The Disappearing Gender Gap: The Impact of Divorce, Wages, and Preferences on Education Choices and Women's Work

Raquel Fernández*
New York University

Joyce Cheng Wong
New York University

September 2011

PRELIMINARY†

Abstract

Women born in 1935 went to college significantly less than their male counterparts and their labor force participation (LFP) as married women averaged 40% between the ages of thirty and forty. The cohort born twenty years later behaved very differently. The education gender gap was eliminated and married women's LFP averaged 70% over the same ages. In order to evaluate the quantitative contributions of the many significant changes in the economic environment, family structure, and social norms that occurred over this period, this paper develops a dynamic life-cycle model calibrated to data relevant to the 1935 cohort. Respecting the asymmetric nature across gender of the economic consequences of divorce, we find that the higher probability of divorce and the changes in wage structure faced by the 1955 cohort are each able to explain, in isolation, a large proportion (60%) of the observed changes in female LFP. Combining all economic and family structure changes, we find that a simple change in preferences towards work can account for the remaining change in LFP. To eliminate the education gender gap requires, on the other hand, an asymmetric change by gender in the psychic cost of obtaining higher education.

*NBER, CEPR, IZA, ESOP.

The author wishes to thank the NSF and the Russell Sage Foundation for financial support.

†We thank Gianluca Violante for helpful comments.

1 Introduction

A white woman born in the US in 1935 experienced a very different work and education trajectory than her male counterpart. In this cohort, 29% of women went on to college and, once they married, only 40% worked on average during the ages of thirty to forty. Men, on the other hand, essentially always worked and over a third more of them (39%) went to college. This gender gap in work and education was dramatically transformed over a twenty year period. For the cohort born in 1955, the same proportion of white women and men attended college (44%) and women's labor force participation increased to the point that 70% of married women worked on average between the ages of thirty to forty. As shown in Figures 1 and 2, the percentage point increase in both education (college) and married women's labor force participation (LFP) in the time interval separating these two cohorts was the largest witnessed over any other twenty year period in US history. Furthermore, as shown in figure 3, the changes in married women's LFP occurred at all stages of their life cycle and these changes were significantly larger than those experienced by the cohorts born either 10 years earlier or later.

A search for candidates capable of explaining why these women's lives were transformed yields an embarrassment of riches. This is a period that witnessed profound changes in the economy, in family structure, and in social norms. During this period the skill premium increased, the gender wage gap fell, and there is evidence that the returns to experience increased, childcare costs fell, and that technological change occurred both in the household and in the wider economy. At the family level there were also important transformations: fertility dropped, people married later and more assortatively, and the probability of a marriage ending in divorce more than doubled. At the cultural level, expectations about women's role changed dramatically: June Cleaver gave way to Madonna's material girl.[1] Opinion polls clearly demonstrate the change in popular opinion towards working women. For example, the proportion of individuals who approved of a wife working if her husband was capable of supporting her went from 17% in 1945 to 63% in 1970.[2]

In theory, any of the above-cited changes could have played an important role in transforming women's lives. The goal of this paper is thus to examine the quantitative contributions of some of these basic forces to the changes in female LFP and education across these two cohorts. In choosing which drivers to focus on, we were particularly interested in the role of divorce. This interest comes from noting two important facts. First, as shown in Figure 4, the 1955 cohort had the highest divorce

[1] June Cleaver – a principal character in the sitcom "Leave it to Beaver" – was the archetypal housewife.
[2] Source: Gallup Poll Data (1945, 1970).

probabilities of any other cohort and continued a trend of high divorce rates that had started ten years later.[3] Hence, it is plausible to argue that women and men in the 1955 cohort anticipated high divorce rates and changed their behavior accordingly.[4]

Second, today's women have overtaken men in educational attainment not only in the US, but also in almost all OECD countries.[5] The cross-country nature of the twin phenomena of women's rapidly increasing LFP and the reversing of the education gender gap points towards drivers of change that operate across many societies and, for education, incentivate women more than men.[6] Divorce is indeed such a force. The overwhelming tendency for children to live with their mothers and the fact that the monetary transfers from ex-husbands are relatively small leaves women significantly worse off than their husbands when they divorce. This would increase women's incentive to work and operate asymmetrically on each gender's desire to obtain more education. Furthermore, the large increase in divorce rates all over the developed world would make this force ubiquitous.[7] While other drivers like increased skill premia have also operated in many countries, there is no evidence that these are consistently higher for women than for men.[8] A consistent quantitative analysis nonetheless requires incorporating all potentially significant drivers of education and work choices.

In order to investigate the quantitative significance of economic and social change on the education gap between men and women and on female LFP, we develop a dynamic life-cycle model and calibrate it using data for the cohort born in 1935. In the multi-period model, agents that differ by gender and ability levels make education decisions, enter marriage markets and, in their working stage of life, make labor, consumption, and savings decisions. They are subject to income, marital and fertility shocks, all conditional on gender and education (as in the data), until retirement whereupon they simply make consumption and savings decisions. Women's disutility of working is allowed to depend on marital status and on the presence and age of children in the household. These disutilities will be used to match women's LFP by age, education, and marital status. Although much of our focus is on the LFP of

[3] Figure 4 plots the cumulative probability of a marriage ending in divorce as a function of the number of years married for couples who initially married in a given decade. As seen in the figure, the failure rate by a given anniversary year of marriage is highest for marriages initiated in the 1975-'84 decade, which is when most of the 1955 cohort married.

[4] This would be less convincing for the cohort born ten years earlier and also saw high divorce rates (but lower than for the 1955 cohort).

[5] Mexico, Turkey and Switzerland are the exceptions.

[6] Indeed, as remarked upon by Goldin, Katz, and Kuziemko (2006), "any explanation of how U.S. women have caught up and surpassed men in college trends should be consistent with this common pattern of international changes."

[7] For example, Kneip and Bauer (2009) show that the divorce rate quadrupled on average for Western European countries between 1960 and 2000.

[8] Across 21 OECD countries, for example, 8 have higher skill premia for women whereas the other 13 have higher skill premia for men (see Strauss and de la Maisonneuve (2009)).

married women, we also target the LFP of divorced women since their welfare is relevant to women's education and labor market decisions.

We find that both changes in family and wage structure had significant effects on the LFP of married women. Each of these changes alone can account for almost 60% of the change in LFP observed between our two cohorts of married college women.[9] These changes account for a smaller proportion of the LFP gap for high-school women (48.2% due to wages and 33.1% due to family). We find that a uniform 10% drop in work disutility across all education and marital status types is able to explain the remaining LFP gap across education and marital status.

Our paper contributes to the literature on education choice and female LFP. While there is a vast literature on schooling decisions, there are very few papers that have investigated why women caught up and then surpassed men in their educational achievement.[10] A primarily theoretical paper by Chiappori, Iyigun, and Weiss (2009) argues that an increased skill premia coupled with the marriage market can motivate an asymmetric reaction in women's versus men's education decisions. Closer to our work in the sense of using a quantitative model, Ge (2010) uses NLSY79 data to structurally estimate a model with endogenous education, female LFP, and marital choices but with no savings-consumption decisions in a setting with a linear utility function. She finds that the marriage market returns (in the form of a higher-earning husband) are an important part of the overall returns to acquiring higher education. There is no analysis, however, of the quantitative significance of marital instability (divorce) on education and women's work decisions in part because her analysis is based only on one cohort. Studying a much longer time period but without a theoretical model, Goldin, Katz, and Kuziemko (2006) conduct an insightful study into the historical evolution of women's versus men's education. They argue that relatively larger economic benefits from college for women coupled with lower psychic education costs than men (e.g. through better preparedness) were the main contributors to this phenomenon.

With respect to work, there is a large and growing literature that examines why women's LFP has changed over different time periods.[11] A recent paper by Eckstein and Lifshitz (2011) provides

[9] Note that the effects of changes are not additive.

[10] Some key papers on the drivers of education choices include, for example, Cameron and Heckman (1998), Cameron and Heckman (2001), and Cameron and Taber (2004). This literature finds that ability and family background explain a large part of education decisions. Notable papers which also look at the education decisions of women include Keane and Wolpin (1997), Keane and Wolpin (2010), and Eckstein and Wolpin (1999). All these papers, however, use the NLSY79 sample, and thus conduct an analysis over only one cohort.

[11] The classic source for an economic history of female labor force participation is Goldin (1990). A more recent overview is given by Costa (2000). Most explanations emphasize technological change in the home or workplace (e.g., Galor and Weil (1996), Greenwood, Seshadri, and Yorukoglu (2005), Attanasio, Low, and Sanchez-Marcos (2008)), or changes in medical/contraceptive technology (e.g., Goldin and Katz (2002), Albanesi and Olivetti (2009a,b), and Knowles (2007a)).

an excellent review of this literature. Their paper, like ours, seeks to provide an explanation of the changing behavior across cohorts of women. They conclude that education and wage changes are largely responsible for the observed evolution of LFP across several cohorts born between 1925 and 1975 and that marital status changes played a small role. Like us, the authors use a dynamic life-cycle model to estimate the contributions of various drivers but, in order to estimate their model, they assume a linear utility function.[12] Our paper, on the other hand, not only endogenizes the schooling decisions of men as well as women but also employs a non-linear utility function. The latter may be key to a proper evaluation of the quantitative importance of the increased risk of divorce.

A notable paper that also employs a life-cycle model with concave utility to study women's choices across several cohorts is Attanasio, Low, and Sanchez-Marcos (2008). Their focus is primarily on the role of changing childcare costs on women's LFP, a factor that our model will also consider. The authors find that a combination of increased female wages and decreased childcare costs is necessary in order to explain the change in LFP of married women between the 1945 and 1955 birth cohorts. Their paper, however, abstracts from changes in marriage markets by assuming that women are married during their whole lives and thus the authors are unable to evaluate the contribution of changes in divorce rates to women's work.

Our paper is organized as follows: the next section presents the basic statistics for our two cohorts, 1935 and 1955. Section 3 develops the dynamic stochastic life-cycle model and section 4 discusses its parametrization. Section 5 presents key features of the benchmark model calibrated to 1935 data. Section 6 investigates the effects of changes in family structure, wage structure, and the possible role of changes in culture in generating the education and female LFP outcomes. It also examines the welfare consequences of these changes. Section 7 shows the results of some robustness checks and section 8 concludes.

2 The Lives of Two Cohorts: 1935 and 1955

In this section we present some key features of the environment faced by our two cohorts. To do this, we mainly use the 1962-2010 waves of the March supplement of the Current Population Survey (CPS), a

Other explanations emphasize changes in wage structure (e.g., Jones, Manuelli, and McGrattan (2003), Gayle and Golan (2010) and Knowles (2007b)).

[12] This linearity, however, has the benefit of a large reduction in computational time making feasible a structural estimation. The presence of concavity, borrowing and savings, and heterogeneity in our model greatly increases the necessary time for each iteration, rendering estimation too burdensome.

cross-sectional survey conducted by the Bureau of the Census. Although this is not a panel, we choose this dataset due to its long time span which allows us to observe the full life span of our cohorts. We construct synthetic cohorts from the cross-sectional data. Our "1935" cohort consists of white men and women born between 1934-1936, while our "1955" are those born between 1954-1956.[13] Note that although our sample is selected based on race, we do not restrict the race of the spouses of the people in our sample. Married people are defined as those "married, with spouse present", singles are those people who report "never married" while divorced people are those who report their marital status to be either "divorced" or "separated".[14]

2.1 LFP and Education

The labor supply of our two cohorts are markedly different at all ages. Figure 5 shows LFP rates for each cohort, by age and education for married and divorced women separately.[15] There are several things to note. First, divorced women (dashed lines) always worked more than their married counterparts (solid lines) at every age, and for both education levels. These LFP differences across marital status, however, are much smaller for the 1955 cohort. Second, college women (circles) work more than high school women (triangles) throughout most of their lives for both cohorts. It is interesting to note, however, that while the difference in LFP across education groups is very small for the 1935 cohort during the first ten years in the figure – corresponding to the ages at which women tend to have young children – this gap increases markedly for the 1955 cohort.[16] This is also the period of life in which the largest increase in LFP (over 100%) occurred between the two cohorts. While married women from the 1935 cohort had strikingly low LFP rates during youth (34.6% for college and 26.4% for high school at in the ages of 25-29), the young married women in 1955 cohort have LFP rates of 70.4% for college and 54.9% for high school. The LFP of divorced women in 1955 is also higher than that of their 1935 counterparts. The increase in LFP is largest at youth, but smaller than the increase seen for married women.

Both men and women increased their education during this time period. The increase for women was almost triple that of men. For the 1935 cohort, 39.2% of the men had some college whereas only 29.1% of women did – a 10 percentage point gap. For the 1955 cohort, the gap disappeared: 44.5% of

[13] We choose to focus on whites as the historical experiences of black men and women have been very different. In particular, black women have worked significantly more than white women throughout. See Potamites (2007).

[14] Also, we do not include widowed people.

[15] The LFP for each age is calculated using data constructed for the two synthetic cohorts from the CPS.

[16] Note that this is not driven by the working behavior of women with children. These large increases in LFP occurred for all groups of married women, both with and without young children.

men went to college compared to 43.7% of the women.[17]

2.2 Marriage Market and Fertility

Over this same period, the divorce probabilities faced by each cohort also changed significantly. Table 1 reports, by gender and education, the proportion of marriages that ended in divorce before their twentieth year.[18] Note that divorce probabilities more than doubled for women and increased by around 70-80 percent for men. Remarriage probabilities stayed fairly constant. It is interesting to note that while women do not show large differences in divorce probabilities by education, college men are less likely to divorce (and more likely to remarry) than their high-school counterparts for both cohorts.

Conditional upon getting married, the assortativeness of marriage increased slightly between the two cohorts. The correlation of spouses in years of education increased for women from 0.61 to 0.64 over these twenty years; for men the increase was a bit smaller – from 0.63 to 0.64. It is probably useful to note though that this is not the most significant indicator of the returns to education in the marital market. Instead, it may be more useful to note that by age 35-39, conditional upon being married, the proportion of women married to college men increased. Reflecting the large increase in the availability of college women, both high school and college men significantly increased the likelihood with which they are married to a college woman, with almost three quarters of higher educated men marrying a college woman. These proportions are reported in Table 2.

Fertility also changed markedly during this period. Women from the 1935 cohort on average had 3.0 children. This statistic differs by education level with college women having, on average, 2.5 children and high school women, 3.2 children. For the 1955 cohort the family size was much smaller: 1.9 children. Again, college women had on average fewer children than their high school counterparts (1.7 versus 2.1). The median age at first birth for the 1935 cohort was 22 for high school women versus 24 for college women. For the cohort born 20 years later, the median age at first birth was not very different: it actually became younger for high school women (age 21) but older for college women (age 25).

[17] Education attainment here is measured at age 30.

[18] Note that in Table 1 we report the divorce and remarriage frequencies for the birth cohort. This is preferable to the marriage cohorts used in Figure 4.

2.3 Wages

The skill premium and the gender wage ratio also changed for our two cohorts. The average ratio of college to high school lifetime hourly wages went from 1.42 for men in the 1935 cohort to 1.54 in the 1955 cohort. For women, the corresponding averages are 1.38 in the 1935 cohort and 1.50 in the 1955 cohort. Thus both sexes saw an increase of some 8% across cohorts. Women's wages also gained ground with respect to those of men. The average ratio of female to male lifetime hourly wages increased from 0.61 for the 1935 cohort to 0.72 in the 1955 cohort.[19]

3 The Model

In this section, we will first outline some concerns with the assumptions employed in the model, then we proceed to explain in detail the economic environment and the household's decision problems.

3.1 Modelling Considerations

We choose a life-cycle model since one of our main variables of interest – female LFP – has a critical life-cycle component. We allow for endogenous consumption and savings decisions (unlike others, e.g., Eckstein and Lifshitz (2011) and van der Klaauw (1996)) particularly because we are interested in the importance of the increased risk posed by divorce. To the extent that households can save, they may be able to buffer this and other risks and ignoring this possibility may quantitatively affect the incentives to work or acquire higher levels of education, particularly for women.

An important question we faced was whether to model an equilibrium marriage market with endogenous marriage/divorce decisions. We chose not to do so for two key reasons. First, in the data, a significant portion of any cohort marries outside its own cohort. For example, for the 1955 cohort, by the age of 25-29 around 69% of women were married whereas only 60% of men. Hence one cannot assume the existence of a simple marriage market which clears in equilibrium by having an equal number of women marry an equal number of men from the same cohort. An alternative would have been to use an overlapping generations structure and solve for its steady-state equilibrium. The data preceding the 1935 cohort and following the 1955 cohort, however, clearly indicates that neither cohort was inhabiting a stationary environment. This is evident in both of our education and female LFP figures depicting the evolution of these variables (see figures 1 and 2) which show a significant upward trend throughout

[19] Wages are calculated from CPS data. See the Appendix for further details.

the twentieth century. Thus, the non-stationary features of the data strongly contradict the assumption of a stationary equilibrium. Solving a non-stationary model requires keeping track of the evolution of the distribution of assets, by gender and education, of divorced people as this affects agents' expected utility of divorce and doing this is beyond our capacity.[20]

Having eliminated the general equilibrium features of the marriage market, one could still opt to keep divorce decisions endogenous. In order to do this one would need to impose assumptions about the distribution of characteristics of prospective spouses in the marriage market (namely, their assets, income, ability, etc) as these affect the payoffs of marriage/divorce. Given the absence of data concerning this distribution for the 1935 cohort, we prefer the transparent and straightforward alternative of modelling marriage and divorce as exogenous shocks and of assuming that the characteristics of agents' (future) spouses are given by a latent marriage market type.

In particular, our simplifying assumption throughout is that agents, in terms of their potential marital partner, are of a given type that is revealed to them once they've made their education decision. Thus, before deciding their education they simply know the conditional probability with which they'll meet different education and ability types. Afterwards although they continue to face uncertainty about marital status, they know the characteristics of their *potential* future partner in terms of ability, education, current income and experience. This is their marital "type". This allows agents to predict the evolution of the asset position of their potential future spouse, radically reducing the computational burden at the cost of decreasing the uncertainty individuals face with respect to their spouse's future income stream. We think this is a good assumption especially as one can argue that individuals are of some latent type in terms of who they marry and that this type gets revealed to them over time.

While treating a critical variable such as divorce as exogenous is problematic in terms of understanding causality, we are fortunate that in our time period of choice there were important legal changes, exogenous to the individual, that affected the risk of divorce. In particular, the introduction of unilateral divorce law significantly increased the risk of divorce since it was no longer necessary to obtain consent from one's spouse in order to divorce. In addition, the social stigma associated with divorce decreased thus also contributing to increased divorce risk, exogenously for the individual.[21]

[20] Note that we cannot reduce the distribution of assets to a few state variables as is done, for example, in Krusell and Smith (1998) or Heathcote, Storesletten, and Violante (2010).

