ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

An International Society for the Advancement of Economic Theory in its Relation to Statistics and Mathematics

http://www.econometricsociety.org/

Econometrica, Vol. 87, No. 1 (January, 2019), 217–253

CAREER AND FAMILY DECISIONS: COHORTS BORN 1935–1975

ZVI ECKSTEIN
Tiomkin School of Economics, IDC Herzliya

MICHAEL KEANE
School of Economics, University of New South Wales

OSNAT LIFSHITZ
Tiomkin School of Economics, IDC Herzliya

Econometrica, Vol. 87, No. 1 (January, 2019), 217–253

CAREER AND FAMILY DECISIONS: COHORTS BORN 1935–1975

ZVI ECKSTEIN
Tiomkin School of Economics, IDC Herzliya

MICHAEL KEANE
School of Economics, University of New South Wales

OSNAT LIFSHITZ
Tiomkin School of Economics, IDC Herzliya

Comparing the 1935 and 1975 U.S. birth cohorts, wages of married women grew twice as fast as for married men, and the wage gap between married and single women turned from negative to positive. The employment rate of married women also increased sharply, while that of other groups remained quite stable. To better understand these diverse patterns, we develop a life-cycle model incorporating individual and household decisions about education, employment, marriage/divorce, and fertility. The model provides an excellent fit to wage and employment patterns, along with changes in education, marriage/divorce rates, and fertility. We assume fixed preferences, but allow for four exogenously changing factors: (i) mother's education, health, and taxes/transfers; (ii) marriage market opportunities and divorce costs; (iii) the wage structure and job offers; (iv) contraception technology. We quantify how each factor contributed to changes across cohorts. We find that factor (iii) was the most important force driving the increase in relative wages of married women, but that all four factors are important for explaining the many socio-economic changes that occurred in the past 50 years. Finally, we use the model to simulate a shift from joint to individual taxation. In a revenue-neutral simulation, we predict this would increase employment of married women by 9% and the marriage rate by 8.1%.

KEYWORDS: Labor supply, marriage market, gender wage gap, education, fertility, life-cycle model, human capital, taxation, contraception, divorce laws.

1. INTRODUCTION

WAGES AND EMPLOYMENT OF MARRIED WOMEN have grown rapidly over the past 50 years, both in absolute terms and relative to other demographic groups. In 1962, married women earned on average 15% less than single women, but their wages caught up to and surpassed those of single women in the mid-1990s, so by 2015, they earned 18% more. The wages of married women also grew much faster than those of either married or single men.[1]

Zvi Eckstein: zeckstein@idc.ac.il
Michael Keane: m.keane@unsw.edu.au
Osnat Lifshitz: osnat.lifshitz@idc.ac.il

Previous versions of this paper were circulated under the titles "Wages and Employment of Married vs. Unmarried Individuals: Cohorts born 1935–1975" and "Sources of Change in the Life-Cycle Decisions of American Men and Women: 1962–2014." The authors thank the referees for their constructive comments. Special thanks to Richard Blundell, Yona Rubinstein, Victor Rios-Rull, and Ken Wolpin and seminar participants at Barcelona, Berkeley, Chicago, IDC-Herzliya, Northwestern, Oxford, Pennsylvania, Princeton, Rice, Stanford, UCL, and Tel Aviv University. Keane's work on this project was supported by Australian Research Council Grants FL110100247 and CE110001029.

[1]The data here refer to Caucasian males and females aged 22–65 in the March CPS survey (see Appendix A of the Supplemental Material (Eckstein, Keane, and Lifshitz (2019))).

 https://doi.org/10.3982/ECTA14474

Broadly speaking, employment rates of single men and women, as well as married men, were all fairly stable over the whole 1962 to 2015 period.[2] In contrast, the employment rate of married women doubled from 30% in 1962 to about 60% in 1995, almost approaching that of single men/women. Since 1995, the employment rate of married women has stabilized.

The increase in relative wages and employment of married women coincided with an increase in their education relative to singles. But we find education alone cannot explain the increase in married women's relative wages. Selection of women into marriage based on *unobserved* labor market skills has also gone from strongly negative to positive.

The focus of this paper is to unravel the causes of these puzzling differences in trends of wages and employment of married women relative to other demographic groups over the 1962–2015 period by analyzing behavior of cohorts born from the 1930s to the 1970s. Potential explanations include: changing selection into marriage, changing incentives to invest in human capital, increased demand for female labor, changing divorce laws, changes in availability of oral contraception, and home production technology. Thus, to understand the changes in wages and employment that occurred over this period, one needs a model that includes marriage market sorting along with individual and household decisions about education, marriage/divorce, employment, and fertility.

Of course, prior literature contains many papers that address some aspects of these phenomena. For example, explanations of why wages/employment of married women have increased include: reduced gender wage discrimination (Jones, Manuelli, and McGrattan (2015)), birth control (Goldin and Katz (2002)), changing divorce laws (Stevenson (2008), Voena (2015), Bronson (2015), increased marriage market returns to female human capital (Low (2014)), growth in "pink collar" occupations (Lee and Wolpin (2006)), improved home technology (Greenwood and Seshadri (2005), Greenwood, Seshadri, and Yorukoglu (2005)), lower child care costs (Attanasio, Low, and Sánchez-Marcos (2008), and changing social norms (Fernandez (2013), Gousse, Jacquemet, and Robin (2017)). Our main contribution is to explain the changing patterns of wages/employment for married women versus other demographic groups in a unified framework where education, marriage, fertility, and labor supply are all treated as endogenous choices.

Specifically, we develop a model that extends the gender-specific life-cycle models of Keane and Wolpin (1997, 2010) into a unified framework of individual and family decisions following the cooperative household model, as in Chiappori (1988, 1992), Mazzocco, Ruis, and Yamaguchi (2007), and Gemici and Laufer (2011). Men and women make decisions as individuals from age 17 to 65, but they also interact with each other in the marriage market, as in Becker (1974, 1981). We incorporate human capital accumulation and exogenous changes in health, taxes, and welfare rules over time. We discipline the analysis by holding preference parameters fixed across cohorts, which enables us to achieve identification via changes in these exogenous factors. The model is estimated using March CPS data on five cohorts born in 1935, 1945, 1955, 1965, and 1975.

Given *fixed* preferences, we show that our model can explain changes across all five cohorts (and demographic groups) in five key outcomes (i.e., labor supply, education, wages, marriage/divorce, and fertility) using four sets of exogenous factors: (i) parental education, health, and taxes/transfers, (ii) marriage market opportunities and divorce costs, (iii) the wage offer distribution, and (iv) birth control. Thus, *given our model*, these four factors are *sufficient* to rationalize the data. Furthermore, we find that dropping any one

[2]Employment rates of single men and women hovered around 70% to 80% for the whole period, while that for married men hovered around 80% to 90% (see Figure 1).

factor leads to substantial deterioration in fit, so all four factors are *necessary* to rationalize the data.[3,4]

Between the 1935 and 1975 cohorts, real full-time wages of married women increased more than twice as much as for married men, and our model captures this well. To assess the contribution of each factor to this dramatic shift, we conduct counterfactual experiments where we allow only some exogenous factors to change across cohorts, while holding other factors fixed. Based on this exercise, our model implies that increasing returns to women's education and experience explain about 75 to 80% of their increase in relative wages. Other factors, including the increase in women's education, played relatively minor roles.

What drives this result is that our model estimates imply that returns to both education and work experience were very low for women in the 1935 cohort. But by the 1975 cohort, the offer wage distribution of women had nearly converged to that for men. The sharp increase in education of married women would have done little in itself to increase their relative wages, given the low returns to education for women that existed in the 1935 cohort.

Our model also captures the fact that, from the 1935 to 1975 cohorts, the employment rate of married women aged 25 to 34 increased sharply from 28% to 63% (while employment of men and single women was fairly stable). Our model attributes roughly 1/2 of this increase to the dramatic improvement in the offer wage distribution facing women (along with that of job offers). The second most important factor, accounting for 1/4 of the increase, was the advent of oral contraception. As for the sharp increase in wages of married versus single women, we find that 2/3 was due to selection of higher ability women into marriage, while 1/3 was due to married women working more and acquiring more human capital.

Our model also captures numerous other major socio-economic changes that occurred over the sample period. Consider the dramatic increase in the fraction of women with at least a college degree, from only 5% in the 1935 cohort to 36% in the 1975 cohort. The model implies that changes in mother's education, changes in the marriage market, and changes in the wage structure each account for roughly 1/3 of the increase in women's education, with contraception playing little role. The drop in marriage rates was largely due to higher mother's education and better labor market returns, but the increase in the divorce rate was mostly driven by changing divorce costs. Availability of oral contraception explains about half the drop in fertility for married women and almost the entire drop for unmarried women.

Finally, we use our model to simulate a scenario where the United States shifts from joint taxation of couples to purely individual taxation. In a revenue-neutral simulation, we predict this would increase employment of married women by 9%, while leaving other groups only slightly affected. It would also increase the marriage rate by 8.1%, reduce the divorce rate by 5.1%, and increase college completion for women by 4.2%. Our ability to predict how such a change in tax policy would affect not just behavior of existing married couples, but also marriage and divorce rates and education, highlights the value of modeling this complex set of endogenous variables in a unified framework.

---

[3]Clearly, fitting the data by age, cohort, and demographic group is equivalent to fitting the aggregate data.

[4]Of course, showing our four factors are necessary and sufficient to fit the data *given our model* does not imply they are necessary and sufficient in a formal logic sense. Hypothetical alternative models might rationalize all these data patterns in other ways, using different factors. But to our knowledge, no existing model has done this.

The outline of the paper is as follows. Section 2 describes data patterns that motivate our analysis. Section 3 presents our model, and Sections 4 and 5 discuss solution, estimation, and identification. Sections 6 and 7 present empirical results, and Section 8 concludes.

2. KEY PATTERNS IN THE DATA

Our data are the March CPS from 1962 to 2015. We discuss the construction of the data in detail in Appendix A of the Supplemental Material (Eckstein, Keane, and Lifshitz (2019)). Here we document employment and wage patterns of white males and females aged 22–65. We define "employment" as working at least 10 hours a week.

2.1. *Employment Rates by Marital Status*

Figure 1A reports employment over time for women. The most striking change is that the employment rate of married women doubled from 30% in 1962 to about 60% in 1995. It then plateaus, and hovers in the low 60% range from 1994 to 2015. In contrast, employment rates of single and divorced women, who behave quite similarly, are remarkably stable. They hover around 70% throughout the whole period.[5] Thus, from 1962 to 1995, roughly 3/4 of the employment gap between married and unmarried women was eliminated.[6]

Figure 1B reports employment rates for men. The dominant picture is one of stability. The employment rate of both single and divorced men hovers around 75% for the whole period, decreasing somewhat after 2006 (due to the Great Recession).[7] Comparing Figures 1A and 1B, the similarity in employment patterns of single men and women is striking.

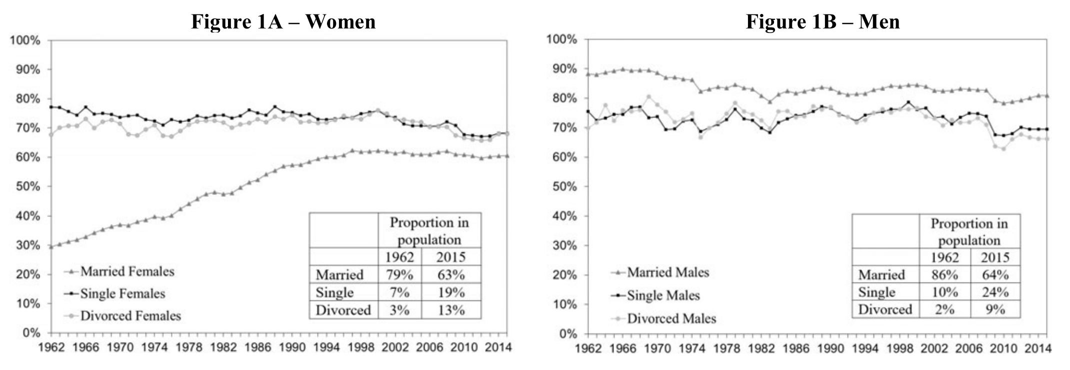

FIGURE 1.—Employment rate by marital status: 1962–2015. *Note*: Fraction employed of the Caucasian population aged 22–65. We define employed as working at least 10 hours a week.

[5]Here, cohabitation is counted as unmarried. We do not have data prior to 1995, but in 1995, 1.6% of the population was cohabiting, and by 2015, it increased to 3.8%. In the cohort of 1975, 4.7% are cohabiting.

[6]The proportion of married women working full time (35+ hours per week) doubled in this period (22% to 44%), while for single (divorced) women, the proportion working full time dropped from 70% (59%) to 50% (53%). For details, see the supplemental web site at http://www1.idc.ac.il/Faculty/Eckstein/EKL.html.

[7]Cyclical patterns are clear in Figure 1. In particular, the Great Recession (2009–2011) caused employment rates to fall for all groups. Cyclical patterns are strongest for single and divorced men, and weakest for married women.

In contrast, married men work more than single men or women. Their employment rate was near 90% in the 1960s. It fell in the recessions of 1974–1975 and 1980–1982. After that, it has hovered in the low 80% range, and never returned to its 1960s levels. Still, if the small decline for married men is compared with the dramatic increase for married women, the employment rate of married men seems relatively quite stable.

The aggregate data are a sum over different cohorts. In Appendix C of the Supplemental Material, we report results separately for the '35, '45, '55, '65, and '75 birth cohorts, defined as including people born 2 years on either side of the birth year. Strikingly, the employment rate increased over these cohorts *only* for married women (see Figure C1). There were no substantial differences across cohorts for single men/women or married men. It is particularly striking that married women in the 1955 (and later) cohorts stay much more attached to the labor force during their prime childbearing years than they did in earlier cohorts. Also notable is that married women in post-1955 cohorts behave very similarly. The historic increase in employment of married women in the 20th century was essentially complete by the 1965 cohort. This is consistent with the flatness in aggregate employment of married women after 1995 (see Figure 1A).

A major challenge for any model seeking to explain the great increase in employment of married women is to *simultaneously* explain the remarkable stability of employment rates of unmarried women and men, and the relatively small decline in employment for married men. Another challenge is to capture the timing, that is, that the increase in married women's employment was complete by the 1965 cohort (or by 1995 in the aggregate data).

2.2. *Wages by Marital Status*

Figure 2 plots annual average earnings of full-time workers by gender and marital status. In 1962, the full-time wage of single women was 15% higher than that of married women. But by 2015, this is reversed, and the full-time wage of married women is 18% higher than for singles. Wages of married women passed those of single women in 1992–1993.[8]

Figure 2B shows earnings patterns for men. In contrast to women, the ranking of married, single, and divorced men by the average annual full-time wage is very stable. But

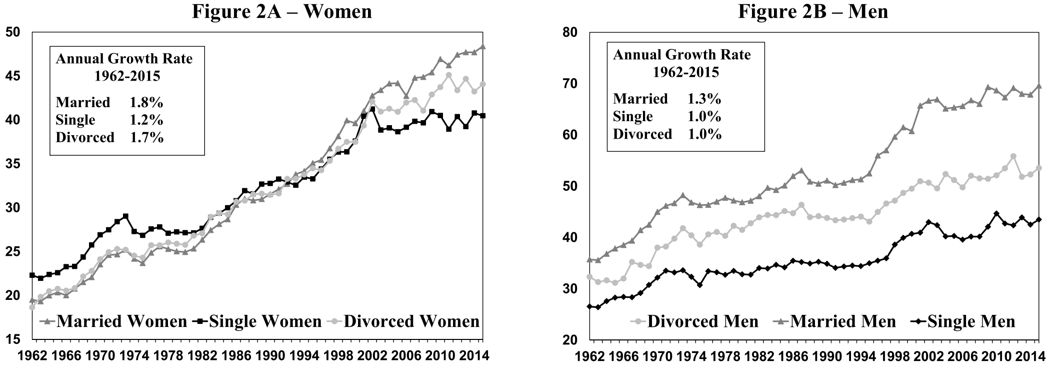

FIGURE 2.—Annual wages by marital status: 1962—2015. *Note*: Real annual wages (in thousands of dollars) of full-time full-year Caucasian workers aged 22–65 (2012 prices). For details, see Appendix A.

