NBER WORKING PAPER SERIES

SOCIAL NORMS, LABOR MARKET OPPORTUNITIES, AND THE MARRIAGE GAP FOR SKILLED WOMEN

Marianne Bertrand
Patricia Cortés
Claudia Olivetti
Jessica Pan

Working Paper 22015
http://www.nber.org/papers/w22015

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2016

Research support from the Global Asia Institute JY Pillay Comparative Asia Research Centre Grant is gratefully acknowledged. We thank seminar participants at the NBER Summer Institute, 2015 SOLE/EALE Meetings, Hitotsubashi Summer Institute in Labor Economics, Hunter College, and the 2016 ASSA Meetings for numerous helpful comments. Maria Dolores Palacios provided excellent research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Social Norms, Labor Market Opportunities, and the Marriage Gap for Skilled Women
Marianne Bertrand, Patricia Cortés, Claudia Olivetti, and Jessica Pan
NBER Working Paper No. 22015
February 2016
JEL No. J0,J01,J11,J12,J16

ABSTRACT

In most of the developed world, skilled women marry at a lower rate than unskilled women. We document heterogeneity across countries in how the marriage gap for skilled women has evolved over time. As labor market opportunities for women have improved, the marriage gap has been growing in some countries but shrinking in others. We discuss a theoretical model in which the (negative) social attitudes towards working women might contribute towards the lower marriage rate of skilled women, and might also induce a non-linear relationship between their labor market prospects and their marriage outcomes. The model is suited to understand the dynamics of the marriage gap for skilled women over time within a country with set social attitudes towards working women. The model also delivers predictions about how the marriage gap for skilled women should react to changes in their labor market opportunities across countries with more or less conservative attitudes towards working women. We test the key predictions of this model in a panel of 23 developed countries, as well as in a panel of US states.

Marianne Bertrand
Booth School of Business
University of Chicago
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
marianne.bertrand@chicagobooth.edu

Claudia Olivetti
Boston College
Department of Economics
Maloney Hall
Chestnut Hill, MA 02467
and NBER
claudia.olivetti@bc.edu

Patricia Cortés
Boston University
pcortes@bu.edu

Jessica Pan
Department of Economics
National University of Singapore
1 Arts Link
Singapore 117570
jesspan@nus.edu.sg

1 Introduction

Marriage rates have been declining throughout most of the industrialized world. In 2010, one in five adults aged 35 to 44 in the US had never married, compared to about 7% in 1970 (Pew Research Center, 2014). The marriage rate in the EU-28 has declined by close to 50 percent between 1965 and 2011, from 7.8 per 1 000 persons to 4.2 (Eurostat, 2015).[1] Similarly, the proportion of single women aged 35 to 39 has increased in East Asia (Jones and Gubhaju, 2009), reaching about one in five in Hong Kong and Japan by the mid to late 2000s compared to one in twenty, at most, in 1970. These declines in marriage rates have been accompanied by a steady rise in age at first marriage in many of these countries. These overall trends in the marriage rate have received widespread attention, and influential work has discussed the marriage market and fertility implications of women's advancement in education and labor markets (Becker, 1973, Goldin, 2006, Stevenson and Wolfers, 2007 and Greenwood et al., 2012).

A somewhat overlooked aspect of the discussion surrounding the overall decline in marriage has been how it has differentially affected skilled and unskilled women. In the US, research has documented a reversal over time of the skilled-unskilled marriage gap, with college-educated women today being as likely, if not more likely, to get married relative to those without a college education (Isen and Stevenson, 2010, Pew Research Center, 2010). In contrast, a number of countries in East Asia have been grappling with the reverse phenomenon, with highly educated women today marrying at a particularly low rate relative to their less educated counterparts (Economist, 2011, Hwang, 2015). For example, 35% and 20% of college-educated women between the ages 35 and 39 in Hong Kong and Korea, respectively, remained single in 2011. Many of these women will likely remain childless given how rare out-of-wedlock births remain in most of Asia (OECD Family Database, 2012).[2]

In this paper, we start by systematically documenting the marriage gap between skilled and unskilled women and its evolution over time for a group of 23 developed countries covering the period from 1995 to 2010. While skilled women overall marry at a lower rate than unskilled women, we show that this gap has been decreasing, in some cases even reversing, in North America, most Nordic countries, and some parts of Western Europe. On the other hand, the gap has remained constant or widened in East Asian countries as well as parts of Southern Europe. This divergence in relative marriage market outcomes for skilled women across developed countries has occurred despite consistent patterns of increased labor market opportunities for skilled women (and men) in all of these countries.

To explain these patterns, we propose a theoretical model in which (negative) social attitudes

[1] `http://ec.europa.eu/eurostat/statistics-explained/index.php/Marriage_and_divorce_statistics`

[2] In 2012, the share of births outside of marriage was just over 2% in Korea and Japan, compared to the OECD average of 39%. The share of out-of-wedlock births in the US is about 40% (OECD Family Database, 2012)

towards working women might contribute towards the lower marriage rate of skilled women, and might also induce a *non-linear* relationship between their labor market prospects and their marriage outcomes. Borrowing from Fernandez, Fogli and Olivetti (FFO, 2004), the key ingredient of the model is that the gender norm generates spousal disagreement over the provision of the household public good. Since skilled women have higher wages and hence are more likely to participate in the labor market, they provide less of the public good relative to unskilled women, which contributes to making them less attractive potential partners in the marriage market.

While FFO (2004) study how this marriage/work trade-off endogenously evolves from one generation to the next, in this paper, we abstract from such intergenerational considerations and explore instead, using a static model, the implications of the model for the interaction between economic opportunities, gender norms and marriage rates across skill groups. In particular, the model predicts a U-shaped relationship between the skilled-unskilled marriage gap for women and their market wage. Intuitively, at low wage levels, the loss in public good consumption due to the wife working, which is increasing in her market wage, is too large relative to the husband's gains from any externality from his wife's own private consumption. However, when the market wage is high enough, working women become increasingly more attractive relative to non-working women as husbands value the externality of their wife's higher income. Thus, assuming fixed or slow-changing social norms, the model predicts that as women's labor market opportunities increase, the marriage "penalty" for skilled women first increases, reaches a peak and then declines, eventually switching sign.

Comparing two societies that only differ in the conservativeness of their gender identity norms, the model predicts that the relative marriage deficit for skilled women will be higher in more traditional societies and that a given increase in skilled women's wages will be more likely to be associated with an *increase* in the marriage deficit for skilled women in the more traditional society, but a *decrease* in that deficit in the more gender equal society.

In the empirical section of the paper, we take the key predictions of the theory to a panel of 23 developed countries. As predicted by the model, we find a strong positive relationship between the relative deficit skilled women experience in the marriage market and the degree of gender-related conservatism. We also show that the relationship between the skilled-unskilled gap in marriage rates and educated women's labor market opportunities appear to differ markedly across groups of countries in a way that is consistent with the prediction of a U-shape relationship derived by our model, with increases in skilled women's wages decreasing their relative likelihood of marrying in more traditional countries but increasing this likelihood in countries with more gender equal norms.

We demonstrate the robustness of these results to alternative definitions of a long-term and stable relationship between men and women, such as definitions that include cohabitation or focus instead on fertility outcomes. We also show that the results are not driven solely by the East Asian countries, and are robust to alternative measures of gender norms.

A final implication of the theory that we bring to the data is with respect to women's educational choice. Our model predicts that in more conservative countries, a lower fraction of women are willing to acquire higher education as they anticipate facing greater barriers in the marriage market as skilled women. This is exactly what we find: the share of females with tertiary education in our sample of developed countries is negatively related to the strength of the gender norms. Moreover, consistent with the underlying mechanism in the model, we find that a relatively higher share of women are educated in countries where the marriage market deficit for skilled women is smaller.

Finally, and importantly, we replicate all of the empirical tests above in a panel of US states covering the 1970 to 2010 period. One key benefit of this replication exercise is that the variation we exploit within a single country is less subject to the unobserved heterogeneity concerns that mire most cross-country analyses. Exploiting differences across 45 states in the strength of gender norms, we show that the key findings in the country panel carry through in the US state panel.

Our paper is related to a small but growing literature that emphasizes how the interaction of economic development (and in particular, women's growing labor market opportunities) and social norms could impact marriage and fertility. Hwang (2015) uses a variant of FFO's (2004) dynamic model of intergenerational transmission of gender attitudes to rationalize the decreasing marriage rates experienced by female college graduates in developed Asian economies. In Hwang's (2015) model, mothers' education affects son's preferences: only men born to college educated mothers acquire non-traditional gender role attitudes.[3] Thus, Hwang shows, in a society with very few college-educated mothers, most men are traditional, and skilled women face a large marriage penalty. Another related paper is Kawaguchi and Lee (2014) who argue that the high demand for foreign brides in developed Asian economies is the consequence of improvements in women's economic status in countries characterized by very traditional gender roles. While both of these papers focus almost exclusively on rationalizing the recent East Asian experience, our paper takes a broader view and proposes a simple model that can reconcile the marriage patterns observed in a large sample of developed countries.

Finally, our paper is closest in spirit to earlier work by Feyrer et al (2008) who propose that cross-country fertility patterns could be explained by the interaction between the increasing status of women in the workforce and their status in the household. Like us, Feyrer et al (2008) suggest the possibility of a non-monotonic relationship between changes over time in the status of women in the workforce and an important social outcome (fertility in their case) due to slow changing norms. At low levels of female wages, women specialize in household production and fertility is high. In an intermediate phase, women have increasing opportunities in the labor market but still shoulder the bulk of household production, resulting in much reduced fertility. Finally, as women's labor market

[3]This is in contrast with FFO (2004), who assume that a man's preference for a stay-home wife depends on his mother's labor force participation, and not education.

opportunities further improve, men begin to share in the burden of child care at home, and fertility is higher than in the intermediate phase. Our specific focus is obviously different (differences in marriage market outcomes for skilled vs unskilled women and implications for women's educational choices) and we propose a formal model to guide our thinking, but at the core, both Feyrer et al (2008) and our paper share an interest in exploring the implications of slow moving gender norms in the face of growing labor market opportunities for women.

The rest of the paper is organized as follows. Section 2 describes the data, and presents some descriptive facts for the cross-country sample. Section 3 presents a static model of marriage, household decision making and education decisions. Section 4 lays out the empirical tests of the model and presents the results using the cross-country panel and the panel of US states. Section 5 concludes.

2 Cross-Country Sample: Data Description and Descriptive Facts

2.1 Data Description

Gender Norms

We use two main sources of data to measure cross-country differences in attitudes toward the role of women in society. The first dataset is the Integrated Values Survey (IVS) from 2006-2013, which is a harmonized dataset that covers both the European Value Survey (EVS) and the World Value Survey (WVS). This dataset provides a range of gender-related questions that were asked consistently across a broad set of countries. For our main analysis, we use the following question: "When jobs are scarce, men have more right to a job than women." The possible responses to the question are agree, disagree or neither. We interpret agreement to this question as expressing the view that it is more important for men to be employed in the labor market relative to women. We chose this question as it provides us with the broadest coverage of countries. For each country, we focus on the responses of individuals age 18 and older. Appendix Table 1 reports the average response to this question for each of the 23 countries in our sample as well as the years in which the survey questions were asked. The table also includes information on the region and sexism group (high, medium and low) that each country belongs to.

While this variable clearly measures some degree of conservatism toward gender roles, one concern is that it does not directly contrast the specific gender roles that are central in our model. The social norm in the model is based on the belief that some individuals may regard a woman's primary sphere to be in the home (provision of the household public good) while a man's primary sphere is in the labor market. To get at this norm more directly, we turn to a second data source, the

International Social Science Program (ISSP). The ISSP is a cross-country collaboration that seeks to build on pre-existing social surveys such as the General Social Survey (GSS) to allow for cross-country comparisons of social trends. Each year, the ISSP rotates a set of topics. Our analysis draws on the questions in the 2002 and 2012 waves of the Family and Changing Gender Roles module. To complement our measure from the IVS, we use the following question from the ISSP: "A man's job is to earn money; a woman's job is to look after the home and family." Respondents indicate their agreement to this statement on a five-point scale: agree strongly, agree, neither agree nor disagree, disagree, and disagree strongly. We code the response "agree" and "strongly agree" as indicating a greater degree of gender conservatism. The downside of using this question is that our sample is reduced by three countries.[4]

Table 1 summarizes the average responses to both questions by region, separately by gender, as well as by educational attainment for men. As expected, there is significant variation in these gender attitude measures across regions. On average, about a third of East Asians agree with both statements, but only between 3 to 8 percent of people in the Nordic countries do. Canada and the United States, as well as parts of Western Europe (UK, France, Belgium and Ireland), have more liberal gender norms compared to countries in Southern Europe and the Eastern part of West Europe (Germany, Austria, Switzerland, and Netherlands). Across countries, the correlation between the average responses to the IVS and ISSP questions is high (0.88).

In most regions, males are more likely to agree with the statements. Interestingly though, the gender gap within country is rarely larger than 5 percentage points.[5] Given that the differences in responses between males and females are not large (particularly in terms of the cross-country rankings), for our main analysis, we will focus on the average response to the questions across males and females. Because of the relatively small sample sizes available for each country in the attitude surveys, focusing on the average response also allows us to construct a more precise measure of gender-related attitudes across countries. Nonetheless, the empirical results are largely similar if we rely only on male responses to the survey questions.

One limitation of the cross-country data on gender attitudes is that the surveys were not conducted in the same year in each country. The survey dates range from 2006 to 2013 for the IVS, and from 2002 to 2012 for the ISSP. Although both the IVS and ISSP include data on attitudes for earlier points in time, we choose to focus on these specific years to ensure comparability and the broadest coverage across countries.[6] Specifically, to construct the country-specific aggregate measure of

[4]See Appendix Table 1 for the average response to this question for each country with ISSP data. Hong Kong, Greece, and Italy are excluded from the 2002 and 2012 ISSP samples.

[5]Larger gaps within countries are observed between skilled and unskilled men: educated men tend to be less conservative.

[6]The earlier years of the surveys cover a relatively small number of countries, making it difficult to distinguish actual differences in responses across countries from differences due to the different survey years. Therefore, rather

gender norms, we first remove the survey year effect by obtaining the residuals from a regression of the individual responses to each gender-related question on a linear term indicating the survey year for the IVS, and a dummy for the 2012 survey year for the ISSP.[7] Next, we use the residuals from the previous regression to create a country-specific measure of "average" gender norms, which is simply the mean across all years of the residual individual-level responses in a given country. For ease of interpretation, we standardize the resulting aggregate gender norm measure to have mean zero and standard deviation one in the sample of 23 countries.

One particular issue with our focus on the reported attitudes in these specific years is that the time period covered does not overlap with the earlier time period for which we are measuring the marriage market and labor market outcomes. For example, one might be concerned that if there are differential trends in gender-related attitudes over time across countries, focusing on attitudes at a single point in time could lead to potentially misleading inferences. Nevertheless, examining trends over time for each gender-related attitude question, plotted separately by the five regions (North America, two groups of countries in Western Europe, Southern Europe, Nordic countries, and East Asia) shows that while all the groups of countries have experienced some decline in the conservativeness of gender attitudes,[8] the relative ranking of countries in terms of average responses to both gender-related attitude questions has been largely constant over time. The cross-country rank correlation in the average response to the question "when jobs are scarce, men have more right to a job than women" across IVS survey waves (1990-1998 and 2007-2013) is approximately 0.89. This suggests that the country-level variation in gender-related attitudes is largely stable over time. For all our empirical analyses, in addition to the continuous measures of gender norms, we will also present results comparing outcomes across three groups of countries with high, medium, and low sexism.

Marriage Rates and Labor Market Outcomes

We use several datasets to construct marriage and labor market outcomes at the country level. Our main datasets are the European Union Labor Force Survey (EU-LFS), the European Union Statistics on Income and Living Conditions (EU-SILC), the European Community Household Panel (ECHP), and the Luxembourg Income Study (LIS), supplemented with Census and Labor Force Surveys for most non-European countries. The exceptions are Japan and Korea, for which we use smaller surveys. A description of data sources by country and year can be found in Appendix Table

than pooling all available data, we focus on the later time period where there is more consistent coverage of a large sample of countries.

[7] We do not control for the full set of year dummies for the IVS as the set of countries available in each year from 2006 to 2013 is quite different (see Appendix Table 1), hence, including the full set of year fixed effects would imply that the year effects would be identified off very different subgroups of countries. For the ISSP, since we only have two time periods, 2002 and 2012, we use a dummy for 2012 to capture the survey year effects.

[8] The only exception are countries in the Eastern part of West Europe in the ISSP.

2. For every country, we work with survey years closest to 1995, 2000, 2005, and 2010.

Our marriage outcomes are based on individuals aged 35-44. This age range was chosen as a compromise between having data for the most recent cohort (individuals aged 35-44 in 2010) and observing completed first marriage decisions among individuals in each cohort. The coding of the marital status variable varies by data source and country.[9] In most countries,"married" individuals include either formal unions or registered partnerships.

For each cohort of individuals in a given country and year, we construct measures of the labor market conditions that they are likely to face based on the labor market outcomes of individuals aged 25 to 54 in that country and year. These measures include the average annual wages of high-skill females and males, the skilled-unskilled wage premium for males and females, and the skilled-unskilled labor force participation difference for males and females.[10] The wage sample is based on full-time employees, defined as those working 35-plus hours per week. For a small number of countries where workers' full-time status is not available, we use the average wages for all employees. To facilitate cross-country comparisons, we convert all the country-specific annual wage measures to 2000 US$ using the PPP conversion factors from the World Bank. Details on the construction of the labor market variables can be found in the Data Appendix. Finally, following the Eurostat education classification, we define skilled individuals as those with completed tertiary education.[11] The data for GDP per capita (PPP) is from the World Bank. The summary statistics for the cross-country data are reported in Appendix Table 3.

2.2 Descriptive Facts

Figure 1 documents the cross-country variation, as of 2010, in the skilled-unskilled marriage rate gap for women and men. We define the skilled-unskilled marriage rate gap among women (y-axis) as the difference between the fraction of women with tertiary education between 35 and 44 years old that were ever married and the fraction of women with less than tertiary education in the same age range that were ever married. Reported on the x-axis is the equivalent skilled-unskilled marriage rate gap among men. It is apparent from Figure 1 that in the majority of the countries in our sample, more educated women marry at a lower rate than their less educated counterparts. In contrast, in the majority of countries, the marriage rate of more educated men is larger than the

[9] Details on the coding for each data source can be found in the Data Appendix.

[10] Ideally, we would have liked to construct labor market conditions for individuals age 35-44 in year t based on the labor market outcomes of those age 25-54 in year $t-10$ as this might be closer to the labor market conditions that were prevalent when making marital status decisions. However, due to data limitations, this is not feasible as this would entail a much smaller set of countries for our analysis. We do, however, use these more accurate measures in our US state panel exercise.

[11] This is defined as individuals who completed ISCED levels 5 or 6.

marriage rate of less educated men.

There are, however, some exceptions. In particular, in most Nordic countries (Norway, Sweden, Finland, Iceland), educated women marry at a higher rate than less educated women; this is also the case in the US, UK, Canada and Ireland. Also, skilled men marry at a lower rate than unskilled men in Italy, Austria, Spain, Switzerland and Greece. Nonetheless, even in these cases, the skilled-unskilled marriage gap is always smaller in absolute value for males relative to females.

Figures 2A and 2B show how the skilled-unskilled marriage market gaps among women and men, respectively, have been evolving over time. In particular, we report the difference in marriage rates between skilled and unskilled individuals aged between 35 to 44 years at four points in time: 1995, 2000, 2005 and 2010. For these figures, we classify countries into six different groups based largely on geography: North America, East Asia, Northern Europe, Southern Europe, and two groups of Western European countries.

In the US and Canada, while skilled women in the earliest cohort married at a lower rate than unskilled women in that same cohort, this gap diminished over time and had fully reversed in sign by the 2005 and 2010 cohorts (Figure 2A). A qualitatively similar picture emerges in the UK, France, Ireland and the Netherlands, with declining skilled-unskilled gaps in female marriage rates, even though only two of these four countries (UK and Ireland) had experienced a reversal of the gap by 2010. Most of the Nordic countries, with the exception of Denmark, have also experienced an overall decline over time in the probability of unskilled women marrying at higher rates than skilled women. In fact, as of 2010, all of the Nordic countries, again with the exception of Denmark, are characterized by a higher marriage rate for skilled women relative to unskilled ones. Interestingly, two of the Nordic countries (Sweden and Finland) display what appears as a U-shape pattern over time: while skilled women married at a higher rate than unskilled ones in 1995, this gap declined between 1995 and 2000, only to rise again after 2000.

The patterns over time for East Asian countries are quite different. While more educated women married at a much lower rate than less educated ones in the earliest cohort, this gap has been largely growing over time. The only exception to this pattern is Hong Kong, where the skilled-unskilled marriage gap is slightly smaller (in absolute value) in the most recent cohort (35 to 44 years old women in 2010) than in the first cohort (35-44 year old women in 1995). Southern Europe is characterized by a fairly large marriage gap between unskilled women and skilled women, and this is true throughout the sample period. While Italy shows only minimal changes over time, Portugal and Spain both appear to have experienced an increase in the gap (in absolute value) followed by a decrease in the last cohort, resulting in weak U-shape patterns.

Figure 2B reports similar trends for the marriage rate of skilled men compared to unskilled men. As indicated above, in 2010, educated men marry at a higher rate than their less educated counterparts

in most countries in our sample. It is apparent from Figure 2B that this pattern was the norm in most countries throughout the sample period. There are, however, a few exceptions, especially in the earlier cohorts in Western and Southern Europe.

A comparison of the trends between men and women across various groups of countries is also interesting. The contrast between men and women in East Asia over time is particularly striking. While the probability that educated women are married in those countries has been declining over time relative to the probability for less educated women (Figure 2A), the overall trend appears to move in the opposite direction among men, with educated men experiencing a higher relative marriage rate in the latest cohort than in the earliest cohort.

