

Divorce Risk, Wages, and Working Wives: A Quantitative Life-Cycle Analysis of Female Labor Force Participation

Raquel Fernández* Joyce Cheng Wong
New York University International Monetary Fund

February 2013^{†‡}

Abstract

This paper examines some of the forces that led to the large increase in married women's labor force participation decisions (LFP) over two cohorts of women in the 20th century. To do this we develop a dynamic quantitative life-cycle model with incomplete markets and risk-averse agents who differ in their education endowments (college or high school) and make work, consumption, and savings decisions. We calibrate the model to match key life-cycle statistics for the 1935 cohort and then change the economic environment (e.g., skill premium and gender wage gap) and the family structure (e.g., the assortativeness of marriage, divorce risk, and fertility) to replicate those for the 1955 cohort. We find that the higher probability of divorce and changes in wage structure faced by the 1955 cohort are each able to explain, in isolation, a large proportion of the observed changes in married women's LFP. Surprisingly, higher divorce risk does not impact LFP through women's desire for greater labor market experience nor does it increase marital assets. Instead, it is the result of the (endogenously) conflicting preferences of husbands and wives towards the adjustment of marital consumption in the face of higher divorce risk.

*NBER, CEPR, IZA, ESOP. The author wishes to thank the NSF and the Russell Sage Foundation for financial support.

[†]We thank Meta Brown, Jeremy Greenwood, and Gianluca Violante for helpful comments on an earlier version as well as seminar audiences at Mannheim University, University of Zurich, Paris School of Economics, Oslo University, IIES at Stockholm University, the NY Federal Reserve, Georgetown University, Hunter College, the World Bank, the IMF, and at various conferences including the 2012 SED annual meetings, LACEA-LAMES, and the Human Capital Conference at the Bank of Italy.

[‡]This is an extensively revised version of the paper previously entitled "The Disappearing Gender Gap: The Impact of Divorce, Wages, and Preferences on Education Choices and Women's Work" (Fernández and Wong (2011)).

1 Introduction

The dramatic increase in the second half of the twentieth century in married women’s labor force participation (LFP) has given rise to a large and growing literature that explores the roles of the multiple forces that led to this evolution. Some explanations have emphasized technological change in the household (e.g. Greenwood, Seshadri, and Yorukoglu (2005)) or in medical/contraceptive technology (e.g., Goldin and Katz (2002), Albanesi and Olivetti (2009a,b), and Knowles (2007a)). Other explanations have emphasized changes in the wage structure such as in the skill premium, the gender wage gap, or in the returns to labor market experience (e.g., Jones, Manuelli, and McGrattan (2003), Gayle and Golan (2010), Olivetti (2006), and Knowles (2007b)). Yet others have focused on changes in culture (Fernández, Fogli, and Olivetti (2004), Fernández (2012)), changes in the structure of the economy (Galor and Weil (1996), Rendall (2010), Goldin (1990), and changes in child-care costs (Attanasio, Low, and Sanchez-Marcos (2008)).

Although several regression-based analyses have examined the relationship between divorce risk and women’s work by exploiting time differences in the introduction of unilateral divorce laws across US states (see, e.g., Peters (1986), Gray (1998), Stevenson (2008)), the dynamic quantitative literature has (with a few recent notable exceptions discussed in the next section) mostly ignored this issue.¹ This is understandable as this strategy allows the researcher to focus on other factors (e.g., child-care cost in Attanasio, Low, and Sanchez-Marcos (2008)) without the burden of the considerable computational complexity inherent in models with multiple marital states. Nonetheless, we believe that the parallel increase in divorce rates and married women’s LFP across many countries, as well as the compelling intuition for why greater divorce risk should lead to higher married women’s LFP (discussed below), warrants a deeper look at this topic.

The goal of our paper is to delve more deeply into the link between married women’s LFP and marital instability using a dynamic quantitative model, without neglecting the

¹Other econometric models have used proxies for individual variation in divorce risk to study the link between marital instability and divorce risk (see, e.g., Johnson and Skinner (1986), Papps (2006), Kneip and Bauer (2007)). There is also evidence that the anticipation of divorce is positively correlated with increased labor supply by married women (see, e.g., Lundberg and Rose (1999)). Overall, the econometric literature has concluded that greater divorce risk led to higher married women’s LFP.

potential importance of several other contributing factors. Since a woman’s motivation to work changes over her life-span and differs by education, we thought it important to conduct our analysis in a life-cycle framework, as has much of the quantitative research in this area (e.g., Attanasio, Low, and Sanchez-Marcos (2008), Eckstein and Lifshitz (2011)). This framework allows for a clear conditioning of the data by gender, education, marital status, and age which allows one to disentangle the differences in behavior across those groups. To this end, we chose to focus on two cohorts for which there is (almost) complete data: the 1935 and the 1955 cohorts.

The 1935 and 1955 cohorts experienced the largest increase in married women’s LFP over any twenty year period. Whereas only 40% of white women from the 1935 cohort worked during the ages of thirty to forty, for the 1955 cohort this proportion increased to 70%.² Over the life spans of these cohorts there were significant changes in the economic environment: the ten percentage point gender gap in college attendance of the 1935 cohort was completely eliminated, the average skill premium faced by each cohort over its working life increased for both men and women by some 12 percentage points, the overall life-time gender wage gap faced by each cohort fell substantially, and there is also evidence that the return to women’s work experience increased and child-care costs fell. At the family level there were also important transformations. Fertility dropped from 3.0 children per woman for the 1935 cohort to 1.9 for the 1955 cohort. Individuals married later and more assortatively, and the probability of a marriage ending in divorce doubled.³

To assess the quantitative contributions of the many factors listed above, we develop a dynamic life-cycle model with incomplete markets and risk-averse agents who differ in their education endowments (college or high school) and make work, consumption, and savings decisions. In the spirit of Cubeddu and Ríos-Rull (2003), marital status and fertility are treated as “shocks” which are conditioned on gender, age, and education.⁴ Agents can be married, single, or divorced. They make consumption, saving, and work decisions in

²The work experience of black women is different from that white women and hence we choose not to mix the two in this analysis. See Potamites (2007) for an analysis of differences in the evolution of LFP across black and white women.

³All quantitative statements in this paragraph are based on authors’ calculations using CPS, Census, and the 2004 SIPP.

⁴See also, e.g., Love (2010) and Ríos-Rull and Sanchez-Marcos (2002) for uses of this approach.

each period and women’s disutility from working may depend on her education, marital status, and on the presence of children. Married agents share household consumption and save/borrow jointly. Their choices maximize the weighted sum of the wife’s and husband’s welfare; unmarried individuals maximize their own welfare. Upon divorce, assets are split between the two ex spouses. Any child belonging to the ex-couple is assumed to live with the mother and the ex-husband provides child-support payments.

