

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

October 2011[†]

Abstract

Women born in 1935 went to college significantly less than their male counterparts and married women's labor force participation (LFP) 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. 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 (about 60%) of the observed changes in female LFP. After 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, for the psychic cost of obtaining higher education to change asymmetrically for women versus men.

*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 reduced for the cohort that followed twenty years later. For the cohort born in 1955, the same proportion of white women and men attended college (44%) and women’s labor force participation increased such 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 workplace. 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

¹June Cleaver – a principal character in the sitcom “Leave it to Beaver” – was the archetypal housewife.

²Source: 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 incentivate women more than men.⁶ Divorce is indeed such a force.⁷ The overwhelming tendency for children to live with their mothers and the fact that 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.⁸ While other drivers such as increased skill premia have also operated in many countries, there is no evidence that these are consistently higher for women than for men.⁹ A proper 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 ages of children in the household. These disutilities will be used to match women's

³Figure 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.

⁴This would be less convincing for the cohort born ten years earlier and also had high divorce rates.

⁵Mexico, Turkey and Switzerland are the exceptions.

⁶Indeed, 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."

⁷Divorce has been shown to matter for a variety of outcomes: asset accumulation (Cubeddu and Ríos-Rull (2003) and Voena (2010)), the gender gap in voting preferences (Edlund and Pande (2002)), female suicide and domestic violence (Stevenson and Wolfers (2006)), and fertility timing (Caucutt, Guner, and Knowles (2002)).

⁸For example, Kneip and Bauer (2009) show that the divorce rate quadrupled on average for Western European countries between 1960 and 2000.

⁹Across 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)).

LFP by age, education, and marital status. Although much of our focus is on the LFP of 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 about 60% of the change in LFP observed between our two cohorts of married college women.¹⁰ These changes account for a smaller proportion of the LFP gap for high-school women (48.3% due to wages and 34.2% due to family). We find that a simple and relatively small change in preferences is able to explain the remaining LFP gap.

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.¹¹ 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 (2011) 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 in high school) 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.¹² A recent paper by Eckstein and Lifshitz (2011) provides

¹⁰Note that the effects of changes are not additive.

¹¹Some 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.

¹²The 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

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.¹³ 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 likewise 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, assumes that all women are married 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

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)). Other explanations emphasize changes in wage structure (e.g., Jones, Manuelli, and McGrattan (2003), Gayle and Golan (2010) and Knowles (2007b)) and changes in culture (Fernández (2011)).

¹³This 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 1954-1956.¹⁴ 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”. Finally, since longevity is deterministic in our model we exclude widows from our sample.

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

¹⁴We 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).

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

¹⁶Note 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.

29.1% of women did – a 10 percentage point gap. For the 1955 cohort, the gap disappeared: 44.5% of men went to college compared to 43.7% of the women.¹⁷

2.2 Marriage Market and Fertility

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.

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 0.64.¹⁹ An important indicator of the returns to education in the marital market is the probability of marrying a college spouse, which increased for both men and women of both education levels. These proportions are reported by gender and education, in Table 2.

Fertility also changed markedly during this period. Using Census data, 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 is not very different: it decreased for high school women (age 21) and increased for college women (age 25).

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

¹⁷Education attainment here is measured at age 30.

¹⁸Note 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. We thank Betsey Stevenson for kindly making the SIPP data and the code she used available to us.

¹⁹Assortativeness data is computed from CPS, for people aged 35-39.

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.²⁰

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 Modeling 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 and concave utility (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; 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 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

²⁰Wages are calculated from CPS data. See the Appendix for further details.

utility of divorce and doing this is beyond our computational capacity.²¹

Even in the absence of a general equilibrium marriage market, one could still opt to keep marriage and divorce decisions endogenous. In order to do this one would need to make 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 these distributions for the 1935 cohort, we prefer the transparent and straightforward alternative of modeling marriage and divorce as exogenous shocks and of assuming that the characteristics of agents' (future) spouses are given by a latent marital type.²²

In particular, our simplifying assumption throughout is that agents are of a given type in terms of their potential marital partner, that is revealed to them once they've made their education decision. Prior to making their education choice they know only the conditional probability with which they'll meet different education and ability types, which is a function solely of their gender and education. Afterwards although they continue to face uncertainty about their marital status, they know the characteristics (ability, education, current income and experience) of their *potential* partner. This is their marital "type". This assumption 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.²³

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

²¹Note 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).

²²Treating divorce as exogenous is problematic in terms of understanding causality but we are fortunate that in our time period 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 also contributing to increased divorce risk. See Stevenson (2008) for evidence of a positive relationship between the introduction of unilateral divorce laws and female LFP.

²³An 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.

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 (**m**), divorced (**d**) or single (**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 marital 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:

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

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, $\psi_{eg}^s(k_t)$, 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

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

units of household consumption. Thus $\mathbf{e}(k_t; s)$ 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).²⁴

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, Remarriage, and Children

Women obtain marital-dependent fertility shocks at the beginning of a period. In keeping with modeling 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_t^{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. 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). 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.

²⁴As explained in the parameterization section, we use an altered McClements scale whose exact specification is outlined in the Appendix.

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 $b_g^s(\bar{y})$ 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 $e \in \{l, h\}$, $h > l$.²⁵ 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 of ability θ and disutility ω chooses an education level so as to solve:

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

where $\mathbb{E}V_{eg, \bar{t}_0}^s(\theta)$ is the expected value of the agent's welfare at the beginning of period zero (denoted $\bar{t}_0$) given ability level θ , sex g , and education level e .²⁶ Expectations here are taken both with respect to 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 (**m**), divorced (**d**), or still single (**s**), i.e.,

²⁵In the data l will be high-school graduates and below and h will be individuals who have more education than a high-school diploma.

