LOWERING STANDARDS TO WED? SPOUSE QUALITY, MARRIAGE, AND LABOR MARKET RESPONSES TO THE GENDER WAGE GAP

Na'ama Shenhav*

Abstract—This paper examines the effect of the female-to-male wage ratio, "relative wage," on women's spouse quality, marriage, and labor supply over three decades. Exploiting task-based demand shifts as a shock to relative pay, I find that a higher relative wage (a) increases the quality of women's mates, as measured by higher spousal education; (b) reduces marriage without substitution to cohabitation; and (c) raises women's hours of work. These effects are consistent with a model in which a higher relative wage increases the minimum nonpecuniary benefits ("quality") women require from a spouse and therefore reduce marriage among low-quality husbands.

I. Introduction

OVER the past three decades, the share of women never married before age 44 rose to 34% from 16%, and the share of births to unmarried women increased to 40% from 20%.[1] This decline in marriage has attracted significant attention; yet, to date, work in this area has neither addressed which marriages did not form or the labor supply implications of the coinciding reduction in the gender wage gap. Both of these are critical for evaluating the social consequences of the marriage decline.

A central concern is that the increasing share of single-mother households adversely affected child outcomes (Autor et al., 2015; Bertrand & Pan, 2013; Kearney & Levine, 2017). However, the ramifications for children depend immensely on the characteristics of the spouses that mothers would have married and the net effect on parental investments.

In the backdrop of these changes in marriage, labor market opportunities for women relative to men showed a marked improvement from 1980 to 2010. Foremost, women's wages relative to men's appreciated by an unprecedented 20%, depicted in figure 1.[2] Paralleling this, women's relative hours worked and share of family income also increased.[3] Thus, it could be that the decline in marriage is simply a by-product of the greater labor productivity of women, which could have improved the living standards of households.

In this paper, I study the effect of the closing of the gender wage gap on women's spouse quality, marriage, and labor market outcomes. I begin with an illustrative model of marriage built on basic economic principles of specialization as in Becker (1973), which generates a set of predictions that guide the empirical analysis: a more equal female-to-male wage ratio should (a) reduce a woman's propensity to marry; (b) particularly a "low-quality" husband, based on his nonpecuniary qualities; (c) raise the average husband quality among married women; and (d) raise women's hours of work. This derives from the fact that women experience fewer pecuniary gains to marriage under a more equal female-to-male wage ratio, and, hence, increase the minimum nonpecuniary benefits that they require from a husband. In aggregate, then, when the gender wage gap decreases, the marriage rate declines, average spouse quality improves, and women's hours of work conditional on working increase.

Since wages are not randomly assigned, empirically I examine the impact of the ratio of the potential female to the potential male wage in the marriage market, or, simply, the "relative wage." I generate a proxy for the relative wage using a novel application of the Bartik approach (Bartik, 1991), which leverages two sources of variation: differential specialization in occupations and industries across sexes and marriage markets and growth in wages across occupations and industries. I show that the addition of occupational variation, which is distinct from earlier work, gives the proxy additional predictive power, consistent with the disparate impacts of computerization across occupations documented during this period.[4] In support of the identifying assumption, I show that the relative wage is not correlated with pre-1980 trends in these outcomes or the average wage in the market and that the potential wage for each sex is not predictive of opposite-sex wages, which rules out many types of confounding shocks in the market.

I identify four key empirical patterns. First, I show that a higher relative female-to-male wage improves husband quality among married couples in terms of observable characteristics. A 10% increase in the relative wage—roughly the variation that I leverage over this period—leads to a 16% increase in the share of women married to a higher-educated spouse. Women are also less likely to be married to an older spouse, which recent surveys find is associated with lower marital satisfaction (Lee & McKinnish, 2017). Further, descriptive evidence suggests a higher relative income is associated with greater marital happiness for couples with a dominant male earner, as in the model.

[1] See, e.g., Lundberg, Pollak, and Stearns (2016) for a recent summary of marriage and cohabitation patterns.

[2] Appendix figure A.1 shows these patterns remain across narrowly defined education groups.

[3] See, e.g., Glynn (2014), Knowles (2013), and Bertrand, Kamenica, and Pan (2015).

[4] See, e.g., Black and Spitz-Oener (2010); Beaudry and Lewis (2014); Weinberg (2000).

FIGURE 1.—WOMEN'S HOURLY WAGE AS A FRACTION OF MEN'S WAGE, AND MARRIAGE RATES

This figure depicts the relative wage, defined as the ratio of average female hourly wage to average male hourly wage, together with marriage rates. Average hourly wages are calculated from the March Annual Demographic files (1962–2012) as annual earnings divided by total hours worked. Sample for wage calculation includes individuals age 18 to 64 with positive hours worked and positive earned income; for marriage, it includes women 22 to 44 years old.

Second, I find that a higher relative wage reduces marriage through a meaningful decline in first marriages and increased divorce. A 10% increase in the relative wage leads to a 3.1 percentage point (p.p.) increase in the share of never-married women and a 1.7 p.p. increase in the share of divorced women. The data do not allow me to distinguish between delay of first marriage and a permanent opting out of marriage; however, suggestively, I find that the first-marriage decisions of low-educated and younger women are particularly responsive to the relative wage, which I interpret as delay of marriage among all young women and possible opting out for low-educated women. These patterns align with the evolving marriage norms among these groups (Lundberg, Pollak, & Stearns, 2016; Edin & Kefalas, 2005). The estimates can explain 20% of the decline in marriage between 1980 and 2010.

Third, I examine whether a higher relative wage affects whether women live with a romantic partner. I find that 65% of women who do not marry under a higher relative wage (and would have lived with their spouse) instead opt to live with a female roommate or live alone. Further, though 23% of these women live with a single man, I do not find a material increase in reported cohabitation, suggesting these may be largely nonromantic living arrangements.

Fourth, I find that a higher relative wage increases women's labor supply, even once I condition on the female wage. A 10% increase in the relative wage induces an additional 1 hour of work per week conditional on working. However, I find no effect on employment or labor force participation. In total, the appreciation of the relative wage causes a 5% increase in women's annual earnings and a 7% rise in women's share of household income.

The results provide a collage of evidence consistent with Beckerian principles, as in the model, and are not consistent with several alternative explanations. Since women do not remain living with their partners, it seems unlikely that women are trying to avoid the costs of marriage. Rather, this aligns with a perception that the marginal spouse is undesirable outside of marriage. Additionally, while the relative wage predicts an increases in the probability that women earn more than men, the effects through the latter "aversion to women earning more than men" channel (Bertrand, Kamenica, & Pan, hereafter BKP) account for the minority of the decline in marriage and none of the rise in spouse quality. I also rule out that the effects are only concentrated in markets with declining men's wages or a high share manufacturing and show that female wages are also influential for marriage decisions. Nevertheless, I find the largest and most precisely-estimated effects for markets where men's wage growth was in the lowest quartile, suggesting that the decline in men's real wages and, hence, "marriageable men," plays an important role in the responses to the relative wage that I document (Autor, Dorn, & Hanson, 2018).

This paper relates to several literatures. First, it adds to the set of studies on the causal determinants of spouse quality, which has thus far focused exclusively on the role of the balance of sexes in the market (Charles & Luoh, 2010; Angrist, 2002; Abramitzky, Delavande, & Vasconcelos, 2011; Ong, Yang, & Zhang, 2020) or individual education (McCrary &

Royer, 2011; Geruso & Royer, 2018). I show that labor market opportunities have a significant impact on partner choice and argue that selection out of less desirable marriages is an important channel for this effect.

Second, it connects to a growing body of work that quantifies the impact of men and women's employment opportunities on marriage using quasi-experimental methods, but has not yet addressed spouse quality (Blau, Kahn, & Waldfogel, 2000), including some concurrent work (Autor et al., 2018; Kearney & Wilson, 2018). Blau et al. (2000) and Autor et al. (2018) find that increasing demand in industries in which women (men) are concentrated leads to reductions (increases) in marriage, while Kearney and Wilson (2018) find no effect of increased demand in the male-dominated fracking industry on marriage. This paper provides the first evidence on spouse quality, grounding the analysis in a conceptual framework and focusing on the impact of wages explicitly.

Moreover, different from Autor et al. (2018) and Kearney and Wilson (2018), who focus on manufacturing and resource extraction, respectively, I obtain variation from all industries and occupations and track outcomes during the period with the greatest growth in the relative wage. This allows me to leverage all combinations of positive and negative wage growth for men and women. I also use up-to-date empirical methods and provide novel evidence of the credibility and robustness of my estimates, improving upon Blau et al. (2000).