A central hypothesis motivating our paper is that the lower marriage rate for educated women relative to less educated women is partly a reflection of men's dislike for having a working wife. Because educated women have better labor market opportunities than less educated women, these women may experience more difficulty in finding a husband; also, these women may tend to be pickier in the marriage market as they have better outside options if they remain single. Both forces, as we will argue in the theoretical model below, contribute to educated women marrying at a lower rate than less educated women, at least until the value of the extra earnings that they can bring to the household start becoming large enough to undo the disutility that men may have when their wives work.

Figure 3 provides graphical evidence of a systematic correlation between the conservativeness of gender norms in a country and the differential marriage rate of skilled women relative to unskilled women in that country. Specifically, in the left panel of Figure 3, we graph the relationship between the gap in marriage rates for high-skilled vs. low-skilled women aged 35-44 in a country in 2010 and the IVS-based measure of gender role attitudes in that country. Figure 3 clearly shows that countries that are more conservative according to this measure are also countries where educated women marry at an especially low rate compared to less educated women. In contrast, we see a much weaker relationship for men (right panel of Figure 3).

In summary, the descriptive evidence in this section confirms that educated women marry at lower rates than their less educated counterparts in the majority of developed countries. There is evidence that the cross-country variation we observe in the marriage gap between educated and less educated women is related to the conservativeness of gender norms in these countries. The time series evidence suggests that there has been heterogeneity across developed countries in how the skilled-unskilled difference in female marriage rate has evolved over time. As labor market opportunities for educated women improved in all of these developed countries, the marriage "deficit" for skilled women has increased in some countries but declined in others, and in a subset of countries, has actually turned into a "surplus."

In the next section, we develop a simple theoretical model that can account for these key facts in the data. Most importantly, the model will be able to explain why the relationship between improvements in skilled women's labor market opportunities and their marriage outcomes might not be linear, either within a country over time, or between countries.

3 Model of Marriage and Household Decision Making

We analyze a static model of education, marriage and household's time allocation decision that ties together the fact that in some countries men overwhelmingly disapprove of married women working with the lower marriage probability faced by skilled women. A marriage gap between skilled and unskilled women emerges endogenously as a consequence of skilled women's higher market wages and the time allocation decisions that these generate.

Borrowing from Fernandez, Fogli and Olivetti (FFO, 2004), the key ingredient in the household model is that gender norm generates spousal disagreement over the provision of the household public good. Since skilled women have higher wages, they provide less of it relative to unskilled women. This makes them less attractive potential partners on the marriage market. Thus a woman's education decision involves a trade-off between a higher wage rate and a lower marriage probability.

In FFO (2004) this trade-off evolves endogenously because the education and marriage decisions of women in one generation affect the skill gap in marriage probability and, consequently, women's education decision, in the next generation. FFO (2004) assume that sons of more educated/working women have more egalitarian gender norms. If skilled women's earnings are sufficiently high so that men raised by educated/working mothers strictly prefer marrying an educated/working wife, an increase in the share of working mothers in one generation causes an increase in the share of educated/working women in the next generation because it improves their marriage prospects. Thus, holding everything else constant, as more women work, gender norms become more liberal and contribute to further increasing the fraction of women who choose to acquire education in the next generation.

In contrast to FFO (2004), our model does not endogenize gender norms but instead chooses to highlight how the presence and strength of the marriage/work trade-off faced by skilled women is a non-monotonic function of gender norms and economic opportunities. Holding gender norms constant, the model predicts a U-shaped relationship between the marriage gap faced by skilled women and their higher market wage. Intuitively, at low wage levels, the loss in public good consumption due to the wife working - the likelihood of which is increasing in her market wage - is too large relative to the husband's gains from the externality from the wife's consumption. However,

when the market wage is high enough, working women become increasingly more attractive relative to non-working women. Thus, assuming fixed or slow-changing social norms within a country, we should see that as women's labor market opportunities increase, the marriage deficit experienced by skilled women increases, reaches its peak and then declines, eventually switching sign.

The model allows us to perform some comparative statistics with respect to the strength of the gender norms. Holding wages constant, the model predicts that the relative deficit skilled women experience in the marriage market is more severe in more traditional countries. Also, the model shows that the range of wages over which the relative deficit skilled women experience in the marriage market increases before peaking is wider the more traditional the gender norms.

3.1 Model Set Up

The timing in the model is as follows. First, women choose whether to become skilled. Men and women then obtain a match at random in the marriage market. Given this match, they decide whether the match quality is high enough to warrant marriage or stay single. Married agents decide how to allocate their time between work and the household production of a public good (for example, children). Single agents do not produce a household public good (they remain childless) and simply consume their labor income.

To keep things simple we assume that there are two types, skilled S or unskilled U, for each sex. Women choose whether to become skilled, while men's type and its distribution in the male population are given. Skilled and unskilled men only differ in terms of market earnings.

3.2 Household Decisions

We model agents' preferences following FFO (2002). In this model the welfare of individual i who is married to another agent j, consists of utility from own private consumption c_i, some spillover α_i from the spouse's consumption c_j, utility from consumption of a household public good c, and utility from the quality of the match with j as perceived by agent i, q_{ij}.

One possible interpretation of the utility function is that individuals obtain utility from their career in a way that is proportional to its status or success as measured by wages. In such an interpretation, α_i is the utility agent i derives from her/his spouse's career. For men, however, the α_i can be especially low because of traditional societal views about women's role. This parameter captures the idea that a working/career wife might challenge the conventional idea of gender roles in a household (e.g. 'identity' as in Akerlof and Kranton, 2000). Note that under this specification the utility penalty, paid by men, associated with a working wife arises endogenously.[12]

[12]The same outcome can be generated under different assumptions on utility transferability or bargaining as long

Match quality $q \in [-\infty, \infty]$ is assumed to be a random draw from a distribution F. Each agent is endowed with a unit of time which is allocated between producing the household public good (t_i) and working in the market ($1 - t_i$).

We assume that given a total time investment of $T = t_i + t_j$, each agent obtains $c = Tn$ units of the public good. One interpretation of this good is children, where n is the (fixed) number of children, and the total time T invested in them determines their quality.

Each agent's private consumption is equal to her earnings, which is the product of the time the agent spends working and her wages, i.e., $c_k = (1 - t_k)w_k$, $k = i, j$.

Agent i's utility when married to agent j, V_i^j, is given by:

$$V_i^j(w_i, w_j, q_i) = \max_{0 \leq t_i \leq 1} [(1 - t_i)w_i + \alpha_i(1 - t_j)w_j + \beta \log(t_i + t_j)n + q_{ij}] \tag{1}$$

where i takes t_j as given and $\beta > 0$, $0 \leq \alpha_i < 1$.

In order to flesh out the properties of the model that are relevant empirically, we assume throughout $w_m > w_f$ so that women have a comparative advantage in home production. We also assume $w_m > \beta$. This assumption implies that men work full time irrespective of their marital status (or skill level). The assumption of linear preferences in private consumption and linear home production technology and the focus on th equilibria where men work full time makes the interpretation of the results more transparent.[13] Under these assumptions the first order condition yields two possible cases: (i) $w_m > \beta > w_f$, then $t_m = 0$ and $t_f = 1$; (ii) $w_m > w_f > \beta$, then $t_m = 0$ and $t_f = \beta / w_f$.

In the first case, the wife does not work and instead dedicates herself full time to the raising of children while the husband works full time. In the second case, the husband's situation is unchanged, but the wife works part time and raises children with the remainder of her time.

The economic intuition for why the relationship between women's economic opportunities and marriage probabilities is non-monotonic can be grasped from the analysis of the solution to the

as, because of disagreement about the allocation over different consumption goods, a married woman works more than would be optimal from her husband's standpoint (see for example, Lundberg and Pollak, 2003, and Mazzocco, 2007). Similarly, the result would be unchanged if we assumed that wives also pay a social penalty from working, as long as this is smaller than their husbands' and there is still spousal disagreement about time allocation decisions. On the other end, a similar outcome could also be generated by a simpler model where the utility penalty is taken as given (see FFO, 2004). We choose this specification because it allows conducting a rich set of comparative static exercises while maintaining a parsimonious theoretical representation.

[13] The gist of the analysis would be unaltered if we were to relax these assumptions. For example, using a CES representation for total time investment or non-linear preferences for private consumption would lead to interior solutions for both husband's and wife's time allocation. Alternatively, if we assume $w_m < \beta$ we could also have an equilibrium where neither of the spouses work on the market.

household's optimal time allocation problem. Recall that in our model both partners equally care about the public good (the share the same β). However, each partner cares more about his/her own private consumption than about the private consumption enjoyed by his or her partner, where each agent's private consumption is equal to her total earnings.

We start from the case $w_f < \beta$, that is women's earnings (i.e. their market productivity but also their private consumption) are below the marginal utility from consumption of the public good. It is then optimal from a woman's standpoint to stay home once married and dedicate all her time to the production of the public good. Husband's preferences are perfectly aligned with his wife's. There is no spousal disagreement. The gender norm does not bind.

In the second case, women's wage is higher than the marginal utility from consumption of the public good ($w_f > \beta$). It is now optimal from the wife's standpoint to split her time between working at home and working on the market. However, at relatively low levels of w_f the positive externality that a husband's receives from wife's increased private consumption is not high enough to compensate for his utility loss coming from the lower production of the public good. As we show below the husband's utility loss is non-monotonic. For low values of w_f his utility drops below the utility that he would obtain from having a nonworking wife. However, as w_f increases the utility loss first peaks and then declines as the externality coming from the increased wife's private consumption increases. Eventually, the utility loss will turn into a premium. That is, for relatively high w_f the positive externality from the wife's private consumption dominates the utility loss in terms of public good consumption.

To formalize this discussion, we characterize the utility of a married man, V_m, as a function of his wife's wage.

First let's consider how V_m varies for $w_f \in [0, \beta)$. The wife does not change her time allocation in response to changes in w_f (she is devoted full-time to household production). Consequently, there is no effect on the husband's utility. Substituting from the first order conditions we get: $V_m(w_f) = w_m + \beta \log n \equiv \overline{V}_m$ where we denote with $\overline{V}_m$ the utility a married man receives if he marries a stay at home wife.

Let's consider the second case when $w_f > \beta$ and therefore a wife works $t_f = \frac{\beta}{w_f}$ at home and its complement to one on the market. Substituting from the first order conditions we get:

$$V_m(w_f) = w_m + \beta \log n + \alpha_m(w_f - \beta) + \beta \log \frac{\beta}{w_f}, \tag{2}$$

which we can re-write as the utility difference between marrying a working or a non-working wife:

$$V_m(w_f) - \overline{V}_m = \alpha_m(w_f - \beta) + \beta \log \frac{\beta}{w_f} \tag{3}$$

with:

$$\frac{\partial(V_m - \overline{V}_m)}{\partial w_f} = \alpha_m - \frac{\beta}{w_f}. \tag{4}$$

First, let's note that since $w_f > \beta$ the second term on the right hand side of equation (3) is negative. Thus, whether a man experiences a utility loss or a premium depends on the sign of $\alpha_m(w_f - \beta) + \beta \log \frac{\beta}{w_f}$. One can show that this difference is negative for $w_f \in (\beta, w_f^*(\alpha_m, \beta)]$ and positive for $w_f > w_f^*(\alpha_m, \beta)$, where $w_f^*(\alpha_m, \beta)$ is the the wage rate that makes a husband indifferent between the two (i.e. $V_m(w_f^*) = \overline{V_m}$).

We analyze how the penalty varies over $w_f \in (\beta, w_f^*(\alpha_m, \beta)]$ by studying equation (4). For $w_f \in [\beta, \frac{\beta}{\alpha_m}]$, the husband's utility is decreasing as a function of w_f; in this interval an increase in w_f leads the wife to decrease the amount of time spent at home thereby decreasing the man's utility by $(\beta - \alpha_m w_f)dw_f/w_f$. For $w_f > \beta/\alpha_m$, the husband's utility is strictly increasing in w_f despite the fact that the wife is putting in less time at home.

The utility loss hits its maximum if $w_f = \frac{\beta}{\alpha_m}$. That is, when the loss in production of the public good due to wife working ($\frac{\beta}{w_f}$) is equal to the gain from the externality deriving from wife's consumption (α_m). Once w_f is past this level the positive externality starts dominating until, eventually, the penalty becomes a premium.

Figure 4 depicts the utility differential for different values of α_m (panel A) and β (panel B). Panel A reveals how the gender norm affects both the sign, size and shape of the husband's utility loss from having a working wife as a function of her economic opportunities. The solid line shows the husband's utility differential when $\alpha = 1$, that is, if there is no spousal disagreement about time allocation decisions. In this case a husband gains from having a working wife and this premium monotonically increases with the wife's wage. The three remaining lines plot the utility differential for alternative values of the gender norm ($\alpha_H = 0.6, \alpha_M = 0.4, \alpha_L = 0.2$), keeping the utility of public good consumption constant. In the most gender equal society (the dashed dotted line) husband's and wife's preferences are more aligned. The range of w_f over which the husband of a working wife experiences a utility loss (relative to a housewife's husband) is relatively small. A husband cares sufficiently about his wife's career. It takes a modest increase in w_f for the positive externality effect to dominate the utility loss. In the intermediate society (the short dash line) the utility loss experienced by the husband is always larger than in the more gender equal society and the utility differential is negative for a wider range of w_f. It takes a large increase in wife's economic opportunities to compensate a husband's for the wife's under-provision of the public good. Finally, in the most conservative society (the long dash line), men care so little about their wife's private

consumption that the utility differential from having a working wife is always negative. It will take a huge increase in wife's wage for the penalty to turn into a premium.

Panel B shows how the utility differential varies for two societies that differ in terms of the value placed on the household good (e.g. time spent with children, home cooked meals etc.). We denote this values as β_L and β_H. The model's predictions are similar to those obtained when comparing societies with different gender norms. In addition, the threshold wage making a woman indifferent between working and specializing in home production varies across societies. It will be thus important as we empirically test the predictions of the model in the following sections to verify that the heterogeneity we observe across countries and over time with respect to their gender norms is not masking correlated heterogeneity with respect to the value placed on the household public good.[14]

It is also useful to characterize the utility of a married woman, V_f, as a function of her husband's wage w_m and her own wage w_f. Since we focus on equilibria where the husband works full time there is no scope for spousal disagreement about the husband's time allocation. So, in the interest of simplifying notation, we set $\alpha_f = 1$, that is, women like husband's career as much as their own. Note that under this assumption, husband's and wife have the same utility when the wife stays at home. Instead, if $\alpha_f < 1$ a housewife's utility is always smaller than that of her husband.

If $w_f \leq \beta$, she specializes in home production $t_f = 1$, her utility when married will only depend on her husband's wage. Substituting from the first order conditions this is given by: $V_f(w_m) = w_m + \beta \log n$.

In the second case, $w_f > \beta$ and $t_f = \frac{\beta}{w_f}$ so that:

$$V_f(w_f, w_m) = w_m + \beta \log n + (w_f - \beta) + \beta \log \frac{\beta}{w_f} \tag{5}$$

Comparing these two equations it is clear that women's utility is higher the higher is their husband's wage, irrespective of their work status. Men's earning opportunities have a premium both on the labor market and on the marriage market. Moreover, for given w_m, we can study the wife's utility differential between working and staying at home as a function of her own wage. Rewriting equation (5) we get $V_f(w_f, w_m) - V_f(w_m) = (w_f - \beta) + \beta \log \frac{\beta}{w_f}$. This function is always positive (and convex) for $w_f > \beta$ indicating that the utility of a working wife is always higher than that of a now-working wife. Note that if $\alpha = 1$, wife's and husband's preferences are perfectly aligned.

[14] We tackle this issue in Section 4.2 for the cross-country analysis, and in Appendix Tables 6 and 7 for the cross-state analysis.

3.3 Marriage Decision: Random Matching and Threshold Quality

We next turn to the matching part of our model. As discussed above, we consider two, exogenously given, types of men, skilled and unskilled, with $w_{mS} > w_{mU} > \beta$. Women choose their type. We assume that female unskilled wages w_{fU} lie in the interval $[0, \beta]$ (i.e., unskilled married women do not work) and skilled wages $w_{fS} > \beta$ (i.e. skilled women work irrespective of their marital status).

The matching process is modeled as a one-period random search in which the probability of a given individual meeting another individual of type j depends on its proportion in the relevant population of the opposite sex. Defining π_i as the share of skilled individuals in the population of gender i. The probability of a woman meeting a skilled man equals π_m, the probability she meets an unskilled man is $1 - \pi_m$.

Matched individuals each obtain a random draw of match quality $q \in [-\infty, \infty]$ from a log-concave distribution F. Individuals then decide whether to stay in a match (marry) and obtain married utility V_i^j as in equation (1) or to remain single whereby her/his utility is given by $U(w_i) = w_i$, that is, there is no household public good nor any externality from another agent's consumption.

Hence agent i will be willing to marry agent j if and only if:

$$\alpha_i(1 - t_j)w_j - w_i t_i + \beta \log(t_i + t_j)n + q_i \geq 0 \tag{6}$$

where i indexes the skill type of the man if i is a male or the skill type of the woman if i is female, and similarly for j.

Thus, we can solve for the reservation qualities, q^*, of males and females. For males, this yields:

$$q_m^* = \begin{cases} -\beta \log n & if \quad w_f < \beta \\ -\beta \log n - \alpha_m(w_f - \beta) - \beta \log \frac{\beta}{w_f} & if \quad w_f > \beta \end{cases} \tag{7}$$

Assuming $w_{fU} \leq \beta$ and $w_{fS} > \beta$, the first line of q_m^* is men's reservation quality if they meet an unskilled woman, whereas the second is the reservation quality when they are matched with a skilled woman. Note that male pickiness is invariant to his own wage (since he works full time whether married or single) and is invariant to female wages if these are below β, that is, for unskilled women. Let's denote with q_{mS}^* and q_{mU}^* men's pickiness when meeting a skilled woman or unskilled woman, respectively. The skilled-unskilled differential in male pickiness is given by $q_{mS}^* - q_{mU}^* = -\alpha_m(w_f - \beta) - \beta \log \frac{\beta}{w_f}$. Note that since we have assumed that skilled and unskilled men do not differ by gender role preferences, men's threshold does not depend on their own type. The comparative statics results for the effect of an increase in female wages on men's reservation quality differential obviously mirrors the analysis of married men's utility. If $\beta < w_{fS} \leq w_f^*(\alpha, \beta)$ this differential is positive. Men experience a utility loss from marrying a skilled working woman

and therefore are pickier when they meet one. The skilled-unskilled reservation quality differential over this interval has an inverted U-shape. However, if $w_{fS} > w_f^*(\alpha, \beta)$ the differential is negative. Men gain from having a working wife when her wage is high enough to compensate for the loss of public good production. It follows that they are less picky towards skilled women. As discussed in the next section, this pattern drives the finding of a U-shaped skilled-unskilled marriage differential as a function of women's economic opportunities.

The comparative static exercise also shows that for any $w_{fS} > \beta$ the threshold quality for skilled women is higher the lower is α and the higher is β.[15]

For women, the same calculation yields:

$$q_f^* = \begin{cases} -\beta \log n - (w_m - w_f) & if \quad w_f \leq \beta \\ -\beta \log n - (w_m - \beta) - \beta \log \frac{\beta}{w_f} if & w_f > \beta \end{cases} \tag{8}$$

Since we assume $w_f^U \leq \beta$ and $w_f^S > \beta$, the first line of q_f^* is the reservation quality of unskilled females, whereas the second is that of skilled females. Skilled women will be pickier than unskilled women since they have a better outside option.[16] Female pickiness is always increasing in her own wage, linearly for $w_{fU} < \beta$ because it only affects her consumption when single, and decreasing in men's wages (see comparative statics in Appendix **??**). That is, both skilled and unskilled women will be pickier towards unskilled men because of their lower earnings.

3.4 The Skilled-Unskilled Marriage Gap

Let's define $\Pi_{f,ij}$ the probability that, for a woman of type i meeting a man of type j, the random draw of the match quality q_{ij} lies above each partner's threshold. This is given by $\Pi_{f,ij} = \int_{q_{m,j}^*}^{\infty} \int_{q_{f,ij}^*}^{\infty} dF(q)dF(q) = \overline{F}(q_{m,ji}^*)\overline{F}(q_{f,ji}^*)$, where $\overline{F} = (1 - F)$ is the complementary cumulative distribution function. Note that under our maintained assumption of common gender preferences across men's type, men's threshold quality for unskilled and skilled women is the same irrespective of their own type. That is, $q_{m,ij} = q_{mj}$. It follows that the skilled-unskilled differential in marriage probability does not depend on the share of skilled men in the population. As we discussed in section 3.6, this assumption can be easily relaxed leaving the bulk of the analysis unchanged.

The probability that a match is formed is given by the product of $\Pi_{f,ij}$ times the probability of meeting a men of type j in the population, $j = S, U$. Having defined π_m as the fraction of skilled

[15] The partial derivatives are as follows: $\frac{\partial q_m^*}{\partial w_m} = 0$, $\frac{\partial q_f^*}{\partial w_m} = -\alpha_f$; $\frac{\partial q_m^*}{\partial w_f} = -\alpha_m + \frac{\beta}{w_f}$ if $w_f > \beta$ and it is otherwise equal to 0; $\frac{\partial q_f^*}{\partial w_f} = \frac{\beta}{w_f}$ if $w_f > \beta$ and equals 1 otherwise. $\frac{\partial q_m^*}{\partial \alpha_m} = \beta - w_f$ if $w_f > \beta$ and is equal to 0 otherwise.