²⁶Henceforth, we use $\bar{t}$ to denote the value of a variable before the realizations of the end-of-period shocks.

$s \in \{\mathbf{m}, \mathbf{d}, \mathbf{s}\}$ and with knowledge of her/his potential-spouse type.²⁷

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.²⁸ Thus we are assuming that the household allocation is Pareto efficient.²⁹

Households are also subject to fertility shocks at the beginning of each period and marital shocks at the end of each period. Asset, consumption and participation choices are made after fertility shocks and income are observed but prior to the realization of the marital shock. Households get divorced with probability $d_{eg,t}$, divorced individuals remarry with probability $r_{eg,t}$, 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:

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:

$$\hat{c}_t(k_t) + a_{t+1}^{\mathbf{m}} = Ra_t^{\mathbf{m}} + (y_{eft} - \kappa(k_t))P_t + y_{emt} \quad (4)$$

²⁷A 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.

²⁸A husband's labor force participation could also be an outcome but since we assume his disutility of working is zero, he will always work

²⁹See Del Boca and Flinn (2010) for an insightful analysis of various modeling approaches one may use to determine household allocations.

where a_t^m 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_{t+1}^m 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_{t+1}^m are allocated to the wife and the remainder to the husband. Thus, the laws of motion for $a_{g,t+1}^s$ are given by:

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

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.³⁰ Recalling that the children reside with their mother, the budget constraint of a divorced woman at time t is given by:

$$\hat{c}_t(k_t) + a_{f,t+1}^d = Ra_{ft}^d + (y_{ft} - \kappa(k_t))P_t + h(k_t, y_{mt}) \quad (6)$$

whereas that of a divorced man is:

$$c_t + a_{m,t+1}^d = Ra_{mt}^d + y_{mt} - h(k_t, y_{mt}) \quad (7)$$

where a_{t+1}^d 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.

³⁰In the model, as long as his child is below the age of twenty.

Thus, a divorced individual i that had saved $a_{i,\bar{t}+1}^{\mathbf{d}}$ faces at the end of period t the following law of motion for the asset position upon entering period $t + 1$.³¹

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

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

$$\hat{c}_t(k_t) + a_{t+1}^s = Ra_t^s + y_{ft} - \kappa(k_t) \quad (9)$$

whereas for a single man:

$$c_t + a_{t+1}^s = Ra_t^s + y_{mt} \quad (10)$$

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.,³²

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

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 $b_g^s(\bar{y})$ that depends on own past earnings history $\bar{y}$, 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 .

³¹Recall that a divorced agent is marrying another divorced agent by assumption, hence i and j both have $\mathbf{d}$ superscripts in equation (8)

³²We assume that singles marry other singles.

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:

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

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.³³ Thus, the value of being a divorced woman at

³³Note 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.

time t is given by:

$$\begin{aligned}
V_{ft}^{\mathbf{d}}(\Omega_t) &= \max_{c_t, P_t, a_{t+1}^{\mathbf{d}}} \frac{(c_t)^{1-\sigma}}{1-\sigma} - \psi_e^{\mathbf{d}}(k_t)P_t + \beta \left\{ (1 - r_{eg,t})\mathbb{E} \left[V_{f,t+1}^{\mathbf{d}}(\Omega_{t+1} | \Omega_t) \right] + r_{eg,t}\mathbb{E} \left[V_{f,t+1}^{\mathbf{m}}(\Omega_{t+1} | \Omega_t) \right] \right\} \\
s.t. \quad &\text{divorced female budget constraint (6) and asset law of motion (8)} \\
\text{and} \quad &x_{t+1} = x_t + P_t
\end{aligned} \tag{13}$$

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:

$$\begin{aligned}
V_{mt}^{\mathbf{d}}(\Omega_t) &= \max_{c_t, a_{t+1}^{\mathbf{d}}} \frac{(c_t)^{1-\sigma}}{1-\sigma} + \beta \left\{ (1 - r_{eg,t})\mathbb{E} \left[V_{m,t+1}^{\mathbf{d}}(\Omega_{t+1} | \Omega_t) \right] + r_{eg,t}\mathbb{E} \left[V_{m,t+1}^{\mathbf{m}}(\Omega_{t+1} | \Omega_t) \right] \right\} \\
s.t. \quad &\text{divorced male budget constraint (7) and asset law of motion (8)}
\end{aligned} \tag{14}$$

Married Household's Problem

A couple that enters the retirement period married solves

$$\begin{aligned}
V_t^{\mathbf{m}}(\Omega_t) &= \max_{c_t, a_{t+1}^{\mathbf{m}}} \frac{c_t^{1-\sigma}}{1-\sigma} + \beta V_{t+1}^{\mathbf{m}}(\Omega_{t+1} | \Omega_t) \\
s.t. \quad &\text{retired married budget constraint (12) and asset law of motion (5)}
\end{aligned} \tag{15}$$

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.

In the working stage, however, a couple that enters a period $t < t^R - 1$ married solves³⁴

$$V_t^{\mathbf{m}}(\Omega_t) = \max_{c_t, P_t, a_{t+1}^{\mathbf{m}}} \chi \left[\frac{c_t^{1-\sigma}}{1-\sigma} - \psi_e^{\mathbf{m}}(k_t) P_t \right] + (1-\chi) \left[\frac{c_t^{1-\sigma}}{1-\sigma} \right] + (1-d_{eg,t})\beta \mathbb{E} [V_{t+1}^{\mathbf{m}}(\Omega_{t+1} | \Omega_t)] \\ + d_{eg,t}\beta \left\{ \chi \mathbb{E} [V_{f,t+1}^{\mathbf{d}}(\Omega_{f,t+1} | \Omega_t)] + (1-\chi) \mathbb{E} [V_{m,t+1}^{\mathbf{d}}(\Omega_{m,t+1} | \Omega_t)] \right\} \quad (16)$$

s.t. the budget constraint for married couples (4) and asset law of motion (5)

and $x_{t+1} = x_t + P_t$

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.³⁵

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:

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

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

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

Single Household's Problem

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

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

s.t. retired single budget constraint (12) and asset law of motion (11)

³⁴As 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.