[8]An interesting detail is that until the mid-1990s, divorced women look very much like married women, but then their relative wage growth slows, and by 2015, their average wage is between those of married and singles.

the magnitudes of the earnings gaps grow over time. In 1962, the annual wage of married men was 33% greater than that of single men, but by 2015, it was roughly 66% greater.

2.3. *Women's Education and the "Marriage Wage Gap"*

During the 1962–2015 period, married women became much better educated relative to unmarried women. In 1962, only 7% had a college degree (or higher), compared to 10% for unmarried women. By 2015, this pattern had reversed, and 36% of married women had a college degree, compared to only 28% of unmarried women.[9] So increased education and work experience (Section 2.1) may explain the growth in relative wages of married women.

To assess this issue, we estimated standard Mincer (1974) earnings equations that control for "potential experience" (i.e., years since leaving school), and its square, as well as education, and a marriage dummy. The coefficient on the latter gives the wage gap between married and single workers *conditional* on potential experience and education—the so-called "marriage wage gap." These results are reported in Appendix D of the Supplemental Material, separately by gender and cohort.

For women, the unconditional married/unmarried wage gap was −12.0% in the 1935 cohort, but it was eliminated in the 1965 cohort, and became +7.8% in the 1975 cohort. The conditional wage gaps obtained from the Mincer equation are −8.9% in the 1935 cohort and +5.2% in the 1975 cohort. Thus, changes in education and potential experience can only account for about 30% of the increase in relative wages of married versus single women.[10]

These results imply that most of the reversal in the marriage wage gap was due to (i) changes in the mapping of potential to actual experience (as married women worked more) and/or (ii) changes in the unobserved characteristics of women who select into marriage. A key focus of this paper is to develop a life-cycle model that can match and interpret these facts, especially the changes in female wages, employment, and selection into marriage.

3. A LIFE-CYCLE MODEL OF EDUCATION, LABOR SUPPLY, MARRIAGE/DIVORCE, AND FERTILITY

In our model, men and women start at age 17 as single individuals in school. They make annual private decisions about school continuation, work, and, in the case of women, fertility. We assume only single people can attend school.[11] Retirement is enforced after age $T = 65$, at which point agents receive a terminal value function. The men and women in the model also interact in a marriage market, so they can choose to form (and later dissolve) couples. Once a couple is formed, decisions about labor supply and fertility are made jointly.

To make marriage decisions, individuals compare the values of the married and single states. Thus, we first describe the problem of married couples, followed by that of single individuals. We are then in a position to explain how we model the marriage market.

[9]Data on education levels for both men and women by marital status can be found in Appendix C, Figure C2.

[10]In contrast to women, the married/unmarried wage gap for men changed little over the past 50 years. It was +20% for the 1935 cohort and +18% for the 1975 cohort. See Appendix D.

[11]School attendance by married people is rare. We rule out school attendance after age 30 for the same reason.

3.1. *The Decisions of a Married Couple*

We assume a collective model of household decision making. Let t denote the annual time period, and let $j = f, m$ denote gender. Individuals have time endowments of one unit per period. This is split between market work (h), and leisure (l), hence, $h_t^j + l_t^j = 1$. Agents can choose to work full-time, part-time, or not at all. Thus, the labor supply choice each period is $h_t^j \in \{1, 0.5, 0\}$, corresponding to a leisure choice of $l_t^j \in \{0, 0.5, 1\}$. We assume that full-time and part-time work correspond to 2000 and 1000 annual hours, respectively.

Conditional on marriage, couples have three choice variables: the levels of leisure for the husband and wife, $\{l_t^m, l_t^f\}$, and pregnancy, $p_t \in \{0, 1\}$. We assume pregnancy leads deterministically to arrival of a child at $t+1$. Letting X_t^j denote work experience, and N_t denote the number of children under 18, the laws of motion for these state variables are $X_{t+1}^j = X_t^j + h_t^j$ for $j = m, f$ and $N_{t+1} = N_t + p_t - p_{t-18}$. In addition, couples make annual decisions about whether to remain married or get divorced. We ignore this for the time being and focus on the joint decisions of couples *conditional* on marriage.

3.1.1. *Preferences and Constraints*

Married couples have total gross income GY_t^M given by the equation

$$GY_t^M = \left(w_t^m h_t^m + w_t^f h_t^f\right) + b_m I\left[h_t^m = 0\right] + b_f I\left[h_t^f = 0\right]. \tag{1}$$

Here w_t^j and h_t^j for $j = f, m$ are annual full-time wage rates and the b_j are unemployment benefits plus values of home production. We will use the M superscript throughout to indicate values for married individuals. Net income is Y_t^M given by the equation

$$Y_t^M = GY_t^M - \tau_t^M\left(\left(w_t^m h_t^m + w_t^f h_t^f\right), N_t\right), \tag{2}$$

where $\tau_t^M(\cdot, \cdot)$ is the tax function for married couples based on the time t tax rules. We model the U.S. federal tax system in detail, including deductions, exemptions, EITC, and the joint taxation of couples (see Appendix B of the Supplemental Material). We assume perfect foresight regarding tax rules.

The household budget constraint takes the form

$$C_t^M = \left(1 - \kappa(N_t)\right) Y_t^M. \tag{3}$$

Here $\kappa(N_t)$ is the fraction of Y_t^M spent on children, based on a square root equivalence scale.[12] We assume a static budget constraint, as it is computationally infeasible to add saving in addition to our other state variables. However, the terminal value function (obtained immediately after age 65) proxies for how labor supply affects Social Security and retirement assets, so these key aspects of savings do enter our model in a reduced form way. Furthermore, we also consider robustness of our results to inclusion of a simple form of short-run consumption insurance.

[12]For a household with two adults, the square root scale implies that $\kappa(N) = 1 - \sqrt{2/(2+N)}$. Thus, $\kappa(N) = 0.194, 0.293, 0.367$ and 0.423 if $N = 1, 2, 3$ or 4, respectively.

The period utility of a married person of age t and gender j in state Ω_{jt} is given by[13]

$$\begin{aligned} U_t^{jM}(\Omega_{jt}) = {} & \frac{1}{\alpha}\left(\psi C_t^M\right)^\alpha + L_{jt}\left(l_t^j\right) + \theta_t^M + \pi_t^M p_t \\ & + A_j^M Q\left(l_t^f, l_t^m, Y_t^M, N_t\right), \quad j = m, f, \end{aligned} \tag{4a}$$

$$L_{jt}\left(l_t^j\right) = \frac{\beta_{jt}}{\gamma}\left(l_t^j\right)^\gamma + \mu_{jt} l_t^j, \quad \gamma < 1, \alpha < 1. \tag{4b}$$

The first term in equation (4a) is a CRRA in consumption with curvature parameter α. We assume household consumption C_t^M is a "public" good. That is, the full amount C_t^M enters the utility of both the husband and wife. The parameter $\psi \in (1/2, 1)$ captures household economies of scale in consumption.[14] The square root equivalence scale gives $\psi = 1/\sqrt{2} = 0.707$, so a couple needs 41% more expenditure than a single person to obtain an equivalent consumption level.

The second term in equation (4a) captures the value of leisure and home production. The third term (θ_t^M) is the utility the individual derives from marriage itself (i.e., the match quality), the fourth term captures the utility (or dis-utility) from pregnancy ($p_t = 1$), and the fifth term captures the utility a couple receives from the quality and quantity of children.

3.1.1.1. *Tastes for Leisure and Value of Home Production.* As shown in equation (4b), the second term in (4a) consists of two parts. The first is a CRRA in leisure with curvature parameter γ. The parameter β_{jt}, which must be positive, shifts tastes for leisure. For women, we allow β_{jt} to depend on p_t, while for both men and women, we allow β_{jt} to depend on education and on health status.

The second term in (4b) captures stochastic variation in the marginal utility of leisure. This is denoted by $\mu_{jt} l_t^j$, where μ_{jt} is a random variable. Our specification of the stochastic process for μ_{jt} is an important and novel aspect of our model. Specifically, we assume that

$$\ln(\mu_{jt}) = \tau_{0j} + \tau_{1j}\ln(\mu_{j,t-1}) + \tau_{2j} p_{t-1} + \varepsilon_{jt}^l \quad \text{where } \varepsilon_{jt}^l \sim \text{iidN}\left(0, \sigma_\varepsilon^l\right), \tag{5}$$

where $0 < \tau_{1j} < 1$. Thus, shocks to tastes for leisure (i.e., home time) follow a stationary AR(1) process. Importantly, the arrival of a new child at time t (i.e., $p_{t-1} = 1$) leads to a shift in tastes for home time (τ_{2j}). This may capture a desire to spend time with children as well as an increase in the time required for home production due to the presence of children.

In practice, we would expect that, particularly for women, the marginal utility of home time will jump up when a newborn arrives (i.e., $\tau_{2f} > 0$). Afterward, provided that no new children arrive, tastes for home time will gradually revert back to normal, as $\tau_{1f} < 1$. This mechanism enables the model to generate relatively large declines in employment for women after childbirth, as well as their subsequent gradual return to the labor force as children grow older. The stochastic terms ε_{jt}^l generate heterogeneity in individual response patterns.

[13]The state vector Ω_{jt} contains four variables that are relevant for $U_t^{jM}(\Omega_{jt})$. These are N_t and μ_{jt}, as well as education and health (which shift β_{jt}). The state vector Ω_{jt} contains several additional variables, whose role will only become clear after the full model is laid out. Thus, we defer giving the complete list of elements of Ω_{jt} until we finish expositing the full model and turn to discussing the DP problem solution (Section 4).

[14]If $\psi = 1/2$ there are no economies of scale, while if $\psi = 1$ then C_t^M is a pure public good for the couple.

The specification in (5) has important advantages over prior approaches to modeling fertility. In dynamic models with endogenous fertility, a person's state vector contains the number of children at each specific age. As a result, the number of possible state vectors is astronomical (see Geweke and Keane (2001) for a discussion). Our approach circumvents this problem, as the state space in our model only contains the scalar variables μ_{jt} and N_t, which leads to tremendous computational savings (see Appendix E2 for further details).

3.1.1.2. *Match Quality, Utility of Pregnancy, Child Production.* In this section, we discuss in detail the last three terms in equation (4a): the utility from marriage, the dis-utility from pregnancy ($p_t = 1$), and the utility from children. First, consider the utility from marriage (θ_t^M), that is, the match quality. We write

$$\theta_t^M = d_1 + d_2 \cdot I\left[E^m - E^f > 0\right] + d_3 \cdot I\left[E^f - E^m > 0\right] + d_4\left(H_t^m - H_t^f\right)^2 + \varepsilon_t^M, \tag{6}$$

where $\varepsilon_t^M \sim \text{iidN}(0, \sigma_\varepsilon^M)$ and E^j denotes education, rank ordered as high school dropout (HSD), high school (HSG), some college (SC), college (CG), and post-college (PC), and $H_t^j \in \{1, 2, 3\}$ denotes health (i.e., good, fair, or poor). The second and third terms capture assortative mating on education. $I[E^m - E^f > 0]$ is an indicator for the man having greater education than the woman, and $I[E^f - E^m > 0]$ is an indicator for the reverse. For example, if $d_3 < 0$, people are averse to matches where the woman has more education. The fourth term captures assortative mating based on health. If $d_4 < 0$, people prefer matches where the partners have similar health. Finally, ε_t^M is a transitory shock to match quality.

Next, consider the utility from pregnancy, π_t. We specify that

$$\pi_t = \pi_1 M_t + \pi_2 H_{ft} + \pi_3 N_t + \pi_4 p_{t-1} + \varepsilon_t^p + \exp\left(\varepsilon_t^{up}\right), \tag{7}$$

where $\varepsilon_t^p \sim \text{iidN}(0, \sigma_\varepsilon^p)$ and $\varepsilon_t^{up} \sim \text{iidN}(pr, 1)$. Here π_t is a function of marital status, where M_t is a 1/0 indicator for marriage, the woman's health, the number of already present children, and lagged pregnancy. The presence of health helps generate that fertility declines with age.

The error term in (7) is designed to help capture changes in contraception technology. It consists of ε_t^p, a Gaussian shock to tastes for pregnancy, along with the positive shock $\exp(\varepsilon_t^{up})$. The latter may lead to unexpected jumps in utility from pregnancy, which in our model is observationally equivalent to random failures of contraception. The cohort-specific parameter pr captures the prevalence of such shocks in each cohort.

Note that equation (7) contains nothing individual-specific. We assume pregnancy decisions are made jointly by the couple, and each party gets the same utility from the decision. Of course, one could imagine individuals in a couple getting different utilities from a pregnancy decision, but we cannot infer such differences from the data, so we ignore them.

Finally, consider the function $Q(\cdot)$ that determines the utility a couple receives from the quality and quantity of children. We assume it is a CES function of the inputs, as follows:

$$Q\left(l_t^f, l_t^m, Y_t^M, N_t\right) = \left(a_f\left(l_t^f\right)^\rho + a_m\left(l_t^m\right)^\rho + a_g\left(\kappa(N_t)Y_t^M\right)^\rho + (1 - a_f - a_m - a_g)N_t^\rho\right)^{1/\rho}. \tag{8}$$

The first three inputs, which are home time of parents and spending per child, $\kappa(N_t)Y_t^M$, all increase child quality. The parameter A_j^M in the utility function (4a) is a scale parameter that multiplies $Q(\cdot)$. This parameter is allowed to differ in the single state (see below).

3.1.2. *Value Function of a Married Couple*

We are now able to write the choice-specific value functions for married *individuals*. These depend on both a person's own state and that of their partner:

$$\begin{aligned} & V_t^{jM}\big(l_t^m, l_t^f, p_t \mid \Omega_{mt}, \Omega_{ft}\big) \\ & \quad = \frac{1}{\alpha}\big(\psi C_t^M\big)^{\alpha} + L\big(l_t^j\big) + \theta_t^M + \pi_t p_t + A_j^M Q\big(l_t^f, l_t^f, Y_t^M, N_t\big) \\ & \quad\quad + \delta E_{\mathrm{MAX}}\big(M_{t+1} V_{t+1}^{jM}(\Omega_{m,t+1}, \Omega_{f,t+1}) + (1 - M_{t+1}) V_{t+1}^{j}(\Omega_{j,t+1})\big), \quad j = f, m. \end{aligned} \tag{9}$$

The current payoff simply reproduces (4a)–(4b). The future component in (9) consists of two parts, corresponding to whether the marriage continues at $t+1$ or not. The term $V_{t+1}^{jM}(\Omega_{m,t+1}, \Omega_{f,t+1})$ is the value of next period's state for partner j if the marriage continues. The newly defined term $V_{t+1}^{j}(\Omega_{jt+1})$ is the value of next period's state for partner j if he/she becomes single (i.e., a divorce occurs). We discuss divorce and the value functions for single persons below.