[21] For example, for most countries surveyed in both the 1991 and the 2005-2008 waves of the World Values Survey, the proportion of people who found divorce "always justifiable" significantly increased from the first to the second wave. In some countries, namely France, Great Britain, the US, Sweden and Russia, this proportion nearly tripled.

Our model incorporates various sources of heterogeneity in addition to gender. Most importantly from the perspective of education choice, we allow individuals to differ in their ability. Ability affects wages differentially by education and thus influences education choices. While wages are not determined entirely within the model in the sense that there are no general equilibrium considerations (e.g. supply and demand), they are nonetheless endogenous to the individual. Wages are a function of the individual's exogenous attributes (gender and ability) and also endogenous choices such as education and experience and thus need to be solved within the model. This is especially important for women since it recognizes that observed wages are a result of selection that operates both in work decisions made that period and also in the past in the form of accumulated experience.

3.2 Some Preliminaries

Agents are born either male m or female f and with an ability endowment θ and an initial asset position a_0 both of which are iid draws from genderless distributions. An individual's life cycle can be divided into three distinct stages of life. In the first stage (the education/spousal-allocation stage), the individual attends school and decides whether to acquire a higher education. That done, the individual gets matched in a education-dependent marriage market. The agent emerges from that initial marriage market belonging to one of three possible households with marital status s: married ($\mathbf{m}$), divorced ($\mathbf{d}$) or single ($\mathbf{s}$) and with their marital type revealed. Next, the individual enters the second stage of life (the working stage), in which they receive fertility shocks, wage draws and martial shocks, and make consumption-savings and female labor force participation decisions. Lastly, in the third stage of life (the retirement stage), individuals no longer work nor suffer shocks to their marital status. Their retirement income is deterministic; they make consumption and savings decisions under certainty.

Preferences, Consumption, and Borrowing Constraints

The instantaneous utility function of an agent with gender g and education level e (with marital status s) is given by:

$$U_{eg}(c_t, P_t; s) = \frac{c_t^{1-\sigma}}{1-\sigma} - \psi^s_{eg}(k_t)P_t \tag{1}$$

where c is consumption and P_t denotes the labor force participation decision, taking the value one if the agent works and zero otherwise. We will henceforth assume that only women suffer disutility from market work, $\psi^s_{eg}(k_t)$, that may depend on education, marital status, and the vector k_t indicating the ages of her children in that period. We normalize men's and single women's work disutility to zero and thus assume that they always work (accordingly, we will not match any empirical LFP moments for them).

Household consumption can be thought of as a public good with congestion. If the household spends $\hat{c}_t$ on consumption goods, this yields

$$c_t = \frac{\hat{c}_t}{\mathbf{e}\left(k_t; s\right)} \tag{2}$$

units of household consumption. Thus $\mathbf{e}\left(k_t; s\right)$ gives the economies of scale that exist which depend on the ages and number of children and whether there are one or two adults in the household (hence the s).[22]

Agents' borrowing is only constrained by the no-bankruptcy condition $a_{T+1} \geq 0$ which imposes that agents must pay off all their debt before they die. Our choice of a utility function with infinite marginal utility of consumption at zero consumption will ensure that the agent is bounded away from the constraint.

Divorce and Children

Women obtain marital-dependent fertility shocks at the beginning of a period. In keeping with modelling one of the major asymmetries between the sexes, we assume that when a couple divorces the children stay with the mother. We use k_t^{ij} to denote a vector that keeps track of the age of each child that woman i had with husband j. If a woman divorces, both she and her ex-husband continue to share the same k^{ij} variable until they remarry. Note that if there are no children, $k_t \equiv 0$. Children remain with the parents/mother until they become adults (at the age of 20 in the model). They make no decisions but deflate household consumption accordingly.

In the advent of divorce, assets are split between the two ex-spouses, with a proportion α of assets going to the wife and $1 - \alpha$ to the husband. Furthermore, we assume that the man pays his ex-wife a proportion of his income as child support as long as she remains both unmarried and with a child under 20. Agents are assumed to remain divorced for at least one model period, i.e. they cannot receive a remarriage and a divorce shock at the same time.

Recall that given our previous assumption, agents know the characteristics of their potential spouse. For computational ease we assume also that when a divorced agent remarries, so does their ex spouse.

[22] As explained in the parametrization section, we use an altered McClemens scale whose exact specification is outlined in the Appendix.

To further simplify matters, upon remarriage each is assumed to marry an agent whose characteristics are the same as those of their ex spouse (in terms of education, ability, fertility, experience, and income shock).[23] Thus, we simply reassign the children to the newly remarried couple. Finally, if the wife does not remarry by the time she enters retirement, she receives a proportion of her husband's retirement benefits.

Income

In each period of the work stage of life, individuals receive wage draws and then decide whether to work. The income process is uncertain. It has an idiosyncratic persistent (z_t) and a transitory component (η_t), and is a gender and education-specific function of experience (x_t) and ability (θ) that takes into account the human capital depreciation that occurs if the agent did not work the prior period, i.e., $y_{egt} = y_{eg}(\theta, x_t, z_t, P_{t-1})$.

In each period during retirement, an individual receives retirement income $b_g^s(\overline{y})$ that is a function of gender, past earnings and marital status.

3.3 The Education and Initial Marital Status Stage

In the initial stage, period 0, an individual goes to school and decides whether to obtain more education given her ability θ. The level of education is a discrete choice with two possible outcomes denoted by $e \in \{l, h\}$, $h > l$.[24] If an agent chooses to obtain more education, the education outcome is h; otherwise it is l. In this initial period there are no consumption, saving, or work decisions.

Obtaining education level h is costly. An individual i is assumed to suffer disutility ω_i from obtaining higher education, where ω_i is a draw from a distribution $C^g(\omega)$ that potentially differs by gender g, $g \in \{m, f\}$. Thus an agent will decide an education level so as to solve:

$$V_{eg,t_0}(\theta) = \max_{e\in\{l,h\}} \left\{ -C^g(\omega) + \mathbb{E}V^s_{hg,\bar{t}_0}(\theta), \quad \mathbb{E}V^s_{\mathrm{lg},\bar{t}_0}(\theta) \right\} \tag{3}$$

where $\mathbb{E}V^s_{eg,\bar{t}_0}(\theta)$ is the expected value of the agent's welfare at the begining of period zero (denoted $\bar{t}_0$) given ability level θ, sex g, and education level e.[25] Expectations here are taken both with respect to

[23] An alternative possibility would be to assume that remarriage is to an agent with the same asset holdings as one's own, as in Voena (2010). Our specification has the advantage of respecting gender differences in asset accumulation behaviors.

[24] In the data l will be high-school graduates and below and h will be individuals who have more education than a high-school diploma.

[25] Henceforth, we use $\bar{t}$ to denote the value of a variable before the realizations of the end-of-period shocks.

the household type s that the individual will draw at the end of this initial period, given her education decision, as well as with respect to all subsequent education-specific shocks to income, marital status, and fertility.

After completing education, individuals enter the marriage market where they face an exogenous, education-and-gender-specific probability vector $p_{eg,t_0}(s)$ of obtaining a given marital status s. An individual emerges from the marriage market either married ($\mathbf{m}$), divorced ($\mathbf{d}$), or still single ($\mathbf{s}$), i.e., $s \in \{\mathbf{m}, \mathbf{d}, \mathbf{s}\}$ and with knowledge of her/his potential-spouse type.[26]

3.4 The Work Stage

Individuals are assumed to spend periods 1 until period t^R in this stage. In each period, every individual receives an education-and-gender-specific wage draw that is a function of the individual's history, in particular her/his work experience and past wage draws. Given these draws, households make both consumption-savings and work decisions. How these decisions are made differs by marital status. Singles and divorced individuals make these decisions to maximize their life-time utility. For married individuals, on the other hand, household consumption and the wife's labor force participation are chosen, as in Chiappori (1988), so as to maximize the weighted sum of the spouses' lifetime utilities.[27] Thus we are assuming that the household allocation is Pareto efficient.[28]

Households are also subject to fertility shocks at the beginning of each period and marital shocks at the end. Asset, consumption and participation choices are made after fertility shocks and income are observed. Households get divorced with probability d_{eg}, divorced individuals remarry with probability r_{eg}, and singles marry with probability $m_{eg,t}$. Women receive fertility shocks ϕ_{est} that differ by education and marital status. A time line showing periods 0 and 1 of an individual's life is given in Figure 6 to clearly illustrate the timing of shocks and decisions.

Budget constraints while working:

Married individuals:

[26] A divorced individual is one who was married at some point and then divorced during this initial period. In the data this will be prior to the age of 30.

[27] A husband's labor force participation could also be an outcome but since we assume his disutility of working to be zero, leaving this outside the maximization problem is innocuous.

[28] See Del Boca and Flinn (2010) for an insightful analysis of the various modelling approaches one can use to determine household allocations.

If married in period t, a household makes its consumption and savings decisions subject to a household budget constraint knowing that at the end of the period its marital status may change as a result of a shock. The married household's budget constraint is given by:

$$\hat{c}_t(k_t) + a^{\mathbf{m}}_{\bar{t}+1} = Ra^{\mathbf{m}}_t + (y_{eft} - \kappa(k_t))P_t + y_{emt} \tag{4}$$

where a^m_t is the married household asset holdings entering period t. Income consists of capital income from last period's asset, where R is the gross return, and y_{egt} is each spouses' labor market income from which expenditures on childcare, $\kappa(k_t)$, are subtracted if a woman works. $\hat{c}$ is the expenditure on consumption and $a^{\mathbf{m}}_{\bar{t}+1}$ is the household's asset position before the realization of the marital shock at the end of period t. At the end of a period t, shocks to marital status may leave the household divorced, which then affects their asset holdings. We assume that a fraction $\alpha \in (0,1)$ of the assets $a^{\mathbf{m}}_{\bar{t}+1}$ are allocated to the wife and the remainder to the husband. Thus, the laws of motion for $a^s_{g,t+1}$ are given by:

$$a^s_{g,t+1} = \begin{cases} a^{\mathbf{m}}_{t+1} = a^{\mathbf{m}}_{\bar{t}+1} & \text{if } s_{t+1} = \mathbf{m},\ g = m, f \quad \text{(i.e. the couple enters } t+1 \text{ still married)} \\ a^{\mathbf{d}}_{f,t+1} = \alpha a^{\mathbf{m}}_{\bar{t}+1} & \text{if } s_{t+1} = \mathbf{d},\ g = f \\ a^{\mathbf{d}}_{m,t+1} = (1-\alpha)\, a^{\mathbf{m}}_{\bar{t}+1} & \text{if } s_{t+1} = \mathbf{d},\ g = m \end{cases} \tag{5}$$

Divorced individuals:

In addition to the reallocation of marital assets right after divorce, we assume that husbands must also make transfer payments to wives in subsequent periods if they have children. In particular, the former husband must pay some fraction of his current income to his ex-wife as child support $h(k_t, y_{mt})$ as long as the child is not an adult.[29] Recalling that the children reside with their mother, the budget constraint of a divorced woman at time t is given by:

$$\hat{c}_t(k_t) + a^{\mathbf{d}}_{f,\bar{t}+1} = Ra^{\mathbf{d}}_{ft} + (y_{ft} - \kappa(k_t))P_t + h(k_t, y_{mt}) \tag{6}$$

[29] In the model, as long as his child is below the age of twenty.

whereas that of a divorced man is:

$$c_t + a^{\mathbf{d}}_{m,\bar{t}+1} = Ra^{\mathbf{d}}_{mt} + y_{mt} - h\left(k_t, y_{mt}\right) \tag{7}$$

where $a^{\mathbf{d}}_{\bar{t}+1}$ is the asset position prior to the realization of the marital shock at the end of period t. Note that a divorced man's consumption equals his expenditure on consumption (since his household consists only of himself).

At the end of period t, the shock to marital status can transit a divorced individual into remarriage. Thus, a divorced individual i that had saved $a^{\mathbf{d}}_{i,\bar{t}+1}$ faces at the end of period t the following law of motion for the asset position upon entering period $t+1$:[30]

$$a^{s}_{i,t+1} = \begin{cases} a^{\mathbf{d}}_{i,t+1} = a^{\mathbf{d}}_{i,\bar{t}+1} & \text{if } s_{t+1} = \mathbf{d} \;\; \text{(i.e., if } i \text{ enters } t+1 \text{ still divorced)} \\ a^{\mathbf{m}}_{t+1} = a^{\mathbf{d}}_{i,\bar{t}+1} + a^{\mathbf{d}}_{j,\bar{t}+1} & \text{if } s_{t+1} = \mathbf{m} \;\; \text{(i.e., if } i \text{ enters } t+1 \text{ remarried to } j\text{)} \\ & \end{cases} \tag{8}$$

Singles:

Single women and men are assumed to always work. Their budget constraint differs only if the woman has a child (whereupon she must pay childcare). Thus, for a single woman

$$\hat{c}_t(k_t) + a^{\mathbf{s}}_{\bar{t}+1} = Ra^{\mathbf{s}}_t + y_{ft} - \kappa\left(k_t\right) \tag{9}$$

whereas for a single man:

$$c_t + a^{\mathbf{s}}_{\bar{t}+1} = Ra^{\mathbf{s}}_t + y_{mt} \tag{10}$$

At the end of a period t, shocks to marital status can transit a single individual i into marriage, whereupon the assets of the spouse i and j are combined, i.e.,[31]

$$a^{s}_{i,t+1} = \begin{cases} a^{\mathbf{s}}_{i,t+1} = a^{\mathbf{s}}_{i,\bar{t}+1} & \text{if } s_{t+1} = \mathbf{s} \text{ (i.e., if } i \text{ enters } t+1 \text{ still single)} \\ a^{\mathbf{m}}_{t+1} = a^{\mathbf{s}}_{i,\bar{t}+1} + \;\; a^{\mathbf{s}}_{j,\bar{t}+1} & \text{if } s_{t+1} = \mathbf{m} \;\; \text{(i.e., if } i \text{ enters } t+1 \text{ married to } j\text{)} \end{cases} \tag{11}$$

[30] Recall that a divorced agent is marrying another divorced agent by assumption, hence i and j both have **d** superscripts in equation (8)

[31] Recall that divorced people married other divorce individuals; hence singles must marry other singles.

3.5 Retirement

In periods t^R through T all individuals are retired and hence do not work. They still make a consumption-savings decision each period that depends on household type and receive a pension $b_g^s(\overline{y})$ that depends on own past earnings history $\overline{y}$, gender, and present household type s. In particular, divorced men may be required to transfer some of their pension to their ex-wife. There are no longer any child support payments at this stage and all individuals are assumed to die at the end of period T. Recall that we assume that individuals cannot die with debt. Thus the budget constraints are given by:

$$\begin{aligned} \hat{c}_t + a_{t+1}^s &= Ra_t^s + b_g^s \\ a_{T+1}^s &\geq 0 \end{aligned} \tag{12}$$

3.6 Optimization Problems

In this section we outline each household type's optimization problem. Before doing so we introduce individual i's state vector $\Omega_{it} = \{a_{it}, x_{i,t-1}, P_{i,t-1}, k_t, z_{it}, e_i, \theta_i, \zeta_{it}\}$. The state variable keeps track of an individual's assets, experience (for a man, his age, whereas for a woman how many periods she worked in the past), whether i worked last period, the number and ages of children, the permanent component of the income shock, an individual's education and ability, and one's marital type, ζ_{it}. In addition to an agent's potential (or actual or ex) spouse's permanent characteristics (ability and education), ζ_{it} also tracks the evolution of the spouse's asset holdings, experience, participation, and permanent income shock, all as of period t.

Divorced agent's problem:

We now characterize the value of being divorced. Upon retirement, the optimization problem is simple given that there is no uncertainty. For an individual of gender g it is given by:

$$V_{gt}^{\mathbf{d}}(\Omega_t) = \max_{c_t, a_{t+1}} \frac{c_t^{1-\sigma}}{1-\sigma} + \beta V_{g,t+1}^{\mathbf{d}}(\Omega_{t+1}|\Omega_t) \text{ for } t \geq t^R$$
$$s.t. \text{ retired divorced budget constraint for } g \text{ in eq. (12)}$$

During the work phase of life, the divorcee has an exogenous education and gender dependent probability of remarrying r_{eg}. Upon remarriage, the maximization problem faced by that individual corresponds to that of a married household's. In any given period $t < t^R - 1$, the divorcee chooses consumption, saving, and, if she's a woman, whether or not to work.[32] Thus, the value of being a divorced woman at time t is given by:

$$V^{\mathbf{d}}_{ft}(\Omega_t) = \max_{c_t, P_t, a^{\mathbf{d}}_{t+1}} \frac{(c_t)^{1-\sigma}}{1-\sigma} - \psi^{\mathbf{d}}_e(k_t)P_t + \beta\left\{(1-r_{eg})\mathbb{E}\left[V^{\mathbf{d}}_{f,t+1}(\Omega_{t+1}|\Omega_t)\right] + r_{eg}\mathbb{E}\left[V^{\mathbf{m}}_{f,t+1}(\Omega_{t+1}|\Omega_t)\right]\right\}$$

$$s.t. \quad \text{divorced female budget constraint (6) and asset law of motion (8)} \tag{13}$$

$$\text{and} \quad x_{t+1} = x_t + P_t$$

where the expectation is taken over future shocks to her income, fertility, and her potential marital partner's income as well as possible future marital status shocks.

The divorced man's problem is identical to that of the divorced female except that there is no participation choice and the budget constraint and value functions are appropriately modified. The value of being a divorced man at time t is:

$$V^{\mathbf{d}}_{mt}(\Omega_t) = \max_{c_t, a^{\mathbf{d}}_{t+1}} \frac{(c_t)^{1-\sigma}}{1-\sigma} + \beta\left\{(1-r_{eg})\mathbb{E}\left[V^{\mathbf{d}}_{m,t+1}(\Omega_{t+1}|\Omega_t)\right] + r_{eg}\mathbb{E}\left[V^{\mathbf{m}}_{m,t+1}(\Omega_{t+1}|\Omega_t)\right]\right\} \tag{14}$$

$$s.t. \text{ divorced male budget constraint (7) and asset law of motion (8)}$$

Married Household's Problem

A couple that enters the retirement period married solves

$$V^{\mathbf{m}}_t(\Omega_t) = \max_{c_t, a^{\mathbf{m}}_{t+1}} \frac{c_t^{1-\sigma}}{1-\sigma} + \beta V^{\mathbf{m}}_{t+1}(\Omega_{t+1}|\Omega_t) \tag{15}$$

$$s.t. \quad \text{retired married budget constraint (12) and asset law of motion (5)}$$

for all $t \geq t^R$. This is relatively simple problem because, in the absence of marital shocks and labor decisions, the spouses agree on the optimal allocation unlike during the working period.