³⁵Note 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:³⁶

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

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 women and the other with the probabilities faced by men from our 1935 cohort. The statistics generated by the model therefore are the right ones to compare with the data which are reported by gender and cohort.³⁷

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 within 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 affect work decisions 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).

³⁶As 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.

³⁷In 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 $\psi_e^s(k)$). 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 is allowed to vary by marital status, education, and children’s age. The disutilities are calibrated within the model in order to match female labor force participation rates by marital status and education for the 1935 cohort. The distinction between working as a single woman versus married or divorced was particularly relevant to the 1935 cohort who grew up thinking of married women primarily as homemakers (the June Cleaver generation). With respect to children, we distinguish between mothers with young children (below the age of 5) and those with older children assuming that only the former bear any additional disutility from working.³⁸ Lastly, these disutilities are allowed to differ by education, reflecting the fact that jobs may have non-monetary rewards (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). 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 the absence of 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 psychic 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.³⁹ Using data to calibrate the monetary portion 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

³⁸Mothers of young children may be especially reluctant to work. See, for example, Bernal (2008) and Bernal and Keane (2009) for some evidence regarding the effects of mother’s working on children’s development.

³⁹The 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.

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:

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

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, and δ is the additional 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, $x_{ft} = \sum_{\tau=1}^t P_{f\tau}$.

The stochastic component to wages w_{et}^e is assumed to be the sum of a (observable to the agent) persistent component (z_{et}) and a transitory component (η_{et}):

$$w_{et} = z_{et} + \eta_{et} \quad (22)$$

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

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

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.⁴⁰

⁴⁰This specification and its estimation is discussed in great detail in, for example, Storesletten, Telmer, and Yaron (2004) and Guvenen (2007), as are the characteristics of the autocovariance functions and the variance growth in the life cycle which motivate the functional form of the stochastic process.

In order to estimate this process directly from the data one needs a panel dataset which includes individual wages across the life-cycle and a measurement of the individual’s ability. One way to measure ability’s 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 both returns to ability and age polynomials did not change over time and use the NLSY79 dataset to estimate these parameters.

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. First, we construct data on hourly wages ($y_{em,t}$) for men of education level e using data on earnings and total hours worked.⁴¹ We then run the following regression in order to estimate the parameters $\hat{\gamma}_{em1}, \hat{\gamma}_{em2}$ of the second degree polynomial on age and the parameter for returns to ability λ_e :

$$\ln y_{em,t}^{NLSY} = \mathbb{D}_{em,t} + \gamma_{em1}age_t + \gamma_{em2}age_t^2 + \lambda_e \ln AFQT + w_{em,t}$$

where $\mathbb{D}_{em,t}$ is a set of year dummies and $AFQT$ denotes the agent’s AFQT score. We use the residuals from this regression to estimate the parameters of the stochastic process $(\sigma_\eta, \sigma_\epsilon, \rho)$ using the minimum distance estimator first proposed by Chamberlain (1984). This method seeks parameters which minimize the distance between the empirical covariance matrix of income residuals and the one obtained from simulating the income process outlined above. This choice of estimator is standard in the literature and its use and identification in this specific income process is described in detail in, for example, Storesletten, Telmer, and Yaron (2004) and Heathcote, Storesletten, and Violante (2010).⁴²

The last step is to compute the time-varying intercepts $\tau_{em,t}$. Note that we cannot simply recover these from the year dummies from the previous regression because those do not span the years pertinent to our two cohorts. Our strategy is to first-difference log wages in a panel dataset as this will eliminate the effect of ability on wages (i.e., they become a fixed effect). We choose the PSID to estimate the time-varying intercepts using first-differenced log wages for men by education level.⁴³ Using data on

⁴¹See the Appendix for details of the sample selection.

⁴²We thank Gianluca Violante for kindly providing us with the estimation code.

⁴³The 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. Although we use all the waves from 1968 to 2009 of the PSID, these years still do not fully span the life-cycles of our two cohorts. Namely, our 1935 cohort was 33 years old in 1968 and our 1955 cohort was 54 years old in 2009. We assume that the wage intercept for the first period of the 1935 cohort (ages 25-29) is the same as the one for the second period (ages 30-34), whereas the last intercept for the 1955 cohort (ages 55-59) is the same as the period before (ages 50-54). As can be seen in the Figures 7 and 8 these are reasonable assumptions.

earnings and hours worked for white males aged between 25 and 65, we construct a measure of log hourly wages. We regress

$$\Delta \ln y_{em,t}^{PSID} - [\hat{\gamma}_{em1}(age_t - age_{t-1}) - \hat{\gamma}_{em2}(age_t^2 - age_{t-1}^2)] = \Delta \mathbb{D}_{em,t} + \xi_{em,t}$$

where $\Delta \ln y_{em,t}^{PSID} = \ln y_{em,t}^{PSID} - \ln y_{em,t-1}^{PSID}$, $\hat{\gamma}_{em1}, \hat{\gamma}_{em2}$, are the estimated values of the age coefficients (from the NLSY) and $\Delta \mathbb{D}_{em,t}$ are the differenced year dummies. To obtain a value for each τ_{em} we average the $\mathbb{D}_{em,t}$ values obtained over the 5 years corresponding to the τ_{em} interval (e.g. $\tau_{em,3}$ is the average over $\mathbb{D}_{em,t}$ estimated for the years 1970 to 1974). This strategy requires a value for the initial year dummy to anchor the sequence. We calibrate the values of these anchors (one for each education level) internally in the model where they will be used to help match the lifetime skill premium for men and the first period gender gap. The time dummies obtained this way are plotted in Figure 7 and 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. The fact that women move in and out of the labor force implies that we cannot use the same method as before to estimate experience and time intercepts for women's wage processes. We calibrated the latter internally so that the model generates the period-by-period gender wage ratio and skill premia for women as seen in the data. The implied values of $\tau_{ef,t}$ for our benchmark cohort are plotted in Figure 7. With respect to returns to experience, 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 a 2% return to wages from an extra year of participation during the ages of 25-40 for women in the 1935 cohort.⁴⁴