It also closely relates to recent work by Bertrand et al. (2015), hereafter BKP, studying the interaction of gender norms with the relative wage in the household: the aversion to women earning more than men. This paper complements BKP in asking about the effects of the relative wage more broadly and the entire distribution of wages to produce the first estimates that are generalizable to any change in the gender wage gap. I also provide new estimates of impacts on spouse quality.

Further, I provide two pieces of evidence distinguishing Beckerian forces from the aversion channel in BKP. First, for women's hours of work and marital satisfaction, I find that the relative wage has the opposite effect from the aversion channel. Second, I perform bounding exercises and directly control for this aversion to quantify the role of each mechanism. I conclude that the aversion channel explains the minority of the impacts on marriage and none of the impact on spouse quality.

Finally, this paper is linked to the small but important literature showing that a higher relative wage can improve the quality of interactions between men and women, most notably by reducing the incidence of emotional and physical abuse against women (Aizer, 2010; Munyo & Rossi, 2015). These studies often emphasize increases in household bargaining as a mechanism. My findings offer two additional channels for these effects: a decline in relationships with low-quality spouses and greater financial independence of women, which can be a source of power against potential spouses (Edin & Kefalas, 2005).

II. Hypothesis Development

To develop hypotheses regarding the effect of the gender pay gap on spouse quality and women's employment, I draw on the basic economic principles of a Beckerian-type model (Becker, 1973). In appendix A.1, I present a model in the spirit of Bertrand et al. (2018) that formalizes these principles. In this section, I provide simple intuition for these principles and summarize the resulting predictions that I take to the data.

My predictions focus on the decisions of single heterosexual women. In deciding whether to marry a potential spouse, women weigh the extrinsic and intrinsic benefits of marriage against the value of remaining single. The extrinsic benefit of marriage consists of the difference between the income of a married household and income of a single woman. This is assumed to be positive since men earn a higher wage than women (despite the fact that married women are expected to work fewer hours than single women). Intuitively, these pecuniary gains from marriage grow as the wages of potential husbands increase and wane as women's own potential wages increase.

The intrinsic benefits of marriage are two-fold. First, since married women work fewer hours than single women, they also spend more time child rearing. This corresponds to a higher quantity of valuable maternal investments enjoyed by the household, which brings positive gains to marriage. Since child rearing is a substitute to work, this benefit is decreasing with women's wages. Second, husbands provide wives with "spouse quality" (and vice versa). This is a subjective, match-specific measure, which includes all nonpecuniary qualities that a woman considers relevant for marriage and can be positive or negative. This could include the strength of a potential husband's preference for paternal investments in children, his sensitivity to his partner and others, or more minor traits such as timeliness and attention to detail. Women are willing to marry if and only if the spouse quality of a potential husband is at least as high as the negative of the sum of the other gains—the threshold spouse quality.

In this framework, when the relative wage increases, holding the average wage constant, the pecuniary and child-rearing gains to marriage decline, and so women's threshold for acceptable husband quality rises.[5] Consequently, my first hypothesis is that average husband quality among married women improves. At the same time, fewer potential husbands surpass the now-higher threshold for husband quality.

In contrast, the implications of a higher relative wage for acceptable wife quality are ambiguous. This is because potential husbands experience two countervailing effects from an increase in women's wages. On the one hand, increases in the wage of potential wives imply that men now experience a greater extrinsic benefit from marriage, since pooled income rises. This raises the gains to marriage and reduces the

[5]Assuming that the average wage is constant does not change any qualitative predictions; but streamlines the discussion, since it implies that both male and female wages are affected by the change in the relative wage. It also matches the empirical specification.

threshold for acceptable wife quality. On the other hand, since a higher wage implies that wives spend less time in child rearing, this reduces the marital gains that come from maternal time with children. This channel tends to push men's threshold for acceptable wife quality upward. It is thus unclear whether more or fewer potential wives will surpass the resulting threshold for wife quality and similarly unclear whether wife quality will increase or decrease.[6]

As a result of these potentially conflicting changes in threshold spouse quality for men and women, the effects of a change in the relative wage on the number of new marriages formed is indeterminant. However, marriage declines with greater certainty in the presence of frictions or asymmetric preferences in the household. For example, if husbands experience some discomfort from wives' careers, then the utility gains from a rise in pooled income would be smaller than the loss in maternal time (see appendix A.1 for details). Recent research documenting marital conflict when women's wages surpass men's suggests that this may be an appropriate assumption in the U.S. context (BKP). A decline in the bargaining power of husbands would similarly attenuate the gains from pooled income. Finally, marriage declines unambiguously if the change in the relative wage is caused by a reduction in the male wage alone—that is, the average wage declines alongside the rise in the relative wage—since this eliminates the uncertainty created by opposing the effects of women's wages.

A third hypothesis is that the marginal marriage, which does not take place when marriage declines following an increase in the relative wage, involves a potential husband with negative spouse quality. This follows from the fact that the pecuniary and child-rearing gains of marriage are positive. Given this, it also seems unlikely that women would remain with this match in the absence of the compensating gains of marriage. Empirically, this implies that cohabitation would also be an unlikely substitute to marriage for these matches.

Fourth, women are expected to work more hours under a higher relative wage. This is because (i) more women will be induced to be single and therefore will be more likely to work full-time work, and (ii) among married women, the higher wage will make it optimal to work more hours. Women's employment may or may not increase, depending the density of would-be-nonemployed wives around the spouse quality threshold.

An extension of these principles is that a higher relative wage will also trigger additional divorces, despite the fact that a higher bar is applied for screening potential spouses. This is because marriages formed under a lower-threshold spouse quality will no longer meet the higher spouse quality. Further, since divorce decisions are based on realized spouse quality after marriage, while screening occurs based on perceived spouse quality prior to marriage, additional couples will be-

come marginal to changes in the relative wage over time. Both of these effects can be expected to be stronger for women who have been married for a longer time, all else equal.

III. Data and Background on the Gender Wage Gap

To create my proxy for potential wages, I use the IPUMS 1970 (1%) Census and the 1980–2011 March Current Population Surveys (King et al., 2010; Ruggles et al., 2010). I restrict the samples to working-age individuals between 18 and 64 years old who have positive reported earnings for the previous year and are not self-employed. The roughly 1 million working-age individuals in the 1970 Census allow me to establish detailed employment shares in each marriage market prior to my period of analysis. I observe approximately 60,000 households in the March CPS that I use to obtain annual information on average hourly wages for each industry and occupation.[7] Having annual wages allows me to test for delayed responses to once-lagged wages; however, the results are nearly identical when I use the decadal wages in the Census (see appendix A.4.2.) See appendix A.3 for more detail on the data construction.

For my outcomes of interest, I turn to the IPUMS 1980–2000 decennial Censuses and the 2010 American Community Survey (ACS) (Ruggles et al., 2010). I observe current marital status, cohabitation (reported from 1990 on by the head of household), employment, and wages for individuals in the sample, as well as spousal age and education and total income in the household. I focus on a sample of almost 5 million women ages 22 to 44, who are likely to be on the margin of marriage. This age restriction also ensures that I can reliably classify whether an individual completed some college or not, since the hazard for completing some college is significantly lower after age 22, and that all women in the sample are of child-bearing age (Bailey, 2006). I include individuals who identify as one of three race-ethnicity groups—white non-Hispanic, black non-Hispanic—and Hispanic; and classify women into two education levels—less than or equal to high school and at least some college. Appendix table A.1 presents descriptive statistics for the sample.

A. Wage Convergence across Sexes

Although many factors contributed to wage growth during this period, a number of recent studies document that technology-induced demand shifts account for the majority of the closing gender wage gap.[8] In particular, the wave of computerization beginning in 1980 reoriented demand for

---

[6]Because of this ambiguity and limited measures of husbands spouse quality, discussed in section VB, I place less emphasis on the predicted effects of the relative wage for wife quality.

[7]The May CPS/Outgoing Rotation Group (MORG) also provides annual information but had a notable redesign in 1994, which makes its measures less comparable over time (Autor, Katz, & Kearney, 2008).

[8]Other factors include the rise in women's educational attainment and workplace experience, the introduction of birth control, the rise in females in professional and managerial occupations, reductions in discrimination, increasingly positive selection into the workforce, and changes in demand. See Blau and Kahn (2016) for discussions of the relative importance of these factors.