We construct LFP, wage, and marital status moments for each of our cohorts at different ages, by gender and education, using a combination of several datasets including the Current Population Survey (CPS), National Longitudinal Survey of Youth (NLSY79), and the Survey of Income and Program Participation (SIPP). We calibrate our model to the life-cycle statistics of the 1935 cohort, conditioning everywhere on gender and education, as well as on age and marital status. We use the calibrated model to assess the quantitative impact of various changes that occurred over the life span of the two cohorts: the increase in assortativeness of marriage, the changing distribution over the age of first marriage, fertility patterns, divorce and remarriage probabilities, the increase in the skill premium and the decline in the gender wage gap, the higher returns to women’s labor market experience, and the declining price of child care. All together we find that these changes are able to account for a large proportion (85%) of the change in married women’s LFP. The increase in divorce risk has a large impact on married women’s LFP, particularly when they are younger. We find that, in isolation, it can account for over 42% of the LFP increase between the two cohorts during the ages of 25-40 for married college women and 49% of the LFP increase for married high-school women.

Our decision to keep marital status exogenous was driven by the desire of developing a reasonably tractable dynamic quantitative model that would permit us to understand the mechanisms that drive the increase in women’s LFP.⁵ Our initial intuition was that working has two potential benefits for married women, both of which become more important in the face of higher divorce risk. First, working increases married women’s experience in the labor market thus allowing them to command higher wages in the future. Since divorced

⁵As discussed in the next section, including endogenous marital status in a model with endogenously evolving state variables vastly increases its computational complexity. That makes it very difficult to understand the quantitative role played by various mechanisms since any changes in modeling assumptions or robustness checks require recalibrating the model, a process that can potentially take months.

women are more likely to work than married women, a higher probability of divorce increases the importance of greater labor market experience. Second, working while married can, by increasing income, potentially increase household savings. Since upon divorce women experience a drop in household income, a higher probability of divorce makes greater marital savings more attractive.

Interestingly, we found that the desire to increase labor market experience did not play a significant quantitative role in the increase in married women’s LFP. Even more surprising, we found that married household savings actually decreased in the face of higher divorce risk. Our analysis indicates instead that married women’s higher LFP is the result of the (endogenously) conflicting preferences of husbands and wives towards the adjustment of marital consumption in the face of higher divorce risk. When confronted with higher divorce risk, married women would prefer to increase saving in order to transfer more assets to the divorced state, thus helping them smooth consumption. Men, on the other hand, would prefer to increase consumption in the married state as this is what allows them to smooth consumption. A “compromise” is reached – the solution to a Pareto weighted utility maximization problem – that involves married women working more which allows married savings to decline by less than what they would otherwise. This outcome is obtained by maintaining fixed the Pareto weights that govern the solution to the household bargaining problem. Thus, in this framework women’s increased LFP is not the result of one party’s increase in bargaining power but rather the solution to opposing preferences in the household regarding how to allocate consumption across marital states.

Our paper is organized as follows: the next section provides a brief review of the quantitative literature in this area. Section 3 develops a dynamic life-cycle model with endogenous labor, consumption and savings. Section 4 presents the data used to parameterize the model for the 1935 cohort. Section 5 discusses some key features of the benchmark model. Section 6 investigates the effects of various changes in family and wage structure in generating women’s LFP outcomes. Section 7 explores the roles of experience, asset accumulation, and concavity in driving the quantitative results regarding divorce. Section 8 conducts various robustness checks and section 9 concludes.

2 A Brief Literature Review

Our paper contributes to the small but growing quantitative dynamic model literature that studies the evolution of female LFP without ignoring the issue of marital transitions. Most closely related to our work is the very interesting recent paper by Eckstein and Lifshitz (2011). The authors also use a dynamic life-cycle model with fertility and marital status determined by shocks (conditioned on education, age, etc.) to quantify the contributions of changes in marital stability, wages, education, and fertility to the evolution of female LFP over several cohorts born between 1925 and 1975.

Surprisingly, Eckstein and Lifshitz (2011) find that changes in marital status (later marriage combined with higher divorce probabilities) account for at most 1% of the changes in female LFP across cohorts born between 1925-1975. The authors find that changes in education and wages played the largest role in the transformation of female LFP. The most obvious key differences with our approach is the assumption of a linear utility function (we use a concave one) and, relatedly, assumptions regarding how married couples make decisions. A linear utility function implies that women do not react to the increase in risk, which may be an important element underlying the response of LFP to a higher divorce probability.⁶ Second, and relatedly, the authors do not allow saving or borrowing and assume that women make unilateral labor market decisions. These two features imply that an increase in the probability of divorce does not increase the level of disagreement between husband and wife: there are no savings to contribute to and the husband is unable to influence his wife's labor market decision. In section 7 we explore further the importance that these assumptions have in generating the differences in results.

Our analysis, like Eckstein and Lifshitz (2011), models marital status as the outcome of an exogenous shock, despite its obvious endogeneity. This modelling choice has some important drawbacks such as not allowing one to determine causality nor conduct a welfare analysis. On the other hand, it has important benefits such as permitting one to include, in a computationally tractable fashion, other endogenously evolving state variables which we consider critical to understanding about the reaction of LFP to divorce risk. The

⁶See too the seminal paper by van der Klaauw (1996) on marital status and female labor supply in a setting with linear utility.

model’s tractability also allows one to conduct sensitivity analyses and to unpack the key mechanisms whereby increased divorce risk might lead to higher LFP. An alternative is to endogenize the marital status outcome. There are pros and cons to both strategies, with the balance depending primarily on the focus of the research question as illustrated by the papers discussed below.

Greenwood, Guner, Kocharkov, and Santos (2012) presents an elegant and insightful analysis of how household technological change and changes in the wage structure affected educational attainment, married female LFP, marriage and divorce rates, and the assortativeness of marriage in the post-war era. In that paper, marital instability and married women’s labor supply arise endogenously as responses to the declining economic importance of marriage (e.g., decreasing household returns to scale and improvements in home production) and to changes in the wage structure. Not surprisingly, however, the ambitious scope of their paper requires other important simplifications such as the assumption that agents are unable to borrow or save and that the economy is in a steady state, both initially in 1960 and in the year 2000.⁷ Their modeling choice makes sense given that their main interest is in producing a unified theory. Our focus, as in Eckstein and Lifshitz (2011), however is on understanding how different forces could impact women’s work decisions and analyzing the mechanisms that lead to different outcomes. In particular, treating marital status transitions as exogenous allows us to focus on the life-cycle features of work decisions and to include other endogenous state variables that we consider critical to thinking about the impact of greater marital instability on women’s labor supply, such as an individual’s labor market experience and a household’s asset holdings.⁸

A very ambitious paper by Mazzocco, Ruiz, and Yamaguchi (2007) examines the interaction between female LFP, savings, and marital decisions in an estimated model with limited commitment where marital status and bargaining power evolve endogenously. This paper is computationally extremely demanding, making it very difficult to investigate the roles

⁷See also Rios-Rull and Regalia (2001) who were among the first to develop a dynamic model with endogenous marital status (but with a continuous labor supply decision) to study the rise in single households. They likewise assume a stationary equilibrium in order to be able to solve the model.

⁸In addition, it is worth noting that there are at least some exogenous factors driving the change in marital stability such as changes in the legal environment (unilateral divorce law was introduced in most of the US in the mid and late 70s) and subsequent changes in culture (divorce became more socially acceptable).

played by various mechanisms. As a result of this (and data limitations), the authors focus their analysis on one 10-year cohort and do not distinguish individuals by education. Although they do not attempt to understand the drivers of LFP over time, their model highlights some of the important interactions between savings, labor supply, and marital status decisions.