[16] Taking the difference between the two, we obtain $q_{fS}^* - q_{fU}^* = (\beta - w_{fU}) - \beta \log \frac{\beta}{w_f}$ which is always positive for $w_{fU} \leq \beta < w_{fS}$.

men in the population, it follows that the marriage probability of a woman of type i will be given by:

$$\Pi_f^i = \overline{F}(q_{mi}^*)\left[\pi_m \overline{F}(q_{f,iS}^*) + (1-\pi_m)\overline{F}(q_{f,iU}^*)\right] \tag{9}$$

The skilled-unskilled difference in marriage probability is then given by $\Pi_f^S - \Pi_f^U$. We analyze how this differential varies for given values of w_{fS}, while keeping the unskilled wage fixed, $w_{fU} \leq \beta$.[17]

Figure 5 provides a graphical representation of the marriage penalty for three different values of α, high, medium and low, having fixed the unskilled and skilled wages in the shaded areas denoted with U and S (or S'), respectively. For a given value of α, the marriage gap mirrors the behavior of men's threshold qualities.[18] For values of $w_{f,S}$ below β (i.e. skilled women are also stay-at-home wives), the difference in marriage probabilities is negative because skilled women have better outside options. For values of $w_{f,S}$ lying in the interval $(\beta, w_f^*(\alpha,\beta)]$ all the terms in (9), which defines a woman's marriage probability, are lower for skilled women. The skilled-unskilled marriage gap is U-shaped in this range. It increases as men's threshold quality for a skilled woman increases, then declining for $\frac{\beta}{\alpha} < w_{f,S} \leq w_f^*(\alpha,\beta))$. Finally, if $w_{f,S}$ is above $w_f^*(\alpha,\beta)$, both types of men prefer skilled women. In this range, the marriage penalty can turn into a premium. The marriage gap is always smallest in the least conservative economy (α_H) and largest in the most conservative economy (α_L). The comparison of the three curves in the figure also reveals that while for the α_H and α_M economy the gap can eventually turn into a premium, it is unlikely that in the α_L economy women's economic opportunities will ever be high enough to compensate husbands for their working wives' under-provision of the public good. The figure also shows how an increase in $w_{f,S}$ might have very different consequences on the skilled-unskilled marriage differential depending on a country's prevailing gender norms. For example, suppose that $w_{f,S}$ increases within the shaded area S. The skilled-unskilled marriage gap would decline in the more gender equal economy but increase in the remaining two economies, with the increase being largest in the most conservative one. Suppose instead that $w_{f,S}$ increases from a value in the shaded area S to a value in S'. In this case we would observe the skilled-unskilled marriage gap turning into a marriage premium in economy H, still negative but decreasing in economy M and negative and further increasing in economy L. Panel B depicts the skill differential in marriage probabilities of men. Because of their higher wage, both types of women have a strict preference for skilled men. It follows that the skill differentials is always positive and increases as $w_{f,S}$ increases.

[17] In this case increasing w_{fS} is equivalent to an increase in the skill premium for women. A similar intuition would apply if we let both w_{fS} and w_{fU} increase, while keeping the skill premium constant. This is true even if $w_{fU} > \beta$ (i.e. unskilled women work), at least for some parameter combinations.

[18] Holding constant $w_{f,U}$, the shape of the skilled-unskilled marriage differential depends on the behavior of the derivative of Π_f^S with respect to $w_{f,S}$. It can be shown that, if F is log-concave and α sufficiently smaller than 1, there is a region of parameters where the marriage differential is U-shaped.

3.5 Education Decision

The expected utility V^i of a woman of type $i = S, U$ given the model parameters $\theta = (w_m, w_f, \alpha, \pi_m)$ is given by:

$$V^i(\theta) = \sum_{j=s,u} \Pi_{f,ij} V_i^j + \left(1 - \sum_{j=s,u} \Pi_{f,ij}\right) U(w_i) \tag{10}$$

We assume that a woman faces an idiosyncratic (utility) cost of becoming skilled of γ where the latter is an iid random draw from a continuous cumulative distribution function $G(\gamma)$ with support $[0, \infty]$.

Let

$$\Delta(\theta) \equiv V^S(\theta) - V^U(\theta) \tag{11}$$

be the relative utility of a skilled female (relative to an unskilled female) given the vector of model parameters θ.

Note that since wages are set, the expected utility differential between skilled and unskilled women is independent of the proportion of women who decide to become skilled. All women with $\gamma \leq \Delta(\theta)$ will decide to become skilled. The equilibrium π_f at any point in time is at the intersection of $\Delta(\theta)$ and $\gamma = G^{-1}(\pi_f)$.

3.6 Discussion

An immediate implication of the search model is that, if we compare two economies that differ by their gender roles, the expected utility of skilled relative to unskilled women would be larger in the less traditional country. This implies that, holding all the other parameters constant, the country with more equal gender norms should have a higher proportion of skilled women and a smaller skilled-unskilled marriage gap.

It is also the case that the expected utility of skilled relative to unskilled women increases in w_{fS}. Thus a higher skilled wage is generally associated with an increase in the proportion of women who choose to become skilled. As we discussed above, at low levels of w_{fS} the increase will be relatively small because of the stronger trade-off between the labor market return and the penalty coming from the higher rejection rates in the marriage market. However, the proportion of skilled women grows faster once the marriage penalty associated with the investment decision declines.

The model also implies that comparing three types of societies that differ in the conservativeness of their gender roles views from very liberal to very conservative, an increase in wages might at the same time increase the marriage penalty faced by skilled women in the more traditional society,

leave the marriage penalty in the middle society mostly unchanged but decreases it in the more gender equal society.

Finally, while we do not formally solve for this case here, we note that the model's predictions are unaltered if we also allow for within country variation in gender norms by education. This additional assumption would be in keeping with the empirical evidence (see Table 1), showing that skilled men tend to be less conservative than unskilled men, though the variation in gender norms across countries tend to be larger than that within a country.

The main difference under such an additional assumption would be that the trade-off faced by skilled women will also be a function of the share of skilled men in the population (as in FFO, 2004). The model would then deliver an additional prediction: an increase in the fraction of educated men reduces the marriage penalty, all else equal. In this case, considering two economies with the same α but different proportion of skilled men, the skilled-unskilled marriage gap should be larger in the country were fewer men are skilled.

4 Empirical Tests of the Model

4.1 Cross-Country Evidence

A first prediction of the theoretical model is that, all else equal, the skilled-unskilled marriage gap should be larger (in absolute value) in countries with more conservative gender attitudes. Suggestive evidence of such a relationship was presented in Figure 3 for the cohort of women that were between 35 and 44 years of age in 2010.

Table 2 shows that this relationship is robust to a variety of controls. The baseline correlation corresponding to the left hand panel of Figure 3 is reported in column (1) of Table 2. The gap in ever-married rates between high and low skilled women in low-sexism countries is 11 percentage points smaller than in high-sexism countries and 4 percentage points smaller than in mid-sexism countries.[19] Columns (2) and (3) of Table 2 show that the relationship is robust to a variety of controls, and in particular controls aimed at proxying for women's labor market opportunities, such as a quadratic in log skilled women's wages (column (2)) or a quadratic term in log(GDP per capita) (column (3)). In addition, columns (2) to (10) include controls for a quadratic in log skilled men's wages, the skilled-unskilled wage premium for men and women, the skilled-unskilled labor force participation premium for men and women, and the fraction of men aged 35 to 44 with tertiary education. Columns (4) and (5), respectively, replicate columns (2) and (3) but rely on

[19] See Appendix Table 1 for the classification of countries into the three sexism groups. This classification is based on the IVS measure of gender attitudes.

the continuous measure of gender attitudes constructed from the IVS instead of the grouping of countries into three categories. Again, we see that more conservative gender norms are correlated with a higher marriage deficit among skilled women.

Columns (6) to (8) replicate the specification in column (3) for the three other cohorts of women that we observe in our data, namely those aged 35 to 44 in 1995, 2000, and 2005. We see that the negative correlation between sexism and the relative marriage rate of skilled women also exists for the 2005 and 2000 cohorts (especially when it comes to high sexism levels); however, there is no statistically significant correlation among women aged 35 to 44 in 1995.

Another prediction of our model is that the relative marriage outcomes of skilled women should be more likely to be affected by the gender norms in their country relative to skilled men (see Figure 5). As shown in Figure 3, while there is a strong cross-country relationship between the marriage rate deficit for skilled women and the conservatism of gender norms, this relationship appears weaker for men. The remaining columns of Table 2 confirm this fact in a multivariate regression. Columns (9) and (10) of Table 2 replicate columns (2) and (4), respectively, but use as an alternative dependent variable the difference in ever married rates between skilled and unskilled men that were between 35 and 44 years old in 2010. As expected based on the visual inspection of data in Figure 3, while the estimates on the sexism measure (either using the three categorical variables or the continuous index) are of the same sign as that for women, they are quantitatively and statistically weaker. In other words, while gender role attitudes in a country appear to be a robust predictor of the relative marriage rate of skilled and unskilled women in that country, the same does not hold for men.

Next, we turn to empirical tests of some of the more subtle predictions of the model regarding the heterogeneity across countries in the marriage gap as labor market conditions change. Specifically, our model predicts that increased labor market opportunities for skilled women may, in fact, negatively impact their marriage rates, especially in countries with more conservative gender attitudes.

Using the pooled sample of countries from 1995 to 2010, in Table 3, we regress the difference in ever-married rates for skilled vs. unskilled women on a proxy for skilled women's opportunities in the labor market: $\log$*(high-skilled female wages)* (columns 1-3) or $\log$*(GDP per capita)* (columns 4-6), and the interaction between the proxy for women's opportunities and country-level sexism. We use $\log$*(GDP per capita)* as an alternative proxy for skilled women's opportunities in the labor market for two reasons - first, compared to the GDP measure, cross-country comparisons of wages are likely to be subject to a higher degree of measurement error. Measurement error is likely to arise since the wage information is drawn from numerous different surveys, in some cases cover different samples of workers (e.g. full-time vs. all workers), and we had to rely on different sources of data to adjust for inflation and purchasing power. Second, wage information is not available in some years for some countries; hence, using the GDP measure, we are able to include five more

country*year observations in our regressions reported in Table 3 (see column (4)). All regressions include country and year fixed effects, and standard errors are clustered at the country level.

The baseline coefficients in columns (1) and (4) indicate that higher labor market opportunities for skilled women are associated with a decrease in the deficit they experience in the marriage market in low sexism countries, but an increase in that deficit in high sexism countries. For example, in column (1), we see that a 10 percent increase in the wage of high skilled women reduces the marriage market "penalty" for high skilled women by 0.6 percentage points in low sexism countries, but increases it by 0.6 percentage points (.06-.12) in high sexism countries.

Columns (2) and (5), respectively, replicate column (1) and (4) but also account for a vector of other time-varying controls: share of males with tertiary education, the female and male wage skill premium, the female and male labor force participation skill premium, and a quadratic in female and male high skilled wages. The patterns are qualitatively unchanged. Finally, in columns (3) and (6), we replace the 3 categorical variables for sexism level with the continuous index. Again, the results are qualitatively unchanged: the relationship between the skilled-unskilled gap in marriage rate and labor market opportunities becomes more negative as the degree of sexism in a country increases. The point estimates for the estimated coefficient on the interaction term between high skilled female wage and sexism is -.057 in column (3) (significant at the 10 percent level), and the coefficient on the interaction term between log(GDP per capita) and sexism is -.169 in column (6) (significant at the 1 percent level).

To assess the magnitude of the results, we apply the observed change by sexism group in high skilled female wages from 1995 to 2010 to our panel model estimates (column (1) of Table 3) and compare these predictions with the observed change in the marriage gap. We find that our model explains between 40 to 55 percent of the observed increase in the marriage gap for the most conservative countries, and 60 to 80 percent of the decline observed in the least sexist countries.[20][21]

Finally, in the last two columns of Table 3, we replicate the analysis of columns (2) and (3) but now use the skilled-unskilled gap in marriage rate among men as the dependent variable. While the point estimates on the interaction terms of interest are of the same sign as that in the regressions using the female sample, their magnitudes are much smaller and none are statistically significant.

A final implication of our model that we can bring to the cross-country data is with respect to

[20]The two numbers per sexism group come from applying either the mean change (55 and 60 percent) or the median change (40 and 82 percent). As an example, the observed mean change in high skilled female wages for the high sexism group was 21 percent. When we multiply 0.21 by -0.057 (-0.12+0.063) we get a predicted change in the gap of -0.012, which accounts for 55 percent of -0.025 (the observed mean change in the marriage gap for this group of countries).

[21]We obtain similar results when we use the GDP model (Table 3, column (4)), though the percentages explained by our model are larger, between 49 and 150 percent.

educational choice. Our model predicts that in more gender conservative countries, a lower fraction of women should decide to become skilled as educated women in those countries expect that they will face greater barriers in the marriage market. This is exactly what we find in the country-level data. Figure 6 plots the gender gap in higher education between men and women that were between 35 and 44 in 2010 in each country as a function of the conservativeness of gender role attitudes in that country. We observe a strong negative relationship. The gender gap in education is smaller in the Nordic countries and larger in the East Asian countries.

Table 4 confirms the negative relationship between female educational attainment and the degree of conservatism of gender norms in a multivariate regression setting. All the regressions control for the share of males with a tertiary education in each country. The basic relationship, as reported in column (1) using the three sexism levels, is robust to controlling for a quadratic in female and male high skilled wages, female and male skill premium, and the female and male labor force participation premium (column (2)). The point estimates in column (2) suggest that the share of females with tertiary education is about 8 percentage points higher in mid sexism countries compared to low sexism ones, and 15 percentage points higher in high sexism countries compared to low sexism ones. We obtain qualitatively similar results when we use the continuous sexism index in column (3).

In column (4), we show that there is a strong positive correlation between the share of females with higher education and the skilled-unskilled marriage gap for females. Similar to column (3), we estimate the regression including the usual set of state-level controls for high-skill men and women's wages and the skill premia. Hence, women's educational achievement is higher in those countries where skilled women marry at a rate that is closer to that of unskilled women. While certainly not definitive, this correlation is consistent with the mechanism that our model predicts linking the gender gap in education to gender norms: the value for women to acquire education tends to be lower in more sexist countries because of the potential negative impact of additional education in the marriage market.

4.2 Robustness

α or β? As discussed in the theory section, cross-country differences in the value placed on the household good could also generate patterns similar to those we have observed in the prior tables. This is potentially a worry for our analysis if there is a strong positive correlation between such preferences for the household good and the strength of gender norms, in which case our empirical findings so far might just be picking up on the value placed on the household good.

To assess this potential threat to our analysis, in Table 5, we investigate whether such a correlation exists across countries. To proxy for the value placed on the household good (β in our model), we compute the total household time devoted to the production of the public good. We rely on data

from the 2012 International Social Survey Program (ISSP) and focus on the following questions: "How many hours spent on household work"; "How many hours spent on family members"; "How many hours spouse, partner spend on household work"; "How many hours spouse, partner spend on family members." We restrict the sample to individuals between 18 and 64 years of age and with at least one own child living with them.[22]

Based on these questions, we construct the following variables (all measured in hours per week): total time spent by parents in (1) total household production (household work + care for family members) (2) household work and (3) care for family members; total time spent by female parent in the three activities; and total time spent by male parent in the three activities. From this micro data, we then estimate, for each variable above, a country dummy based on regressing the variable on country fixed effects, number of children, a dummy for the presence of children under 5 years of age, household size, education and age of the respondent, and a dummy indicating if the respondent has a partner living in the household. In the final step, as reported in Table 5, we regress the estimated country dummies for each of the nine variables above on the continuous index of sexism in each country. Figure 7 provides a visual inspection of this analysis for time spent on household production (e.g. household work + care of family members).

As is apparent from Table 5 and Figure 7, there appears to be no systematic correlation between total (e.g. husband plus wife) time spent on household production and gender norms. Not surprisingly, gender norms correlate strongly with the distribution of time spent in household production between spouses: women do a greater share of the non-market work in more sexist countries. But the sum of the household work across both spouses does not correlate with gender norms. Finally, in columns (1) to (3) of Table 6, we re-estimate the main specifications reported in Table 2 including the proxy for β (total time spent on household production) as a control. Reassuringly, the results remain unchanged. In sum, we find little support for the possibility that our results might be picking up on differential valuation of the public good across countries.

Cohabitation. Another concern with our analysis so far is that marriage might not be a core outcome of interest, or more precisely, it might be a poor proxy for the prevalence of long-term partnerships within a country. This could be especially problematic for countries where long-term cohabitation is common, and couples live in a marriage-like relationship without actually entering a formal union. To account for this, we complement our analysis in Tables 2 and 3 with alternative dependent variables that might better capture the existence of such long-term partnerships. We construct an alternative measure of marriage that also includes cohabitation. We compute the skilled-unskilled gap in ever married or currently cohabiting rates for women that are between

[22]The 16 countries for which we have data to conduct this analysis include Austria, Canada, Switzerland, Denmark, Germany, Spain, Finland, France, Ireland, Iceland, Japan, Korea, Norway, Sweden, Taiwan, and the UK.

35 and 44 years old in 2010.[23] Unfortunately, cohabitation is only available in the EU-SILC, and therefore this analysis must be restricted to the subset of European countries and to a cross-sectional analysis. As is apparent from Appendix Figure 2, there is a strong positive correlation between the cohabitation rate for skilled and unskilled women (Panel A) and a strong correlation between our standard measure of the skilled-unskilled gap in marriage rates and this more inclusive measure that also captures cohabitation (Panel B).

Columns (4) to (6) of Table 6 replicate columns (1), (2) and (4) of Table 2, respectively, for this alternative measure of the skilled-unskilled gap in long-term partnerships. We see that more conservative gender norms reduce the likelihood of marriage or cohabitation for skilled women relative to unskilled women. In other words, our results in Table 2 do not appear sensitive to the possible miscoding of some long-term partnerships.[24]

Fertility. Yet another way to proxy for the existence of a long-term partnership between a man and woman is to look for the presence of children. While fertility is an important outcome in itself, it is also measurable for a large number of countries and for a longer period of time than cohabitation. A weakness of fertility as a measure of partnership is that it also includes children that may have resulted from short-term and less stable relationships (e.g., teen pregnancy). Moreover, the likelihood of such fertility outside of a stable relationship might be more relevant in some countries than others, and more importantly for us, differentially relevant across skill groups. With these caveats in mind, we construct as a variable for fertility, defined as a dummy variable that equals to 1 if a woman between 35 and 44 years of age lists an own child within her household.[25,26]

Appendix Figure 3 shows that a strong positive correlation exists between the skilled-unskilled gap in marriage rates and the skilled-unskilled gap in fertility, as just defined. Columns (7) to (9) of Table 6 replicate columns (1), (2) and (4) of Table 2, respectively, using the skilled-unskilled gap in fertility among women as the outcome. We see that more conservative gender norms reduce the likelihood that skilled women have children as compared to less skilled women. Columns (1) and (2) of Table 7, which replicate columns (3) and (4) of Table 3, also establish that our results in Table 3 continue to hold when we define the dependent variable of interest as the skilled-unskilled

[23] We cannot measure whether a person was cohabiting at some point in the past but is currently not.

[24] Unfortunately, we are not able to estimate a version of Table 3 for this outcome as this measure is not available in the earlier periods.

[25] In the EU-LFS, we derive this dummy from a variable that links children to their parents. The fertility measure for many European countries is available for 2005 and 2010 both from the LFS and the SILC. No fertility measures can be constructed for some EU countries (mostly Nordic countries) using the LFS. Given the much larger samples of the LFS our preferred specification uses the LFS measures when available and the SILC only when they are not.

[26] Given our age range of interest (35 to 44) we are likely to have measurement error for the older women in the range if their children have already left the household. Given education differences in age at first birth, the measurement error is likely to be larger for low educated women.

gap in fertility: the relationship between the skilled-unskilled gap in fertility and labor market opportunities for skilled women is more negative in high sexism countries.[27]

Excluding East Asia. Another concern with our analysis is the extent to which our key results are solely driven by the East Asian experience. The findings in columns (4) to (6) of Table 6 already establish that our findings in Table 2 are not driven by the East Asian countries as we do not measure cohabitation for these countries and hence they are excluded from the robustness analysis in those columns of Table 6. Columns (3) and (4) of Table 7 also establish that our results in Table 3 are not entirely driven by East Asia. There, we replicate columns (3) and (4) of Table 3 excluding the East Asian countries. We find that qualitatively similar patterns hold: the relationship between the skilled-unskilled gap in marriage rate and women's labor market opportunities becomes more negative as the sexism level in a country increases.

Alternative Measure of Gender Norms. As discussed above, we also constructed an alternative measure of the strength of gender norms within a country based on the answer to the following question in the ISSP: "A man's job is to earn money; a woman's to look after the home." While this question might arguably be more directly related to "α" in our theoretical model, we do not use this variable in our main analysis as answers to this question are available for a smaller set of countries.

In columns (10) to (12) of Table 6, and (5) and (6) of Table 7, we confirm that our core findings hold using this alternative measure. The skilled-unskilled gap in marriage rate is larger in more sexist countries, and more sexist countries experience a smaller decrease in the skilled-unskilled gap in marriage rate when labor market opportunities for skilled women increase. We note, however, that the estimated coefficient in column (6) of Table 7 is not statistically significant.

4.3 Evidence from US States

While the main goal of this paper is to provide a theoretical explanation (and empirical tests) for the divergent experiences over time across developed countries in the marriage market penalty for skilled women, the empirical tests that we have presented in the previous section can also be performed on a panel of US states. The first advantage of focusing on the US context is that there is considerably less unobserved heterogeneity across US states than across countries. However, this also comes at cost. Compared to the cross-country analysis, there is less variation across states in

[27]The variable linking mothers to children is included in the EU-LFS only starting in 1998, thus we do not have fertility measures for Europe in 1995. We do for other countries, so the panel covers 1995-2010 but has fewer observations in total than in Table 3.

the US than across developed countries in the key variable that drives our theory, the strength of gender norms. The second advantage of focusing on the US context is data availability: we can construct longer panels of all the key state-level variables that are required for our analysis and provide a cleaner mapping of labor market conditions at the time when individuals are making their marriage decisions. By combining 1970 to 2000 US Census and American Community Survey (2008 to 2011) data, our US state panel covers four decades, from 1970 to 2010.

We use the 1973 to 2014 waves of the General Social Survey (GSS) to measure gender norms at the state level in the US. We rely on the answers to the question "It is better if the man is the achiever outside the home and the woman takes care of home and family." Respondents indicate their agreement to this statement on a four-point scale - agree strongly, agree, disagree and disagree strongly. We code the response "agree strongly" and "agree" as indicating a greater degree of gender conservatism. Since the GSS did not include the "jobs scarce" question from the IVS that we used in the cross-country analysis, we focus on this specific question as it was most comparable to the ISSP question that was also used in the cross-country analysis. Moreover, this question appears closest in spirit to proxying for the gender roles that are central in the model, which is based on the belief that some may regard a woman's primary role to be in the home.