FIGURE 2.—CHANGE IN FEMALE, MALE ln WAGE FROM 1980

This figure depicts the growth in log female hourly wage and log male hourly wage relative to 1980. Average hourly wages are calculated from the March Annual Demographic files (1962–2012) as annual earnings divided by total hours worked. Data include individuals age 18 to 64 with positive hours worked and positive earned income.

occupations based on their complementarity with technology, increasing the need for workers in cognitive- and people-oriented occupations while incrementally eliminating physically intensive occupations (Autor, Levy, & Murnane, 2003; Bacolod & Blum, 2010; Autor & Price, 2013; Deming, 2017). These shifts favored women due to their historical specialization and innate advantage in the former category of work (Beaudry & Lewis, 2014; Weinberg, 2000). The sharp break from previous trends comes across in a simple plot of wage growth by skill group and gender, as in figure 2. It also highlights the particular significance of 1980 for the acceleration of wages of high-skilled women and the depreciation of the wages of low-skilled men (Autor et al., 2008; Katz & Murphy, 1992; Blau & Kahn, 1997).

Given the role of mechanization in this convergence, it is interesting to observe how the growth in the relative wage varies across states with distinct industrial concentrations. Figure 3 depicts the geographic variation in the change in the log relative wage from 1980–2010. The convergence of male and female wages ranges from 5% to 30%.[9] The majority of states experienced convergence above 18%, with a higher rate of convergence seen in the Rust Belt, the center of American manufacturing, and lower convergence primarily in the South and New England. Employment of women in the health sector, which had rapidly growing wages over this period, may have also influenced this pattern. This could have contributed to the high rates of convergence in Min-

nesota, for example, where 15% of low-skilled women were employed as doctors, nurses, or health assistants in the 1970 Census.

In the next section, I return to these occupational differences across sexes and geographic areas and the ensuing wage convergence as an important building block for my empirical strategy.

IV. Estimation Strategy

Conceptually, I would like to estimate how a woman's outcomes change when the log of the ratio between her potential wage and the potential wage of her likely spouses—the log relative wage in her marriage market—rises. I focus on the impact of potential wages, not actual wages, because they have been argued to be the relevant measure for marriage decisions (Pollak, 2005). For this reason, I follow BKP and Aizer (2010) in running reduced-form regressions of outcomes on a proxy for potential wages, rather than instrument observed wages with the proxy, which allows for potential wages to have an impact on marriage decisions through multiple channels, such as through higher bargaining power as well as through a higher realized wage. To implement this, I follow patterns of assortative mating and define a marriage market, μ, as men and women who share a common education level e (up to high school or at least some college), race r (white non-Hispanic, black non-Hispanic, Hispanic), and who live in the same state s.

Using these parameters to determine the marriage market is not perfect, since it imposes some assumptions about the

---

[9] Appendix figure A.2 shows the variation in male and female wages across states.

FIGURE 3.—VARIATION ACROSS STATES IN CLOSING THE GAP: CHANGE IN LOG($\frac{\text{FEMALE WAGE}}{\text{MALE WAGE}}$), 1980–2010

- (.23,.28]
- (.17,.23]
- (.11,.17]
- [.05,.11]
- No data

This figure shows the state-level growth in the log female-to-male wage ratio between 1980 and 2010. Source: 1980–2000 decennial censuses, 2010 ACS; author's calculation.

set of relevant mates for a given woman, but it has several clear benefits. First, I maintain continuity with BKP, who define marriage markets using the same set of education, race, and geographic cells. This allows for greater comparability across estimates. Second, defining education in two coarse bins, rather than finer categories, follows historical patterns of intermixing across education categories. In particular, among 1980 married women, 20% of the college-educated were married to a man with only some college, and 32% of women who had only some college married a man with a college degree.[10] An additional advantage is that these coarse bins allow me to categorize young men and women who may be in transition between some college and college to their permanent marriage market with greater confidence. Third, defining the marriage market to one's state of residence, rather than a more granular geographic level, allows for matching to occur outside one's locality. If these parameters are misspecified, such that the relative wage I assign to individuals is not reflective of own wages or potential partners' wages, I expect that my results will be attenuated by measurement error.

A.   Construction of Relative Potential Wage

Since wages and potential wages are not randomly assigned, I follow the general approach of the literature and set the potential wage as the weighted average of national industry-specific log wages $w_{jt,-s}$, excluding the state in

which the market is located ($-s$), where the weights reflect the share of the market μ and sex g employed in industry j prior to the period of analysis (1970) (Aizer, 2010; BKP Diamond, 2016):

$$w_{\mu gt} = \sum_j \frac{E_{j\mu g,1970}}{E_{\mu g,1970}} \times w_{jt,-s}.$$

To gain intuition for this approach, one can consider the simple case of two industries, j_1 and j_2, and that wages vary from low to high over time only for j_1. In that case, the the logic of identification simplifies to a difference-in-difference design, comparing the effect of the wage change across markets that have a greater presence of j_1 to those with a lesser presence of j_1.

I introduce two adjustments to this setup. First, I incorporate variation in shares and wages across occupations o within an industry, drawing on recent work that highlights the importance of occupation-specific tasks in the closing of the gender gap. This requires knowledge of the historical industry and occupation of men and women in a marriage market. Since these cells are quite small in the data, I approximate this by assuming that men and women in each race and skill nationally work in the same occupations within each industry. Then I estimate the initial employment share of each marriage market and sex in o and j, $\frac{E_{oj\mu g,1970}}{E_{\mu g,1970}}$, by multiplying the 1970 share of a sex and marriage market that I observe working in each industry by the share of that sex with the same skill and race in 1970 United States that worked in a given occupation in that industry.[11] The use of subindustry wages has become a standard way to gain additional precision in this

[10]A slightly lower share of 1980 married women with some college married men with only a high school diploma (29%). By 2010, women with exactly some college were less likely to marry men with a college degree than men with high school or less (22% versus 33%); however, since this is an endogenous outcome, I prefer to follow historical patterns by grouping some college and college together. The results are not sensitive to an alternative education grouping that separates college-educated from lower-educated, as shown in appendix A.4.

[11]I calculate this as $\frac{E_{jersg,1970}}{E_{ersg,1970}} \times \frac{E_{ojerg,1970}}{E_{jerg,1970}}$. Appendix figure A.3 shows a high correspondence between observed shares and this prediction.

prediction, but typically relies on demographic-level wages (BKP). I discuss the additional power gained from this approach in section IVC.

Second, I allow systematic updating of the initial marriage-market weights following national trends in occupation growth. This helps to account for the large changes in the distribution of employment over this period. However, rather than rely on overall growth in occupations, which was influenced by changes in women's labor supply behavior (Black & Juhn, 2000), I take advantage of the differential growth in the importance of occupations across industries. This is reflected in deviations in the growth of the within-industry employment share of the occupation, $\frac{E_{ojt,-s}}{E_{jt,-s}}$, from the growth in the national employment share of the occupation, $\frac{E_{ot,-s}}{E_{t,-s}}$.[12] This source of growth is more likely to reflect industry-productivity or industry-technology than labor supply decisions and contributes additional power to the potential wage, which I quantify in the next section.

Letting the ratio of the within-industry share in relative to 1970 be $\pi^W_{ojt,-s}$, and the ratio of the national occupation share in relative to 1970 be $\pi_{ot,-s}$, this updating term is written as

$$\pi^{W*}_{ojt,-s} = \left(\pi^W_{ojt,-s}\right)\left(\frac{1}{\pi_{ot,-s}}\right).$$

The updating term is then normalized such that the weights sum to 1. Updating the weights in this manner allows the weights for each marriage market to more realistically reflect the contemporaneous local employment conditions without compromising the validity of the potential wage.[13]

The resulting potential wage is

$$\widehat{w_{\mu gt}} = \sum_j \underbrace{\frac{E_{j\mu g,1970}}{E_{\mu g,1970}}}_{\substack{\text{Between-industry}\\\text{exposure, 1970}}} \times \sum_o \underbrace{\frac{E_{oj\mu g,1970}}{E_{j\mu g,1970}}(\pi^{W*}_{ojt,-s})}_{\substack{\text{Within-industry}\\\text{exposure, t*}}} \times w_{ojt,-s}.$$

Intuitively, the key variation in the log relative potential wage ($\widehat{w_{\mu female t}} - \widehat{w_{\mu male t}}$) comes from the extent of the segregation of men and women in occupations and industries within a marriage market interacted with the change in returns to occupation over time. All else equal, marriage markets experience more growth in the relative wage when men and women in the market have less overlap in their occupations and industries and when the occupations and industries that women are in experience relatively more growth in wages. As illustration of this specialization, appendix table A.2 shows that there is little overlap in the ten most common occupations for men and women.