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.⁴⁵ 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

⁴⁴For example, Attanasio, Low, and Sanchez-Marcos (2008)'s calibrated returns to experience imply a return of about 2% for the cohort of women born around 1945. Olivetti (2006) estimates the return 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.

⁴⁵For 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 child support equivalent to 10% of his current income (unless she remarries). We do robustness checks using other values for this parameter.⁴⁶

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.⁴⁷

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 spousal 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 mimics 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 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

⁴⁶In 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.

⁴⁷The 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)). In our robustness check we investigated other proportions as well.

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.⁴⁸ 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/remarriage conditional on both year and duration of marriage/divorce due to small sample size in each year-of-marriage bin for each of our cohorts. Given these considerations, we choose a simple alternative. Recall that in our model, the first period corresponds to the ages of 25-29, the second period, to the ages of 30-34, and so on. Conditional on gender and education, we compute the proportion of people who begin age 25 as married but who undergo at least one divorce between the ages of 25-29. This gives us the probabilities of divorce between periods 1 and 2 in the model. We proceed analogously for the remaining periods, deriving 4 age-varying divorce probabilities by education and gender. Thus, our model reproduces, period-by-period, the correct proportion of people who divorced in each age bracket.

For remarriage rates, due to sample size restrictions, we proceed in a different way than above.⁴⁹ We calculate, conditional on gender and education, the proportion of people who remarry before the twentieth anniversary of their last divorce and assume a uniform remarriage rate over this twenty year interval.⁵⁰ The resulting divorce and remarriage rates are reported in Tables 7 and 8.⁵¹

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 a woman has in her lifetime, by education level, as in the PSID data for each cohort.⁵² In particular, for the 1935 cohort, single women and college

⁴⁸The SIPP is a series of short panels (from 2.5 to 4 years) of approximately 14,000 to 36,700 households in the US.

⁴⁹The number of people in the SIPP who entered ages 25 and 30 as divorced is too small – around 30 for men and 15 for women – to use the same procedure.

⁵⁰We augment our birth cohorts by 2 years due to sample size (i.e. the 1935 cohort is defined as all people born between 1933-1937). To calculate the remarriage rate we proceed in the following manner: individuals who divorced after the age of 30 contribute to the remarriage rate if they remarried within 20 years of their divorce and prior to the age of 60, if not, they are considered divorced. For those individuals who enter the age of 30 already divorced, the 20 year window in which to count as “remarried” commences as of age 30. This is the correct procedure since the initial proportion of individuals across marital states already accounts for divorces and remarriages that took place before the age of 30.

⁵¹Recall 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. Note that the median duration of divorce is 5 years for the 1935 cohort and 6 years for the 1955 cohort.

⁵²We choose to match the proportion of women (by education and marital status) who are mothers because LFP behavior is driven more 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 on the age of the youngest child in the household, we chose to focus on the distinction between mothers and others (this is the strategy followed by Attanasio, Low, and Sanchez-Marcos (2008)).

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 9. 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 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.⁵³

For the 1955 cohort the structure and values of fertility shocks are the same as those for the 1935 cohort. The frequencies in the initial period are adjusted so as to match the proportions in the data for the 1955 cohort between the ages of 25-29 as reported in Table 9. 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.⁵⁴ This generates the following numbers of children over their lifespan: 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.⁵⁵ Its exact numbers (by child's age) are reported in the Appendix.

⁵³The 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.

⁵⁴Once 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.

⁵⁵This scale is very similar to the OECD scale, but it has the advantage that it was computed based on expenditure data from families.

Women who have children under the age of 10 at home are assumed to incur childcare costs if they work. These costs are assumed to depend only 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.

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 reasonable benchmark. In our robustness section we explore values of $\alpha = 0.3$ and $\alpha = 0.7$ 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.⁵⁶ Although this mapping is only approximate,

⁵⁶Note 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.

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 match the proportions of men and women who go to college, while the remaining parameters are used to match LFP.

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 work disutility 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

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

where $\bar{c}_{fe}^{\mathbf{m}}$ is the average per-period consumption of married women with education level e .⁵⁷ 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 as these have very low LFP rates, requiring the model to generate a large disutility number in order to “explain” this. 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 68.8% of average per-period female wages, and those of an older child are 24.0%. We can compare these values to those obtained by

⁵⁷To calculate $\bar{c}_{fe}^{\mathbf{m}}$ 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.

Attanasio, Low, and Sanchez-Marcos (2008). In their calibrated model they find that childcare costs are 66% of a woman's mean earnings.⁵⁸

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

Turning next to education (college) costs, since these 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.⁵⁹ To do so we solve for the z_g such that:

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

where $\bar{\omega}_g$ is the average cost of higher education (by gender) implied by the calibrated parameter μ_g and $\bar{c}_{fh,t}^{\mathbf{m}}$ is the average consumption of married college women at time t . Using this metric, women's education costs correspond, on average, to a 16.5% decrease in per-period consumption, whereas the consumption reduction implied by the average cost of higher education for men is 18.2%.⁶⁰ 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.