To combine the individual responses from different GSS waves into a single state-specific measure, we first regress the individual-level responses to the question on a full set of year dummies. Throughout, we use responses for all individuals age 18 and older. Next, we use the residuals from this regression to create a measure of the average gender conservatism in a particular state, which is simply the mean across all years of the residual individual-level response in a given state. For ease of interpretation, for the main analysis, we standardized the state-level sexism measure to have a mean of zero and a standard deviation of one in the full sample of states. We are able to construct a state level measure of sexism based on answers to this statement for 45 states, including DC.[28]

Summary statistics for the key labor, marriage market and educational variables for US states are presented in Appendix Table 5. As in the cross-country data, we measure outcomes for women (and men) who are between 35 and 44 years old in a given year. As we already noted (for a shorter time period) in the cross-country data, the deficit that skilled women used to experience in the marriage market in the US has been shrinking over time. While men used to complete college at a much higher rate than women in 1970, this male advantage had disappeared by 2000 and by 2010, a greater share of women than men between 35 and 44 years of age had completed a college degree.

Figure 8 replicates Figure 1 for the US state-level analysis. We plot the difference in marriage rates between skilled and unskilled women in 2010 (y-axis) against the difference in marriage rates between skilled and unskilled men in the same year (x-axis). As of 2010, in all states except Missouri

[28]The six states that were not included in earlier waves of the GSS are Hawaii, Idaho, Maine, Nebraska, Nevada and New Mexico. See Appendix Table 4 for the classification of states into three sexism groups.

and Wyoming, skilled men marry at a higher rate than unskilled ones. While we observe no deficit for skilled women in the marriage market in 2010 in the US at the aggregate level (Appendix Table 5), this statistic hides substantial variation across states, which we propose to explain through the lens of our theory. We also see that the majority of the data is below the 45-degree line. In other words, in the majority of states, skilled women experience a greater deficit in the marriage market compared to skilled men. In the figure, DC stands out as a clear outlier in the size of the marriage gaps for both women and men. For our subsequent analysis, we present results that include DC; however, we have also estimated specifications where we omit DC and the results are very similar.[29]

Figure 9 shows that the variation across states in the "penalty" skilled women experience in the marriage market correlates with the strength of gender norms at the state-level. As in the cross-country analysis, we classify states into three groups according to their degree of sexism – high, medium and low. The skilled-unskilled gap in marriage is more likely to be positive in low-sexism states and more likely to be negative in high-sexism states. There also appears to be a negative relationship between the high-skilled marriage rate gap among men and state-level gender norms, even though, as already discussed above, all but two states are characterized by a higher marriage rate for skilled men compared to unskilled men.

Table 8 confirms the correlation in Figure 9 in a regression framework which allows us to assess the sensitivity of the correlation to the inclusion of additional controls. The sample in Table 8 is restricted to the 45 states for which we can construct the sexism measure. The univariate correlation in column (1) between the skilled-unskilled gap in marriage rates in 2010 and the high-sexism dummy is, if anything, strengthened when control for a quadratic in high skilled female and male wages, as well as other state-specific controls such as the share of males with tertiary education, the female and male skill premium, and the female and male labor force participation skill premium (column (2)). Note that, unlike the cross-country data, the labor market controls are measured among individuals age 25 to 54 in the preceding decade. This ensures that the labor market controls proxy for the relevant labor market conditions that individuals are likely to face when making their marriage decisions (when they are age 25 to 34).[30] Column (3) shows that the relationship between the marriage gap by education among women and sexism is also robust to using the continuous measure of sexism.[31]

[29] These results are available on request.

[30] Due to data limitations, we were not able to construct the variables in the preceding time period in the cross-country analysis.

[31] Appendix Table 6 replicates the analysis in Table 5. Just like in the cross-country case, we find no systematic correlation across US states between total (e.g. husband plus wife) time spent on household production (even though gender norms correlate in the expected direction to the allocation of that time between husband and wife). Appendix Table 7 replicates the analysis in columns (1) to (3) of Table 6. Controlling for total time spent on household production (as a proxy for β) does not qualitatively alter the results in columns (1) to (4) of Table 8, even though the difference between mid-sexism states and low-sexism states is no longer statistically significant.

Columns (4) to (7) replicate column (2) for the following years, respectively: 2000, 1990, 1980 and 1970. What is most remarkable about this analysis of the variation in the relationship over time is that the relationship appears to be most pronounced in the last 3 decades (1990 to 2010) but is not present in the earlier decades (1980 and 1970). Recall that one prediction of our model is that educated women may only start experiencing a decline in their relative marriage rate compared to less educated ones when: a) they become less appealing to men who dislike having a working wife and b) they become more picky in the marriage market as their labor market opportunities and utility if they remain single increases. When labor market opportunities for skilled women are relatively low, we would not expect to see large differences in their relative marriage rates based on the sexism of the environment. Hence, the lack of a relationship between our two key variables in the earlier decades (1970 and 1980) could be rationalized through the lens of our model.

The two remaining columns of Table 8 replicate the analysis in columns (2) and (3) for men. While the visual inspection of the data in Figure 8 suggested, if anything, a stronger correlation between gender norms and the skilled-unskilled marriage gap for men, the multivariate evidence in the last two columns of Table 8 weakly suggest that this correlation is less robust. In particular, the relationship for men is not statistically significant when we use the continuous measure of sexism (column (9)).

Table 9 mirrors Table 3. In particular, we test for the finer prediction of our model: increases in labor market opportunities for skilled women are more likely to be associated with an increase in their relative marriage rate in less sexist states but a decrease in more sexist states. However, given the higher quality of the wage data in this context, we do not present results where we use state GDP per capita as an alternative proxy of labor market opportunities for skilled women. The analysis is performed on a 1970 to 2010 panel of US states. As indicated above, all the labor market variables (including high-skilled female wages) are measured in the prior decade to proxy for the labor market conditions that might have been most relevant at the time marriage market decisions were made. All regressions in Table 9 include state and year fixed effects and standard errors are clustered at the state level.

Recall that the main prediction of the model we want to test in this table is whether the relationship between the marriage market premium for skilled women and their labor market opportunities is mediated by the level of sexism in their place of residence. The evidence in Table 9 appears consistent with that in the country-level panel in Table 3. Specifically, we find that increases in labor market opportunities for skilled women are associated with an increase in the deficit they experience in the marriage market (compared to less skilled ones) in the more sexist states (.073-.089). The point estimate on log (female high skilled wage) is positive and statistically significant for the least sexist states (.073) and also positive but smaller in magnitude (.073-.047) for the mid-sexism states, as predicted by the theory. Column (2) shows that the patterns in column (1) are

robust to including a battery of time-varying state controls: share of males with tertiary education, the female and male skill premium, the female and male labor force participation premium, and a quadratic in female and male high skilled wages. Finally, column (3) shows that the patterns in column (2) are robust to using the continuous sexism measure: increased labor market opportunities are associated with relatively fewer marriages for skilled women as the sexism of the state increases.

Columns (4) and (5) replicate columns (3) and (4), respectively, for males. We also observe that skilled men's relative marriage rates are relatively more negatively affected by higher labor market opportunities for women in more sexist states; however, the relationship is not statistically robust to using the continuous sexism measure.

Tables 10 and 11 test the predictions of our model for women's educational outcomes. Table 10 replicates in the US states context the analysis of Table 4 where we observed a negative correlation between sexism and women's educational attainment. The dependent variable in Table 10 is the share of women in a state between 35 and 44 years of age that have completed a college degree. Each regression is a cross-section of states, at different points in time. All the regressions control for the share of men with a college degree in each state. Columns (1) to (4) present the results for the cohort of women aged 35 to 44 in 2010. In column (1), we observe that the share of women having completed a college degree is highest in low sexism states. Column (2) shows that the coefficient on the dummy for high sexism states is slightly more negative when we include the same set of state-level controls as in the previous table (a quadratic in male and female high skill wages, skill premium, labor force participation skill premium). The negative relationship between sexism and women's educational attainment is also robust to using the continuous sexism measure (column (3)).

As we had observed in the cross-country data, in column (4), we show that there is a positive relationship between the share of women having completed a college degree in a state and the difference in marriage rate between high and low skilled women, consistent with the mechanism articulated in the theory. However, we note here that, unlike in the cross-country analysis, this relationship is not statistically significant.

The remaining columns of Table 10 replicate column (2) for the other time periods: 2000, 1990, 1980, and 1970. While sometimes imprecise, the estimates suggest a negative correlation between sexism and female educational attainment from 1980 onwards. Sexism does not appear to relate to women's educational outcomes for the 1970 cohort (e.g. women that were between 35 and 44 years old in 1970). Again, our theory provides a way to rationalize this pattern. The explanation we propose for why norms may lead to worse female educational outcomes in more sexist states is that women may actively choose to be less educated in those states because they anticipate having to pay a cost in the marriage market for the increased labor market opportunities that are available to skilled individuals. Assuming that in the earlier period, education is less likely to translate into

greater labor market opportunities, then the relative sexism of the place is less likely to matter. Nonetheless, we can only conjecture as to whether this is indeed the correct explanation.

Table 11 fully leverages on the longer time series nature of the state-level data to ask whether the ambient level of sexism in a state mediates the sensitivity of women's education decision to their improving labor market opportunities. The dependent variable in all regressions is the share of women that are between 35 and 44 years old in a given year (2010, 2000, 1990, 1980 or 1970) and have completed a college degree. All regressions in Table 11 control for state and year fixed effects, as well for the share of males in the same age bracket that have completed a college education. Standard errors are clustered at the state level. Also, as in earlier tables, labor market opportunities for skilled women (log(female high skilled wage)) are measured in the prior decade to reflect as much as possible on the conditions that were relevant when the 35 to 44 year old women were making their education decision.

Column (1) shows that a given increase in labor market opportunities for skilled women increase their schooling achievement more in the least sexist states than in the more sexist states. While statistically insignificant, the estimated coefficient on the interaction term between labor market opportunities and mid-level sexism is of the expected sign (negative) and magnitude (e.g. smaller than the coefficient on the interaction term between labor market opportunities and high level sexism). When we include controls for a quadratic term in high-skilled female and male wages as well as the usual set of state-time varying controls (columns (2)), the results become weaker. Only in the specification where we use the continuous sexism measure (column (3)) do we find statistical significance at the 10 percent level.

Overall though, the evidence in Table 11 suggests that women's educational choices in the face of improving labor market opportunities are affected by the sexism of the environment that they face. This, combined with the panel evidence in Table 9 regarding marriage market outcomes, is consistent with the predictions of our theoretical model. Improvement in labor market opportunities for skilled women may translate into fewer women deciding to skill up via education in more sexist states as the benefits that women can reap in the labor market due to furthering their education are counteracted by worse prospects in the marriage market for skilled women in those states.

In summary, the state-level evidence appears largely consistent with the findings from the cross-country analysis. The fact that similar patterns are observed within a single country context is reassuring and suggests that the cross-country patterns are unlikely to be driven entirely by unobserved heterogeneity.

5 Conclusion

We develop a simple theoretical model that helps to rationalize the relatively lower marriage rate of educated women through the lens of gender identity norms. The model can also provide an explanation as to why educated women's relative deficit in the marriage market might react in opposite directions to improvements in their labor market opportunities in more vs. less gender conservative societies, and hence why women's educational choices in the face of growing labor market opportunities might depend on the strength of these gender norms as they balance labor and marriage markets considerations. We verify some of the key predictions of our model both in a panel of developed nations as well as a panel of US states.

The model we propose obviously abstracts from various relevant factors, and could be enriched by future work. In particular, to emphasize how better labor market opportunities for skilled women interact with slower moving social norms, we have taken the extreme perspective of fixed and exogenous gender role attitudes. A richer model would account for the fact that, even if slowly, gender role attitudes have been converging towards less conservative views, and would endogenize this process of change. A richer model might also allow for the quality of the household public good (such as raising a child) to differ based on the educational attainment of the wife.

Our analysis has implications for the expected long run trend of what is today a troubling phenomenon in many gender conservative countries, in particular East Asia and Southern Europe: the increasing singleness rate of college educated women. Given that non-traditional family structures and out-of-wedlock births continue to be quite rare in these societies (particularly East Asia), this "flight from marriage" among highly educated women is also likely to translate in a decline in fertility for this educational group, only reinforcing the already low fertility rate in this part of the world. The fact that highly skilled women are disproportionately foregoing childbearing could also result in lower social returns to education in these societies; and may further slow down the dynamics of adjustment of gender norms to the new labor market reality if only children of educated and/or working women develop more liberal gender attitudes (FFO, 2004).

Even if gender role attitudes do not change in East Asia, or change very slowly, our analysis suggests that further improvements in the labor market opportunities for skilled women should ultimately result in improving their relative attractiveness in the marriage market. While predicting when this will happen is beyond the mainly qualitative nature of the exercise we have performed here, the fact that the marriage rate of educated women has caught up to (and in some cases surpassed) that of less educated women in more gender-equal societies should give East Asian and Southern European countries hope about the transitional nature of the phenomenon they are experiencing.

References

[1] Akerlof, George and Rachel Kranton. 2000. "Economics and Identity." *Quarterly Journal of Economics*. Vol 115(3): 715-753.

[2] Becker, Gary S. 1973. "A Theory of Marriage: Part I." *Journal of Political Economy*. Vol. 81(4): 813-46.

[3] Fernandez, Raquel, Alessandra Fogli and Claudia Olivetti. 2004. "Mothers and Sons: Preference Formation and Female Labor Force Dynamics." *Quarterly Journal of Economics*. Vol. 119(4): 1249-1299.

[4] Fernandez, Raquel, Alessandra Fogli and Claudia Olivetti. 2002. "Marrying Your Mom: Preference Transmission and Women's Labor and Education Choices." NBER Working Paper No. 9234 (September).

[5] Feyrer, James, Bruce Sacerdote and Ariel Dora Stern. 2008. "Will the Stork Returns to Japan? Understanding Fertility within Developed Nations." *Journal of Economic Perspectives*. Vol. 22(3): 3-22.

[6] Goldin, Claudia. 2006. "The Quiet Revolution that Transformed Women's Employment, Education, and Family." *American Economic Review*. Vol. 96(2): 1-21.

[7] Greenwood, Jeremy, Ananth Seshadri and Mehmet Yorukoglu. 2005. "Engines of Liberation." *Review of Economic Studies*. Vol. 72(1): 109-133.

[8] Hwang, Jisoo. 2015. "Housewife, "Gold Miss," and Equal: The Evolution of Educated Women's Role in Asia and the US." Working Paper.

[9] Isen, Adam and Betsey Stevenson. 2010. "Women's Education and Family Behavior: Trends in Marriage, Divorce and Fertility," John Shoven, eds., Demography and the Economy, University of Chicago Press.

[10] Jones, Gavin W. and Bina Gubhaju. 2009. "Trends in Marriage in the Low Fertility Countries of East and Southeast Asia." *Asian Population Studies*. Vol. 5(3): 297-265.

[11] Kawaguchi, Daiji and Soohyung Lee. 2014. "Brides for Sale: Cross-Border Marriages and Female Immigration." Working Paper.

[12] Lundberg, Shelly and Robert A. Pollak. 2003. "Efficiency in Marriage." *Review of Economics of the Household*. Vol 1(3): 153-67.

[13] Mazzocco, Maurizio. 2007. "Household Intertemporal Behavior: A Collective Characterization and a Test of Commitment." *Review of Economic Studies*, Vol. 74 (3): 857-895.

[14] OECD Family Database. 2012. http://www.oecd.org/social/family/SF_2_4_Share_births_outside_marriage.pdf

[15] Pew Research Center. 2010. "The Reversal of the College Marriage Gap." Available online at: http://www.pewresearch.org/2010/10/07/the-reversal-of-the-college-marriage-gap/

[16] Pew Research Center. 2014. "Record Share of Americans Have Never Married." Available online at: http://www.pewsocialtrends.org/2014/09/24/record-share-of-americans-have-never-married/

[17] Stevenson, Betsey and Justin Wolfers. 2007. "Marriage and Divorce: Changes and their Driving Forces." *Journal of Economic Perspectives*. Vol. 21(2): 27-52.

[18] The Economist. 2011. "The Flight from Marriage." Available online at: http://www.economist.com/node/21526329

Figure 1. Cross-country Variation in the Marriage Gap in 2010 by Gender

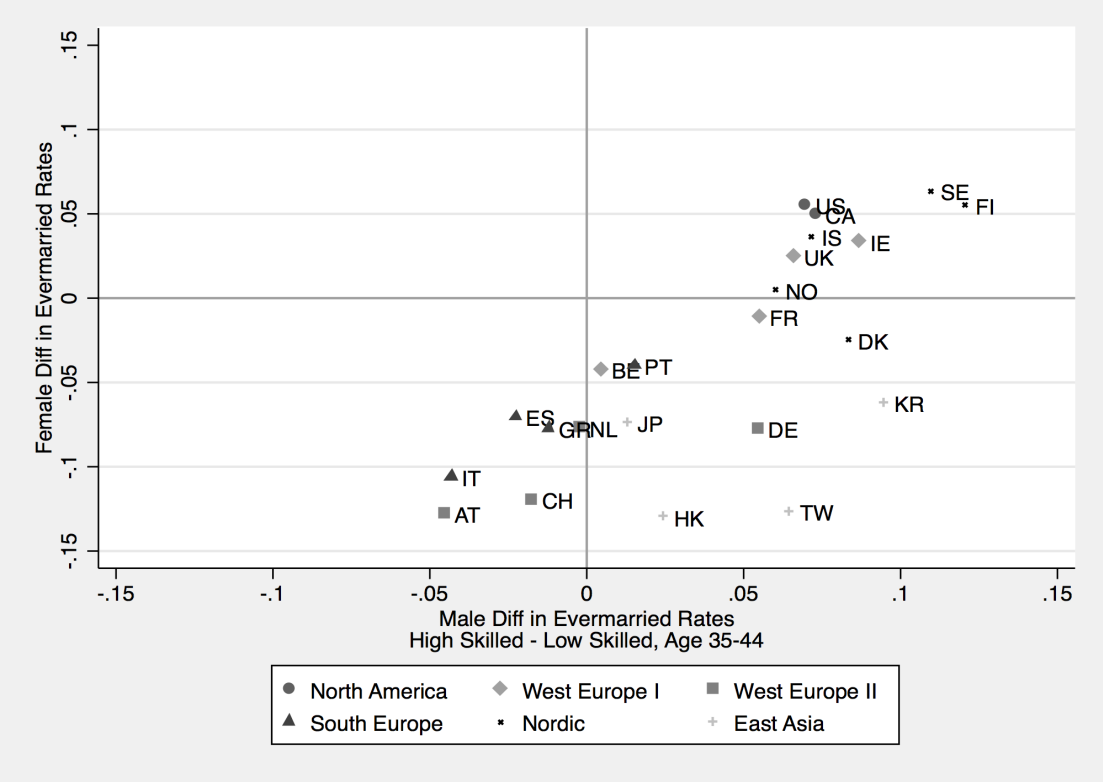

Note: The skilled-unskilled marriage rate gap among women (y-axis) is defined as the difference between the fraction of women with a tertiary education between 35 and 44 years old that were ever married and the fraction of women with less than a tertiary education in the same age range that were evermarried. Reported on the x-axis is the equivalent skilled-unskilled marriage rate gap among men. Refer to Appendix Table 2 for the data sources used to construct the marriage gaps for each country.

Figure 2A. Difference in Ever Married Rates (High Skilled - Low Skilled) from 1995 to 2010 by Country
Females, 35-44

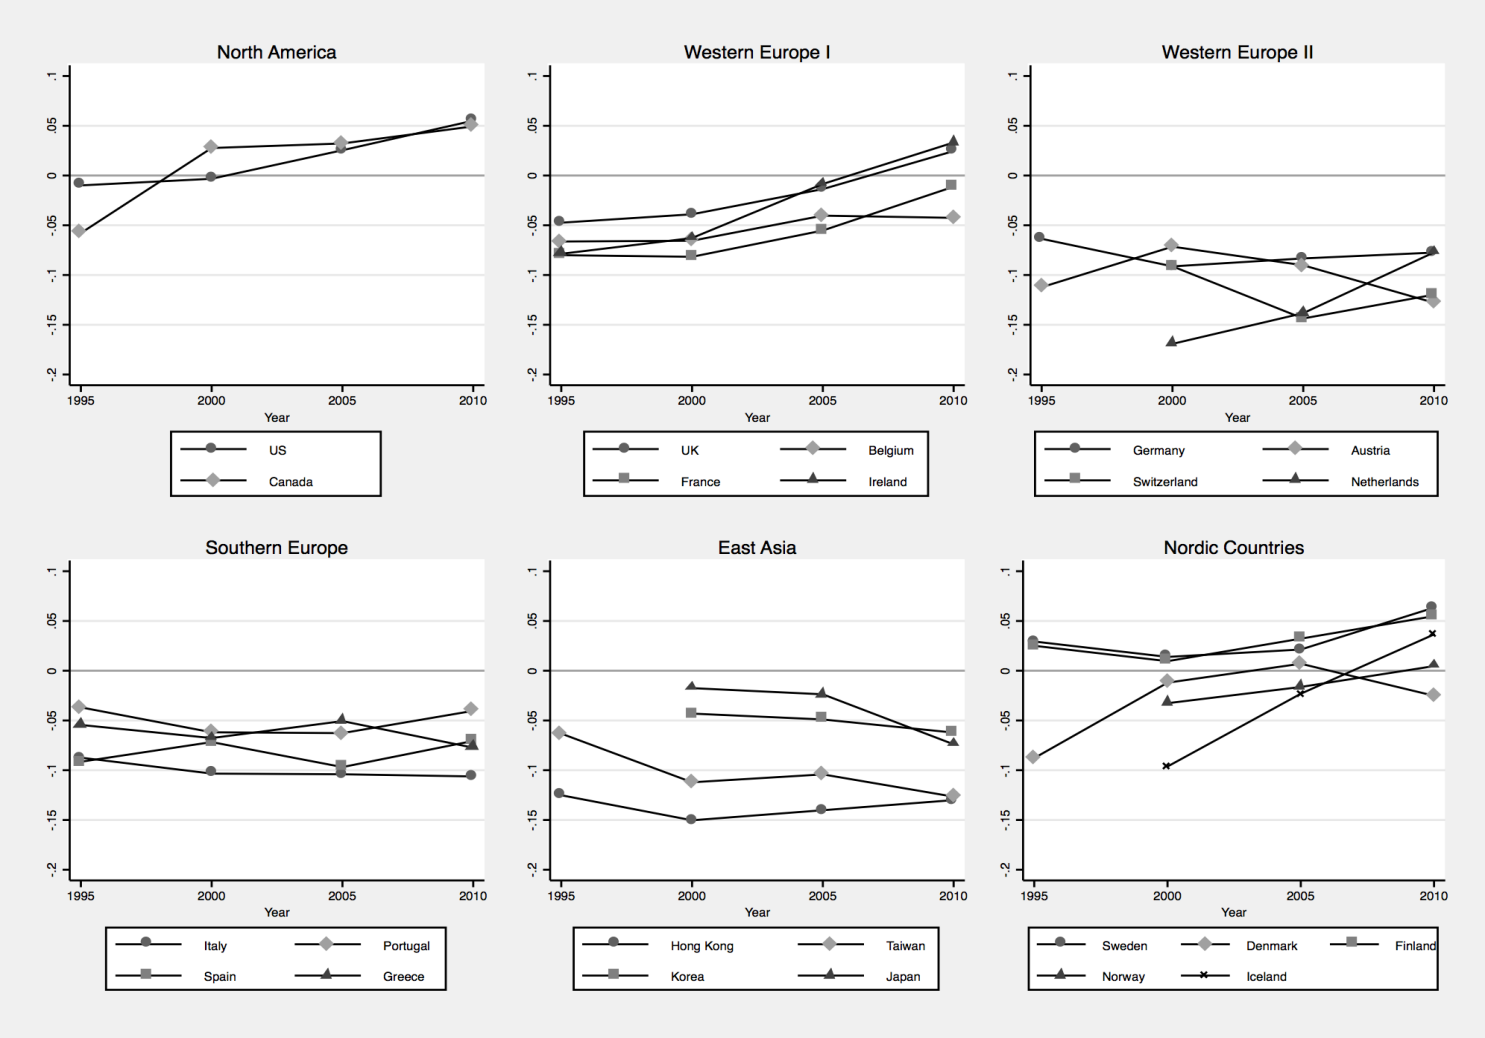

Note: See Appendix Table 2 for a description of the data sources. The figure shows the trends in the skilled-unskilled difference in evermarried rates for females aged 35-44 from 1995 to 2010 for each of the 23 countries in our sample.