B. Estimating Equation

With this in hand, I use data collapsed to cells defined by a marriage market, μ, cohort, c, and decade, t, to estimate

$$Y_{\mu ct} = \beta ln\widehat{RelativePotentialWage}_{\mu t} + \alpha_\mu + \delta_{rt}$$
$$+ \chi_{et} + \gamma_{st} + \xi_{ct} + \rho_{rs} * t + X_{\mu t}\phi + \upsilon_{\mu ct}. \quad (1)$$

The potential wages I construct circumvent the concerns of using local wages in this equation, but this approach remains vulnerable to potential biases that could induce a spurious correlation between potential wages and family outcomes. The most severe of these concerns stems from the initial patterns of sorting into occupations and industries across marriage markets. Returning to the simple example above, if markets with a greater weight given to j_1 are unobservably different, this approach would misattribute those differences to the potential wage. To address this, I include a vector of marriage market fixed effects, α_μ, which fully absorbs the cross-market variation in initial occupation and industry choices as well as in marital preferences and expectations regarding family and labor market work.

This solution would not be sufficient, however, if occupation choices are correlated with trends in behavior rather than fixed differences. This could happen, for instance, if states with a larger hospital sector, and, hence, a larger share of nurses, also tended to have a strong downward trajectory in marriage or if low-educated men and women concentrated in a few occupations and increasingly tended toward delaying marriage. Therefore, I add a rich set of fixed effects to control for differential trends for each race, education level, and state: δ_{rt}, χ_{et}, γ_{st}. I also allow for varying trends by state and race $\rho_{rs} \times t$. The inclusion of this extensive set of controls ensures that I will not spuriously attribute the effect of secular trends to the relative wage and significantly reduces the scope for systematic bias to affect my estimates. As I discuss in the next section, as one might hope, the variation in the relative potential wage is largely orthogonal to these fixed effects. Nevertheless, there is insufficient variation in wages to be able to include all of the three-way interactions between state, education, race, and year in this specification, which remains a limitation of this approach.[14]

---

[12]To fix ideas, in appendix table A.3 I present $\pi^W_{oj,2010}$ and $\pi_{o,2010}$ for management (columns 2 and 3, respectively). I update the weights using the ratio of columns 2 and 3.

[13]Appendix table A.22 shows that the results are the same without updating shares.

[14]Although in theory one could also include fixed effects for each year-education-race cell or each year-state-education cells, empirically there is insufficient variation in the potential wage to include these. This suggests that an important source of variation in potential wages comes from differences in employment shares across education groups within a state and across education-race groups. This would be a concern, for example, if there are unobserved trends or shocks to marital outcomes within states across education groups that are correlated with potential wages. Since many of the salient changes in marriage occur across states or across education groups over this time period, this seems less likely. Further, such unobserved changes might be expected to affect both men and women; however, if that was the case, I would expect that potential wages for men and women would be correlated due to these unobserved common shocks. I show that this is not the case in section IVC.

Table 1.—Relationship between Potential Wages and Observed Wages

| | Correlation with Actual | | | Cross-Effects? | |
	Relative	Female	Male	Female	Male
ln Rel. Wage	0.833***				
(Potential)	(0.225)				
ln Female Wage		0.426**		0.571***	−0.232
(Potential)		(0.202)		(0.186)	(0.255)
ln Male Wage			0.481***	−0.228	0.613***
(Potential)			(0.156)	(0.277)	(0.203)
Partial R-squared	0.067	0.039	0.041		
Observations	1,064	1,064	1,064	1,064	1,064

This table shows the coefficients from estimating equation (1), omitting cohort controls. The dependent variable is shown in the column heading. Standard errors are clustered at the state level, and cells are weighted by the female population in the cell. * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

Sources: Potential wage: 1970 Decennial Census, 1980–2011 March CPS, Wages: 1980–2000 Decennial Censuses, 2010 ACS.

To address remaining confounds in the marriage market, I include $X_{\mu t}$, a vector of mean educational attainment for men and women and the sex ratio, which could also influence marriage decisions and matching. In most specifications, I also control for the market-level average potential wage to account for changes in the level of wages that affect the combined earning power of households. The inclusion of the average wage does not alter the results, as I discuss in section VA. Although potential wages do not vary by cohort, I include cohort by year fixed effects, ξ_{ct}, to increase the precision of the estimates. All regressions are weighted by population, and I cluster standard errors at the state level.

Finally, for ease of interpretation, I rescale the coefficients from all regressions to represent the effect of a 10% increase in the relative wage. Since this is roughly the variation that I leverage between 1980 and 2010, the estimates can be interpreted as the impact of the relative wage on outcomes over this period.

C. Evidence on Identifying Assumptions

This estimating equation produces unbiased estimates of the effect of the potential relative wage under key identifying assumptions. First, the potential wage should be correlated with the observed wage, which I used as a stand-in for unobserved potential wages. Appendix figure A.5 presents descriptive evidence of this correlation using observed wages in the Census/ACS. It shows a positive correlation between the long change (1980–2010) in log relative, log female, and log male wages and the equivalent change in the potential wage.

The first three columns of table 1 show the estimated correlations from regressing observed wages on potential wages using a modified version of equation (1), which omits the unnecessary cohort controls. The estimates show sizable correspondences between potential wages and observed wages that are precisely estimated. The coefficients are 0.426 ($p < 0.05$) for women, 0.481 (male, $p < 0.01$) for men, and 0.833 ($p < 0.01$) for the relative wage. In appendix A.4 I show that the addition of variation by occupation increases the magnitude

of this correlation four-fold and that updating reduces the estimated standard errors by 10%.[15] I also show that the results are robust to using a more succinct list of occupations, using alternative definitions of the marriage market or using Census wages.

Second, the proxy for potential wages must be plausibly exogenous to decisions in the marriage market. As mentioned earlier, it is not necessary for the potential wage to have an impact on outcomes only through wages. To increase the plausibility of this assumption, I include a rich set of fixed effects, which absorb a large number of potentially worrisome sources of variation (see discussion above.) As shown in appendix table A.4, the variation in the potential wage is uncorrelated with many of these fixed effects. Not surprisingly, adding the marriage market fixed effect in column 2 has the largest impact on the coefficient, which is consistent with a large portion of the variation coming from the initial, and persistent, employment patterns in each marriage market.

Nonetheless, one might be concerned that the residual variation in the potential wage could be correlated with contemporaneous or lagged shocks to the marriage market, which could result in reverse causation. To address these concerns, I perform two types of placebo tests. First, I ask: Is the current potential wage for men or women predictive of opposite-sex wages (Autor et al., 2018)? This allows me to rule out the possibility that the potential wage is spuriously correlated with a shift in general marriage market conditions, such as a local resource boom or greater enthusiasm for working. Reassuringly, columns 4 and 5 of table 1 show that only the coefficient on the same-sex potential wage is significant. Second, I check whether the future relative wage is predictive of past outcomes in the market. This allows me to assess whether the relative wage I construct is correlated with past trends in outcomes, as I describe in further detail in section VI. I find no significant relationships across many outcomes, consistent with the identifying assumption.

Finally, following Goldsmith-Pinkham, Sorkin, and Swift (2018), I verify that no single occupation-industry is contributing the majority of the identifying variation for the potential wage. This ensures that the occupation-industry space is sufficiently large, which is, in theory, a sufficient condition for exogeneity under the "many invalid instruments"

---

[15]See appendix table A.21. Panels A and B show the correlation between observed wages and (a) a proxy constructed with shares and wages at the demographic-industry level and (b) a proxy constructed with shares at the demographic-industry-occupation level and wages at the industry-occupation level, respectively. The panel B proxy is equivalent to my preferred proxy without any dynamic updating. Once the full set of controls for my preferred specification is added in column 7, the correlation in panel B is four times higher than the correlation in panel A. Further, comparing across panel C, which has my preferred proxy, and panel B, the standard errors are roughly 10% smaller in panel C. For further detail, see the discussion in appendix A.4. In results not reported, I find that a potential wage using only occupation variation suffers from lack of power similar to the industry-only potential wage, which is sensible given the variance in occupation-specific wages across industries.

TABLE 2.—IMPACT OF RELATIVE WAGE ON MARRIAGE

	(1) Married	(2) Divorced	(3) Never Married
A: Relative only			
Effect of 10% increase in relative wage	−0.051***	0.019***	0.032***
	(0.014)	(0.006)	(0.010)
Observations	23,573	23,573	23,573
B: Relative controlling for average			
Effect of 10% increase in relative wage	−0.048***	0.017***	0.031***
	(0.009)	(0.006)	(0.008)
Effect of 10% increase in average wage	0.079***	−0.044***	−0.027***
	(0.011)	(0.007)	(0.009)
Mean Y	0.645	0.102	0.245
Observations	23,573	23,573	23,573

This table shows the coefficients from estimating equation (1) rescaled to represent the effect of a 10% increase in the relative (potential) wage. The dependent variable is shown in the column heading. Standard errors clustered at the state level, and cells are weighted by the female population in the cell. *$p < 0.10$, **$p < 0.05$, and ***$p < 0.01$.
Source: Census 1980, 1990, 2000, and ACS 2010.

argument (Borusyak, Hull, & Jaravel, 2018; Goldsmith-Pinkham et al., 2018). Appendix figure A.4 shows that the top five occupation-industries contribute 17% of the variation, in terms of Rotemberg weights, for men and 50% of the variation for women, the majority of which is accounted by schoolteachers (18%). Further, excluding female schoolteachers, the top shares are dispersed across different industries and occupations, indicating that endogenous shares are less likely to be problematic for this setting.

