

GENDER IDENTITY AND RELATIVE INCOME WITHIN HOUSEHOLDS*

MARIANNE BERTRAND
EMIR KAMENICA
JESSICA PAN

We examine causes and consequences of relative income within households. We show that the distribution of the share of income earned by the wife exhibits a sharp drop to the right of $\frac{1}{2}$, where the wife's income exceeds the husband's income. We argue that this pattern is best explained by gender identity norms, which induce an aversion to a situation where the wife earns more than her husband. We present evidence that this aversion also impacts marriage formation, the wife's labor force participation, the wife's income conditional on working, marriage satisfaction, likelihood of divorce, and the division of home production. Within marriage markets, when a randomly chosen woman becomes more likely to earn more than a randomly chosen man, marriage rates decline. In couples where the wife's potential income is likely to exceed the husband's, the wife is less likely to be in the labor force and earns less than her potential if she does work. In couples where the wife earns more than the husband, the wife spends more time on household chores; moreover, those couples are less satisfied with their marriage and are more likely to divorce. These patterns hold both cross-sectionally and within couples over time. *JEL* Codes: D10, J12, J16.

I. INTRODUCTION

We begin by establishing the following fact: among married couples in the United States, the distribution of the share of household income earned by the wife drops sharply at $\frac{1}{2}$ —where the wife starts to earn more than the husband. Standard economic models of the marriage market cannot account for this pattern. Instead, we argue that *gender identity norms* play an important role in marriage.

Akerlof and Kranton (2000, 2010) import ideas about identity from sociology and social psychology into economics. They define identity as a sense of belonging to a social category, coupled with

*We thank John Abowd for pointing us to the SIPP/SSA/IRS Public Use File Project. Amanda Chuan and David Toniatti provided excellent research assistance. We thank the George J. Stigler Center for the Study of the Economy and the State and the Initiative on Global Markets, both at the University of Chicago Booth School of Business, and the Sloan Research Fellowship, for financial support. We thank the editors, anonymous referees, and numerous seminar participants for helpful comments and suggestions.

© The Author(s) 2015. Published by Oxford University Press, on behalf of President and Fellows of Harvard College. All rights reserved. For Permissions, please email: journals.permissions@oup.com

The Quarterly Journal of Economics (2015), 571–614. doi:10.1093/qje/qjv001.
Advance Access publication on January 29, 2015.

a view on how people in that category should behave. They propose that identity influences economic outcomes because deviating from the prescribed behavior is inherently costly. In one application of this model, the social categories are *man* and *woman*, and these categories are associated with specific behavioral prescriptions, such as “a man should earn more than his wife.”¹ Survey responses indeed suggest the prevalence of such prescriptions. For example, 38 percent of the U.S. respondents to the 1995 World Value Survey agree with the claim “If a woman earns more money than her husband, it’s almost certain to cause problems.” Moreover, this attitude persists in the more recent cohorts.²

In this article, we analyze how the behavioral prescription that “a man should earn more than his wife” affects social and economic outcomes. First we demonstrate that considerations of relative income affect whether people get married. Using a Bartik-style instrument, we show that when a randomly chosen woman within a marriage market becomes more likely to outearn a randomly chosen man, the marriage rate declines.³ This result suggests a potential link between two important social developments over the past several decades: the increase in women’s income relative to that of men and the decline in the prevalence of marriage. Our estimates imply that aversion to having the wife earn more than the husband explains 29 percent of the decline in marriage rates over the last thirty years.

We then assess the impact of gender identity on women’s labor supply. For each married woman in our sample, we estimate the

1. Baker and Jacobsen (2007) attempt to endogenize such social norms. They point out that a rule that mandates the customary division of labor within the household can be Pareto improving if it eliminates wasteful acquisition of human capital aimed at improving one’s bargaining position.

2. The overall sample includes U.S. respondents, both male and female, to the 1995 World Values Survey Question: “If a woman earns more money than her husband, it’s almost certain to cause problems. Do you agree strongly, agree, disagree, or disagree strongly?” We exclude those with missing education (since we later examine how attitudes vary by education) and those who responded “Don’t know” to the question. Out of the remaining 1,483 respondents, 38 percent “agree” or “agree strongly.” Among those born since 1965, 37 percent “agree” or “agree strongly.”

3. We flexibly control for the distribution of men’s income and the distribution of women’s income. Hence, we are not simply picking up the fact that women with higher incomes are less likely to get married.

distribution of potential earnings based on her demographics. We show that if the probability that her income would exceed the husband's income is high, she is less likely to participate in the labor force. Moreover, if she does work, the gap between her realized and potential income is greater (in part due to fewer hours of work). These patterns suggest that women reduce their labor supply so as to avoid gender role reversals in earnings. Of course, an important concern is that high-skill women who marry low-income men may have unobservable characteristics that keep them out of the labor force. We consider several approaches to deal with this issue. First, we show that the key coefficient is stable as we include various controls for characteristics of the couple. Second, controlling for (a proxy for) relative income at marriage does not affect our estimates. Most important, we document the same patterns within couples over time.

Even though it seems that couples try to avoid the situation where the wife earns more than the husband, this situation has nonetheless become more common. For example, in the American Community Survey 2008–2011, the wife earns more than the husband in 27 percent of the couples.⁴ Among these couples, the violation of the gender identity norm seems to influence the quality of marriage. Using data from the National Survey of Families and Households, we find that the couples where the wife earns more than the husband are less happy, report greater strife in their marriage, and are ultimately more likely to get a divorce.

We also examine how the violation of the gender identity norm affects the division of home production. We find that the gender gap in home production—how much more time the wife spends on non-market work than the husband—is *larger* in couples where she earns more than he does. This suggests that a “threatening” wife takes on a greater share of housework so as to assuage the husband's unease with the situation. The wife, of course, may ultimately get tired of working this “second shift” (Hochschild and Machung 1989), which might be one of the mechanisms behind our results on divorce.

Most of these results rely on cross-sectional variation across couples. To further alleviate concerns about omitted variable bias, we draw on data from the Panel Study of Income Dynamics (PSID) to examine how changes in relative income

4. Throughout the article, we focus exclusively on couples where both the husband and the wife are between 18 and 65 years of age.

within a couple affect the aforementioned outcomes. We replicate all of our results: when the wife's income exceeds the husband's, the wife becomes more likely to exit the labor force and she takes on more chores; moreover, there is suggestive evidence that divorce becomes more likely.

The fact that traditional gender role attitudes still influence behavior has important consequences. Over the last half a century, women have experienced substantial labor market gains; the gender gap in labor force participation and the gender gap in earnings have both declined.⁵ Despite these gains, substantial gender gaps remain, both in labor force participation and in earnings. Female labor force participation appears to have plateaued since the early to mid-1990s (Blau and Kahn 2006). Among full-time, full-year workers, the gender gap in earnings remains at 25 percent. This halted progress has led researchers to consider less traditional (within economics, at least) factors that might influence the gender gap in labor market outcomes (Bertrand 2010). The results we present in this article support the view that slow-moving identity norms are an important factor that limits further convergence in labor market outcomes. Women are bringing personal glass ceilings from home to the workplace.

Since the initial work by Becker (1973, 1974), the economic analysis of marriage markets has made great strides by developing tractable models that abstract from issues like tradition and identity. Consequently, while the economics literature on marriage markets is vast, little of it examines the role of gender identity. A few papers (e.g., Fernandez, Fogli, and Olivetti 2004; Fortin 2005, 2009;) have examined how the variation in gender attitudes (across countries, across time, or across couples) correlates with women's labor force participation. In contrast, this article examines the extent to which the overall prevalence of traditional attitudes affects a wide range of outcomes. In addition to women's labor force participation and the gender gap in

5. Several factors have been identified as contributing to these gains. First and foremost has been the reversal of the gender gap in education (Blau and Kahn 2006; Goldin, Katz, and Kuziemko 2006). Various technological innovations, such as the contraceptive pill, have favored women (Goldin and Katz 2002; Greenwood, Seshadri, and Yorukoglu 2005). Labor demand has shifted toward industries where female skills are overrepresented (Black and Juhn 2000; Weinberg 2000). Finally, better regulatory controls and greater competitiveness have reduced labor market discrimination against women (Black and Strahan 2001; Black and Brainerd 2004).

income, we study the distribution of relative income within households, marriage rates, division of home production, marriage satisfaction, and divorce.⁶ Some of these latter issues—especially the impact of relative income on household production and divorce—have been explored in the sociology literature (West and Zimmerman 1987). We discuss the relationship of our results to the sociology literature in the relevant parts of the article. Fisman et al. (2006) document a pattern related to our results in a dating context: men value women’s intelligence or ambition if and only if it does not exceed their own.

II. RELATIVE INCOME WITHIN HOUSEHOLDS

II.A. *Distribution of Relative Income*

Figure I depicts the distribution, across married couples in the United States, of the share of the household income earned by the wife. Specifically, we use the Survey of Income and Program Participation (SIPP), which is linked to administrative data on income from the Social Security Administration and the Internal Revenue Service. The SIPP consists of a series of national panels, each representative of the U.S. civilian, noninstitutionalized population. For each married couple, we use the observation from the first year that the couple is in the panel, which ranges from 1990 to 2004. This leaves us with 73,654 couple-level observations.

We define $relativeIncome_i$ as $\frac{wifeIncome_i}{wifeIncome_i + husbIncome_i}$ where i indexes the couple, and $wifeIncome_i$ and $husbIncome_i$ are the total labor and self-employment income of the wife and the husband, respectively. Figure I depicts the frequency distribution of $relativeIncome_i$ grouped in 20 bins, along with a lowess (locally weighted scatterplot smoothing) estimate of the distribution on each side of $relativeIncome_i = \frac{1}{2}$.

The distribution exhibits a sharp drop at the point where the wife starts to earn more than the husband. The McCrary (2008)

6. Using administrative data from Denmark, Pierce, Dahl, and Nielsen (2012) employ a regression discontinuity design to argue that a husband is more likely to use erectile dysfunction medication if he earns less than his wife. Stuart, Moon, and Casciaro (2011) find that winning an Academy Award is associated with a greater risk of divorce for Best Actresses but not for Best Actors.




FIGURE I

Distribution of Relative Income (SIPP Administrative Data)

The data are from the 1990 to 2004 SIPP/SSA/IRS gold standard files. The sample includes married couples where both the husband and wife earn positive income and are between 18 and 65 years of age. For each couple, we use the observation from the first year that the couple is in the panel. Each dot is the fraction of couples in a 0.05 relative income bin. The vertical line indicates the relative income share=0.5. The dashed line is the lowest smoother applied to the distribution allowing for a break at 0.5.

test for the discontinuity of the distribution function estimates that the distribution drops by 12.3 percent ($p < .01$).⁷

We next turn to examining how the distribution of relative income varies by decade, by whether the couple has any children, and by how long the couple has been married. We cannot study these questions using the administrative data from the SIPP because there are too few observations. Instead we turn to data from the U.S. Census Bureau. Specifically, we use the 1970–2000 U.S. Censuses and the American Community

7. Since there is a positive—albeit small, about 0.0026—fraction of couples with relative income exactly at $\frac{1}{2}$, we estimate the discontinuity to the right of $\frac{1}{2}$. We also conduct a placebo exercise, testing for the presence of a discontinuity to the right of each point in the set $\{0.100, 0.105, \dots, 0.445\} \cup \{0.555, 0.560, \dots, 0.900\}$. We find that there is a significant discontinuity ($p < .05$) at roughly 5%, namely, $\frac{8}{140} = 0.057$, of these points.