Figure 2B. Difference in Ever Married Rates (High Skilled - Low Skilled) from 1995 to 2010 by Country
Males, 35-44

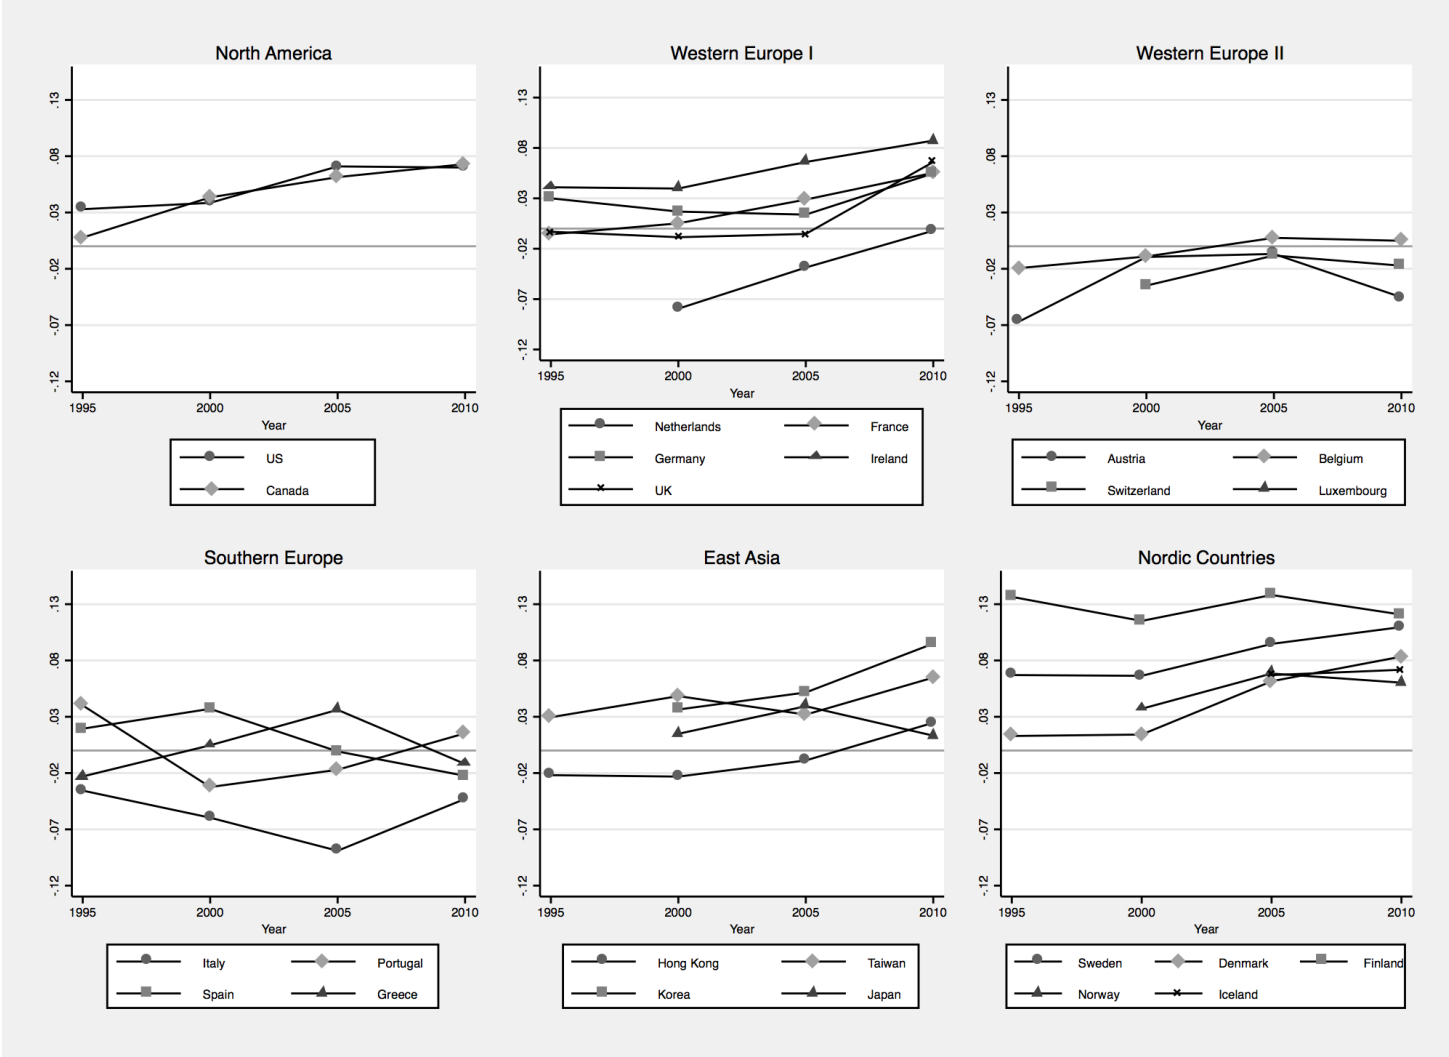

Note: See Appendix Table 2 for a description of the data sources. The figure shows the trends in the skilled-unskilled difference in evermarried rates for males aged 35-44 from 1995 to 2010 for each of the 23 countries in our sample.

Figure 3. Correlation between Marriage Gaps and Social Norms by Gender in 2010

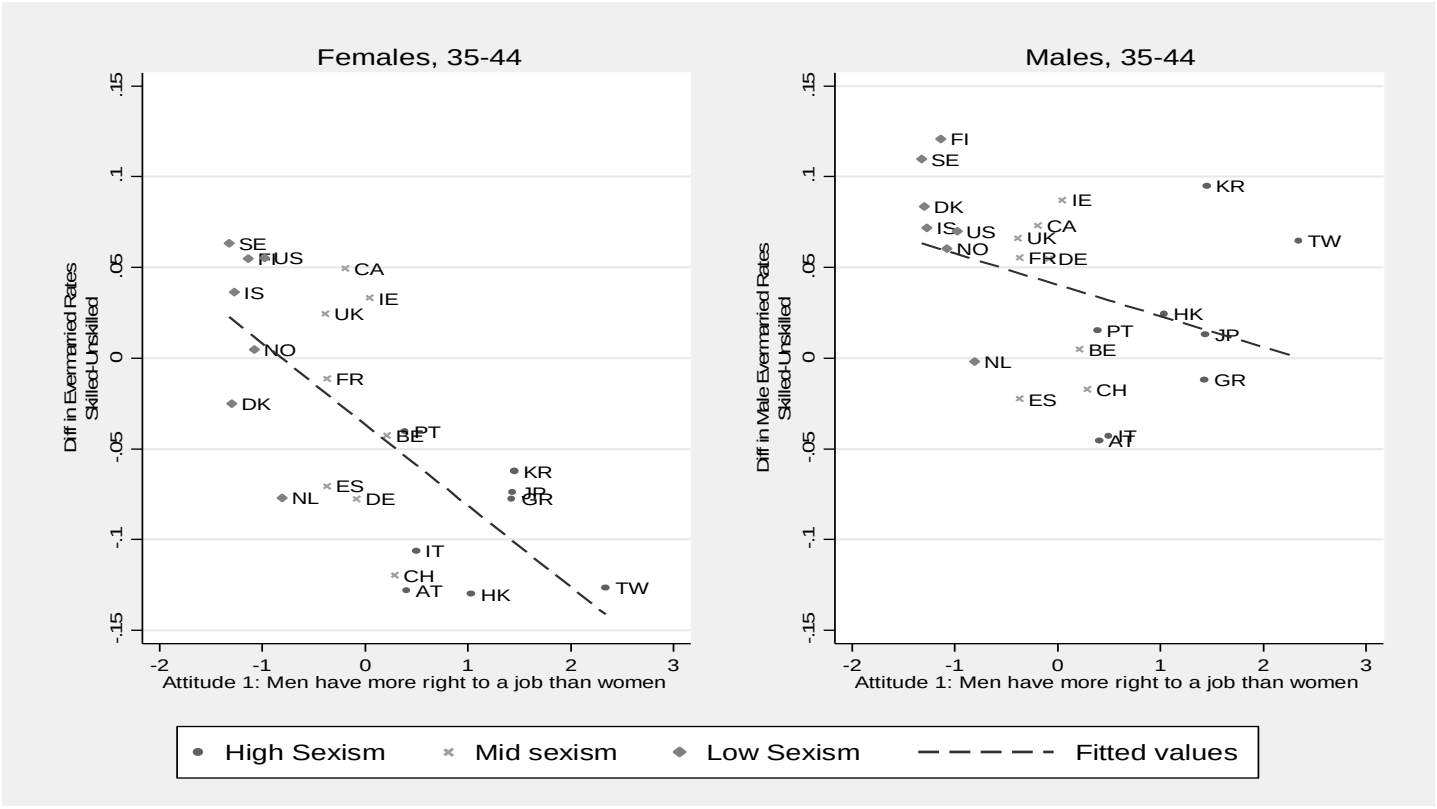

Notes: See Appendix Table 2 for a description of the data sources. The figure shows the cross-country relationship between the skilled-unskilled difference in evermarried rates for females (left panel) and males (right panel) aged 35 to 44 and the IVS-based measure of gender attitudes.

Figure 4: Husband's Utility Differential from Marrying a Working Woman as a Function of Wife's Wage

A. Different values of α

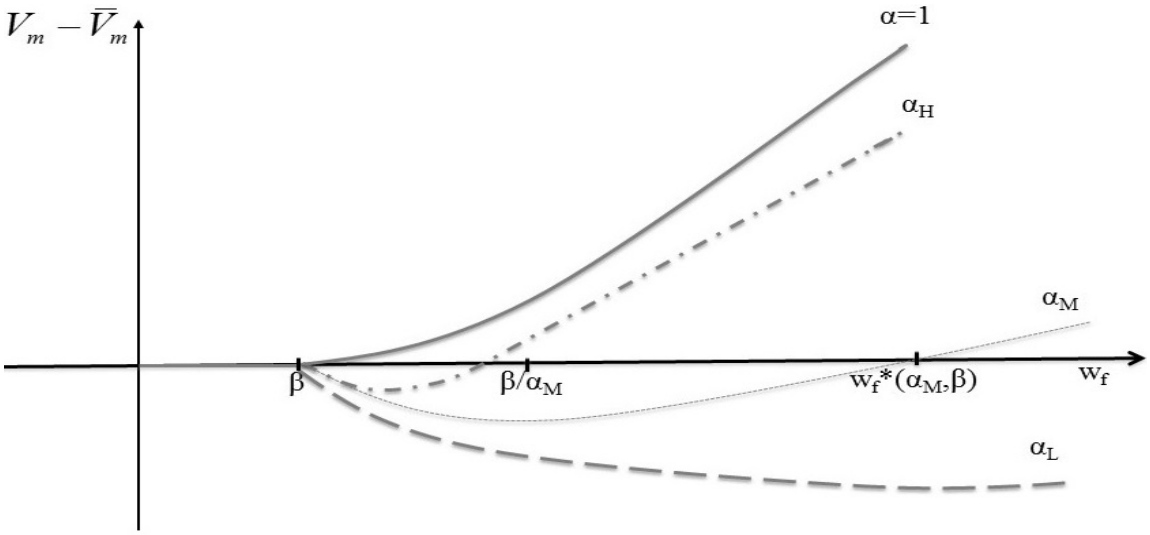

B. Different values of β

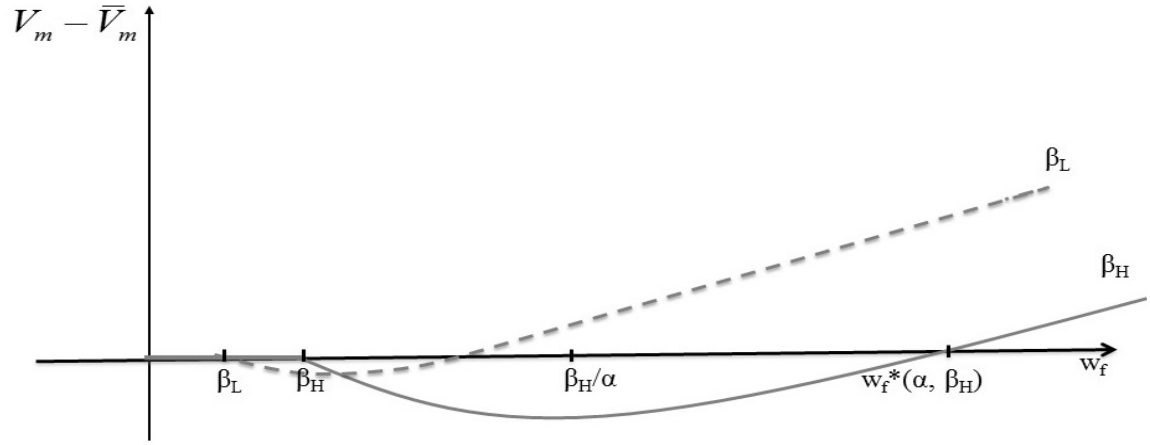

Figure 5: Marriage Gaps as a Function of Wife's Wage and Gender Norms

A. Female Marriage Gap

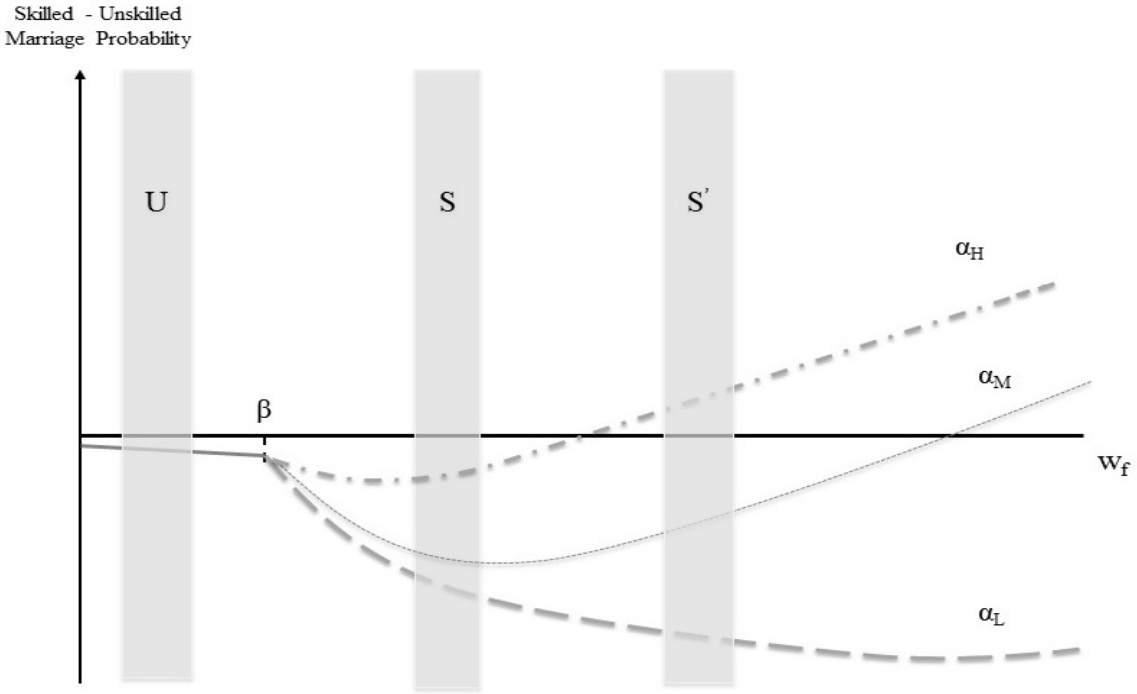

B. Male Marriage Gap

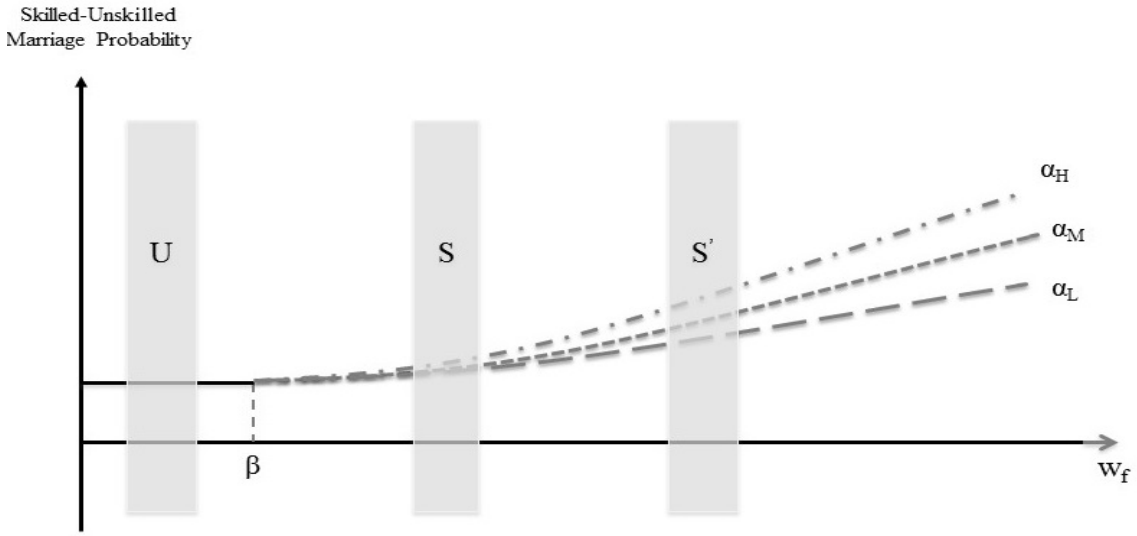

Figure 6. Gender Gap (Female-Male) in Education and Gender Norms, 2010

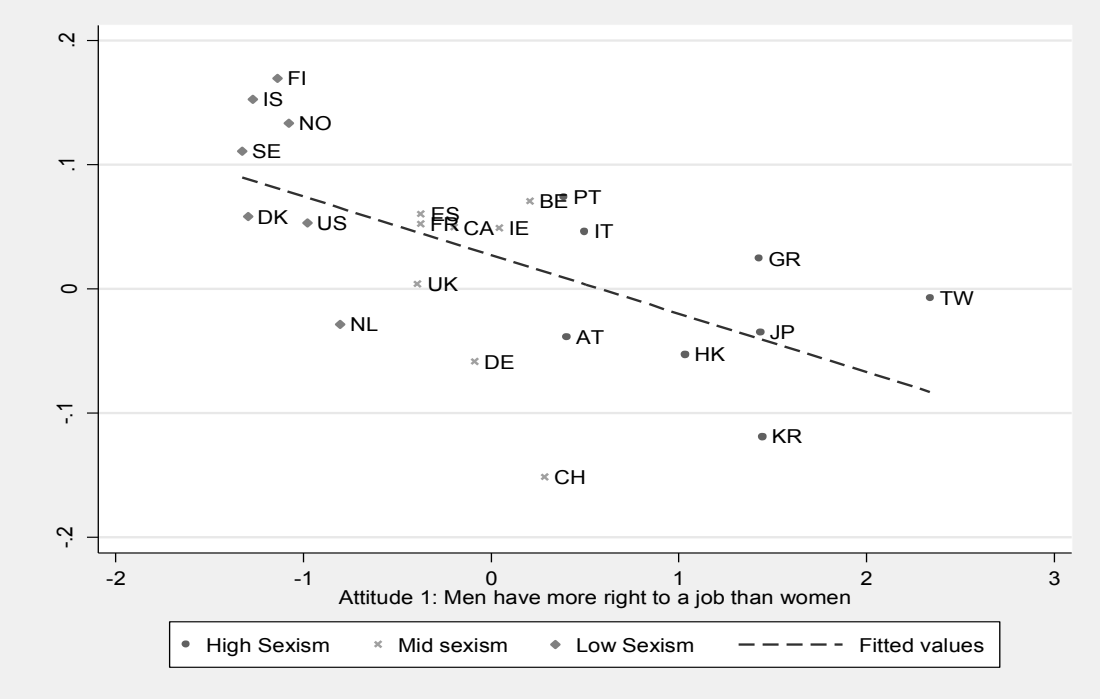

Note: See Appendix Table 2 for a description of the data sources. The figure shows the cross-country relationship between the gender gap in tertiary education (defined as the female-male difference in the share with a tertiary education among those aged 35 to 44) and the IVS-based measure of gender attitudes.