The $t+1$ state depends on the current state $\{\Omega_{mt}, \Omega_{ft}\}$ and current choices $\{l_t^m, l_t^f, p_t\}$ via the laws of motion of the state variables. δ is the discount rate and $E_{\mathrm{MAX}}(\cdot)$ is the expectation taken over elements of the $t+1$ state that are unknown at t. These include M_{t+1}, $\{\varepsilon_{jt+1}^{l}\}$ for $j = m, f$, ε_{t+1}^{M}, ε_{t+1}^{p} and ε_{t+1}^{up}, as well as realizations of wage shocks and job offers. We defer a detailed discussion of these until Section 3.3, which describes the labor market.

3.1.3. *Household Decision Making for Married Couples*

In our collective model, the household value function is given by

$$\begin{aligned} V_t^M\big(l_t^m, l_t^f, p_t \mid \Omega_{mt}, \Omega_{ft}\big) = {} & \lambda V_t^{fM}\big(l_t^m, l_t^f, p_t \mid \Omega_{mt}, \Omega_{ft}\big) \\ & + (1-\lambda) V_t^{mM}\big(l_t^m, l_t^f, p_t \mid \Omega_{mt}, \Omega_{ft}\big). \end{aligned} \tag{10}$$

Here λ and $(1-\lambda)$ are Pareto weights. We set $\lambda = 0.5$ for simplicity.[15] The V_t^{jM} for $j = f, m$ are the choice-specific value functions of the *individual* married partners. The Ω_{jt} for $j = f, m$ are the state vectors of these individuals. Couples seek a choice vector $\{l_t^m, l_t^f, p_t\}$ to maximize (10).

The maximization of (10) is subject to the constraint that both parties prefer marriage over the outside option of divorce.[16] Let $V_t^m(\Omega_{mt})$ and $V_t^f(\Omega_{ft})$ denote the maximized value functions of single males and females in period t. Utility is not transferable, so a divorce occurs if the value of the outside (single) option exceeds the value of marriage for

[15]In theory and empirical applications, there are several ways to model the Pareto weight (e.g., Mazzocco (2007)). Our simple specification here is similar to that of Voena (2015), who considered a household planning problem with a unilateral divorce regime. We discuss our decision to stay with this simple specification in Section 5.

[16]If we take the *unconstrained* maximum of (10), we might obtain a solution for $\{l_t^m, l_t^f, p_t\}$ where only one party prefers to stay married. In a transferable utility framework, marriage may be supportable in such a case, using transfers between partners. We do not adopt this approach, as it is not clear how such transfers would be enforceable after a couple agrees to remain in the married state.

either party. Let $\mathcal{F}$ denote the feasible set of choice options. A choice vector $\{l_t^m, l_t^f, p_t\} \in \mathcal{F}$ if

$$V_t^{jM}\big(l_t^m, l_t^f, p_t \mid \Omega_{mt}, \Omega_{ft}\big) \geq V_t^j(\Omega_{jt}) - \Delta_{jt} \quad \text{for } j = f, m, \tag{11}$$

where Δ_{jt} is the cost of divorce.[17] If $\mathcal{F} = \emptyset$, no choice vector $\{l_t^m, l_t^f, p_t\}$ satisfies (11).[18,19]

We can now formally define the solution to the maximization problem. Denote the vector of household choices that maximize equation (10) as $\{l_t^{m*}, l_t^{f*}, p_t^*\}$. That is,

$$\{l_t^{m*}, l_t^{f*}, p_t^*\} = \begin{cases} \arg \max\limits_{\{l_t^m, l_t^f, p_t\} \in \mathcal{F}} V_t^M\big(l_t^m, l_t^f, p_t \mid \Omega_{mt}, \Omega_{ft}\big) & \text{if } \mathcal{F} \neq \emptyset, \\ \emptyset & \text{if } \mathcal{F} = \emptyset. \end{cases}$$

The form of (10) ensures that $\{l_t^{m*}, l_t^{f*}, p_t^*\}$ is a Pareto efficient allocation. If one or more parties prefer to remain single for all possible $\{l_t^m, l_t^f, p_t\}$, then $\mathcal{F} = \emptyset$ and a divorce occurs.

The maximized value function of a married *individual* in state Ω_{jt} is given by

$$V_t^{jM}(\Omega_{mt}, \Omega_{ft}) \equiv \begin{cases} V_t^{jM}\big(l_t^{m*}, l_t^{f*}, p_t^* \mid \Omega_{mt}, \Omega_{ft}\big) & \text{for } j = f, m \text{ if } \mathcal{F} \neq \emptyset, \\ -\infty & \text{for } j = f, m \text{ if } \mathcal{F} = \emptyset. \end{cases} \tag{12}$$

The maximized value function depends on both the own state Ω_{jt} and that of the partner. This dependence arises from two sources: (i) *individual* choice-specific value functions of married people $V_t^{jM}(l_t^m, l_t^f, p_t \mid \Omega_{mt}, \Omega_{ft})$ depend on both own and partner states, and (ii) the vector $\{l_t^{m*}, l_t^{f*}, p_t^*\}$ is a joint decision made by the couple, so it depends on $\{\Omega_{mt}, \Omega_{ft}\}$. Note also that if $\mathcal{F} = \emptyset$, then no action exists such that person j can be married at time t, so a divorce occurs. Then $V_t^{jM} = -\infty$, so behavior is governed solely by the single value function $V_t^j(\Omega_{jt})$.

3.2. *The Decisions of Single Households*

In this section, we describe the optimization problems of single (i.e., unmarried) men and women. To begin, note that the gross income of a single person is simply

$$GY_t^j = w_t^j h_t^j + b_j \cdot I\big[h_t^j = 0\big] + cb_t(N_t) \cdot I(j = f, N_t > 0) \quad \text{for } j = m, f. \tag{13}$$

As in equation (1), the w_t^j for $j = f, m$ are annual full-time wage rates, $h_t^j \in \{0, 0.5, 1\}$ are hours of work levels, and the b_j are unemployment benefits plus values of home time. The term $cb_t(N_t)$ is a function designed to capture the array of social benefits targeted at single mothers in the United States. These include AFDC/TANF benefits, public housing, child care subsidies, etc.

Modeling welfare benefits and take-up is extremely complex (see Keane and Moffitt (1998), Keane and Wolpin (2010)). Hence, we treat $cb_t(N_t)$ as an exogenous stochastic

---

[17] The cost of divorce depends on the number of children, $\Delta_{jt} = \alpha_4^j + \alpha_5^j \mathrm{N}_t$, where parameters α_4^j and α_5^j may change over time due to changing divorce laws.

[18] Formally, we can write $\mathcal{F} = \{l_t^m, l_t^f, p_t \mid V_t^M(l_t^m, l_t^f, p_t \mid \Omega_{mt}, \Omega_{ft}) \geq V_t^j(\Omega_{jt}) - \Delta_{jt} \text{ for } j = m, f\}$.

[19] In Section 3.3, we discuss additional constraints on the feasible set $\mathcal{F}$ that arise via the labor market. Every period, workers receive job offers probabilistically. Hence, not every choice $\{l_t^m, l_t^f\}$ is available. These labor market constraints on the choice set can affect marriage and divorce probabilities.

process that we fit from data prior to estimation (see Appendix A of the Supplemental Material for details). Importantly, the benefit rule $cb_t(N_t)$ provides a natural exclusion restriction. It affects behavior of single women *directly* through the budget constraint, but it only affects behavior of married women, and all men, *indirectly* through the marriage market and the household bargaining process.

The net income of a single person is given by

$$Y_t^j = GY_t^j - \tau_t^S\big(w_t^j h_t^j, N_t\big), \quad j = f, m, \tag{14}$$

where $\tau_t^S(w_t^j h_t^j, N_t)$ is the time t tax function for single individuals calculated using the tax rules described in Appendix B. Thus, the budget constraint for a single person is simply

$$C_t^j = \big(1 - \kappa(N_t)\big)Y_t^j. \tag{15}$$

Note that both single men and women may have children ($N_t > 0$). These may be children from a previous marriage or, in the case of single women, children born outside of marriage.

In our model, utility functions exist at the *individual* level, and are not fundamentally altered by marriage. Within marriage, collective household decisions are made by constrained maximization of a weighted average of the individual partners' utility functions, as in (10). Consistent with this, we specify the utility functions of singles to be as similar as possible to those of married individuals. Consider the per-period utility function of a single female:

$$U_t^f(\Omega_{ft}) = \left(\frac{1}{\alpha}(C_t)^\alpha + L_j(l_t)\right)(1 - s_t) + \vartheta_{ft}s_t + \pi_t p_t + A_f^s Q(l_t, 0, Y_t, N_t), \tag{16}$$

where s_t is a 1/0 indicator for school attendance. Provided the single woman is not in school ($s_t = 0$), her utility function is fundamentally identical to that of a married woman, as one can see by comparing (4a)–(4b) and (16). The only differences are that, in (16), consumption is individual-specific (i.e., the equivalence scale $\psi = 1$), utility from marriage is (of course) dropped, the utility from children is allowed to differ from the married state ($A_f^s \neq A_j^M$), and, of course, the home-time of the husband is set to zero in the Q function.

As we noted earlier, we assume only single people can attend school ($s_t = 1$). To simplify our model, we do not model how consumption is financed by students.[20] Nor can we measure "leisure" time for students in a way comparable to workers.[21] Thus, consistent with prior work like Keane and Wolpin (1997), we simply define a "utility while in school" variable, which we denote by ϑ_{jt} for $j = m, f$. As we see in (16), students receive ϑ_{jt} as a current payoff, rather than the function of (C_t, l_t) that workers receive. This payoff is given by

$$\vartheta_{jt} = \vartheta_{0j} + TC\cdot I(E_t > HSG) + \vartheta_{1j}PE + \vartheta_{2j}\mu_j^W \quad \text{for } j = m, f. \tag{17}$$

Here ϑ_{jt} is a function of tuition cost TC, which is only relevant for higher education, the skill endowment μ_j^W, and parents' education, denoted PE. As Keane and Wolpin (2001) discussed, PE affects utility while in school in three key ways: (i) it affects consumption

[20]Consumption while in school may be financed by a combination of parental transfers, financial aid, part-time work, etc. (see Keane and Wolpin (2001)). It is beyond the scope of our paper to model all these possibilities.

[21]While we can see hours of market work, we cannot measure hours spent on school work.

while in school because more educated parents make larger financial transfers, (ii) it affects tastes for school, and (iii) it affects ability at school. These are very strong causal factors driving children's education. They help to generate changes in education levels across cohorts, as cohorts with better educated parents tend to themselves get more education.[22]

Education evolves as follows: At age 17, all people are in school, at education level "HSD." Two more years of school are required to become a high school graduate (HSG). For the next 1 to 3 years of school, the person is at the some college (SC) level, while 4 more years are needed to become a college graduate (CG). Any additional years lead to the PC level.

We can now write the choice-specific value functions for single females:

$$V_t^f(l_t, p_t, s_t \mid \Omega_{ft}) = \left(\frac{1}{\alpha}(C_t)^\alpha + L_f(l_t)\right)(1-s_t) + \vartheta_{ft}s_t + \pi_t p_t + A_f^s Q(l_t, 0, Y_t, N_t) + \delta E_{\mathrm{MAX}} V(\Omega_{f,t+1}), \tag{18a}$$

$$E_{\mathrm{MAX}} V(\Omega_{f,t+1}) = E_{\mathrm{MAX}}\big(M_{t+1} V_{t+1}^{fM}(\Omega_{m,t+1}, \Omega_{f,t+1}) + (1 - M_{t+1}) V_{t+1}^f(\Omega_{f,t+1})\big), \tag{18b}$$

where $E_{\mathrm{MAX}} V(\Omega_{f,t+1})$ takes into account that the person may get married at $t+1$.

Similarly, for single males, we have the choice-specific value function:

$$V_t^m(l_t, s_t \mid \Omega_{mt}) = \left(\frac{1}{\alpha}(C_t)^\alpha + L_m(l_t)\right)(1-s_t) + \vartheta_{mt}s_t + A_m^s Q(0, l_t, Y_t, N_t) + \delta E_{\mathrm{MAX}} V(\Omega_{m,t+1}). \tag{19}$$

Equations (18a)–(18b) and (19) are symmetric, except that the latter does not include a pregnancy option. The future component in (19) is defined analogously to that for women.

Now we consider the optimization problem of singles. In Section 3.5, we discuss the marriage market, but we must first consider decision making *conditional* on being single, that is, the state where no marriage offer is available or where it has already been declined.

Let $V_t^m(\Omega_{mt})$ and $V_t^f(\Omega_{ft})$ denote the maximized value functions of single males and females in period t. Let $\mathcal{S}_t^m$ and $\mathcal{S}_t^f$ denote the feasible set of choice options for a single male and female in period t, respectively. As we will see in Section 3.3, workers receive job offers probabilistically, so $\mathcal{S}_t^m$ and $\mathcal{S}_t^f$ may not include all possible levels of work hours and leisure. To proceed, for women and men we have, respectively,

$$V_t^f(\Omega_{ft}) = \max_{\{l_t, p_t, s_t\} \in \mathcal{S}_t^f} V_t^f(l_t, p_t, s_t \mid \Omega_{ft}), \tag{20}$$

$$V_t^m(\Omega_{mt}) = \max_{\{l_t, s_t\} \in \mathcal{S}_t^m} V_t^f(l_t, s_t \mid \Omega_{mt}). \tag{21}$$

These value functions appear in (11) and (27) that govern divorce and marriage decisions.

[22]Note that ϑ_{ft} captures utility of school net of costs. Without data on costs, these cannot be identified separately.

3.3. *The Labor Market—Wage Offers and Job Offers*

The wage offer functions have a standard Ben-Porath (1967), Mincer (1974) form:

$$\ln w_{et}^{j} = \omega_{1e}^{j} + \omega_{2e}^{j} X_t - \omega_{3e}^{j} X_t^2 + \varepsilon_{jt}^{W} \quad \text{for } j = f, m, \tag{22}$$

where X_t is work experience (years) and $e \in \{HSD, HSG, SC, CG, PC\}$ is education level. We let the wage function parameters $\{\omega_{ke}^{j}\}_{k=1,3}$ vary freely by cohort, gender, and education. Thus, at a given education level, both starting wages and returns to experience may differ between males and females, capturing two potential dimensions of discrimination.[23] Our specification allows returns to experience to differ by education, as recent studies (e.g., Imai and Keane (2004), Blundell, Costas-Dias, Meghir, and Shaw (2016)) find greater experience returns for more educated workers. We let parameters vary by cohort to allow for changes in the wage structure over time, particularly changes in returns to education/experience and relative wages between men and women.