[32] Note that the problem at time $t = t^R - 1$ is slightly different since the continuation value is given by the solution to the retirement stage problem.

In the working stage, however, a couple that enters a period $t < t^R - 1$ married solves[33]

$$
\begin{aligned}
V_t^{\mathbf{m}}(\Omega_t) = \max_{c_t, P_t, a_{t+1}^{\mathbf{m}}} \quad & \chi \left[\frac{c_t^{1-\sigma}}{1-\sigma} - \psi_e^{\mathbf{m}}(k_t)\, P_t\right] + (1-\chi)\left[\frac{c_t^{1-\sigma}}{1-\sigma}\right] + (1-d_{eg})\beta\mathbb{E}\left[V_{t+1}^{\mathbf{m}}(\Omega_{t+1}\,|\Omega_t)\right] \\
& + d_{eg}\beta\left\{\chi\mathbb{E}\left[V_{f,t+1}^{\mathbf{d}}(\Omega_{f,t+1}\,|\Omega_t\right] + (1-\chi)\mathbb{E}\left[V_{m,t+1}^{\mathbf{d}}(\Omega_{m,t+1}\,|\Omega_t)\right]\right\} \qquad (16) \\
& s.t. \quad \text{the budget constraint for married couples (4) and asset law of motion (5)} \\
& \text{and} \quad x_{t+1} = x_t + P_t
\end{aligned}
$$

where χ denotes the Pareto weight of the wife. Note that $\Omega_{f,t+1}$ and $\Omega_{m,t+1}$ are the state variables of the wife and husband, respectively.[34]

We use an asterisk to denote the resulting outcomes from the optimization problem above. Thus, the value of being married for a man is given by:

$$
V_{mt}^{\mathbf{m}}(\Omega_t) = \frac{c_t^{*1-\sigma}}{1-\sigma} + (1-d_{eg})\beta\mathbb{E}(V_{m,t+1}^{\mathbf{m}}(\Omega_{t+1}^*\,|\Omega_t) + d_{eg}\beta\left\{\mathbb{E}\left[V_{m,t+1}^{\mathbf{d}}(\Omega_{m,t+1}^*\,|\Omega_t)\right]\right\} \qquad (17)
$$

whereas the value of being married for a woman is given by

$$
V_{ft}^{\mathbf{m}}(\Omega_t) = \frac{c_t^{*1-\sigma}}{1-\sigma} - \psi_e^{\mathbf{m}}(k_t)\, P_t^* + (1-d_{eg})\beta\mathbb{E}\left[V_{f,t+1}^{\mathbf{m}}(\Omega_{t+1}^*\,|\Omega_t)\right] + d_{eg}\beta\left\{\mathbb{E}\left[V_{f,t+1}^{\mathbf{d}}(\Omega_{f,t+1}^*\,|\Omega_t)\right]\right\} \qquad (18)
$$

Single Household's Problem

An individual of gender g who enters the retirement period single solves:

$$
V_{gt}^{\mathbf{s}}(\Omega_t) = \max_{c_t, a_{t+1}^{\mathbf{s}}} \frac{c_t^{1-\sigma}}{1-\sigma} + \beta V_{t+1}^{\mathbf{s}}(\Omega_{t+1}\,|\Omega_t) \text{ for } t \geq t^R \qquad (19)
$$

$$
s.t. \quad \text{retired single budget constraint (12) and asset law of motion (11)}
$$

[33] As before, that the problem at time $t = t^R - 1$ is slightly different since the continuation value is given by the solution to the retirement stage problem.

[34] Note that we are slightly abusing notation by using Ω_t as the household's state vector. Since both Ω_{ft} and Ω_{mt} have the same information, however, this is just a question of formal notation.

An individual of gender g who enters period $t < t^R - 1$ single solves:[35]

$$
\begin{aligned}
V^{\mathbf{s}}_{gt}(\Omega_t) = \max_{c_t, P_t, a^{\mathbf{s}}_{t+1}} \; & \frac{c_t^{1-\sigma}}{1-\sigma} - \psi^{\mathbf{s}}_{eg}(k_t)\, P_t \\
& + (1 - m_{eg,t})\beta \mathbb{E}\left[V^{\mathbf{s}}_{t+1}(\Omega_{t+1} \,|\Omega_t)\right] + m_{eg,t}\beta \mathbb{E}\left[V^{\mathbf{m}}_{g,t+1}(\Omega_{f,t+1} \,|\Omega_t)\right] \qquad (20) \\
& s.t. \quad \text{the budget constraint for singles (9) or (10) and asset law of motion (11)} \\
& \text{and} \quad x_{t+1} = x_t + P_t
\end{aligned}
$$

4 Parametrization

In this section we describe the calibration of the model. In order to respect the fact that men and women marry outside their cohorts we solve the model twice, once with the marriage market probabilities faced by the women and the other with these given by the probabilities faced by men of our 1935 cohort. The statistics generated by the model therefore are the right ones to compare with the data which are reported for each gender of the cohort.[36]

Some model parameter values are taken from preexisting estimates from the literature, others are estimated directly from the data using model restrictions. The remaining set of parameters are calibrated inside the model in order to match certain moments in the data. The reasoning guiding different choices is explained below. Table 3 shows the parameters estimated "outside" the model and Tables 4 and 5 displays the "internally" calibrated parameters.

Our key parameters of interest are those that govern work discipline for women and education decisions for both sexes. These parameters include those which govern disutility from labor and education, childcare costs and several parameters which affect wage dynamics. We next proceed to explain the choices of functional form and their calibration in detail.

Demographics and Preferences

The model period is 5 years. Individuals begin the working stage of life at age 25 (period $t = 1$) where they remain for 7 periods. Retirement begins in the model period $t^R = 8$ (thus at age 60) and death occurs at the end of model period $T = 12$ (at age 85).

[35] As before, that the problem at time $t = t^R - 1$ is slightly different since the continuation value is given by the solution to the retirement stage problem.

[36] In other words, we are not reporting what the spouses of our agents do but rather what the women and men of our cohort do.

Parametrizing our utility function requires specifying the parameter σ governing the coefficient of relative risk aversion, the discount factor β, and the disutility from work for various categories of women (the $\psi_e^{\mathbf{s}}(k)$). Since most estimates for the relative risk aversion parameter in the literature vary between one and two, we chose $\sigma = 1.5$ for our utility function specification. This value is in line with the values found by Attanasio and Weber (1995) using US consumption data. We set the discount factor $\beta = 0.90$ (for a five year period) which corresponds to a conventional yearly discount factor of 0.98.

The disutility of labor varies by marital status, education, and number of children. They are calibrated internally in the model in order to match female labor force participation rates by marital status and education for the 1935 cohort. With respect to children, we distinguish between mothers with young children (below the age of 5) and those with older children. These categories are meant to capture various distinctions across different women's disutility from working. The distinction between working as a single versus as a married or as a divorced woman was particularly relevant to the 1935 cohort who grew up thinking of married women primarily as homemakers (the June Cleaver generation). One may also wish to distinguish between working with a pre-school child at home versus without.[37] Lastly, these costs could differ by education, reflecting the fact that jobs may have non-monetary rewards that affect women's disutility of working (e.g. the distinction made between a "job" and a "career"). We normalize the disutility of working for single women to zero (the same as that of all men) and assume that there is no additional child-related disutility of working once a woman's children are over the age of 5. The results for the internal calibration of these disutility parameters are reported in Table 4 and discussed in the benchmark model section.

Education Costs

Parametrizing the cost of education is a challenge. First, given that there is not data on the ability distribution of individuals with some college relative to the rest of the population for either of our two cohorts, we assume a cost that is invariant to ability though the monetary reward is not. Second, although the cost of acquiring a higher education is modeled as a psychic one, it can also be seen as consisting of both a monetary and psychic component.[38] Using data to calibrate the monetary portion

[37] In addition to foregoing maternal bonding time with a child, there is also some evidence of negative effects from the mother choice to work on the child's intellectual and emotional development. See, for example, Bernal (2008) and Bernal and Keane (2009) for some of the findings in this literature.

[38] The computational model shows that individuals are not borrowing constrained and hence they would be able to cover (reasonable) financial costs of higher education via borrowing. Doing this correctly, however, requires modifying the budget constraint to account for the cost of education.

of the cost would not be helpful, however, as it still leaves a free parameter. Given the absence of data, we choose instead to pin down the cost of a higher education by requiring the model to generate the correct proportion of individuals attending college by gender. We assume that these (psychic) costs are a random draw from a gender specific log-normal distribution LN $\sim (\mu_g, \sigma^2)$. We normalize the standard deviation for both genders to $\sigma = 1$ and we internally calibrate μ_f, μ_m to generate the correct proportions as reported in Table 5. The economic meaning of these parameters is discussed in the benchmark model section.

Income Process

For an individual of gender g with education level e, her/his wage at time t is given by $y_{eg,t}$ such that:

$$\ln y_{eg}(\theta, x_t, z_t, P_{t-1}) = \tau_{eg,t} + \gamma_{eg1}x_t + \gamma_{eg2}x_t^2 + \lambda_e \ln \theta - \delta(1 - P_{t-1}) + w_{et} \tag{21}$$

where τ_{egt} captures a time varying component in aggregate wages, by education and gender, and $\gamma_{eg1}, \gamma_{eg2}$ are education and gender specific experience polynomials. λ_e captures the returns to individuals ability θ by education group, while δ is human capital depreciation incurred from not working in the previous period. Note that since men always work, $x_{m,t} = t$ and $P_{m,t-1} = 1$, $\forall t < t_R$, whereas for women, $x_{ft} = \sum_{\tau=1}^{t} P_{f\tau}$.

There is also a stochastic component to wages w_t^e that is assumed to be the sum of an (observable to the agent) persistent component (z_{et}) and a transitory component (η_{et}):

$$w_{et} = z_{et} + \eta_{et} \tag{22}$$

The z_{et} persistent shock is modeled as an AR(1) process

$$\begin{aligned} z_{et} &= \rho_e z_{e,t-1} + \epsilon_{et}, \\ \epsilon_{et} &\sim N(0, \sigma_{\epsilon_e}) \end{aligned} \tag{23}$$

while the transitory shock is distributed $\eta_{et} \sim N(0, \sigma_{\eta_e})$. This choice of model for the stochastic process is standard in the literature and it is consistent with both the sharp drop in the autocovariance function for wages between lags 0 and 1 and also with the large increase in the variance of wages observed in the data over the life-cycle.[39]

Ability is a key variable in the model as it affects the return to education. One way to measure its contribution to wages is to proxy it with an individual's AFQT scores and then estimate how wages differ by ability and education. Since there is no dataset which reports AFQT scores for either of our cohorts, we assume that returns to ability did not change over time and use the NLSY79 dataset to estimate these returns. The NLSY79 is a panel survey of a representative sample of 12,686 American young men and women who were 14-22 years old when they were first surveyed in 1979. This is an ideal dataset since its panel structure permits one to follow the evolution of individual earnings and includes individual scores on the AFQT test. The fact that one can only observe individuals' wages until 2009, however, implies that the median age is 49 which makes it difficult to estimate age profiles $\gamma_{eg1}x_t + \gamma_{eg2}x_t^2$ with precision. To deal with this, we first use wage data from PSID to estimate age polynomials. Although there is no ability information in that data, its panel nature implies that we can estimate the effect of age/experience by education since ability is assumed to only affect the level of log wages. We can then estimate the return to the AFQT score in the NLSY wages, having first subtracted from them the estimated age coefficients. This is the strategy used by Gallipoli, Meghir, and Violante (2010) in order to estimate λ_e for each education level.

Ideally we would like to estimate the parameters above separately for men and women. Since there is selection in observed female wages we instead estimate returns to ability using men's data and assume they take the same value for women. We split the AFQT89 scores into quintiles and assume that agents draw ability from a uniform distribution over those quintiles. Thus each agent's ability corresponds to the median value of AFQT89 score in one of the ability quintile bins.

We use all the waves from 1968 to 2007 of the PSID in order to estimate age-earnings profile for the men in our two education groups.[40] These parameters are assumed to not change across our two cohorts. Using data on earnings and hours worked for white males aged between 18 and 65 we estimate a 2nd degree polynomial for each education group. The estimated values for the $\gamma_{em1}, \gamma_{em2},\ e = \{h, l\}$ parameters are reported in Table 3 and details of our estimation procedure and sample selection are described in the appendix. Finally, we use the residuals from the wage equations from the NLSY (free

[39] This specification and its estimation is discussed in great detail in, for example Storesletten, Telmer, and Yaron (2004) and Guvenen (2007), as are the characterstics of the autocovariance functions and the variance growth in the life cycle which motivate the functional form of the stochastic process.

[40] The Panel Study of Income Dynamics (PSID) is the longest panel survey conducted in the US, starting in 1968. Interviews were conducted on an annual basis until 1997, and from then onwards, biennially.

of age and ability effects) to estimate the stochastic component of wages w_{et} specified in equation (22).[41]

The time varying intercepts $\tau_{em,t}$ for men are calculated from the data using first differences such that each education level's wage growth rates are consistent with the growth rates observed in the data. Assuming that men always work and that education choices are made by age 25 implies that men's ability and education are fixed effects which disappear when one first differences the data. This procedure yields all the time intercepts except for that in period one; accordingly, we calibrate $\tau_{em,1}$ internally within the model. The implied intercepts for 1935 men are plotted in Figure 7 and the first differences calculated from the data for the 1955 cohort are shown in Figure 8.

We assume that women share the same returns to ability by education (λ_e) and the same stochastic wage process (w_{et}) as men. Concerning the parameters for returns to experience $\gamma_{ef1}, \gamma_{ef2}$, the literature has found values in the range of 2 to 5% returns to wages from one year of participation for women born in later cohorts (1940s onwards). Since there is evidence that the returns to experience has increased over time, we choose parameters which imply returns to wages from an extra year of participation, during the ages of 25-40, of 2% for women in the 1935 cohort.[42] Finally, we do not assume that the time varying intercepts for women ($\tau_{ef,t}$) are the same as men's and instead calibrate them and the male $\tau_{em,1}$,$e = h, l$ internally in order to respect the selection into education by ability in our model. Estimating these parameters from the data would necessitate a dataset that included measures of wages and ability for the 1935 cohort which, as mentioned previously, does not exist. These parameters are instead calibrated such that the model generates the correct period by period gender wage ratio and skill premia for women. The implied values of $\tau_{eg,t}$ are plotted in Figure 7.

Our model abstracts from alimony since the evidence in the data shows that both the proportion of divorced people who receive it and the monetary amounts are very small.[43] Child support is a more common and substantial payment. For example, Del Boca and Flinn (1995) find it to be about 20% of the father's income. The rate of non-compliance, however, is fairly high at 37%. Beller and Graham

[41] We use the Minimum Distance Estimator (MDE) (Chamberlain (1984)) to estimate the parameters for the stochastic wage process; this procedure is standard in the literature. Heathcote, Storesletten, and Violante (2010) is an excellent source for further details and identification. We thank Gianluca Violante for kindly providing us with the estimation code.

[42] For example, Attanasio, Low, and Sanchez-Marcos (2008)'s calibrated returns to experience imply a returns of about 2% for the cohort of women born around 1945. Olivetti (2006) estimates the returns to one year of full time work to be between 3 and 5% using data in the 1970 Census and she also finds an increase of almost 90% in returns to experience using the 1990 Census.

[43] For example, Voena (2010) uses the National Longitudinal Survey of Young and Mature Women to show that only 10% of divorced women report receiving alimony between 1977 and 1999 and the monetary amounts correspond to only about 15% of the divorced woman's household income. Using 1978 CPS data, Peters (1986) reports alimony payments which correspond to under 3% of the average male earnings that year.

(1988) report an average child support payment of $1115 in 1978. Given an average male wage that year of around $13,000, this amounts to 8.7% of the male wage; these authors also find a high rate of non-compliance (over 50%). In the light of this evidence, we assume that as long as his ex-wife has children under the age of 20, the man pays her child support equivalent to 10% of his current income (unless she remarries). We do robustness checks using other values for this parameter.[44]

After retirement, for computational simplicity (as in Guvenen (2007)), individuals in our model receive a constant pension which is a function of her/his last observed earnings. The exact functional form of the pension system mimics the US Social Security bendpoints (following Heathcote, Storesletten, and Violante (2010)) and it is outlined in the appendix. Married couples receive either the sum of the husband's and wife's pensions or 1.5 times the husband's pension (whichever one is higher). A divorced woman receives, in addition to her own pension, 10% of her ex-husband's pension.[45]

Family Formation and Fertility

In the model, after agents choose their education level, they receive marital status shocks and then enter the working stage as married, single or divorced with the proportions found in CPS data for our cohorts at age 25-29. These proportions are given in Figure 9, and they vary by gender, education and cohort.

As discussed previously, all agents are assigned a marital "type" that is revealed in period zero after the agent's choice of education level. That type permanently determines the characteristics of one's spouse independently of whether the agent enters period 1 as married, single, or divorced. For each agent, given her/his education, we assign a spouse (type) so as to match the conditional distribution of spouses' education as seen in the data for people aged between 35-39 for each respective cohort. These proportions (conditional upon one's education), are reported in Table 2. Thus, the pattern of marriages by education will mimic the degree of assortativeness found in the data. The ability level of one's spouse is assumed to be a random draw from the (endogenous) ability levels within the spouse's education group.

The probability of marriage for agents who enter a period single is calculated directly from the

[44]In our model, we abstract from the risk of non-compliance by the father in the payment of child support by making these payments certain. Uncertainty in these payments would increase the effects of a higher divorce rate.

[45]The laws governing an ex-wife's claim to the man's pension have evolved over time. Before 1980 unvested pensions were not considered part of marital property. Currently, pensions are divided as part of marital property and they are frequently the most valuable portion of the marital real estate (see Oldham (2008)). From our reading, it is not clear what proportion of the pension to grant to the ex wife (it depends on years married, etc.) but our assumption of 10% seemed appropriate. In our robustness check we investigated other proportions as well.

evolution over time of the proportion of people who are never married in the CPS data for the 1935 and 1955 cohorts. These probabilities are reported in Table 6, by age, education and gender.

