portion of the expected residential sexism for migrants moving from states with sexism $\overline{S}_b$ that is attributable to the exogenous relative spatial location of various states.

The existing literature suggests that where previous generations of migrants from given initial location have historically settled represents a second fixed cost of migration for the subsequent movers. Various papers, mostly studying international migration, posit that the "enclaves" formed by previous migrants affect the ease with which later movers can adjust to different possible destinations (see, for example (Card, 2001)). These historical migration patterns thus represent a plausibly exogenous driver of the location choices of current migrants, since the factors that drove previous waves of migrants to settle where they did are arguably otherwise unrelated to later migration decisions.

The internal migrants studied in this paper moved from their states of birth after 1970.

Figure 3: Distance of States Residence from States of Birth Among U.S. Migrants

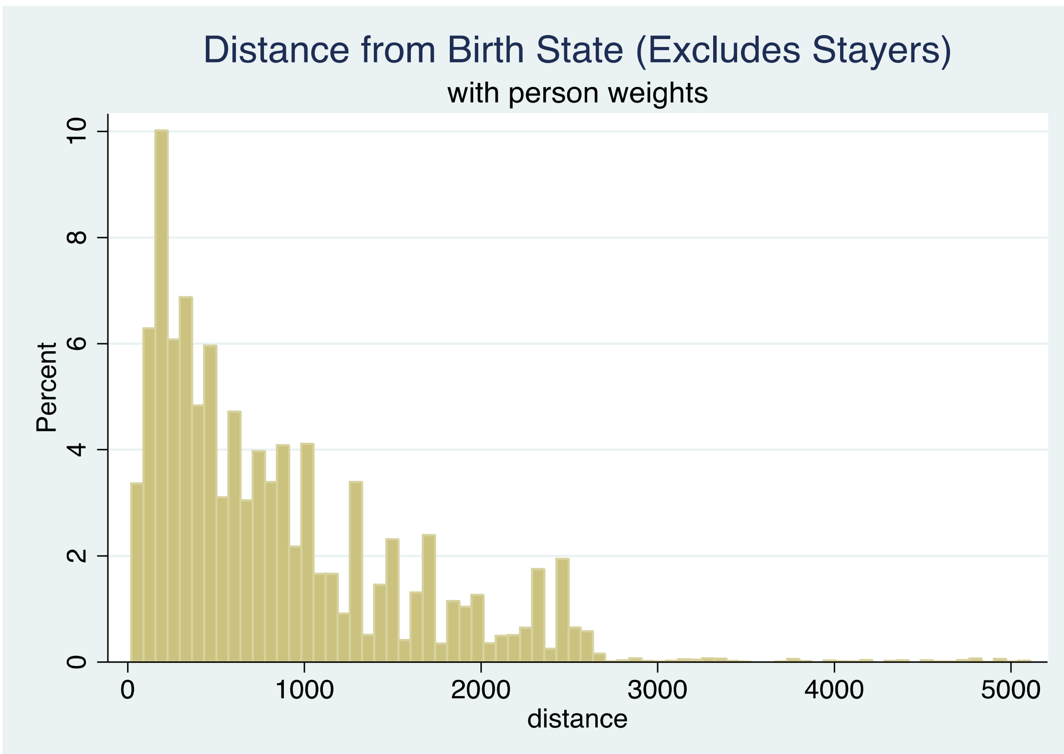

Note: Data are from several years of the Census and ACS. See text for details.

We proxy for pre-existing enclaves these migrants confronted using the share of people born in state b who lived in a different state j in 1970, or κ_{bj}. Using the same logic used for relative location of a birth state, the portion of current location decisions of migrants' attributable to exogenous pre-existing migration patterns is given by

$$Z_2(\kappa_{bj}) = \sum_{b \neq j} \overline{S}_j \times \kappa_{bj}. \tag{6}$$

Given the foregoing, equation (4) is a first stage regression for the endogenous variable $\overline{S}_r$ in equation (6), where $Z_1(\Delta_{bj})$ and $Z_2(f_{bj})$ are instrumental variables whose "strength" is captured by the parameters ρ_1 and ρ_2. Two Stage Least Squares (TSLS) regression performed on this system allows us to estimate the causal effect of residential sexism - $\beta + \delta^r$.

4.2 First-Stage Results

The two panels of Figure 4 illustrate the variation the two instrumental variables isolate. The x-axis of Figure 4a is mean sexism in different states of birth. Although there is substantial variation in background sexism, many states have very similar levels of $\overline{S}_b$. For example, migrants from Massachusetts and Maryland were born into very similar background sexism. The y-axis in the figure plots the value of the instrumental variable based on distance, $Z_1(\Delta_{bj})$.

The figure shows that the component of migrants' destination decision attributable exclusively to the exogenous accident of where their state of birth is located geographically relative to all other states predicts that migrants from Maryland and Massachusetts would move to markets with very different levels of sexism. Looking across the figure, one sees similar variation for different pairs of states with similar background sexism - Pennsylvania/Louisiana, Tennessee/West Virginia and New York/Ohio to name some examples. Figure 4b repeats this exercise for the instrumental variable $Z_2(\kappa_{bj})$, which is based on the amount of historical migration flows from the origin state (pre-existing "enclaves"). Here also one sees substantial variation in predicted residential market sexism, conditional on the level of background sexism from which the migrant comes.

Whether $Z_1(\Delta_{bj})$ and $Z_2(\kappa_{bj})$ are valid instruments for the TSLS analysis depends on how much of the variation in $\overline{S}_r(b)$ they explain, holding constant $\overline{S}_b$. The first stage results in Table 3 answers this question. We conduct this analysis on the main sample of internal migrants in the Census/ACS, born in states of birth b and living in different states r. We collapse the more than 6 million individual observations on adult male and female migrants down to (state of birth $\times$ state of residence $\times$ gender $\times$ year) cells, yielding slightly more than 15000 observations once empty br cells with no reliable sexism data are dropped.[15]

The regressions in Table 3 relate sexism in the market to which a given group of migrants move to the two instrumental variables. All the models in the table control for the migrant group's background sexism, and include fixed effects for the Census region of state b, interactions between region effects and gender fixed effects, and interactions between year fixed effects and gender dummies. The results in the first two columns show that each of the IVs is separately strongly related to residential sexism, and explains a sufficiently large portion of the variation in the potentially endogenous regressor to remove any "weak instrument" concern. This is especially true of the "enclave" or "historical settlement pattern" measure, for which the F-stat of 33 far exceeds the conventional threshold of 10. When the two variables are entered simultaneously, only the enclave IV remains strongly significant. This is not surprising; where migrants settled in the past was likely based on temporary factors from that time period *plus* the historically important propensity to not move too far from home. Again the F-stat on the pair of IVs is comfortably above 10.

One might be concerned that the IVs could be correlated with the labor market traits or family and marriage preferences of migrants. For example, perhaps being (exogenously) surrounded by a particular level of average sexism in nearby states causes only certain types

[15]It should be noted that we use different samples for the labor market and non-labor market results. The latter uses only women, so has about half the number of observations after collapsing the individual-level data. We conduct a separate first-stage analyses using the sample on which the subsequent TSLS analysis is performed. The first stage results are very similar across samples.

Figure 4: Predicted Residential Sexism based on: (a) Distance Between State of Birth and Other States and, (b) Historical Destination of Previous Migrants from State of Birth