The error term ε_{jt}^{W} in equation (22) has a permanent/transitory structure:

$$\varepsilon_{jt}^{W} = \mu_j^{W}(PE) + \tilde{\varepsilon}_{jt}^{W} \quad \text{where } \tilde{\varepsilon}_{jt}^{W} \sim \text{iidN}\big(0, \sigma_{\varepsilon}^{W}\big). \tag{23}$$

The time-invariant error component μ_j^{W} is the person's skill endowment (as in Keane and Wolpin (1997)). Recall from Section 3.2 that the skill endowment may also affect taste for school. We allow for three initial skill endowment levels (low, medium, high). The probability a person is each type is allowed to depend on parents' education (PE), as people with more educated parents tend to have higher skill endowments, and there is a strong intergenerational correlation in education (see Keane and Wolpin (2001), Eckstein and Lifshitz (2011)).

In each time period, people who were unemployed at the start of the period ($h_{t-1} = 0$) may receive full- and/or part-time job offers probabilistically. Thus, their possible choice sets for hours are $D_t = \{0\}, \{0, 0.5\}, \{0, 1\}$, or $\{0, 0.5, 1\}$. The probabilities of getting a full-time offer and a part-time offer are each determined by a logit model of the form

$$P_j(k \in D_t) = \frac{\exp\big(\phi_{j0k} + \phi_{j1k} e_t^r + \phi_{j2k} X_t + \phi_{j3k} H_t\big)}{1 + \exp\big(\phi_{j0k} + \phi_{j1k} e_t^r + \phi_{j2k} X_t + \phi_{j3k} H_t\big)} \quad \text{for } k = 1, 2, \tag{24}$$

where $k = 1, 2$ denote full- and part-time, respectively, and $j = f, m$. Here $e_t^r = 1, \dots, 5$ corresponds to the five education levels in ascending order, X_t is work experience, and H_t is health. An employed individual ($h_{t-1} > 0$) has the option to keep his/her previous job, unless an exogenous separation occurs. The separation probability obeys a similar logit function that also depends on e_t^r, X_t, and H_t. Thus, a person may be involuntarily unemployed due to exogenous separation or because he/she draws the empty choice set (i.e., $D = \{0\}$). And a person may be voluntarily unemployed because he/she quits their prior job or has an offer and rejects it.

By introducing job offer probabilities we can capture the idea that women who leave the labor force (perhaps after marriage or childbirth) may have a difficult time obtaining job offers later, as their work experience lags behind that of women who continue working. Hence, the impact of not working on experience (and wages) can accumulate over

[23]Different intercepts or slopes between males and females may also capture that males and females of a given age and education are not perfect substitutes in production, causing rental rates on male/female labor to differ.

time. This is important as a potential source of persistence in women's employment (or unemployment), as in Eckstein and Wolpin (1989). At older ages, we also have that poor health may limit one's ability to get job offers, which may reduce employment and even encourage retirement.

3.4. *Health Status*

Health obeys a three-state Markov chain, where $H_t^j \in \{1, 2, 3\}$ indicate good, fair, and poor, respectively. The transition probabilities differ by age and cohort. We assume health is an *exogenous* process, so it can be estimated outside the model. See Appendix A for details.

Health is important in our model. For example, we require people to retire after age 65, but declining health may induce them to retire earlier, as health affects both tastes for work (Eq. (4b)) and job offer probabilities (Eq. (24)). Health is also a dimension of match quality (Eq. (6)) on which people sort in the marriage market, and it shifts tastes for pregnancy (Eq. (7)). Furthermore, as we assume health is not affected by employment, marriage, or fertility decisions, it generates exogenous variation in these decisions (*given* our model).

3.5. *The Marriage Market*

The final component of the model is the marriage market. Single people may receive marriage offers, and they choose to become married if they draw a good enough match. To make this decision, they must compare the value of remaining single to the value of entering the married state. This section describes how the matching process works.

3.5.1. *Marriage Offers*

At the start of a period, a single individual may receive a marriage offer. Denote the probability of receiving an offer as $p_j^H(\Omega_{jt})$ for $j = f, m$. We assume the probability is given by a binomial logit model that depends on age and age-squared, whether a person is below 18, and whether a person is in school. The age effects differ by gender.

A marriage offer is characterized by a vector of attributes of a potential spouse, denoted by $\mathcal{M}_{jt}$. It is convenient to describe the construction of marriage offers in three steps:

First, we assume marriage offers always come from a potential spouse of the same age (t). This is necessitated by technical issues that arise in solving the dynamic programming problem. See Appendix E of the Supplemental Material for details. We do not think this assumption will have too great an effect on the results, because the large majority of married couples are in fact close in age.[24]

Second, we draw the education of the potential spouse. We assume potential spouses have three possible education levels: high school and below (HS, $ed = 0$), some college (SC, $ed = 1$), or college or above (C, $ed = 2$). The probability of receiving an offer from a potential spouse of the HS, SC, or C type depends on a person's own education.[25]

[24]For the cohorts of 1935–1975, the age gap between partners is below 5 years for 78% of all couples. It is below 7 years for 87% of couples, and below 10 years for 94% of couples.

[25]To simplify the MNL, we combine the HSD and HSG levels into "HS," and the CG and PC levels into "C." Then, if a person draws "HS," we assign education level HSD or HSG to the potential partner according to the actual fraction in the data (by cohort and age). We do the same to convert "C" draws to CG and PC offers.

Specifically, if the individual gets a marriage offer, we draw the potential partner's education using a multinomial logit (MNL) with the following latent indices:

$$\begin{aligned} \nu_{jt}^{C} &= \eta_{0j}^{C} + \eta_{1j}^{C} \cdot I\left[ed^{m} - ed^{f} = 2\right] + \eta_{2j}^{C} \cdot I\left[ed^{m} - ed^{f} = 1\right] + \epsilon_{jt}^{C}, \\ \nu_{jt}^{SC} &= \eta_{0j}^{SC} + \eta_{1j}^{SC} \cdot I\left[ed^{m} - ed^{f} = 1\right] + \epsilon_{jt}^{SC}. \end{aligned} \tag{25}$$

High school is the base case with $\nu_{jt}^{HS} = 0$. The parameters η govern the extent to which a person is more or less likely to receive offers from potential partners whose education differs from their own. Changes in η across cohorts reflect changes in supply of potential partners of different education levels, as well as changing tastes for partners of different types.

Third, we draw the remaining elements of $\mathcal{M}_{jt}$. The five remaining *observed* elements are drawn from the population distribution of all potential partners within a person's own age/education cell.[26] These distributions are not conditional on unobservables, so we can obtain them from the data. The four *unobserved* elements are drawn from their population distributions as specified in the model. These are the potential partner's tastes for leisure μ_{jt}, labor market ability μ_{j}^{W}, transitory wage shock $\tilde{\varepsilon}_{jt}^{W}$, and the taste for marriage ε_{t}^{M}. The stochastic terms μ_{jt}, μ_{j}^{W}, $\tilde{\varepsilon}_{jt}^{W}$, and ε_{t}^{M} are observed by both parties as part of the marriage offer. Both parties also understand which terms are permanent and which terms are only transitory.

Putting this all together, the marriage offer for a single female consists of the vector

$$\mathcal{M}_{ft} = \left(E^{m}, H^{m}, X^{m}, N^{m}, PE^{m}, h_{t-1}^{m}, \mu_{mt}, \mu_{m}^{W}, \tilde{\varepsilon}_{mt}^{W}, \varepsilon_{t}^{M}\right). \tag{26}$$

Marriage offers for males ($\mathcal{M}_{mt}$) have an analogous form.

3.5.2. *Marriage Decisions*

Given a marriage offer $\mathcal{M}_{jt}$, a single person can construct the vector $(\Omega_{ft}, \Omega_{mt})$ that characterizes the state of the couple if they marry. That is, $(\Omega_{jt}, \mathcal{M}_{jt}) \to (\Omega_{ft}, \Omega_{mt})$ for $j = f, m$. The potential partner also knows $(\Omega_{ft}, \Omega_{mt})$. Both parties calculate the value of marriage, denoted by $V_{t}^{jM}(\Omega_{mt}, \Omega_{ft})$ for $j = f, m$ in equation (12). A marriage is formed if and only if

$$V_{t}^{fM}(\Omega_{mt}, \Omega_{ft}) > V_{t}^{f}(\Omega_{ft}) \quad \text{and} \quad V_{t}^{mM}(\Omega_{mt}, \Omega_{ft}) > V_{t}^{m}(\Omega_{mt}). \tag{27}$$

If the pair decides to marry, they proceed to make collective decisions about work and fertility as described in Section 3.1. If the pair decides to remain single, they make decisions about work, school, and (for women) fertility as described in Section 3.2.

3.6. *Terminal Period and Retirement*

The terminal period in the model is fixed at age 65, after which everyone must retire. Of course, people can choose to stop working earlier if desired. By setting the terminal period at 65, we avoid the complications of modeling Social Security and the accumulation

[26] These elements of $\mathcal{M}_{jt}$ are the partner's health, work experience, number of children, PE, and lagged work.

of retirement savings.[27] To reduce computational burden, we assume the terminal value function $V^j_{T+1}(\Omega_{j,T+1})$ at $T = 65$ is a simple function of the state variables (dated at the end of period $T = 65$). Thus, the terminal value function accounts for retirement savings in a reduced form way. See Section 6.4 and Appendix G for details.

3.7. *Summary*

This completes the exposition of the model. Note that the choice set of a married couple is $\{l^m_t, l^f_t, p_t\}$, as well as whether to stay married, while that of a single person is work hours, school attendance, whether to marry, and, for women, pregnancy. We have stochastic terms in tastes for leisure, pregnancy, school, and marriage. As there is a stochastic term associated with every choice, the model will produce a non-degenerate likelihood.

Finally, it is useful to discuss the mechanisms that drive marriage in the model. First, there is the public good nature of couples' consumption. Each partner consumes 71% of total household consumption. Thus, marriage may increase consumption of both parties. However, if a person has much higher earning capacity than their potential spouse, his/her consumption may fall with marriage. Thus, a person with higher earning capacity will tend to have a higher reservation earning capacity for their spouse (ceteris paribus). This occurs for two reasons: (i) the higher a person's income, the higher the income of his/her spouse must be to prevent consumption from falling after marriage, and (ii) a person with higher earning capacity will have a higher probability that his/her offers are accepted, enabling them to be more selective. These mechanisms help to generate assortative mating in the model.

Second, people get utility from marriage itself (see Eq. (6)). But the magnitude of this utility differs across potential spouses. This gives the individual an incentive to search over marriage offers (i.e., an option value for waiting). In (6), we specified that people prefer spouses with similar education. This helps to generate assortative mating on education. Interestingly, there is a trade-off between θ^M_t and w^j_t (as noted earlier). This means a person is willing to accept a larger education difference if it is compensated by a higher wage.

4. SOLUTION OF THE MODEL

We back-solve the model from age 65 to 17.[28] We stress that we solve the dynamic programming (DP) problems of *individual* males and females. The individual solves his/her problem understanding the probabilities of marriage and divorce, and how decisions will be made when married (i.e., that couples solve the joint problem described in Section 3.1).

The state space Ω_{jt} of our DP problem is discrete. One state variable is marital status (M_t). The set of remaining state variables depends on marital status. We now list all the state variables in Ω_{jt}, along with the number of possible values (or grid points) for each:

For a single person, the state variables are: Gender ($j = m, f$); Age (t), $t = 17, \dots, 65$; education (E) with 5 levels; experience (X) with 5 grid points 0,1.5, 3.5, 7.5, and 15;[29]

[27] For our five cohorts, the "normal" age for claiming Social Security (SS) benefits was gradually increased from 65 to 67. But workers can also opt to receive "early" SS benefits at age 62, subject to a penalty in the form of a reduced benefit level. To avoid having to model this decision, we chose to fit our model only to data up through age 61. By setting the terminal value function at age 65, we implicitly assume away the early SS option.

[28] We assume everyone lives to at least age 65, and abstract from marriages that end via the death of a spouse.

[29] Note that we reduce the size of the state space by adopting a grid of only 5 possible experience levels, and interpolating the value functions between those points. We concentrate the grid points at low experience levels so the approximation can capture the fact that returns to experience are higher at low levels of experience.

children (N) with 4 levels (0, 1, 2, 3+);[30] health (H) with 3 levels; taste for leisure (μ_{jt}) with 3 levels;[31] the skill endowment (μ_j^w) with 3 levels; lagged work ($h_{t-1} \geq 0$) with 2 levels; lagged pregnancy (p_{t-1}) and parent education with 2 levels (college or not). For a married person, or a single with a marriage offer, the state variables are all of the above, plus the characteristics of the spouse or potential spouse ($\mathcal{M}_{jt}$).[32] Match quality is the same for both partners.

The number of children is a state variable regardless of whether a person is married. The law of motion for children conditional on a new marriage at age $t+1$ is $N_{t+1} = N_t + p_t - p_{t-18} + M_{t+1}(1 - M_t)N^s$, where N^s is number of children of the potential spouse. We do not distinguish own from spouse's children after a marriage is formed. Such a distinction is not possible in the data, and it would expand the state space dramatically. If a divorce occurs, the number of children remains a state variable for both partners. It continues to enter the $Q(\cdot)$ function in (18a)–(18b) and (19) and the budget constraint in (15).[33] The values of A_j^s for $j = m, f$ capture that men/women may have different responsibility/concern for children of a prior marriage.

Crucially, starting at age 17, a single person must make choices taking into account how they affect his/her marriage market opportunities. This means being able to predict the distribution of potential spouses *conditional* on own age and education in *future* periods. We assume people have perfect foresight about these distributions. This is imposed implicitly in estimation by: (i) using the same offer distribution that we fit within the estimation as the distribution that people use to forecast offers, and (ii) requiring that the model based on this assumed distribution provides a good fit to realized assortative mating patterns. This assumption circumvents the need to solve for the offer distribution as an endogenous object that emerges from the marriage market equilibrium, which would be infeasible in a dynamic context. Appendix E contains additional technical details about the marriage market.

5. ESTIMATION AND IDENTIFICATION

We estimate our model using repeated cross-section CPS data, as in Eckstein and Lifshitz (2011); specifically, using annual data from the March CPS for the period 1962 to 2015. The sample is restricted to white civilian adults over age 16. We divide the sample into five cohorts born within two years of the reference years 1935, 1945, 1955, 1965, and 1975.[34]

The estimation method is the Method of Simulated Moments (MSM), as proposed by McFadden (1989) and Pakes and Pollard (1989). See Appendix F for details. Our procedure involves simulating hypothetical life-cycle data for 1000 men and women for each

[30] We avoid including ages of children in the state space when we solve the model. This is discussed in detail in Appendix E Part 2. But we keep track of children's age when we simulate the model. We impose that 4 is the maximum number of children, as few people in the data have more, and it reduces the size of the state space.

[31] We adopt a three-state Markov approximation to the AR(1) process for tastes for leisure.

[32] In the case of a single person with a marriage offer, the children the potential spouse would bring to the marriage is also relevant. We denote this by N^s and assume it has two possible levels.

[33] We assume that both parents contribute part of their income to their children after divorce. But we assume no alimony payments. This is plausible, as during our sample period the enforcement of alimony was weak.

[34] We did not use the 1925 cohort as the data does not start until age 37, and we did not use the 1985 cohort as the data exists only up to age 31. Additional sample restrictions are discussed in Appendices A and H.

cohort. We observe cohorts for different age ranges; the last observation is age 61 for the 1935, 1945, and 1955 cohorts, 51 for the 1965 cohort, and 41 for the 1975 cohort.[35]

We estimate "benchmark" and "full" specifications of the model. In the benchmark, the only differences between cohorts are the initial conditions, which are parent education, the health process, and the tax/welfare rules. We call these "benchmark factors." In the full specification, we let three additional sets of parameters differ by cohort: (a) marriage offer probabilities and divorce costs, (b) wage/job offer distributions, and (c) fertility shocks.