Next we need divorce and remarriage rates. The main data set which records individuals' marital histories is the Survey of Income and Program Participation (SIPP) conducted by the Census. Ideally one would keep track of each agent's marriage duration conditional upon year of marriage (age). Doing so, however, would add significantly to the computational complexity since it would increase the state space. Moreover, it is difficult to estimate with precision the probabilities of divorce conditional on both year of marriage and duration due to small sample size in each year-of-marriage/duration bin for each of our cohorts. Given these considerations, we choose a simple transparent alternative. From the 2004 SIPP data we calculate, conditional on gender and education, the proportion of marriages that fail before their twentieth anniversary and assume a uniform divorce rate over this twenty year interval.[46] We allow this rate to differ by gender and education. The remarriage rates are calculated analogously. The resulting divorce and remarriage rates are reported in Table 7.[47]

Fertility shocks are education and marital status dependent and are calibrated to yield both the proportion of women who are mothers during the ages of 25-29 for each cohort, by education level and marital status, and to generate the average number of children, by education level, as in the PSID data for each cohort.[48] In particular, for the 1935 cohort, single women and college women of all marital status receive a fertility shock in the first period that takes the value of zero or one. Divorced women and married high school women receive a fertility shock that takes the value of zero or two. The probability of receiving a non-zero value is calibrated so as to match the initial proportions (age 25-29) in the data by marital status and education as reported in Table 8. In period 2, divorced and single women are not hit by fertility shocks whereas all women who are married in period 2 are assumed to have an additional

[46] We do this for all individuals from each of our birth cohorts augmented by 2 years due to sample size (i.e. the 1935 cohort is defined as all people born between 1933-1937). For individuals who marry after the age of 29 (but before the age of 60), we count the fraction of marriages which end in divorce before their 20th anniversary, provided that the anniversary takes place before the individual reaches the age of 60. For those individuals who enter the age of 30 already married, their contribution to marriage duration is counted as years of marriage after the age of 29. We do this because, for all those people who enter the age of 30 as married, we already took into account their marital status when we calculated our initial proportions of people in each marital state category.

[47] Recall that we assume that our agents cannot divorce and remarry at the end of the same period and thus they effectively remain divorced for at least 5 years. This fits well with the data because the median duration of divorce is 5 years for the 1935 cohort and 6 years for the 1955 cohort.

[48] We choose to respect the proportion of women (by education and marital status) who are mothers in the data because LFP behavior is mostly driven by the presence of a young child at home rather than by the number of children women have. Moreover, given that our childcare costs and additional disutility from labor depend only the presence and age of the youngest child in the household, we chose to focus on the distinction between being a mother versus not (this is the strategy followed by Attanasio, Low, and Sanchez-Marcos (2008)).

child. Lastly, all women who were married in both periods 1 and 2 receive an additional fertility shock in the second period that can take the value of zero or two for high school women, and zero or one for college women. The frequency of shocks is calibrated to generate the following numbers: 2.54 children per college woman, 3.20 for high school women and an overall average of 3.00 children per woman for the 1935 cohort.[49]

For the 1955 cohort the structure and values of fertility shocks are the same as those for the 1935 cohort, with the frequencies adjusted so as to match the proportions in the data for the 1955 cohort between the ages of 25-29 as reported in Table 8. In the second period, once again, no divorced or single women receive any additional children. All women who were not married in period zero and got married at the end of period one receive a child. Lastly, once again, all women who were married in both periods are hit by a fertility shock that can take the value of zero or one.[50] This generates the following numbers: college women have 1.74 children on average, high school women have 2.11 and, overall there are 1.91 children per woman for the 1955 cohort.

Consumption deflator and child care costs

Children are assumed to live with their parents (or mother, if parents are divorced) until the age of 20. They make no decisions in the household but deflate consumption according to their age. We choose to use the McClements scale to calculate the economies of scale in consumption.[51] Its exact numbers (by child's age) are reported in the Appendix.

Women who have children under the age of 10 at home are assumed to incur childcare costs if they work. These costs are modelled as only depending on the age of the youngest child, i.e. if a household has a young and an old child, they only incur the childcare cost once, for the younger child. We calibrate the childcare costs for young children (aged 0-4), κ_{young} and for old children (aged 5-9), κ_{old} internally in the model and their values are reported in Table 5.

[49] The numbers for average number of children were computed using the PSID by calculating the average number of children ever born to women from each of our cohorts by the time they reached aged 40. Due to sample size constraints, we define the 1935 cohort as women born in 1933-1937 and the 1955 cohort those born in 1953-1957. In order to generate the correct number of average children, we assign a 98% probability of a non-zero fertility outcome during the second period for women who have been married for two periods if they have college education and 43% probability if they have high school.

[50] Once again, to generate the correct number of average children seen in the data, we assign the probabilities of receiving an extra child during the second period to women who have been married for 2 periods of 73% if they are college and 52% if they are high school.

[51] This scale is very similar to the OECD scale, but it has the advantage that it was computed based on expenditure data from families.

Other External Parameters

An important parameter in the model is the Pareto weight on a woman's welfare in the household allocation problem. There is no real guidance as to what this should be. A recent paper by Voena (2010), using variations in savings behavior and divorce laws, estimates a value of 0.25. In our benchmark calibration we set this value, χ, to 0.3; the robustness section investigates the effect of changing this weight.

Upon divorce, we assume that the woman is responsible for the children and that assets are split equally between the husband and wife (i.e., $\alpha = 0.5$). In the data, at the time when most of the divorces were occurring for our 1935 cohort, most states either had equitable distribution laws or community property laws. In the former, asset division is dictated by court of law, which may impose an equal split or favor either the spouse who contributed more towards the asset or the one who has higher needs. Under community property law, assets (and debts) are divided equally across the spouses. Thus, an equal split is a good benchmark. In our robustness section we explore slightly higher and lower values and find similar quantitative results. Finally, the gross interest rate is set to $R = 1.077$ which, in this five-year-period model corresponds to an annual interest rate of 1.5%. This is the average real return on a 3 month t-bill over the period of 1935-2008.

To summarize, we have a total of 29 parameters which we calibrated internally: 8 parameters which govern the disutility of labor for married and divorced women, 2 parameters for childcare costs, 14 time-varying wage-intercept parameters for women and 2 for men, one wage depreciation parameter and 2 education cost parameters. We choose to match a total of 45 statistics for our 1935 cohort: 28 average LFP rates for married and divorced women, by age and education, 7 time-varying gender wage ratios and the proportions of men and women who go to college. We also match 7 time-varying skill premia for women and the average lifetime skill premium for men.[52] Although this mapping is only approximate, it may be useful to think of the time-varying wage-intercept parameters as mainly targeting the skill premia and gender wage ratio statistics. The education cost parameters help us match the proportions of men and women who go to college, while the remaining parameters are used to match LFP.

[52]Note that we only target one skill premium statistic for men using $\tau_{em,1}, e = \{h, l\}$, since the remainder values are calculated from first differences in the data on male wages.

5 The 1935 Benchmark

As seen in Figure 10, the benchmark model does an excellent job of reproducing the LFP profiles for both married and divorced women according to their education and also matches the proportions of men and women who choose to go to college (see Table 5). Furthermore, as shown in Figure 11 the model is also quite successful at matching the period-by-period gender wage ratio and skill premia for women.

The disutilities of working in various states are reported in Table 4. They are higher, across all marital-fertility categories, for high-school women than for college women. The ratios of costs across categories is fairly similar for both education groups.

To make economic sense of the disutility from work numbers we can calculate their equivalent consumption cost. A simple way to translate the cost into consumption units is to calculate the decrease in average consumption that women would be willing to bear to avoid these costs, i.e., to find the z such that

$$u(z\bar{c}^{\mathbf{m}}_{fe}) = u(\bar{c}^{\mathbf{m}}_{fe}) - \psi_{es}(k) \tag{24}$$

where $\bar{c}^{\mathbf{m}}_{fe}$ is the average per-period consumption of married women with education level e.[53] These proportions (z) are shown in the lower panel of Table 4. The percentage loss in consumption due to disutility from labor is similar, within each education category, for married and divorced women without children. It increases markedly for women with children, particularly for married high-school women. This large number is generated by the model in order to "explain" why these women have such low LFP rates. Note from figure 5 that fewer than 35% of high-school women are working during the ages of 25-34 despite the fact that their consumption is lower than their college counterparts' who have, on average, higher-earning husbands.

The calibrated childcare costs for a young child correspond to around 69.3% of average per-period female wages, and those of an older child are 23.7%. We can compare these values to those obtained by Attanasio, Low, and Sanchez-Marcos (2008). In their calibrated model they find that their childcare costs are 66% of a woman's mean earnings.[54]

The mean of the (lognormal distribution) of education costs are similar for men and women, though

[53] To calculate $\bar{c}^{M}_{fe}$ we find the average consumption of married women of education e in each period and then average across periods. We choose to express all percentages in terms of married women's average consumption as the pool of divorced women is constantly changing.

[54] This number is reported in terms of the mean earnings of a thirty-year-old woman who worked continuously prior to childbirth.

somewhat lower for the latter. This reflects the finding that, for all ability levels, the ex-ante welfare gap between college and high school is larger for men than it is for women. Why is this the case? This is mostly due to a slightly higher skill premium for men combined with the fact that they work considerably more than women rendering this skill premium more important. Recall that in our model men (and single women) are assumed to always work whereas women's average LFP over their lifetime is only 53.3%.

Since education (college) costs are incurred only once, it is easier to interpret their economic significance by solving for their equivalent cost in terms of the average per period consumption of married college women.[55] To do so we solve for the z_g such that:

$$\sum_{t=1}^{T} u(z_g \overline{c}^{\mathbf{m}}_{fh,t}) = \sum_{t=1}^{T} \beta^t u(\overline{c}^{\mathbf{m}}_{fh,t}) - \bar{C}_g \tag{25}$$

where $\bar{C}_g$ is the average cost of higher education (by gender) implied by the calibrated parameter μ_g and $\overline{c}^{\mathbf{m}}_{fh,t}$ is the average consumption of married college women at time t. Using this metric, women's education costs correspond, on average, to a 16.0% decrease in per period consumption, whereas the consumption reduction implied by the average cost of higher education for men is 17.7%.[56] Note that it is possible to think of this psychic cost as having a monetary component (tuition and foregone wages) in addition to a purely psychic component. This would simplify the interpretation of their magnitudes but would require us to introduce an additional borrowing decision in period zero. Thus producing a monetary plus psychic estimate is not an immediate calculation and would still leave (as in Gallipoli, Meghir, and Violante (2010)) a free psychic component to interpret.

Some Implications of the Calibrated Model

The model generates statistics which were not directly targeted in the calibration and thus provide additional checks of the model. Recall that we did not directly target the LFP of women with children and thus a comparison with the data is informative. Due to the CPS data's lack of information regarding the presence of young children in the household for our time period, we instead use Census data. Using

[55] We choose to express the education costs in terms of a fixed consumption category (married college women) since the shape of preferences imply that the consumption equivalence of a fixed cost C is increasing in consumption, i.e., the z that solves $u(zc) = u(c) - C$ is an increasing function of c.

[56] An alternative is to consider the average cost incurred by those who actually choose higher education (denote this by $\hat{C}_g$). Computing z_g yields an equivalence of 11.7% of consumption for women and 13.8% of consumption for men.

the 1960 and 1970 Census, we can observe the cohort at the age of 25 and 35, respectively, and obtain an average LFP rate across those years of 28.5% for women with children under the age of 10.[57] The analogous statistic in our model would be the LFP of women with children under 10 during the first three periods of life.[58] This yields an LFP of 26.7%, i.e. very close to the data. The LFP for women without children under 10 during this same time period is 63.8% in the data versus 65.0% in the model.

Next, it may be useful to report the proportion that attend college by aptitude level and gender (see Table 9). Although we do not have data with which to compare these numbers, it is nonetheless reassuring that the model yields an increasing proportion of each gender going to college as a function of aptitude. While we do not have a formal proof that the welfare gap between college and high school is increasing in aptitude level, a simple two period version of the model yields $\lambda_h - \lambda_l > 0$ as a sufficient condition for a monotonically increasing response in aptitude (θ), which is satisfied in the data.[59]

The quantitative model allows us to evaluate the welfare costs of divorce for women and men by education. Given divorced women retain sole custody of the children, they suffer a large drop in consumption upon the dissolution of marriage. Comparing the average consumption of women who are married in both periods one and two with women who are married only in period one, the average consumption of divorced women in period 2 is 19.0% lower than their married counterpart's for college women and 9.2% smaller for high-school women.

An alternative calculation is the percentage of average consumption a married woman would be willing to sacrifice in period 2 in order to remain in her married state. This is the z_e that solves

$$u(z_e\bar{c}_e^{\mathbf{m}}) - \bar{\psi}_e^{\mathbf{m}} = u(\bar{c}_e^{\mathbf{d}}) - \bar{\psi}_e^{\mathbf{d}}, \tag{26}$$

where $\bar{\psi}_e^s$ denotes the disutility from labor weighted by the proportion of working women of education e in marital status $s \in \{\mathbf{m}, \mathbf{d}\}$.[60] This yields a consumption loss of 24.9% for a college women and 19.1% for her high school counterpart. The smaller gap in consumption loss between education groups comes

[57] The variable which denotes the existence and age of the youngest child in the household does not begin until 1968 in the CPS. For the Census data, we define women with children under 10 as those who report the existence of an own child under that age in the household.

[58] After the first three periods, given the structure of fertility shocks in the model all children would be older than ten.

[59] To see this, consider a two period version of the model in which, in the first period agents choose whether to go to college and in the second the agent consumes her/his earnings. Note that according to our model specification, the wage of an individual of ability θ can be written as $w_e(\theta) = w_e(\theta_0)\exp\left(\lambda_e\left(\theta - \theta_0\right)\right)$ where θ_0 is some given base ability level. Next, let C denote the cost of college. We want to show $\frac{\partial C}{\partial \theta} > 0$. Using CRRA preferences and assuming $w_h\left(\theta\right) - w_l\left(\theta\right) > 0$ (a necessary condition for individuals to attend college if costs are positive) yields $\lambda_h > \lambda_l$ as a sufficient condition.

[60] Note that this is also weighted by the proportion of women of have young children at home.

from taking labor disutility into account, which is higher for high-school women, and from the latter's higher marginal utility resulting from lower consumption.

Men, on the other hand, gain from divorce. Comparing the average consumption (by education) of men who are married in both periods one and two with men who are married only in period one, the average consumption of divorced men in period 2 is 22.0% higher than that of their married counterpart's for college men and 21.1% higher for high-school men. Note that the difference in consumption gain for men across the two education levels is significantly smaller than the gap between college and high-school women. This is because men's wives do not significantly participate in the labor market and thus, from most men's perspective, the economic consequences of divorce arise mostly from a decrease in household size.

We use our calibrated model to investigate the quantitative importance of various economic and social changes in determining education and female LFP for the 1955 cohort.

6 From 1935 to 1955

In this section we investigate the consequences of first imposing the 1955 family structure on the 1935 benchmark followed by the 1955 wages structure and lastly, the combination of the two. Next we explore the role of preferences. We conclude with a welfare analysis.

6.1 1955 Family Structure

1955 Divorce Rates

In order to gauge the importance of the greatly increased divorce rate faced by 1955 cohort, we start by examining the impact of these changes on outcomes. To do this, we increase the per period probability of divorce from the levels faced in 1935 (around 5.5% on average) to the levels faced in 1955 which average about 12%. We hold all other parameters constant at their 1935 levels, including the initial proportions by marital status, the per period marriage and remarriage rates, and the conditional probabilities of marrying by education in the marriage market.

As one can see from Figure 12, married women's LFP significantly increases as a response to the heightened risk of divorce.[61] This increase is concentrated mainly during the childbearing ages for women of both education levels. The magnitude of the increase for college women during their first

[61] There is also an increase in the LFP of divorced women (not shown).

two periods of work is of around 25 percentage points while that of high school women is around 20 percentage points.

The higher divorce probabilities give married women greater incentives to work in order to mitigate the negative economic consequences of divorce. Women benefit by working more both because the increased participation when young increases wages through returns to experience and because the higher household income translates into a higher asset level in the married household, increasing the amount that the wife will retain in the case of divorce.

The increase in divorce probabilities also affected the proportion of men and women who choose higher education and, most importantly, this effect is asymmetric across genders: there is a 2 percentage point increase in the proportion of women who choose to enroll in higher education (29.1% in the benchmark vs 31.1%) versus an almost 3 percentage point drop in the proportion of men who choose to do the same (39.2% vs. 36.3% in this experiment). This confirms our earlier speculation that increases in the divorce rate worked to decrease the education gap between men and women in a quantitatively important fashion. In particular, had no other forces been at work, the 1955 divorce rate would have closed almost half (49.2%) of the initial education gender gap.

One can shed light on this asymmetric reaction by noting that a higher divorce risk increases the incentive for women to work and thus increases a woman's return to a college education in the form of higher wages. In the case of men, on the other hand, high-school men benefit more from divorce than their college counterparts because they have higher divorce rates and lower remarriage rates. Furthermore, the wives of high-school men increased their LFP in reaction to higher divorce rates, by more than the wives of college men.

1955 Marriage Market

Next, we proceed to evaluate the effects from the remaining changes in the marriage market which occurred between the 1935 and the 1955 cohorts. Using the 1935 benchmark model, we now change not only the divorce probabilities but also the conditional probabilities of marrying a spouse of a given education level, the marriage and remarriage probabilities, and the initial distribution of marital states all so as to match the 1955 marriage market.

The changes in female LFP across education and marital status are very similar to the ones obtained when we changed only the divorce probabilities (shown in Figure 13). The education choices, however, react significantly more, becoming positive for men as well. The proportion of men with college goes from 39.2% in the benchmark to 42.5%; for women, the increase is from 29.1% to 34.8%. This is mostly a result of the increase in the proportion of individuals who emerge from period zero in a one-person household – either single or divorced. These marital states have low consumption and low assets make borrowing less attractive. Going to college increases earnings which helps to mitigate the adverse consequences of this change in initial proportions.

1955 Fertility, Childcare, and Marriage Market

Finally, in addition to the changes in the marriage market, we allow (i) fertility patterns to change, (ii) childcare costs to decrease by 20%, and (iii) both the preceding changes to happen simultaneously. We are interested in the effect of childcare costs since Attanasio, Low, and Sanchez-Marcos (2008) conclude that 50% of the increase in married women's LFP between the cohort born in 1945 and the one born in 1955 is due to the fall in childcare of 15%. We investigate a 20% decline in childcare costs in order to account for the fact that our cohorts are further apart.