A very nice recent paper by Voena (2012) builds on Mazzocco et al.’s framework in order to study how the introduction of unilateral divorce differentially affected the asset accumulation behavior of married households and the labor supply of married women across US states with different property rights laws. Her model also endogenizes marital status but, since the life-cycle aspect is not central to the paper’s main focus (she is not comparing across cohorts), it considerably lightens its computational demands by assuming that periods are lengthy (10 years) and by not distinguishing across women of different education levels. For our research, on the other hand, it is important to respect educational differences and to distinguish among different points in the life cycle (e.g., those with or without young children) as these differences show up clearly in the data.

3 The Model

In this section we present the economic environment and the households’ decision problems.

3.1 Some Preliminaries

Agents are born with gender g (either male m or female f). They enter life (at age 25 in the calibration) endowed with a given education level e that can be either high or low (h or l) and a marital type ζ .⁹ They are also endowed with an initial asset a_0 draw and an initial marital status s_0 . The marital status can be either married (**m**), divorced (**d**) or single (**s**). The distribution of the initial marital status is both gender and education dependent.

An individual’s life cycle can be divided into two distinct stages of life. The first stage

⁹The prior NBER wp version of this paper endogenized the education choices of agents. In this version we omit this choice as we found that simply including ability differentials did not provide for a sufficiently rich educational sector for the calibration. In any case, we obtained similar quantitative results for LFP.

corresponds to working life. In that stage, individuals receive fertility and marital shocks, wage draws, and make consumption-savings and work decisions. The second stage of life corresponds to retirement. In that stage, individuals face a non-stochastic environment and make consumption-savings decisions.

Before presenting the economic environment in greater detail, it is worth spending some time explaining our assumption of a marital type. Although we have assumed that marital status is exogenous, that assumption in and of itself does not greatly simplify the computational burden associated with an evolving marital status. The possibility that an agent may marry (either for the first time or after a divorce) in the future implies that agents need to know the probability with which they will be matched with spouses of different characteristics. In particular, they need to know the joint distribution over prospective spouses' assets, education, persistent wage shocks, labor market experience, and prior children. These characteristics affect the payoffs of marriage/divorce and hence calculating an agent's expected utility in order to solve for optimal consumption, savings and LFP requires, at each point in time, assessing the welfare associated with each possible outcome. One modelling choice often used to reduce the computational complexity is to assume a stationary distribution of characteristics (this is the assumption, for example, in Greenwood, Guner, Kocharkov, and Santos (2012) and Jacquemet and Robin (2011)). Our interest in an environment that is obviously experiencing considerable change (in education, divorce, women's LFP, etc.) rules out this avenue. Instead we assume that agents are endowed with a marital "type" which gives them information about the individual that they may marry. That is, we assume that individuals know both the education and the *current* value of the time-varying characteristics of their potential spouse (e.g., that agent's asset level, wage shock, labor market experience, etc., in that period but not in the future). This assumption significantly reduces the computational burden at the cost of decreasing the uncertainty individuals face.¹⁰

¹⁰An alternative possibility would be to assume that remarriage is to an agent with the same asset holdings as one's own, as in Voena (2012). Mazzocco, Ruiz, and Yamaguchi (2007) assumes that an unmarried agent meets others whose asset holdings are in some interval around one's own. Our specification has the advantage of respecting gender differences in asset accumulation behaviors while unmarried.

Preferences, Consumption, and Borrowing Constraints

The instantaneous utility function of an agent with gender g , education e , and 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 LFP decision, taking the value one if the agent works and zero otherwise. We henceforth assume that only women suffer disutility from market work, $\psi_e^s(k_t)$, that depends on e, s and on a vector k_t indicating the ages of her children in that period. Men's and single women's work disutility is normalized to zero; they thus always work and 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).

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 status 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. For computational ease we assume that divorced ex-spouses remarry at the same time so that the children are simply reassigned to the newly remarried couple.

Upon 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, the man pays his ex-wife 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. Finally, if the wife does not remarry by the time she enters retirement, she receives a proportion of her husband's retirement benefits.

Income

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

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.

3.2 The Work Stage

Individuals are assumed to spend periods 1 until period t^R in this stage. In each period, individuals receive an e, g -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 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. Note that married individuals, on the other hand, will in general not agree on the optimal choice of female LFP and savings. As in Chiappori (1988), we assume that married individuals choose household consumption and the wife's labor force participation

so as to maximize the weighted sum of the spouses' lifetime utilities. Thus the household allocation is constrained Pareto efficient.¹¹

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 1 to clearly illustrate the timing of shocks and decisions.

Budget Constraints in the Working Phase

Married individuals:

If married in period t , a household makes decisions subject to a 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:

$$\hat{c}_t(k_t) + a_{\bar{t}+1}^m = Ra_t^m + (y_{eft} - \kappa(k_t))P_t + y_{emt} \quad (3)$$

where a_t^m is the married household asset holdings entering period t .¹² Income consists of capital income from last period's asset, where R is the gross return, and y_{egt} is each spouse's labor market income from which expenditures on childcare, $\kappa(k_t)$, are subtracted if the wife works. \hat{c} is the expenditure on consumption and $a_{\bar{t}+1}^m$ is the household's asset position before the realization of the marital shock at the end of period t . Henceforth, we use \bar{t} to denote the value of a variable before the realizations of the end-of-period shocks.

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 $a_{\bar{t}+1}^m$ are allocated to the wife and the remainder to the husband. Thus, the laws of motion

¹¹See Del Boca and Flinn (2010) for an insightful analysis of various modeling approaches one may use to determine household allocations.

¹²Note that we are assuming, as in Mazzocco, Ruiz, and Yamaguchi (2007) that married agents are not allowed to have their own private savings or borrowing. This seems like the more realistic alternative.

for $a_{g,t+1}^s$ are given by:

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

Divorced individuals:

We assume that if a couple divorces, the ex-husband makes transfer payments to his ex-wife in subsequent periods if they have children under the age of twenty and she is not remarried. In particular, the former husband must pay a fraction of his current income to his ex-wife as child support $h(k_t, y_{mt})$. Recalling that the children reside with their mother, the budget constraints of a divorced woman and a divorced man at time t are thus given by:

$$\text{Woman: } \hat{c}_t(k_t) + a_{f,\bar{t}+1}^{\mathbf{d}} = Ra_{ft}^{\mathbf{d}} + (y_{ft} - \kappa(k_t))P_t + h(k_t, y_{mt}) \quad (4)$$

$$\text{Man: } c_t + a_{m,\bar{t}+1}^{\mathbf{d}} = Ra_{mt}^{\mathbf{d}} + y_{mt} - h(k_t, y_{mt}) \quad (5)$$

where $a_{\bar{t}+1}^{\mathbf{d}}$ 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. 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 $t+1$:¹³

$$a_{i,t+1}^s = \begin{cases} a_{i,t+1}^{\mathbf{d}} = a_{i,\bar{t}+1}^{\mathbf{d}} & \text{if } s_{t+1} = \mathbf{d} \text{ (i.e., if } i \text{ enters } t+1 \text{ still divorced)} \\ a_{t+1}^{\mathbf{m}} = a_{i,\bar{t}+1}^{\mathbf{d}} + a_{j,\bar{t}+1}^{\mathbf{d}} & \text{if } s_{t+1} = \mathbf{m} \text{ (i.e., if } i \text{ enters } t+1 \text{ remarried to } j) \end{cases} \quad (6)$$