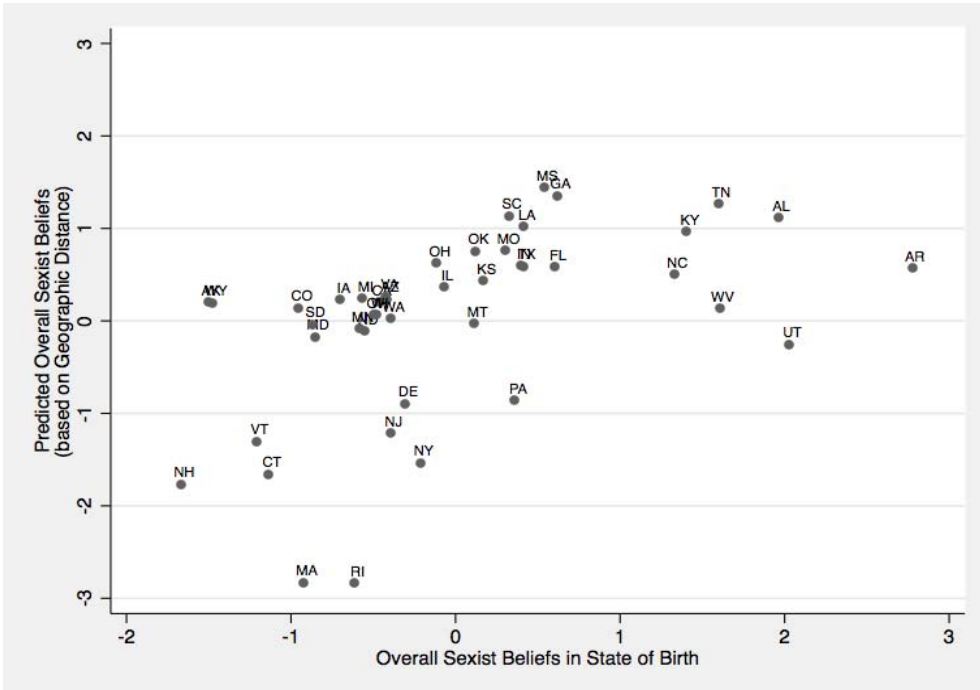

(a) Predicted Residential Sexism Given Relative Location

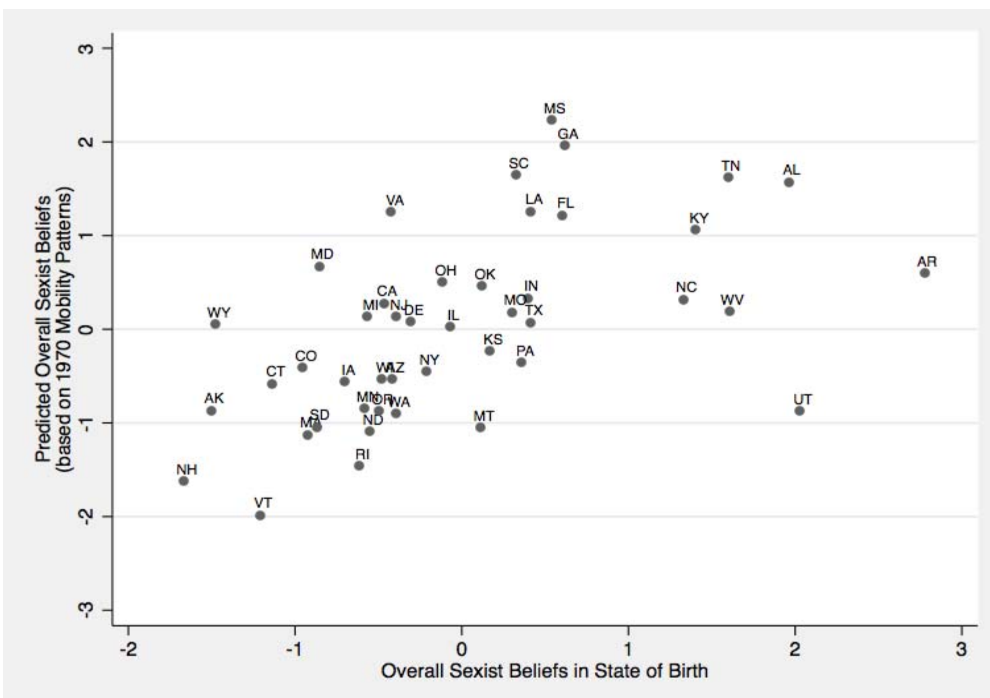

(b) Predicted Residential Sexism Given Historical Migration

Table 3: First Stage Results: Effect of Relative Distance and Enclave Instrumental Variables on Residential Sexism

	LFP sample			
	Outcome: Overall Sexism in the State of Residence			
	(1)	(2)	(3)	(4)
Distance IV	0.258***		0.066	0.067
	[0.068]		[0.064]	[0.065]
Enclave IV		0.264***	0.232***	0.232***
		[0.046]	[0.041]	[0.040]
F-stat of excluded instrument	14.6	33.2	19.0	19.5
Controls:				
Outmigration Rate				X
Overall Sexism in State of Birth*Female	X	X	X	X
Overall Sexism in State of Birth	X	X	X	X
Region of Birth FE	X	X	X	X
Female*Year FE	X	X	X	X
Observations	15,103	15,103	15,103	15,103
R-squared	0.074	0.083	0.083	0.083

Notes: The data are from the 1980, 1990, 2000 US Census and 2010-2012 ACS (3-year aggregate). The sample is restricted to whites age 25 to 64 for the labor market outcomes and age 20 to 40 for the non-labor market outcomes (females only) who are not currently living in their state of birth. Sexist beliefs are normalized to have a mean of 0 and standard deviation 1 across the 44 states with available information to construct the index of sexist beliefs in the GSS (excluding DC). The Distance and Mobility IV are standardized to have mean 0 and standard deviation 1 in the sample of states.

of people to move in the first place compared to people in places where the nature of sexism in nearby states is different. Our analysis compares migrants who grew up in the same level of sexism. Suppose that persons from states with similar background sexism have the same distribution of preferences regarding the sexism they would like to live under as adults, all else equal. Then, to the degree that the propensity to migrate in the first place is driven by sexism in neighboring states, the fraction of people migrating from two states with similar sexism but with different levels of the IVs should differ systematically. The regressions in the fourth column of the table adds the out-migration rate from each origin state b as a control to account for at least some of the latent traits of migrants that could be correlated with the IVs. The results show that this control does not change the main conclusion of the first stage analysis.[16]

5 Benchmark Causal Effects

This section presents the main benchmark estimates of the effects of the two types of sexism on adults' various outcomes. We first show results for background sexism, then present and discuss TSLS estimate of the effect of exposure to higher residential sexism.

5.1 Fixed Effects Estimates for Mean Background Sexism

Our estimates of the causal effects of background sexism come from fixed-effects regressions which are variants of regression (3) and of the form:

$$y_{br}^{k} = \delta^{b}(f_{br} \times \overline{S}_{b}) + \Gamma_{r} + \theta_{b} + u_{br} \tag{7}$$

where Γ_r is a vector of state of residence fixed effects. Under our assumption that the community into which someone is born is randomly assigned, $\overline{S}_b$ is uncorrelated with the

[16] In the TSLS below, the results are essentially unaffected by which of the four first-stage specification shown in the four columns of Table 3 the analysis uses. Results for these various exercises available upon request.

error term in the regression. Obviously, any feature of the state of residence, including its level of sexism and fixed features φ_r are subsumed in the state of residence fixed effects in regression (7). Below, we estimate the causal effect of $\overline{S}_r$ using the TSLS approach described previously.

In (7), δ^b measures the causal effect on migrants' adult outcomes where they live now of the level of overall sexism into which they were born. Table 4 presents the regression results. Recall, because background potentially partly determines the skills and preferences that a person brings into adulthood it could affect both men's and women's adult outcomes. The coefficient for the level sexism variable in the table is the estimated effect for men, while the coefficient on sexism term that is interacted with the female dummy variable measures the differential effect for women. The first two columns show results for labor market outcomes. As seen in Column (1), labor force participation rates are lower for men born in states with higher levels of overall sexism. Higher background sexism has a significantly larger impact on women's labor force participation relative to the effect for men. The results imply that, for groups of migrants currently living in the same state, gender gaps in labor market outcomes are bigger the higher the amount of sexism in their states of birth.

For a sense of the magnitude of these estimated effects, a standard deviation increase in sexism in one's state of birth (approximately the difference between having born in California rather than Mississippi, or in Minnesota rather than Texas) reduces men's labor force participation rate by 0.9 percentage points. For women, the same increase in the amount of sexism in their state of birth reduces labor force participation rates as adults by an additional 0.7 percentage points. The decline in relative female labor force participation associated with a one-standard deviation increase in sexism in one's state of birth is about 3.5 percent of the mean unconditional female-male labor force participation gap across migrants' states of birth and close to 40 percent of the standard deviation of gaps across states. A standard deviation increase in average sexism is roughly the same size as the interquartile range across states. The results for the wage gap among workers shown in the second column are qualitatively similar, albeit not statistically significant.

For both labor force participation and wages, we estimate versions of (7) that control for own education, age, and the mean NAEP reading and math test scores in the state of birth.[17] Controlling for these measures does not affect estimates effect of background sexism on women's labor force participation relative to men's, but causes the estimated effect for men to go to zero. For men's wages too, we find no effect of background sexism after accounting for education. These results indicate that the effect of higher background sexism for men in Table 4 operate entirely through the schooling people are provided because they were born and raised in a more sexist place. However, higher background sexism worsens women's labor market outcomes beyond anything having to do with this general education effect.