Survey (ACS) 2008 to 2011. Our sample includes household heads and their spouses.⁸

One important issue with these data is that there is a large number of couples—a few percent—with relative income exactly at $\frac{1}{2}$. We know from the administrative data that the actual distribution of relative income does not exhibit such a spike. There are at least four reasons the spike might be present in the Census Bureau data. First, incomes are top-coded, so any couple where both individuals earn more than the upper bound will have a relative income of $\frac{1}{2}$. Second, a nontrivial fraction of incomes in the census data are imputed. Third, income for both spouses can be reported by a single respondent. Fourth, except for the 1990 census, individual incomes are rounded for confidentiality reasons, which increases the chance that two incomes exactly coincide. To deal with these issues, we apply the following steps. First, we drop couples where either individual's income is imputed. Second, we drop couples where both spouses have top-coded incomes; for couples where one spouse has a top-coded income, we multiply his or her income by 1.5 and leave the other individual's income as is. Even after following these steps, a substantial mass at $\frac{1}{2}$ remains (around 3 percent of couples), which could be a reflection of either rounding or misreporting. Hence, in a final step, we “de-round” the data, save for 0.26 percent of couples (the fraction of couples with relative income of $\frac{1}{2}$ in the administrative data). To de-round the data, we replace each individual's income with a uniform draw from the range of incomes that would have been rounded to that income.⁹ All of the figures that draw on the census and ACS data have been “cleaned” in the sense that the spike at $\frac{1}{2}$ has been dealt with through the procedure described above.

Figure II shows the cleaned distribution of relative income in the census, restricting the data to the 1990 and 2000 censuses and the 2008–2011 ACS to roughly match the years represented in the SIPP. The distribution is quite similar to the one based on the administrative data. Most important, the extent of the discontinuity at $\frac{1}{2}$ is also similar. The drop in the distribution in the

8. We do not include other household members who are married because the spouse links for these are based on IPUMS imputed values.

9. Our results are similar if we instead assign incomes based on demographics or based on a triangular kernel.




FIGURE II

Distribution of Relative Income (Census Bureau Data)

The data are from the 1990 and 2000 censuses and 2008 to 2011 ACS single-year files. The sample includes married couples where both the husband and the wife earn positive income and are between 18 and 65 years of age. Each dot is the fraction of couples in a 0.05 relative income bin. The vertical line indicates the relative income share=0.5. The dashed line is the lowest smoother applied to the distribution allowing for a break at 0.5.

Census Bureau data is 13.6% ($p < .01$), compared to the 12.3% drop in the SIPP.¹⁰

Using the Census Bureau data, we can examine how the distribution of relative income has evolved over time. As Figure III shows, the drop in the distribution at equal incomes is present in each decade. That said, the size of the discontinuity seems to have gotten smaller in the past 20 years. In 1970, the drop is 20.6%; in 1980, 26.2%; in 1990, 24.4%; in 2000, 10.8%; finally in the ACS 2008–2011 aggregate, the drop in the distribution is 10.0%. All estimates are significant ($p < .01$).

We next turn to the presence of children. There are a number of reasons the presence of children in the household might affect

10. We investigated the potential influence of the tax code on the distribution of relative income. We focused on the 1990 census (where income data are not rounded) and couples without children (whose tax liability is easier to impute). We computed the fraction of couples within \$200 of a tax bracket cutoff across levels of relative income. There is no evidence that bunching at tax bracket cutoffs is more prevalent when relative income is 0.5.

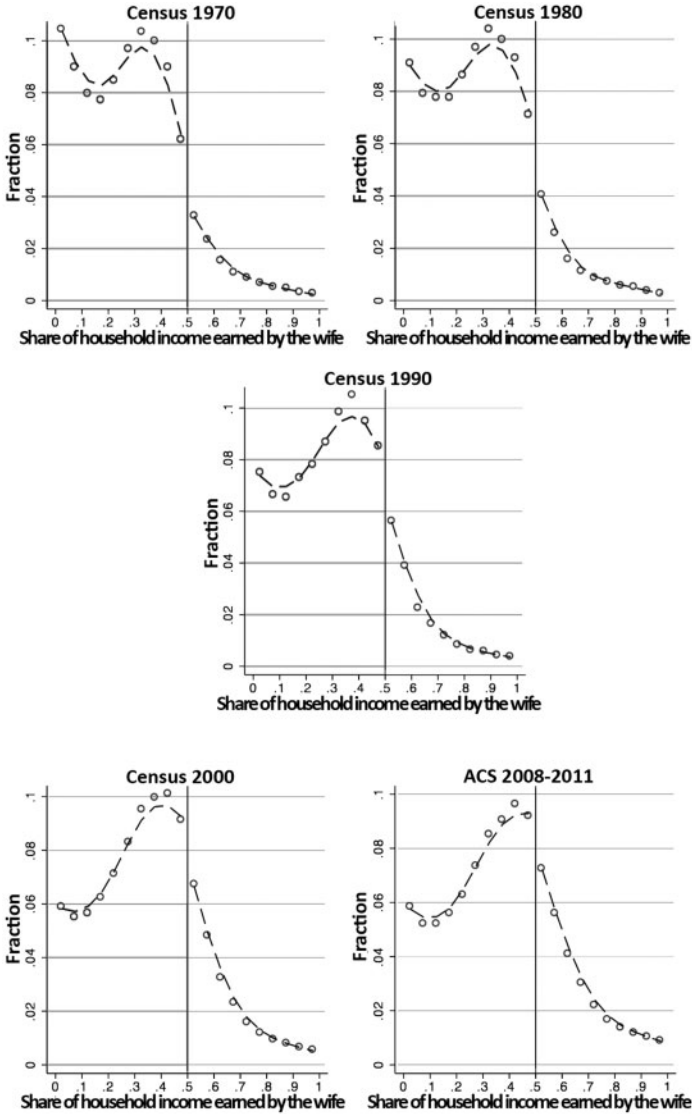


FIGURE III

Distribution of Relative Income over Time (Census Bureau Data)

The data are from the 1970 to 2000 censuses and 2008 to 2011 ACS single-year files. The sample includes married couples where both the husband and the wife earn positive income and are between 18 and 65 years of age. Each dot is the fraction of couples in a 0.05 relative income bin. The vertical line indicates the relative income share=0.5. The dashed line is the lowest smoother applied to the distribution allowing for a break at 0.5.

the overall distribution of relative income, but the discontinuity at equal incomes turns out to be present both for couples with and without a child: among those with a child, the drop is 15.2% ($p < .01$); among those without, the drop is 11.6% ($p < .01$).¹¹ These distributions are depicted in Figure A.1 in the Online Appendix.¹²

Overall the observed discontinuity could be due to the fact that couples avoid getting married if she earns more than he does, due to evolution of relative income during marriage, or due to the impact of relative income on divorce. Examining how the size of the discontinuity varies with marriage tenure sheds some light on the relative importance of these channels. For this exercise, we restrict our attention to the 2008–2011 ACS since earlier census data do not include information on how long the couple has been married.

We find that the discontinuity is present, albeit with a smaller drop of 8.4%, among couples who have been married for one year or less. Among couples married 2–5 years, the drop is 10.1%; among those married 6–10 years, 12.9%; and finally, among couples who have been married for more than 10 years, the drop in the distribution at equal incomes is 10.5%. All estimates are highly statistically significant ($p < .01$).¹³ In the next subsection, we argue that the discontinuity at equal incomes is driven by gender identity norms. Under this interpretation, the discontinuity among the newlyweds implies that gender identity affects who marries whom, while the fact that the discontinuity grows with marriage tenure (over the first decade of marriage) suggests that identity considerations also influence the evolution of relative income within a couple and/or the likelihood of divorce. In Sections IV and V we explore these two channels in greater detail.

II.B. Relative Income in Standard Models of the Marriage Market

In standard models of the marriage market, marriages are partnerships formed for the purpose of joint production and joint

11. For these calculations, we restrict attention to the 1990 and 2000 censuses and the 2008–2011 ACS for ease of comparison with the discontinuity in the overall distribution.

12. The Online Appendix contains all supplemental figures and tables and the detailed description of the data.

13. Figure A.2 in the Online Appendix depicts the distributions of relative income by marriage tenure.

consumption (see Weiss 1997 for a survey). A widely studied class of models is one where men and women are endowed with a single-dimensional attribute that positively affects the family output. If utility is nontransferable, equilibrium induces positive assortative matching. If utility is transferable, equilibrium is consistent with either positive or negative assortative matching depending on whether the individuals' attributes are complements or substitutes. Whether positive or negative, assortative matching relates the *ranks* of individuals in their gender-specific distribution of attributes. For example, if matching is positive and income is the attribute of interest, a woman in the 30th percentile of women's income distribution will match with the man who is in the 30th percentile of men's income distribution. Whether the woman earns more or less than the man in absolute terms has no significance in these models.¹⁴

A second important class of models is one where marriages enable division of labor to exploit comparative advantage or increasing returns. In case of increasing returns, it is optimal for only one spouse to work, so for purposes of studying relative income we focus on comparative advantage. With decreasing returns and comparative advantage, both spouses may participate in the labor force; their contributions to household income and to household production are determined by their relative productivity in those two activities. In these models, women tend to do more chores and men tend to earn higher incomes because women are relatively more productive at home while men are relatively more productive in the labor force. That said, the point where the wife and the husband earn the same income again plays no special role—the gains from marrying a man are continuous in his labor market productivity with no discrete jump at the point where endogenous incomes would be equal. Moreover, in Section VI we show that the amount of time a wife spends on household production increases when she earns more than the husband, which stands in sharp contrast with the basic prediction of this class of models.

Existing literature also distinguishes modes of decision making within households. A key distinction is between models with a common objective where all decisions are taken to

14. As we discuss in Section V, difference in income percentiles between the spouses does not predict divorce, whereas wife earning more than the husband in absolute terms does.

maximize a shared utility function, and models where the husband and the wife have distinct preferences and bargain their way to a Pareto efficient outcome. Relative income could clearly influence the bargaining positions within the couple, but once again equilibrium outcomes are continuous in the outside option with no special significance of the point where the two incomes are equal.¹⁵

Standard models of the marriage market thus cannot account for the observed distribution of relative income with its discontinuous drop at the point where the wife earns more than the husband.¹⁶ One natural interpretation of this discontinuity is simply that some couples try to avoid the circumstance where the wife earns more than the husband. Recall that a substantial fraction of U.S. respondents to the World Values Survey agrees with the claim that “If a woman earns more money than her husband, it’s almost certain to cause problems.” The tendency to agree with this claim turns out to be much stronger among less educated respondents: only 28% of the couples where both the husband and the wife have at least some college education agree with the claim, compared with 45% of the couples where neither spouse went beyond high school. Hence, if gender role attitudes are indeed the source of the cliff in the distribution of relative income, we should expect the discontinuity to be greater among less-educated couples. This is indeed the case. Among less-educated couples, the distribution drops by 20.1% ($p < .01$). Among more educated couples, there is only a (statistically insignificant) drop of 5.53% (see Figure A.3 in the Online Appendix). This difference in the discontinuity across these education groups is significant ($p < .05$).

III. MARRIAGE RATES AND RELATIVE INCOME

For the past 40 years, marriage rates in the United States have been steadily declining. Between 1970 and 2008, the fraction of young adults who are currently married decreased by 30 to

15. See Lundberg and Pollak (1996) for a survey of bargaining models of the household.

16. We discuss here only a fraction of the vast literature on marriage markets, but similar arguments can be used to show that other motives for marriage formation—credit constraints, public consumption, risk pooling, and so on—also cannot account for the observed distribution of relative income.

50 percentage points among all race, gender, and education groups (Autor 2010).¹⁷ Over the same period, women's income has greatly increased relative to that of men. Results from the previous section suggest a potential link between these two trends: if couples dislike unions where the husband earns less than the wife, then, as women start to command a greater share of labor income, marriages may become less appealing.