V. Results

A. Marriage and Cohabitation

Panel A of table 2 presents the estimated effect of the relative wage on marriage decisions before controlling for the average potential wage. Column 1 shows that a 10% increase in the relative wage leads to a 5.1 p.p. (8%) decline in the probability that a woman is married. To put this in perspective, this effect size is of similar magnitude to the increase in the fraction of never-married women following the introduction of the birth control pill and twice as large as the increase attributed to the rise in incarceration (Goldin & Katz, 2002; Charles & Luoh, 2010).[16] My estimate implies that approximately 20% of the decline in marriage during the past three decades is attributable to the increase in the relative wage.[17] The subsequent columns decompose this result into its two main components: dissolution of existing marriages and decline in first marriage. I find that a 10% increase in the

relative wage leads to a 1.9 percentage point increase in the likelihood of divorce, accounting for one-third of the change in marriage rates. The 3.2 p.p. rise in never-married women, an 18% increase, then accounts for the remainder of the marriage decline. This could either reflect a delay of first marriage or a more permanent opting out of marriage. I return to this question in section VD, where I analyze responses separately for younger and older women.

Panel B examines the sensitivity of my results to controlling for the mean inflation-adjusted potential earnings in the market, which I construct as the average of the male and female potential wage in the market. This allows me to separate the effect of the relative wage from absolute wages, which appear to be slightly correlated in figure 2. The estimated effects of the relative wage are only marginally reduced by the introduction of this control variable. A 10% increase in the relative wage leads to a 4.8 p.p. decline in marriage, a 1.7 p.p increase in the likelihood of divorce, and a 3.1 p.p. increase in the likelihood of being never married. The insensitivity of the point estimates to this control indicates that there is substantial variation in the relative wage measure independent of the average wage measure; nevertheless, I include this control going forward to allow for the cleanest interpretation of the results, and the results are not affected when I do not include it.

Having shown that higher relative wages leads to a decline in marriage, I turn to considering which couples respond to the change in this incentive. As a first step, I look at whether women substitute away from marriage toward cohabitation, a potentially less costly form of commitment. Since the data on cohabitation are incomplete, I use three complementary measures to gain evidence on whether women are living with a romantic partner. First, I use the official report of cohabitation, which is available only for the head of household from 1990 on. Second, I look at whether a woman lives only with an unmarried man, who may be a romantic partner or a platonic roommate. I consider these two measures as providing approximately lower and upper bounds of effects on the stock of cohabitating couples. Third, I observe whether a woman lives with only another woman, as certain evidence of living in a platonic arrangement.

Column 1 of table 3 shows that a higher relative wage does not meaningfully increase official reports of cohabitation. The point estimate indicates that a 10% increase in the relative wage reduces cohabitation by 1% and is statistically insignificant. Consistent with this, in the following columns, I find a rise in women living in a platonic arrangement. The estimates imply that 30% of women who would have married and resided with a husband instead live with another woman (a 1.2 p.p. increase relative to the 4 p.p. decline in living with a husband). There is a similar increase in the share of women living alone, although less precisely estimated, accounting for 35% of the decline in living with a husband. By comparison, I find a smaller increase (by 0.9 p.p., or 22.5% of the previously married women) in the propensity to live with an unmarried man.

---

[16]The availability of the pill led to a 6 p.p. increase in the likelihood of being never married, which was mirrored in a similar-size decline in divorce rates (Goldin & Katz, 2002). The rise in imprisonment from 1980 to 2000 increased the share of never-married women by 2.3 p.p. and reduced divorce rates by 0.9 p.p (Charles & Luoh, 2010, with calculations based on their tables 2 and 3).

[17]Calculated as $\frac{5.1*.8}{73.83-55.98} = 22.9\%$. See appendix table A.1 for source numbers.

TABLE 3.—IMPACT OF RELATIVE WAGE ON COHABITATION AND PLATONIC LIVING ARRANGEMENTS

	Cohabitation	Only Other Adult in HH Is:			Live Alone	2+ Other Adults
	(1) Official Report	(2) Husband	(3) Single Male	(4) Female	(5)	(6)
Effect of 10% increase in relative wage	−0.004 (0.018)	−0.040*** (0.011)	0.009** (0.004)	0.012*** (0.003)	0.014 (0.013)	0.006 (0.012)
Mean Y	0.159	0.457	0.066	0.059	0.259	0.159
Observations	16,925	23,573	23,573	23,573	23,573	23,573
Average wage	Yes	Yes	Yes	Yes	Yes	Yes

This table shows the coefficients from estimating equation (1) rescaled to represent the effect of a 10% increase in the relative (potential) wage. The dependent variable is shown in the column heading, and the unit of observation is a cell defined by an education × race × state × birth cohort × year. The official report of cohabitation in column 1 is an indicator for whether the head of household reports living with an unmarried partner, and is only available from 1990 on. The outcomes in columns 2, 3, and 4 are indicator variables that take on the value of 1 if an individual lives in a two-adult household (i.e., there are two individuals who are at least 18 years old) where the other adult is her husband, a single male, or a female, respectively. The outcomes in columns 5 and 6 are indicators for living alone or with two or more other adults, respectively. Standard errors are clustered at the state level, and cells are weighted by the female population in cell. * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.
Source: Census 1980, 1990, 2000, and ACS 2010.

TABLE 4.—IMPACT OF RELATIVE WAGE ON SPOUSAL EDUCATION

	Spouse Education, Relative to Own			Spouse Minus Own Education
	(1) Less	(2) Same	(3) More	(4)
Effect of 10% increase in relative wage	−0.023* (0.012)	−0.028 (0.017)	0.051*** (0.013)	0.258*** (0.064)
Mean Y	0.319	0.357	0.324	0.017
Observations	22,663	22,663	22,663	22,663
Average wage	Yes	Yes	Yes	Yes

This table shows the coefficients from estimating equation (1) rescaled to represent the effect of a 10% increase in the relative (potential) wage. The dependent variable is an indicator for a spousal characteristic, is shown in the column heading. Standard errors are clustered at the state level, and cells are weighted by the female population in the cell. * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.
Source: Census 1980, 1990, 2000, and ACS 2010.

B. Spousal Matching

Second, I analyze whether these is an improvement in the attributes of women's spouses. The primary outcome of interest is spousal education, which has frequently been cited as influential in women's matching decision (Charles & Luoh, 2010; Fisman, Iyengar, Kamenica, & Simonson, 2006). Education also correlates with multiple other dimensions of skill, such as social skills, health, and parenting style, which may independently be valued in the marriage market. I also analyze impacts on spousal age, another common measure of spouse quality, though I place less emphasis on these as a measure of spousal quality since the empirical evidence of women's preferences over age are mixed.[18]

Table 4 presents effects for spousal education outcomes. The table shows a consistent pattern: a higher relative wage causes women to be more likely to be married to spouses who are more educated, conditional on marrying. The first three columns measure the share of women who have a spouse with fewer years of education, the same level of education, or more years of education than they do, respectively. I find that a 10% increase in the relative wage leads to a 15% increase in the probability of marrying a partner more educated than oneself (5.1 p.p.), which results from equal declines in the

probability of marrying a partner less educated and same educated relative to oneself. In the final column, I find that a 10% increase in the relative wage leads to a 0.26 year increase in the difference between husbands and wives.[19]

Appendix table A.6 indicates that spousal age also responds to the relative wage. A higher relative wage reduces the likelihood a woman has a husband a spouse older than she is ($p < 0.1$), and increases the likelihood she has a same-age or younger husband, though the latter effects are not significant. While there is not a consensus in the literature, previous research suggests that same-age and younger spouses are associated with having a more attractive spouse and greater initial marital satisfaction for women, respectively (Mansour & McKinnish, 2013; Lee & McKinnish, 2017).

By construction, this implies that husbands are more likely to have a less educated wife and, suggestively, more likely to have an older wife. Relative to women, men place less weight on a partner's intelligence (Fisman et al., 2006) but have a strong preference for having a younger wife (Low, 2016). Thus, these results could be consistent with a decline in the spouse quality for men; however, the lack of precision in the estimates over age makes this more uncertain.