We find that the further changes in education beyond those already attributed to the 1955 marriage market are relatively small for both experiments (i) and (ii). A larger than linear effect, however, is achieved when both fertility and childcare are changed simultaneously. In that case, although all women (and men) gain from these changes, high-school women gain by more and the final proportion of women who choose to go to college decreases from the 34.8% achieved in the previous experiment to 32.1%. For men, on the other hand, the combined effect of lower fertility and childcare is to increase the relative attractiveness of college mostly because their (college) wives increase their LFP (see Figure 13); 43.9% of them choose to enroll in college.

The changes in LFP resulting from the lower fertility and childcare costs are shown in Figure 13 (the lines with the caption "MM + children"). As seen in the figure, these changes increase the LFP of married college women in the first period of life, almost to parity with the value seen in the 1955 data, and have a small positive effect for them thereafter. They also decrease somewhat the LFP of married high-school women when they are younger but increase it thereafter.

To summarize, changes in the marriage market, fertility and childcare costs account for a significant fraction of the change in LFP across the two cohorts of married women, especially for those with a college education. For example, these changes account for 64.6% of the average LFP gap across the two cohorts of college married women from age 25 to 50 and 37.2% of the average LFP gap for the equivalent high-school women. Interestingly, they generate too much of an increase in the LFP of divorced women for both education groups (see Figure 14). With respect to education, these changes can explain almost all of the increase in the proportion of men who choose college whereas for women it can account for 20.6% of the gap between 1935 cohort and 1955 cohort.

6.2 1955 Wage Structure

We next turn to quantifying the contribution of changes in the wage structure. As is commonly recognized, two key changes in wages that took place in the second half of the 20th century were: (i) the skill premium increased for both men and women, and (ii) the ratio of female to male wages also increased. Quantifying the effect of these changes, given the endogeneity of the ability distribution in each education category and women's selection into work, requires changing certain parameters which govern wages, in particular the $\tau_{eg,t}$ which discipline the time path of wages. Recall that the first differences of the sequence of $\{\tau_{em,t}\}_{t=2,\ldots,t^R}$ for men were calculated directly from wage data for each cohort. The ones calculated for the 1955 cohort (see Figure 8) will now be used for this experiment. That leaves us with $\tau_{em,1}$ and $\{\tau_{ef,t}\}_{t=1,\ldots,t^R}$ which we calibrate internally in order to match certain wage statistics as described below. The returns to ability, the parameters for the stochastic process, and the returns to experience for men, on the other hand are left at their 1935 values.[62]

1955 Skill Premia and Gender Gap

We now proceed to investigate the effects of the 1955 wage gender ratios and skill premia for both genders.. To do this we use the sequences of $\{\tau_{ef,t}\}_{t=1,\ldots,T}$ for women and $\tau_{em,1}$ for men to match the period-by-period gender wage ratio and skill premia for women, the average lifetime skill premium for men, while keeping average male wages at their 1935 levels.[63]

As shown in Figure 15 in the graph with the caption "Skill Premium + Gender Gap", these changes in wage structure result in significant LFP changes for college women (almost a 10 percentage point average increase across the lifecycle) and a smaller increase for high-school women (6 percentage points). College becomes significantly more attractive both for women and men, with 35.6% of women choosing it compared to 47.47% of men. Both men and women respond to the increased skill premia by

[62] Recall that we estimated the returns to ability and the parameters for the stochastic process using NLSY79 data because we do not have specific data pertaining to each of our two cohorts. We maintain the returns to experience for men unchanged as these were estimated from the data as pure age/experience effects, after controlling for year effects.

[63] Note that in the data, average real lifetime wages for men are unchanged between these two cohorts. However, an increase in the gender wage ratio introduces a level effect. Given that we have a non-homothetic model, it is not clear that one should eliminate all level effects as this would imply that the disutility of labor is proportional to income.

becoming more educated, but men respond more strongly largely as a result of their greater labor market participation which increases the returns from college.

1955 Returns to Experience

Given that the literature in this field has found that there has been an increase in the returns to experience for women and that it matters for their LFP (see footnote 42 for a discussion), an additional important exercise is to compound the changes in the wage structure (increased skill premia and gender wage ratio) with a higher returns to experience. We do this by changing the parameters $\gamma_{ef1}, \gamma_{ef2}$ to 0.1593 and -0.003, respectively, so that an extra year of experience translates to approximately a 3% increase in wages for the 1955 cohort, a reasonable number given Olivetti (2006)'s finding, using the 1970 Census, of a 3-5% return to working full-time for an additional year.[64]

As seen in Figure 15 in the graph labelled "All Wages", the higher returns to experience results in married women working more when young, with a more prolonged effect for high-school women. It has basically no effect on married women's LFP once they are older. The proportions of college men and women also increases minutely (less than 0.2 percentage points).

1955 Wage and Family Structure

Combining all the changes in wage structure from the previous section with the changes in family structure from section 6.1, we obtain the LFP path for married women shown in the graph in Figure 15 labelled "wages + family". As is clear from the figure, married women's LFP is still below the levels in the data for women of both education groups in the 1955 cohort. The changes in family structure, when layered on top of the changes in wage structure, impacted LFP mostly during fertility years for both married high school and college women, with an additional smaller effect later in life for the former. The early life increase in LFP is particularly strong for college women, so much so that LFP during the ages of 25-29 (period 1) implied by the model is basically at the same level as the one seen in the data.

The changes in wage and family structures combined can account for a bit over three quarters of the average LFP gap between the two cohorts during the ages of 30-49 for married women of both education groups. The effect of these changes for divorced women is shown in Figure 16. As shown in the figure, the predicted change is so large that both education groups, but especially college women, work

[64] In the robustness checks we experimented with 5% returns to experience with very similar results.

substantially more than in the data.

It is important to note that the combination of family and wage changes are not able to explain the education gap. Although the proportion of women in college increases by a large amount – from 29.1% to 38.9% – so that they are much closer to the 1955 data point of 43.7, men's college proportion also increases by 10 percentage points (from 39.2% to 49.2%) leading a 5 percentage point overprediction with respect to the data and to the same 10 percentage point education gender gap that existed earlier. Thus, we next turn to examining the role of education costs and preferences in explaining the remaining LFP gap and the failure to account for the education gender gap.

6.3 Preferences

It is very difficult to deny that social attitudes towards women changed dramatically over the life spans of the two cohorts. Rather than treat preference changes as a residual, we investigate whether a simple change in work preferences is able to account for the remaining LFP. Prior to this, we investigate what changes in education costs are needed to match the education outcomes in the data.

The Psychic Cost of College

In order to match the education proportions seen in the data for the 1955 cohort we need to change the cost of college for both genders. As explained above, the combination of all the 1955 changes generated too large a proportion of college men and too small a proportion of college women. We can now recalibrate the model to include not only the changes in family and wage structure but also any change in education costs required to generate the college proportions seen in the data.

Performing the exercise above yields new means for the distributions of psychic costs of $\mu_f = 0.5217$ and $\mu_m = 1.1642$. This translates into a decrease in female average education costs of 26.6% whereas male education costs increase by a similar proportion (28.9%). The predicted LFP rates for women of every education and marital status are very similar to the ones obtained without the change in education cost and hence are not shown here.[65]

How can one interpret the fact that the model requires an asymmetric reaction by gender on education costs? One way to think about the education cost reactions is that they combine an increase in monetary costs for both genders (as has been well-documented in the literature) with a sufficiently large decrease

[65] It is always true that simply changing the ability composition of women who attend college does not make any appreciable difference to female LFP behavior.

in women's psychic cost of college so as to undo the monetary rise.[66] As argued by Goldin, Katz, and Kuziemko (2006), girls may have a lower nonpecuniary cost of college preparation and attendance than boys. They point to the fact that at the turn of last century females graduated high school at a higher rate than males, mimicking what is found today for college graduation rates. They advance the hypothesis, which accords with our finding of a higher psychic cost for men, that boys may have a lower level of non-cognitive skills (e.g. they higher incidence of behavioral problems) leading them to be less prepared, on average, to attend college.

Work Preferences

It is difficult to ignore the large change in societal preferences that took place in the twenty years separating the two cohorts. There is by now a large literature demonstrating the importance of culture for various outcomes including female LFP.[67] Given that changes in education costs, wage and family structures are not able to fully account for the changes in the female LFP that took place between the the 1935 and 1955 cohorts, we now ask whether changes in preference towards work are able to explain the gaps between predicted LFP and the data.[68]

Unfortunately, there is little quantitative evidence to discipline how preferences should be changed.[69] In light of this we will ask whether a simple proportional decrease of 10% in all married women's work disutilities is able to generate the work behavior observed in the data. In addition to this change, given that the model is significantly overpredicting the LFP of divorced women, we will assume that the preferences of the latter now resemble those of their married counterparts by education. Why should this be the case? One plausible argument is that once divorce became much more common – an experience that any woman might encounter – the typical divorced woman no longer belonged to some selected pool in terms of her preferences and instead had the same average preferences as those prevalent in her education group.

The left panel of Figure 17 plots the model's predictions for the LFP of married women by education

[66] Although the literature has also found that changing family background (i.e. more parents being college educated) has led to changes in the nonpecuniary costs of education, there is no reason for this to affect genders asymmetrically.

[67] See Fernández and Fogli (2009) for the influence of culture on female LFP and Fernández (2010) for an overview of this literature.

[68] An alternative to interpreting the changes in ψ as resulting from preference change is to attribute them to technological change that made it easier for women to work outside the home (e.g. washing machines as in Greenwood, Seshadri, and Yorukoglu (2005) or changes in brain-biased technological change as in Rendall (2010)).

[69] See Fernández (2011) for a calibrated learning model of female LFP that makes use of longitudinal poll evidence to govern how preferences evolved.

and the right panel does the same for the LFP of divorced women. As can be seen in the resulting graphs, this simple change in preferences does a remarkable job in generating LFP series that are very close to the data. We conclude that the combination of changes in economic environment and family structure in conjunction with a very simple change in female work preferences can account for basically all the changes in female LFP across cohorts. To account for the changes in education, however, requires an asymmetric reaction by gender in the cost of education as explained previously.

6.4 Welfare

How did women and men in the 1955 cohort fare relative to their 1935 counterparts? We will answer this question in a few steps. First, we keep education decisions fixed at the levels obtained after changing the economic environment, family structure, and college costs to reflect those in 1955 (i.e. as obtained in section 6.3) but keeping female work preferences unchanged from the 1935 benchmark. Second, we allow preferences to change. Lastly, we do a ful-fledged ex ante welfare analysis.

Given the the distribution of agents into college and high school implied by the 1955 model and after they have sunk their college costs, we can ask whether agents were better off in the 1955 environment or in the 1935 environment. To do this, we find the proportion z by which consumption across all states would have to be scaled such that, conditional on education and gender, the expected utility in $\tau' = 1955$ is the same as $\tau^0 = 1935$, i.e.,

$$\sum_{t=t_0}^{T} \sum_{\theta} \pi_{\tau'}^{eg}(\theta)\, \mathbb{E}V_{egt}\left(z, \theta; \tau'\right) = \sum_{t=t_0}^{T} \sum_{\theta} \pi_{\tau'}^{eg}(\theta)\, \mathbb{E}V_{egt}\left(\theta; \tau^0\right) \tag{27}$$

where $\pi_{\tau'}^{eg}(\theta)$ is the proportion of agents of ability θ and gender g that have education level e in 1955. $\mathbb{E}V_{egt}(z, \theta; \tau')$ is the expected utility in period t of an agent of gender g, education e, and ability θ, given that her/his consumption level in each possible state is multiplied by z and that they face the 1955 environment. A z greater than 1 indicates that agents were, conditional on education and gender, better off in 1935 than in 1955.

Conditional on education, we find a sharp gender divide in how agents fared. Women of both education levels were made worse off from living in the 1955 environment relative to the 1935 one. Men, on the other hand, were made better off across education levels. In consumption equivalent terms z, college women in 1955 would require a 0.5% increase in all consumption levels to leave them indifferent between 1955 and 1935 whereas high school women would require a larger increase of 2.9%. Both college and high school men, however, would need large decreases in consumption levels in order to be indifferent between the two cohort structures. College men in 1955 would require a 10.8% decrease in all levels of consumption whereas high-school men would be willing to suffer a 4.0% decrease. The asymmetry across genders stems from the fact that men were made better off from the change in family structure, particularly the increase in divorce probabilities, whereas women lost from the same change. For both genders, the relative welfare gains from 1955 are greater for college than for high school individuals. This is mainly a result of the increased skill premium for both genders.

Next we can allow female preferences toward work to change as in section 6.3. We add these to the other changes in the economic and family structures. The female work preferences in 1935 are kept at their original values. We find that the change in preferences for work improves women's welfare sufficiently such that college women now prefer the 1955 environment. High school women, however, are still worse off. College women in 1955 would now be required to decrease their consumption across all states by 2.5% in order to be indifferent between 1955 and 1935; high school women, on the other hand, would require a 2.1% increase in order to be indifferent. For men there are additional gains from the change in female work preferences as their wives now work more. College men in 1955 would require a 16.7% decrease in consumption whereas high-school men would need a 5.4% decrease in order to be indifferent.

It is important to understand that these conditional statements do not take into account the welfare gains generated from the large increase in women's educational attainment. In order to do that, we need to perform an ex ante welfare analysis and not condition on education. Hence, we now condition only on gender and find z such that the expected utility of being gender g is the same in 1935 as in 1955, i.e., we solve:

$$\sum_{t=t_0}^{T}\sum_{\theta}\sum_{e}\pi_{\tau'}^{eg}\left(\theta\right)\mathbb{E}V_{egt}\left(z,\theta;\tau'\right)=\sum_{t=t_0}^{T}\sum_{\theta}\sum_{e}\pi_{\tau^0}^{eg}\left(\theta\right)\mathbb{E}V_{egt}\left(\theta;\tau^0\right) \tag{28}$$

Note that, in contrast with equation (27), this expression takes into account the fact that the proportion of individuals in college was different for the 1955 cohort versus the 1935 one, i.e., the right-hand-side of the equation now has $\pi_{\tau^0}^{eg}\left(\theta\right)$ rather than $\pi_{\tau'}^{eg}\left(\theta\right)$.

The ex ante welfare analysis indicates that all agents gained from the 1955 environment (both with and without changes in female work preferences). With changed preferences, women in 1955 would require a 5.3% decrease in consumption across all states in order to be indifferent between the 1955 environment and the 1935 environment; men would require a 14.4% decrease. Not surprisingly given the preceding analysis, the gain for men from 1955 relative to 1935 is considerably larger than that for women.[70]

7 Robustness Checks

In this section we discuss the robustness of our model's main findings to alternative values of parameters whose empirical foundations are less well grounded.

A key parameter that governs the welfare of wives versus husbands is the pareto weight χ used to obtain a solution to the married household's allocation problem. We chose $\chi = 0.3$, implying that the wife had lower bargaining power than her husband. Another plausible choice would have been $\chi = 0.5$, where husband and wife receive equal weight in the household allocation problem. With this in mind, we recalibrate the benchmark 1935 model using $\chi = 0.5$.

While most wages parameters remain similar (albeit slightly higher), all parameters for disutility from labor become smaller than they were in the original benchmark in order to explain why women worked as much as they did. Repeating the experiment of changing only divorce probabilities to those for the 1955 cohort (see beginning of section 6.1), we find equivalent results although the reactions are slightly more muted. As before, the response in LFP to increased divorce remain focused on their first years of working life. While in the original benchmark the increase in LFP during the ages of 25-34 was around 25 percentage points for college and 20 for high school, those same statistics are now around 20 and 15 percentage points, respectively. With respect to education choices, women respond to higher divorce probabilities by going more to college (30.3% vs. 31.1% in the same experiment in the original benchmark) and men choose to go less (35.9% versus 36.3% before).

The recalibrated model also yields the same implications as the original benchmark when we redo the experiment in which we change all wage and family structures (see end of section 6.2), although, once again, the effects are smaller. These changes now account for around 61% of the LFP gap between the two cohorts of married women.

Next we can investigate whether a change in women's bargaining power can help explain the cohort changes. This is motivated by research that argues that, given the changes in female wages, woman's bargaining power within the household has increased over time. For example, Knowles (2007a) finds

[70] If female work preferences are not allowed to change, women in 1955 would require a 3.3% decrease in consumption levels whereas men would require an 11.9% decrease.

that $\chi = 0.34$ in 1970 but due to rise in wages of women this value increased to $\chi = 0.41$ in the 1990s. Since we chose $\chi = 0.3$ for 1955 as well, this possibly underestimates the bargaining power of the wife then. Below we explore the effect of compounding all wage and family structure changes as in section 6.3 with an increased pareto weight of the wife to $\chi = 0.5$.

A feature of any simple model with disutility of labor for women is that an increase in the wife's pareto weight will, ceteris paribus, lower her participation rate. This is indeed what happens in this model. Compared to our specification in which we change wages and family structure, the average lifetime LFP of married college women drops by 10 percentage points from 72.9% to 62.9% and their high school counterparts' participation is 12 percentage points lower (from 60.1% to 48.3%), although their life cycle participation profiles remain for the most part higher than the ones seen for the 1935 cohort. The labor supply of divorced women during their youth is also lower than in the specification with $\chi = 0.3$ as they take into account the possibility of future marriage and decreased participation. With respect to education decisions and comparing with the specification with $\chi = 0.3$, fewer women choose to go to college (34.35%). The proportion of men who choose to go to college is almost unaffected, with a drop of less than 0.5 percentage points.

Additional robustness checks included changing the proportion of a man's income which must be paid in child support and in pension support for ex wives. We recalibrated the model for 5% and 20% of income. The parameter values implied by these alternative proportions are generally similar to the ones implied by the benchmark model and the effects on labor and education decisions across the different experiments were also very similar.

Finally, given that we did not have firm evidence as to the proportion of household assets obtained by a wife upon divorce, we recalibrate the model with asset splits of $\alpha = 0.3$ and $\alpha = 0.7$. Implied labor disutility parameters are generally lower than the benchmark in the former and higher in the latter (although these differences are less than 5%). Implied childcare costs also change in a similar way and by a similar magnitude. The experiments on these recalibrated benchmarks also yielded similar results as with the original benchmark.[71]

[71] All results from the robustness checks are available upon request.

8 Concluding Remarks

This paper developed a dynamic stochastic life-cycle model to evaluate how changes in family structure, economic environment, and cultural norms contributed to changes in the education choices and women's LFP. The model was calibrated to match key statistics of the 1935 cohort. We then proceeded to change characteristics of the environment in order to approach the one faced by the 1955 cohort. The model was successful in predicting increased LFP for women, accounting for over three quarters of the LFP gap of married women between the ages of 30-49. However, the model also implied that preferences towards work and education needed to have evolved in order to fully account for the changes in LFP across the two cohorts and the closing of the education gender gap.