In this section, we analyze how the share of individuals who are currently married varies with the distributions of men's and women's (potential) income. Throughout the section, we use 1980 to 2000 data from the U.S. Census and the 2008–2011 American Community Survey.¹⁸ We assign individuals to marriage markets based on the pattern of homophily: most marriages occur between men and women who are of the same race and are of similar age and education.¹⁹ Moreover, marriages tend to form between individuals who live close to each other. Accordingly, we define marriage markets based on race, age group, education group, and the state of residence. The three race groups we consider are (non-Hispanic) whites, (non-Hispanic) blacks, and Hispanics.²⁰ The three age groups are (i) 22–31 for women and 24–33 for men; (ii) 32–41 for women and 34–43 for men; and (iii) 42–51 for women and 44–53 for men. The two education groups are (i) high school degree or less, and (ii) some college or more. Online Appendix Table A.1 documents sorting along these dimensions. For example, 98% of wives who are white are married to a husband who is white,²¹ 72% of wives with a high school degree or less are married to a husband with similar educational qualifications, and 76% of wives aged 22–31 are married to a husband aged 24–33. Overall, 59% of all marriages are between a man and woman from the same marriage market.²²

Given a particular marriage market, we wish to know how the changes in women's income relative to that of men affect

17. Part of this decrease is due to delay in marriage, but the fraction of older adults who are married has also been declining.

18. For convenience, we refer to the 2008–2011 ACS as 2010 data.

19. Our approach to assigning individuals to marriage markets is similar to those used in Charles and Luoh (2010) and Loughran (2002).

20. We drop individuals of other races.

21. For a broader discussion of same-race marriages in the United States, see Fisman et al. (2008).

22. This share has not changed much over the time period we consider (it ranges from a minimum of 57% in 2000 to a maximum of 61% in 1980).

marriage rates. For each marriage market m and year $t \in \{1980, 1990, 2000, 2010\}$ we compute how likely it is, when a woman encounters a man, that her income exceeds his. Specifically, given woman i and man j , consider a binary variable that takes value 1 if i 's income exceeds j 's. We define $PrWomanEarnsMore_{mt}$ as the mean of this variable taken across all possible couples. Operationally, we construct this variable by randomly drawing 50,000 women and men with replacement and computing the share of couples where the woman earns more than the man.

We consider several measures of income. First we use individuals' actual earnings, where we code an individual as having zero income if he or she is not in the labor force. Second we construct a measure of predicted earnings based on demographic characteristics. In particular, we assign each woman and man in a census year to a demographic group defined based on race, age,²³ education,²⁴ and state of residence. We then assign potential income to each individual by drawing from the earnings distribution of those in the demographic group who have positive income. Finally, for our preferred specification, we construct distributions of income based on a Bartik-style instrument that isolates the variation in relative income, which is plausibly unrelated to the factors that directly affect the marriage market.

Across all census years and marriage markets, the likelihood that a randomly chosen woman earns more than a randomly chosen man is about 0.25. This likelihood has increased steadily over time, going from 17–20% in 1980 to about 31–33% in 2010.²⁵ More important for our purposes, these dynamics have varied across marriage markets.²⁶ Thus, there is ample variation in $PrWomanEarnsMore_{mt}$ even when we include marriage market and year fixed effects. Note that this residual variation stems from both compositional shifts within a marriage market over time and from shocks that differentially affect men and women within a marriage market. When we turn to our instrumental variables approach, we isolate the component of the latter variation which stems from nationwide changes in labor demand across industries.

23. We group age into three-year intervals.

24. We group education into the following categories: less than high school, high school, some college, college, more than college.

25. See Online Appendix Table A.2 for summary statistics.

26. See Online Appendix Table A.3.

Our baseline OLS specification regresses $shareMarried_{mt}$ —the share of males who are currently married²⁷—on $PrWomanEarnsMore_{mt}$, controlling for the logs of the average female and male income and the logs of female and male income at each income decile in the marriage market that year. All specifications include marriage market fixed effects as well as year fixed effects interacted with race, age group, education group, and the state of residence. We include these interactions with year fixed effects because the relationship between demographic variables and marriage rates may have changed over time. The unit of observation is a marriage market in a census year. Standard errors are clustered by state, and each observation is weighted by the number of women in the marriage market.

This baseline specification is in column (1) of Table I. The estimated impact of $PrWomanEarnsMore_{mt}$ on $shareMarried_{mt}$ is -0.080 , but is not statistically significant. Column (2) adds a control for average women's income divided by the sum of average men's and women's income. With this specification, the estimated effect remains negative but is smaller and still not significant. In column (3), we include several additional marriage market by year controls: the sex ratio, male and female incarceration rates, average years of schooling for men and women, and the number of men and women in the market. The estimated effect becomes stronger and significant.

In columns (4)–(6) we consider the same three specifications, but we construct the variable $PrWomanEarnsMore_{mt}$ (as well as all other relevant income variables in the regression) using predicted income. As we show later, relative income considerations have a direct impact on the realized income of males and females, so predicted income is likely to be a cleaner measure of the relevant distributions. In these specifications, estimated impact of $PrWomanEarnsMore_{mt}$ on $shareMarried_{mt}$ is consistently negative, stable (ranging from -0.236 to -0.266 across the three specifications), and always highly significant ($p < .01$).

The identifying assumption behind these specifications is that $PrWomanEarnsMore_{mt}$ is uncorrelated with unobserved shocks that influence marriage rates. The fact that the coefficient of interest is stable across columns (4)–(6) somewhat ameliorates concerns about omitted variables, but to provide further support

27. We get similar results if we use the share of women who are married.

TABLE I
POTENTIAL RELATIVE INCOME AND MARRIAGE RATES

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
		Actual			Predicted			Bartik	
Income measure:									
<i>PrWomanEarnsMore</i>	-0.080 [0.075]	-0.046 [0.080]	-0.209*** [0.074]	-0.266*** [0.068]	-0.252*** [0.066]	-0.236*** [0.062]	-0.515*** [0.189]	-0.343* [0.183]	-0.351* [0.181]
In Average Women's Income	0.055* [0.030]	0.171** [0.071]	0.088 [0.074]	0.066* [0.036]	0.266*** [0.108]	0.151 [0.108]	0.270 [0.177]	0.943*** [0.333]	0.461 [0.371]
In Average Men's Income	0.023 [0.032]	-0.092 [0.070]	0.005 [0.073]	-0.001 [0.053]	-0.201** [0.084]	-0.063 [0.093]	0.114 [0.140]	-0.558* [0.292]	-0.097 [0.348]
Sex Ratio			-0.030*** [0.007]			-0.027*** [0.007]			-0.006 [0.007]
Female Incarceration Rate			-0.369 [0.241]			-0.292 [0.232]			-0.048 [0.172]
Male Incarceration Rate			0.433*** [0.089]			0.210*** [0.071]			0.056 [0.069]
Female Average Years of Education			0.009 [0.008]			0.005 [0.007]			-0.002 [0.007]
Male Average Years of Education			-0.031*** [0.010]			-0.023*** [0.008]			-0.010 [0.007]
Number of Females (per million)			0.001 [0.005]			0.003 [0.006]			-0.003 [0.008]
Number of Males (per million)			0.004 [0.005]			0.002 [0.006]			0.005 [0.007]

TABLE I
(CONTINUED)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Actual		Predicted						
	Dependent variable: <i>shareMarried</i>								
Observations	3,635	3,635	3,635	3,578	3,578	3,578	3,636	3,636	3,636
R-squared	0.990	0.990	0.991	0.990	0.990	0.991	0.991	0.991	0.991
Control for relative income	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Control for deciles of men's and women's income	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Data are from the 1980 to 2000 U.S. Censuses and 2008 and 2011 ACS single-year files. The unit of observation is a marriage market by decade (1980, 1990, 2000, and 2010, using 2008–2011). A marriage market is defined by the interaction of age group, education group, race, and state of residence. See text for further details on the construction of the marriage market. *shareMarried* is the share of males who are married. *PWomanEarnsMore* is the probability that a randomly chosen woman earns more than a randomly chosen man. *Sex Ratio* is the ratio of the male population to the female population in a marriage market. All specifications include marriage market fixed effects, decade fixed effects, and the decade interacted with the age group, the education group, the race, the state of residence and deciles of men's and women's income. "Control for relative income" is defined as the ratio of average women's income to the sum of average men's and women's income. Columns (1)–(3) use actual earnings to calculate *PWomanEarnsMore* (and all other relevant income variables used in the regression); columns (4)–(6) use predicted earnings to calculate *PWomanEarnsMore* (and all other relevant income variables in the regression) and columns (7)–(9) use the Bartik specification to calculate *PWomanEarnsMore* (and all other relevant income variables used in the regression). Regressions are weighted by the number of women in the marriage market. Standard errors clustered at the state level are in brackets. ***, ***, ** significant at 1%, ** at 5%, * at 10%.

for our causal interpretation, we now turn to an instrumental variables approach.

Historically men and women have tended to work in different industries (e.g., women are overrepresented in services and men in construction and manufacturing). Based on the industry composition of the state and the industry-wide wage changes at the national level, we can thus isolate gender-specific variation in local wages that is driven solely by aggregate labor demand (which is presumably uncorrelated with the characteristics of workers in a given marriage market). This approach builds on previous work by Bartik (1991) and Aizer (2010). In contrast to previous uses of the “Bartik instrument,” which focus on changes in *average* wages, we construct an instrument for the entire *distribution* of potential income in each marriage market.

We begin by instrumenting for average yearly wages by gender and marriage market as follows:

$$\bar{w}_{mt}^g \equiv \sum_j \gamma_{rejs,1980}^g \times w_{reajt,-s}^g$$

where g indexes gender, r race, e education group, a age group, j industry,²⁸ t census year, and s state. Variable $w_{reajt,-s}^g$ is the average wage in year t in industry j for workers of a given gender, race, education, and age group in the nation, excluding state s . Variable $\gamma_{rejs,1980}^g$ is the fraction of individuals with gender g , race r , and education e in state s who are working in industry j , as of the base year 1980.²⁹

Variable \bar{w}_{mt}^g is strongly correlated with the actual mean income of gender g in marriage market m in year t : states that initially had relatively more women in industries that subsequently experienced wage growth at the national level tend to have more growth in women’s income relative to that of men.

28. We consider 12 industry groups: agriculture; mining; construction; manufacturing; transportation; wholesale trade; retail trade; finance, insurance, and real estate; business, personal, and repair services; entertainment and recreation services; professional services; and public administration.

29. We choose to use 1980 as the base year instead of an earlier decade (e.g., 1970) as the earlier census data sets are only 1% samples. Thus, compared to the 1980 census data, there are many fewer observations, which would result in noisier estimates of γ_{rejs}^g .

But unlike the variation in actual income, variation in \bar{w}_{mt}^g over time is driven by aggregate shocks and is thus plausibly orthogonal to factors that might directly influence marriage rates in market m .

Similarly, we wish to construct a measure of the entire distribution of income by gender which is driven solely by aggregate shocks. We modify the standard Bartik instrument to compute predicted yearly wages at the $p = \{5\text{th}, 10\text{th}, 15\text{th}, \dots, 90\text{th}, 95\text{th}\}$ percentile.

Specifically, let

$$\bar{w}_{mt}^{g,p} \equiv \sum_j \gamma_{rejs,1980}^g \times w_{reajt,-s}^{g,p}$$

where $w_{reajt,-s}^{g,p}$ is the p th percentile of the national income distribution in year t in industry j for workers of a given gender, race, education, and age group, excluding state s . A priori it is not clear that $\bar{w}_{mt}^{g,p}$ will be correlated with the p th percentile of gender g 's distribution in market mt . For example, if half the women in a demographic group m work in some industry j^{high} where the minimum income is y^{high} and half the women in m work in some industry j^{low} where the maximum income is $y^{low} < y^{high}$, increase in the 5th percentile of wages in industry j^{high} will not raise the 5th percentile of wages of women in market m . This example, however, has little empirical relevance—a posteriori, the distributions defined by $\{\bar{w}_{mt}^{g,p}\}_p$ indeed correlate with the actual distributions of income. In other words, the Bartik instrument has a strong “first stage” when it is used to predict how the distribution of income varies across markets (see Online Appendix Table A.4).