To gain insight into unobservable spouse quality, I turn to the National Survey of Families and Households (NSFH), described in appendix A.5. Unfortunately, the small number of young, married households in the survey (3,000) provide too little power to support an identification strategy with the proxy I construct. Therefore, I follow BKP and provide descriptive evidence of the relationship between relative wife-to-husband income and the measures of interest. Motivated by the model, I focus on households where men are the dominant earner (80% of these households). The results show that a higher relative income is associated with greater happiness with the marriage among women, which is mirrored by a decline in reporting of marriage trouble by women and men. This provides speculative evidence that a higher relative wage could also raise the unobserved quality of husbands, on top of the improvement in spousal attributes documented above.

---

[18] See, e.g., Mansour and McKinnish (2013), Hitsch, Hortasu, and Ariely (2010), Low (2016), and Lee and McKinnish (2017).

[19] Nonetheless, in appendix table A.5, I check whether relative education responds to the relative wage and find no effect.

TABLE 5.—IMPACT OF RELATIVE WAGE ON WOMEN'S LABOR MARKET OUTCOMES

	Conditional on Working				Unconditional	
	(1) Weekly Hours	(2) Weeks Worked	(3) ln(Wkly Inc)	(4) ln(Inc)	(5) Any Earnings	(6) In Labor Force
Effect of 10% increase in relative wage	1.021***	−0.274	0.059**	0.054*	−0.011	0.001
	(0.224)	(0.441)	(0.023)	(0.030)	(0.015)	(0.013)
Mean Y	36.823	43.487	5.747	9.420	0.723	0.718
Observations	23,222	23,222	23,186	23,188	23,573	23,573
Average wage	Yes	Yes	Yes	Yes	Yes	Yes

This table shows the coefficients from estimating equation (1) rescaled to represent the effect of a 10% increase in the relative (potential) wage. The dependent variable is a labor market outcome, shown in the column heading. Standard errors are clustered at the state level, and cells are weighted by the female population in the cell. * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.
Source: Census 1980, 1990, 2000, and ACS 2010.

C. Labor Market

Table 5 analyzes women's labor market outcomes. In the first column, a 10% increase in the relative wage causes women to work 1 additional hour per week. This is a small change in hours relative to the mean but explains the entirety of the rise in weekly hours of work over this period. This increase in hours worked could be caused by the change in the composition of women, since single women tend to work more hours, or by changes in working behavior conditional on marital status. A back-of-the-envelope calculation suggests that the change in the marital composition of women can explain at most 20% of the increase in hours of work, and the remaining 80% is changes in working behavior among married and single women.[20]

The remainder of table 5 shows a statistically insignificant effect on weeks worked, but that women earn 5% more weekly income and 5% more in annual income ($p < .1$). In the last column, I find a statistically insignificant effect on employment. This is consistent with the model's prediction that if there is a sufficiently high lower bound of spouse quality, women who are marginal on the extensive labor margin may be inframarginal to marriage. These estimates suggest that on average, a lower relative wage does not deter women from working, but it does reduce hours worked, consistent with specialization within households.

Financial independence is an additional outcome of interest because it may provide a channel by which a higher relative wage can reduce women's susceptibility to emotionally or physically abusive relationships. Appendix table A.7 shows that a 10% increase in the relative wage leads to a 3.4 p.p. increase in a woman's share of household earnings, a 7% increase. There are also substantial increases in the likelihood that women control all of the income in the household. A 10% increase in the relative wage leads to a 2.4 p.p. decline in the probability that a woman is in a male breadwinner household, a 20% decline, accompanied by an equivalent increase in households where there is a female breadwinner. These effects indicate that a higher relative wage reduces women's reliance on a male earner, which could be a mechanism for previously documented reductions in domestic violence.

D. Heterogeneous Responses

Having shown that the relative wage reduces marriage on average, I examine whether differential responses could explain the steeper decline in marriage among low-skilled women shown in figure 1. Appendix table A.8 shows that the effects across low- and high skilled women are qualitatively similar and not statistically different from one another.[21] However, these results may mask heterogeneity in marital responses if differential behavior emerges at a later age. For example, low- and high-skilled women may place equal importance on the pecuniary incentive to marry at younger ages but respond differently to the incentive at later ages as marriage norms and joint child rearing become greater concerns.

In line with this, table 6 shows that for low-skilled women, a higher relative wage leads to postponement of marriage both at younger (22–30) and older (31–44) ages, while for high-skilled women, a higher relative wage appears to only delay marriage at younger ages. This pattern of responses can help to explain two salient trends across these groups: increasing age at first marriage for all women and reduction in marriage between ages 33 to 44 among low-skilled women (Lundberg et al., 2016). As might be expected, a higher relative wage leads to increased divorce between ages 31 and 44, but not between ages 22 and 30, for both groups of women.

Additionally, I check in appendix table A.9 whether the effects on marriage and divorce dissipate over time, possibly due to reductions in the stock of marriages with undesirable spouses (Rotz, 2016). I do not find any decline in the coefficients across decades, suggesting that there continue to be marriages that are marginal to the relative wage even as screening standards for spouses have increased over time.

---

[20]Conservatively, I assume that the difference in working hours between single and married women over this period is the maximum of the period (4 hours in 1981) and multiply this by the 4.8 p.p. increase in single women. This would imply that the increase in single women could increase the number of working hours by 0.2 hours.

[21]Examining effects across races, I find more precise and larger responses among whites, although I find significant effects for blacks for impacts on first marriage, spousal quality, and hours of work. The confidence intervals for both blacks and Hispanics allow for meaningful impacts across outcomes, though.

TABLE 6.—POSTPONING OR OPTING OUT? DIFFERENTIAL RESPONSES BY AGE AND EDUCATION

	22–30	31–44
A: Never Married		
Effect of 10% Increase in Rel. Wage	0.033**	0.038***
× Low Skill	(0.013)	(0.011)
Effect of 10% Increase in Rel. Wage	0.091***	−0.029*
× High Skill	(0.017)	(0.016)
Observations	9,286	14,287
Mean Y - Low Skill	0.322	0.125
Mean Y - High Skill	0.469	0.145
B: Divorce		
Effect of 10% Increase in Rel. Wage	−0.005	0.022**
× Low Skill	(0.007)	(0.010)
Effect of 10% Increase in Rel. Wage	0.009	0.042***
× High Skill	(0.009)	(0.015)
Observations	9,286	14,287
Mean Y - Low Skill	0.076	0.130
Mean Y - High Skill	0.050	0.129

This table shows the coefficients from estimating equation (1) interacted with indicators for education rescaled to represent the effect of a 10% increase in the relative (potential) wage. Interactions of the average potential wage with education are also included. The age of the sample is shown in the column heading. Standard errors are clustered at the state level, and cells are weighted by the female population in the cell. $^* p < 0.10$, $^{**} p < 0.05$, and $^{***} p < 0.01$.
Source: Census 1980, 1990, 2000, and ACS 2010.

VI. Robustness

A. Reverse Causation from Unobserved Trends

As discussed earlier, one might be concerned that the effects are driven by reverse causation, which in theory could happen if initial employment shares were correlated with trends in unobserved marital preferences. To test for this, I ask whether markets that are predicted to have an increase in relative wages historically had dissimilar trends from other markets. Following Autor, Dorn, and Hanson (2013), I estimate a first-differenced (ten-year differences) specification, using lagged outcomes from 1960 to 1970, and potential wages and controls collapsed to a marriage-market-by-year-pair panel.[22]

Appendix table A.10 presents the results for regressions of the relative wage for 1980 to 1990 (panel A), 1990 to 2000 (panel B), or 2000 to 2010 (panel C) on 1960–1970 outcomes. Of the fifteen marriage outcomes analyzed in the first five columns, none of the relationships are statistically significant, and the direction of the relationship is typically inconsistent over time or with the the main results. While some of the coefficients approach the size of the main estimates, the standard error is typically several times larger, such that the 95% confidence interval includes an effect size of comparable magnitude in the opposite direction as the main estimate. The subsequent columns show results for the remainder of the key outcomes, as well as for poverty and single motherhood. I continue to find no statistically meaningful evidence for reverse causation.

B. Lagged Effects and Migration

Appendix table A.11 examines the sensitivity of the results to including the one-year lag of the relative and the average wage in addition to the contemporaneous wages. Reassuringly, the main (contemporaneous) effects are relatively unchanged from the main specification, which suggests that present shocks have a distinct and strong impact on marriage and labor market decisions. The effect of the lagged relative wage is usually half the size and statistically less precise than the contemporaneous effect.