The third and fourth columns of Table 4 show results for migrant women's non-labor market outcomes. The samples on which these regressions are run contain only women aged 40 or less. Again, the estimated effects of background sexism are large and strongly statistically significant. For two adult women living in the same state but born elsewhere, the one from a state where sexism one standard deviation higher is 2.1 percentage points more likely to have been married and is bear her first child approximately 0.2 years sooner. These very strong effects are not significantly affected by controlling for own or state level education (see last four columns of Appendix Table A4). This suggests that the sexism where women were born (and likely raised) affects their outcomes through channels other than their schooling. As we have argued, a good candidate explanation is that being exposed to sexist norms earlier in life affects her preferences, with consequences for her decision-making throughout her life.

[17] See Appendix Table A4.

Table 4: Effect of Background Sexism on Labor Market and Non-Labor Market Outcomes

	Labor Force Participation	Log Wages, conditional on working	Never Married	Age at First Birth
	(1)	(2)	(3)	(4)
	All Migrants		Female Migrants Only	
Sexist Beliefs in State of Birth:				
Overall*Female dummy	-0.007***	-0.003		
	[0.002]	[0.003]		
Overall	-0.009***	-0.008	-0.021***	-0.196***
	[0.003]	[0.006]	[0.003]	[0.045]
Controls:				
State of Residence*Female FE	X	X		
Female*Year FE	X	X		
State of Residence FE	X	X	X	X
Year FE	X	X	X	X
Region of Birth FE	X	X	X	X
Observations	15,103	15,063	7,522	7,453
R-squared	0.928	0.927	0.801	0.769

Note: The table reports coefficients (standard errors) from OLS regressions of residual state-level female-male employment and log hourly wage gaps on sexist beliefs in the state of birth. Residual female-male employment and wage gaps are estimated using the sample of whites age 25 to 64 from the 1977-2017 May/ORG CPS data and control for years of schooling, age dummies, gender-specific year effects and state fixed effects. The non-labor market outcomes are from a sample of white females age 20 to 40 from the 1980 to 2000 Census and 2010-2012 ACS (3-year aggregate) and control for the number of years of schooling, age dummies and year fixed effects. Sexist beliefs are normalized to have a mean of 0 and standard deviation 1 across the 44 states with available information to construct the index of sexist beliefs in the GSS. Robust standard errors weighted by the inverse of the variance of the outcome reported in parentheses. ***significant at 1%, **5%, *1%.

5.2 TSLS Estimates for Mean Residential Sexism

We next present TSLS estimates of the effect of residential sexism on migrants' outcomes. In principle, these results account for endogeneity and measurement error problems in naive OLS models, which are also shown in the table for easy comparison. The TSLS models use both the relative-distance and historical settlement pattern instrumental variables discussed above.[18] Analogous to Table 4, the coefficient on the un-interacted residential sexism variable is the estimated effect of higher residential sexism on men, while the coefficient on the sexism variable that is interacted with gender is the differential effect for women. In this TSLS framework, these causal effects measure how their outcomes are affected for migrants induced to live in a particular level of residential sexism because of the instruments. Recall that the regressions control for mean sexism in the state of birth so the estimated residential sexism are from comparisons of persons born into similar background sexism.

The labor market results in Table 5 indicate that living in a more sexist state has no effect on the likelihood that a man participates in the labor force. This result holds for both the OLS estimates in column (1) and for the TSLS results in columns (2) and (3). By contrast, we find that higher residential sexism has a strong negative effect on a woman's labor force participation. Together, these estimates imply that higher residential sexism decreases the relative labor force participation rate of women compared to men who were born into the same level of background sexism. The results in columns (4)-(6) of Table 5 show that higher residential sexism significantly lowers wages of men, with OLS and TSLS estimates that are nearly identical. Interestingly, unlike the participation results, we find no larger adverse effect on women's wages compared to men's from higher sexism where they live. Thus, our results imply that higher residential sexism does not increase gender wage gaps conditional on employment.

Table 6 presents TSLS estimates of the effect of residential sexism on women's non-labor

[18] TSLS results from models that use only one of the instruments are qualitatively similar and are available on request.

Table 5: TSLS Estimates of Effect of Residential Sexism on Labor Market Outcomes

	Labor Force Participation			Log Wages, conditional on working		
	OLS	TSLS		OLS	TSLS	
	(1)	(2)	(3)	(4)	(5)	(6)
Sexist Beliefs in State of Residence:						
Overall*Female Dummy	-0.011**	-0.033***	-0.033***	-0.004	0.023	0.023
	[0.004]	[0.010]	[0.010]	[0.008]	[0.022]	[0.022]
Overall	-0.005	-0.003	-0.002	-0.067***	-0.068**	-0.062**
	[0.003]	[0.005]	[0.006]	[0.014]	[0.028]	[0.027]
Controls:						
Outmigration rate in state of birth			X			X
Overall Sexism in State of Birth*Female	X	X	X	X	X	X
Overall Sexism in State of Birth	X	X	X	X	X	X
Region of Birth FE*Female	X	X	X	X	X	X
Female*Year FE	X	X	X	X	X	X
Outcomes residualized of:						
Education, Age dummies	X	X	X	X	X	X
Observations	15,103	15,103	15,103	15,103	15,103	15,103
R-squared	0.917	0.906	0.907	0.823	0.819	0.818

Note: The data from the 1980, 1990, 2000 US Census and 2000-2012 ACS (3-year aggregate). The sample is restricted to white individuals age 25 to 64 who are not currently living in their state of birth. Sexist beliefs are normalized to have a mean of 0 and standard deviation 1 across the 44 states with available information to construct the index of sexist beliefs in the GSS (excluding DC). In Columns (2) and (3), sexist beliefs in the state of residence (Overall*Female and Overall) are instrumented using both the Distance and Mobility IVs, respectively. ***significant at 1%, **5%, *1%.

market outcomes. Recall that the regression in this table are estimated on the sample of women aged 40 or less. The results show that, compared to other women who were born into similar levels of sexism, female migrants living in more sexist states are more likely to have ever married and bear their first child at appreciably younger ages. These patterns are found both in the OLS regressions and in the preferred TSLS models. The OLS and TSLS estimates are both strongly statistically significant and do not appreciably differ in magnitude. Compared to the fixed effects results in Table 4 the estimates in Table 6 indicate that, whereas the two types of sexism act in the same direction on a woman's marriage and age of first childbearing, the effect of residential sexism on these outcomes is substantially larger - nearly two and a half times as big as the effect of the sexism into which she was born as a child.

Overall, our benchmark causal estimates indicate that both background and residential

Table 6: TSLS Estimates of Effect of Residential Sexism on Non-Labor Market Outcomes

	Never Married (Age 20 to 40)			Female Age at First Child (Age 20 to 40)		
	OLS	TSLS		OLS	TSLS	
	(1)	(2)	(3)	(4)	(5)	(6)
Sexist Beliefs in State of Residence:						
Overall	-0.029***	-0.039***	-0.037***	-0.197***	-0.197**	-0.242**
	[0.006]	[0.011]	[0.012]	[0.035]	[0.090]	[0.090]
Controls:						
Outmigration Rate in State of Birth			X			X
Overall Sexism in State of Birth	X	X	X	X	X	X
Region of Birth FE	X	X	X	X	X	X
Year FE	X	X	X	X	X	X
Outcomes residualized of:						
Education, Age dummies	X	X	X	X	X	X
Observations	7,522	7,522	7,522	7,453	7,453	7,453
R-squared	0.698	0.690	0.693	0.481	0.481	0.480

Note: The data from the 1980, 1990, 2000 US Census and 2000-2012 ACS (3-year aggregate). The sample is restricted to white females age 20 to 40 who are not currently living in their state of birth. Sexist beliefs are normalized to have a mean of 0 and standard deviation 1 across the 44 states with available information to construct the index of sexist beliefs in the GSS (excluding DC). In Columns (2) and (3), sexist beliefs in the state of residence (Overall*Female and Overall) are instrumented using both the Distance and Mobility IVs, respectively. ***significant at 1%, **5%, *1%.

sexism strongly affect adult women's outcomes. Being born into a higher level of sexism reduces a woman's relative labor market outcomes and leads her to marry and have children sooner, wherever she ends up living as an adult. The substantial long-term effects of sexism into which a woman happens to have been born are mostly dwarfed by the even bigger effects of higher residential sexism, which act in the same direction on her outcomes. Wages are the exception to this pattern. While larger background sexism reduces wages, we find that higher residential sexism has no statistically significant effect on the gender log wage difference. Residential sexism seems to chiefly affect a woman's labor market outcomes on the extensive, work/non-work margin with scant effect on relative wages.