⁵⁸This number is reported in terms of the mean earnings of a thirty-year-old woman who worked continuously prior to childbirth.

⁵⁹We 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 .

⁶⁰An alternative is to consider the average cost incurred by those who actually choose higher education. Computing z_g yields an equivalence of 11.9% of consumption for women and 14.1% of consumption for men.

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 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 10.⁶¹ The analogous statistic in our model would be the LFP of women with children under 10 during the first three periods of life.⁶² This yields an LFP of 26.9%, 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 64.7% in the model.

Next, it may be useful to report college attendance by aptitude level and gender (see Table 10). Although we do not have data relevant to the 1935 cohort 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 as this was the case in the NLSY79 data. 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.⁶³

The quantitative model allows us to evaluate the welfare costs of divorce for women and men by education level. Given that 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 18.7% 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

⁶¹The 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.

⁶²After the first three periods, given the structure of fertility shocks in the model all children would be older than ten.

⁶³To 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 $w_e(\theta) = w_e(\theta_0) \exp(\lambda_e(\theta - \theta_0))$ 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 $w_h(\theta) - w_l(\theta) > 0$ (a necessary condition for individuals to attend college if costs are positive) yields $\lambda_h > \lambda_l$ as a sufficient condition.

willing to sacrifice in period 2 in order to remain in her married state. This is the z_e that solves

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

where $\bar{\psi}_e^s$ denotes the average disutility from labor of working women of education e in marital status $s \in \{\mathbf{m}, \mathbf{d}\}$. Note that this average includes the disutility of working with children at home. This calculation yields a consumption loss of 24.6% for a college women and 19.2% for her high-school counterpart. The smaller gap in consumption loss between education groups comes from their differences in labor disutility as this 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 21.7% higher than that of their married counterpart's for college men and 20.9% higher for high-school men. Note that the difference in consumption gains for men across the two education levels is significantly smaller than the gap between college and high-school women. This is because married women have low LFP rates and thus, on average, the economic consequences of divorce for men arise mostly from a decrease in household size.⁶⁴

6 From 1935 to 1955

In this section we investigate the consequences of first introducing the 1955 family structure on the 1935 benchmark, followed by imposing the 1955 wage structure and lastly the combination of the two. Having changed the family and economic structures, we then explore the role of preferences. We conclude with a welfare analysis.

⁶⁴Peterson (1996) examines the economic consequences of divorce. He finds, using 1976-'77 data, that women's income-to-needs ratio (using a household-equivalence scale) fell by 27 percent whereas men's increased by 10 percent. All the literature surveyed by him agrees that the economic consequences of divorce were to make women worse off and men better off.

6.1 1955 Family Structure

1955 Divorce Rates

In order to gauge the importance of the markedly higher 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 from the ages of 25 to 44, in 1935 (around 5.5% on average) to the levels faced in 1955 (around 9.1% on average). A few things to note concerning the divorce probabilities faced by the 1955 cohort. First, the divorce probabilities faced by the 1955 cohort are, on average, higher for both genders and education groups. Second, the age profile of the divorce probabilities changes. Whereas in 1935 the divorce rate is more or less uniform across the ages of 25-44, in 1955 most of the divorce risk is concentrated in the ages of 25-34.

In this exercise, 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 a college spouse. As one can see from Figure 12, married women's LFP dramatically increases in the first two to three periods.⁶⁵ This reflects the much higher divorce risk faced by women in those periods leading them to work more both to accumulate greater experience and more assets. Both will improve a woman's welfare if she divorces. The magnitude of the LFP increase for college women during their first two periods of work is of around 33 percentage points while that of high school women is around 28 percentage points.

The increase in divorce probabilities also affects the proportion of men and women who choose higher education and, most importantly, this effect is asymmetric across genders: there is an almost 2 percentage point increase in the proportion of women who choose to enroll in higher education (29.1% in the benchmark vs 30.7%) versus a greater than 3 percentage point drop in the proportion of men who choose to do the same (39.2% vs. 35.9% 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 significant fashion. In particular, had no other forces been at work, the 1955 divorce rate would have closed almost half (49.0%) 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

⁶⁵There is also an increase in the LFP of divorced women (not shown).

their college counterparts because they have higher divorce rates and lower remarriage rates.

1955 Marriage Market

Next, we proceed to evaluate the effects of the remaining changes in the marriage market. 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 (i.e., the assortativeness of marriage), 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 with the label “Marriage Market”). 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 43.6%; for women, the increase is from 29.1% to 34.0%. 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. 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 a fall in childcare costs 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 costs 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 34.0% 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”). Interestingly, they decrease somewhat the LFP of married high-school women when they are younger (but increase the LFP of married college women), and increase it thereafter. This decrease in LFP during early working life is mainly because the total number of children for high school women drops significantly across the two cohorts (from 3.2 children per woman to 2.1). This leads to an increase in household consumption which then dampens women’s incentives to work.