We found both changes in the family and wage structure important in explaining the profile of women's work. In particular, the increased probability of divorce faced by the 1955 cohort is a key driver of the increase in women's work and it produced the desired asymmetric reaction in the education choices of men and women, helping to reduce the education gender gap significantly. Furthermore, changes in divorce probabilities alone account for around 60% of the LFP increase during the ages of 25-40 for married women. The changes in wage structure, however, particularly the increased skill premium for men, undo a large part of the elimination of the education gender gap and thus ultimately required an asymmetric change in education costs to match the data for the 1955 cohort. Thus it is fair to say that both the asymmetric nature of the economic consequences of divorce and the asymmetric change in education costs are required to produce the much more symmetric work and education outcomes across the sexes.

The model also had interesting welfare predictions. Conditional on education level, men greatly benefited from the changing economic environment, whereas both high school and college women lost from those changes. However, once we change female preferences for work, college women also enjoyed small welfare gains in the 1955 world. High-school women on the other hand remained worse off. An ex ante welfare analysis (i.e., unconditional on education) revealed that both women and men fared better in 1955. Men's welfare gains, however, are substantially larger than those of women reflecting, in large part, the asymmetric gender costs of divorce.

This paper points to the importance of changes in marital status as a driver of education and women's LFP. While the paper takes an significant first step, much work remains to be done in order to understand how family structure, education, and work choices interact. In the future, it would be important to endogenize household marriage, divorce decisions, and women's bargaining power. This is required in order to address any phenomenon in which the assumption of a stable unitary household leads one to ignore important consequences stemming from a changing family structure.

References

Albanesi, S., and C. Olivetti (2009a): "Gender Roles and Medical Progress," working paper.

——— (2009b): "Home Production, Market Production and the Gender Wage Gap: Incentives and Expectations," *Review of Economics Dynamics*, 12, 80–107.

Attanasio, O., H. Low, and V. Sanchez-Marcos (2008): "Explaining Changes in Female Labor Supply in a Life-Cycle Model," *The American Economic Review*, 98(4), 1517–1552.

Attanasio, O., and G. Weber (1995): "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," *Journal of political Economy*, 103(6), 1121–1157.

Beller, A. H., and J. W. Graham (1988): "Child Support Payments: Evidence from Repeated Cross Sections," *American Economic Review: Papers and Proceedings*, 78(2), 81–85.

Bernal, R. (2008): "The Effect of Maternal Employment and Child Care on Children's Cognitive Development," *International Economic Review*, 49(4), 1173– 1209.

Bernal, R., and M. P. Keane (2009): "Child Care Choices and Children's Cognitive Achievement: The Case of Single Mothers," working paper.

Cameron, S., and J. Heckman (1998): "Life cycle schooling and dynamic selection bias: models and evidence for five cohorts of American males," *Journal of Political Economy*, 106(2), 262–333.

——— (2001): "The dynamics of educational attainment for black, Hispanic, and white males," *Journal of Political Economy*, 109, 455–499.

Cameron, S., and C. Taber (2004): "Borrowing Constraints and the Returns to Schooling,," *Journal of Political Economy*, 112(1), 132–182.

Chamberlain, G. (1984): "Panel Data," in *Handbook of Econometrics*, ed. by Z. Griliches, and M. Intriligator, vol. 2, chap. 22. North Holland.

Chiappori, P.-A. (1988): "Rational Household Labor Supply," *Econometrica*, 56(1), 63–90.

Chiappori, P.-A., M. Iyigun, and Y. Weiss (2009): "Investment in Schooling and the Marriage Market," *American Economic Review*, 99(5), 1689–1713.

Costa, D. (2000): "From Mill Town to Board Room: The Rise of Women's Paid Labor," *Journal of Economic Perspectives*, 14(4), 101–122.

Del Boca, D., and C. Flinn (1995): "Rationalizing Child-Support Decisions," *American Economic Review*, 95(5), 1241–1262.

——— (2010): "Endogeneous Household Interaction," working paper.

Eckstein, Z., and O. Lifshitz (2011): "Dynamic Female Labor Supply," *Econometrica*, forthcoming.

Eckstein, Z., and K. Wolpin (1999): "Why Youths Drop Out of High School: The Impact of Preferences, Opportunities and Abilities," *Econometrica*, 67(6), 1295–1339.

Fernández, R. (2010): "Does Culture Matter?," in *Handbook of Social Economics*, ed. by A. Bisin, J. Benhabib, and M. Jackson. Elsevier.

——— (2011): "Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century," working paper.

Fernández, R., and A. Fogli (2009): "Culture: An Empirical Inves- tigation of Beliefs, Work, and Fertility," *American Economic Journal: Macroeconomics*, 1, 146–77.

Gallipoli, G., C. Meghir, and G. L. Violante (2010): "Equilibrium Effects of Education Policies: A Quantitative Evaluation," working paper.

Galor, O., and D. Weil (1996): "The Gender Gap, Fertility, and Growth," *American Economic Review*, 86(3), 374–387.

Gayle, G.-L., and L. Golan (2010): "Estimating a Dynamic Adverse Selection Model: Labor Force Experience and the Changing Gender Earnings Gap," working paper.

Ge, S. (2010): "Women's College Decisions: How Much Does Marriage Matter?," *Economic Inquiry*, forthcoming.

Goldin, C. (1990): *Understanding the Gender Gap: An Economic History of American Women*. New York: Oxford University Press.

Goldin, C., and L. F. Katz (2002): "The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions," *Journal of Political Economy*, 110(4), 730–770.

Goldin, C., L. F. 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.

Greenwood, J., A. Seshadri, and M. Yorukoglu (2005): “Engines of Liberation,” *Review of Economic Studies*, 72(1), 109–133.

Guvenen, F. (2007): “Learning Your Earning: Are Labor Income Shocks Really Very Persistent?,” *The American Economic Review*, 97(3), 687–712.

Heathcote, J., K. Storesletten, and G. L. Violante (2010): “The Macroeconomic Implications of Rising Wage Inequality in the United States,” *Journal of Political Economy*, 118(4), 681–722.

Jones, L., R. Manuelli, and E. McGrattan (2003): “Why are Married Women Working so much?,” working paper.

Keane, M. P., and K. Wolpin (1997): “The Career Decisions of Young Men,” *Journal of Political Economy*, 105(3), 473–522.

——— (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.

Kneip, T., and G. Bauer (2009): “Did Unilateral Divorce Laws Raise Divorce Rates in Western Europe?,” *Journal of Marriage and Family*, 71, 592–607.

Knowles, J. (2007a): “High-Powered Jobs: Can Contraception Technology Explain Trends in Women’s Occupational Choice?,” working paper.

——— (2007b): “Why Are Married Men Working So Much? Home Production, Household Bargaining and Per-Capita Hours,” IZA Discussion Papers 2909, Institute for the Study of Labor (IZA).

Krusell, P., and A. A. Smith (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106(5), 867–896.

Oldham, J. T. (2008): “Changes in the Economic Consequences of Divorces, 1958-2008,” *Family Law Quarterly*, 42(3), 419–448.

Olivetti, C. (2006): "Changes in Women's Hours of Work: The Role of Returns to Experience," *Review of Economics Dynamics*, 9(4), 557–587.

Peters, E. (1986): "Marriage and Divorce: Informational Constraints and Private Contracting," *American Economic Review*, 76(3), 437–454.

Potamites, E. (2007): "Why Do Black Women Work More? A Comparison of White and Black Married Women's Labor Supply," working paper.

Rendall, M. (2010): "Brain versus Brawn: The Realization of Women's Comparative Advantage," working paper.

Storesletten, K., C. I. Telmer, and A. Yaron (2004): "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy*, 112(3), 695–717.

Strauss, H., and C. de la Maisonneuve (2009): "The Wage Premium on Tertiary Education: New Estimates for 21 OECD Countries," *OECD Journal: Economic Studies*, 2009.

van der Klaauw, W. (1996): "Female Labour Supply and Marital Status Decisions: A Life-Cycle Model," *Review of Economic studies*, 63(2).

Voena, A. (2010): "Yours, Mine and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?," mimeo.

Figures and Tables

Figure 1: Proportions of men and women with least some college education, by cohort

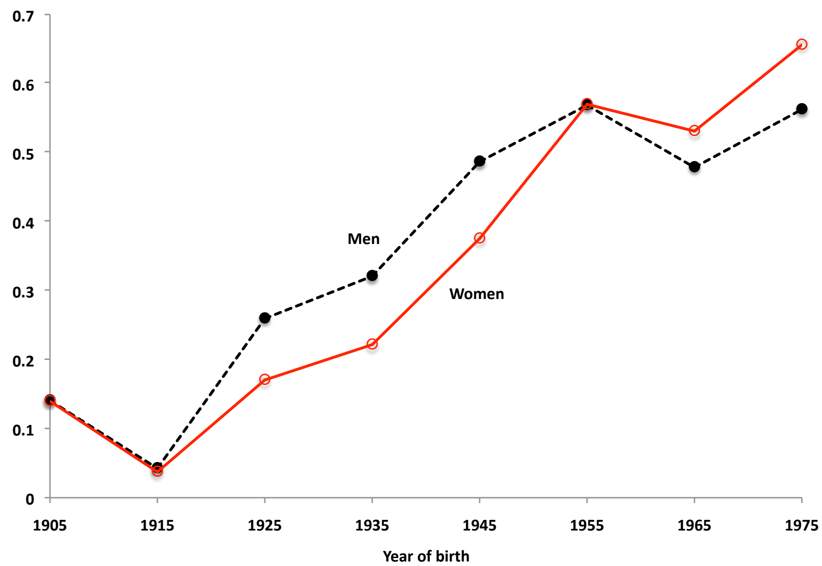

Note: Census data from 1940 to 2000 and American Community Survey data from 2009. Sample includes all white men and women, between the ages of 34 and 36. People with at least some college education are those with at least 1 year of college. Those with high school are defined as people with a high school diploma or no more than 12 years of education.

Figure 2: LFP rates of married white women ages 34-36

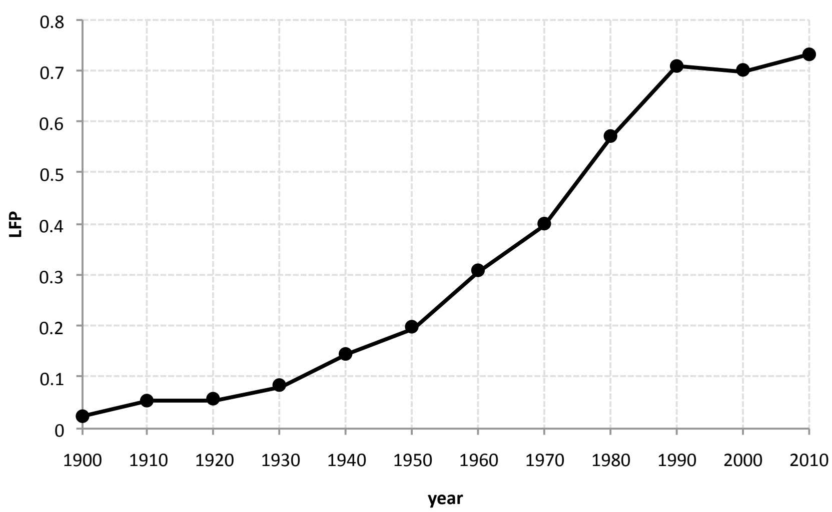

Note: Census data from 1900 to 2000 and American Community Survey data from 2009. Sample includes all married white women, between the ages of 34 and 36, not living in group quarters. The proportion plotted corresponds to those who report being in the labor force.

Figure 3: LFP rates of married white women, by birth cohort and age

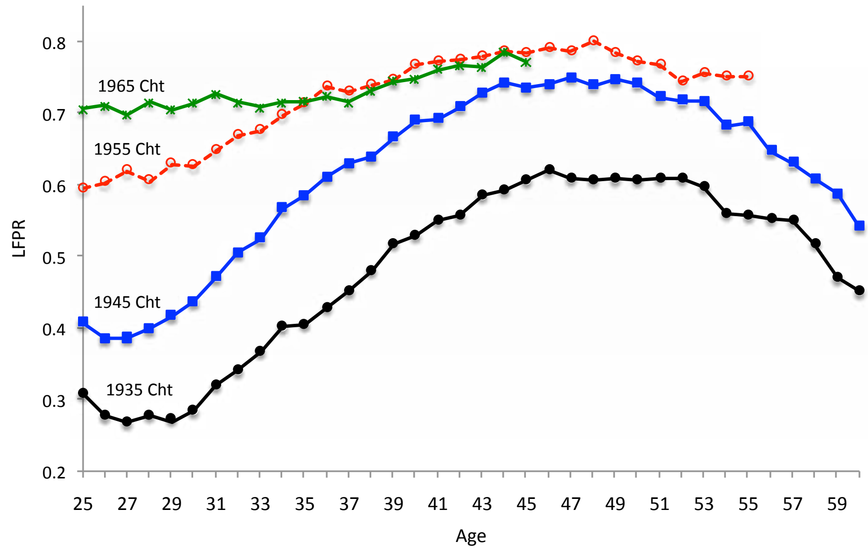

Note: Current Population Survey data from 1962-2008. Each birth cohort is defined as a 3 year interval around its median, e.g. those who belong in the 1965 cohort are all women born between 1963 and 1966. Figure plots the proportion of all married white women by birth cohort, not living in group quarters, who report being in the labor force.

Figure 4: Proportions of marriages ending in divorce, by duration and decade of marriage.

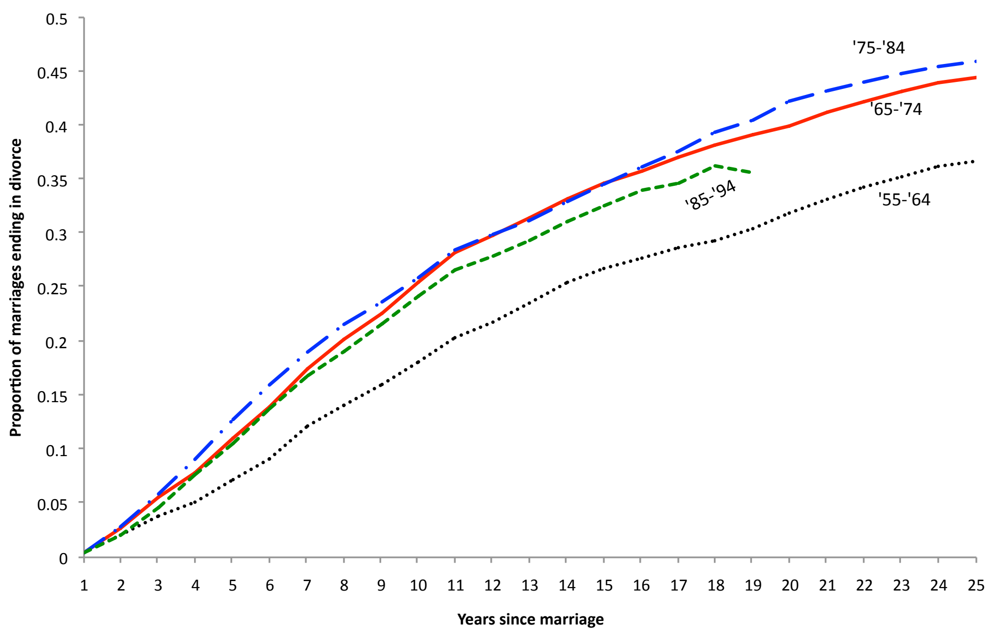

Note: Survey of Income and Program Participation 2004. The figure plots the proportion of marriages (whites only) that end in divorce as a function of the number of years married, for marriages initiated in each given 10 year interval.

Figure 5: LFP of married and divorced women at each age, by education and birth cohort.

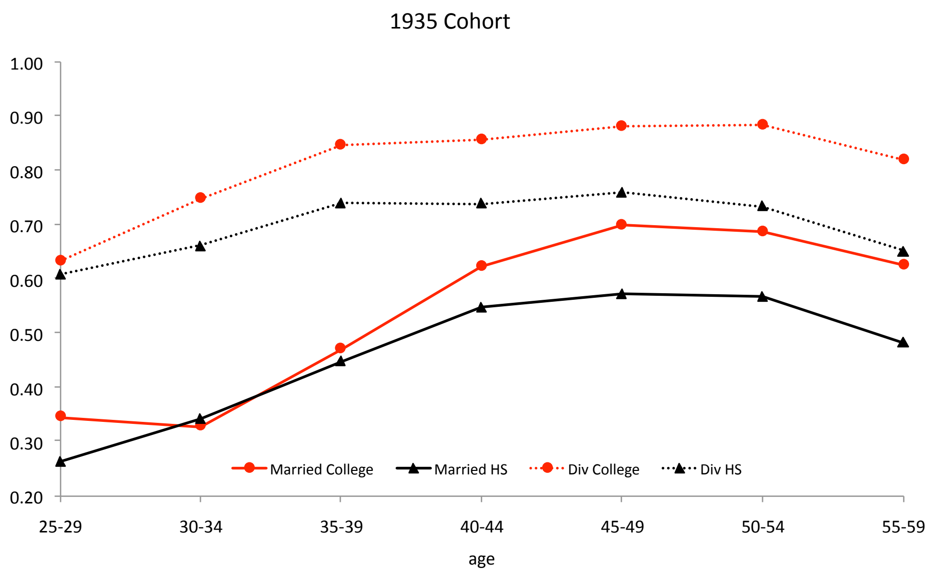

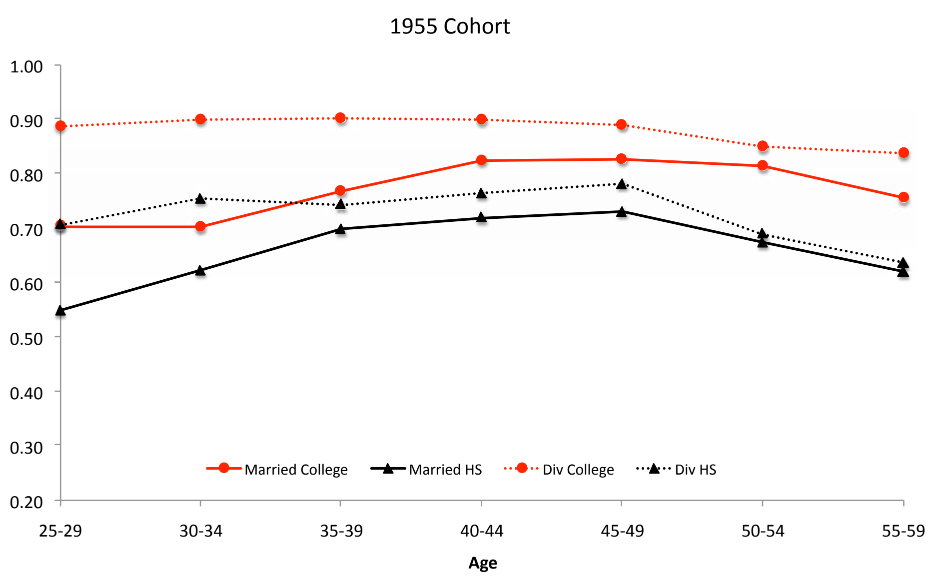

Note: Current Population Survey 1962-2008. LFP for each birth cohort is defined as all white women born in a 3 year interval around the defined year, who report being in the labor force. "Married" is defined as as "married, with spouse present"; "Divorced" is defined as either "divorced" or "separated". See Figure 1 for education level definitions.