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 child care). Thus, for a single

¹³Recall that a divorced agent is marrying another divorced agent by assumption, hence i and j both have \mathbf{d} superscripts in equation (6)

woman and a single man:

$$\text{Woman: } \hat{c}_t(k_t) + a_{t+1}^s = Ra_t^s + y_{ft} - \kappa(k_t) \quad (7)$$

$$\text{Man: } c_t + a_{t+1}^s = Ra_t^s + y_{mt} \quad (8)$$

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.,¹⁴

$$a_{i,t+1}^s = \begin{cases} a_{i,t+1}^s = a_{i,\bar{t}+1}^s & \text{if } s_{t+1} = \mathbf{s} \text{ (i.e., if } i \text{ enters } t+1 \text{ still single)} \\ a_{t+1}^{\mathbf{m}} = a_{i,\bar{t}+1}^s + a_{j,\bar{t}+1}^s & \text{if } s_{t+1} = \mathbf{m} \text{ (i.e., if } i \text{ enters } t+1 \text{ married to } j) \end{cases} \quad (9)$$

3.3 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 . Recall that we assume that individuals cannot die with debt. Thus the budget constraints are given by:

$$\hat{c}_t + a_{t+1}^s = Ra_t^s + b_g^s \quad \text{and} \quad a_{T+1}^s \geq 0 \quad (10)$$

3.4 Optimization Problems

In this section we outline each household type's optimization problem. Before doing so we introduce individual i 's state vector $\Omega_{it} = \{a_{it}, x_{i,t-1}, P_{i,t-1}, k_t, z_{it}, e_i, \zeta_{it}\}$ which 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 persistent component of the income shock, an individual's education, and one's marital type, ζ_{it} . As discussed previously, in addition to an agent's potential (or actual or ex) spouse's education, ζ_{it} also tracks the spouse's asset holdings, experience, participation,

¹⁴We assume that singles marry other singles.

and persistent income shock, all at period t . Finally, define the law of motion for experience as:

$$x_{t+1} = x_t + P_t \quad (11)$$

Divorced Agent's Maximization Problem

Upon retirement, the optimization problem is simple given that there is no uncertainty. For an individual of gender g it is given by:

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

s.t. retired divorced budget constraint for g in eq. (10)

During the work phase of life, the divorcee has an exogenous eg -specific probability of remarrying r_{eg} . Upon remarriage, the maximization problem transits 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 to work.¹⁵ Thus, the value of being a divorced woman at time t is given by:

$$V_{ft}^{\mathbf{d}}(\Omega_t) = \max_{c_t, P_t, a_{t+1}^{\mathbf{d}}} \frac{(c_t)^{1-\sigma}}{1-\sigma} - \psi_e^{\mathbf{d}}(k_t)P_t \quad (12)$$

$$+ \beta \left\{ (1 - r_{eg,t}) \mathbb{E} \left[V_{f,t+1}^{\mathbf{d}}(\Omega_{t+1} | \Omega_t) \right] + r_{eg,t} \mathbb{E} \left[V_{f,t+1}^{\mathbf{m}}(\Omega_{t+1} | \Omega_t) \right] \right\}$$

s.t. budget constraint (4) and laws of motion (6) and (11)

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.

Similarly, the value of being a divorced man at time t is:

$$V_{mt}^{\mathbf{d}}(\Omega_t) = \max_{c_t, a_{t+1}^{\mathbf{d}}} \frac{(c_t)^{1-\sigma}}{1-\sigma} \quad (13)$$

$$+ \beta \left\{ (1 - r_{eg,t}) \mathbb{E} \left[V_{m,t+1}^{\mathbf{d}}(\Omega_{t+1} | \Omega_t) \right] + r_{eg,t} \mathbb{E} \left[V_{m,t+1}^{\mathbf{m}}(\Omega_{t+1} | \Omega_t) \right] \right\}$$

s.t. budget constraint (5) and laws of motion (6) and (11)

¹⁵Note 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. This is true for all other maximization problems that involve $t = t^R - 1$

Married Household's Maximization Problem

A couple that enters the retirement period married solves

$$V_t^{\mathbf{m}}(\Omega_t) = \max_{c_t, a_{t+1}^{\mathbf{m}}} \frac{c_t^{1-\sigma}}{1-\sigma} + \beta V_{t+1}^{\mathbf{m}}(\Omega_{t+1} | \Omega_t) \quad (14)$$

s.t. retired married budget constraint (10) and asset law of motion (3.2)

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

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

$$\begin{aligned} V_t^{\mathbf{m}}(\Omega_t) = \max_{c_t, P_t, a_{t+1}^{\mathbf{m}}} & \chi \left[\frac{c_t^{1-\sigma}}{1-\sigma} - \psi_e^{\mathbf{m}}(k_t) P_t \right] + (1-\chi) \left[\frac{c_t^{1-\sigma}}{1-\sigma} \right] + (1-d_{eg,t})\beta \mathbb{E} [V_{t+1}^{\mathbf{m}}(\Omega_{t+1} | \Omega_t)] \\ & + d_{eg,t}\beta \left\{ \chi \mathbb{E} [V_{f,t+1}^{\mathbf{d}}(\Omega_{f,t+1} | \Omega_t)] + (1-\chi) \mathbb{E} [V_{m,t+1}^{\mathbf{d}}(\Omega_{m,t+1} | \Omega_t)] \right\} \end{aligned} \quad (15)$$

s.t. budget constraint (3) and laws of motion (6) and (11)

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.¹⁶ Using an asterisk to denote the resulting outcomes from the optimization problem above, the value of being married for a man is given by:

$$\begin{aligned} V_{mt}^{\mathbf{m}}(\Omega_t) = & \frac{c_t^{*1-\sigma}}{1-\sigma} \\ & + (1-d_{eg,t})\beta \mathbb{E}(V_{m,t+1}^{\mathbf{m}}(\Omega_{t+1}^* | \Omega_t) + d_{eg,t}\beta \left\{ \mathbb{E} [V_{m,t+1}^{\mathbf{d}}(\Omega_{m,t+1}^* | \Omega_t)] \right\} \end{aligned} \quad (16)$$

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

$$\begin{aligned} V_{ft}^{\mathbf{m}}(\Omega_t) = & \frac{c_t^{*1-\sigma}}{1-\sigma} - \psi_e^{\mathbf{m}}(k_t) P_t^* \\ & + (1-d_{eg,t})\beta \mathbb{E} [V_{f,t+1}^{\mathbf{m}}(\Omega_{t+1}^* | \Omega_t)] + d_{eg,t}\beta \left\{ \mathbb{E} [V_{f,t+1}^{\mathbf{d}}(\Omega_{f,t+1}^* | \Omega_t)] \right\} \end{aligned} \quad (17)$$

¹⁶Note 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.