A noteworthy feature of the TSLS results in Table 5 is that the participation point estimates are about three times as large as the corresponding OLS results. Strikingly, there is no similar difference between OLS and TSLS for the non-labor market estimates in Table 6. What type of endogenous sorting would cause OLS estimates of the negative effect of residential sexism on the likelihood of participation, in particular, to be biased towards zero?[19] This type of bias implies that, for some reason, female migrants who are observed living in more sexist states have relatively high levels of latent labor force attachments. Something that might generate this pattern is if women who move to states with higher sexism are systematically more likely to have been made to do so *by their jobs*. Because these women would be disproportionately likely to be working after moving, their presence in the set of all female migrants living in a state imparts a downward bias to estimates of the negative effect residential sexism on labor force participation. By isolating variation in location that is attributable only to exogenous relative distance or historical settlement patterns, the TSLS models hopefully remove this bias.

6 Mechanisms for Causal Effect of Residential Sexism

We have argued that the causal effects of background sexism reflect the long-term effect of norms one was exposed to earlier in life. By contrast, the effect of higher residential sexism reflects not only the effect of her internalization of norms that currently surround her where she lives now, but also discrimination imposed upon her in the market. How much, if any, of the effect of higher residential sexism is due to discrimination rather than norms? In this section, we conduct a variety of analyses to shed light on this question.

6.1 Average Male and Female Sexism for Different Outcomes

We begin this portion of our analysis with an assessment of the differential effect on the two types of outcomes of average sexism in a market among men and women. Two ideas motivate this exercise. First, the portion of the overall effect of residential sexism attributable to discrimination should be larger for outcomes like wages and labor force participation which labor market discrimination can directly affect, compared to outcomes like the ages

[19]OLS regressions might also suffer from attenuation bias caused by the fact that the use of the coarse geographic area of the state as the labor market mis-measures the market in which women interact and work. One might have expected any such mis-measurement to also cause substantial attenuation bias in the non-labor market results, yet the OLS and TSLS for these regression are very similar.

of marriage or fertility which are affected only indirectly by market discrimination, if at all. Second, should there be discriminatory action in the labor market, one would expect it to more closely tied to sexism among men than among women. This is because men's historically greater prevalence in positions of decision-making authority in the market likely provides them with greater ability to engage in discriminatory action compared to sexist women.

Using GSS reports of respondent gender, we create separate measures of sexism for men and women in each state. Figure 5 plots the mean of these male and female sexism indices for the different states. The figure shows that although the two measures are very strongly positively correlated, they are far from perfectly so. Men are, on average, more sexist than women in most states, but the extent to which this is true varies substantially from state to state and there are several states for which mean sexism among women is higher than that for men. From the figure, it is clear that there is significant variation in one of the mean sexism measures conditional on information about the other - an important requirement for the analysis we conduct below.

We estimate a series of regressions relating different outcomes in a state simultaneously to average male and average female sexism in the state. This part of our analysis does not require the migrant sample used above, with information about state of birth and state of residence. Therefore, instead of the Census/ACS, we use data from the CPS. As discussed above, the CPS is the preferred source of data on labor market information in the U.S. Data on two non-labor market outcomes in the table still comes from the Census/ACS. All the regressions in the table control for education and age dummies and for observed state-level characteristics.

The results in column (1) of Table 7 show that the gender labor force participation gap in a state is significantly larger the more sexist the average man in the state, but exhibits no statistically significant relationship with the state's average female sexism. Mirroring the wage results found previously for overall sexism with the Census wage data, the results in

Figure 5: Sexism Among Men and Among Women in Each State

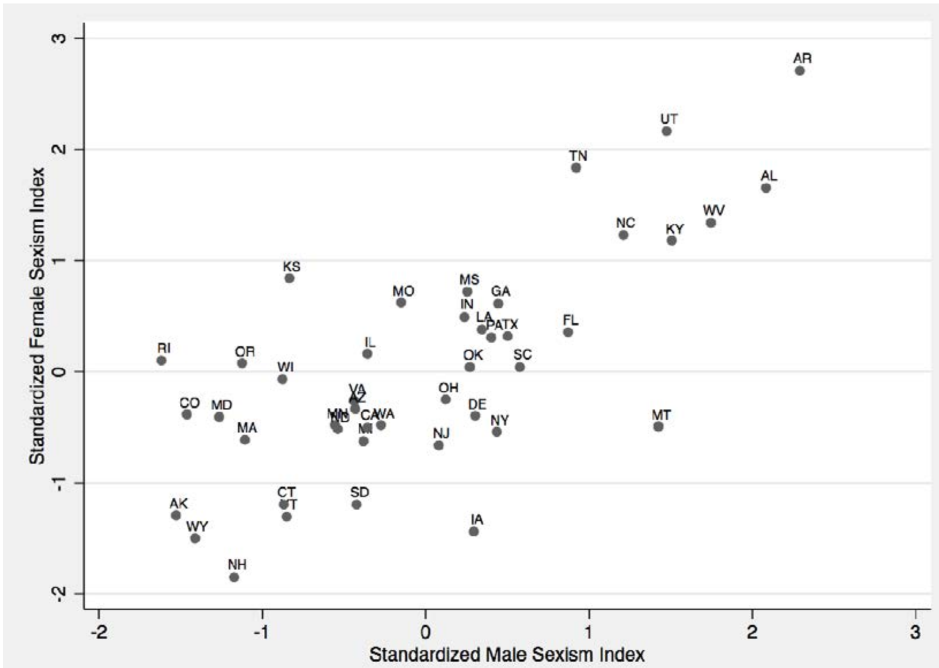

Note: Data are from the General Social Survey. The standardized female and male sexism indices have mean 0 and standard deviation 1 in the cross-state sample (44 states).

column (2) of the table show that the wage gap in a state measure using CPS data does not vary at all with mean sexism among either men or women in the state.

In column (3) we study the gender difference not in the wages received by workers but rather in the wages that men and women could *command* in the labor market. Our examination of the "offer" (or "potential") wage gap follows a long tradition that attempts to account for missing wage data among those who do not work - a particularly important problem among women. The offer wage is what a person would be paid if they *did* work, but people only work if their offer wage exceeds their reservation utility from non-work. Presumably, discriminatory treatment against women lowers their offer wages compared to those of men, which also lowers their relative likelihood of working.

We measure the gender difference in offer wages in a state as the median difference in selection-corrected wages between men and women. Our selection correction is based on a simple imputation procedure whereby non-fulltime employed women and men in the top

Table 7: Mean Male vs. Mean Female Sexism and Women's Outcomes

	LFP Gap (Female - Male)	Log Wage Gaps, conditional on working (Female-Male)	Selection-Corrected Log Wage Gaps (Female-Male)	Share of Females Nevermarried	Average Female Age at First Child
	(1)	(2)	(3)	(4)	(5)
Sexist Beliefs in State of Residence					
Male average	-0.011*	-0.004	-0.030***	-0.010	-0.093
	[0.006]	[0.006]	[0.010]	[0.011]	[0.091]
Female average	-0.007	-0.001	-0.004	-0.032**	-0.204*
	[0.006]	[0.006]	[0.010]	[0.012]	[0.107]
Sample	CPS 1977-2017, Age 25 to 64			Census/ACS 1980-2012, Age 20 to 40	
Individual level controls in first-stage regression:					
Education, Age dummies	X	X	X	X	X
R-squared	44	44	44	44	44
Observations	0.298	0.017	0.254	0.456	0.345

Notes: The table reports coefficients (standard errors) from OLS regressions of residual state-level female-male employment and wage gaps on various measures of sexist beliefs. The residual female-male employment and wage gaps are estimated using the sample of white age 25 to 64 from the 1977-2017 May/ORG CPS data and control for the number of years of schooling, age dummies, gender-specific year effects and state fixed effects. The non-labor market outcomes are estimated using the sample of white females age 20 to 40 from the 1980 to 2000 Census and 2010-2012 ACS (3-year aggregate) and control for the number of years of schooling, age dummies and year fixed effects. Sexist beliefs are normalized to have a mean of 0 and standard deviation 1 across the 44 states with available information to construct the index of sexist beliefs in the GSS. Robust standard errors weighted by the inverse of the variance of the outcome reported in parentheses. ***significant at 1%, **5%, *1%.

quartile of the gender-specific education distribution are assigned wages above the median and women and men in the bottom quartile of the gender-specific education distribution are assigned wages below the median. Our procedure follows the approach of Brown (1984), Neal (2004) and Olivetti and Petrongolo (2008).[20] Using these selection-corrected wages, we estimate median regressions, controlling for education, age fixed effects, marital status, gender specific year effects, state effects and state×female effects. The coefficients on the state×gender dummies measure the median offer wage gaps in a state.