This modification to the standard Bartik approach allows us to construct a measure of *PrWomanEarnsMore* $_{mt}$ whose variation over time is orthogonal to local labor market conditions. Specifically, we draw from the distributions defined by $\{\bar{w}_{mt}^{g,p}\}_p$, and calculate the likelihood that a randomly chosen women earns more than a randomly chosen man. Column (7) repeats the baseline specification from column (1), but using \bar{w}_{mt}^g and $\bar{w}_{mt}^{g,p}$ to construct all measures of income. The estimated impact of *PrWomanEarnsMore* $_{mt}$ on the marriage rate is negative (-0.515) and significant ($p < .01$). In column (8) we control for relative income. The estimate declines to -0.343 and becomes significant only at the 10% level. Finally, in column (9), we

include a set of marriage market by year controls. The estimate is similar at -0.351 and remains significant at the 10% level.³⁰

We also estimate the specification in column (9) separately for the less educated (high school or less) and more educated (some college or more) marriage markets. The impact of *PrWomanEarnsMore_{mt}* on the marriage rate is -0.631 (std. err. = 0.174) in the less educated markets, compared to 0.141 (std. err. = 0.177) in the more educated markets. This pattern further supports the view that the observed relationship is driven by traditional attitudes toward gender roles.

Taken together, these results highlight the importance of the relative distribution of men and women's income in marriage markets. The estimate from our preferred specification (column (9)) implies that the secular increase in the income of women relative to that of man explains 29 percent of the overall decline in marriage rates from 1980 to 2010.³¹

Note that the relative distribution of men and women's income might influence the formation of marriage even in the absence of gender identity considerations. In Beckerian models of the marriage market, one of the key benefits of marriage is specialization. Specialization, in turn, is more valuable if a man and a woman have different opportunities in the labor market. As *PrWomanEarnsMore* increases, there are smaller "gains from trade" that can be achieved through marriage.³² This force alone might account for our results. That said, the evidence we present in other sections of this article is in direct conflict with the standard models of the marriage market and, as we already discussed, standard models cannot easily explain the observed distribution of relative income. Thus, the view that couples have an aversion to the wife earning more than the husband

30. Variability in the income of potential spouses can also lead to a lower marriage rate. This fact, however, cannot account for our findings because we include controls for the deciles of the distributions of men and women's income in the marriage market in all specifications. In other words, all specifications in Table I already control for any linear effect of variability in income on marriage rates.

31. The coefficient -0.351 multiplied by the 14 percentage point increase in *PrWomanEarnsMore* is 29% of the 17 percentage point decrease in marriage rates (Online Appendix Table A.2).

32. Oppenheimer (1997) argues against the view that reduced gains from trade account for the decrease in marriage rates, but she focuses on the impact of women's income per se rather than the relative income of men and women.

provides a more parsimonious explanation of the various patterns we present herein.

IV. WOMEN'S LABOR SUPPLY AND RELATIVE INCOME

The previous sections establish that couples are less likely to form if the wife's income would exceed the husband's. When such couples do form, we might expect gender identity to distort labor market outcomes. A wife whose income would exceed her husband's may choose to stay at home so as to be less threatening. Or she may distort her labor supply in other ways—for example, work fewer hours or take a job that is less demanding and pays less. In this section we analyze such potential distortions in the wife's labor force participation and labor market outcomes.

This analysis complements existing work on how the variation in attitudes toward gender influences women's labor supply. Fortin (2005) uses data from the World Values Surveys in 25 Organisation for Economic Co-operation and Development (OECD) countries over a 10-year period and shows that the social representation of women as homemakers and men as breadwinners is associated with a low labor force participation by women and a large gender gap in income. Fortin (2009) examines a similar question in a single country (the United States) over a longer time period (1977–2006). She shows that the evolution of gender role attitudes over time correlates with the evolution of female labor force participation. In particular, while women's gender role attitudes steadily became less traditional until the mid-1990s (e.g., more and more women came to disagree with the notion that husbands should be the breadwinners and wives should be the homemakers), these trends reversed in the mid-1990s, precisely at the time that coincides with the slowdown in the closing of the gender gap in labor force participation.

Throughout this section, we use data on married couples from the 1970 to 2000 U.S. Censuses and the ACS single-year files from 2008 to 2011 (which we refer to as the 2010 ACS for simplicity). We restrict the sample to those couples where the husband is working. For each couple i , we estimate the distribution of the wife's potential earnings as follows. For $p \in \{5, \dots, 95\}$, we define w_i^p as the p th percentile of earnings among working women in the wife's demographic group that year. We assign the demographic group based on race, age group, education

level, and state of residence.³³ We then define a variable *PrWife EarnsMore*_{*i*} = $\frac{1}{19} \sum_p \mathbf{1}_{\{w_i^p > \text{husbIncome}_i\}}$ where *husbIncome*_{*i*} is the husband's income.³⁴ Thus, whether the wife works or not, *PrWifeEarnsMore*_{*i*} captures the likelihood that she would earn more than her husband if her income were a random draw from the population of working women in her demographic group.³⁵

Summary statistics for this sample are presented in Online Appendix Table A.5. Across all census years, the mean of *PrWifeEarnsMore*_{*i*} is 0.18. Not surprisingly, this probability has increased over time, from 0.09 in the 1970 census to 0.27 in the 2010 ACS. Across all years, about 67% of wives participate in the labor force. As we mentioned in the introduction, wives' labor force participation increased steeply between 1970 and 1990 (going from 43% to 70%), but has essentially plateaued since 1990, with 74% of wives in the labor force in the 2010 ACS.

IV.A. Labor Force Participation

We first examine wives' labor force participation. One of the strongest ways to conform to traditional gender roles is for the wife to stay at home while the husband plays the role of breadwinner. Might it be the case that some couples retreat to traditional gender roles when gender identity is threatened by the possibility that the wife would be the primary provider?

Given a couple *i*, let *wifeLFP*_{*i*} be a binary variable equal to 1 if the wife is in the labor force. In column (1) of Table II, we consider, as the baseline specification, a linear probability model:

$$\begin{aligned} \text{wifeLFP}_i = & \beta_0 + \beta_1 \times \text{PrWifeEarnsMore}_i \\ & + w_i^p + \beta_2 \times \ln\text{HusbIncome}_i + \beta_3 \times X_i + \varepsilon_i, \end{aligned}$$

where $\ln\text{HusbIncome}_i$ is the log of husband's income, w_i^p are controls for the wife's potential income at each of the vigintiles, and X_i represents nonincome controls: year fixed effects, state fixed

33. Throughout this section, by "age group" we mean five-year intervals and by "education level" we mean the following categories: less than high school, high school, some college, college, more than college.

34. Since we estimate the distribution based on every 5th percentile from the 5th to the 95th, *PrWifeEarnsMore* is based on 19 percentiles.

35. As we discuss in the next subsection, the income of the women who do work may also be distorted by gender identity considerations. By "potential income" we mean the (possibly distorted) income that the wife would likely earn were she to join the labor force.

TABLE II
POTENTIAL RELATIVE INCOME AND WIFE'S LABOR FORCE PARTICIPATION

	(1)	(2)	(3)	(4)	(5)	(6)
<i>PrWifeEarnsMore</i>	-0.178*** [0.004]	-0.142*** [0.004]	-0.139*** [0.004]	-0.143*** [0.004]	-0.148*** [0.005]	-0.152*** [0.005]
Observations		7,384,176			1,375,121	1,375,121
R-squared	0.097	0.103	0.104	0.145	0.087	0.090
Additional controls:						
Cubic in <i>lnHusbIncome</i>	no	yes	yes	yes	yes	yes
<i>lnMedianWifePotential</i> × <i>lnHusbIncome</i>	no	no	yes	yes	no	no
<i>anyChildren</i>	no	no	no	yes	no	no
Wife's demographic group × Husband's demographic group	no	no	no	yes	no	no
<i>PrWifeEarnsMore AtMarriage</i>	no	no	no	no	no	yes
Vintiles of the wife's and the husband's potential income at marriage	no	no	no	no	no	yes
Marriage duration fixed effects	no	no	no	no	no	yes
Sample restriction	none	none	none	none	2010sub	2010sub

Notes: Data are from the 1970 to 2000 U.S. Census and 2008 to 2011 ACS single-year files. The sample consists of couples where both the wife and the husband are between 18 and 65 years old and the husband is working. The dependent variable *Wife in the labor force* is a binary variable that equals 1 if the wife is in the labor force, 0 otherwise. The key independent variable *PrWifeEarnsMore* is the probability that the wife's income would exceed the husband's if her income were drawn from the distribution of positive earnings in her demographic group. *lnHusbIncome* is the log of husband's income. *lnMedianWifePotential* is the log of the median of the distribution of positive earnings in the wife's demographic group. *anyChildren* is a binary variable that equals 1 if the wife reports having any children, 0 otherwise. *PrWifeEarnsMoreAtMarriage* is the probability that income drawn from the distribution of positive earnings in the wife's demographic group exceeds income drawn from the distribution of positive earnings in the husband's demographic group in the closest census year to the year of their marriage. Marriage duration fixed effects are dummies for the number of years the couple has been married. All regressions include controls for log husband's income, vintiles of the wife's potential income, wife's and husband's education (five categories), wife's and husband's five-year age group, wife's and husband's race (three categories), year, and state fixed effects. Sample restriction "2010sub" indicates a subsample of ACS where the difference between year of marriage and the closest census year is no more than five years and *PrWifeEarnsMoreAtMarriage* could be computed. Standard errors are reported in brackets. ***significant at 1% level, **at 5%, *at 10%.

effects, the wife's and the husband's race, the wife and the husband's age group, and the wife's and the husband's education level. Standard errors are clustered by the wife's demographic group (which pins down the distribution of her potential income). The baseline estimate of β_1 is -0.178 ($p < .01$).

The husband's income might affect the marginal utility of household income non-linearly, so in column (2) we include a cubic polynomial in $\ln \text{HusbIncome}_i$. The estimate of β_1 falls somewhat to -0.142 but remains statistically significant ($p < .01$). Another issue is that the impact of the wife's potential income on her labor supply might interact with the husband's income for reasons that are separate from the couple's concern that she will earn more than he does. To deal with this, in column (3) we add a control for the median of the wife's predicted income (w_i^{50}) interacted with the income of the husband. The estimate of β_1 is stable at -0.139 .

This estimated effect is economically significant: a 10 percentage point increase in the probability that a wife would earn more than her husband reduces the likelihood that she participates in the labor force by around 1.4 percentage points. Put differently, a 1 standard deviation increase (across all years) in the probability that a wife would earn more than her husband reduces the likelihood that she participates in the labor force by about 3.5 percentage points.

The main concern in this subsection is that a woman who is willing to marry a man whose income is below her potential income might have unobservable characteristics that keep her out of the labor force. For example, highly educated women that marry men with lower education and low earnings might be systematic underachievers or systematically lack the confidence to participate in the labor market; such women might be relatively more drawn toward home production and child-rearing activities. We consider two approaches to deal with this concern.