Single Agent's Maximization Problem

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

$$V_{gt}^s(\Omega_t) = \max_{c_t, a_{t+1}^s} \frac{c_t^{1-\sigma}}{1-\sigma} + \beta V_{t+1}^s(\Omega_{t+1} | \Omega_t) \text{ for } t \geq t^R \quad (18)$$

s.t. retired single budget constraint (10) and asset law of motion (9)

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

$$V_{gt}^s(\Omega_t) = \max_{c_t, P_t, a_{t+1}^s} \frac{c_t^{1-\sigma}}{1-\sigma} - \psi_{eg}^s(k_t) P_t \quad (19)$$

$$+ (1 - m_{eg,t}) \beta \mathbb{E} [V_{t+1}^s(\Omega_{t+1} | \Omega_t)] + m_{eg,t} \beta \mathbb{E} [V_{g,t+1}^m(\Omega_{f,t+1} | \Omega_t)]$$

s.t. budget constraint (7) or (8) and laws of motion (9) and (11)

4 Parametrization

In this section we describe the calibration of the model.¹⁷ Some model parameter values are taken from preexisting estimates and others are estimated directly from the data using model restrictions. A 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 1 reports the parameters estimated “outside” the model and Tables 2, 3, and 4 report the “internally” calibrated parameters and their targets. 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, child-care costs and several parameters which affect wage dynamics.

To construct our key statistics we mainly use the 1962-2010 waves of the March supplement of the Current Population Survey (CPS), a 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:

¹⁷In order to capture the fact that men and women marry outside their cohorts, we solve the model twice, once for the marriage market probabilities faced by women and a second time for the probabilities faced by men. See Appendix for further details.

the “1935” cohort consists of white men and women born between 1934-1936, while the “1955” are those born between 1954-1956.¹⁸ 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. 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).

To parametrize the utility function we set $\sigma = 1.5$ since most estimates for the relative risk aversion parameter in the literature vary between one and two. 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 and are calibrated within the model in order to match female LFP rates by marital status and education for the 1935 cohort.¹⁹ 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.²⁰ 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”). Single women’s and all men’s disutility from work is normalized to zero.

¹⁸As explained previously, we choose to focus on whites as the historical experiences of black men and women have been very different. In particular, black women have worked significantly more than white women throughout. 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.

¹⁹The 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.

²⁰Mothers of young children may be especially reluctant to work. See, for example, Bernal (2008) and Bernal and Keane (2009) for some evidence regarding the effect of a mother’s working on a child’s development.

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:

$$\ln y_{eg}(, x_t, z_t, P_{t-1}) = \tau_{eg,t} + \gamma_{eg1}x_t + \gamma_{eg2}x_t^2 - \delta(1 - P_{t-1}) + w_{et} \quad (20)$$

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. 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 of wages w_{et} is assumed to be the sum of a (observable to the agent) persistent component (z_{et}) modeled as an AR(1) process and a transitory component $\eta_{et} \sim N(0, \sigma_{\eta_e})$.

$$w_{et} = z_{et} + \eta_{et}, \quad z_{et} = \rho_e z_{e,t-1} + \epsilon_{et}, \quad \epsilon_{et} \sim N(0, \sigma_{\epsilon_e}) \quad (21)$$

This choice of model for the stochastic process is standard in the literature and 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.²¹

We estimate the parameters of the income processing using the PSID. First, we construct data on hourly wages ($y_{em,t}$) for men of education level e using data on earnings and total hours worked.²² 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:

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

²¹This 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.

²²See the Appendix for details of the sample selection.

where \mathbb{D}_{emt} is a set of year dummies. 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).²³ Finally, the time-varying intercepts $\tau_{em,t}$ are given by the estimates of \mathbb{D}_{emt} . These are shown in figure A.1 in the Appendix.

We assume that women share 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. Instead, the latter are calibrated internally so as to generate 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 A.1 in the Appendix. 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.²⁴

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.²⁵ 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,

²³We thank Gianluca Violante for kindly providing us with the estimation code.

²⁴For 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.

²⁵For example, Voena (2012) 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.

however, is fairly high at 37%. Beller and Graham (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.²⁶

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 bend points (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.²⁷

Family Formation and Fertility

Agents’ initial marital status is assigned in proportions to match the ones in the CPS data for our cohorts at age 25-29. These proportions are given in Figure A.2 in the Appendix, and they vary by gender, education, and cohort.

Recall that agents are assigned a marital “type” which permanently determines the characteristics of one’s spouse. 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 individuals between 35-39 for each respective cohort. These proportions, conditional on education and gender, are reported in Table 5. Thus, the pattern of marriages mimics the degree of assortativeness in education found in the data.

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

²⁶In 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.

²⁷The 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.

data for each cohort. These probabilities are reported in Table A.1, by age, education and gender. Next we need divorce and remarriage rates which we compute using the Survey of Income and Program Participation (SIPP) conducted by the Census.²⁸ These rates are reported in Table 6 and the details of the computation are in the Appendix.

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.²⁹ The exact timing of shocks by cohort and education are given in the Appendix.

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. We use the McClements scale to calculate the economies of scale in consumption.³⁰ Its exact numbers (by child's age) are reported in the Appendix.

Women who have children under the age of 10 at home are assumed to incur child-care costs if they work which depend only on the age of the youngest child, i.e. if a household has a young and an old child, child-care costs are incurred once and for the younger child. We calibrate the child-care 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 2.

Other External Parameters

There is little guidance as to what the Pareto weight on a woman's welfare in the household allocation problem should be. A recent paper by Voena (2012), using variations in savings behavior and divorce laws, estimates a value of 0.25. In our benchmark calibration we set

²⁸The SIPP is a series of short panels (from 2.5 to 4 years) of approximately 14,000 to 36,700 households in the US.

²⁹We 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)).

³⁰This scale is very similar to the OECD scale, but it has the advantage that it was computed based on expenditure data from families.

this value, χ , to 0.3; the robustness section investigates the effect of changing this weight and finds small effects.

Upon divorce, we assume that the woman is responsible for the children and that assets are split equally between the ex-spouses (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. The robustness section explores values of $\alpha = 0.3$ and $\alpha = 0.7$ and finds 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.

4.1 Solution Method

The households face a known finite horizon which implies that the dynamic problem can be solved numerically by backwards recursion from the last period of life using value-function iteration. At each age, the households solve for their consumption-savings rule and LFP decisions taking as given their state variables that period and next period's value function. In addition to household assets, our model has two other potentially continuous state variables: the husband's and wife's persistent components of earnings, w_{met} and w_{fet} , respectively. As including more than one continuous state variable, while possible, is computationally costly, we choose to discretize these two variables, leaving assets as the only continuous state.

During the working stages of the life cycle, the model combines a discrete decision (whether the woman participates in the market) and a continuous decision (the level of savings). This combination may lead to non-concavities in the value function. Given enough uncertainty, the value function conditional on today's participation is concave; we numerically check that this holds for our model. We follow Attanasio, Low, and Sanchez-Marcos (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 (working versus not working) intersect only once. This is where the woman’s participation decision switches from not working to working. Thus, conditional on all other states, for all values $a_t < a_t^*$ the woman works and for all values $a_t > a_t^*$ the woman does not work. We numerically check both the concavity of the conditional value functions and the uniqueness of the reservation asset level. We then solve for each household’s optimal level of asset accumulation, conditional on the LFP decision, i.e., we solve for $a_{t+1}^{lfp} = \text{argmax}_{a_{t+1}} V^{lfp}(\Omega_t)$, $lfp = \{work, not\ work\}$. The Appendix presents further details about the solution, including number of grid points and the algorithm used to calculate the optimal level of assets.