Figure 6: Timeline for periods 0 and 1 of agent's life

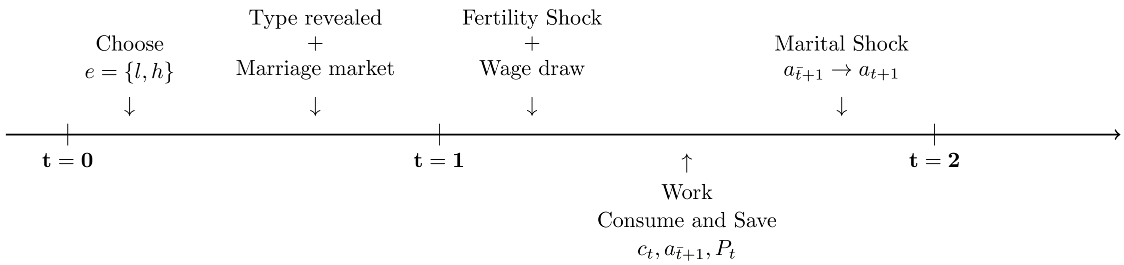

Figure 7: Implied log wage intercepts for men and women, by education.

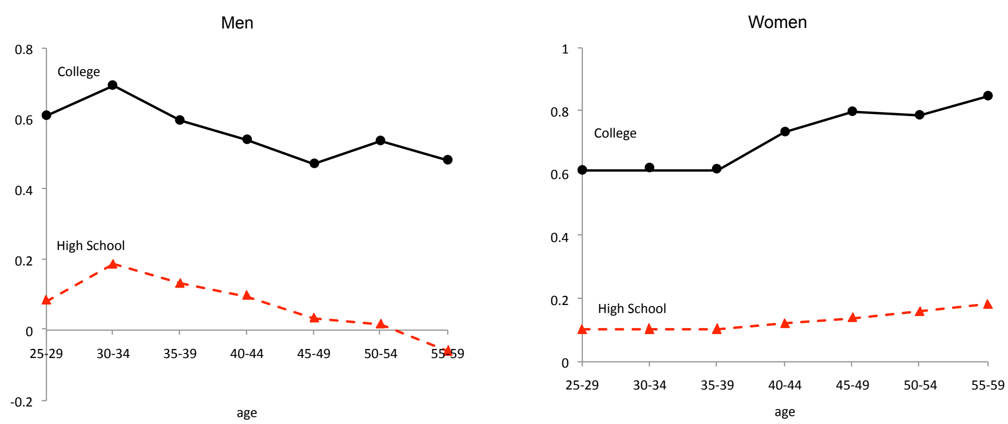

Figure 8: Implied first differences in log wage intercepts for men of the 1955 cohort

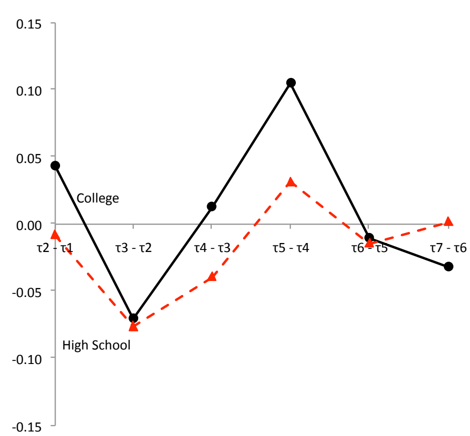

Figure 9: Proportion of each gender-education group in each marital state during the ages of 25-29, by cohort.

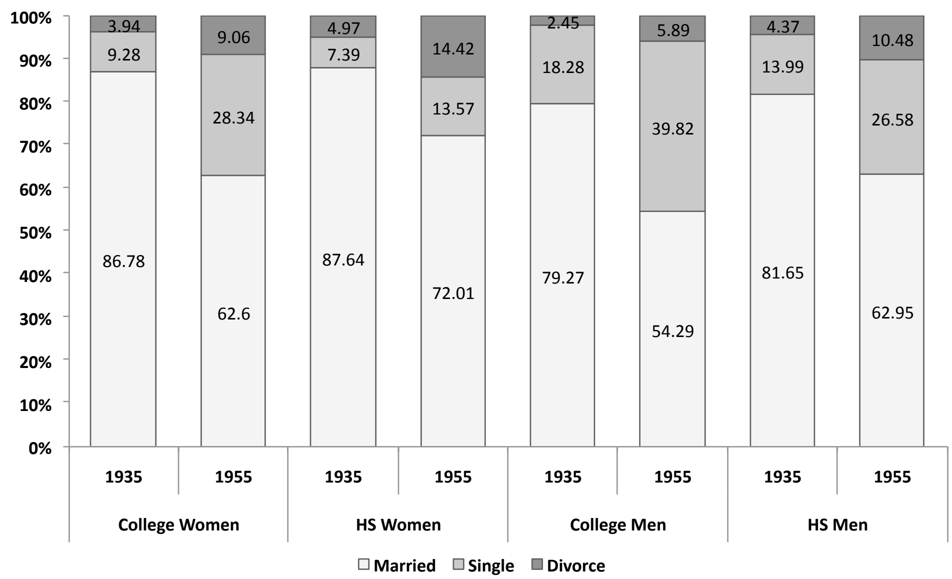

Note: Current Population Survey. Sample consists of all white men and women in each birth cohort. See Figure 5 for definitions of married and divorced; "Single" is defined as "never married". See Figure 1 for education definition.

Figure 10: LFP for married and divorced women by education and age for the 1935 cohort, model vs data.

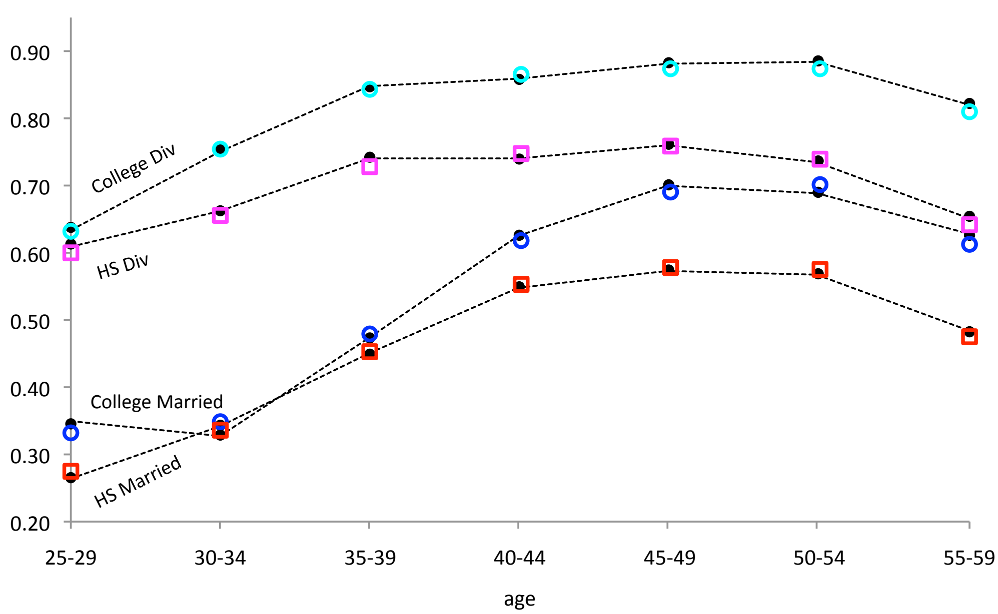

Note: • = data. (□, ◯) = model predictions

Figure 11: Skill premium and the gender wage ratio for the 1935 cohort by age, model vs. data.

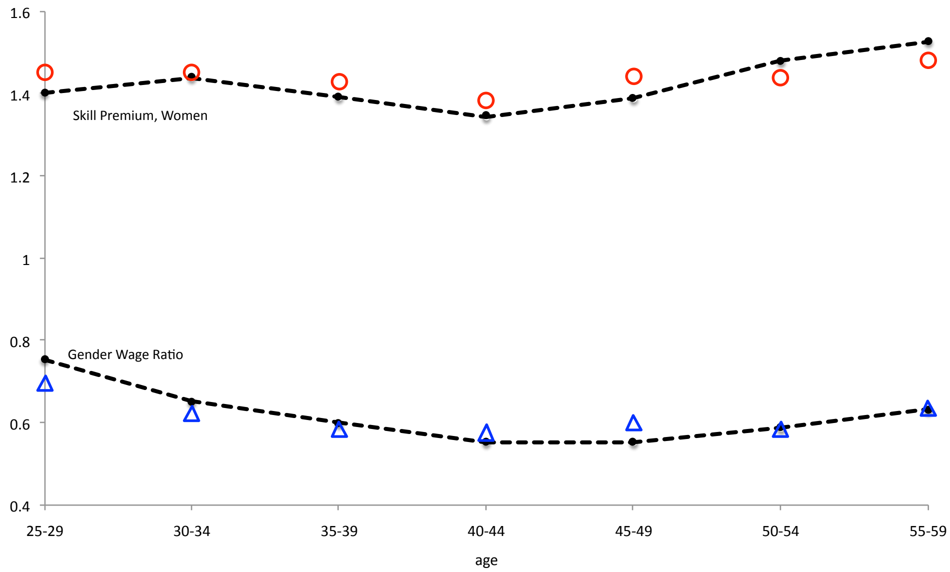

Note: Current Population Survey. • = data. (□, △) = model predictions

Figure 12: Model predictions for married women with 1955 divorce rates

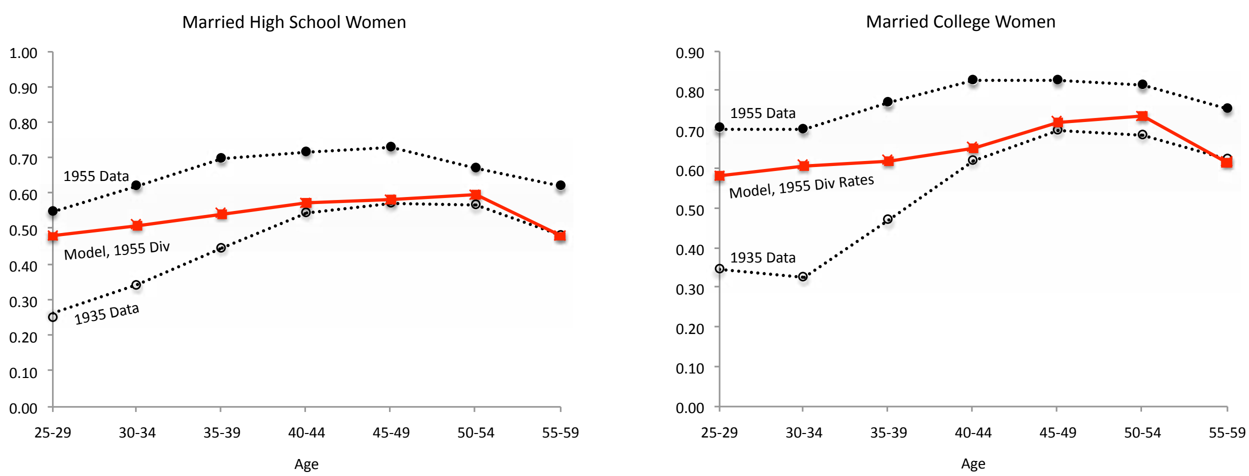

Figure 13: Model predictions for married women with various 1955 family structure changes

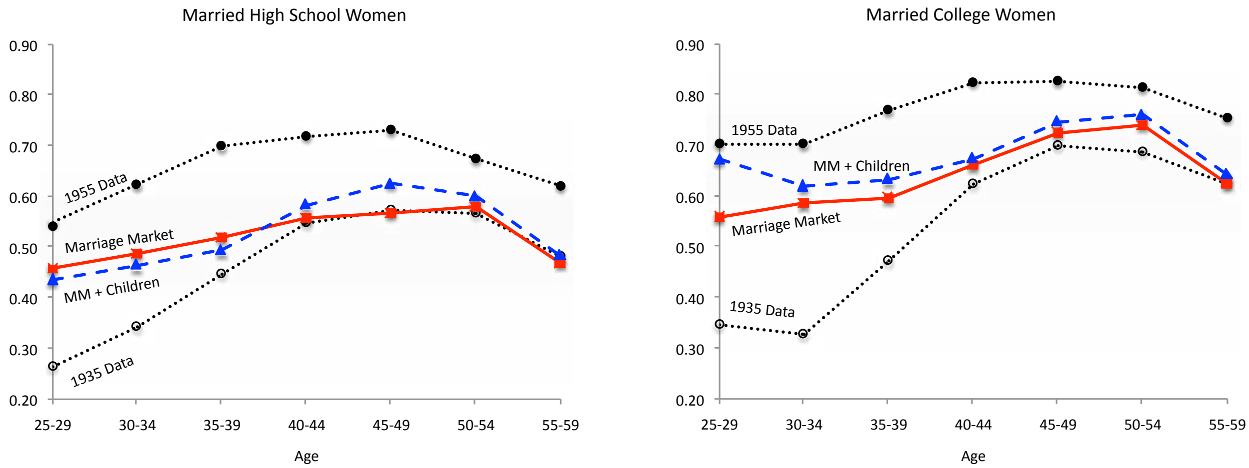

Notes: "Marriage Market" = Benchmark model with changes in the marriage market. "MM+Children" = benchmark model with all changes in family structure

Figure 14: Model predictions for divorced women with various 1955 family structure changes

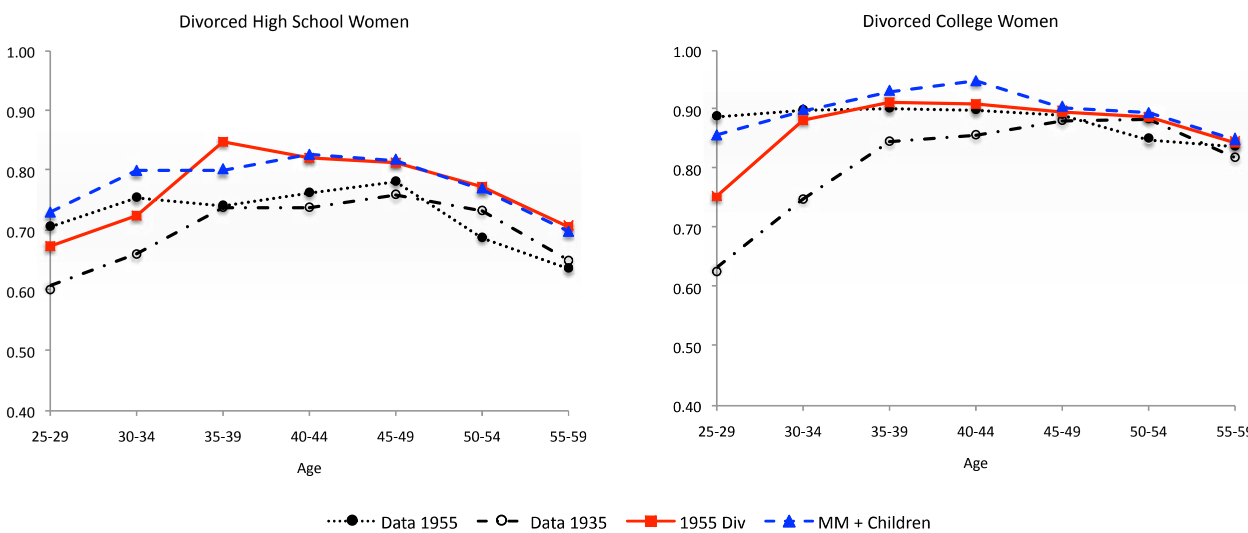

Notes: "1955 Div" = Benchmark model with changes in the divorce probabilities, "MM + Children" = benchmark model with all changes in family structure

Figure 15: Model predictions for married women with various 1955 wage and family structure changes

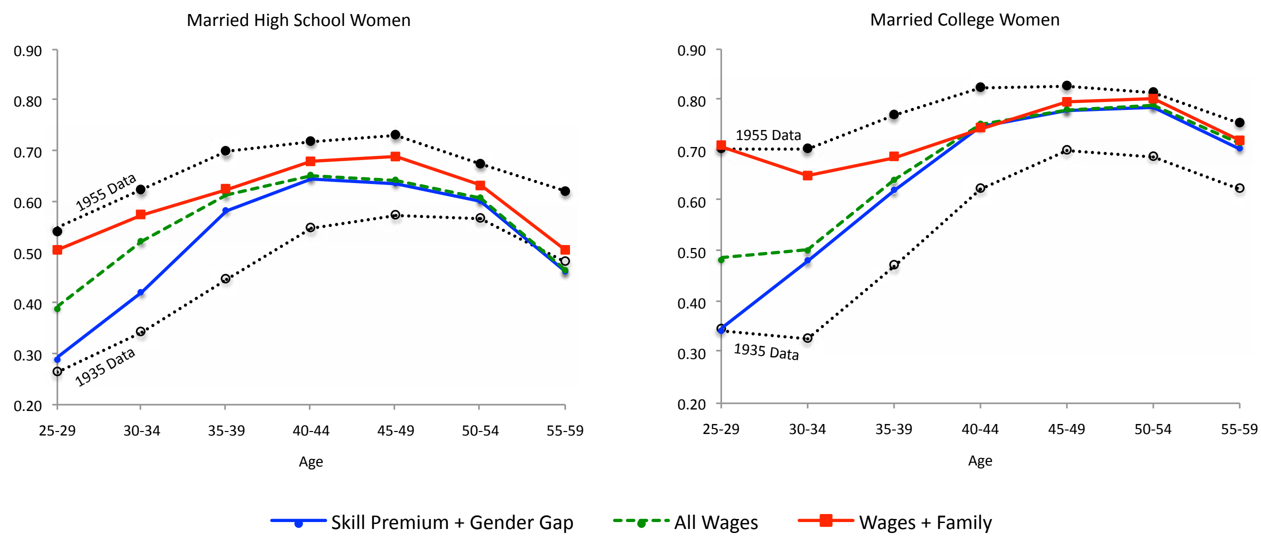

Figure 16: Model predictions for divorced women with all 1955 wage and family structure changes