To estimate the benchmark model, we use only data for the cohorts of 1945, 1955, and 1965, leaving the cohorts of 1935 and 1975 to test the time invariance of the preference parameters. There are 1505 moments for the 1945 and 1955 cohorts, and 1181 moments for the 1965 cohort, so the benchmark model is fit to 4191 moments in total. The complete set of moments is listed in Appendix H, and the parameter estimates are presented in Appendix I.

Note that we set $b_m = b_f = 0$ in equations (1) and (13), as it was not feasible to identify both the value of home time parameters in the budget constraint and the preference for leisure parameters in the utility function (i.e., the β and τ_0 parameters in equations (4a)–(4b) and (5)).

5.1. *Identification*

In complex dynamic models, it is well known that formal proofs of semiparametric identification are infeasible. So it is standard to rely on heuristic arguments. Our model can be viewed as a dynamic version of Heckman's (1974) labor supply model, but where offer wage functions are combined with a far more elaborate selection mechanism. Intuitively, just as in his static model, identification relies on exclusion restrictions of two types. First, to identify offer wage functions given selection, we need variables that exogenously shift the decision rule for work (e.g., by shifting preferences or values of leisure) but do not enter the offer wage function directly. Second, to identify utility parameters, we need variables that exogenously shift offer wages but do not alter preferences or values of leisure.[36]

Consider first the offer wage function. As wages are only observed for those who choose to work, and the wage function contains a latent variable for unobserved skill, we have selection and endogeneity bias in estimating returns to education and experience. But our model contains several variables that affect the decision to work (through preferences, the value of leisure, or other channels), but that do not enter the offer wage function, and that are exogenous from the perspective of the agents.[37] These are health, welfare benefit rules, $cb_t(N_t)$, and the distributions of potential partners and competitors in the marriage market:

[35]For the 1935–55 cohorts data from 62–65 is available. However, as we noted earlier, we chose to fit our model only to data up through age 61 to avoid having to model the possible early receipt of Social Security at age 62.

[36]These heuristic arguments mimic identification arguments used in static labor supply models like Heckman (1974). There, the identification argument consists of two parts: First, in order to identify offer wage function parameters, one requires variables that enter the decision rule for working but do not enter the offer wage function. Second, in order to identify labor supply elasticities, one also needs variables that enter the offer wage function but not the utility function (and hence not the decision rule for work).

[37]Children have often been used as an instrument in this sense. But completed fertility (i.e., the number of children) is presumably correlated with the skill endowment, so it is not a valid exclusion. On the other hand, the arrival of newborn children shifts tastes for leisure but does not affect wages.

(i) Health affects work decisions directly through tastes for leisure and job offer probabilities. But it does not affect offer wages directly, generating an exclusion restriction.[38]

(ii) Welfare benefit rules provide an obvious exclusion restriction, as they affect decisions about work, marriage, and fertility but do not enter the offer wage function.

(iii) The distribution of potential partners (marriage market opportunities conditional on education and experience) is an important source of returns to education/experience—and hence an important influence on work decisions—but it does not affect offer wages directly.

Second, we also have variables that shift offer wages but do not alter preferences or values of leisure, thus enabling us to identify utility parameters. For example, tax rules vary across cohorts and calendar years within cohorts. These tax rule changes alter after-tax offer wages, allowing us to identify labor supply elasticities with respect to after-tax wages.

Given offer wage and labor supply functions, a heuristic argument for identification of the decision rule for school attendance is straightforward. Following Willis and Rosen (1979), we can, in principle, construct expected values of lifetime earnings corresponding to different education levels, and then, in principle, back out the tastes for school (and other preference parameters) required to match the observed schooling distribution.[39]

Similarly, to identify the decision rule for marriage, we can use the offer wage and labor supply functions (both for marrieds and singles) to calculate expected values of future earnings and consumption for married versus single individuals of both genders. We could then, in principle, back out preferences for marriage that rationalize observed marriage decisions. Of course, *in practice*, such calculations would be incredibly complex, which is precisely why the literature resorts to heuristic arguments.[40]

In the literature using method-of-moments estimation, an alternative type of heuristic argument for identification is often used. This mode of argument focuses on which data moments pin down certain model parameters. Of course, in complex nonlinear models, all moments potentially influence all parameters. But it can still be useful to give an intuitive discussion of which moments are most informative about certain parts of the model.

For example, moments involving starting wages and the distribution of completed schooling help pin down wage function intercepts (conditional on education) and parameters of tastes for schooling.[41] As individuals only leave school to enter work if a wage offer is high enough, the model's predicted starting wages incorporate a selection correction. Similarly, moments involving wages conditional on experience convey information about returns to experience (corrected for selection based on the model's decision rule for work).

Moments involving employment/wages pin down tastes for leisure and consumption, as well as parameters of the job offer function. Consider a model with a 100% job offer

[38]While a decline in health does not reduce wages immediately, it will reduce wages over time through its effect on accumulated work experience. This mechanism was found to be important by Capatina, Keane, and Maruyama (2018).

[39]Both future and present marriage market opportunities matter for schooling decisions. The expected quality of future marriage offers depends on one's education, and this is an important incentive to stay in school. But students also get marriage offers, and the distribution of these offers depends on their current education level.

[40]Of course, it is straightforward to test whether any particular parametric model—no matter how complex—is identified simply by checking the invertibility of the Hessian.

[41]Parents' education affects tastes for school but it also affects the probabilities of the initial skill level. So it is endogenous in the sense that it enters the wage function directly.

probability versus one with a lower probability. In the former, all unemployment is due to rejected offers, so a higher variance of tastes for leisure is needed to generate a given level of unemployment. And the composition of the unemployed generally differs between the two settings (as rejecting offers is a choice, but job destruction is not).[42] Changes in employment with arrival of a newborn also help to pin down some parameters of tastes for leisure.

Marriage/divorce rates and marital sorting help identify tastes for marriage. For instance, if tastes for marriage have small variance, one would choose partners solely based on income and education. More variation in tastes leads to less perfect sorting. Similarly, the nature of sorting is affected by who one is likely to meet (e.g., a lower proportion of college types in the population lowers the chance of marrying that type). As a final example, variation of fertility with income/employment helps pin down parameters of the child quality production function.

Dynamic models can be estimated from cross-sectional data, provided one observes the endogenous state variables.[43] But we face the added complication that experience is not observed in the CPS (we only see age). We rely on the structure of our model to deal with this problem. Most importantly, we assume age does not enter the wage function directly, so all life-cycle wage growth is due to experience. We do not attempt to disentangle experience versus age effects on wages, as this could not be done credibly without data on both variables.

Furthermore, our model generates a work decision rule that implies a mapping from age to experience. Many observable quantities are affected by that mapping. As one example, the magnitude of returns to work experience affects labor supply behavior, and in particular the relationship between the wage-age profile and the wage-hours profile (see Imai and Keane (2004)). Another example is the observed married/unmarried wage gap. The higher the participation rate, the better is age as a proxy for experience. Specifically, as we will see below, changes in the marriage wage gap over time can be explained (in part) by changes in the age-experience mapping induced by rising employment rates of married women.

Finally, as indicated earlier, we assume Pareto weights are constant across time and cohorts. Allowing them to vary (e.g., to preserve a marriage when total surplus is positive but one party prefers divorce) would create two problems: First, it enlarges the state space, which is infeasible. Second, strong identification of the weights requires data on private consumption/leisure/housework time by members of the household, but these data are not available. As we show below, our model provides a good fit to marriage and divorce rates, which are directly affected by the Pareto weights. Thus, we feel that allowing them to vary would not add much to the model or significantly improve its fit (see Gihleb and Lifshitz (2016)).

6. ESTIMATION RESULTS AND INTERPRETATION

We first estimate our model on pooled data from the 1945, 1955, and 1965 cohorts, assuming common parameters across cohorts, and with no fertility control shock in (7). We refer to this as the "benchmark" specification. Here, all that differs across cohorts are

[42]For instance, given their high wage rates, it is difficult to generate the observed unemployment among male college graduates without some job destruction.

[43]Unobserved exogenous state variables can always be integrated out of choice probabilities (see Rust (1987)).

the initial conditions: mother's education, the tax/welfare rules, and the health process. We refer to these as "benchmark factors." Of these factors, mother's education clearly has the primary impact on cohort differences, as it increased greatly from the '35 to the '75 cohort.[44]

A subset of parameter estimates of the benchmark model are presented in Appendix Table I-I and Appendix G of the Supplemental Material. These are the parameters of preferences (for leisure, consumption, school, marriage, pregnancy, and children), the child quality production function, the shock variances, the initial ability distribution, and the terminal value function. Almost all are significant and with the expected signs. We do not yet report the labor and marriage market parameters in equations (22) and (24)–(25), as these are discussed below for the "full" model.

Not surprisingly, the benchmark model cannot fit the data for all five cohorts (see Appendix J). Thus, our empirical strategy is to allow some exogenous factors to differ across cohorts in order to obtain a "full" specification that does fit the data well. But, rather than conduct an arbitrary specification search, we discipline the analysis in two important ways:

First, we hold preference parameters fixed for all cohorts, including 1935 and 1975, at their estimated values in the benchmark specification. Second, based on the literature and our own judgment, we chose three additional factors that we deemed a priori to be the most important changes across the cohorts. We hypothesized that, given our model, these three additional factors are both necessary and sufficient to explain the main cohort differences. Specifically, in the "full" specification, we allow for the following four exogenous factors to differ by cohort, and we estimate their parameters for each cohort separately:

(i) Benchmark factors: mother's education, health, and tax/welfare rules.

(ii) Marriage market parameters: the marriage market matching function (25) and the divorce cost parameters, α_4^j and α_5^j.

(iii) Labor market parameters: the offer wage equation (22) and offer probabilities (24).

(iv) Availability of oral contraception: imperfect fertility control is captured by including in (7) the positive shock to taste for pregnancy $\exp(\varepsilon_t^{up})$ where $\varepsilon_t^{up} \sim \text{iidN}(pr, 1)$, and allowing pr to differ freely across cohorts.

Appendix Tables I-II and I-III report the cohort-specific parameter estimates for (ii)–(iv). These are almost all significant with expected signs and magnitudes. The next three subsections discuss how factors (i)–(iv) influence the fit of the model to (A) wages/employment; (B) the married versus unmarried wage gap, and (C) education, marriage, divorce, assortative mating, and fertility. We also discuss how each factor contributes to changes in behavior across cohorts.

6.1. *Wages and Employment by Cohort*

In Appendix J, Table J-I, we present the fit of both the "benchmark" and "full" specifications to mean wages and employment rates by cohort, broken down by gender, marital status, and age. Obviously, the benchmark model fits poorly for all cohorts except 1955. In particular, it fits the data on mean wages poorly for all other cohorts for almost all gender/marriage/age cells. The benchmark model does fit employment for men rather well in all cohorts. But it provides a very poor fit to employment for women. In particular, it

[44]The actual rates of college graduate mothers are: 6% for 1935 cohort, 6% for the 1945 cohort, 11% for the 1955 cohort, 20% for the 1965 cohort, and 27% for the 1975 cohort. For the tax schedule, see Appendix B; for the health transition process and the tax and welfare rules, see Appendix A.

greatly over-predicts both wages and employment for married women in the 1935–1945 cohorts, and it greatly under-predicts these quantities for the 1965–1975 cohorts.

In contrast, the full specification provides a very good fit to wages and employment in all cohorts for almost all gender/marriage/age cells. This supports our hypothesis that factors (i)–(iv) are sufficient to explain all the main cohort differences in wages and employment. In particular, we can explain (a) the rapid employment growth for married women at the same time that employment of single women and all men is fairly stable, (b) the very rapid wage growth for married women relative to other groups, and (c) the stabilization of employment of married women in the 1965–1975 cohorts. Furthermore, the model fits untargeted moments for annual transition rates of employment and marital status extremely well (see Appendix J, Table J-III).

We have shown that, with fixed preferences across cohorts from 1935 to 1975, factors (i)–(iv) are *sufficient* to fit the key facts regarding changes in wages and employment for married versus unmarried women (and all men) across the 1935–1975 cohorts. But are they all *necessary*? Next, we assess the contribution of each factor to the observed changes across cohorts.

In Table I, we use the full specification to decompose the contribution of each of the exogenous factors (i)–(iv) to changes across cohorts from 1935 to 1975. The predicted values of the full specification are very close to the data, so this decomposition also summarizes the contribution of each factor to the *actual* changes over time. We calculated the contribution of each factor by adding them in stages, in the order (i)–(iv), and asking how much of the total change from 1935 to 1975 is explained at each stage. Results are reported in Table I.[45]

TABLE I

DECOMPOSING SOURCES OF COHORT DIFFERENCES—WAGES AND EMPLOYMENT

				Contribution of Each Factor			
	1935 Fitted	1975 Fitted	Total % Change	Benchmark	Marriage Market	Labor Market	Contra-ception
Wages (Thousands of $)							
Married Women—Ages 25–34	20.5	39.0	90%	11%	7%	65%	8%
Married Women—Ages 35–44	25.1	51.2	104%	12%	5%	81%	5%
Unmarried Women—Ages 25–34	23.3	37.7	62%	4%	1%	55%	1%
Unmarried Women—Ages 35–44	28.4	43.5	53%	3%	1%	49%	0%
Married Men—Ages 25–34	36.2	51.3	42%	1%	1%	40%	0%
Married Men—Ages 35–44	52.2	69.8	34%	1%	1%	32%	0%
Unmarried Men—Ages 25–34	30.0	42.9	43%	3%	1%	39%	0%
Unmarried Men—Ages 35–44	42.9	56.3	31%	2%	1%	28%	0%
Employment							
Married Women—Ages 25–34	0.27	0.63	130%	13%	13%	67%	36%
Married Women—Ages 35–44	0.44	0.66	50%	4%	5%	35%	6%
Unmarried Women—Ages 25–34	0.68	0.75	11%	1%	0%	8%	1%
Unmarried Women—Ages 35–44	0.70	0.72	2%	0%	0%	2%	0%
Married Men—Ages 25–34	0.91	0.89	−2%	0%	−1%	−1%	0%
Married Men—Ages 35–44	0.92	0.90	−2%	0%	−1%	−2%	0%
Unmarried Men—Ages 25–34	0.78	0.79	2%	0%	0%	2%	0%
Unmarried Men—Ages 35–44	0.79	0.75	−5%	0%	0%	−5%	0%

[45]The step-by-step process has the drawback that the contribution of each factor may be sensitive to the order in which they are added. An alternative method is to change the factors one at a time holding other

According to Table I, top row, the benchmark model can explain only 11 points of the 90% increase in wages of married women aged 25–34 that occurred from the 1935 to 1975 cohorts. This 11 point increase was largely driven by the increase in mother's education over this period, which increased their daughters' skill endowment and tastes for school. When we allow factor (ii) to differ by cohort (i.e., marriage market opportunities and divorce costs), it explains another 7 points. But, when we introduce factor (iii), changes in labor market opportunities, it explains a substantial 65 points of the wage increase. Finally, we find that improved fertility control (oral contraception) contributed 8 points to the increase in wages.[46]

Thus, factor (iii), which captures changing labor market opportunities (i.e., offer wage functions and job offer probabilities), explains 3/4 of the 90% increase in wages for married women aged 25–34, with the other three factors explaining the remaining 1/4. Similarly, for married women aged 35–44, factor (iii) explains 4/5 of the 104% increase in wages. For other groups, changes in labor market opportunities explain almost all of the increase in wages.