The results in column (3) of Table 7 show that when related simultaneously to mean sexism among men and women living there, the selection-corrected wage gap in a state loads exclusively on the mean male sexism and exhibits no statistically significant relationship with average residential female sexism. This difference is not a matter of statistical imprecision; the estimated effect for male sexism is more than five times as large as that for residential female sexism.

Columns (4) and (5) show results for the two non-labor market outcomes. In a striking reversal of the participation and wage results, we find that both the likelihood of ever having married and the age of first child bearing are associated exclusively with mean sexism among women in the market, but not at all with mean sexism among men. Again, these effects are not the result of one parameter being estimated with particularly large standard errors. Rather, the two point estimates are of very different sizes, with the estimated effect for mean female sexism being roughly three times as big as the insignificant associations for male sexism.

[20] The advantage of this imputation approach is that it only requires us to make assumptions on the position of the imputed wage observations relative to the median (i.e. whether the unobserved offer wage is above or below the median). As noted in Olivetti and Petrongolo (2008), this procedure does not require assumptions on the actual level of missing wages, nor does it require us to identify an instrument that satisfies the exclusion restriction as required in the valid implementation of the Heckman two-stage sample selection correction model. Note also that this selection correction procedure does not impute wages for women with missing wages in the middle two quartiles on the education distribution under the assumption that one cannot, with reasonable confidence, predict whether the wages these women could command in the market would be above or below the median. Our results are robust to alternative selection-correction imputations such as imputing wages for women and men in the top and bottom ten percent of the education distribution. These results available upon request.

6.2 Percentile Tests

The results in Table 7 are consistent with the interpretation that the effect of residential sexism on labor market outcomes is chiefly due to discrimination from sexist men, while residential sexism's effect on non-labor outcomes reflects the influence of internalized sexist norms from other women. We use perhaps the main implication of Becker (1957) to further probe the validity of this conclusion.

Becker's model of prejudice- or taste-based discrimination implies that if there were gender discrimination in the labor market based on men's sexist beliefs, market forces would sort women away from interactions with the most sexist men. Women's equilibrium outcomes would be determined by the sexism of the most sexist men with whom they were forced to interact after this sorting - the "marginal discriminator". Since women constitute roughly half of the population, this cost-minimizing sorting should result in the marginal discriminator having sexism somewhere close to the middle of the male sexism distribution and *not* far away from the median, such as around the 10th or 90th percentiles. Further, if gender discrimination is due to men's sexism, labor market gaps should not be systematically related to *any* moment of the female residential sexism distribution.

We estimate a series of regressions relating labor market outcomes in a state to the 10th, median and 90th percentiles of the distributions of male and female sexism in the state. Table 8 shows the results. Column (1) shows that when regressed on all three percentiles of male sexism simultaneously, the labor force participation gap is strongly statistically related to the median, but not either the 10th or 90th percentiles. By contrast, column (2) shows that when regressed only on percentiles of female residential sexism, a state's participation gap does not vary in a statistically significant way with *any* part of the female sexism distribution in the state, although the point estimate for the median is bigger than those for the other two percentiles. The regression in column (3) includes the three percentiles for both female and male sexism. In this full regression, the participation gap is related only with median male sexism - the point estimate for which is twice as large at that for the statistically insignificant

Table 8: Relationship Between Alternative Percentiles of Sexist Beliefs and Women's Labor Market Outcomes

	Labor Force Participation Gap (Female - Male)			Selection-Corrected Log Wage Gaps (Female-Male)		
	(1)	(2)	(3)	(4)	(5)	(6)
Male Sexist Beliefs in State of Residence:						
10th Percentile	-0.000		0.001	-0.002		0.000
	[0.008]		[0.007]	[0.009]		[0.011]
Median	-0.017**		-0.014*	-0.030**		-0.028**
	[0.007]		[0.007]	[0.013]		[0.014]
90th Percentile	0.001		0.001	-0.004		-0.005
	[0.006]		[0.007]	[0.011]		[0.011]
Female Sexist Beliefs in State of Residence:						
10th Percentile		-0.005	-0.003		-0.010	-0.007
		[0.008]	[0.011]		[0.013]	[0.014]
Median		-0.011	-0.007		-0.007	0.005
		[0.009]	[0.013]		[0.018]	[0.017]
90th Percentile		-0.000	0.004		-0.013	-0.002
		[0.005]	[0.006]		[0.010]	[0.011]
R-squared	44	44	44	44	44	44
Observations	0.275	0.243	0.330	0.267	0.143	0.270

Note: The data are from the 1977 to 2017 MORG/May CPS. Sample is restricted to whites age 25 to 64. Sexist beliefs are normalized to have a mean of 0 and standard deviation 1 across the 44 states with available information to construct the index of sexist beliefs in the GSS (excluding DC). Robust standard errors weighted by the inverse of the variance of the outcome reported in parentheses. ***significant at 1%, **5%, *1%.

Table 9: Relationship Between Alternative Percentiles of Sexist Beliefs and Women's Non-Labor Market Outcomes

	Share of Females Nevermarried (Age 20 to 40)			Average Female Age at First Birth (Age 20 to 40)		
	(1)	(2)	(3)	(4)	(5)	(6)
Male Sexist Beliefs in State of Residence:						
10th Percentile	-0.016		0.002	-0.088		0.037
	[0.010]		[0.012]	[0.084]		[0.088]
Median	-0.009		0.008	-0.128		-0.007
	[0.014]		[0.013]	[0.097]		[0.101]
90th Percentile	-0.016		-0.024	-0.073		-0.133
	[0.014]		[0.015]	[0.098]		[0.107]
Female Sexist Beliefs in State of Residence:						
10th Percentile		-0.015	-0.027**		-0.144	-0.199*
		[0.012]	[0.012]		[0.107]	[0.115]
Median		-0.011	0.005		-0.011	0.084
		[0.013]	[0.016]		[0.120]	[0.125]
90th Percentile		-0.021**	-0.021**		-0.182**	-0.181**
		[0.009]	[0.010]		[0.072]	[0.082]
R-squared	44	44	44	44	44	44
Observations	0.356	0.449	0.516	0.284	0.358	0.407

Note: The data are from the 1980, 1990, 2000 US Census and 2012 ACS (3-year aggregate). Sample is restricted to white women age 20 to 40. Sexist beliefs are normalized to have a mean of 0 and standard deviation 1 across the 44 states with available information to construct the index of sexist beliefs in the GSS (excluding DC). Robust standard errors weighted by the inverse of the variance of the outcome reported in parentheses. ***significant at 1%, **5%, *1%.

estimate for median female.

The next three columns in the table show how the offer wage gap in a state is related to the different percentiles of male and female sexism. These results mirror the participation findings and are, in fact, slightly stronger. The selection-corrected wage gap in a state does not systematically vary with any percentile of female sexism. However, the wage gap is related in a strong and a highly statistically significant way only with the median of the male residential sexism. The point estimate for the male sexism is more than ten times larger than the estimated association for any other percentile point in the table.

While taste-based discrimination models make sharp predictions about the relative importance of different quantiles of the relevant sexism for labor market outcomes, there is no such prediction from accounts describing the operation of norms. If they are strongly affected by norms, non-labor market outcomes could thus vary systematically with any part or parts of female sexism distribution and be consistent with norms influencing outcomes.

Table 9 shows the results of the same percentile exercise for the two non-labor market outcomes that we do for wage and employment gaps in Table 8. The difference between the two sets of outcomes is striking. All of the specifications in Table 9 find that neither the likelihood that women in the state married by age 40 nor the average age at which they have their first child is related in a statistically significant fashion with any percentile of male sexism. Both outcomes are, however, strongly related to female sexism - although in neither case is the association with the median. Instead, as best illustrated in the specifications in columns (3) and (6), which include all the percentiles of both male and female sexism, the two outcomes are significantly related to the top (90th percentile) and bottom (10th percentile) of the distribution of female sexism but not with the median.