First, we examine the sensitivity of $\hat{\beta}_1$ to the inclusion of other controls. In column (4), we include as controls an indicator variable for whether the wife ever had a child and indicator variables for the full interaction of the wife's demographic group and the husband's demographic group. The inclusion of these additional controls barely affects the estimate of β_1 . The fact that our estimate appears stable across specifications suggests that to the extent that the observable characteristics in our data are

representative of unobservables, the negative value of $\hat{\beta}_1$ is not due to an omitted variable bias (Altonji, Elder, and Taber 2005).³⁶

Second, in the 2010 ACS, we can partly isolate the variation in $PrWifeEarnsMore_i$ that is driven by changes in relative income that took place after the couple got married. Unlike the earlier census data, the ACS contains information on the year the current marriage started. We can thus proxy for relative income of spouses at the time of marriage as follows. For each couple i , let m_i be the census year that is the closest to the year of marriage (we drop couples for whom the difference between the year of marriage and the closest available census year is more than five years). We then construct the distribution of potential income for both the husband $\{h_i^p\}$ and the wife $\{w_i^p\}$ in year m_i , using the same procedure as before. Based on these two distributions, we define a variable $PrWifeEarnsMoreAtMarriage_i = \frac{1}{361} \sum_p \sum_q \mathbf{1}_{\{w_i^p > h_i^q\}}$. This variable captures the probability that based on the couples' demographics, the wife's potential income exceeded that of the husband at the time of marriage.³⁷

We first show that the focus on the new subsample does not make a difference: column (5) replicates the specification from column (2) but using only data on couples for whom we could construct relative earnings at marriage. The effect of $PrWifeEarnsMore_i$ on labor force participation is very similar as in the overall sample.³⁸ Our main specification is in column (6), where we include as controls $PrWifeEarnsMoreAtMarriage_i$ and the vigintiles of the potential earnings for both the husband and the wife at marriage.³⁹ In this specification, we are relying only on the variation that stems from *changes* in the wife's and the husband's relative income since marriage.⁴⁰ Yet the estimate of β_1 is unaffected.

36. The coefficient is also stable if we include the additional controls one by one.

37. Since we estimate the distribution based on every 5th percentile from the 5th to the 95th, there are 19 percentiles for each spouse and thus $19 \times 19 = 361$ overall comparisons in the expression $\sum_p \sum_q \mathbf{1}_{\{w_i^p > h_i^q\}}$.

38. In Online Appendix Table A.6, we consider the specification separately decade by decade. In each period we find a negative and statistically significant effect.

39. We also control nonparametrically for marriage duration in years.

40. Two concerns remain. First, our proxy for relative income at marriage is imperfect as we do not know the couples' actual income at the time. Second, to the extent that couples can predict how their relative income will evolve after marriage, some concerns about selection are still present.

Overall, although we do not have an exogenous source of variation in *PrWifeEarnsMore_i*, the data suggest that married women sometimes stay out of the labor force so as to avoid a situation where they would become the primary breadwinner. Moreover, we find that the effect is stronger among less educated couples, which is consistent with the view that it reflects gender identity norms.⁴¹ Finally, in Section VII, we document a similar effect using panel data: within couples, when a wife outearns her husband, she becomes less likely to be in the labor force the following year.

IV.B. Gap between Potential and Realized Income

Having the wife leave the labor force is a financially costly way to restore traditional gender roles. It would be less costly for the wife to simply reduce her earnings to a level that does not threaten the husband's status as the primary breadwinner. In this subsection, we present evidence for such behavior.

Given a couple i , let $incomeGap_i = \frac{wifeIncome_i - wifePotential_i}{wifePotential_i}$ where $wifePotential_i$ is simply the mean of the distribution of potential earnings for the wife, as defined in the previous subsection. To emphasize the distortions in income for women who do not leave the workforce, we focus on the sample of couples where the woman is working. These results are reported in Table III, which follows exactly the same structure as Table II, but with $incomeGap_i$ rather than $wifeLFP_i$ as the dependent variable.

Column (1) of Table III considers the baseline specification and yields an estimate of β_1 equal to -0.031 ($p < .01$). Including a cubic polynomial of $lnHusbIncome_i$ in column (2) strengthens the effect: $\hat{\beta}_1 = -0.095$ ($p < .01$). A 10 percentage point increase in the probability that a wife would earn more than her husband increases the gap between her actual earnings and her potential earnings by about 1 percentage point. Columns (3)–(6) consider the same robustness checks as in the previous subsection. The effect persists and the estimate is reasonably stable across specifications, though somewhat larger in the recent years.⁴²

41. We consider the specification in column (2) separately for less educated couples (both the wife and the husband have only a high school education or less) and more educated couples (both the wife and the husband have at least some college education). Among the former, the estimate of β_1 is -0.18 (std. err. = 0.007); among the latter, it is -0.129 (std. err. = 0.004).

42. Online Appendix Table A.7 considers the specification separately decade by decade. In each period we find a negative and statistically significant effect.

TABLE III
POTENTIAL RELATIVE INCOME AND WIFE'S REALIZED EARNINGS

	(1)	(2)	(3)	(4)	(5)	(6)
	Dependent variable: <i>incomeGap</i>					
<i>PrWifeEarnsMore</i>	-0.031*** [0.007]	-0.095*** [0.006]	-0.095*** [0.006]	-0.109*** [0.007]	-0.168*** [0.009]	-0.176*** [0.009]
Observations	5,306,682	5,306,682	5,306,664	5,306,664	1,049,793	1,049,793
R-squared	0.004	0.006	0.006	0.050	0.007	0.013
Additional controls:						
Cubic in <i>lnHusbIncome</i>	no	yes	yes	yes	yes	yes
<i>lnMedianWifePotential</i> × <i>lnHusbIncome</i>	no	no	yes	yes	no	no
<i>anyChildren</i>	no	no	no	yes	no	no
Wife's demographic group × Husband's demographic group	no	no	no	yes	no	no
<i>PrWifeEarnsMoreAtMarriage</i>	no	no	no	no	no	yes
Vintiles of the wife's and husband's potential income at marriage	no	no	no	no	no	yes
Marriage duration fixed effects	no	no	no	no	no	yes
Sample restriction	none	none	none	none	2010sub	2010sub

Notes. Data are from the 1970 to 2000 U.S. Census and 2008 to 2011 ACS single-year files. The sample consists of couples where both the wife and the husband are between 18 and 65 years old and the husband is working. The dependent variable *incomeGap* measures the difference between the wife's realized and potential earnings. The key independent variable *PrWifeEarnsMore* is the probability that wife's income would exceed the husband's if her income were drawn from the distribution of positive earnings in the wife's demographic group. *lnHusbIncome* is the log of husband's income. *lnMedianWifePotential* is the log of the median of the distribution of positive earnings in the wife's demographic group. *anyChildren* is a binary variable that equals 1 if the wife reports having any children, 0 otherwise. *PrWifeEarnsMoreAtMarriage* is the probability that income drawn from the distribution of positive earnings in the wife's demographic group exceeds income drawn from the distribution of positive earnings in the husband's demographic group in the closest census year to the year of their marriage. Marriage duration fixed effects are dummies for number of years the couple has been married. All regressions include controls for the log husband's income, vintiles of the wife's potential income, wife's and husband's education (five categories), wife's and husband's five-year age group, wife's and husband's race, year, and state fixed effects. Sample restriction "2010sub" indicates a subsample of ACS where the difference between year of marriage and the closest census year is no more than five years and *PrWifeEarnsMoreAtMarriage* could be computed. Standard errors are reported in brackets. ***, ** and * at 1%, 5%, and 10%.

In Online Appendix Table A.8, we show that part of the effect is moderated through the number of hours the wife works. This channel accounts for about a third of the negative relationship between $incomeGap_i$ and $PrWifeEarnsMore_i$.

In summary, women's labor supply decisions seem to be distorted in situations where there is a threat that they might become the primary breadwinner. In the next section we document some of the costs that arise when the woman does earn more than the husband. The presence of those costs provides a potential rationalization for the labor market distortions that we document here.

V. MARITAL STABILITY AND RELATIVE INCOME

Does relative income affect marital stability? To address this question, we exploit the rich information on marital satisfaction and marital outcomes from the National Survey of Families and Households (NSFH). The NSFH is a nationally representative survey of U.S. households and includes approximately 9,500 households that were followed over three waves from 1988 to 2002. We use data from the first two waves (1987–88 and 1992–94).⁴³ We restrict our analysis to couples where at least one person in the household has positive labor income. Our sample consists of approximately 4,000 married couples.

The NSFH has three questions on marital stability. The first question asks: "Taking things all together, how would you describe your marriage?" Respondents can choose answers on a scale of 1 (very unhappy) to 7 (very happy). Close to 50 percent report being "very happy," however, so we define a binary variable $happyMarriage_i$ that simply indicates whether the answer is "very happy." The second question asks: "During the past year, have you ever thought that your marriage might be in trouble?" We define a binary variable $marriageTrouble_i$ that indicates an affirmative response. The third question asks: "During the past year, have you and your husband/wife discussed the idea of separating?" We define a binary variable $discussSeparation_i$ that indicates the answer is affirmative.

43. We do not use Wave 3 since, due to budgetary constraints, it is based on a different sampling rule that substantially alters the profile of the respondents.

The NSFH also provides information on the wife's and the husband's annual labor income, on the basis of which we define self-explanatory variables $\ln WifeIncome_i$, $\ln HusbandIncome_i$, and $\ln TotalIncome_i$.⁴⁴ For each couple we also compute $relativeIncome_i$, the share of the household income earned by the wife. In Wave 1, the mean of $relativeIncome_i$ is 0.27 and it exceeds $\frac{1}{2}$ in 15 percent of households. We define $wifeEarnsMore_i$ as a binary variable equal to 1 if $relativeIncome_i > \frac{1}{2}$. Summary statistics for the main variables used in the analysis are in Online Appendix Table A.9.

In Table IV, we examine how the relative income within the household affects answers to survey questions about marriage quality. Our baseline specification is a linear probability model

$$Y_i = \beta_0 + \beta_1 \times wifeEarnsMore_i \\ + \beta_2 \times \ln WifeIncome_i + \beta_3 \times \ln HusbandIncome_i \\ + \beta_4 \times \ln TotalIncome_i + \beta_5 \times X_i + \varepsilon_i,$$

where Y_i is the answer to the survey question and X_i represents nonincome controls: region fixed effects,⁴⁵ indicator variables for whether the wife is working, whether the husband is working, the wife and the husband's race and education groups, and a quadratic in the wife's and the husband's age.⁴⁶

We pool the wives' and the husbands' responses (and include as a control an indicator variable for whether the respondent is the wife or the husband). We cluster standard errors at the level of the couple. In Online Appendix Table A.10, we examine the wives' and the husbands' responses separately. In short, we find the same patterns as here for each spouse (with somewhat less consistent precision). One might be tempted to use the difference between the coefficients on the wives' and the husbands' responses to determine whether it is the wife or the husband who dislikes the reversal of traditional gender roles. We suspect, however, that such a comparison is not particularly useful. If, say, the

44. In this and the next section, we set $\ln WifeIncome = 0$ if the wife's income is equal to zero, and in all regressions we include an indicator variable for whether the wife's income was zero. We apply the same procedure for the husband's income.

45. State identifiers are not available in the public-use version of the NSFH.

46. We weight the observations using the couple-level weights. The NSFH provides two sets of weights: a person-level weight and a couple-level weight. Results are similar whether we use no weights, person-level weights, or couple-level weights.