Next, in appendix table A.12, I examine to what extent selective interstate migration could influence the results using data from 1980 to 2000.[23] Panel A shows the baseline estimates for this sample, together with the estimated effect of a higher relative wage on migration. I find that a 10% increase in the relative wage is associated with a 2.3 p.p. increase in interstate migration, suggesting that markets with a higher relative wage are more attractive to women.

In panels B and C, I separate the effect of the relative wage into the impact on new arrivals and women who have not migrated in the last five years. The change in outcomes among recent migrants gives an upper bound of the role of selective migration in my measured effects.[24] Using this conservative measure, migration could account for up to 52% of the estimated increase in never-married women, 20% of the increase in divorce, 28% of the pairing with higher educated spouses, and 38% of the increase in female breadwinners. This ceiling on the role of migration implies that the majority of the impacts are due to changes in behavior among long-standing residents.

VII. Discussion

In this section, I examine alternative theories that could generate some of these patterns and then look to quantify the spillover impacts of the decline in marriage on children.

A. Alternative Explanations

A first alternative rationalization of the marriage decline is that women do not leave their partners. Rather, because marriage is costly, when the relative wage increases, women prefer not to incur these costs and substitute toward cohabitation, a less costly form of commitment. Under this hypothesis, the number of marriages can fall even when the marginal spouse quality is positive. Counter to this theory, I find little, if any, substitution toward cohabitation in section VA.

A second possibility is that these effects reflect reactions to the absolute decline in male wages, which reflects fewer

---

[22]Specifically, I estimate $\Delta Y_{erst'} = \beta_1 \Delta ln\widehat{RelativePotentialWage}_{erst} + \beta_1 \Delta ln\widehat{AveragePotentialWage}_{erst} + \delta_{rt} + \chi_{et} + \gamma_{st} + \rho_{rs} + \Delta X_{erst} \phi + \upsilon_{erst'}$.

[23]Migration is defined by the presence of individuals who moved across states in the past five years, which is available in the 1980, 1990, and 2000 Censuses.
[24]To prevent concerns of endogenous stratification, the outcomes are a combination of migration status and marriage status: for example, moved in the last 5 years and never married.

"marriageable men," rather than the rise in the relative wage. This theory would predict that (a) the effects depend only on the male wage and (b) are concentrated in areas with declining men's wages. I first separate the effects of the relative wage into effects of the male wage and female wage. It shows that men's and women's wages have opposite-signed effects on the main outcomes. The effect of men's wages is typically larger and more precisely estimated. Nonetheless, women's wages have a statistically significant impact on the never-married rate, spousal education, and women's hours of work. This indicates that women's wages are influential for spouse quality and first-marriage decisions, contrary to a theory that emphasizes male wages alone.

In appendix table A.14, I probe the importance of declining male wages by testing whether the effects are concentrated in markets where men experienced the least growth in potential wages, in the first quartile of potential wages. It shows that the effects of the relative wage are dispersed throughout the distribution of growth in male potential wages; however, the largest and most precisely estimated effects are in the lowest quartile of men's wage growth. This suggests that the decline in men's real wages plays an outsize role in the responses to the relative wage that I document but does not appear to fully explain the impacts.

A third possibility is that the effect of the relative wage is simply picking up the "aversion to women earning more than men" (BKP). I address this concern using three complementary analyses. To bound the impacts of this aversion, I regress the observed probability that a woman earns more than a man in a marriage market, *PrWomanEarnsMore*, on the potential relative wage, using my main estimating equation.[25] As shown in the first column of appendix table A.15, a 10% increase in the potential relative wage correlates with a 2.9 p.p. increase in the probability that a woman earns more than a man in the market. Pairing this estimate with the marriage coefficients from BKP, I calculate that the aversion channel could account for 1 p.p. of my 4.8 p.p. effect.[26] This indicates that the majority of the effect that I estimate cannot be explained by the aversion channel.

Next, I directly test the sensitivity of my main outcomes to controlling for *PrWomanEarnsMore*. Appendix table A.16 shows that the coefficient on the relative wage is essentially unchanged by the inclusion of this control. As in BKP, the coefficient on *PrWomanEarnsMore* indicates that a rise in the probability that a woman earns more reduces the likelihood of marriage.[27] I do not find a statistically significant effect of

PrWomanEarnsMore on spouse quality and the coefficient is negative, which goes in the opposite direction of the effect of the relative wage. Third, more suggestively, I test whether the aversion channel would predict the asymmetry that I find in the effect of male and female wages. The second column of appendix table A.15 shows that the increase in *PrWomanEarnsMore* from male and female wages is in fact symmetric. Thus, if the only channel by which the male or female wage affected marriage was through *PrWomanEarnsMore*, we would expect symmetric effects of these wages.

Finally, to address other potential theories, in appendix table A.17 I introduce controls for other potential mechanisms highlighted in the literature, such as the rise in incarceration (panel A) and male wage inequality (panel B), or the decline of manufacturing (panel C) (Loughran, 2002; Gould & Paserman, 2003; Charles & Luoh, 2010; Autor et al., 2018). The addition of these controls does not change the qualitative results, although the effects in panel C are less precisely estimated.

B. Spillovers to Children's Family Structure

Understanding the impact of the relative wage on the well-being of children is a complex question that would require evidence on a multitude of outcomes outside the scope of this paper. Nevertheless, an important piece of this is determining to what extent the effects on marriage that I have documented bear on children's family structure. I answer this by examining impacts on the total number of children in a market, as well as separating the impacts on being married/never-married/divorced by the presence of children.[28]

Column 1 of appendix table A.18 shows that a 10% increase in the relative wage reduces the share of women with children present by 3 p.p. or 5%. Decomposing this effect, I find that the decline in children is primarily driven by a substantial 5.6 p.p. ($p < 0.05$) reduction in women who are married with children, an 11% effect, and, to a lesser extent, by an increase in the share of women who divorce prior to having children (1.2 p.p., $p < 0.05$) and who remain childless in marriage. On the other hand, there is also a rise in never-married mothers, which offsets 40% of the decline in married mothers. Hence, a higher relative wage leads fewer women to marry and have children but also increases the share of single mothers.[29] The net effect on children, then,

---

[25]Following BKP, I construct the probability that a woman earns more than a man by taking random draws of 100,000 men and women in a marriage market. The means of this variable align with reported means by BKP for the observed probability that a woman earns more.

[26]Calculated as 0.35 p.p. (SE: 0.181 p.p.) decline in marriage for a 1 p.p. increase in the probability that a woman earns more than a man (column 9 of table 1 of BKP) times 3. Using the upper bound of the BKP 95% confidence interval allows for at most a 2 p.p. effect, or 41% of my estimate.

[27]BKP also report estimates for the impact of the potential female and potential male wage in the market; however, the reported coefficients are somewhat difficult to interpret or compare to my estimates since BKP also

control for the average relative income and deciles of the female and male wage in the market, which are likely to share some collinearity with the average male and female wage.

[28]Outcomes are additive so that the sum of the impact on, say, being never-married with children and never-married without children sums to the total impact on being never-married.

[29]This is in contrast to recent studies that utilize changes in the sex ratio as an instrument for changes in women's bargaining power and find that increases in the male-to-female sex ratio increase women's marriage and reduce single motherhood (Brainerd, 2017; Abramitzky et al., 2011). However, unlike those settings, an increase in the relative wage raises women's bargaining power while also reducing the pecuniary gains to marriage.

will depend on the relative welfare implications of each of these circumstances, which is left for future work.

VIII.   Conclusion

The role of women in the American household has undergone a substantial shift over the past three decades. Despite the significant attention given to this trend, particularly the decline in marriage, the counterfactual to this transition has been under-explored. This paper presents new theoretical predictions and evidence on the significance of the declining gender wage gap for spouse quality, marriage, and women's labor market outcomes. For identification, I rely on variation from shifts in demand over this period, which favored traditionally "female" occupations and affected marriage markets differentially due to historical patterns of industry location.

I provide empirical support for each of four theoretical predictions. First, I find that a 10% increase in the relative (potential) wage increases the likelihood that women marry a higher-educated spouse by 16%. I also find suggestive evidence that marital satisfaction increases when wives earn a higher share of household income. Second, I show that the marriage rate declines by 7%. The reduction in marriage is explained by reduction in first marriages as well as increased propensity to divorce. Third, consistent with the marginal husband being undesirable, I find that 65% of women who do not marry opt to live with a female roommate or alone, and I do not find strong evidence for substitution toward cohabitation. Fourth, I find that a higher relative wage increases women's labor supply.