That women's non-labor market outcomes are not related to median female residential sexism but are related to the tails of that distribution is perfectly consistent with norms accounting for sexism's effect on these outcomes. As noted, existing accounts about social norms make no sharp theoretical prediction about which percentiles of underlying beliefs are particularly important for the operation of these effects. The results in Table 9 show that state-level outcomes do not always, for some mechanical reason, load most strongly on to median sexism in a horse-race with other percentiles. This raises confidence that the finding in Table 8 concerning the male median and labor market outcomes truly captures the effect of marginal discriminator as predicted by theory.

6.3 Robustness

Before concluding this section, we present some results intended to assess the sensitivity of the results in Tables 8 and Table 9 to two features of our estimation.

The first of these features is the set of percentiles we use to measure different points in the distributions of sexism. The tests we conduct are based on the predictions of a Becker model of discrimination that, since women represent roughly half the labor market, the sorting mechanism that determines equilibrium outcomes in the presence of taste-based discrimination should cause the marginal discriminator to be drawn from center of the sexism distribution rather than the tails. There being no rule about precisely which percentiles represent the "middle" versus the "tails" of the distributions, our main results settle upon the 10th, 90th and the median as natural candidates. Plus, these are the percentile points examined by Charles and Guryan (2008) in their work showing that black-white wage gaps are related, as a taste-based model would predict given blacks' small share of the overall population, *only* with the left-tail of the distribution of racial animus.[21] Yet, there are two reasons to suppose that the marginal discriminator should have sexism slightly less than the 50th percentile. For one thing, whereas women now constitute half of the labor force, over the entire period we study, the labor force was roughly 40% female. Further, Becker's model argues that because of the cost savings that non-discriminating firms enjoy, they

[21] 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 10th percentile and not to the median and 90th 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.

Figure 6: R-Squared Statistics from Bivariate Regressions of Each Outcome on Separate Quantiles of Male and Female Residential Sexism

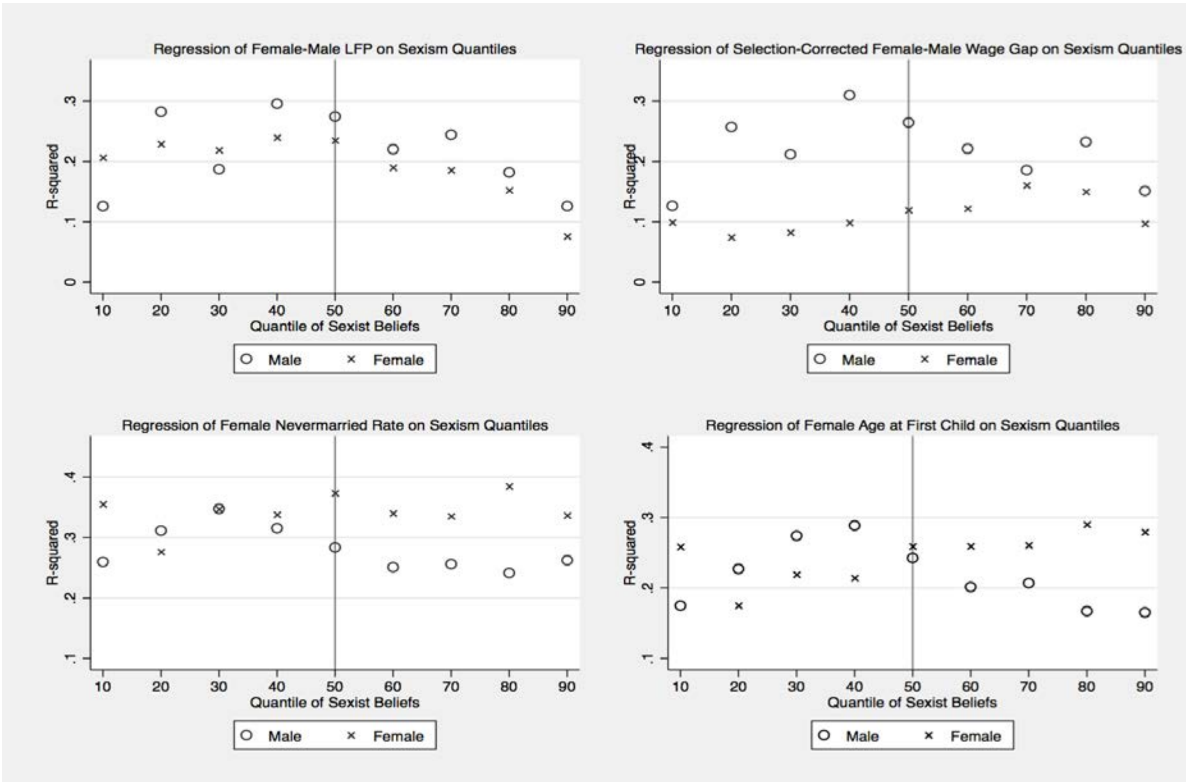

Note: Each circle or cross reports the R-squared from a bivariate regression of the outcome indicated in the title of each graph on the 10th, 20th, ..., 90th percentile of the male (circle) or female (cross) residential sexism distribution within each state.

should grow slightly larger than their discriminating counterparts, with the result that they should account for a disproportionate share of hiring of persons from the disadvantaged group (women in our case).

To examine whether our results hinge in an important way on the percentile points we have chosen to examine, we conduct a robustness exercise in which we estimate a series of simple bivariate regressions relating each outcomes to different percentiles of male and female sexism, ranging from the 10th to the 90th percentile. Figure 6 plots the R-squares from each of these eighteen pairwise associations. The figure shows the point in male and female residential sexism distributions where the association with the particular outcome is strongest. Two things stand out clearly in the figure. One is that for the labor market

Figure 7: Sensitity of Regression Estimates to Omitting Observations from Different Individual States: Coefficient on Male and Female Sexism from 43 Regressions

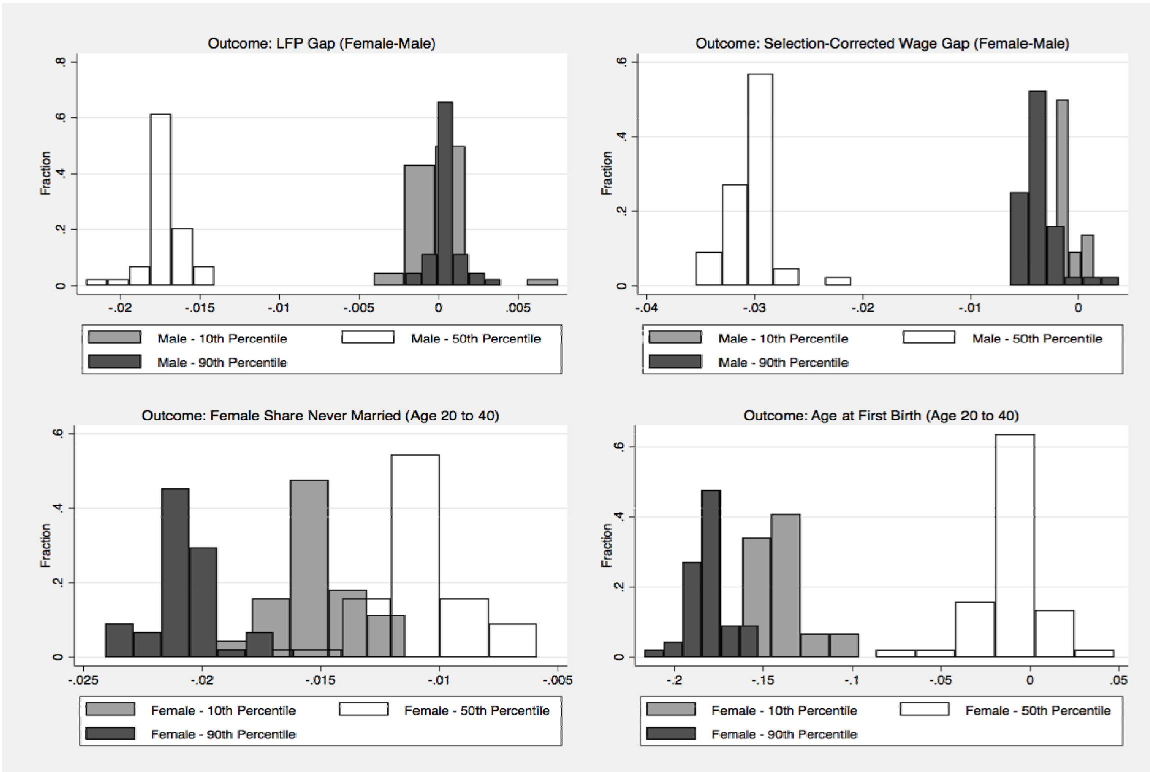

Note: The data for the labor market outcomes are from the 1977 to 2017 MORG/May CPS while the data for the non-labor market outcomes are from the 1980, 1990, 2000 US Census and 2012 ACS (3-year aggregate). Sample is restricted to whites age 25 to 64 for labor market outcomes and white women age 20 to 40 for non-labor market outcomes. Data on Sexism are from the General Social Survey. The standardized female and male sexism indices have mean 0 and standard deviation 1 in the cross-state sample (44 states).

outcomes the pairwise associations for males yield the highest R-squared, whereas for the non-labor market outcomes associations R-squared statistics are highest with percentiles of female sexism. The other noteworthy pattern in the figure is which percentiles show the biggest R-squared. For labor market outcomes, the association is highest somewhere between the 40th and median of male sexism. In sharp contrast, for neither non-labor market outcomes is the association maximized for median female sexism. It is worth emphasizing that for these results no special restrictions or specifications have been imposed on the data. We therefore find it reassuring that these associational findings are so strongly consistent with our main results.