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 70.6% of the average LFP gap across the two cohorts of college married women from age 25 to 50 and 60.5% of the average LFP gap for the equivalent high-school women. Interestingly, they generate too much of an increase in the LFP of divorced high school women (see Figure 14). With respect to education, these changes can explain almost all of the increase in the proportion of men who choose college (43.9% in this model versus 44.2% in the data for 1955) 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 $\{\tau_{em,t}\}_{t=2,\dots,t^R}$ for men were calculated directly from wage data for each cohort. The ones calculated for the years corresponding to the 1955 cohort (see Figure 8) will now be used for this experiment. That leaves us with the sequences of $\{\tau_{ef,t}\}_{t=1,\dots,t^R}$ for women 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.⁶⁶

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 $\{\tau_{ef,t}\}_{t=1,\dots,T}$ for women to match the period-by-period gender wage ratio and skill premia for women.⁶⁷

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 life-cycle) and somewhat smaller increase for high-school women (6 percentage points). The increase in female LFP is fairly uniform across most periods, as contrasted with the effect of divorce which increased LFP primarily in the first few work periods. College becomes significantly more attractive both for women and men, with 35.7% of women choosing it compared to 47.2% of men. Both men and women respond to the increased skill premia by 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 an increase in the returns to experience for women increased their LFP (see footnote 44 for a discussion), an additional exercise is to compound the changes in the wage structure above (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 a 3% increase in wages for the 1955 cohort (relative to the 2% for the 1935 cohort), a reasonable number given Olivetti’s (2006) finding, using the 1970 Census, of a 3-5% return to working full-time for an additional year.⁶⁸

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

⁶⁶Recall 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.

⁶⁷Note 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.

⁶⁸In the robustness checks we experimented with 5% returns to experience with very similar results.

basically no effect on married women’s LFP once they are older. The proportions of college men and women also increases minutely (less than 0.4 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”. Except for the first and second periods, 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, impacts 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 6 percentage points higher than the value in the data for the 1955 cohort.

The changes in wage and family structures combined can account for 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 work 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 38.7% – 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 49.5%) leading a 5 percentage point over-prediction 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 for married women, the too high LFP for divorced women, and the failure to account for the education gender gap.

6.3 Preferences

It is difficult to deny that social attitudes towards women changed dramatically over the lives of the two cohorts. Rather than treat preference changes as a residual that explains the remaining LFP gap, we investigate whether a simple change in work preferences is able to account for it. Prior to this, we

examine the changes required in education cost in order to match the education outcomes in the data.

The Psychic Cost of College

As noted previously, the combination of all the 1955 changes generated too large a proportion of college men and too small a proportion of college women. We 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 $\mu_f = 0.5540$ and $\mu_m = 1.2146$. This translates into a decrease in female average education costs of 28.1% whereas male education costs increase by a similar proportion (30.6%). 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.⁶⁹

How can one interpret the fact that the model requires an asymmetric reaction by gender on education costs? One way to think about the changed education costs 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 in women's psychic cost of college so as to undo the economic increase in costs.⁷⁰ 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. Thus, the decrease in female psychic costs may simply reflect their true lower value once discrimination is removed.

⁶⁹Across all our experiments, the additional change in LFP that comes from the changed ability compositions across education groups is miniscule.

⁷⁰Although the literature has also found that changing family background (i.e. more parents being college educated) has led to changes in the non-pecuniary costs of education, there is no reason for this to affect genders asymmetrically.

Work Preferences

There exists a growing literature that demonstrates the importance of culture for various outcomes including female LFP.⁷¹ 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 1935 and 1955 cohorts, we now ask whether changes in preferences (cultural change) towards work are able to explain the gaps between predicted LFP and the data.⁷²

Unfortunately, there is little quantitative evidence to discipline how preferences should be changed.⁷³ 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 over-predicting the LFP of divorced women, we will assume that the preferences of the latter now resemble those of their married counterparts. 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 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 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

⁷¹See Fernández and Fogli (2009) for the influence of culture on female LFP and Fernández (2010) for an overview of this literature.

⁷²An 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 in brain-biased technological change as in Rendall (2010)).

⁷³See Fernández (2011) for a calibrated learning model of female LFP that makes use of longitudinal poll evidence to govern how preferences evolved.

section 6.3) but keeping female work preferences unchanged from the 1935 benchmark. Second, we allow preferences to change. Lastly, we do a full-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, expected utility in $\tau' = 1955$ is the same as $\tau^0 = 1935$, i.e.,

$$\sum_{t=t_1}^T \sum_{\theta} \pi_{\tau'}^{eg}(\theta) \mathbb{E}V_{egt}(z, \theta; \tau') = \sum_{t=t_1}^T \sum_{\theta} \pi_{\tau'}^{eg}(\theta) \mathbb{E}V_{egt}(\theta; \tau^0) \quad (27)$$

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 she/he faces the 1955 environment. A z greater than 1 indicates that agents were 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 2.7%. 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 10.9% decrease in all levels of consumption whereas high-school men would be willing to suffer a 4.1% 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 allow female preferences toward work to change as in section 6.3, in addition to the other changes in the economic and family structures. 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 1.9% 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 17.0% decrease in consumption whereas high-school men would need a 5.5% 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, not condition on education, and take into account the cost of college. 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, taken into that the education costs for 1955, $\omega(\tau')$ are not the same for 1935, $\omega(\tau^0)$ (i.e., we solve:

$$\sum_{t=t_0}^T \sum_{\theta} \sum_e \pi_{\tau'}^{eg}(\theta) \mathbb{E}V_{egt}(z, \theta, \omega(\tau'); \tau') = \sum_{t=t_0}^T \sum_{\theta} \sum_e \pi_{\tau^0}^{eg}(\theta) \mathbb{E}V_{egt}(\theta, \omega(\tau^0); \tau^0) \quad (28)$$

Note that 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 has $\pi_{\tau^0}^{eg}(\theta)$ rather than $\pi_{\tau'}^{eg}(\theta)$.

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 5.4% decrease in consumption across all states in order to be indifferent between the 1955 environment and the 1935 environment; men would require a 14.7% decrease. Not surprisingly given the preceding analysis, the gain for men from 1955 relative to 1935 is considerably larger than that for women.⁷⁴

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 $\chi = 0.5$,

⁷⁴If female work preferences are not allowed to change, women in 1955 would require a 3.4% decrease in consumption levels whereas men would require an 12.1% decrease.