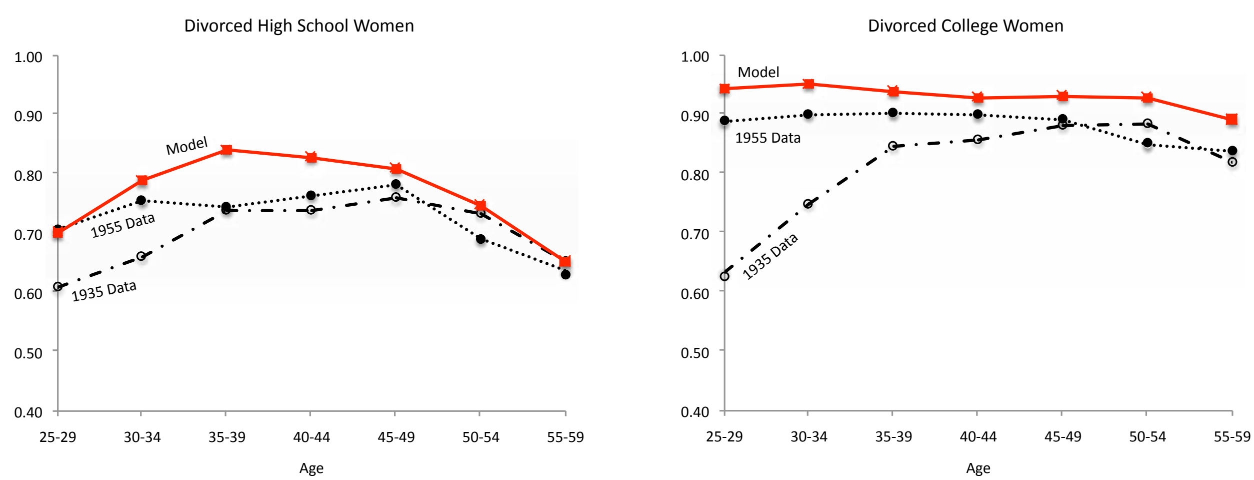

Figure 17: Model predictions with all 1955 wage and family structure and preference changes

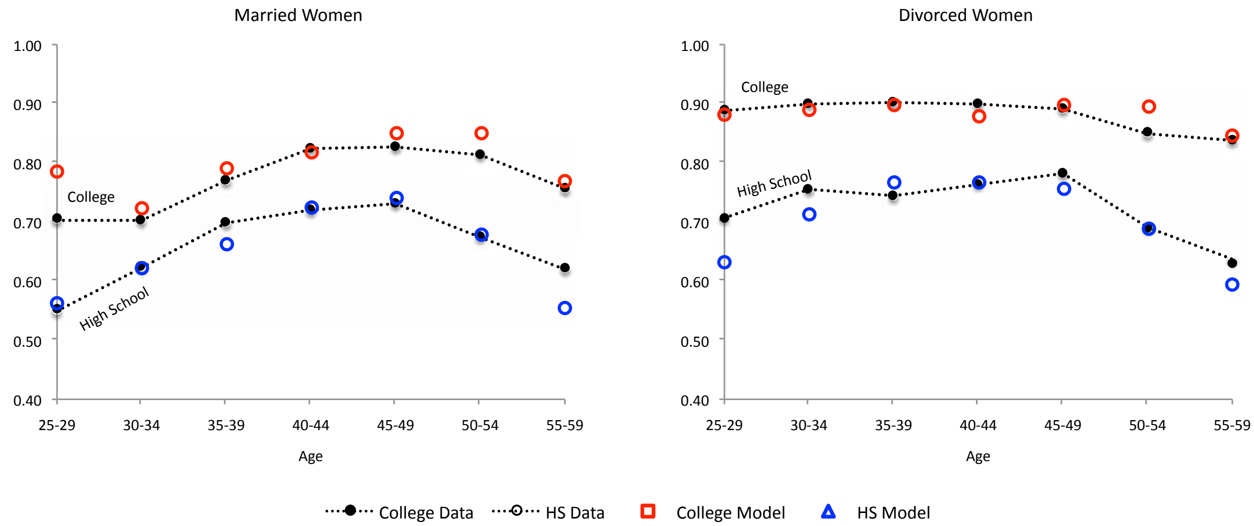

Table 1: Divorce and remarriage rates after 20 years by gender, education, and cohort.

		Women		Men	
		HS	College	HS	College
1935 Cohort	Divorce	20.96	17.11	25.02	18.96
	Remarriage	87.62	86.06	87.61	91.27
1955 Cohort	Divorce	40.41	39.89	42.15	35.42
	Remarriage	86.75	85.46	85.86	92.18

Notes: SIPP 2004. Divorce rates are calculated as the proportion of marriages which end in divorce before the 20th wedding anniversary. Remarriage rates are calculated as the proportion of people who remarry before they reach the 20th anniversary of their last divorce. Birth cohorts are defined as people born in a five year interval around the year of the birth cohort. See Figure 1 for education definitions.

Table 2: Proportions of marriages to college spouse conditional on own gender and education, by cohort.

	College Spouse	
	1935 Cohort	1955 Cohort
College Woman	75.20	78.41
HS Woman	21.39	29.21
College Man	59.72	74.06
HS Man	12.07	24.24

Note: CPS. The proportions of individuals with college spouse are calculated for married people between the ages of 35-39, by gender and education, in each 3 year birth cohort.

Table 3: External Parameters

Parameter		Value
Micro estimates of Intertemporal Elasticity of Substitution	σ	1.5
Discount Factor	β	0.90
Risk Free Interest Rate	R	1.16
Regression log wage on age and age^2, HS men	γ_{lm1}	0.14851
	γ_{lm2}	-0.000457
Regression log wage on age and age^2, College men	γ_{hm1}	0.20334
	γ_{hm2}	-0.000992
Gradient of log AFQT scores on log wage, by educ.	$\gamma_{l\theta}$	0.11157
	$\gamma_{h\theta}$	0.11874
Returns to experience, women	γ_{ef1}	0.1041
	γ_{ef2}	-0.003
Persistence of wage residuals, by educ.	ρ_l	0.90037
	ρ_h	0.86546
Std. Dev. of transitory error of wage residuals,by educ.	σ_{η_l}	0.10787
	σ_{η_h}	0.07653
Std. Dev. of persistent error of wage residuals, by educ.	σ_{ϵ_l}	0.02850
	σ_{ϵ_h}	0.03417
Time varying log wage intercepts, men	$\{\tau_{lm,t}\}_{t=2,t^R}$	see Figure 7
	$\{\tau_{hm,t}\}_{t=2,t^R}$	
Probability of divorce, remarriage and first marriage	$d_{eg}, r_{eg}, \{m_{eg,t}\}_{t=1,\ldots 4}$	see Tables 7 and 6

Table 4: Disutility of labor parameters, $\psi_{em}(k)$

		Married	Married with Child	Divorced	Divorced with Child
Parameter values	High School	0.0853	0.2217	0.0964	0.1193
	College	0.0406	0.0947	0.0450	0.0688
Consumption equivalence	High School	11.2%	33.3%	12.7%	16.1%
	College	5.1%	12.5%	5.7%	8.9%

The consumption equivalence numbers give the fraction of average consumption of a married woman of eduation e that a woman of the same education level would be willing to sacrifice in order to avoid the disutility of labor associated with a particular marital and fertility state. See text for exact calculation.

Table 5: Parameters Calibrated Internally

Parameter		1935 Bench
Disutility of labor for women	$\psi_e^s(k_t)$	see Table 4
Childcare costs for children aged 0-4 and 5-9	κ_{young} κ_{old}	1.1470 0.3922
Initial intercept log wages for men, by education	$\tau_{lm,1}$ $\tau_{hm,1}$	0.0835 0.6098
Time varying log wage intercepts, women	$\{\tau_{lf,t}\}_{t=1,t^R}$ $\{\tau_{hf,t}\}_{t=1,t^R}$	see Figure 7
Wage depreciation from not working	δ	0.0842
Education cost	μ_f μ_m	0.8311 0.9102
Moments	**Data**	**1935 Bench**
LFP of married women, by age and education (14)	see Fig 5	see Fig 10
LFP of divorced women, by age and education (14)	see Fig 5	see Fig 10
Proportion of women who go to college	0.2905	0.2905
Proportion of men who go to college	0.3924	0.3922
Skill Premium by age, women (7)	see Fig 11	see Fig 11
Ratio of male to female wages by age (7)	see Fig 11	see Fig 11
Skill Premium (lifetime), men	1.43	1.43

Table 6: Transition probabilities from single to first marriage

		Men		Women	
	age	College	High School	College	High School
1935 Cohort	30	51.24	44.19	18.14	41.01
	35	26.06	18.90	19.52	03.84
	40	22.65	15.91	23.64	05.63
	45	17.32	6.18	1.63	1.00
1955 Cohort	30	41.89	32.27	41.26	21.84
	35	33.17	11.42	24.39	14.99
	40	21.78	21.28	18.91	15.52
	45	24.68	9.48	14.35	2.28

Note: CPS 1962-2008. Probabilities are calculated using the evolution of the proportion of people who are "never married" between the age shown and 5 years before. These probabilities are conditional on being single.

Table 7: Per period divorce and remarriage probabilities

		Women		Men	
		HS	College	HS	College
1935 Cohort	Divorce	5.90	4.84	6.95	5.22
	Remarriage	38.84	38.49	37.23	42.48
1955 Cohort	Divorce	12.25	12.11	12.60	10.46
	Remarriage	36.11	35.38	35.21	40.18

Note: These numbers were derived using 2004 SIPP data as reported in the text and assuming a uniform 5 years probability of divorce. See text for how they were calculated.

Table 8: Proportion of women who are mothers during the ages of 25-29, by cohort, education and marital status

		Single	Married	Divorced
1935 Cohort	High School	8.22	90.81	62.46
	College	2.07	90.98	43.14
1955 Cohort	High School	10.39	81.25	52.37
	College	3.70	59.87	32.59

Note: CPS 1962-2008. Proportion of women between the ages of 25-29 who report having at least one own child in the household.

Table 9: Proportion of men and women in the model who choose to go to college, by ability quintile

Quintile	Women	Men
1	4.95	11.27
2	18.04	28.08
3	30.55	42.79
4	41.26	52.02
5	49.98	61.22

A Appendix: Data

A.1 Wages

Hourly wages are computed from the CPS using the individuals' reported labor income and hours and weeks worked last year. Prior to 1977, for hours per week, we use the variable which reports the hours worked in the previous week. Also for that period, the weeks worked in the previous year are reported in intervals; we use the midpoint of the interval. From 1977 onwards, we use the variables for usual hours worked per week (last year) and continuous variable for number of weeks worked last year. The lifetime averages reported are averaged over year-by-year averages. Sample weights are used throughout (PERWT). Concerning top-coded observations, we follow the procedure in Katz and Autor (1999). We mutiply all top-coded observations until 1996 by 1.5. After 1996, top-coded observations in the CPS correspond to the average value of all top-coded observations, thus we do not impose further treatment.

We compute the gender wage ratio as the ratio of the average wage of women versus men. The skill premium is computed analogously, using the average wage of college versus high school.

A.2 Income Process

$$\ln y_{eg}(\theta, x_t, z_t, P_{t-1}) = \tau_{eg,t} + \gamma_{eg1} x_t + \gamma_{eg2} x_t^2 + \lambda_e \ln\theta - \delta(1 - P_{t-1}) + w_{et}$$

$$\begin{aligned} w_{et} &= z_{et} + \eta_{et}, \qquad \eta \sim N(0, \sigma^2_{\eta,e}) \\ z_{et} &= \rho_e z_{e,t-1} + \epsilon_{et}, \qquad \epsilon \sim N(0, \sigma^2_{\epsilon,e}) \end{aligned}$$

Age Profiles: The coefficients on age earning profiles for men, $\Gamma_1^{em}, \Gamma_2^{em}$ are estimated using the pooled sample of PSID for the years 1968-2007, restricted to white males who are heads of households. We regress the log earnings measure on age and year dummies, separately by education. The variable for earnings we use is total labor income of the head and we deflate it to 1992 dollars using CPI for all urban consumers. We use the variable on highest grade completed in order to assign an education level to each individual and estimate the parameters separately for people with high school or less ($e = l$) and people with more than high school ($e = h$).

We exclude individuals in the the Latino, SEO and immigrant samples. We also drop observations from people younger than 25 and people older than 65 years old and those who report being self-employed. We choose only individuals with at least 8 (not necessarily consecutive) observations. Furthermore, we drop individuals with missing, top-coded and zero earnings those with zero, missing or more than 5840 annual work hours. Individuals with changes in log earnings greater than 4 or less than -2 are also eliminated from the sample. This leaves us with 2652 individuals in the "low" education group and 1903 in the "high" education group.

Ability Level Gradients: We do not observe individuals in the NLSY79 for a long enough stretch of their lifetime in order to get a precise estimate for age polynomials. Following Gallipoli, Meghir and Violante (2010), we instead use the age polynomials calculated using the PSID and generate log wage measures from the NLSY79 which are "age-free". We use an individual's AFQT89 score as a measure for θ. In order to estimate the ability gradient r_e we regress an the log wage on the log of the individual's AFQT89 score, under the assumption that the error term is uncorrelated with this score. r_e is estimated separately for each education group $e = \{l, h\}$. We estimate these gradients on the cross-sectional sample of the NLSY79, restricted to white males only.

We construct an hourly wage measure by dividing total earnings by total hours worked the previous calendar year, deflated by CPI-U into 1992 dollars. We use only people who have finished their education and we exclude people who change their highest grade completed after the age of 25. Using the highest grade completed variable, we divide individuals into the two education groups: people with 12 years of schooling or less ($e = l$) and people with more than 12 years of completed schooling($e = h$).

Individuals with missing or top-coded earnings in at least one year are dropped from the sample; those people who report being unemployed, out of the labor force or in the military are also dropped. We keep only observations with positive earnings, and we further drop observations in which the individual worked less than 400 or more than 5840 hours. We further drop individuals who report real wages below \$1 and above \$400, in 1992 dollars, and those who have changes in log earnings greater than 4 and less than -2. We are left with 1852 individuals in the low education group and 449 in the high education group.

Estimates of the Labor Shock Processes: Given the residuals from the regression of log wages on ability gradients using the NLSY79 data, we estimate the parameters for the persistent and transitory shocks using the Minimum Distance Estimator (Chamberlain (1984))[72]. The methods of estimating this

[72]We are grateful to Gianluca Violante for providing us with the code used.

process are standard in the literature (see e.g. Heathcote, Storesletten and Violante (2004) for a detailed description of the method). Note that we allow for time-varying $\sigma^2_{\epsilon,e}$ and $\sigma^2_{\eta,e}$ during the estimation process. In the model, we use as inputs the average value across the sample and this is what we report.

A.3 Consumption Deflator:

We use an altered McClements scale ($\mathbf{e}(k_t; s)$) in order to deflate household consumption. Table 10 reproduces the original McClements scale in normalized for one adult.

Table 10: McClements Scale

			+1 child, by age:						
1 adult	2 adults	+ 1 adult	0 - 1	2- 4	5-7	8-10	11-12	13-15	16-18
1	1.64	+0.75	+ 0.148	+0.295	+ 0.344	+0.377	+0.41	+0.443	+0.59

Since we have 5 year periods, and our children are aged 0-4, 5-9, 10-14, 15-20, we weigh the scale accordingly. For example, a child aged 0-4 will add: $0.4(0.148) + 0.6(0.295) = 0.2362$.

The scale $\mathbf{e}(k_t; s)$ is constructed using the number of adults in the household (1 if $s = S, D$ and 2 if $s = M$) and the number of children and their respective ages (k_t).

A.4 Pensions:

To compute retirement benefits for a model household, we modify the approach used in Heathcote, Storesletten and Violante (2010) in order to avoid keeping track of an individual's average earnings over their lifecycle. More specifically, we take each individual's last observed earnings y_T and compute social security benefits as follows: 90% of y_T up to a first threshold equal to $0.38\bar{y}_T$ where $\bar{y}_T$ is the average observed earnings in the economy, plus 32% of y_T from this bendpoint to a higher bendpoint equal to $1.59\bar{y}_T$, plus 15 percent of the remaining y_T exceeding this last bendpoint. For married households, this process is done for both the husband and the wife and the household total benefits are either sum of their benefits or 1.5 times the husband's benefits, whichever one is highest.

B Appendix: Model Solution

In our model, households have a finite horizon, so the dynamic problem is solved numerically by backwards recursion from the last period of life. At each age, the households solve for their consumption savings rule and participation decisions taking as given their state variables that period and next period's value function. Handling more than one continuous state variable is possible but computationally costly. In addition to assets, our model has four other potential continuous state variables: the persistent component of earnings of the husband, w_t^{em} and of the wife, w_t^{ef} and the ability draws of the husband θ^m and of the wife θ^f. We discretize all of these four variables, leaving assets as our only continuous state. We choose 10 nodes whose location are age-dependent for each of the earnings components and use the five quintiles of ability distribution outlined in the main text.

During the working stages of the lifecycle, our model combines a discrete decision (whether the woman participates in the market) and a continuous decision (the amount of savings). This combination may lead to non-concavities in the value function. Furthermore, the existence of transitions across marital states also requires some attention.

For all periods $t > 4$, since there are no longer any transitions across different marital status, the household problem for the single men and women and divorced men is a trivial consumption savings problem. The problem of the married couple and of the divorced woman becomes a similar problem to the one encountered by Attanasio et al. (2008) which also combined the discrete participation choice of the wife together with the continuous choice of assets. The crux of the maximization problem lies in the fact that one cannot guarantee concavity of the value function even if one controls for the participation decision that period. Given enough uncertainty, however, the conditional value function will be concave. We follow Attanasio et al. (2008) and impose (and check) a unique level of reservation assets a_t^* at which, given all other state variables, the conditional (on participation) value functions intersect only once and thus the woman's participation decision switches at that point from not working to working. We numerically check both the concavity of the conditional value functions and the uniqueness of the reservation asset level.

The problem of the household who enters each period $t \leq 4$ as a married couple needs to take into account the continuation values of the husband and the wife, which could be different. A note must be made about the solution of the problem for the divorced and single (men and women) for all periods $t \leq 4$. Given our assumption that each individual has perfect information about their spouses' contemporaneous state variables, she/he will solve her/his own optimization problem taken as given their potential/ex spouses' asset allocation for next period as given. Then, the optimization problems are solved for each potential spouse separately and the final optimal allocation is given by the solution to the fixed point problem (imposing rational expectations).