The other robustness exercise we conduct is motivated by the fact that state-level variation plays an important role in our analysis. We are therefore naturally concerned that a small set of states - or even one particular state - might drive our main results. To assess how important an issue this is, we estimate the main regressions in Tables 8 and 9 many times, each time dropping observations from a particular state from the estimation sample. We then plot histograms of points estimates forthcoming from this series of "leave-out" regressions. Since there are 44 states in our sample, we run 43 regressions for each outcome. Having already shown the men's sexism drives labor market outcomes, the baseline regressions for labor market results only includes the three percentiles of male sexism. For non-labor market outcomes, the baseline regression is only estimated on the three different percentiles of female sexism.

Figure 7 presents the histograms of point estimates across the various regressions. The three colors in the figure represent point estimates for the three different percentiles from the series of regressions. The figure shows that for the labor market outcomes, the point estimates for median male sexism from the series of regressions are clustered around a negative value. The point estimates for the other two percentile of male sexism all cluster around 0 across all of the regressions. The results for the non-labor market outcomes are very different. For these outcomes, the histograms show that point estimates for median female sexism are clustered around 0, while those for the other two percentiles both cluster around negative values. These results provide strong graphical evidence that the main results we show above are not driven by any particular state, but instead represent something fundamental about the economy overall.

7 Conclusion

This paper studies how prevailing sexist beliefs about women's abilities and appropriate roles affect American women's socioeconomic outcomes. Studying adults who live in one state but who were born in another, and using fixed effects and instrumental variables methods, we show that sexism in a woman's state of birth and in her current state of residence both lower her wages and likelihood of labor force participation, and lead her to marry and bear her first child sooner. We argue that sexism where she was born, which we call background sexism, affects a woman's outcomes even after she is an adult living in another market through the influence of norms that she internalized during her formative years.

Our findings concerning background sexism extend the growing body of evidence concerning how important that exposure to others beliefs about the role that they should play in society or within the family can affect women over their entire lives by potentially altering their own preferences. Whereas previous work has demonstrated the importance of norms on some outcomes for persons who have immigrated to the U.S. from different cultural backgrounds, or examined observed differences across countries, we show that similar forces affect women born and raised in the U.S.

We present various pieces of evidence which, collectively, suggest the sexism where a woman lives (residential sexism) affects her non-labor market outcomes through the influence of prevailing sexist beliefs of other women where she lives. By contrast, residential sexism's effects on her labor market outcomes seem to operate chiefly through the mechanism of market discrimination by sexist men. These results extend the massive existing literature on discrimination. "Taste-based" models - along with models of statistical discrimination one of the two main classes of theoretical accounts of market discrimination - posit that discriminatory actions in the market are undertaken by persons holding negative or aversive sentiments towards the disadvantaged group in question. Our finding that quantiles of sexism in a market are associated with worse labor market outcomes for women in the very specific way that a model of gender discrimination arising from underlying sexist sentiments would predict, provides sound evidence that prejudice-based discrimination, undergirded by prevailing sexist beliefs may be an important driver of women's outcomes in the U.S.

References

Acemoglu, Daron and Matthew O. Jackson, "Social Norms and the Enforcement of Laws," *Journal of the European Economic Association*, 2017, *15* (2), 245–295.

Akerlof, George A. and Rachel E. Kranton, "Economics and Identity," *The Quarterly Journal of Economics*, August 2000, *115* (3), 715–753.

Alesina, Alberto, Paola Giuliano, and Nathan Nunn, "On the Origins of Gender Roles: Women and the Plough," *The Quarterly Journal of Economics*, May 2013, *128* (2), 469–527.

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney, "Trends in U.S. Wage Inequality: Revising the Revisionists," *Review of Economics and Statistics*, 2008, *90*, 300–323.

Bailey, Martha, Melanie Guldi, and Brad J. Hershbein, "Is There a Case for a 'Second Demographic Transition?' Three Distinctive Features of the post-1960 U.S. Fertility Decline," in Carola Frydman Leah P. Boustan and Robert A. Margo, eds., *Human Capital in History: The American Record*, Chicago, IL: The University of Chicago Press, 2014, pp. 273–312.

Beaudry, Paul and Ethan Lewis, "Do Male-Female Wage Differentials Reflect Differences in Return to Skill? Cross-City Evidence from 1980-2000," *American Economic Journal: Applied Economics*, 2014, *6* (2), 178–194.

Becker, Gary, *The Economics of Discrimination*, The University of Chicago Press, 1957.

Benabou, Roland and Jean Tirole, "Incentives and Prosocial Behavior," *American Economic Review*, August 2006, *96* (5), 1652–1678.

Bertrand, Marianne, Emir Kamenica, and Jessica Pan, "Gender Identity and Relative Income Within Households," *Quarterly Journal of Economics*, 2015, *130* (2), 571–614.

Black, Dan A., Natalia Kolesnikova, and Lowell J. Taylor, "Why Do So Few Women Work in New York and So Many in Minneapolis? Labor Supply of Married women Across U.S. Cities," *Journal of Labor Economics*, 2014, *79.*

Black, Sandy E. and Alexandra Spitz-Oener, "Technological Change and the Skill Content of Women's Work," *Review of Economics and Statistics*, 2010, *92* (1), 187–194.

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

_ **and** _ , "The Gender Wage Gap: Extent, Trends, and Explanations," *Journal of Economic Literature*, 2017, *55* (3), 789–865.

— , **Lawrence M. Kahn, and Kerry L. Papps**, "Gender, Source Country Characteristics, and Labor Market Assimilation Among Immigrants," *Review of Economics and Statistics*, 2011, *93* (1), 43–58.

Boustan, Leah P., "Was Postwar Suburbanization White Flight? Evidence from the Black Migration," *Quarterly Journal of Economics*, 2010, *125* (1), 417–443.

— , **Price V. Fishback, and Shawn Kantor**, "The Effect of Internal Migration on Local Labor Markets: American Cities During the Great Depression," *Journal of Labor Economics*, 2010, *28* (4), 719–746.

Brown, Charles, "Black-White Earnings Ratios since the Civil Rights Act of 1964: The Importance of Labor Market Dropouts," *Quarterly Journal of Economics*, February 1984, *99* (1), 33–44.

Bursztyn, Leonardo, Thomas Fujiwara, and Amanda Pallais, "Acting Wife: Marriage Market Incentives and Labor Market Investments," *American Economic Review*, Nov 2017, *107 (11)*, 3288–3319.

Burt, Keith B. and Jacqueline Scott, "Parent and Adolescent Gender Role Attitudes in 1990s in Great Britain," *Sex Roles*, 2002, *46*, 239–45.

Card, David, "Immigration Inflows, Native Outflows, and the Local Labor Market Impacts of Higher immigration," *Journal of Labor Economics*, January 2001, *19* (1), 22–64.

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

— **and Ming-Ching Luoh**, "Gender Differences in Completed Schooling," *Review of Economics and Statistics*, 2003, *85* (3), 559–577.

Cortes, Patricia, "The Effect of Low-Skilled Immigration on United States Prices: Evidence from CPI Data," *Journal of Political Economy*, 2008, *116* (3), 381–422.

— **and Jose Tesseda**, "Low-Skilled Immigration and the Labor Supply of Highly Skilled Women," *American Economic Journal: Applied Economics*, 2011, *3* (3), 88–123.

Fernandez, Raquel, Alessandra Fogli, and Claudi Olivetti, "Mother and Sons: Preference Development and Female Labour Force Dynamics," *Quarterly Journal of Economics*, 2004, *119*, 1249–1299.