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 concentrated on their first periods of working life. While in the original benchmark the increase in LFP during the ages of 25-34 was around 30 percentage points for both college and high school, those same statistics are now around 20 and 15 percentage points, respectively. With respect to education choices, the responses are similar (albeit a bit smaller than in the original benchmark). 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 58% 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 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 71.3% to 62.7% and their high school counterparts' participation is 12 percentage points lower (from 63.7% to 49.6%), 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 (34.7%). 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. The experiments on these recalibrated benchmarks also yielded similar results as the original benchmark.⁷⁵

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. Each in isolation is able to account for a large proportion (60%) of the observed changes in female LFP. The increased probability of divorce faced by the 1955 cohort, in particular, 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

⁷⁵All results from the robustness checks are available upon request.

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 allow female preferences for work to change, 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 a significant first step by incorporating interactions between family structure, education, and work choices in a dynamic life-cycle setting, much work remains to be done. In the future, it may also be important to endogenize marriage, divorce, and women's bargaining power.

References

- Albanesi, S. and C. Olivetti (2009a). Gender roles and medical progress. working paper.
- Albanesi, S. and C. Olivetti (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.
- Cameron, S. and J. Heckman (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.
- Caucutt, E. M., N. Guner, and J. Knowles (2002). Why do women wait? matching, wage inequality, and the incentives for fertility delay. *Review of Economic Dynamics* 5, 815–855.
- Chamberlain, G. (1984). Panel data. In Z. Griliches and M. Intriligator (Eds.), *Handbook of Econometrics*, Volume 2, Chapter 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.
- Cubeddu, L. and J.-V. Ríos-Rull (2003). Families as shocks. *Journal of the European Economic Association* 1(2/3).
- Del Boca, D. and C. Flinn (1995). Rationalizing child-support decisions. *American Economic Review* 95(5), 1241–1262.
- Del Boca, D. and C. Flinn (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.
- Edlund, L. and R. Pande (2002). Why have women become left-wing: the political gender gap and the decline in marriage. *The Quarterly Journal of Economics* 117(917-961).
- Fernández, R. (2010). Does culture matter? In A. Bisin, J. Benhabib, and M. Jackson (Eds.), *Handbook of Social Economics*. Elsevier.
- Fernández, R. (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 investigation 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. (2011). Women's college decisions: How much does marriage matter? *Journal of Labor Economics* 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.
- Güvenen, 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.
- Katz, L. F. and D. H. Autor (1999). Changes in the wage structure and earnings inequality. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, Volume 3A. Amsterdam: North-Holland.
- Keane, M. P. and K. Wolpin (1997). The career decisions of young men. *Journal of Political Economy* 105(3), 473–522.
- Keane, M. P. and K. Wolpin (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.
- Knowles, J. (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.
- Peterson, R. R. (1996). A re-evaluation of the economic consequences of divorce. *American Sociological Review* 61(3), 528–536.
- 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.
- Stevenson, B. (2008). Divorce law and women's labor supply. *Journal of Empirical Legal Studies* 5(4).
- Stevenson, B. and J. Wolfers (2006). Bargaining in the shadow of the law: Divorce laws and family distress. *The Quarterly Journal of Economics* 121(1), 267–288.
- 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 Maisonnette (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 at 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 white married 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 white married women, by birth cohort and age




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





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 birth cohort year, who report being in the labor force. "Married" is defined 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 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.




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




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




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




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




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	19.81	16.74	23.94	18.21
	Remarriage	87.62	86.06	87.61	91.27
1955 Cohort	Divorce	38.22	36.53	40.79	31.31
	Remarriage	86.75	85.46	85.86	92.18

Notes: SIPP 2004. Divorce rates are calculated as the proportion of first 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 first divorce. Birth cohorts are defined as people born in a five year interval centered on 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}	0.1650
	γ_{lm2}	-0.00308
Regression log wage on age and age ² , College men	γ_{hm1}	0.2276
	γ_{hm2}	-0.00285
Gradient of log AFQT scores on log wage, by educ.	λ_l	0.11157
	λ_h	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	$\{d_{eg,t}, r_{eg,t}, m_{eg,t}\}_{t=1,\dots,4}$	see Tables 6, 7, 8

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

		Married	Married with Child	Divorced	Divorced with Child
Parameter values	High School	0.0840	0.2233	0.0962	0.1182
	College	0.0420	0.0959	0.0448	0.0667
Consumption equivalence	High School	11.1%	33.6%	12.7%	15.9%
	College	5.3%	12.7%	5.7%	8.6%

The consumption equivalence numbers give the fraction of average consumption of a married woman of education 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}	1.1450
	κ_{old}	0.4029
Initial intercept log wages for men, by education	$\tau_{lm,1}$	0.0841
	$\tau_{hm,1}$	0.6128
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	0.0839
Education cost	μ_f	0.8409
	μ_m	0.9298
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	29.05	29.05
Proportion of men who go to college	39.24	39.24
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

	age	Men		Women	
		College	High School	College	High School
1935 Cohort	30	51.24	44.19	18.14	41.01
	35	26.06	18.90	19.52	3.84
	40	22.65	15.91	23.64	5.63
	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 probabilities

		d_1	d_2	d_3	d_4
1935 Cohort					
Men	HS	5.87	8.60	8.49	6.44
	College	6.45	6.03	4.67	5.67
Women	HS	5.39	4.93	5.34	5.38
	College	3.44	3.84	4.62	5.46
1955 Cohort					
Men	HS	18.04	11.60	7.85	4.70
	College	17.42	9.17	5.65	5.06
Women	HS	14.48	9.90	5.36	3.61
	College	13.99	9.90	6.17	2.93

Note: d_1 refers to ages 25-29, d_2 to ages 30-34, d_3 to ages 35-39 and d_4 to ages 40-44. These numbers were derived using 2004 SIPP data. See text for details about their calculation.