In contrast to wages, we only see substantial increases in employment for married women aged 25–34 (130%) and married women aged 35–44 (50%). As we emphasized in Section 2, employment was rather stable for the other groups. Focusing on the younger group of married women, we find that changing labor market opportunities account for 67 points of the increase, that is, roughly half. Contraception accounts for roughly 1/4, and the baseline factors (primarily increased mothers' education) and changing marriage market conditions account for 1/8 each.[47] Thus, while labor market conditions are the main factor, we cannot explain the entire increase in employment of married women without also considering changes in mothers' education, the marriage market, and especially contraception.

The role of contraception is worth emphasizing. The estimates of equation (7) imply that fertility control improved greatly from the 1935 to 1945 cohorts, and that it was nearly perfect from the 1955 cohort onward.[48] The model implies that of the 36 point increase in employment of young married women (from 27% to 63%) that occurred from the 1935 to the 1975 cohorts, 9 points were due to improved contraception. We found it impossible to fit employment of young women in the 1935–1945 cohorts without accounting for this factor (see Appendix J and Eckstein, Keane, and Lifshitz (2016) for further details).

As changes in labor market opportunities play such an important role in explaining wage changes for all groups, and changes in employment for married women, it is interesting to examine how the wage structure changed over cohorts. We allow a very flexible

factors fixed. The one-at-a-time method has the problem that, if there are important interactions between factors, the individual contributions may not add up to anything close to 100%. Remarkably, we found the estimated contribution of each factor was nearly identical regardless of the method we used. The one-at-a-time results are in Appendix K.

[46]Note that, at the final stage, we re-estimate all the model parameters (except preferences, which are always held fixed at benchmark values) to get the final fit of the full specification.

[47]Recall that "marriage market conditions" include both the marriage offer distribution and divorce costs. Chiappori, Fortin, and Lacroix (2002) found a positive effect of easier divorce laws on women's employment. Fernández and Wong (2011) find large effects of divorce probabilities on married women's employment.

[48]Recall $\varepsilon_t^{up} \sim \text{iidN}(pr, 1)$. For the 1935 and 1945 cohorts, we estimate $pr = -0.18$ and $pr = -0.79$, respectively, while for later cohorts, we estimate $pr \ll 0$ (see Appendix I). These results conform with historical patterns. The oral contraceptive pill was invented in the late 1950s, and approved by the FDA in 1960. It became available to married women in all states after *Griswold v. Connecticut*, 1965, and to unmarried women in all states after *Eisenstadt v. Baird*, 1972. Thus, oral contraceptives were not available to the 1935 cohort until late in their reproductive years, and availability to the 1945 cohort was mixed. But the later cohorts had unrestrictive access.

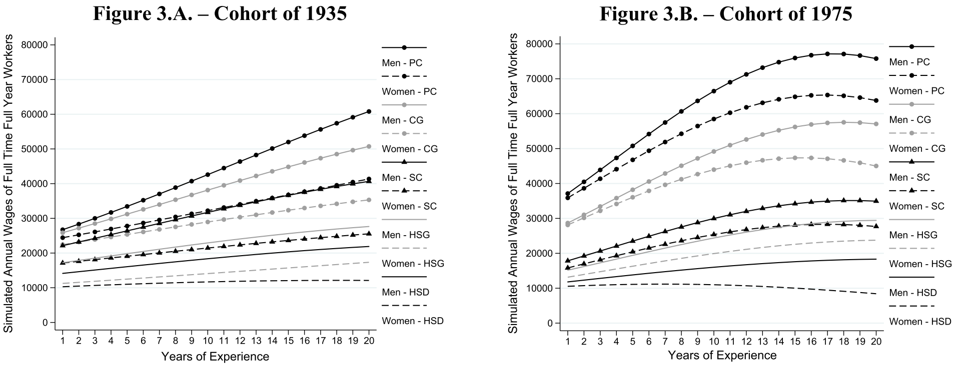

FIGURE 3.—Simulated annual wages by education and years of experience.

wage offer function (22) in which parameters are allowed to vary freely by cohort, gender, and education. For given education, both starting wages and returns to experience differ between males and females, in ways that differ by cohort. Given this flexible structure, individual coefficients are not very informative, and it is more useful to look at graphs of the offer wage/experience profiles. We present graphs for the 1935 and 1975 cohorts in Figure 3.

Figure 3A presents estimated offer wage profiles for the 1935 cohort. Each color/symbol combination represents an education group, and the dashed lines are for women while solid lines are for men. The inferior labor market opportunities for women are evident. At each of the five levels of education, the starting offer wage for women is well below that for men. Even more striking is that the slopes of the offer functions in experience are much less for women than men. For example, by 11 years of experience, a woman with a post-graduate degree (dashed line with circles) receives offers no better than a man with only some college (solid line with triangles).

Figure 3B shows that offer wage functions for men and women became much more similar in the 1975 cohort.[49] The starting wage gap at each education level was sharply reduced. It is still true that the men's offer wage function at each level of education lies slightly above that for women, and the experience slopes for men are still slightly greater than for women. But the improvement from 1935 is striking. Clearly, returns to education and experience increased greatly for women over this period.

One obvious potential explanation for these patterns is that discrimination against women in the labor market has been reduced (see Jones, Manuelli, and McGrattan (2015)). But two other possibilities are that: (a) returns to female skills have increased, perhaps due to the growth of the service sector (see Lee and Wolpin (2006), Johnson and Keane (2013)), and (b) human capital investments in girls prior to age 17 have improved. Obviously, our model is not designed to disentangle these scenarios. A key point, however, is that the complex changes we estimate cannot be captured by a reduction in the traditional measure of discrimination, that is, the so-called "male/female wage gap" (or wage function intercept). Rather, we find that the male/female wage gap changed differ-

[49]If offer wages are similar for men and women, the dashed and solid lines of a given color/symbol combination should lie close together in Figure 3. This is far from true in 1935, where we often see a dashed line (women) of color/symbol X lying below a solid line (men) of color/symbol Y, when color/symbol Y actually represents a lower level of education than color/symbol X. But in the right panel, for the 1975 cohort, we see a nearly perfect grouping of the offer wage functions by color/symbol combination.

entially at different levels of education/experience. This pattern is hard to explain by a reduction in pure discrimination.

Finally, we also consider how job offer rates have varied over time. According to our estimates, offer rates for already employed individuals are always 96 to 97%. So it is more interesting to focus on the non-employed. In the 1935 cohort, a non-working woman had a 1/3 chance of getting a full- or a part-time offer, regardless of her education level. For high-school men in the 1935 cohort, the probability of a full-time offer was much higher (58%), and this increased to 68% for college men. However, the offer probabilities faced by women and men converge substantially. In the 1975 cohort, high school women are much more likely to get full-time offers (60%), and the chance is even better for college women (68%). These figures are still below those for men (71% and 79%, respectively), but the convergence is striking.[50]

6.2. *The Marriage Wage Gap*

As the full specification fits mean wage differences between married and unmarried women and men nearly perfectly for all cohorts, in this section we ask whether the model is also able to fit the dramatic changes in the marriage wage gap (see Appendix D). This was an untargeted moment in the estimation. Thus, our ability to fit the marriage wage gap is a test of the external validity of the model. To proceed, we estimated Mincer wage equations for each cohort of women, using (i) the actual CPS data, (ii) simulated data from the benchmark model, and (iii) simulated data from the full model. Table II reports the results.

As we noted in Section 2, in the CPS the marriage wage gap shifts from −8.9% in the 1935 cohort to 5.2% in the 1975 cohort. Thus, the "negative" selection into marriage of the 1935 cohort is reversed to "positive" selection in the 1975 cohort. The benchmark model generates a 4.9 percentage point increase in the marriage wage gap from the 1935 cohort to the 1975 cohort—roughly 35% of the observed change. But the full model predicts a 12.8 point increase, so it accounts for 90% of the total observed change.[51] We regard this as an impressive validation of the model, as fitting the nature of selection into marriage is a more subtle challenge than, say, fitting levels of wages and employment.

In the row of Table II labeled "control for experience," we use simulated data from the full model and include true work experience rather than age in the Mincer equation. This

TABLE II

MARRIAGE WAGE GAP BY GENDER AND COHORT

	Women Marriage Wage Gap					Men Marriage Wage Gap				
	1935	1945	1955	1965	1975	1935	1945	1955	1965	1975
Data	−8.9%	−6.8%	−1.7%	2.0%	5.2%	19.7%	18.7%	19.5%	19.7%	18.3%
Benchmark Model	−3.6%	−3.7%	−1.1%	0.8%	1.3%	11.9%	12.3%	12.0%	12.9%	12.3%
Full Model	−8.4%	−6.4%	−1.0%	2.3%	4.4%	12.9%	13.8%	13.6%	13.8%	13.7%
Control for Experience	−3.3%	−2.8%	2.0%	3.2%	5.0%	4.3%	4.4%	5.5%	6.5%	6.4%
Control for Ability	0.8%	0.8%	1.1%	0.7%	1.0%	1.2%	0.8%	0.9%	1.4%	0.9%

[50]The figures in the text are calculated using equation (24) and the offer probability parameters in Appendix I.

[51]Recall that our analysis is disciplined by the assumption that preferences are invariant across cohorts. The remaining 10% change of the marriage wage gap that we do not explain could be due to changing preferences.

is only possible in simulated data, as we do not observe actual experience in the CPS. With this change, the model predicts an 8.3 percentage point increase in the marriage wage gap. Thus, the model implies that $12.8 - 8.3 = 4.5$ points (or one-third) of the 12.8 point predicted increase in the marriage wage gap is due to changes in the mapping from age to experience (i.e., married women now work more and acquire more human capital).

Mulligan and Rubinstein (2008) find that amongst married women, selection on latent ability into employment became more favorable over this period. In the last column of Table II, we also control for unobserved ability, which we observe in the simulated data. In principle, if the mapping from education, experience, experience-squared, and latent ability to wages was exactly log-linear, this equation should control for all differences between married and single women, and the marriage wage gap should vanish. In fact, the equation implies a small (but statistically insignificant) positive marriage wage gap for all cohorts. But it increases by only 0.2 percentage points from 1935 to 1975. Thus, we estimate that $8.3 - 0.2 = 8.1$ points or about 2/3 of the 12.8 point predicted increase of the marriage wage gap is due to selection of higher ability women into marriage.

The right side of Table II shows that, in contrast to women, there is no clear trend in the marriage wage gap for men, which hovers around 19%. Thus, the challenge is to explain its absolute level. The benchmark model generates a marriage wage gap that hovers around 12%, while in the full model it hovers around 13.5%. Thus, our full model generates roughly 70% of the marriage wage gap for males. The remaining 30% is due to factors not included in model. For example, the responsibilities created by marriage and children may cause married men to "work harder" (which we cannot capture given invariant individual preferences).

If we run the Mincer regression on simulated data from our full model, substituting true experience for age, the marriage wage gap for men drops to about 4 to 6%. Thus, the model implies roughly half the (fitted) marriage wage gap is accounted for by married men accumulating more work experience per unit of age. If we also control for latent ability, the marriage gap becomes small and insignificant. Thus, the other half of the marriage wage gap is explained by selection into marriage of men with high unobserved ability.

6.3. *Marriage, Divorce, Assortative Mating, Fertility, and Education*

So far, we have focused on the central issue of the growth in married women's wages and employment. Here, we examine how the model fits other demographic outcomes. Appendix J, Table J-II presents the fit of both the "benchmark" and "full" specifications to (a) marriage and divorce rates, (b) assortative mating, (c) fertility, and (d) education. Clearly, the benchmark model fits poorly for all cohorts except 1955 (and perhaps 1965). On the other hand, the full model matches *all key* moments for all five cohorts quite accurately. This establishes that adding only our three exogenous factors is *sufficient* for the full specification to provide a good fit to all endogenous variables in nearly all gender/marital-status/age cells.

We now turn to the question of whether all three factors are necessary to fit the data, and, more generally, how each factor contributes to changes between the 1935 and 1975 cohorts. The results are presented in Table III, where, similarly to Table I, we assess the separate contribution of each factor by adding them one-by-one to the baseline model.

As we see in Table III, the percentage of women with at least a college degree increased from 5% in the 1935 cohort to 36% in the 1975 cohort (see Goldin, Katz, and Kuziemko (2006) for a discussion). The benchmark model explains less than 1/3 of this increase. It does so primarily via increased mother's education, which raises daughter's skill endowments and tastes for school (as in Keane and Wolpin (2010)). When we add factors

TABLE III

DECOMPOSING SOURCES OF COHORT DIFFERENCES—MARRIAGE, CHILDREN, EDUCATION

	1935 Fitted	1975 Fitted	Total % Change	Contribution of Each Factor: Benchmark	Marriage Market	Labor Market	Contra-ception
Family moments							
Marriage Rate—Ages 25–34	0.86	0.60	−30%	−20%	−7%	−3%	0%
Marriage Rate—Ages 35–44	0.84	0.70	−16%	−7%	−7%	−2%	−1%
Divorce Rate—Ages 25–34	0.03	0.09	206%	31%	144%	13%	17%
Divorce Rate—Ages 35–44	0.08	0.12	62%	3%	54%	5%	0%
Married Women # of Children—Ages 25–34	2.54	1.51	−41%	−8%	−12%	0%	−20%
Married Women # of Children—Ages 35–44	2.24	1.94	−14%	−2%	−4%	0%	−6%
Unmarried Women # of Children—Ages 25–34	0.92	0.32	−66%	−6%	−6%	−1%	−53%
Unmarried Women # of Children—Ages 35–44	0.75	0.51	−32%	−3%	−4%	−1%	−24%
Education Distribution at 30							
Women's CG + PC Rate	0.05	0.36	620%	180%	220%	200%	20%
Men's CG + PC Rate	0.20	0.29	45%	5%	10%	30%	0%
Assortative Mating							
HSD With HSD	0.55	0.56	2%	0%	2%	2%	−2%
HSG With HSG	0.64	0.49	−23%	−9%	−8%	−5%	−2%
SC With SC	0.24	0.53	121%	−4%	25%	100%	0%
CG With CG	0.33	0.49	48%	6%	15%	27%	0%
PC With PC	0.12	0.43	258%	33%	33%	183%	8%
HSG Women With CG Men	0.34	0.08	−76%	−9%	−21%	−47%	0%
CG Women With HSG Men	0.02	0.12	500%	100%	150%	250%	0%

(ii)–(iv), the most important is changes in marriage market conditions, which account for over 1/3 of the increase. The second most important factor was changing labor market conditions (i.e., higher returns to education), which we already discussed in Section 6.1. Finally, we find that improved contraception played a minor role in increasing women's education.[52]

Why were changing marriage market conditions so important for the education level of women? One reason is that, according to the full model, a college educated woman in the 1935 cohort had little chance of getting marriage offers from college educated men, but by the 1975 cohort, the chance was much greater.[53] Keane and Wolpin (2010) found that roughly half the return to education for women comes through improved marriage market prospects,[54] and our model implies that the benefits of a college degree arising *through the marriage market* improved greatly for women between the 1935 and 1975 cohorts.