To summarize, there are a total of 27 parameters calibrated internally: 8 parameters which govern the disutility of labor for married and divorced women, 2 parameters for child-care costs, 14 time-varying wage-intercept parameters for women and 2 for men and 1 wage depreciation parameter. We choose to match a total of 43 statistics for our 1935 cohort: 28 average LFP rates for married and divorced women, by age and education, 7 time-varying gender wage ratios, 7 time-varying skill premia for women, and the average lifetime skill premium for men.³¹ Although this mapping is only approximate, it may be useful to think of the time-varying wage-intercept parameters as targeting the skill premia and gender wage ratio statistics. The remaining parameters are mainly used to match the life-cycle pattern of LFP by education and marital status.

We calibrate the internal parameters by minimizing the distance between implied model moments and their data targets. The total distance between moments and targets is computed as a weighted average of the squared difference between each moment and target. The weights are such that the 28 moments for LFP receive half of the total weight and the 15 moments for wages receive the other half of the weight. Within each group, the moments are equally weighted.

³¹Note that we only target one skill premium statistic for men using $\tau_{em,1,e} = \{h,l\}$, since the remainder values are calculated from first differences in the data on male wages.

5 The 1935 Benchmark

As seen in Figure 2, the benchmark model does an excellent job of reproducing the LFP profiles for both married and divorced women according to their education. Furthermore, as shown in Figure 2 the model does a good job of matching the period-by-period gender wage ratio and skill premia for women. The LFP and wage targets as well as the model predictions are reported in Table 3.

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 cost of working when divorced is slightly lower than when married. The highest cost is borne by working mothers with children. The ratios of costs across categories is similar for both education groups. One way to think about the differences in these numbers is that they also capture differences in the average attractiveness of jobs available to women, by education, as well as differences in beliefs about the psychic or developmental costs to children from having a working mother.³²

To make economic sense of the work disutility numbers we can calculate their equivalent consumption cost. To do this we compute the decrease in average consumption that women would be willing to bear over one period in order to avoid incurring the disutility cost of working that period, i.e., we find the z such that

$$u(z\bar{c}_{fe}^m) = u(\bar{c}_{fe}^m) - \psi_{es}(k) \quad (22)$$

where \bar{c}_{fe}^m is the average per-period consumption of married women (calculated over their lifetime) with education level e .³³ These proportions (z) are reported in 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 women have very low LFP rates, a large disutility number is required in order to “explain” this. Note

³²Fernández (forthcoming) provides a simple model of married women’s LFP in which women’s beliefs about the cost of working evolve endogenously over time through a learning process.

³³To calculate \bar{c}_{fe}^M we find the average consumption of married women of education e in each period (with or without children) 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.

from figure 2 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 since these have, on average, higher-earning husbands.

The internally calibrated child-care costs for a young child correspond to around 68.9% of average per-period female wages; the equivalent for an older child is 24.0% of average per-period female wages.³⁴ By way of comparison, Attanasio, Low, and Sanchez-Marcos (2008) find child-care costs to be 66% of a woman's mean earnings in their calibrated model.³⁵ Thus, the two numbers are remarkably close.

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. First, although data limitations do not allow us to target asset levels for our two cohorts, we find that the ratio of average wealth to average earnings generated by the model is 3.87, which is close to the value of 3.94 for the US found by Heathcote, Storesletten, and Violante (2010).³⁶

Second, recall that we did not directly target the LFP of women with children and thus a comparison with the data is informative. Using the 1960 and 1970 Census (the CPS does not report the presence of young children in the household for our time period), we observe the 1935 cohort at the ages of 25 and 35, respectively, and obtain an average LFP rate across those years of 28.5% for married and divorced women with children under the age of 10.³⁷ The analogous statistic in our model would be the LFP of women with children under 10 during the first three periods of life.³⁸ This yields an LFP of 27.0%, i.e. very close to the data. The LFP for married and divorced women without children under 10 during this same time period is 63.8% in the data versus 64.8% in the model.

³⁴The wages used to create the average are those from age 25-45.

³⁵This number is reported in terms of the mean earnings of a thirty-year-old woman who worked continuously prior to childbirth.

³⁶This value is computed as the average of period-by-period ratios between average observed assets and labor income in the model, across all household types.

³⁷The 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.

³⁸After the first three periods, given the structure of fertility shocks in the model, all children would be older than ten.

Third, the quantitative model allows one to evaluate the welfare cost of divorce by gender and education. Given that divorced women retain sole custody of the children, they suffer a large drop in consumption upon the dissolution of marriage as child-support and work do not, on average, compensate for the loss of the ex-husband's income. To illustrate the cost of divorce we can compare the average consumption of women who are married in both periods one and two, say, with women who are married only in period one (ignoring the remaining periods). The average consumption of divorced women in period 2 is 18.8% lower than their married counterpart's for college women and 9.3% lower for high-school women.

An alternative illustration that takes into account as well the fact that divorced women also work more is to calculate the proportion of average consumption a married woman would be willing to sacrifice in period 2 in order to remain in her married state that period. This is the z_e that solves

$$u(z_e \bar{c}_e^{\mathbf{m}}) - \bar{\psi}_e^{\mathbf{m}}(k) = u(\bar{c}_e^{\mathbf{d}}) - \bar{\psi}_e^{\mathbf{d}}(k), \quad (23)$$

where $\bar{\psi}_e^s(k)$ denotes the average disutility from work decisions of women of education e in marital status $s \in \{\mathbf{m}, \mathbf{d}\}$ and with children of ages k . The average is thus the weighted sum of zeroes for the women who don't work plus the weighted disutilities (varying by children's age) for women who do work. \bar{c}_e^s is the average consumption in period 2 of a woman with education e and marital state s . This alternative calculation yields a consumption loss of 24.6% for a college women and 19.3% for her high-school counterpart. The smaller gap in consumption loss between education groups comes from their differences in labor disutility (the latter is higher for high-school women).

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 equivalent one for 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 and not from a change in household income.

Although the empirical literature in this area does not have numbers that allow a direct comparison of consumption costs, Peterson (1996) provides a survey of the literature in this field. Using data from 1976-'77, he finds 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 are to make women worse off and men better off.

6 From 1935 to 1955

In this section we investigate the consequences of modifying the 1935 environment so that it mimics the 1955 one. We proceed by first changing the assortativeness of marriage to match that in 1955. We then change the other features of family structure to match that in 1955. Since the impact of various changes are non-linear, we return to the 1935 family structure and impose various features of the 1955 wage structure. Lastly, we examine at the impact of the combination of changing the entire family and wage environment to that of 1955.