Table 8: Per period remarriage probabilities

	Women		Men	
	HS	College	HS	College
1935 Cohort	38.84	38.49	37.23	42.48
1955 Cohort	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 probability of remarriage. See text for details about their calculation.

Table 9: 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 10: Proportion of men and women in the benchmark 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 used for the skill premium and the gender wage ratios are computed from the CPS using the individuals' reported labor income and hours and weeks worked last year. We use the sample of white men and women who do not live in group quarters. Prior to 1977, for hours per week, we use the variable which reports the hours worked in the previous week, by intervals; we use the midpoint of the interval. From 1977 onwards, we use the variables for "usual hours worked per week" (last year) and the continuous variable for number of weeks worked last year. Whenever we compute lifetime averages for a variable, we first compute the average of the variable over each year and then average across years. Sample weights are used throughout (PERWT). Concerning top-coded observations, we follow the procedure in Katz and Autor (1999). We multiply 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

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

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

Age Profiles and Ability Level Gradients:

We construct an hourly wage measure from NLSY79 data by dividing total earnings by total hours worked the previous calendar year, deflated by CPI-U into 1992 dollars. We use only white men 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.

The coefficients on age earning profiles for men, $\gamma_{em1}, \gamma_{em2}$ and the ability gradient λ_e are computed using data from the NLSY79. We use an individual's AFQT89 score as a measure for θ . In order to estimate the ability gradient λ_e we regress the log wage on the log of the individual's AFQT89 score under the assumption that the error term is uncorrelated with this score. λ_e is estimated separately for each education group $e = \{l, h\}$.

Estimates of the Labor Shock Processes:

Given the residuals from the regression of log wages on age and ability gradients using the NLSY79 data, we estimate the parameters for the persistent and transitory shocks using the Minimum Distance Estimator (Chamberlain (1984)). The methods of estimating this 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 $\sigma_{\epsilon,e}^2$ and $\sigma_{\eta,e}^2$ during the estimation process. In the model, we use as inputs the average value across the sample and this is what we report.

Time-Changing Wage Intercepts:

In order to estimate the time-varying age intercepts for men $\tau_{em,t}$ we use the pooled sample of PSID for the years 1968-2009, restricted to white males who are heads of households. We exclude individuals in 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 1645 individuals in the "low" education group and 1261 in the "high" education group.

In order to account for the fact that the PSID does not include information on an individual's

ability, we use first-differenced earnings in the estimation; note that this implies that we recover only first differences in $\tau_{em,t}$. We run the following regression separately by education:

$$\Delta \ln y_{em,s}^{PSID} - [\hat{\gamma}_{em1}(age_s - age_{s-1}) - \hat{\gamma}_{em2}(age_s^2 - age_{s-1}^2)] = \Delta \mathbb{D}_{em,s} + \xi_{em,s}$$

where the left hand side is the log earnings measure free of the age effects estimated from NLSY, and $\Delta \mathbb{D}_{em,t}$ are the differenced year dummies. The earnings variable is total labor income of the head deflated it to 1992 dollars using CPI for all urban consumers. Education levels are assigned as the highest grade completed.

In order to obtain a value for each $\tau_{em,t}$ in the model, we average the $D_{em,s}$ values obtained over the 5 years corresponding to the $\tau_{em,t}$ interval. Recalling that our earliest data is from 1968 (when our 1935 cohort is 33 years old), our assumptions (see footnote 43) imply that $\tau_{em,1} = \mathbb{D}_{em,1968}$. We calibrate $\tau_{em,1}$ internally. For the 1955 cohort, on the other hand, we are missing the last intercept $\tau_{em,5}$, since the cohort is only observed up to the median age of 55 in 2009. We thus assume that $\tau_{em,5} = \tau_{em,4}$

A.3 Consumption Deflator:

We use an altered McClements scale ($\mathbf{e}(k_t; s)$) in order to deflate household consumption. Table 11 reproduces the original McClements scale in normalized for one adult.

Table 11: McClements Scale

1 adult	2 adults	+ 1 adult	+1 child, by age:						
			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 life-cycle. 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 during the last period, plus 32% of y_T from this bend-point to a higher bend-point equal to $1.59\bar{y}_T$, plus 15% of the remaining y_T exceeding this last bend-point. For married households, this process is done for both the husband and the wife; the household total benefits are either the 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. In addition to assets, our model has four other potential continuous state variables: the persistent components of earnings 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 . Handling more than one continuous state variable is possible but computationally costly. We discretize these four variables, leaving assets as our only continuous state. We choose 10 nodes whose locations are age-dependent for each of the earnings components and the quintiles of ability distribution outlined in the main text.

During the working stages of the life-cycle, 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 as explained below.

For all periods $t > 4$, since there are no longer any transitions across different marital status, the maximization problem for the single men and women and divorced men is a straightforward consumption-savings problem. The problem of the married couple and of the divorced woman combines the discrete participation choice of the wife together with the continuous choice of assets. The combination implies that concavity of the value function is not guaranteed even if one controls for the participation decision that period. Given enough uncertainty the value function conditional on today's participation is concave

We follow Attanasio et al. (2008) and impose (and check) a unique level of reservation assets a_t^* at which, given the values of all other state variables, the conditional 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 optimization problem of the household who enters a period $t \leq 4$ as a married couple must take into account the continuation values of the husband and wife, which are different. It's important to note that the solution of the optimization problem for the divorced and single (men and women) involves the calculation of a fixed point. Given that each individual has perfect information about their potential spouses' contemporaneous state variables, she/he will solve her/his own optimization problem taken as given their potential spouses' asset allocation for next period.