TABLE IV
RELATIVE INCOME AND MARITAL SATISFACTION

	(1)	(2)	(3)	(4)
Panel A: dependent variable: <i>happyMarriage</i>				
<i>wifeEarnsMore</i>	-0.068**	-0.060*	-0.070*	-0.065*
	[0.031]	[0.032]	[0.036]	[0.037]
Observations	7,659	7,659	7,659	7,659
R-squared	0.025	0.026	0.025	0.025
Panel B: dependent variable: <i>marriageTrouble</i>				
<i>wifeEarnsMore</i>	0.082***	0.078***	0.079**	0.086**
	[0.027]	[0.029]	[0.033]	[0.034]
Observations	7,520	7,520	7,520	7,520
R-squared	0.047	0.048	0.047	0.048
Panel C: dependent variable: <i>discussSeparation</i>				
<i>wifeEarnsMore</i>	0.068***	0.064***	0.060**	0.065**
	[0.024]	[0.024]	[0.028]	[0.028]
Observations	7,507	7,507	7,507	7,507
R-squared	0.034	0.034	0.034	0.034
Additional controls:				
Cubic in <i>lnWifeIncome</i> and <i>lnHusbIncome</i>	no	yes	no	no
<i>relativeIncome</i>	no	no	yes	yes
<i> Wife-Husb Income Rank </i>	no	no	no	yes

Notes. The data are from Wave 1 of the National Survey of Family and Households (NSFH). The sample is restricted to couples where both the wife and the husband are between 18 and 65 years old and at least one person in the household has positive income. The sample includes observations from both husbands and wives. The dependent variables, *happyMarriage*, *marriageTrouble*, and *discussSeparation* are binary variables based on respondents' answers about their marriage (details in the text). *relativeIncome* is the share of the household income earned by the wife. *wifeEarnsMore* is an indicator variable for whether *relativeIncome* > 0.5. *lnWifeIncome* and *lnHusbIncome* are the logs of the wife's and husband's income, respectively. *|Wife-Husb Income Rank|* is the absolute difference in the income ranks of husbands and wives in their respective income distributions. All regressions include the log of the wife's income, log of the husband's income, log of the total household income, a quadratic in wife and husband's age, indicator variables for wife and husband's race and education (five categories), region fixed effects, and an indicator variable for whether only the wife is working or only the husband is working. Each regression includes an indicator variable for whether the wife or the husband is the respondent and has standard errors clustered at the level of the couple. All regressions are weighted using the Wave 1 couple weights from NSFH. Robust standard errors are reported in brackets. ***significant at 1%, **at 5%, *at 10%.

husband is initially the one who is unhappy, he may start to behave in ways that make the wife unhappy, perhaps even more so.⁴⁷ Such a possibility echoes Al Roth's Iron Law of Marriage: you cannot be happier than your spouse (Roth 2008).

As column (1) of Table IV shows, spouses tend to be less happy with the marriage, are more likely to report that the

47. Sayer et al. (2011) analyze a postdivorce survey that ascertains which spouse wanted the divorce more, but they do not examine the role of relative income.

marriage is in trouble, and are more likely to have discussed separation if the wife earns more than the husband. In column (2), we add more flexible income controls, namely, cubic polynomials of $\ln WifeIncome_i$ and $\ln HusbandIncome_i$. The estimates of β_1 are mostly unaffected. Since gender identity is more plausibly associated with a prescription that “the husband should earn more than the wife” than with a prescription that “it is better for the wife to earn 20% rather than 30% of the household income,” in column (3) we include $relativeIncome_i$ as a control. This variable does not predict survey responses and the estimates of β_1 are again unaffected. Hence, relative income within a household matters if and only if it makes the wife the primary breadwinner.

Finally, in column (4) we include a control for the absolute difference in income *rank* of the wife and the husband (defined within gender-specific distributions of income). Recall that standard models with assortative matching predict this is the variable that should matter. We again find no impact on $\hat{\beta}_1$: differences in income ranks have no predictive power; what matters is whether the wife earns more than the husband in absolute terms.

The effects are economically significant. In our preferred specification (column (3)), we find that if the wife earns more than the husband, spouses are 7 percentage points (15%) less likely to report that their marriage is very happy, 8 percentage points (32%) more likely to report marital troubles in the past year, and 6 percentage points (46%) more likely to have discussed separating in the past year.

Next we turn away from survey data to the revealed stability of marriage. For each couple in Wave 1 (1987–88), we construct a binary variable $divorced_i$ which is equal to 1 if the couple is separated or divorced when they are reinterviewed in Wave 2 (1992–94).⁴⁸ In Table V, we consider the same specifications as in Table IV but with $divorced_i$ rather than the survey responses as the dependent variable.⁴⁹ Column (1) considers the baseline specification. Column (2) controls more flexibly for the wife’s and the husband’s earnings. In both specifications we find that when

48. One concern is that there may be selective attrition by divorce status. If divorced couples are less likely to remain in the panel, we would underestimate the overall tendency to divorce, but the estimate of our key coefficient would be unaffected. Moreover, we find that there is no relationship between attrition and measures of marital stability. The overall attrition rate is about 10 percent.

49. All of the independent variables are measured in Wave 1. We use Wave 2 only to identify whether the couple is still together.

TABLE V
RELATIVE INCOME AND DIVORCE

	(1)	(2)	(3)	(4)
	Dependent variable: <i>divorced</i>			
<i>wifeEarnsMore</i>	0.062** [0.025]	0.060** [0.026]	0.048 [0.030]	0.051* [0.030]
Observations	3,439	3,439	3,439	3,439
R-squared	0.080	0.086	0.080	0.080
Additional controls:				
Cubic in <i>lnWifeIncome</i> and <i>lnHusbIncome</i>	no	yes	no	no
<i>relativeIncome</i>	no	no	yes	yes
<i> Wife-Husb Income Rank </i>	no	no	no	yes

Notes. The data are from Waves 1 and 2 of the National Survey of Family and Households (NSFH). The sample is restricted to couples where both the wife and the husband are between 18 and 65 years old (in Wave 1) and at least one person in the household has positive income. The dependent variable *divorced* is an indicator for whether the couple is divorced or separated as of Wave 2. *relativeIncome* is the share of the household income earned by the wife. *wifeEarnsMore* is an indicator variable for whether *relativeIncome* > 0.5. *lnWifeIncome* and *lnHusbIncome* are the logs of the wife's and husband's income, respectively. *|Wife-Husb Income Rank|* is the absolute difference in the income ranks of husbands and wives in their respective income distributions. All regressions include the log of the wife's income, log of the husband's income, log of the total household income, a quadratic in wife and husband's age, indicator variables for wife and husband's race and education (five categories), region fixed effects, and an indicator variable for whether only the wife is working or only the husband is working. All regressions are weighted using the Wave 1 couple weights from NSFH. Robust standard errors are reported in brackets. ***significant at 1%, **at 5%, *at 10%.

the wife earns more than the husband, the likelihood of divorce increases by about 6 percentage points ($p < .05$). Since 12 percent of couples in the sample get divorced, this estimate implies that having the wife earn more than the husband increases the likelihood of divorce by 50 percent. In column (3), we include a control for relative income. The estimate decreases slightly to about 5 percentage points and becomes statistically insignificant ($p = .11$). Including a control for the absolute difference in wife and husband's income rank does not change the estimate appreciably (column (4)).⁵⁰ Hence, as before, what seems to matter is not relative income per se or the difference in ranks but whether the norm that the husband should earn more than the wife has been violated.

50. Separation or divorce occurs when the marriage fully breaks down and can be regarded as the end-point of marital instability. Among couples who remain married in both survey waves, we can also examine whether *wifeEarnsMore* in Wave 1 is associated with a deterioration in reported marital stability. We find some suggestive evidence in this direction but most of the point estimates are not statistically significant (see Online Appendix Table A.11).

The main concern with the interpretation of these results is the heterogeneity of the effect by education. When we consider the specification in column (3) separately by couples' education, we find that *wifeEarnsMore* increases divorce only among the *more* educated couples (where both the husband and the wife have at least some college education).

Our results in this section contribute to an existing literature in sociology that examines the relationship between relative income and divorce. The results in this literature are somewhat mixed. For example, Heckert, Nowak, and Snyder (1998), Jalovaara (2003), and Liu and Vikat (2004) report that couples are more likely to divorce if the wife earns more than the husband, while Rogers (2004) argues that divorce is most likely when the couples earn around the same amount of money.⁵¹ In contrast to this literature, we examine marital satisfaction in addition to divorce and, more important, we analyze the impact of relative income conditional on each individual's income (and the total household income.) None of the aforementioned papers considers such specifications.⁵²

VI. HOME PRODUCTION AND RELATIVE INCOME

Traditional gender roles also contain prescriptions about the division of chores within the household. In this section, we explore whether, when the wife earns more than the husband, the couple changes behavior at home so as to alleviate the sense of gender role reversal.

We use data from the American Time Use Survey (ATUS) and the Current Population Survey (CPS), covering the years 2003 to 2011. As in the previous section, we restrict our analysis to couples where at least one person has positive labor income. Following Aguiar and Hurst (2007), we define chores as "core" nonmarket work (meal preparation and cleanup, doing laundry, vacuuming, etc.), time spent "obtaining goods and services" (e.g., grocery shopping), and time spent in "other" home production (home maintenance, vehicle repair, gardening, etc.). Our measure of *childcare* includes primary child care (e.g., changing diapers and feeding the child), educational child care (such as helping a child with

51. Kalmijn, Loeve, and Manting (2007) argue that the impact of relative income on dissolution differs across married and cohabiting couples.

52. Jalovaara (2003) reports divorce rates across all possible combinations of the wife's and the husband's income, but she employs a very rough measure that distinguishes between only five levels of income.

homework), and recreational child care (playing games with the children, taking them to the zoo, etc.). For each individual in the sample, we define *totalNonMarketWork_i* as the weekly number of hours spent on *chores* and *childcare*.

Variables *lnWifeIncome_i*, *lnHusbIncome_i*, and *lnTotIncome_i* are based on the weekly earnings at the main job reported in the CPS or ATUS interviews (see Online Appendix for details).⁵³ Based on these earnings we define *relativeIncome_i* and *wifeEarnsMore_i* as before. Summary statistics are presented in Online Appendix Table A.12. Wives spend an average of 36 hours a week on chores and childcare, compared with 22.5 hours for husbands. Mean relative income is 0.34 and the wife earns more than the husband in 25 percent of the couples.

Ideally, we would like to compare the wife–husband gap in time spent on home production across couples where the wife earns more and those where she does not. Unfortunately, the ATUS/CPS only includes one respondent per household. Thus, to analyze how relative income impacts the division of home production, we focus on the interaction between the impact of gender and the impact of relative income on time use. Specifically, in column (1) of Table VI, we consider the baseline OLS model:

$$\begin{aligned}
 \text{totalNonMarketWork}_i &= \beta_0 + \beta_1 \times \text{female}_i \times \text{wifeEarnsMore}_i \\
 &+ \beta_2 \times \text{female}_i + \beta_3 \times \text{wifeEarnsMore}_i \\
 &+ \beta_4 \times \ln\text{WifeIncome}_i + \beta_5 \times \text{female}_i \times \ln\text{WifeIncome}_i \\
 &+ \beta_6 \times \ln\text{HusbIncome}_i + \beta_7 \times \text{female}_i \times \ln\text{HusbIncome}_i \\
 &+ \beta_8 \times \ln\text{TotIncome}_i + \beta_9 \times \text{female}_i \times \ln\text{TotIncome}_i \\
 &+ \beta_{10} \times X_i + \beta_{11} \times \text{female}_i \times X_i + \varepsilon_i,
 \end{aligned}$$

where X_i includes year, state, and day of the week fixed effects; indicator variables for whether only the wife is working; whether only the husband is working; the wife and the husband's race and education groups; and a quadratic in the wife's and the husband's age. Our coefficient of interest is β_1 . A positive estimate of β_1

53. The reliance on a weekly earnings measure here, which deviates from the annual labor market earnings measures we use in other sections, is due to the lack of annual earnings for the substantive subsample of ATUS respondents who did not complete the March CPS.