Figure 7. Time Spent on Household Production and Gender Norms

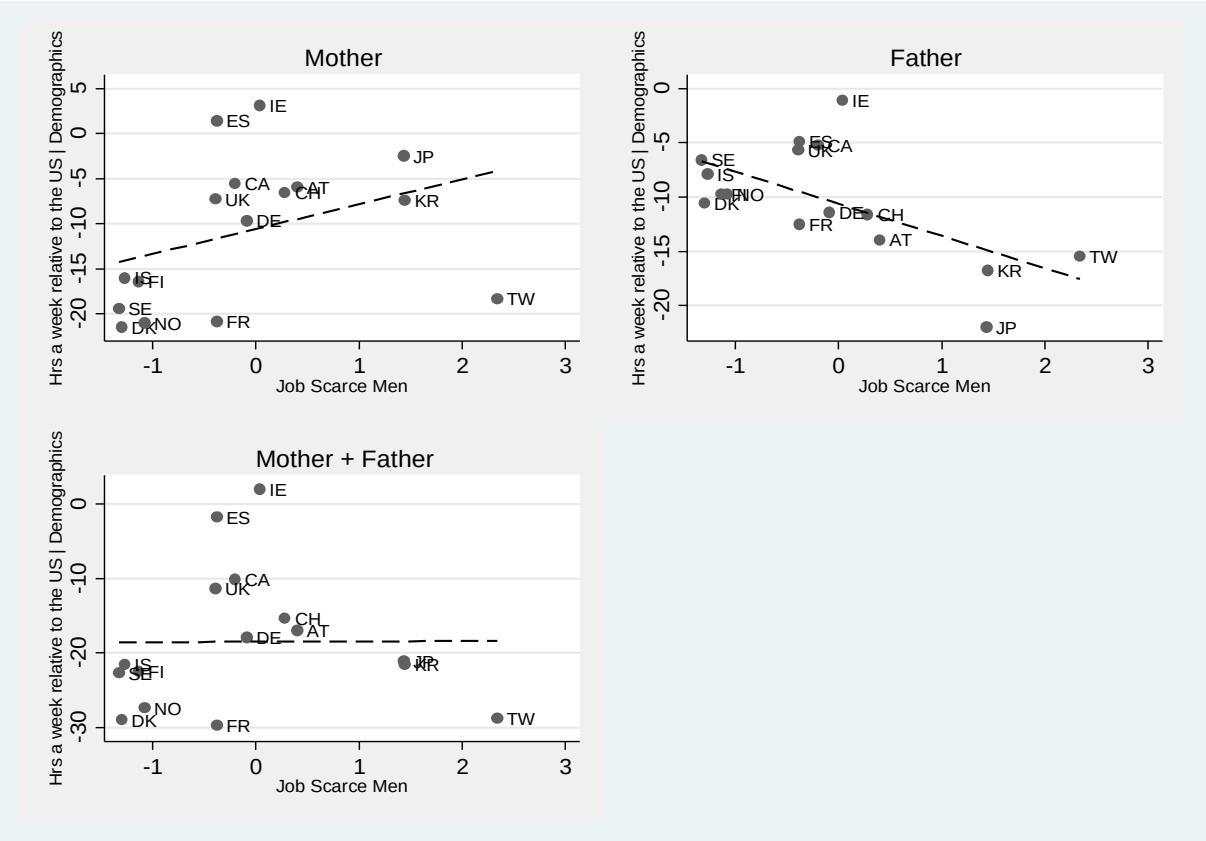

Note: The y axis is the country dummy coefficient of a regression of hours per week spent in total household production (family care plus household work) on country fixed effects, number of children, dummy for children under 5 years of age, household size, education and age of the respondent, and if she has a partner living in the household. Reference country is the US. Data comes from the 2012 ISSP, with the sample restricted to individuals aged 18-64 with at least one own child living in the household.

Figure 8. Cross-State Variation in the Marriage Gap in 2010 by Gender

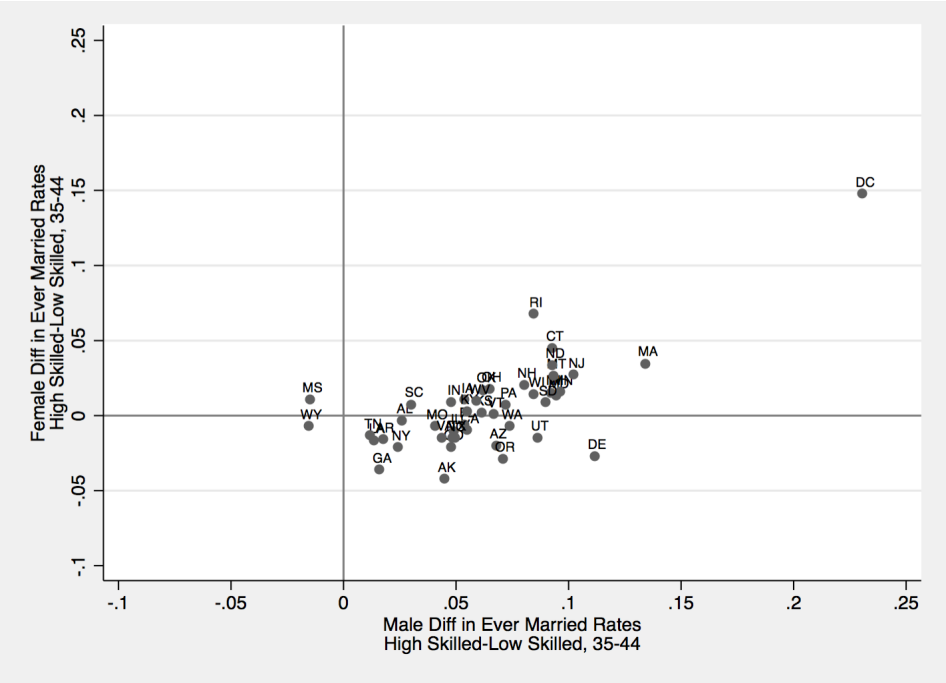

Note: The data is from the 1970 to 2000 US Census and the 2008 to 2011 ACS. The skilled-unskilled marrate rate gap among women (y-axis) is defined as the difference between the fraction of women with a college degree between 35 and 44 years that were ever married adn the fraction of women with less than a college degree in the same age range that were ever married. Reported on the x-axis is the equivalent skilled-unskilled marriage gap among men.

Figure 9. Correlation between Marriage Gaps and Social Norms by Gender in 2010 (US State Level Data)

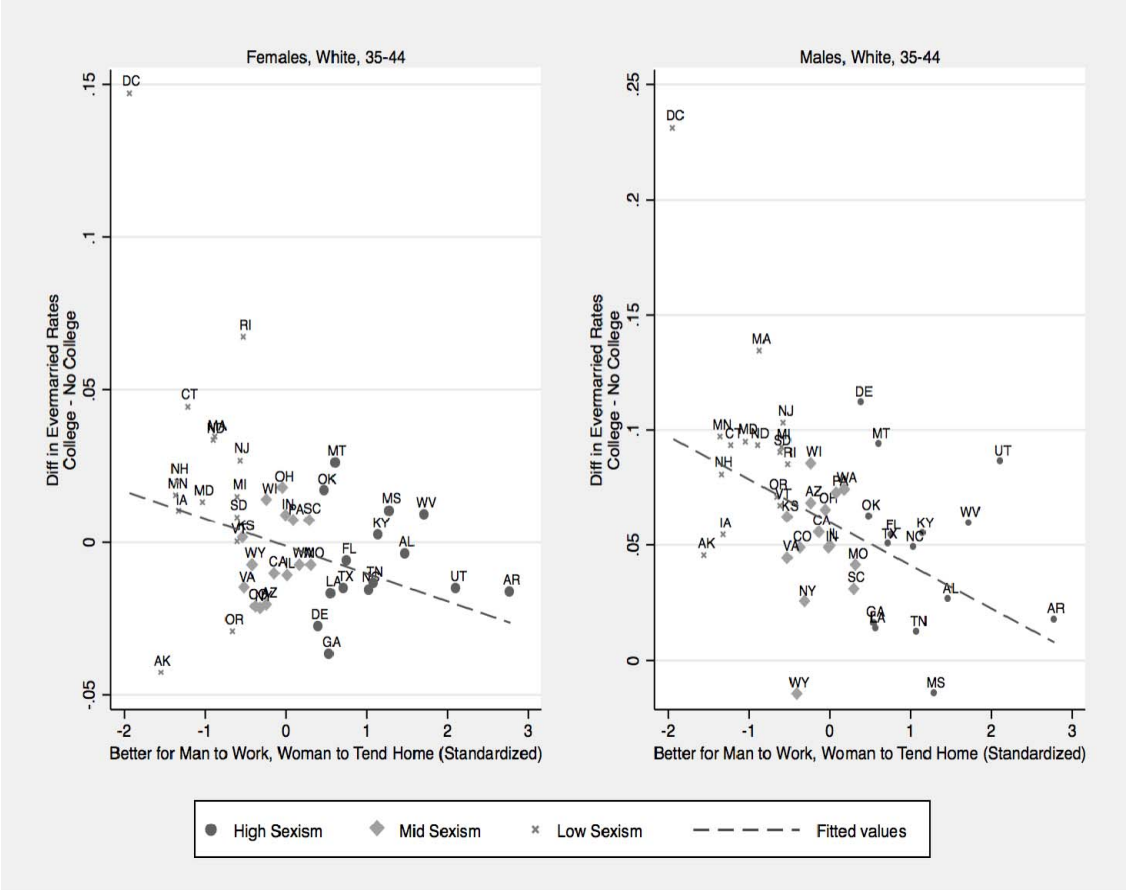

Note: The data is from the 1970 to 2000 US Census, the 2008 to 2011 ACS and the 1973-2014 GSS. The figure shows the cross-state relationship between the skilled-unskilled difference in ever married rates for females (left panel) and males (right panel) aged 35 to 44 and the state-level measure of gender attitudes.

Table 1. Gender-Related Attitudes Across Countries

	IVS: Men have more right to a job than women (% Agree)					ISSP: Man's job to earn money, woman's job to look after family (% Strongly Agree / Agree)				
				Males					Males	
	All	Males	Females	< Post-sec	Post-sec or more	All	Males	Females	< Post-sec	Post-sec or more
Full Sample	0.171	0.180	0.162	0.205	0.112	0.201	0.227	0.178	0.282	0.133
	(0.109)	(0.116)	(0.105)	(0.128)	(0.095)	(0.105)	(0.108)	(0.104)	(0.132)	(0.071)
No. of countries	23	23	23	23	23	20	20	20	20	20
East Asia	0.330	0.348	0.313	0.392	0.274	0.352	0.381	0.326	0.483	0.241
South Europe	0.225	0.249	0.201	0.276	0.118	0.291	0.307	0.277	0.347	0.068
West Europe I	0.161	0.158	0.163	0.183	0.096	0.183	0.217	0.150	0.276	0.112
West Europe II	0.171	0.174	0.167	0.197	0.095	0.220	0.242	0.199	0.304	0.170
North America	0.102	0.108	0.097	0.130	0.064	0.175	0.208	0.146	0.261	0.161
Nordic Countries	0.037	0.042	0.032	0.052	0.025	0.085	0.107	0.066	0.129	0.071
Correlation between average responses to IVS and ISSP questions: 0.88										

Note: The first row reports the sample mean of the average male and female response (standard deviation) across countries for each of the two statements from IVS (2006 to 2013) and ISSP (2002 and 2012). The sample mean of the average male response across countries is also reported separately by education level. In the remaining rows, we report the mean of the average response across five broad country groupings. East Asia includes Hong Kong, Japan, Korea, and Taiwan; South Europe includes Italy, Portugal, Greece, and Spain; West Europe I includes UK, France, Belgium, and Ireland; West Europe II includes Germany, Austria, Switzerland, and Netherlands; North America includes Canada and the US; Nordic countries include Sweden, Finland, Denmark, Norway, and Iceland.

Table 2. Correlation between the Marriage Gap and Social Norms among Individuals age 35-44

	Dep Var. Difference in Ever Married Rates (High Skilled - Low Skilled)									
	A. Females								B. Males	
	Year=2010					2005	2000	1995	Year=2010	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
High Sexism Dummy	-0.109***	-0.137***	-0.136**			-0.082*	-0.078*	0.010	-0.045	
	[0.023]	[0.036]	[0.048]			[0.042]	[0.040]	[0.090]	[0.031]	
Mid Sexism Dummy	-0.043	-0.042	-0.057			-0.029	-0.041	-0.013	-0.034	
	[0.029]	[0.045]	[0.056]			[0.041]	[0.032]	[0.066]	[0.029]	
Attitude measure: Men have more right to a job than women (mean 0, var 1)				-0.055***	-0.058***					-0.018
				[0.015]	[0.017]					[0.013]
Log(Female high-skill wage)		-1.706		-1.493		-6.247	-10.966	-11.672	0.070	-0.031
		[5.648]		[5.874]		[9.818]	[8.552]	[60.129]	[4.228]	[4.096]
Log(Female high-skill wage)^2		0.073		0.067		0.304	0.517	0.555	0.003	0.007
		[0.283]		[0.293]		[0.485]	[0.420]	[2.927]	[0.209]	[0.201]
Log(GDP per capita)			-1.729		-3.104					
			[3.078]		[3.808]					
Log(GDP per capita)^2			0.078		0.143					
			[0.144]		[0.177]					
Other Controls	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	23	23	23	23	23	21	20	14	23	23
R-squared	0.477	0.691	0.672	0.686	0.686	0.567	0.617	0.741	0.615	0.590

Note: The unit of observation is a country. Each column is a separate regression with the difference in evermarried rates between high-skilled and low-skilled for females (Panel A) and males (Panel B) for each of the years indicated as the dependent variable. The high (mid) sexism dummy refers to the top (middle) tertile of countries in terms of conservativeness of gender norms as measured using the IVS. The attitude measure is the continuous measure of gender norms from the IVS standardized to have mean 0 and standard deviation 1 in the sample of 23 countries. Log(Female high-skill wage) is the log average wage of full-time (working 35 or more hours per week) females age 25 to 54 with a tertiary education in each country in 2000 US dollars. Log(GDP per capita) is the GDP per capita based on purchasing power parity for each country. Other controls include the share of males in each country with tertiary education, the female and male skilled-unskilled wage premium, the female and male skilled-unskilled labor force participation premium, and a quadratic in female and male high skilled wages. Robust standard errors are reported in brackets.*** p<0.01, ** p<0.05, * p<0.1.

Table 3. Marriage Gap and the Interaction between Social Norms and Women's Labor Market Opportunities

	Dep Var. Difference in Ever Married Rates (High Skilled - Low Skilled)							
	A. Females						B. Males	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log(Female high-skill wage)	0.063	-1.878	-3.065		-3.254	-1.765	0.364	-0.127
	[0.039]	[2.952]	[3.459]		[3.001]	[2.986]	[1.921]	[1.987]
Log(Female high-skill wage)*High sexism	-0.120**	-0.190***					-0.048	
	[0.049]	[0.060]					[0.057]	
Log(Female high-skill wage)*Mid sexism	0.015	0.043					0.020	
	[0.054]	[0.067]					[0.039]	
Log(Female high-skill wage)*Sexism Index			-0.057*					-0.023
			[0.030]					[0.024]
Log(GDP per capita)				0.152	0.340**	0.120**		
				[0.106]	[0.123]	[0.053]		
Log(GDP per capita)*High sexism				-0.242**	-0.455***			
				[0.092]	[0.093]			
Log(GDP per capita)*Mid sexism				-0.014	-0.234**			
				[0.098]	[0.102]			
Log(GDP per capita)*Sexism Index						-0.169***		
						[0.025]		
Other Controls	No	Yes	Yes	No	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	80	78	78	85	78	78	78	78
R-squared	0.885	0.902	0.883	0.868	0.933	0.934	0.908	0.908

Note: The unit of observation is a country by year. The sample includes 23 countries across four time periods (1995, 2000, 2005, 2010). Each column is a separate regression with the difference in evermarried rates between high-skilled and low-skilled for females (Panel A) and males (Panel B) as the dependent variable. All regressions include year and country fixed effects. The number of observations in each column is fewer than 92 due to missing information on wages or marraige gaps for some countries in some years. The high (mid) sexism dummy refers to the top (middle) tertile of countries in terms of conservativeness of gender norms as measured using the IVS. The Sexism Index is the continuous measure of the conservatism of gender norms from the IVS standardized to have mean 0 and standard deviation 1 in the sample of 23 countries (higher values indicate countries with more conservative gender norms). Log(Female high-skill wage) is the log average wage of full-time (working 35 or more hours per week) females age 25 to 54 with a tertiary education in each country in 2000 US dollars. Log(GDP per capita) is the GDP per capita based on purchasing power parity for each country. Other controls include the share of males in each country with tertiary education, the female and male skilled-unskilled wage premium, the female and male skilled-unskilled labor force participation premium, and a quadratic in female and male high skilled wages. Robust standard errors clustered at the country level are reported in brackets.*** p<0.01, ** p<0.05, * p<0.1.

Table 4. Correlation between Female Tertiary Attainment and Social Norms Across Countries

	Dep Var. Share of Females with Tertiary Education						
	Year=2010				2005	2000	1995
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
High Sexism Dummy	-0.116***	-0.152***			-0.127***	-0.129***	-0.143**
	[0.029]	[0.037]			[0.023]	[0.029]	[0.052]
Mid Sexism Dummy	-0.082**	-0.082			-0.074**	-0.067**	-0.087**
	[0.036]	[0.048]			[0.031]	[0.028]	[0.036]
Attitude measure: Men have more right to a job than women (mean 0, var 1)			-0.061***				
			[0.016]				
Diff. in Ever Married Rates (high-skill - low-skill) - Females				0.806***			
				[0.195]			
Share of Males with Tertiary Education	0.627***	0.778***	0.949***	0.893***	0.778***	0.690***	0.383*
	[0.097]	[0.150]	[0.166]	[0.179]	[0.086]	[0.139]	[0.210]
Other Controls	No	Yes	Yes	Yes	No	No	No
Observations	23	23	23	23	21	20	16
R-squared	0.690	0.889	0.869	0.870	0.867	0.838	0.873

Note: The unit of observation is a country. Each column is a separate regression with the share of females with tertiary education in each time period (2010, 2005, 2000, 1995) as the dependent variable. The high (mid) sexism dummy refers to the top (middle) tertile of countries in terms of conservativeness of gender norms as measured using the IVS. The attitude measure is the continuous measure of gender norms from the IVS standardized to have mean 0 and standard deviation 1 in the sample of 23 countries. All regressions include a control for the share of men with tertiary education in each country. Other controls include a quadratic in female and male high-skilled wages, the female and male skilled-unskilled wage premium, and the female and male skilled-unskilled labor force participation premium. Robust standard errors are reported in brackets.*** p<0.01, ** p<0.05, * p<0.1

Table 5. Household Production and Gender Norms

	Hours a week spent in activity relative to the US | Demographics								
	Total Household Production			Care for Family Members			Household Work		
	(1) Female	(2) Male	(3) Total	(4) Female	(5) Male	(6) Total	(7) Female	(8) Male	(9) Total
Attitude measure: Men have more right to a job than women (mean 0, var 1)	2.756	-2.967***	0.043	0.602	-2.051***	-1.164	2.131	-0.931**	1.174
	[2.302]	[0.841]	[1.924]	[1.123]	[0.687]	[1.221]	[1.385]	[0.320]	[1.259]
Mean of Dep. Var	-10.86	-10.32	-18.50	-12.19	-8.61	-18.19	1.12	-1.70	-0.43
Observations	16	16	16	16	16	16	16	16	16
R-squared	0.133	0.383	0.000	0.015	0.319	0.030	0.259	0.350	0.086

Note: The outcome variable is the country dummy coefficient of a regression of hours a week spent in the activity (by the respondent or by her/his spouse) on country fixed effects, number of children, dummy for children under 5 years of age, household size, education and age of the respondent, and if she has a partner living in the household. Note that total household production is the sum of care for family members and household work. The excluded country dummy is for the US. Data comes from the 2012 ISSP, with the sample restricted to individuals aged 18-64 with at least one own child living in the household. The countries included in the regression are: AT, CA, CH, DK, DE, ES, FI, FR, IS, JP, KR, SE. We also have data for the US, which we use to construct the outcome variable. Robust standard errors are reported in brackets. *** p<0.01, ** p<0.05, * p<0.1.