6.1 1955 Family Structure

1955 Assortativeness

The first change we make, and keep throughout all the subsequent exercises, is to impose the 1955 degree of marital sorting by changing the conditional probabilities with which women and men marry each other as a function of own education (see Table 5).³⁹ This corresponds to a change in the correlation coefficient in education from 0.650 and 0.624 for women and men of the 1935 cohort, respectively to 0.654 and 0.641 for women and men of the 1955 cohort, respectively. Although the literature (e.g., Weiss (1997), Ge (2011), Iyigun and Walsh (2007)) has found this to be an important factor in explaining the rise in women's college education, as shown in Figure 3 this factor is not quantitatively significant

³⁹We include this change throughout to incorporate the fact that household incomes have changed for both college and high school women since they are both more likely to have a college spouse.

for married women’s LFP. In fact, the increased assortativeness reduces slightly the LFP of married women, particularly that of college women. This is not surprising as the increase in proportion of college-educated men means that women in both education categories have a higher probability of marrying one and thus having higher household income.

1955 Divorce Rates

Next, in order to gauge the importance of the markedly higher divorce rate, we substitute the 1935 profile of divorce rates with the one faced by the 1955 cohort. In this exercise we hold all other parameters constant at their 1935 levels except for the change in marital assortativeness above. Thus, the initial distribution of individuals at age 25 by marital status, the per period marriage and remarriage rates, and fertility is held at 1935 levels. Note from table 6 that the divorce probabilities faced by the 1955 cohort are, on average, higher for both genders and education groups. Also, the age profile of divorce changes across cohorts: 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.

As can be seen in Figure 3, in response to the higher divorce probabilities, married women’s LFP dramatically increases in the first two to three periods.⁴⁰ The average magnitude of the LFP increase for college women over their first two periods of working life is of around 30 percentage points while that of high school women is around 25 percentage points.

1955 Marriage Market, Fertility, and Child-Care Costs

Next, we evaluate the effects of the remaining changes in the marriage market. Using the 1935 benchmark model, we change not only the assortativeness of marriage and the divorce probabilities, but also the marriage and remarriage probabilities as well as the initial distribution of marital states at age 25, all so as to match the 1955 marriage market. As can be seen from Figure A.2 in the Appendix, between the two cohorts there was a significant decline in the proportion of individuals who married by age 25-29. There was also an increase in the proportion divorced by that age. Furthermore, as shown by table A.1,

⁴⁰The LFP of divorced women (see line with squares in Figure 5) also increases but not as dramatically.

the probability of transitioning from single to married changed such that these probabilities look more similar across genders (by education).

The changes in female LFP obtained from incorporating all changes in the marriage market are shown in Figure 4 (given by the line with squares). As can be seen, they look very similar to the one obtained when we changed only the divorce probabilities, though there is also an increase in LFP at later ages now as well.

An additional change that occurred during this time period was a decline in fertility and, potentially, a decrease in child-care cost. We incorporate these changes, in addition to the changes in the marriage market, by allowing fertility patterns to change as in the data and by assuming that child-care costs decreased by 20%. Attanasio, Low, and Sanchez-Marcos (2008) investigate a decrease of 1-5% in child-care costs for the cohorts born in 1945 and 1955. Since our cohorts are further apart, we examine a 20% decline in these costs.

The additional changes in LFP resulting from the lower fertility and child-care costs are shown in Figure 4 (the line with triangles). The LFP reaction of married high-school women differs from that of married college women. The former decreases slightly when younger whereas the LFP of married college women increases, and both increase thereafter (all relative to the 1955 divorce only case). This difference in reaction is mainly due to the more significant drop in the total number of children for high-school women (from 3.2 children per woman to 2.1), which *ceteris paribus* increases their household consumption and thus dampens their incentive to work.

Attanasio, Low, and Sanchez-Marcos (2008) found that 50% of the increase in married women's LFP between the 1945 cohort and the 1955 cohort could be explained by the decline in child-care costs. To contrast our findings with theirs we examine the effect of a decrease in child-care costs on its own with no other change in family structure (other than the marital sorting). The line with circles in Figures 4 and 5 shows the impact of 20% lower child-care costs on the 1935 environment.⁴¹ As can be seen from the figures (the line with small circles), a lower cost of child-care increases female LFP, for both married and divorced women. The increase in LFP for married women is significant (but not as large as found in ? it accounts for about 22% of married women's LFP increase for high-school

⁴¹To make comparisons easier, we have also allowed the degree of assortativeness to change to the 1955 level, as in the prior exercises.

women and 12% for college women during the ages of 25-44.

To summarize, changes in the marriage market, fertility and child-care costs account for a significant proportion of the change in LFP across the two cohorts of married women, especially for those with a college education. These changes account for 72% of the average LFP gap across the two cohorts of college married women from age 25 to 44 and for 67% of the same LFP gap across the entire working life. The corresponding numbers for high-school women are 63% (ages 25-44) and 58% (ages 25-59). Interestingly, they generate too much of an increase in the LFP of divorced high-school women (see Figure 5).

6.2 1955 Wage Structure

Next we quantify the contribution of changes in the wage structure independently of any changes that occurred in family structure (other than marital sorting). As is commonly recognized, two key changes in wages took place over this period: (i) the skill premium increased for both men and women, and (ii) the ratio of female to male wages increased (the gender wage ratio). Averaging over the working lifetime of the two cohorts, the skill premium increased by 0.12 for both men and women. Over the same period the average gender wage ratio went from 0.61 to 0.72.

Quantifying the effect of the changes, given women's selection into work, requires changing the parameters which govern wages, in particular the $\tau_{eg,t}$ which discipline their time path. Recall that the sequence of $\{\tau_{em,t}\}_{t=2,\dots,t^R}$ for men were calculated directly from wage data for each cohort (the year dummies). The ones that correspond to the 1955 cohort will now be used for this experiment. The sequences of $\{\tau_{ef,t}\}_{t=1,\dots,t^R}$ for women are calibrated internally in order to match certain wage statistics as described below. The parameters for the stochastic process and the returns to experience for men (which were estimated from the data as pure age/experience effects, after controlling for year effects), on the other hand, are left at their 1935 values. A subsequent exercise examines the consequences of changing women's returns to experience.

1955 Skill Premia and the Gender Wage Ratio

To examine the effects of the 1955 wage gender ratio and skill premia for both genders, we calibrate the sequences of $\{\tau_{ef,t}\}_{t=1,\dots,T}$ (i.e., the wage intercepts for women) to match, period-by-period, the gender wage ratio and the female skill premium.⁴²

As shown in Figure 6 in the graph with the caption “Skill Premium + Gender Gap”, these changes in wage structure result in significant LFP changes for both college women (a 10 percentage point average increase across the life cycle) and high-school women (12 percentage points on average). 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.

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 24 for a discussion), an additional exercise is to compound the changes in the wage structure above with a higher returns to women’s labor market experience. We do this by changing the parameters $\gamma_{ef1}, \gamma_{ef2}$ to 0.1061 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.⁴³

As seen in Figure 6 in the graph labelled “All Wages”, the higher returns to experience results in married women working more when young, with a more prolonged effect for high-school women. It has basically no effect on married women’s LFP once they are older.

To summarize, combining all the changes in wage structure we find that these changes account for 58% of the average LFP gap across the two cohorts of college married women across the entire working life. The corresponding number for high-school women is 51% of the average LFP gap. For divorced women, the model predicts an increase which is larger than what is seen in the data, as can be seen in Figure 7 (line with small circles).

⁴²Note 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 whether one should eliminate all level effects as this would imply that the disutility of labor is proportional to income.

⁴³In the robustness checks we experimented with 5% returns to experience with very similar results.