TABLE VI
RELATIVE INCOME AND THE GENDER GAP IN NONMARKET WORK

	(1)	(2)	(3)	(4)	(5)	(6)
		Dependent variable: <i>Total nonmarket work (hours per week)</i>				
<i>female</i> × <i>wifeEarnsMore</i>	1.087 [0.740]	1.263* [0.762]	2.183*** [0.782]	2.297*** [0.756]	2.961*** [0.844]	2.757*** [0.998]
<i>wifeEarnsMore</i>	0.460 [0.523]	0.132 [0.544]	-0.031 [0.557]	-0.147 [0.588]	-0.546 [0.600]	-0.439 [0.707]
Observations	37,665	37,665	37,665	37,665	22,390	18,227
R-squared	0.233	0.233	0.234	0.285	0.224	0.375
Additional controls:						
Cubic in <i>lnWifeIncome</i> and <i>lnHusbandIncome</i>	no	yes	yes	yes	yes	yes
<i>relativeIncome</i>	no	no	yes	yes	yes	yes
Children controls	no	no	no	yes	yes	yes
Sample restriction	none	none	none	none	both spouses have positive income	week-day only

Notes: Data source are ATUS/CPS from 2003 to 2011. The sample is restricted to married individuals who are between 18 and 65 years old and whose spouse is also between 18 and 65 years old. The sample is further restricted to couples where both spouses are present at the time of the ATUS interview, and where at least one of the spouses is employed at the time of the ATUS interview. *wifeIncome (HusbandIncome)* is the wife's (husband's) weekly earnings at main job. Earnings are drawn from the ATUS interview for the spouse that completes that interview and from the CPS interview for the other spouse; earnings from the CPS interview are changed to 0 if that spouse is not working at the time of the ATUS interview. *relativeIncome* is the share of the household income earned by the wife. The key independent variable *wifeEarnsMore* is an indicator variable for whether *relativeIncome* < 0.5. "Children controls" include indicator variables for whether the respondent has no children, whether the youngest child is 3 or younger, between 4 and 6, or is older than 6. All regressions include the log of the wife's income, log of the husband's income, log of the total household income, year, state, and day of the week fixed effects, the wife and the husband's race, a quadratic in wife and husband's age, indicator variables for the wife's and the husband's education groups (five categories), and indicator variables for whether only the husband is working, and whether only the wife is working. The specifications also include the interaction of all these controls with an indicator variable for whether the ATUS respondent is female. The estimated coefficients on female ATUS respondent in a regression that includes the same controls as in column (2) but only interacts *wifeEarnsMore* with female ATUS respondents are, respectively: 16.70 [0.26] (nonmarket work + childcare); 11.00 [0.22] (nonmarket work). Each observation is weighted using the ATUS/CPS weight. Standard errors are reported in brackets. ***significant at 1% level, **at 5%, *at 10%.

would indicate that, *ceteris paribus*, in couples where the wife earns more than the husband, she also spends more time on chores and childcare. In the baseline specification, the estimate of β_1 is 1.09 (statistically insignificant). In column (2) we include more flexible cubic polynomial controls for the wife's and the husband's income. The estimate of β_1 is similar at 1.26.

In column (3) we include a control for *relativeIncome_i*. That is particularly important in this context because standard Beckerian forces will lead the wife to do less housework when her contribution to family income increases. With this control, the estimate of β_1 increases to 2.18 and becomes highly significant ($p < .01$): Beckerian forces are present, but once a wife earns more than her husband she starts to compensate for it by spending more time on chores and childcare.⁵⁴

In column (4), we further control for the presence of children of various ages.⁵⁵ The estimate remains largely unchanged. Finally, our results are robust to restricting the sample to couples where both the wife and the husband have strictly positive earnings (column (5)) as well as time-use during weekdays only (column (6)). In Online Appendix Table A.13, we separately analyze the impact on various components of total nonmarket work. In general we find positive coefficients, but only the effects on chores are statistically significant. When we examine the effect by education group, we find a larger estimate among the couples that are more educated but standard errors are much too large to draw any strong inference from this comparison.

In summary, it seems that women who are overly successful in the labor force pay for this success at home to abate the reversal of the traditional gender roles. Akerlof and Kranton (2000) report that women do not undertake less than half of the housework even if they work or earn more than the husband.

54. One potential source of an omitted variable bias would be that in couples that are "more traditional," women are less likely to earn more than their husbands and are more likely to take on a larger share of housework. This force would bias the estimate of β_1 downward, in the opposite direction of our finding. Similarly, one might be concerned that there are unobservable factors that lower the husband's income (below his potential and below the wife's income) and simultaneously lower his ability to do household chores. This is unlikely to be important, however: when the husband's realized income is further below his potential income, he spends more time on chores.

55. Specifically, we add indicator variables for whether there is no child, the youngest child is younger than 3, the youngest child is between 4 and 6, or the youngest child is older than 6.

Our finding is even more striking: the gender gap in nonmarket work is greater when the wife earns more than the husband.

As we mentioned earlier, the idea that individuals modify their behavior so as to fulfill gender roles has been present in the sociology literature at least since West and Zimmerman (1987). This idea was first applied to the division of household chores by Fenstermaker Berk (1985). More recently, Bittman et al. (2003) report that the extent of the wife's housework decreases in relative income when she makes less than the husband, but this relationship reverses when relative income exceeds one half. In contrast, Gupta (2007) and Gupta and Ash (2008) argue that the number of hours the wife spends on chores is solely determined by her level of income, without any regard for her relative income in the household.⁵⁶ Our findings are more in line with Bittman et al. (2003), though our approach varies from theirs on two dimensions: we use a different data set⁵⁷ and, more important, our specification includes controls for each individual's income.⁵⁸ Finally, in the next section we build on the existing evidence by documenting how relative income impacts household production within couples over time.

VII. PANEL DATA EVIDENCE

In this final section, we analyze labor force participation, marital stability, and home production using *within-couple* variation in relative income. We combine data from the family files (1968–2009) and the marital history file (2009) of the PSID (see Online Appendix for details of sample construction). We focus our analysis on couple-year observations where at least one of the spouses has positive earnings in the year(s) where relative earnings measures are used as independent variables. Summary statistics are in Online Appendix Table A.14.

56. Cooke (2006) finds that among U.S. couples where the wife earns more than the husband, the likelihood of divorce is lower if the wife engages in “compensatory behavior”—that is, if she does a greater share of the housework. Greenstein (2000) argues that both men and women adjust their household production to neutralize deviance from traditional gender roles.

57. Bittman et al. (2003) draw on an Australian time-use survey from 1992.

58. Including these controls is particularly important given the results in Gupta (2007) and Gupta and Ash (2008).

VII.A. *Wife's Labor Force Participation*

Our analysis of wives' labor force participation in Section IV was based on an imputed likelihood that the wife would earn more than her husband. The panel setting offers an alternative to this imputation exercise. Because we observe couples over time, we can ask whether an actual realization of the wife earning more than her husband is predictive of the wife subsequently leaving the workforce. Specifically, we estimate a linear probability model with couple fixed effects regressing the wife's labor force participation in year t on whether the wife earned more than the husband in year $t-1$ (controlling for $\ln\text{WifeIncome}$, $\ln\text{HusbandIncome}$, and $\ln\text{TotalIncome}$ in $t-1$ as well as for a number of other covariates).⁵⁹

As shown in Table VII, Panel A, column (1), if a wife earns more than her husband, she is 1.9 percentage points less likely be in the labor force the following year ($p < .01$). Moreover, in results not reported here, we find that if a wife earns more than her husband, she starts to work fewer hours the following year. These results echo Winslow-Bowe's (2006) finding that when the wife earns more than the husband, that situation typically does not persist for more than a year or two.

In column (2), we include cubics of each spouse's income as controls; the estimate is unaffected. When we add *relativeIncome* in $t-1$ as a control in column (3), the effect falls down to 1.3 percentage points but remains highly significant ($p < .01$). In column (4) we add a vector of controls for the presence of children; the estimate is barely affected. Finally, as column (5) shows, the specifications with couple fixed effects yields the same estimate as the cross-sectional relationship within the PSID. This further supports the view that selection is not driving the results in Section IV.

VII.B. *Divorce*

We study the impact of relative income on divorce using similar specifications as in the previous subsection, but we allow for a longer lag structure since divorce may take time to be completed.

59. Namely, year and state fixed effects, an indicator variable for only the wife working in $t-1$, an indicator variable for only the husband working in $t-1$, and quadratic functions of the wife's and the husband's age.

Hence, we regress divorce in year t on *wifeEarnsMore* in both $t - 1$ and $t - 2$.⁶⁰ When we construct the dependent variable, we code the couple as having divorced if they are no longer listed as married and there is no explicit mention of widowhood.⁶¹

Panel B of Table VII shows the results. In all of the specifications, the estimated effect of *wifeEarnsMore* on divorce is positive but typically imprecise. The lack of statistical power is also evident in the fact that the relationship is not significant even in the cross section. With this important caveat in mind, the point estimates are substantial: the likelihood of divorce increases by 25 percent if the wife earned more than her husband two years ago.

VII.C. Home Production

The PSID asks the following question of both husbands and wives: "About how much time do you spend on housework in an average week—I mean time spent cooking, cleaning, and other work around the house?" Relative to ATUS, this question is more vague about the exact content of the home production activities, but the PSID has two important advantages: (i) both spouses are asked to answer the question, which allows us to directly measure the gender gap in home production within a household, and (ii) the panel nature of the data allows us to examine how changes in relative income affect this gap.

We estimate the exact same specifications as in subsection VII.A but with the gap in housework as the dependent variable. As Panel C of Table VII shows, within a given couple, when the wife starts to earn more than the husband she takes on relatively more housework. In Online Appendix Table A.15, we show that the effect is driven by the wife doing more and the husband doing less housework, though the former channel plays a larger role.

60. Accordingly, we control for each spouse's income and total income in $t - 2$ as well as $t - 1$, and for year and state fixed effects, indicator variables for only the wife working in $t - 1$ and $t - 2$, indicator variables for only the husband working in $t - 1$ and $t - 2$, and quadratic functions of the wife's and the husband's age.

61. See the Online Appendix for a detailed discussion of how we construct the divorce variable. Our results are similar if we alternatively code the couple as having divorced only if divorce or separation was explicitly reported. We prefer our measure since it yields an average divorce rate of 2% per year, roughly similar to that observed in other data sets. (Under the alternative measure, average divorce rate is only 0.8% per year.)

TABLE VII
RELATIVE INCOME AND LABOR FORCE PARTICIPATION, DIVORCE, AND GENDER GAP IN HOUSEWORK, PANEL DATA EVIDENCE

	(1)	(2)	(3)	(4)	(5)
Panel A: relative income and labor force participation					
<i>wifeEarnsMore</i> in $t-1$	-0.019*** [0.004]	-0.018*** [0.004]	-0.013*** [0.004]	-0.012*** [0.004]	-0.014*** [0.003]
Observations	110,123	110,123	110,123	110,123	110,123
R-squared	0.718	0.720	0.720	0.721	0.662
Panel B: relative income and marital stability					
	Dependent variable: <i>Divorced</i> in t				
<i>wifeEarnsMore</i> in $t-1$	0.003 [0.002]	0.003 [0.003]	0.004 [0.003]	0.004 [0.003]	0.000 [0.003]
<i>wifeEarnsMore</i> in $t-2$	0.005* [0.003]	0.005* [0.003]	0.005 [0.003]	0.005 [0.003]	0.002 [0.003]
Observations	72,169	72,169	72,169	72,169	72,169
R-squared	0.418	0.418	0.418	0.418	0.013

TABLE VII
(CONTINUED)

	(1)	(2)	(3)	(4)	(5)
Panel C: relative income and gender gap in housework					
<i>wifeEarningsMore</i> in <i>t</i>	0.059 [0.250]	0.464* [0.261]	1.048*** [0.275]	0.904*** [0.273]	1.141*** [0.310]
Observations	67,603	67,603	67,603	67,603	67,603
R-squared	0.586	0.587	0.587	0.595	0.376
Additional controls:					
Couple fixed effects	yes	yes	yes	yes	no
Cubic in <i>lnWifeIncome</i> and <i>lnHusbandIncome</i>	no	yes	yes	yes	yes
<i>relativeIncome</i>	no	no	yes	yes	yes
Children controls	no	no	no	yes	yes