Table 6. Correlation between the Female Marriage Gap and Social Norms among Females age 35-44: Robustness Tests

	Dep Var. Difference in Outcome (High Skilled - Low Skilled) in 2010 - Females											
	A. Evermarried Rates, includes controls for household production			B. Ever Married or Currently Cohabitating			C. Own Child at Home			D. Evermarried Rates, alternative measure of social norms		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
I. IVS attitude measure: Men have more right to a job than women												
High Sexism Dummy	-0.129***	-0.127***		-0.084***	-0.083***		-0.098***	-0.099***				
	[0.022]	[0.022]		[0.022]	[0.025]		[0.018]	[0.020]				
Mid Sexism Dummy	-0.075*	-0.062*		-0.007	-0.001		-0.057**	-0.050**				
	[0.037]	[0.032]		[0.030]	[0.030]		[0.021]	[0.022]				
Sexism Index (mean 0, sd 1)			-0.046***			-0.037*			-0.041**			
			[0.008]			[0.020]			[0.015]			
II. ISSP attitude measure: A man's job is to earn money; a woman's to look after the home												
High Sexism Dummy										-0.107***	-0.110***	
										[0.025]	[0.027]	
Mid Sexism Dummy										-0.028	-0.025	
										[0.028]	[0.031]	
Sexism Index (mean 0, sd 1)												-0.056***
												[0.013]
Log(Female high-skill wage)		2.924	1.183		2.768	-1.115		5.351**	2.424		-0.020	-4.155
		[4.783]	[4.258]		[6.089]	[6.840]		[2.208]	[3.504]		[3.247]	[6.424]
Log(Female high-skill wage)^2		-0.145	-0.060		-0.135	0.055		-0.263**	-0.121		-0.000	0.212
		[0.236]	[0.210]		[0.295]	[0.331]		[0.108]	[0.170]		[0.159]	[0.323]
Controls for Total Hhld Production (Mother + Father hrs per week)	Yes	Yes	Yes	No	No	No	No	No	No	No	No	No
Other Controls	No	No	No	No	No	No	No	Yes	Yes	No	Yes	Yes
Observations	16	16	16	16	16	16	22	22	22	20	20	20
R-squared	0.539	0.581	0.538	0.361	0.378	0.224	0.470	0.659	0.532	0.497	0.507	0.681

Note: The unit of observation is a country. The dependent variable is the difference in the outcome indicated in each panel for high-skilled and low-skilled females for females in 2010. Panel A examines the relationship between the high-skill minus low-skill difference in female evermarried rates and gender norms as measured using the IVS, including controls for total household production. Panel B uses evermarried + currently cohabiting rates as the outcome. The cohabitation data comes from the EU-SILC and is only available for the European countries. Panel C uses fertility, as proxied by the share of females with an own-child present at home, as the outcome. The fertility measure is not available for Iceland. Panel D examines the relationship bewteen the difference in female evermarried rates and gender norms using an alternative measure of gender norms based on the ISSP survey. The ISSP survey is not available for Italy, Greece, and Hong Kong. Given the reduced number of observations in Panels A and B, we do not include other controls in those panels. Other controls in Panels C and D include the share of males with tertiary education, the female and male skill premium, the female and male labor force participation premium, and a quadratic in female and male high skilled wages. See Table 5 for the construction of Total Hhld production. Robust standard errors are reported in brackets. *** p<0.01, ** p<0.05, * p<0.1.

Table 7. Female Marriage Gap and the Interaction between Social Norms and Women's Labor Market Opportunities: Robustness Tests

	Dep Var. Difference in Outcome (High Skilled - Low Skilled) - Females					
	A. Own Child at Home		B. Evermarried Rates, Excludes East Asia		C. Evermarried Rates, Alternative measure of social norms	
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Female high-skill wage)	0.288	0.499	-3.851	-4.688	-2.464	-4.508
	[6.637]	[6.877]	[2.847]	[3.457]	[5.011]	[5.553]
I. Attitude measure: Men have more right to a job than women (mean 0, var 1)						
Log(Female high-skill wage)*High sexism	-0.147*		-0.162**			
	[0.082]		[0.067]			
Log(Female high-skill wage)*Mid sexism	-0.074		0.018			
	[0.068]		[0.057]			
Log(Female high-skill wage)*Sexism Index		-0.094**		-0.073**		
		[0.043]		[0.029]		
II. Attitude measure: A man's job is to earn money; a woman's to look after the home.						
Log(Female high-skill wage)*High sexism					-0.186*	
					[0.105]	
Log(Female high-skill wage)*Mid sexism					0.077	
					[0.080]	
Log(Female high-skill wage)*Sexism Index						-0.060
						[0.047]
Other Controls	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	62	62	65	65	66	66
R-squared	0.930	0.932	0.931	0.924	0.891	0.876

Note: The unit of observation is a country by year (1995, 2000, 2005, 2010). The dependent variable is the difference in the outcome indicated in each panel for high-skilled and low-skilled females for females age 35 to 44. The regressions follow the same structure as those in Table 3. Panel A uses fertility, as proxied by the share of women with an own child at home, as the outcome. The number of observations in Panel A. is smaller than in Table 3 as the fertility variable is not available for European countries in the earlier years. Sexism groups in Panel B are redefined dropping East Asia. Panel C reports similar results to that in Table 3, using an alternative measure of social norms from the ISSP. The ISSP question is not available for Italy, Greece, and Hong Kong. Other controls include the share of males with tertiary education, the female and male skill premium, the female and male labor force participation premium, and a quadratic in female and male high skilled wages. Robust standard errors clustered at the country level are reported in brackets.*** p<0.01, ** p<0.05, * p<0.1.

Table 8. Correlation between the Marriage Penalty and Social Norms among Individuals age 35-44 across US States

	Dep Var. Difference in Ever Married Rates (High Skilled - Low Skilled)								
	A. Females							B. Males	
	Year=2010			2000	1990	1980	1970	Year=2010	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
High Sexism Dummy	-0.031**	-0.041***		-0.007	-0.015**	0.012	0.006	-0.011	
	[0.012]	[0.012]		[0.007]	[0.006]	[0.007]	[0.018]	[0.013]	
Mid Sexism Dummy	-0.029**	-0.023*		-0.011*	-0.007	0.001	0.006	-0.024**	
	[0.012]	[0.012]		[0.006]	[0.005]	[0.007]	[0.015]	[0.010]	
Attitude measure: Better if man is achiever outside home and women take care of home and family (mean 0, var 1)			-0.012***						0.000
			[0.004]						[0.007]
Log(Female high-skill wage)		-13.623**	-9.903	6.021**	-8.388*	-4.162**	-12.318**	-14.065	-11.914
		[5.733]	[6.624]	[2.606]	[4.479]	[1.726]	[5.314]	[9.504]	[10.497]
Log(Female high-skill wage)^2		0.633**	0.460	-0.290**	0.402*	0.201**	0.604**	0.671	0.572
		[0.269]	[0.311]	[0.124]	[0.218]	[0.085]	[0.263]	[0.445]	[0.492]
Other Controls	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	45	45	45	45	45	45	45	45	45
R-squared	0.215	0.635	0.573	0.787	0.613	0.674	0.460	0.742	0.702

Note: The unit of observation is a state (including DC). The data is from the 1970 to 2000 US Census and 2008 to 2011 ACS. The sample is restricted to whites. Attitudes are measured using the 1973 to 2014 General Social Survey and are not available for 5 states (Hawaii, Idaho, Maine, Nebraska, Nevada, and New Mexico). Each column is a separate regression with the difference in evermarried rates between high-skilled and low-skilled for females (Panel A) and males (Panel B) age 35 to 44 for each of the years indicated as the dependent variable. The high (mid) sexism dummy refers to the top (middle) tertile of states in terms of conservativeness of gender norms as measured using the GSS. The attitude measure is the continuous measure of gender norms from the GSS standardized to have mean 0 and standard deviation 1 in the sample of 45 states. Log(female high-skill wage) is the log average wage of full-time (working 35 or more hours per week) females age 25 to 54 with a college degree in the preceding decade (t-10) in each state. Other controls include the share of males with college education, the female and male skill premium, the female and male labor force participation skill premium, and a quadratic in female and male high skilled wages. The labor market controls are measured among individuals age 25 to 54 in year t-10. Robust standard errors are reported in brackets. *** p<0.01, ** p<0.05, * p<0.1.

Table 9. Marriage Gap and the Interaction between Sexism and Women's Labor Market Opportunities Across US States

	Dep Var. Difference in Ever Married Rates (High Skilled - Low Skilled)				
	A. Females			B. Males	
	(1)	(2)	(3)	(4)	(5)
Log(Female high-skill wage)	0.073**	-0.577	-0.894	-2.772**	-2.854**
	[0.036]	[1.720]	[1.699]	[1.248]	[1.309]
Log(Female high-skill wage)*High sexism	-0.089***	-0.082***		-0.032	
	[0.023]	[0.025]		[0.020]	
Log(Female high-skill wage)*Mid sexism	-0.047*	-0.040		-0.031*	
	[0.024]	[0.026]		[0.015]	
Log(Female high-skill wage)*Sexism Index			-0.022**		-0.008
			[0.010]		[0.011]
Other Controls	No	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes
Observations	225	225	225	225	225
R-squared	0.593	0.604	0.584	0.835	0.831

Note: The unit of observation is a state by decade (1970, 1980, 1990, 2000, and 2010). The data is from the 1970 to 2000 US Census and 2008 to 2011 ACS. The sample is restricted to whites. Attitudes are measured using the 1973 to 2014 General Social Survey and are not available for 5 states (Hawaii, Idaho, Maine, Nebraska, Nevada, and New Mexico). Each column is a separate regression with the difference in evermarried rates between high-skilled and low-skilled for females (Panel A) and males (Panel B) age 35 to 44 as the dependent variable. All regressions include year and state fixed effects. The high (mid) sexism dummy refers to the top (middle) tertile of states in terms of conservativeness of gender norms as measured using the GSS. The attitude measure is the continuous measure of gender norms from the GSS standardized to have mean 0 and standard deviation 1 in the sample of 45 states. Log(female high-skill wage) is the log average wage of full-time (working 35 or more hours per week) females age 25 to 54 with a college degree in the preceding decade (t-10) in each state. Other controls include the share of males with college education, the female and male skill premium, the female and male labor force participation skill premium, and a quadratic in female and male high skilled wages. The labor market controls are measured among individuals age 25 to 54 in year t-10. Robust standard errors clustered at the state level are reported in brackets. *** p<0.01, ** p<0.05, * p<0.1

Table 10. Correlation between Female College Attainment and Social Norms Across US States

	Dep Var. Share of Females with College Degree (Age 35 to 44)							
	Year=2010				2000	1990	1980	1970
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
High Sexism Dummy	-0.020**	-0.025**			-0.022*	-0.009	-0.016***	-0.008
	[0.010]	[0.012]			[0.013]	[0.011]	[0.004]	[0.006]
Mid Sexism Dummy	-0.014**	-0.013*			-0.013	-0.005	-0.007	-0.003
	[0.005]	[0.007]			[0.009]	[0.007]	[0.005]	[0.005]
Attitude measure: Better if man is achiever outside home and women take care of home and family (mean 0, var 1)			-0.018***					
			[0.006]					
Diff. in Ever Married Rates (high-skill - low-skill) - Females				0.146				
				[0.132]				
Share of males with college degree	0.986***	0.879***	0.869***	0.857***	0.813***	0.900***	0.712***	0.554***
	[0.034]	[0.075]	[0.057]	[0.082]	[0.105]	[0.060]	[0.036]	[0.052]
Other Controls	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	45	45	45	45	45	45	45	45
R-squared	0.975	0.982	0.987	0.980	0.974	0.980	0.989	0.935

Note: The unit of observation is a state. Each column is a separate regression with the share of females age 35 to 44 with a college degree in each time period as indicated in the columns as the dependent variable.The high (mid) sexism dummy refers to the top (middle) tertile of states in terms of conservativeness of gender norms as measured using the GSS. The attitude measure is the continuous measure of gender norms from the GSS standardized to have mean 0 and standard deviation 1 in the sample of 45 states. All regressions include a control for the share of men age 35 to 44 with a college degree in each country. Other controls include a quadratic in female and male high-skilled wages, the female and male skill premium, and the female and male labor force participation skill premium. The labor market controls are measured among individuals age 25 to 54 in year t-10. Robust standard errors reported in brackets. *** p<0.01, ** p<0.05, * p<0.1.

Table 11. Female College Attainment and the Interaction between Sexism and Women's Labor Market Opportunities Across US States

	Dep Var. Share of Females with College Education		
	(1)	(2)	(3)
Log(Female high-skill wage)	0.006	-1.934**	-1.637*
	[0.030]	[0.800]	[0.914]
Log(Female high-skill wage)*High sexism	-0.037*	-0.018	
	[0.020]	[0.016]	
Log(Female high-skill wage)*Mid sexism	-0.020	-0.017	
	[0.017]	[0.015]	
Log(Female high-skill wage)*Sexism Index			-0.015*
			[0.008]
Share of males with college education	1.053***	0.940***	0.910***
	[0.182]	[0.130]	[0.122]
Other Controls	No	Yes	Yes
Year FE	Yes	Yes	Yes
State FE	Yes	Yes	Yes
Observations	225	225	225
R-squared	0.987	0.989	0.990

Note: The unit of observation is a state by decade (1970, 1980, 1990, 2000, and 2010). The data is from the 1970 to 2000 US Census and 2008 to 2011 ACS. The sample is restricted to whites. Attitudes are measured using the 1973 to 2014 General Social Survey and are not available for 5 states (Hawaii, Idaho, Maine, Nebraska, Nevada, and New Mexico). Each column is a separate regression with the share of females age 35 to 44 with a college degre as the dependent variable. All regressions include year and state fixed effects. The high (mid) sexism dummy refers to the top (middle) tertile of states in terms of conservativeness of gender norms as measured using the GSS. The sexism index is the continuous measure of gender norms from the GSS standardized to have mean 0 and standard deviation 1 in the sample of 45 states. All regressions control for the share of males age 35 to 44 with college education. Other controls include the female and male skill premium, the female and male labor force participation skill premium, and a quadratic in female and male high skilled wages. The labor market controls are measured among individuals age 25 to 54 in year t-10. Robust standard errors clustered at the state level are reported in brackets. *** p<0.01, ** p<0.05, * p<0.1.

Appendix Figure 1: Convergence in Gender-Related Attitudes

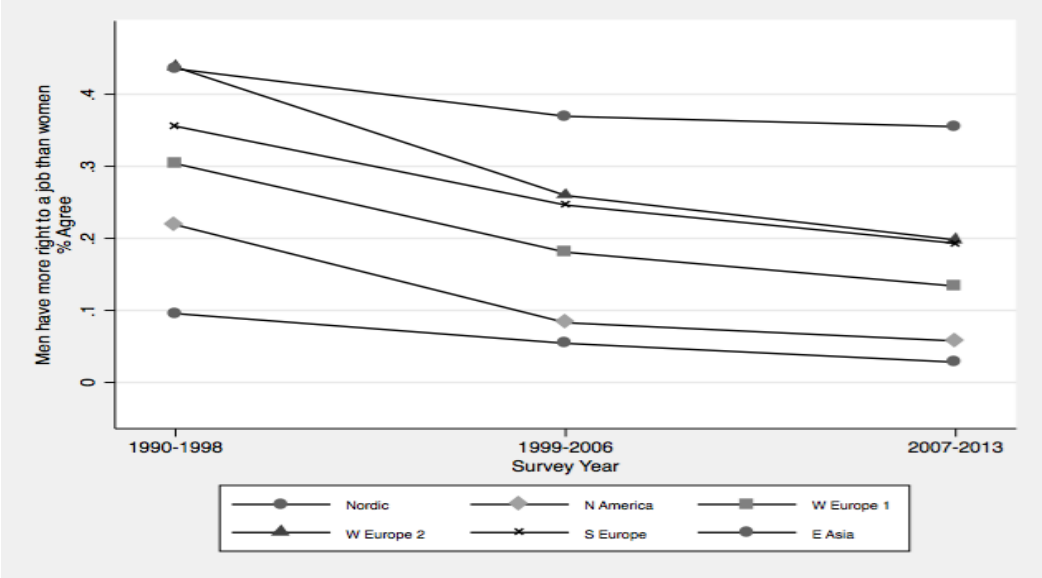

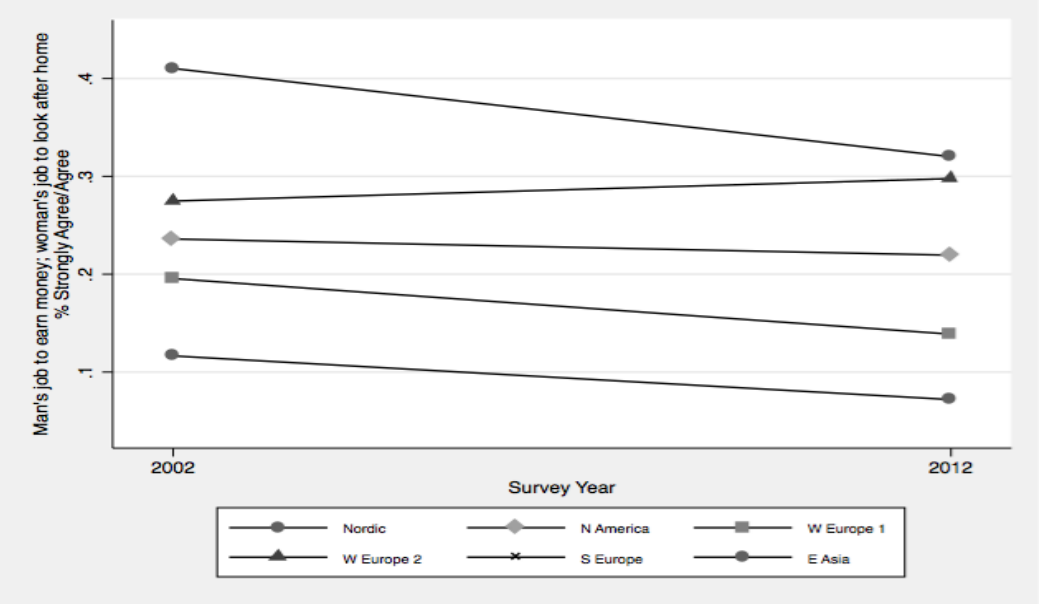

Note: The graphs plot the average of the mean responses of individuals age 18 and older for each country by country group and survey year. The top figure uses the gender-related question from the IVS while the bottom figure uses the question from the ISSP. The sample of countries is restricted to those with available data in each of the survey years.

Appendix Figure 2. Cohabitation Rates by Skill and Correlation between Marriage Gap Measures (2010)

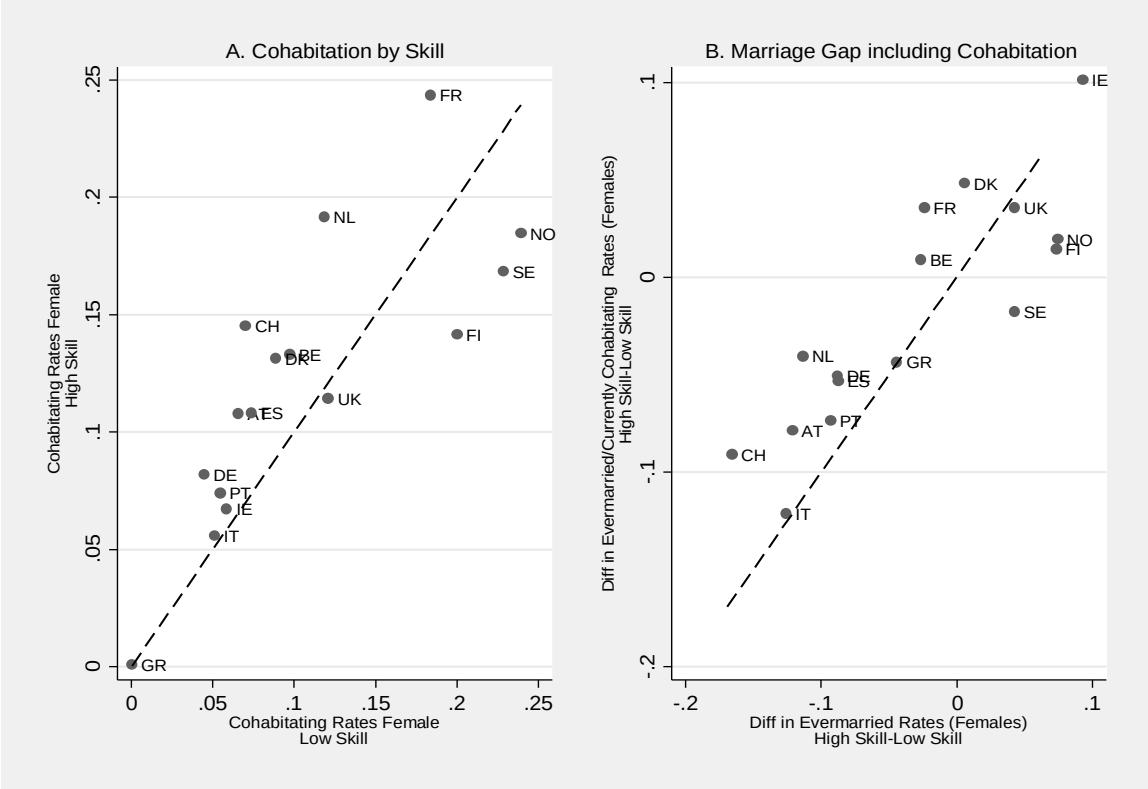

Note: The data is from the 2010-2011 EU-SILC. Panel A plots cohabitation rates of tertiary-educated females aged 35 to 44 against the cohabitation rates of non-tertiary educated females of the same age range. Panel B plots the skilled-unskilled gap in marriage+cohabitation rates against the skilled-unksilled gap in marriage rates among females aged 35 to 44.