— **and** — , "Culture: An Empirical Investigation of Beliefs, Work, and Fertility," *American Economic Journal: Macroeconomics*, 2009, *1* (1), 146–177.

Flabbi, Luca, "Prejudice and Gender Differentials in the U.S. Labor Market in the Last Twenty Years," *Journal of Econometrics*, 2010, *156*, 190–200.

Fortin, Nicole, "Gender Role Attitudes and Women?s Labour Market Outcomes Across OECD Countries," *Oxford Review of Economic Policy*, 2005, *21*, 416–438.

— , "Gender Role Attitudes and Women's Labor Market Participation: Opting-Out, AIDS, and the Persistent Appeal of Housewifery," Technical Report 2010.

Goldin, Claudia and Lawrence F. Katz, "The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions," *Journal of Political Economy*, 2002, *110*, 730–770.

— , — , and Ilyana Kuziemko, "The Homecoming of American College Women: The Reversal of the College Gender Gap," *The Journal of Economic Perspectives*, 2006, *20* (4), 133.

Greenwood, Jeremy, Ananth Sheshadri, and Mehmet Yorukoglu, "Engines of Liberation," *The Review of Economic Studies*, 2005, *72*, 109–133.

Hsieh, Chang-Tai., Erik Hurst, Charles I. Jones, and Peter J. Klenow, "The Allocation of Talent and U.S. Economic Growth," Working Paper 18693, National Bureau of Economic Research January 2013.

Isen, Adam and Betsey Stevenson, "Women's Education and Family Behavior: Trends in Marriage, Divorce and Fertility," in "Topics in Demography and the Economy National Bureau of Economic Research, Inc.," National Bureau of Economics Research, 2008.

Kiecolt, K. Jill and Alan C. Acock, "The Long-Term Effects of Family Structure on Gender-Role Attitudes," *Journal of Marriage and the Family*, 1988, *50*, 709–17.

Munshi, Kaivan, "Networks in the Modern Economy: Mexican Migrants in the U.S. Labor Market," *Quarterly Journal of Economics*, 2003, *118* (2), 549–599.

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

Olivetti, Claudia and Barbara Petrongolo, "Unequal Pay or Unequal Employment? A Cross-Country Analysis of Gender Gaps," *Journal of Labor Economics*, oct 2008, *Vol. 26, No. 4*, 621–654.

Ortega, Francesc and Giovanni Peri, "Openness and Income: The Role of Trade and Migration," *Journal of International Economics*, 2014, *92* (2), 231–251.

Smith, Adam, *The Theory Of Moral Sentiments*, A. Millar, 1759.

Weinburg, Bruce, "Computer Use and the Demand for Female Workers," *Industrial and Labor Relations Review*, 2000, *53* (1), 290–308.

Data Appendix

A. Construction of Hourly Wages using the CPS Data

We use data from the May CPS for 1977 to 1978 and the CPS Merged Outgoing Rotation Groups (MORG) for 1979 to 2017. The main analysis sample includes whites between the ages of 25 to 64. Following the procedure used by Autor, Katz, and Kearney (2008), hourly wages are defined as the log of reported hourly earnings for those paid by the hour and the log of usual weekly earnings divided by hours worked last week for non-hourly workers. The wage sample excludes those who are self-employed. Top-coded earnings are multiplied by 1.5. Hourly earnings that fall below $1.675/hour in 1982 dollars are dropped, as are hourly wages exceeding 1/35th the top-coded value of weekly earnings. All the earnings measures are deflated to 2009$ using the price deflator for personal consumption expenditures (PCE) from the BEA. We exclude allocated earnings observation in all years, except where allocation flags are not available (Jan 1994 to Aug 1995). Between 1989 and 1993, nonflagged allocated observations are identified and dropped by using the unedited earnings values provided in the source data.

B. Construction of Hourly Wages using the Census/ACS

We use the Census IPUMS 5% data for the 1980, 1990, and 2000 Census and the 2012 (2010-2012) three-year aggregate data from the ACS. The sample includes whites between the ages of 25 to 64 who are not living in group quarters (i.e. non-military and non-institutionalized). Hourly wages are calculated as total annual wage and salary income divided by the product of weeks worked and usual hours worked in the previous year. The wage sample excludes the self-employed. Following Autor, Katz, and Kearney (2008), we drop the bottom 1% of hourly earners and multiply hourly wages of top-coded earners by 1.5. The maximum hourly wage is limited to 1.5 times the maximum annual income amount divided by 1,750 (35 hours per week for 50 hours per year). The earnings measures are deflated to 2000$ using the price deflator for personal consumption expenditures (PCE) from the BEA.

C. Coding of Education in the CPS and Census/ACS Samples

We follow steps outlined in Autor, Katz, and Kearney (2008) to create a comparable measure of the number of years of schooling across years in the CPS. In 1992, the CPS changed the education question. Figures from Park (1994) are used to assign years of completed education to each worked based on gender, race, and highest degree held. For the Census and ACS samples, we impute the number of years of education based on the variable indicating the individual's highest year of school or degree completed (*educd*).

Table A1: Convergence in Mean State Level Labor and Non-Labor Market Outcomes Over Time

	Unconditional		*Residualized*	
	State FE only	State FE and Year FE	State FE only	State FE and Year FE
	(1)	(2)	(3)	(4)
		Female-Male LFP Gap		
R-squared	0.185	0.952	0.194	0.944
		Female-Male Wage Gap Among Workers		
R-squared	0.129	0.949	0.166	0.929
No. of Obs (44 states by 7 time periods)	352	352	352	352
		Female Share Nevermarried (Age 20 to 40)		
R-squared	0.287	0.985	0.067	0.997
		Female Age at First Birth (Age 20 to 40)		
R-squared	0.410	0.959	0.248	0.976
No. of Obs (44 states by 4 time periods)	176	176	176	176

Note: Regressions use data from several years of Current Population Surveys (CPS) for labor market outcomes and from several years of Census/ACS fro non-labor market outcomes. See text for details.

Table A2: GSS Questions Used to to Construct Sexism Index

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)

Table A3: Convergence Over Time in Average State-Level Sexism

	State FE only (1)	State FE and Year FE (2)
	Overall Sexism	
R-squared	0.498	0.869
	Average Male Sexism	
R-squared	0.535	0.794
	Average Female Sexism	
R-squared	0.473	0.870
No. of Obs (44 states by 3 time periods)	107	107
Includes:		
State Fixed Effects	yes	yes
Year Fixed Effects	no	yes

Note: The three time periods are (1) 1977 to 1988 (2) 1989 to 1993 (3) 1994 to 1998. The number of observations do not add up to 44*3 = 132 due to missing observations in some states in some time periods. We restrict the analysis to state*time period cells with at least 15 male and female respondents.

Table A4: Estimates of Background Sexism, Controlling on Individual X's

	Labor Force Participation		Log Wages, conditional on working		Never Married		Age at First Birth	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	All Migrants				Female Migrants Only			
Sexist Beliefs in State of Birth:								
Overall*Female dummy	-0.007***	-0.008***	-0.003	-0.011***				
	[0.002]	[0.002]	[0.003]	[0.002]				
Overall	-0.009***	0.001	-0.008	0.002	-0.021***	-0.012***	-0.196***	-0.115***
	[0.003]	[0.001]	[0.006]	[0.003]	[0.003]	[0.003]	[0.045]	[0.021]
Controls:								
Average NAEP Reading and Math Scores in State of Birth		X		X		X		X
Female*Year FE	X	X	X	X		X		X
State of Residence FE	X	X	X	X	X	X	X	X
Year FE	X	X	X	X	X	X	X	X
Region of Birth FE	X	X	X	X	X	X	X	X
Outcomes residualized of:								
Education, Age dummies		X		X		X		X
Observations	15,103	14,083	15,063	14,055	7,522	7,016	7,453	6,962
R-squared	0.928	0.938	0.927	0.931	0.801	0.839	0.769	0.587

Note: Regressions use data from several years of Current Population Surveys (CPS) for labor market outcomes and from several years of Census/ACS for non-labor market outcomes. See text for details. Columns (2), (4), (6), and (8) use residual outcomes estimated from individual-level regressions that include controls for the number of years of education and age fixed effects. These columns also control for average reading and math scores for males and females in each state of birth. The test score data is from the NAEP Long Term Trend (NAEP-LTT), which was administered to a random sample of 9, 13, and 17 year olds in various years between 1971 and 2004. The number of observations in these columns are slightly smaller than that in the baseline specification as math scores are missing for two states (New Hampshire and Vermont).