Notes. The data are from the PSID, 1968 to 2009. See Online Appendix for sample construction for each panel. The sample is restricted to couples where both the wife and the husband are between 18 and 65 years of age and where at least one of the spouses is working in the year income and relative income measures are used in the regression. *Wife in the labor force* in *t* is a dummy variable that equals 1 if the wife is in the labor force at time *t*, 0 otherwise. *Divorced* in *t* is a dummy variable that equals 1 if the couple is divorced or separated at time *t*, 0 otherwise. *Wife housework* in *t* - *Husband housework* in *t* is the difference between the wife's and the husband's weekly hours of home production at time *t*. *Wife housework* in *t* is the wife's weekly hours of home production at time *t*. *Husband housework* in *t* is the husband's weekly hours of home production at time *t*. *wifeEarningsMore* in *t* is an indicator variable that equals 1 if *relativeIncome* > 0.5 at time *t*. *relativeIncome* is the share of the household income earned by the wife. *lnWifeIncome* is the log of the wife's labor market earnings; *lnHusbandIncome* is the log of the husband's labor market earnings. "Children controls" include indicator variables for whether the respondent has no children, whether the youngest child is 3 or younger, between 4 and 6, older than 6. All regressions include the log of the wife's income, the log of the husband's income, the log of the couple's income, an indicator for whether only the wife is working, an indicator for whether only the husband is working, a quadratic in wife's and husband's age, year fixed effects and state fixed effects. All independent variables related to income in a regression (including whether only the wife is working or whether only the husband is working) are measured in the same year or years as *WifeEarningsMore* in that regression. Standard errors are clustered at the couple level and are reported in brackets. ***,**,*significant at 1% level, **at 5%, *at 10%.

VIII. CONCLUSION

We show that the social norm “a man should earn more than his wife” influences the distribution of relative income within households, the patterns of marriage and divorce, women’s labor supply, and the division of home production activities between husbands and wives.

By definition, the norm that we focus on here would be of no practical relevance in a world where a woman could never earn more than her husband. The relative gains in women’s labor market opportunities over the past half century have thus made this aspect of gender identity increasingly relevant. We suspect that these changes were particularly important because they happened quickly in comparison to the slow-moving social norms and concepts of gender.

While our empirical work focuses on the United States, rapid gains in women’s labor market opportunities are not unique to this country. Even more rapid changes have taken place in developed Asian countries, such as Korea and Japan. At the same time, these Asian countries have experienced large declines in marriage rates and fertility among educated women. As suggested by Hwang (2012), the interaction of economic growth and inter-generational transmission of gender attitudes might play an important part in these developments.

In future work, we would like to better understand the long-run determinants of gender identity. While the evidence in this article suggests that the behavioral prescription that “a man should earn more than his wife” helps explain economic and social outcomes even in the most recent decade, this does not imply that this prescription is as strong today as it was in the past. How are gender identity norms evolving in the face of market forces that are making those norms increasingly costly?

UNIVERSITY OF CHICAGO
UNIVERSITY OF CHICAGO
NATIONAL UNIVERSITY OF SINGAPORE

SUPPLEMENTARY MATERIAL

An Online Appendix for this article can be found at QJE online (qje.oxfordjournal.org).

REFERENCES

- Aguiar, Mark, and Erik Hurst, "Measuring Trends in Leisure: The Allocation of Time Over Five Decades," *Quarterly Journal of Economics*, 122 (2007), 969–1006.
- Aizer, Anna, "The Gender Wage Gap and Domestic Violence," *American Economic Review*, 100 (2010), 1847–1859.
- Akerlof, George A., and Rachel E. Kranton, "Economics and Identity," *Quarterly Journal of Economics*, 115 (2000), 715–753.
- , *Identity Economics: How Our Identities Affect Our Work, Wages, and Well-Being* (Princeton, NJ: Princeton University Press, 2010).
- Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber, "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools," *Journal of Political Economy*, 113 (2005), 151–184.
- Autor, David, "U.S. Labor Market Challenges over the Longer Term," Paper prepared for the Federal Reserve Board of Governors (2010).
- Baker, Matthew J., and Joyce P. Jacobsen, "Marriage, Specialization, and the Gender Division of Labor," *Journal of Labor Economics*, 25 (2007), 763–793.
- Bartik, Timothy, *Who Benefits From State and Local Economic Development Policies?* (Kalamazoo, MI: W. E. Upjohn Institute for Employment Research, 1991).
- Becker, Gary S., "A Theory of Marriage: Part I," *Journal of Political Economy*, 81 (1973), 813–846.
- , "A Theory of Marriage: Part II," *Journal of Political Economy*, 82 (1974), S11–S26.
- Bertrand, Marianne, "New Perspectives on Gender," in *Handbook of Labor Economics*, O. Ashenfelter, and D. Card, eds. Amsterdam: North-Holland, 1545–1592.
- Bittman, Michael, Paula England, Liana Sayer, Nancy Folbre, and George Matheson, "When Does Gender Trump Money? Bargaining and Time in Household Work," *American Journal of Sociology*, 109 (2003), 186–214.
- Black, Sandra E., and Elizabeth Brainerd, "The Impact of Globalization on Gender Discrimination," *Industrial and Labor Relations Review*, 57 (2004), 540–559.
- Black, Sandra E., and Chinhui Juhn, "The Rise of Female Professionals: Women's Response to Rising Skill Demand," *American Economic Review*, 90 (2000), 450–455.
- Black, Sandra E., and Philip E. Strahan, "The Division of Spoils: Rent-Sharing and Discrimination in a Regulated Industry," *American Economic Review*, 91 (2001), 814–831.
- Blau, Francine D., and Lawrence M. Kahn, "The U.S. Gender Pay Gap in the 1990s: Slowing Convergence," *Industrial and Labor Relations Review*, 60 (2006), 45–66.
- Charles, Kerwin K., and Ming Ching Luoh, "Male Incarceration, the Marriage Market, and Female Outcomes," *Review of Economics and Statistics*, 92 (2010), 614–627.
- Cooke, Lynn Prince, "'Doing' Gender in Context: Household Bargaining and Risk of Divorce in Germany and the United States," *American Journal of Sociology*, 112 (2006), 442–472.
- Fenstermaker Berk, Sarah, *The Gender Factory: The Apportionment of Work in American Households* (New York: Plenum Press, 1985).
- Fernandez, Raquel, Alessandra Fogli, and Claudia Olivetti, "Mothers and Sons: Preference Development and Female Labor Force Dynamics," *Quarterly Journal of Economics*, 119 (2004), 1249–1299.
- Fisman, Raymond, Sheena S. Iyengar, Emir Kamenica, and Itamar Simonson, "Gender Differences in Mate Selection: Evidence from a Speed Dating Experiment," *Quarterly Journal of Economics*, 121 (2006), 673–697.
- , "Racial Preferences in Dating," *Review of Economic Studies*, 75 (2008), 117–132.
- Fortin, Nicole, "Gender Role Attitudes and Women's Labour Market Outcomes across OECD Countries," *Oxford Review of Economic Policy*, 21 (2005), 416–438.
- Fortin, Nicole, "Gender Role Attitudes and Women's Labor Market Participation: Opting Out, AIDS, and the Persistent Appeal of Housewifery," University of British Columbia Working Paper, 2009.

- Goldin, Claudia, and Lawrence F. Katz, "The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions," *Journal of Political Economy*, 110 (2002), 730–770.
- Goldin, Claudia, Lawrence F. Katz, and Ilyana Kuziemko, "The Homecoming of American College Women: The Reversal of the Gender Gap in College," *Journal of Economic Perspectives*, 20 (2006), 133–156.
- Greenstein, Theodore, "Economic Dependence, Gender, and the Division of Labor in the Home: A Replication and Extension," *Journal of Marriage and Family*, 62 (2000), 322–335.
- Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu, "Engines of Liberalization," *Review of Economic Studies*, 72 (2005), 109–133.
- Gupta, Sanjiv, "Autonomy, Dependence, or Display," *Journal of Marriage and Family*, 69 (2007), 399–417.
- Gupta, Sanjiv, and Michael Ash, "Whose Money, Whose Time? A Nonparametric Approach to Modeling Time Spent on Housework in the United States," *Feminist Economics*, 14 (2008), 93–120.
- Heckert, D. Alex, Thomas C. Nowak, and Kay A. Snyder, "The Impact of Husbands' and Wives' Relative Earnings on Marital Disruption," *Journal of Marriage and the Family*, 60 (1998), 690–703.
- Hochschild, Arlie, and Anne Machung, *The Second Shift* (New York: Viking, 1989).
- Hwang, Jisoo "Housewife, 'Gold Miss,' and Equal: The Evolution of Educated Women's Role in Asia and the U.S.," Working Paper, Harvard University, 2012.
- Jalovaara, Marika, "The Joint Effects of Marriage Partners' Socioeconomic Positions on the Risk of Divorce," *Demography*, 40 (2003), 67–81.
- Kalmijn, Matthijs, Anneke Loeve, and Dorien Manting, "Income Dynamics in Couples and the Dissolution of Marriage and Cohabitation," *Demography*, 44 (2007), 159–179.
- Liu, Guiping, and Andres Vikat, "Does Divorce Risk Depend on Spouses' Relative Income? A Register-Based Study of First Marriages in Sweden in 1981–1998," MPIDR Working Paper WP 2004-010, 2004.
- Loughran, David S., "The Effect of Male Wage Inequality on Female Age at First Marriage," *Review of Economics and Statistics*, 84 (2002), 237–251.
- Lundberg, Shelly, and Robert A. Pollak, "Bargaining and Distribution in Marriage," *Journal of Economic Perspectives*, 10 (1996), 139–158.
- McCrary, Justin, "Testing for Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test," *Journal of Econometrics*, 142 (2008), 698–714.
- Oppenheimer, Valerie K., "Women's Employment and the Gain to Marriage: The Specialization and Trading Model," *Annual Review of Sociology*, 23 (1997), 431–453.
- Pierce, Lamar, Michael S. Dahl, and Jimmi Nielsen, "In Sickness and in Wealth: Psychological and Sexual Costs of Income Comparison in Marriage," *Personality and Social Psychology Bulletin*, 39 (2013), 359–374.
- Rogers, Stacy J., "Dollars, Dependency, and Divorce: Four Perspectives on the Role of Wives' Income," *Journal of Marriage and Family*, 66 (2004), 59–74.
- Roth, Alvin E., "Deferred Acceptance Algorithms: History, Theory, Practice, and Open Questions," *International Journal of Game Theory*, 36 (2008), 537–569.
- Sayer, Liana C., Paula England, Paul Allison, and Nicole Kangas, "She Left, He Left: How Employment and Satisfaction Affect Men's and Women's Decisions to Leave Marriage," *American Journal of Sociology*, 116 (2011), 1982–2018.
- Stuart, Colleen H., Sue Moon, and Tiziana Casciaro, "The Oscar Curse: Status Dynamics and Gender Differences in Marital Survival," Working Paper, Carnegie Mellon University, 2011.
- Weinberg, Bruce A., "Computer Use and the Demand for Female Workers," *Industrial and Labor Relations Review*, 53 (2000), 290–308.
- Weiss, Yoram, "The Formation and Dissolution of Families: Why Marry? Who Marries Whom?" in *Handbook of Population and Family Economics*, M. R. Rosenzweig, and O. Stark, eds. Amsterdam: North-Holland, 81–123.
- West, Candace, and Don H. Zimmerman, "Doing Gender," *Gender & Society*, 1 (1987), 125–151.
- Winslow-Bowe, Sarah, "The Persistence of Wives' Income Advantage," *Journal of Marriage and Family*, 68 (2006), 824–842.