These results suggest that relative earning power is an influential factor in the family formation and labor market decisions of women. This indicates that reducing the gender wage gap is not only a matter of "fairness in the workplace," but also places men and women on more equal footing in the marriage market. Moreover, this paper provides an important first step toward building an understanding of the welfare effects of the gender wage gap by highlighting empirical channels for improvements in welfare. I leave advancing the evidence on this question for future research.

REFERENCES

Abramitzky, R., A. Delavande, and L. Vasconcelos, "Marrying Up: The Role of Sex Ratio in Assortative Matching," *American Economic Journal: Applied Economics* 3 (2011), 124–157.

Aizer, A., "The Gender Wage Gap and Domestic Violence," *American Economic Review* 100 (2010), 1847–1859.

Angrist, J., "How Do Sex Ratios Affect Marriage and Labor Markets? Evidence from America's Second Generation," *Quarterly Journal of Economics* 117 (2002), 997–1038.

Autor, D. H., D. Dorn, and G. H. Hanson, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review* 103 (2013), 2121–2168.

——— "When Work Disappears: Manufacturing Decline and the Falling Marriage Market Value of Young Men," *American Economic Review: Insights* 1: 2 (2018), 161–178.

Autor, D., D. Figlio, K. Karbownik, J. Roth, and M. Wasserman, "Family Disadvantage and the Gender Gap in Behavioral and Educational Outcomes," Institute for Policy Research at Northwestern University working paper (2015).

Autor, D. H., L. F. Katz, and M. S. Kearney, "Trends in U.S. Wage Inequality: Revising the Revisionists," this REVIEW 90 (2008), 300–323.

Autor, D. H., F. Levy, and R. J. Murnane, "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics* 118 (2003), 1279–1333.

Autor, D. H., and B. Price, "The Changing Task Composition of the US Labor Market: An Update of Autor, Levy, and Murnane (2003)," unpublished manuscript, MIT (2013).

Bacolod, M. P., and B. S. Blum, "Two Sides of the Same Coin U.S. Residual Inequality and the Gender Gap," *Journal of Human Resources* 45 (2010), 197–242.

Bailey, M. J., "More Power to the Pill: The Impact of Contraceptive Freedom on Women's Life Cycle Labor Supply," *e Quarterly Journal of Economics* 121 (2006), 289–320.

Bartik, T. J., *Who Benefits from State and Local Economic Development Policies?* (Kalamazoo, MI: W. E. Upjohn Institute for Employment Research, 1991).

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

Becker, G. S., "A Theory of Marriage: Part I," *Journal of Political Economy* 81 (1973), 813–846.

Bertrand, M., P. Cortes, C. Olivetti, and J. Pan, "Social Norms, Labor Market Opportunities, and the Marriage Gap for Skilled Women," IZA Institute for Labor Economics working paper 11382 (2018).

Bertrand, M., E. Kamenica, and J. Pan, "Gender Identity and Relative Income within Households," *Quarterly Journal of Economics* 130 (2015), 571–614.

Bertrand, M., and J. Pan, "The Trouble with Boys: Social Influences and the Gender Gap in Disruptive Behavior," *American Economic Journal: Applied Economics* 5 (2013), 32–64.

Black, S. E., and C. Juhn, "The Rise of Female Professionals: Are Women Responding to Skill Demand?" *American Economic Review* 90 (2000), 450–455.

Black, S. E., and A. Spitz-Oener, "Explaining Women's Success: Technological Change and the Skill Content of Women's Work," this REVIEW 92 (2010), 187–194.

Blau, F. D., and L. M. Kahn, "Swimming Upstream: Trends in the Gender Wage Differential in 1980s," *Journal of Labor Economics* 15 (1997), 1–42.

——— "The Gender Wage Gap: Extent, Trends, and Explanations," NBER working paper 21913 (2016). doi:10.3386/w21913.

Blau, F. D., L. M. Kahn, and J. Waldfogel, "Understanding Young Women's Marriage Decisions: The Role of Labor and Marriage Market Conditions," *Industrial and Labor Relations Review* 53 (2000), 624–647.

Borusyak, K., P. Hull, and X. Jaravel, "Quasi-Experimental Shift-Share Research Designs," NBER working paper 24997 (2018).

Brainerd, E., "The Lasting Effect of Sex Ratio Imbalance on Marriage and Family: Evidence from World War II in Russia," this REVIEW 99 (2017), 229–242.

Charles, K. K., and M. C. Luoh, "Male Incarceration, the Marriage Market, and Female Outcomes," this REVIEW 92 (2010), 614–627.

Deming D. J., "The Growing Importance of Social Skills in the Labor Market," *Quarterly Journal of Economics* 132 (2017), 1593–1640.

Diamond, R., "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980–2000," *American Economic Review* 106 (2016), 479–524.

Edin, K., and M. J. Kefalas, *Promises I Can Keep: Why Poor Women Put Motherhood before Marriage* (Berkeley: University of California Press, 2005).

Fisman, R., S. S. Iyengar, E. Kamenica, and I. Simonson, "Gender Differences in Mate Selection: Evidence From a Speed Dating Experiment," *Quarterly Journal of Economics* 121 (2006), 673–697.

Geruso, M., and H. Roye, "The Impact of Education on Family Formation: Quasi-Experimental Evidence from the UK," NBER working paper 24332 (2018).

Glynn, S. J., "Breadwinning Mothers, Then and Now," Center for American Progress technical report (2014).

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

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift, "Bartik Instruments: What, When, Why, and How," NBER working paper 24408 (2018).

Gould, E. D., and M. D. Paserman, "Waiting for Mr. Right: Rising Inequality and Declining Marriage Rates," *Journal of Urban Economics* 53 (2003), 257–281.

Hitsch, G. J., A. Hortasu, and D. Ariely, "Matching and Sorting in Online Dating," *American Economic Review* 100 (2010), 130–163.

Katz, L. F., and K. M. Murphy, "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *Quarterly Journal of Economics* 107 (1992), 35–78.

Kearney, M. S., and P. B. Levine, "The Economics of Nonmarital Childbearing and the Marriage Premium for Children," *Annual Review of Economics* 9 (2017), 327–352.

Kearney, M. S., and R. Wilson, "Male Earnings, Marriageable Men, and Nonmarital Fertility: Evidence from the Fracking Boom," this REVIEW 100 (2018), 678–690.

King, M., S. Ruggles, J. T. Alexander, S. Flood, K. Genadek, M. B. Schroeder, B. Trampe, and R. Vick, "Integrated Public Use Microdata Series, Current Population Survey: Version 3.0" (2010). Machine-readable database.

Knowles, J., "Why Are Married Men Working So Much? An Aggregate Analysis of Intra-Household Bargaining and Labor Supply," *Review of Economic Studies* 80 (2013), 1055–1085.

Lee, W.-S., and T. McKinnish, "The Marital Satisfaction of Differently-Aged Couples," *Journal of Population Economics* 31 (2017), 337–362.

Loughran, D. S., "The Effect of Male Wage Inequality on Female Age at First Marriage," this REVIEW 84 (2002), 237–250.

Low, C., "Pricing the Biological Clock: Reproductive Capital on the US Marriage Market," unpublished manuscript, University of Pennsylvania (2016).

Lundberg, S., R. A. Pollak, and J. Stearns, "Family Inequality: Diverging Patterns in Marriage, Cohabitation, and Childbearing," *Journal of Economic Perspectives* 30 (2016), 79–102.

Mansour, H., and T. McKinnish, "Who Marries Differently Aged Spouses? Ability, Education, Occupation, Earnings, and Appearance," this REVIEW 96 (2013), 577–580.

McCrary, J., and H. Royer, "The Effect of Female Education on Fertility and Infant Health: Evidence from School Entry Policies Using Exact Date of Birth," *American Economic Review* 101 (2011), 158–195.

Munyo, I., and M. Rossi, "Real Exchange Rate, the Wage Gender Gap and Domestic Violence," unpublished manuscript, Inter-American Development Bank (2015).

Ong, David, Yu (Alan) Yang, and Junsen Zhang, "Hard to Get: The Scarcity of Women and the Competition for High-Income Men in Urban China," *Journal of Development Economics* 144 (2020), 102434, https://doi.org/10.1016/j.jdeveco.2019.102434.

Pollak, R. A., "Bargaining Power in Marriage: Earnings, Wage Rates and Household Production," NBER working paper 11239 (2005).

Rotz, D., "Why Have Divorce Rates Fallen?" *Journal of Human Resources* 51 (2016), 961–1002.

Ruggles, S., J. T. Alexander, S. Flood, R. Goeken, M. B. Schroeder, and M. Sobek, "Integrated Public Use Microdata Series: Version 5.0" (2010), machine-readable database.

Weinberg, B. A., "Computer Use and the Demand for Female Workers," *Industrial and Labor Relations Review* 53 (2000), 290–308.