It is notable that assortative mating increased substantially. The chance a college graduate married another college graduate increased from 33% in the 1935 cohort to 49% in the 1975 cohort. And the chance a person with a graduate degree married another person with a graduate degree increased from 12% to 43% (see also Low (2014)). The model

[52]The fraction of men with a college or post-graduate degree rose much less than that of women (i.e., from 20% to 29%). By far the largest contributing factor to this increase was increasing returns to education.

[53]This is both because the number of college men was lower, and due to changing tastes for partners.

[54]Ge (2011) predicted female college graduation would drop 8% without the marriage market return to education. Chiappori, Iyigun, and Weiss (2009), Chiappori, Costa-Dias, and Meghir (2018), Fernández and Wong (2011) and Low (2014) also studied the interaction of the marriage market and education.

implies that changing labor market conditions explain 56% and 71% of these increases, respectively.

Why did changing labor market conditions account for such an increase in assortative mating? Consider a highly educated man in the 1935 cohort. His lifetime consumption would be little different if he married a high school versus a college educated woman, as (a) the education wage premium for women was small, and (b) women had low employment rates. By the 1975 cohort, highly educated women had become much more economically attractive to highly educated men. Of course, this marriage market effect would have caused forward-looking women to invest more in education in the first place.[55]

Recall we include changes in divorce costs as part of marriage market conditions. We estimate the fixed cost of divorce for women went down by 45% from the 1935 to 1975 cohorts, with most of that drop completed by the 1955 cohort (see Appendix I, Table I-III).[56,57] It is interesting to decompose the changes in women's education due to changes in overall marriage market conditions into (a) the part due to changes in the marriage offer distribution versus (b) the part due to the fall in divorce costs. We find that only 1/6 of the effect was due to the changes in the marriage offer distribution, while 5/6 was due to changes in divorce laws.

As Bronson (2015) noted, the reduction of divorce costs in the United States occurred mostly via changes in state laws in the early 1970s that made divorce easier (no fault unilateral divorce). She analyzed the impact of these changes on women's education using a model fairly similar to ours, albeit with many different modeling choices.[58] She estimated that more lenient divorce laws increased the college graduation rate of women in 1970 from 18% to 21%. Our closest comparison is our 16% to 23% predicted increase between the '45 and '55 cohorts (Appendix J, Table J-II).

The marriage rate fell and age at marriage was delayed from the 1935 to 1975 cohorts. According to our model, factor (i)— primarily increased mother's education—was the main factor driving down marriage at early ages. It accounts for 2/3 of the drop in the marriage rate at ages 24–35. As we have seen, mother's education has a strong positive effect on daughter's education, so it is not surprising it causes delayed marriage. In contrast, mothers' education and changing marriage market conditions contribute about equally to the smaller decline in the marriage rate at older ages. The model implies the substantial rise in the divorce rate was almost entirely due to changed marriage market conditions (including falling divorce costs).

[55]Our argument is consistent with Low (2014). She noted that over time, highly educated women have become much more likely to form matches with highly educated men. Hence, the cost of delaying marriage to pursue a higher degree has fallen. We adopt the parallel language that the return to educational investment for women has increased because their marriage market prospects at the end of the investment period have improved.

[56]Note that major changes in the divorce laws occurred in the 1970s and 1980s (see Voena (2015) and Matouschek and Rasul (2008)). This is consistent with our estimates of the drop in the cost of divorce for the 1945 and 1955 cohorts, as most of them would have gotten married in the 1970s and 1980s. For men, we find a small increase in divorce costs (see Appendix I, Table I-III).

[57]The cost per child (α_5^j) fell only slightly across cohorts for both women and men. In the 1975 cohort, the constant cost of divorce is about the same for women and men, but the cost per child is 37% higher for women.

[58]On the one hand, she had savings and two occupations, and, on the other end, she assumed exogenous fertility, kept track of tenure rather than work experience, did not estimate wage functions jointly with the model, and did not fit data on multiple cohorts. Her model is estimated on the 1955–1960 birth cohort, which would have made school choices under the lenient divorce law regime in 1970. She then simulated a shift to a strict regime.

Table III shows sharp drops in fertility for married women. We find half of this drop was due to the availability of oral contraception. The rest was split between factor (i)— primarily higher mother's education—and changing marriage market conditions. Strikingly, changes in labor market opportunities play essentially no role in explaining the drop in fertility at younger or older ages. Fertility of unmarried women fell even more sharply. The model implies this was almost entirely due to oral contraception. Viewed another way, the model implies most births to unwed mothers are unplanned (i.e., induced by the ε_t^{up} shocks). However, while our results imply contraception was a key factor driving down fertility, and that it also led to about 1/4 of the substantial increase in employment for younger women (ages 25–34), we find it had little effect on women's education, marriage/divorce rates, or marriage market matching, and a limited role regarding employment (see Eckstein, Keane, and Lifshitz (2016) for further details).[59]

Summarizing the results of Sections 6.1 and 6.3, we see that all four exogenous cohort-varying factors in the full model are necessary to explain the data. If we exclude any one of these factors, the model gives a poor fit to at least one key aspect of the data.

6.4. *Robustness Checks: Home Production and Savings*

Greenwood, Guner, Kocharkov, and Santos (2016) argued that costs of household production may provide an alternative explanation (aside from contraception) for high fertility and low employment of young married women in the pre-WWII period. In Appendix L, we report results where we let the cost of home production decline across cohorts. But we conclude this does not improve the fit of our model, or substantively alter its predictions. This does not necessarily contradict Greenwood et al. (2016), as most improvements in home production that they emphasized had already occurred by 1950. Technologies like washing machines, refrigerators, disposable diapers, etc. were already widely available for the cohorts we study. Similarly, Albanesi and Olivetti (2016) find better maternal health care led to substantial increases in employment of young women in the 1930-50 period (prior to our data).

We also examined whether including savings might lead to significant changes in the behavior of our model. Including assets as an additional state variable, and consumption as an additional choice, would render solution of the model infeasible. But we are able to incorporate a very simple form of buffer stock saving (see Appendix L). When we added this feature, we found no evidence of improvement in fit, or any significant changes in behavior.

Notably, our model accounts for retirement savings in a reduced form way, as the terminal value function implies agents place large values on terminal work experience and education. As agents cannot work past age 65, these quantities only matter as proxies for retirement assets. If V_{66} did not depend on work experience, labor supply would drop off (too) precipitously before the $T = 65$ terminal period (as older workers do not need to save for retirement).

7. POLICY ANALYSIS: TAX REFORM AND LABOR SUPPLY

As our model successfully predicts labor supply, marriage, fertility, wages, and other key demographic outcomes for five cohorts, all of which faced different tax structures, it

[59] Goldin and Katz (2002) argued that the oral contraception pill was a key driver of increased education/employment of women.

seems credible to use it to predict the impact of changes in tax rules. Here, we use the model to simulate the effect of changing the U.S. tax code to eliminate joint taxation of married couples.

Most countries tax incomes of married individuals as if they were single, a policy known as individual taxation. The United States is one of the few countries that taxes couples' joint income, using a different tax schedule for married and single households (see Appendix B). This system of joint taxation, combined with the progressivity of the tax schedule, generates the so-called "marriage tax." For example, consider a married woman whose husband has high earnings. If the woman chooses to start working, she will face a high marginal tax rate on her first dollar earned (i.e., she is in the same marginal tax bracket as her husband). This can create a strong disincentive for married women to work (see Apps and Rees (2009)).

The impact of joint taxation can be substantial. For example, consider the 2011 tax rules. A married woman (with no children) whose husband earned \$88k would pay a tax rate of 25% on her very first dollar of income. In contrast, if she were single, she would only begin to pay a marginal rate of 25% once her income reached \$44k, which is roughly the mean earnings of 45 year old married women in the 1965 cohort (see Appendix C).[60,61]

In Table IV, we report the results of an experiment where we eliminate joint taxation of income for married couples in the 1965 cohort.[62] That is, we assume that each person in that cohort knows with perfect foresight, from age 17 to 65, that their earnings in each calendar year will be taxed according to the *individual* tax schedule in effect in that year, regardless of their marital status (or their spouse's earnings).[63] According to our model, elimination of joint taxation would have increased the labor supply of married women in the 1965 cohort by 8.3% over the ages from 25 to 55.[64] Married men increase their labor supply only slightly (0.6%), consistent with the view that joint taxation primarily impacts the labor supply of women.

The elimination of joint taxation also increases the incentive for young women to acquire human capital, as they are more likely to work when they are older. Thus, the college completion rate for women increases by 4.2%, the pre-tax wage rate of married women increases by 1.3%, and even single women work (0.6%) and earn (0.9%) a bit more.

Moreover, elimination of joint taxation increases the marriage rate by 8%, reduces the divorce rate by 4.3%, reduces fertility of married women by roughly 4%, and reduces fertility of single women by roughly 1%. Our ability to predict not only how elimination of joint taxation would affect behavior of existing married couples, but also marriage rates, education, and other demographics, highlights the value of our unified modeling framework.

A shift to individual taxation increases revenue by roughly 9%, as individuals face higher marginal rates at lower income cutoffs. We can render the policy revenue-neutral

[60]There are cases where a household faces lower taxes under a joint system. A good example is if both partners have high earning capacity, and both work. Alm, Diert-Conlin, and Whitington (1999) gave a nice discussion of joint taxation in the U.S.

[61]Bick and Fuchs-Schündeln (2018) showed how a typical married secondary earner in the U.S. faces tax rates as high as a typical single person in Sweden (which has much higher rates in general, but has individual taxation). They showed cross-country variation in joint taxation helps explain variation in married women's employment.

[62]This is the latest cohort with data over the main working years, until the age of 51.

[63]For years after 2017, agents assume the tax rules of 2017 will continue to apply (i.e., static expectations).

[64]We focus on ages 25 to 55 because we want to focus on how joint taxation affects behavior of prime age married women. But results would be similar if we looked at ages 17–65. Note that the elimination of the marriage tax penalty closes 45% of the gap of the employment rates between married and unmarried women.

TABLE IV

IMPLEMENTING INDIVIDUAL TAXATION OF INCOME FOR 1965 COHORT[a]

	1965				
	Baseline	Individual Tax	Percentage Change	Ind. Tax Revenue Neutral	Percentage Change
Gross Wages (Thousands of $)					
Married Women	41.9	42.4	1.3%	42.4	1.2%
Unmarried Women	42.0	42.3	0.6%	42.3	0.7%
Married Men	63.4	63.3	−0.2%	63.3	−0.2%
Unmarried Men	47.6	47.7	0.0%	47.7	0.1%
Employment					
Married Women	0.65	0.70	8.3%	0.71	9.0%
Unmarried Women	0.75	0.76	0.9%	0.76	1.2%
Married Men	0.89	0.89	0.6%	0.89	0.9%
Unmarried Men	0.76	0.76	−0.1%	0.76	0.2%
Family Moments					
Marriage Rate	0.68	0.73	8.0%	0.73	8.1%
Divorce Rate	0.12	0.12	−4.3%	0.12	−5.1%
Married Women # of Children	1.66	1.60	−3.9%	1.59	−4.0%
Unmarried Women # of Children	0.40	0.40	−1.1%	0.40	−1.3%
Education					
Women's CG + PC Rate	0.24	0.25	4.2%	0.25	4.2%
Men's CG + PC Rate	0.26	0.26	0.0%	0.26	0.0%

[a]Gross Wages—Average simulated annual wages of full-time workers aged 25 to 55. Employment—Average simulated employment rate of workers aged 25 to 55. Family moments—Average simulated rates for people aged 25 to 55. Education—Simulated college and post-college graduation rates at age 30.

by cutting all rates by 9.3% (see Table IV, right columns). Labor supply of married women now increases by an additional 0.7%, giving a 9.0% increase overall, but other quantities are little affected.

The large employment responses of married women in our tax experiments (8.3% or 9.0%) may appear consistent with the traditional view that labor supply elasticities of married women are large (see Keane (2011)). But our estimates of preference parameters α and γ imply that in a static single agent model with a linear budget constraint and continuous hours, the Marshallian elasticity is about −0.30 (see Appendix I). As $\alpha < 0$, the income effect dominates the substitution effect. Yet as Keane and Rogerson (2012) discussed, in a complex structural model like ours, labor supply elasticities are not simple functions of preference parameters, but depend on all aspects of the structure. Hence, they can only be calculated by simulation.

In Table V, we use our model to simulate Marshallian labor supply elasticities broken down by gender, marital status, age, and cohort. The results are obtained by simulating permanent 5% increases in offer wages. One clear result in Table V is that Marshallian elasticities for married men are modest in size and vary little by age or cohort. For example, in the 1965 cohort, for the three age ranges 25–34, 35–44, and 45–54, they range from 0.15 to 0.17. For single men, the elasticities are only slightly higher and they are again stable over cohorts. And the elasticities for single women are very similar to those of single men.

In contrast, Marshallian elasticities for married women are much greater (exceeding 1.0 for all cohorts/age groups). They are fairly stable across cohorts for older women (35–54),

TABLE V

LABOR SUPPLY ELASTICITIES BY GENDER, MARITAL STATUS, AGE, AND COHORT

Elasticities	1935	1945	1955	1965	1975
Married Women—Ages 25–34	1.80	1.84	1.27	1.25	1.13
Married Women—Ages 35–44	1.12	1.32	1.13	1.12	1.18
Married Women—Ages 45–54	1.20	1.10	1.04	1.06	
Unmarried Women—Ages 25–34	0.21	0.23	0.19	0.18	0.22
Unmarried Women—Ages 35–44	0.19	0.28	0.21	0.21	0.17
Unmarried Women—Ages 45–54	0.16	0.16	0.20	0.20	
Married Men—Ages 25–34	0.15	0.15	0.20	0.17	0.19
Married Men—Ages 35–44	0.14	0.17	0.20	0.15	0.17
Married Men—Ages 45–54	0.16	0.19	0.20	0.15	
Unmarried Men—Ages 25–34	0.16	0.16	0.20	0.18	0.23
Unmarried Men—Ages 35–44	0.17	0.20	0.21	0.16	0.16
Unmarried Men—Ages 45–54	0.21	0.18	0.16	0.22	

but drop substantially for younger women (25–34). This drop is concentrated between the 1945 and 1955 cohorts, when the elasticity for young married women fell from 1.84 to 1.27. As can be seen in Appendix Table J-I, employment of 25–34-year-old married women increased sharply from 38% to 55% from the 1945 to 1955 cohorts, with more modest increases afterwards.[65]

The large labor supply elasticities for married women suggest that reform of the tax code to eliminate joint taxation is a desirable strategy if one's goal is to enhance aggregate labor supply. Two other recent papers, Guner, Kaygusuz, and Ventura (2012) and Borella, De Nardi, and Fang (2017), also predict that the elimination of joint taxation would have large positive effects on labor supply of married women. These papers use quite different modeling strategies from ours, so this result appears to be robust to a variety of modeling choices.[66]

8. CONCLUSION

We studied the life-cycle decisions of five cohorts of U.S. men and women born from 1935 to 1975, a period of dramatic socio-economic change. We develop a life-cycle model that captures many of the key changes, and allows us to quantify the exogenous factors that drove them. Our model treats labor supply, education, marriage, and fertility as endogenous decisions. It accounts for human capital accumulation, the evolution of health,

[65] Estimating static labor supply functions for married women in the CPS, Blau and Kahn (2007) found elasticities dropped steadily and dramatically from the 1935 to 1975 cohorts for women of all ages. This contrasts sharply with our results where we find general stability, except for younger women between the 1945 and 1955 cohorts.