Appendix Figure 3. Correlation between Marriage Gap and Fertility Gap for Females (2010)

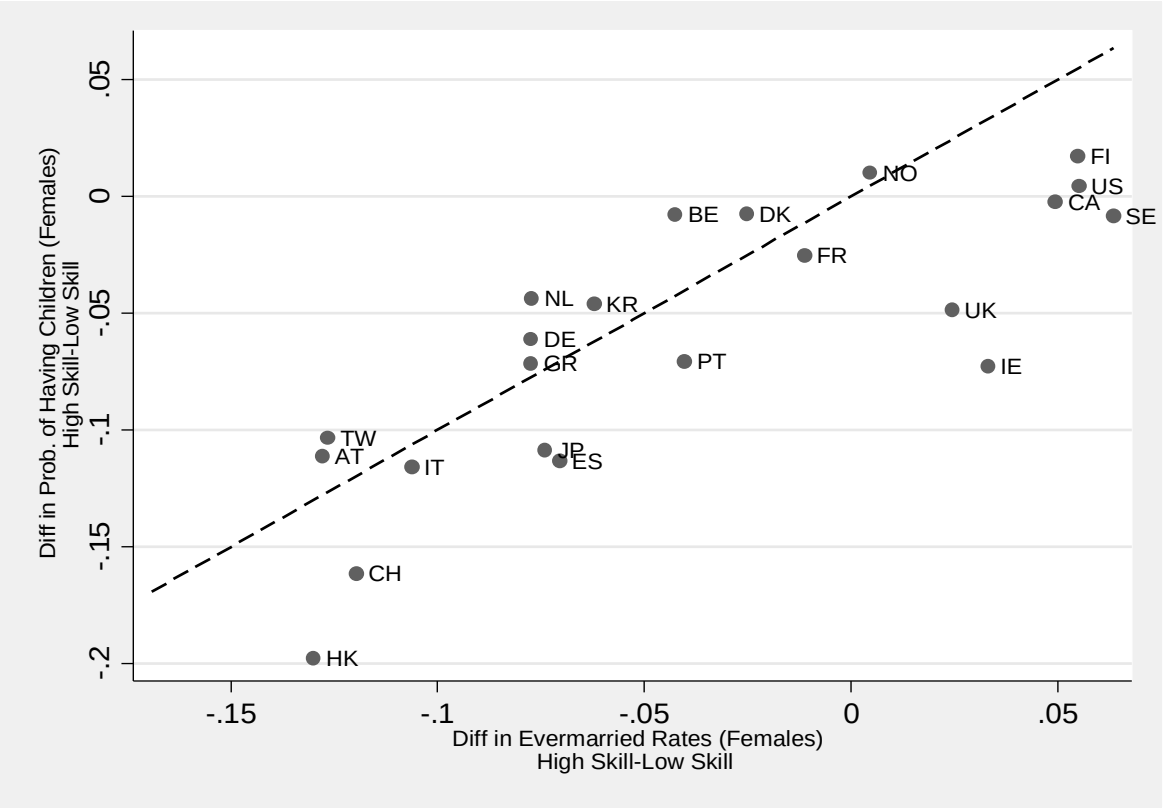

Note: See Appendix Table 2 for a description of the data sources. The figure shows the cross-country relationship between the skilled-unskilled fertility gap and the skilled-unskilled marriage gap for females aged 35 to 44 in 2010.

Appendix Table 1: Classification of Countries by Sexism Groups

			IVS			ISSP		
Country	Country code	Region	Survey years	% Agree: When jobs are scarce, men have more right to a job than women	Sexism Group (1: Low, 2: Mid, 3: High)	Survey years	% Agree: A man's job is to earn money; a woman's job is to look after the home and family.	Sexism Group (1: Low, 2: Mid, 3: High)
Sweden	SE	Nordic	2006, 2009, 2011	0.022	1	2002	0.076	1
Denmark	DK	Nordic	2008	0.023	1	2002, 2012	0.101	1
Iceland	IS	Nordic	2009	0.026	1	2012	0.067	1
Finland	FI	Nordic	2005, 2009	0.067	1	2002, 2012	0.108	1
Norway	NO	Nordic	2008	0.047	1	2002, 2012	0.074	1
USA	US	N. America	2006, 2011	0.061	1	2002, 2012	0.227	2
Netherlands	NL	W. Europe II	2006, 2008, 2012	0.096	1	2002	0.125	1
UK	UK	W. Europe I	2006, 2009	0.139	2	2002, 2012	0.183	2
France	FR	W. Europe I	2006, 2008	0.147	2	2002, 2012	0.166	2
Spain	ES	S. Europe	2007, 2008, 2011	0.153	2	2002, 2012	0.215	2
Canada	CA	N. America	2006	0.143	2	2012	0.123	1
Germany	DE	W. Europe II	2006, 2008, 2013	0.171	2	2002, 2012	0.184	2
Ireland	IE	W. Europe I	2008	0.169	2	2002, 2012	0.157	2
Belgium	BE	W. Europe I	2009	0.187	2	2002	0.246	2
Switzerland	CH	W. Europe II	2007, 2008	0.208	2	2002, 2012	0.242	3
Portugal	PT	S. Europe	2008	0.207	3	2002	0.336	3
Austria	AT	W. Europe II	2008	0.208	3	2002, 2012	0.328	3
Italy	IT	S. Europe	2005, 2009	0.219	3			
Hong Kong	HK	East Asia	2005, 2013	0.243	3			
Greece	GR	S. Europe	2008	0.321	3			
Japan	JP	East Asia	2005, 2010	0.306	3	2002, 2012	0.276	3
Korea	KR	East Asia	2005, 2010	0.344	3	2012	0.333	3
Taiwan	TW	East Asia	2006, 2012	0.428	3	2002, 2012	0.448	3

Note: The data is from the 2006 to 2013 Integrated Values Survey (IVS) and the 2002 and 2012 International Social Survey Program (ISSP). The gender-related questions are based on the responses of individuals age 18 or older. The table reports the fraction of respondents in each country who agree with each of the gender-related statements in the IVS and ISSP (across all available survey years). The sexism group classifications (high, mid, low) are based on country-specific average of the residual individual-responses after netting out the survey year effects.

Appendix Table 2. Data Sources for Cross-country Panel, by Year

	1995*	2000*	2005*	2010*
UK	EU-LFS/ECHP	EU-LFS/ECHP	EU-LFS/EU-SILC	EU-LFS/EU-SILC
France	EU-LFS/ECHP	EU-LFS/ECHP	EU-LFS/EU-SILC	EU-LFS/EU-SILC
Germany	EU-LFS/ECHP	EU-LFS/ECHP	EU-LFS/EU-SILC	EU-LFS/EU-SILC
Netherlands		EU-LFS/ECHP	EU-LFS/EU-SILC	EU-LFS/EU-SILC
Ireland	EU-LFS/ECHP	EU-LFS/ECHP	EU-LFS/EU-SILC	EU-LFS/EU-SILC
Austria	EU-LFS/ECHP	EU-LFS/ECHP	EU-LFS/EU-SILC	EU-LFS/EU-SILC
Belgium	EU-LFS/ECHP	EU-LFS/ECHP	EU-LFS/EU-SILC	EU-LFS/EU-SILC
Switzerland		EU-LFS/LIS	EU-LFS/LIS	EU-LFS/EU-SILC
Italy	EU-LFS/ECHP	EU-LFS/ECHP	EU-LFS/EU-SILC	EU-LFS/EU-SILC
Spain	EU-LFS/ECHP	EU-LFS/ECHP	EU-LFS/EU-SILC	EU-LFS/EU-SILC
Portugal	EU-LFS/ECHP	EU-LFS/ECHP	EU-LFS/EU-SILC	EU-LFS/EU-SILC
Greece	EU-LFS/ECHP	EU-LFS/ECHP	EU-LFS/EU-SILC	EU-LFS/EU-SILC
Sweden	EU-LFS/LIS	EU-LFS/ECHP	EU-LFS/EU-SILC	EU-LFS/EU-SILC
Finland	EU-LFS/LIS	EU-LFS/ECHP	EU-LFS/EU-SILC	EU-LFS/EU-SILC
Denmark	EU-LFS/ECHP	EU-LFS/ECHP	EU-LFS/EU-SILC	EU-LFS/EU-SILC
Iceland			EU-LFS/EU-SILC	EU-LFS/EU-SILC
Norway		EU-LFS/LIS	EU-LFS/EU-SILC	EU-LFS/EU-SILC
US	CPS	CPS	CPS	CPS
Canada	LFS	LFS	LFS	LFS
Taiwan	MPUS	MPUS	MPUS	MPUS
Hong Kong	Census	Census	Census	Census
Japan		JGSS	JGSS	JGSS
South Korea	KLIPS	KLIPS	KLIPS	KLIPS

* or closest year with available data

Notes: EU-LFS: European Union Labor Force Survey

ECHP: European Community Household Panel

EU-SILC: European Union Statistics on Income and Living Conditions

LIS: Luxembourg Income Study

CPS: US Current Population Survey

LFS: Canadian Labor Force Survey

MPUS: Taiwan Man Power Utilization Survey

JGSS: Japanese General Social Survey

KLIPS: Korean Labor and Income Panel Study

Appendix Table 3. Descriptive Statistics - Cross-country Data

	1995	2000	2005	2010
Difference in Evermarried Rates (high-skilled - low-skilled), Females age 35-44	-0.0718	-0.0586	-0.0446	-0.0368
	(0.0299)	(0.0526)	(0.0532)	(0.0660)
Difference in Evermarried Rates (high-skilled - low-skilled), Males age 35-44	0.00217	0.0130	0.0305	0.0404
	(0.0326)	(0.0440)	(0.0494)	(0.0487)
Share of Males with Tertiary Education	0.218	0.257	0.291	0.339
	(0.0767)	(0.0873)	(0.0882)	(0.0996)
Gender Gap in Share with Tertiary Education (Female - Male)	-0.0271	-0.0165	0.0101	0.0270
	(0.0460)	(0.0643)	(0.0756)	(0.0806)
GDP per capita (2014 dollars)	29,885	34,482	38,816	40,299
	(5,691)	(6,428)	(8,260)	(8,173)
Annual wages (2000 US$), high-skill females age 25 to 54	26,863	26,452	29,031	31,672
	(7,334)	(8,386)	(10,252)	(8,978)
Skill premium, males age 25 to 54	0.490	0.435	0.407	0.418
	(0.249)	(0.253)	(0.162)	(0.171)
Skill premium, females age 25 to 54	0.441	0.396	0.374	0.385
	(0.234)	(0.219)	(0.172)	(0.184)
LFP Gap (College - No College), females age 25 to 54	0.211	0.157	0.136	0.117
	(0.0745)	(0.110)	(0.0722)	(0.0794)
LFP Gap (College - No College), males age 25 to 54	0.0338	0.0387	0.0297	0.0318
	(0.0284)	(0.0273)	(0.0379)	(0.0365)
No. Countries	17	23	22	23

Source: See Appendix Table A2. Please refer to the text and data appendix for details on the construction of the variables.

Appendix Table 4: Classification of US States by Sexism Group

	GSS Attitude Measure: Better if man is achiever outside home and women take care of home and family		
	% Agree (average from 1973 to 2014)	Standardized residual after taking out year effects (mean 0, sd 1)	Sexism Group (1: Low, 2: Mid, 3: High)
Arkansas	0.61	2.78	3
Utah	0.57	2.12	3
West Virginia	0.56	1.72	3
Alabama	0.53	1.48	3
Mississippi	0.47	1.29	3
Kentucky	0.47	1.16	3
Tennessee	0.47	1.08	3
North Carolina	0.47	1.03	3
Florida	0.44	0.76	3
Texas	0.43	0.72	3
Montana	0.47	0.61	3
Louisiana	0.42	0.57	3
Georgia	0.42	0.54	3
Oklahoma	0.42	0.49	3
Delaware	0.37	0.40	3
Missouri	0.40	0.33	2
South Carolina	0.42	0.31	2
Washington	0.38	0.19	2
Pennsylvania	0.41	0.10	2
Illinois	0.39	0.03	2
Indiana	0.38	0.00	2
Ohio	0.39	-0.04	2
California	0.38	-0.12	2
Wisconsin	0.38	-0.22	2
Arizona	0.36	-0.22	2
New York	0.36	-0.30	2
Colorado	0.35	-0.36	2
Wyoming	0.36	-0.40	2
Virginia	0.35	-0.50	2
Kansas	0.35	-0.51	2
Rhode Island	0.37	-0.52	1
New Jersey	0.35	-0.56	1
Michigan	0.35	-0.59	1
Vermont	0.31	-0.60	1
South Dakota	0.29	-0.60	1
Oregon	0.34	-0.66	1
Massachusetts	0.31	-0.87	1
North Dakota	0.33	-0.89	1
Maryland	0.30	-1.03	1
Connecticut	0.31	-1.21	1
Iowa	0.30	-1.32	1
New Hampshire	0.33	-1.34	1
Minnesota	0.30	-1.35	1
Alaska	0.25	-1.54	1
District of Columbia	0.22	-1.94	1

Note: The data is from the 1973 to 2014 General Social Survey. The first column reports the fraction of individuals in each state who agree (strongly agree, agree) with the statement from 1977 to 2014. The second column reports the standardized value of the state-average of the residual individual-level responses to the gender-related question in the GSS after removing year effects. The state classifications (high, mid, low) reported in the last column are based on the standardized sexism measure in column (2). Refer to text for details on the construction of the attitude variable.

Appendix Table 5. Descriptive Statistics - US States

	1970	1980	1990	2000	2010
Female Marriage Gap (age 35-44)	-0.051	-0.051	-0.051	-0.045	0.004
	(0.044)	(0.022)	(0.017)	(0.024)	(0.031)
Male Marriage Gap (age 35-44)	0.006	-0.006	-0.001	0.018	0.064
	(0.034)	(0.016)	(0.017)	(0.029)	(0.041)
College Share - Males (age 35-44)	0.187	0.265	0.324	0.279	0.329
	(0.067)	(0.098)	(0.091)	(0.098)	(0.106)
College Share - Females (age 35-44)	0.098	0.168	0.273	0.290	0.374
	(0.044)	(0.091)	(0.098)	(0.102)	(0.110)
Log annual full-time wages - College females age 25-54 in *t-10*	9.594	9.860	9.893	10.115	10.342
	(0.171)	(0.157)	(0.158)	(0.205)	(0.189)
Skill premium for males age 25-54 in *t-10*	0.374	0.375	0.310	0.457	0.494
	(0.111)	(0.097)	(0.081)	(0.064)	(0.074)
Skill premium females age 25-54 in *t-10*	0.529	0.510	0.399	0.475	0.478
	(0.123)	(0.133)	(0.064)	(0.072)	(0.076)
No. of States	45	45	45	45	45

Note: The data is from the 1970 to 2000 US Census and the 2008 to 2011 ACS. The marriage gap is defined as the difference in evermarried rates between college-educated and non-college educated females/males. The sample is limited to 45 states due to data availability in the General Social Survey (used to create the state-level gender attitude measure).

Appendix Table 6. Household Production and Gender Norms - US States

	Minutes a day spent in activity relative to California | Demographics					
	Child Care			Household Work		
	(1)	(2)	(3)	(4)	(5)	(6)
	Females	Males	Both	Females	Males	Both
Better if man achiever outside home, woman take care of home (mean 0, var 1)	0.427	-4.171**	-2.049	2.599	-3.376**	-0.761
	[2.161]	[1.660]	[1.294]	[2.477]	[1.657]	[1.612]
Mean of Dep. Var	1.30	0.50	0.66	-19.56	-2.42	-11.99
Observations	44	44	44	44	44	44
R-squared	0.001	0.175	0.058	0.042	0.079	0.008

Note: The outcome variable is the state dummy coefficient of a regression of minutes spent a day in the activity on state fixed effects, day of the week dummies, a dummy for whether the day is a holiday, number of children, household size, age of the youngest child, education and age of the respondent, and if she has a partner living in the household. The construction of the outcome variable in columns (3) and (6) also controls for gender. The excluded state is for California. Data is from the 2003 to 2012 ATUS, with the sample restricted to individuals aged 25-54 with at least one own child living in the household. Robust standard errors in brackets, *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 7. Correlation between the Marriage Gap and Social Norms among Individuals age 35-44 across US States - Controlling for Total Household Work (Beta)

	Dep Var. Difference in Ever Married Rates (High Skilled - Low Skilled)							
	A. Females				B. Males			
	Year=2010				Year=2010			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
High Sexism Dummy	-0.025**		-0.023**		-0.021		-0.020	
	[0.010]		[0.010]		[0.013]		[0.013]	
Mid Sexism Dummy	-0.007		-0.011		-0.022**		-0.020**	
	[0.010]		[0.007]		[0.009]		[0.008]	
Attitude measure: Better if man is achiever outside home and women take care of home and family (mean 0, var 1)		-0.008**		-0.008**		0.000		0.000
		[0.003]		[0.003]		[0.005]		[0.005]
Total Household Work	-0.001	-0.000			0.000	-0.000		
	[0.000]	[0.000]			[0.001]	[0.001]		
Child Care Time			0.001	0.001*			0.000	0.000
			[0.001]	[0.001]			[0.001]	[0.001]
Other Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	44	44	44	44	44	44	44	44
R-squared	0.540	0.479	0.549	0.521	0.640	0.595	0.636	0.595

Note: See notes to Table 8. Total household work and Child care time are measured using the ATUS (see notes to Appendix Table 6). Robust standard errors are reported in brackets. *** p<0.01, ** p<0.05, * p<0.1.

Data Appendix for "Social Norms, Labor Market Opportunities, and the Marriage Gap for Skilled Women" by Bertrand, Cortes, Olivetti and Pan

1. Construction of Key Variables

Educational Classification

For consistency across countries, individuals are classified as "skilled" if they have completed tertiary education and "unskilled" if they did not complete tertiary education.

- EU-LFS data (UK, France, Germany, Netherlands, Ireland, Austria, Belgium, Switzerland, Italy, Spain, Portugal, Greece, Sweden, Finland, Denmark, Iceland, Norway)

  Skill levels were coded based on the level of education reported in the variable HATLEVID which summarizes the highest level of education or training successfully completed. Three classifications were provided: L, Lower secondary, M, Upper secondary and H, Third level. The L level corresponds to ISCED levels 1 to 3c. The M level corresponds to ISCED levels 3-4, and the H level corresponds to ISCED levels 5-6.

  Individuals were coded as having completed tertiary education if they were classified as "H" i.e. having completed ISCED levels 5-6.

Based on the ISCED description of the education programs, we construct similar measures of skill for the remaining countries. The details of how these variables are constructed for each country are as follows:

- US: CPS data – skilled individuals are classified as those whose highest educational attainment is an associate's degree or more.
  - Unskilled: None-Some college no degree (Codes 0-90, 100, 110)
  - Skilled: Associate degree (91, 92), Bachelor's degree (111), 5+ year of college or graduate degree (120-125).
- Canada: LFS data – skilled individuals are classified as those whose highest educational attainment is a bachelor's degree or more.
  - Unskilled: 0 to 8 years (0), Some secondary (1), Grade 11 to 13, graduate (2) Some postsecondary (3), Postsecondary degree or diploma (4)
  - Skilled: Bachelor's degree (5), Graduate degree (6)
- Taiwan: Manpower Utilization Survey (MPUS) – skilled individuals are classified as those whose highest educational attainment is junior college or more.
  - Unskilled: Illiterate (1), Self-educated (2), Primary school (3). Junior high school (4), Senior high school (5), Vocational school (6)
  - Skilled: Junior college (7), University (8), Masters (9) and PhD (10)
- Hong Kong: Census – skilled individuals are classified as those whose highest educational attainment is a diploma or more
  - Unskilled: No schooling, Kindergarten, Primary 1-6, Secondary 1-5, Grade 12, Secondary 6 and 7, Craft level (including apprenticeship)

- Skilled: Diploma/Vocational education, Higher Diploma, Associateship or equivalent courses in Universities, Diploma or certificate courses in post-secondary colleges, Nurse training courses, sub-degree distance learning courses, first-degree, postgraduate

- South Korea: Korean Labor and Income Panel Study– skilled individuals are classified as those whose highest educational attainment is a 2-year program with a vocational, technical, or associate degree.
  - Unskilled: No schooling (2), Elementary school (3), Lower secondary (4) and (5) Upper secondary.
  - Skilled: 2-years college, vocational, technical, associate degree (6), University (4 years or more) (7), Graduate school (master's) (8), graduate school (doctoral) (9)

Marriage market outcomes

For each country in our sample, we construct the ever marriage rates of women and men age 35-44 by education. Individuals are coded as 0 if they report being single and 1 if they report any other marital status classification (single, divorced, widowed, legally separated).

The definition of marriage for each country is as follows:

For countries in the EU-LFS data (UK, France, Germany, Netherlands, Ireland, Austria, Belgium, Switzerland, Italy, Spain, Portugal, Greece, Sweden, Finland, Denmark, Iceland, Norway), marital status was coded based on the following:

Marital status is the conjugal status of each individual in relation to the marriage laws of the country (i.e. de jure status). Some countries (i.e. the Nordic, the Netherlands, France) have a legal framework for registering partnerships (in most countries these are same-sex partnership and they have a legal status parallel to married couples). Such information has also to be treated in a harmonised way and it is proposed to treat them as married and classify them under group 2 when they still exists, else as 0 as appropriate (legal separation or death of one of the partners).

For the remaining countries (US, Canada, Taiwan, Hong Kong, Japan and South Korea), marital status was defined as legal marriage.

Labor market variables

Construction of wage series:

- Earnings sample restricted to full-time employees who work 35 or more hours per week
- Annual wages for full-time employees age 25 to 54

The country-specific annual wage measures are converted to 2000 USD dollars using the following procedure:

- Country-level wage measures are first deflated to 2000$ using the country-specific CPI from the World Bank

  World Bank country-specific CPI data:
  http://data.worldbank.org/indicator/FP.CPI.TOTL

- The European local currencies are then converted to Euros (for the EU member nations) using exchange rates in 2000.

  Historical local currency to Euro exchange rates: http://ec.europa.eu/eurostat/web/exchange-rates/data/database

- Finally, all wage measures are converted to 2000$ using PPP conversion factors for private consumption from the World Bank.

  World Bank PPP conversion factors, private consumption:
  http://data.worldbank.org/indicator/PA.NUS.PRVT.PP?page=3

For countries that were not included in the World Bank data, we used the following data sources to implement the conversions:

- Taiwan
  - CPI data from the Taiwan National Statistics Bureau
    http://eng.stat.gov.tw/ct.asp?xItem=12092&ctNode=1558&mp=5
  - PPP conversion rate for GDP at current prices (national currency per current USD) from: http://www.econstats.com/weo/V013.htm

- Hong Kong
  - CPI data from the Hong Kong Census and Statistics Department
    http://www.censtatd.gov.hk/hkstat/sub/so60.jsp

<u>European Community Household Panel (1994 to 2001)</u>

- Annual wage measures for 15 countries from 1994 to 2001
- Countries included: UK, Sweden, Finland, Denmark, Germany, Netherlands, Belgium, Austria, Ireland, France, Italy, Spain, Portugal, Greece

LIS data used to fill in the gaps from ECHP data for the following countries:
- Switzerland 1992 (recoded to 1995)
- Sweden, Norway, Finland 1995
- Switzerland 2000
- Switzerland 2004 (recode to 2005)