[66] Guner, Kaygusuz, and Ventura (2012) used an overlapping generations equilibrium model where marriage, fertility, and education are exogenous, and human capital is endogenous for women but exogenous for men. They predicted elimination of joint taxation would cause employment of married women to increase by 9.5% (compared to our 9.0% figure). The model in Borella, De Nardi, and Fang (2017) is a dynamic life-cycle model like ours, except they chose to make marriage, fertility, and education exogenous while introducing saving and a richer model of retirement behavior. They predicted that the elimination of joint taxation would increase married women's employment by roughly 10%. The agreement among our three studies is rather striking. Of course, the other studies cannot make predictions about marriage rates, education, and fertility as we do here—because they allow for saving, they must assume these variables are exogenous to maintain tractability (as we must assume equation (25) is invariant to our experiment).

and changes in tax and welfare rules. Prior work has not included *all* these features simultaneously.

We discipline our model by assuming fixed preferences, and show that it nevertheless provides an excellent fit to the observed changes in all the above listed behaviors across all five cohorts. The model is able to fit these changes using changes in four exogenous factors: (i) parental education, health, and taxes/transfers, (ii) changing marriage market conditions, (iii) changes in wage and job offer distributions, and (iv) changes in birth control technology.

Our estimates show the offer wage distribution has changed radically across cohorts. In the 1935 cohort, the wage structure facing women was inferior to that for men, as both (a) starting wages given education, and (b) returns to experience, were much lower. By the 1975 cohort, starting wages had almost converged, and experience returns for women had greatly improved (but still lagged men). Also, in the 1935 cohort, women were less likely to receive full-time job offers, but by the 1975 cohort, offer probabilities had almost converged.

A key fact we seek to explain is that, from the 1935 to 1975 cohorts, the employment rate of married women aged 25 to 34 increased sharply from 28% to 63%, while employment of other groups was fairly stable. Our model attributes roughly 1/2 of this increase to the dramatic improvement in the offer wage and job offer distributions for women. The second most important factor, accounting for 1/4 of the increase, was the advent of oral contraception.

Another important socio-economic change is that women passed men in educational attainment in recent cohorts. Our model implies that improvements in mothers' education (which enhance daughters' skill endowments and tastes for school), changes in the marriage market, and changes in the wage structure each account for roughly 1/3 of the increase in women's education. The benefits of a college degree arising *through the marriage market* increased greatly for women over time, because (i) divorce became easier, and education provides insurance in the event of divorce, (ii) college women became more likely to match with college men, and (iii) women became more likely to work after marriage.

Another key fact is that between the 1935 and 1975 cohorts, real wages of married women nearly doubled, and increased at twice the rate of wages of married men. Our model implies that changes in the wage structure *alone* account for 75–80% of this change. An interesting implication is that the increase in women's education did not, by itself, lead to much improvement in their wages. Only after women's returns to education/experience began to catch up to those of men did higher education translate into higher wages and employment.

The marriage rate fell substantially among young women, and we find better maternal education was a key driver of this change. But for women aged 35–44, improved labor market conditions were also an important factor. In contrast, the model implies the substantial rise in divorce rates was mostly due to changes in divorce costs. Our model captures that the wage rate of married relative to single women went from −9% in the 1935 cohort to +5% in the 1975 cohort. We find that 2/3 of this change was due to selection of higher ability women into marriage, while 1/3 was due to women acquiring more human capital after marriage.

Fertility dropped substantially over this period. Our model implies that availability of oral contraception explains about half the drop for married women and almost the entire drop for unmarried women. It also explains about 1/4 of the substantial increase in employment for younger women. But, in contrast to Goldin and Katz (2002), we find contraception had little or no effect on women's education or other demographic outcomes.

In summary, we see that no one factor can explain the multiplicity of socio-economic changes we observe over the past 50 years. Changes in all four exogenous factors we have considered were necessary in one or more dimensions to provide a good fit to all the key facts we sought to explain for the changes from the 1935 to 1975 cohorts.

Finally, we used our model to simulate a U.S. shift from joint taxation of couples to individual taxation. In a revenue-neutral scenario, this increases labor supply of married women by 9% (with small effects on other groups). Furthermore, as marriage and fertility are endogenous in our model, we predict the policy would increase the marriage rate by 8.1%, reduce the divorce rate by 5.1%, and increase the college completion rate of women by 4.2%.

In our model, generations are linked by the fact that educational attainment of parents affects preference and skill distributions of children. Our model fits the distribution of completed education very well for *every* cohort. Thus, while each cohort treats their parents' education as predetermined, our model explains the evolution of the education distribution across cohorts. As an extension, we could use the education choices predicted by our model for the 1965 cohort to generate initial conditions for the 1985 cohort, and so on. We can then predict how demographics like married women's employment rate will evolve in the future.

REFERENCES

ALBANESI, S., AND C. OLIVETTI (2016): "Gender Roles and Medical Progress," *Journal of Political Economy*, 124 (3), 650–695. [246]

ALM, J., S. DIERT-CONLIN, AND L. WHITINGTON (1999): "The Marriage Penalty," *Journal of Economic Perspectives*, 23 (3), 193–204. [247]

APPS, P., AND R. REES (2009): *Public Economics and the Household*. Cambridge Books. Cambridge: Cambridge University Press. [247]

ATTANASIO, O., H. LOW, AND V. SÁNCHEZ-MARCOS (2008): "Explaining Changes in Female Labor Supply in a Life-Cycle Model," *The American Economic Review*, 98 (4), 1517–1552. [218]

BECKER, G. (1974): "A Theory of Marriage: Part I," *Journal of Political Economy*, 81, 813–846. [218]

——— (1981): *A Treatise on the Family*. Cambridge: Harvard University Press. [218]

BEN-PORATH, Y. (1967): "The Production of Human Capital and the Life Cycle of Earnings," *Journal of Political Economy*, 75 (4, Part 1), 352–365. [230]

BICK, A., AND N. FUCHS-SCHÜNDELN (2018): "Taxation and Labor Supply of Married Couples Across Countries: A Macroeconomic Analysis," *Review of Economic Studies*, 85 (3), 1543–1576. [247]

BLAU, R., AND L. KAHN (2007): "Changes in the Labor Supply Behavior of Married Women: 1980–2000," *Journal of Labor Economics*, 25 (3), 393–438. [249]

BLUNDELL, R., M. COSTAS-DIAS, C. MEGHIR, AND J. SHAW (2016): "Female Labour Supply, Human Capital and Welfare Reform," *Econometrica*, 84 (5), 1705–1753. [230]

BORELLA, M., M. DE NARDI, AND Y. FANG (2017): "Marriage-Related Policies in an Estimated Life-Cycle Model of Households' Labor Supply and Savings for Two Cohorts," Report. [249]

BRONSON, M. A. (2015): "Degrees Are Forever: Marriage, Educational Investment and Life Cycle Labor Decisions of Men and Women," Manuscript, Georgetown University. [218,245]

CAPATINA, E., M. KEANE, AND S. MARUYAMA (2018): "Health Shocks and the Evolution of Consumption and Income Over the Life-Cycle," UNSW Working Paper. [236]

CHIAPPORI, P. A. (1988): "Rational Household Labor Supply," *Econometrica*, 56 (1), 63–90. [218]

——— (1992): "Collective Labor Supply and Welfare," *Journal of Political Economy*, 100 (3), 437–467. [218]

CHIAPPORI, P. A., M. COSTA-DIAS, AND C. MEGHIR (2018): "The Marriage Market, Labor Supply and Education Choice," *Journal of Political Economy*, 126 (S1), S26–S72. [244]

CHIAPPORI, P. A., B. FORTIN, AND G. LACROIX (2002): "Marriage Market, Divorce Legislation and Household Labor Supply," *Journal of Political Economy*, 110, 37–72. [240]

CHIAPPORI, P. A., M. IYIGUN, AND Y. WEISS (2009): "Investment in Schooling and the Marriage Market," *American Economic Review*, 99 (5), 1689–1713. [244]

ECKSTEIN, Z., AND O. LIFSHITZ (2011): "Dynamic Female Labour Supply," *Econometrica*, 79 (6), 1675–1726. [230,234]

ECKSTEIN, Z., AND K. WOLPIN (1989): "Dynamic Labour Force Participation of Married Women and Endogenous Work Experience," *Review of Economic Studies*, 56 (3), 375–390. [231]

ECKSTEIN, Z., M. KEANE, AND O. LIFSHITZ (2016): "Sources of Change in the Life-Cycle Decisions of American Men and Women: 1962–2014," CEPR DP 11393. [240,246]

——— (2019): "Supplement to 'Career and Family Decisions: Cohorts Born 1935–1975'," *Econometrica Supplemental Material*, 87, https://doi.org/10.3982/ECTA14474. [217,220]

FERNÁNDEZ, R. (2013): "Culture as Learning: The Evolution of Female Labor Force Participation Over a Century," *American Economic Review*, 103 (1), 472–500. [218]

FERNÁNDEZ, R., AND J. WONG (2011): "The Disappearing Gender Gap: The Impact of Divorce, Wages, and Preferences on Education Choices and Women's Work," Working Paper 17508, NBER. [240,244]

GE, S. (2011): "Women's College Decisions: how Much Does Marriage Matter?," *Journal of Labor Economics*, 29 (4), 773–818. [244]

GEMICI, A., AND S. LAUFER (2011): "Marriage and Cohabitation," Manuscript. [218]

GEWEKE, J., AND M. KEANE (2001): "Computationally Intensive Methods for Integration in Econometrics," *Handbook of Econometrics*, 5, 3463–3568. [225]

GIHLEB, R., AND O. LIFSHITZ (2016): "Dynamic Effects of Educational Assortative Mating on Labor Supply," Manuscript. [237]

GOLDIN, C., AND L. KATZ (2002): "The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions," *Journal of Political Economy*, 110 (4), 730–770. [218,246,250]

GOLDIN, C., L. KATZ, AND I. KUZIEMKO (2006): "The Homecoming of American College Women: The Reversal of the College Gender Gap," *Journal of Economic Perspectives*, 20 (4), 133–156. [243]

GOUSSE, M., N. JACQUEMET, AND J. ROBIN (2017): "Marriage, Labour Supply and Home Production," *Econometrica*, 85 (6), 1873–1919. [218]

GREENWOOD, J., AND A. SESHADRI (2005): "Technological Progress and Economic Transformation," in *The Handbook of Economic Growth*, Vol. 1B, ed. by P. Aghion and S. Durlauf. Amsterdam: Elsevier North-Holland, 1225–1273. [218]

GREENWOOD, J., N. GUNER, G. KOCHARKOV, AND C. SANTOS (2016): "Technology and the Changing Family: A Unified Model of Marriage, Divorce, Educational Attainment, and Married Female Labor-Force Participation," *American Economic Journal: Macroeconomics*, 8 (1), 1–41. [246]

GREENWOOD, J., A. SESHADRI, AND M. YORUKOGLU (2005): "Engines of Liberation," *Review of Economic Studies*, 72 (1), 109–133. [218]

GUNER, N., R. KAYGUSUZ, AND G. VENTURA (2012): "Taxation and Household Labour Supply," *The Review of Economic Studies*, 79 (3), 1113–1149. [249]

HECKMAN, J. (1974): "Shadow Prices, Market Wages, and Labor Supply," *Journal of Political Economy*, 42 (4), 679–694. [235]

IMAI, S., AND M. KEANE (2004): "Intertemporal Labor Supply and Human Capital Accumulation," *International Economic Review*, 45 (2), 601–642. [230,237]

JOHNSON, M., AND M. KEANE (2013): "A Dynamic Equilibrium Model of the U.S. Wage Structure, 1968–1996," *Journal of Labor Economics*, 31 (1), 1–49. [241]

JONES, L., R. MANUELLI, AND E. MCGRATTAN (2015): "Why Are Married Women Working so Much?," *Journal of Demographic Economics*, 81 (1), 75–114. [218,241]

KEANE, M. (2011): "Labor Supply and Taxes: A Survey," *Journal of Economic Literature*, 4, 961–1075. [248]

KEANE, M., AND R. MOFFITT (1998): "A Structural Model of Multiple Welfare Program Participation and Labor Supply," *International Economic Review*, 39 (3), 553–589. [227]

KEANE, M., AND R. ROGERSON (2012): "Micro and Macro Labor Supply Elasticities: A Reassessment of Conventional Wisdom," *Journal of Economic Literature*, 50 (2), 464–476. [248]

KEANE, M., AND K. WOLPIN (1997): "The Career Decisions of Young Men," *Journal of Political Economy*, 105 (3), 473–522. [218,228,230]

——— (2001): "The Effect of Parental Transfers and Borrowing Constraints on Educational Attainment," *International Economic Review*, 42 (4), 1051–1103. [228,230]

——— (2010): "The Role of Labor and Marriage Markets, Preference Heterogeneity and the Welfare System in the Life-Cycle Decisions of Black, Hispanic and White Women," *International Economic Review*, 51 (3), 851–892. [218,227,243,244]

LEE, D., AND K. WOLPIN (2006): "Intersectoral Labor Mobility and the Growth of the Service Sector," *Econometrica*, 74 (1), 1–46. [218,241]

LOW, C. (2014): "Pricing the Biological Clock: Reproductive Capital on the US Marriage Market," Manuscript, Wharton School of Business. [218,244,245]

MATOUSCHEK, N., AND I. RASUL (2008): "The Economics of the Marriage Contract: Theories and Evidence," *The Journal of Law and Economics*, 51 (1), 59–110. [245]

MAZZOCCO, M. (2007): "Household Inter Temporal Behaviour: A Collective Characterization and a Test of Commitment," *The Review of Economic Studies*, 74 (3), 857–895. [226]

MAZZOCCO, M., C. RUIS, AND S. YAMAGUCHI (2007): "Labor Supply, Wealth Dynamics and Marriage Decisions," Manuscript. [218]

MCFADDEN, D. (1989): "A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration," *Econometrica*, 57 (5), 995–1026. [234]

MINCER, J. (1974): *Schooling, Experience and Earnings*. New York: National Bureau of Economic Research. [222,230]

MULLIGAN, C., AND Y. RUBINSTEIN (2008): "Selection, Investment, and Women's Relative Wages Over Time," *The Quarterly Journal of Economics*, 123 (3), 1061–1110. [243]

PAKES, A., AND D. POLLARD (1989): "Simulation and the Asymptotics of Optimization Estimators," *Econometrica*, 57, 1027–1057. [234]

RUST, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55 (5), 999–1033. [237]

STEVENSON, B. (2008): "Divorce Law and Women's Labor Supply," *Journal of Empirical Legal Studies*, 5 (4), 853–873. [218]

VOENA, A. (2015): "Yours, Mine and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?," *The American Economic Review*, 105 (8), 2295–2332. [218,226,245]

WILLIS, J., AND S. ROSEN (1979): "Education and Self-Selection," *Journal of Political Economy*, 87 (5), S7–S36. [236]

Co-editor Liran Einav handled this manuscript.

Manuscript received 14 June, 2016; final version accepted 20 July, 2018; available online 8 August, 2018.