6.3 1955 Wage and Family Structures Combined

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 6 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, but not by a large amount. When the changes in family structure are introduced after those in wage structure, as done here, the former impacts married women’s LFP mostly during their fertility years. The changes in wage and family structures combined can account for 93% of the average LFP gap between the two cohorts during the ages of 25-44 for married college women and for more than the entire gap between those same ages for high school women. Across the life cycle, those changes account for over 90% of the average LFP gap for married college women and over three quarters for married high school women. The effect of these changes for divorced women is shown in Figure 7 where it can be seen that the predicted response has a greater proportion of divorced women working than in the data.

7 Understanding the Role of Divorce Risk

As shown previously, the profile of higher divorce rates faced by the 1955 cohort was an important factor in the increase in married women’s LFP. As earlier work by Eckstein and Lifshitz (2011) found that the higher marital instability played a very small role in this increase, it is particularly important to investigate the factors that give rise to our results. In this section we investigate the mechanisms by which higher divorce risk delivers higher LFP and the role of concavity in the quantitative analysis.

7.1 The Role of Experience and Asset Accumulation

Why does a higher divorce rate increase married women’s LFP? A plausible reasoning would be that, since greater divorce risk implies a higher probability of being in a low consumption state (i.e., becoming a divorced mother with children), married women face a greater incentive to work in order to potentially increase their consumption should they

divorce. The channels by which working while married increases consumption when divorced are twofold. First, working increases a woman’s labor market experience and thus also her future earnings. This allows her to have higher consumption should she divorce and need to work (recall that the LFP of divorced women is, endogenously, very high). Second, working while married allows married couples to increase their savings. If the married couple subsequently divorces, the larger asset transfer would allow the divorced woman to increase her consumption (recall that, upon divorce, marital assets are split fifty-fifty).

We now turn to examining the validity and quantitative importance of each of these two channels. First, to gauge how much labor market experience matters, we can have women face an exogenous schedule of wage increases over time that mimics perfectly the increase in wages that come from the return to experience. This is equivalent to obtaining an additional return to age. In this way, married women’s LFP choices are rendered independent of their desire to have a higher return to working in the divorced state since this return is assured independently of their prior LFP choices.

Column 1 in table 7 reports the LFP of married high school (panel A) and college women (panel B) over their working lives under the assumption that the calibrated 2% increase in annual labor income is exogenous, i.e., independent of a women’s work history. All other parameters are the same as in the benchmark 1935 model except for assortativeness. We use the 1955 degree of assortativeness throughout, just as in the prior section, to allow for easier comparisons. Column 2 reports the LFP results from running the 1955 divorce experiment in this new setting. Column 3 reports the percentage point *change* in LFP from divorce (i.e., column 2 minus column 1). Column 4 reports the percentage point *change* in married women’s LFP from the 1955 divorce experiment but with *endogenous* experience (i.e., it replicates the original 1955 divorce experiment of section 6.1).

We are now set to evaluate the role played by the desire to accumulate experience in generating the results associated with marital instability. As can be easily seen from comparing columns 3 and 4, the desire to gain labor market experience played basically no role in increasing married women’s LFP. Had the return to experience been exogenous, i.e., solely a function of age, there would have been virtually the same changes in LFP in

response to higher divorce risk as when this return is endogenous. Overall, if experience was exogenous, married college women’s LFP would have averaged 0.54 percentage points less over their working lives; high-school women would have had on average a 0.03 percentage points smaller increase in LFP.

Next we analyze the role played by women’s desire to accumulate more marital assets. Table 8 shows the evolution of assets over the lifetime of married high-school (panel A) and college (panel B) women. Column 1 reports the evolution of assets for the benchmark model (with 1955 assortativeness). Column 2 shows how assets react to introducing the 1955 divorce profile. Note that the higher divorce risk in fact induces households to save *less* in the first 4 periods (recall that divorce risk, by assumption, is present only in these periods), rather than more.⁴⁴ Why is this? To understand this outcome, it is important to recall that, for a given divorce risk, men and women have different incentives to save. Since, *ceteris paribus*, men are better off divorced whereas women are worse off, consumption smoothing implies women would prefer to shift more consumption to the divorced state whereas men would prefer to shift more consumption to the married state. The effect on savings of an increase in divorce risk is thus to increase married women’s desire to save and to decrease married men’s desire to do so. The ultimate effect on household saving depends on each party’s bargaining power in the household (i.e., on their Pareto weight). With the benchmark weights, men have greater decision-making power and succeed in reducing savings (i.e., they borrow even more than before the divorce rate increase).⁴⁵

To summarize, the effect of the increase in divorce risk is to decrease the savings of married couples and to increase married women’s LFP. Both changes *increase* the consumption of women in their married state without increasing their ability to consume in the divorced state (modulo the experience effect which is small). Hence, the explanation originally proposed for why increased divorce risk leads to higher married women’s LFP is

⁴⁴This finding is in agreement with the empirical evidence presented in Mazzocco, Ruiz, and Yamaguchi (2007). The authors document, using the PSID, that couples who divorce have significantly lower assets the year prior to the divorce than couples who continue married. Of course, selection on various characteristics may be an important driver of this fact.

⁴⁵Note that household savings drop for all couples. In particular, considering only couples in which i. the wife did not work in the 1935 environment but, ii. faced with same wage and fertility shocks, the wife work when faced with 1955 divorce probabilities, we find that household savings drop for these couples, though naturally by less than for those married households in which women did not work in either environment.

wrong. Why then are married women working more? The answer is that married women work more because, if they didn't, the decline in married household savings would be even larger, leading to even lower consumption for a woman in the divorced state.⁴⁶ Working while married allows the increase in marital consumption (desired by the husband) to be achieved in part by increasing married household income rather than solely by decreasing married household savings.

It is useful to understand in greater depth the role played by the Pareto weights in the analysis above. To do so, we examine how the responses of savings and married women's LFP to the increase in divorce risk vary with the Pareto weight (χ) on a woman's welfare. For ease of comparison, Column 3 of Table 8 reports the net *change* in savings resulting from the 1955 divorce profile in an otherwise 1935 environment with the original Pareto weight of $\chi = 0.3$ (i.e., the difference between column 2 and column 1). Column 4 reports the net change in savings from the same divorce experiment but for the case of $\chi = 0.5$; column 5 does the same for $\chi = 0.7$.⁴⁷ As can be seen clearly from this table, the greater the bargaining strength of married women, the greater is the amount by which saving increases in periods 1 through 4 as a reaction to the higher divorce rate.

It is important to note that, as χ increases, the increase in household savings is not accompanied by a decrease in the reaction of married women's LFP. On the contrary, returning to Table 7, column 5 shows the percentage point *change* in married women's LFP given the increase in divorce rates for $\chi = 0.5$; column 6 does the same for $\chi = 0.7$. Thus, as shown, the greater is women's bargaining strength, the greater is the reaction of married women's LFP on average and of savings. This reflects women's desire to shift consumption to the divorced state: working serves both to increase savings and to optimally balance (given the Pareto weights) the opposing desires of wives and husbands regarding the distribution of consumption over different states. If married women's LFP did not increase, the increase in household savings would be smaller (or even negative as for our

⁴⁶Indeed, if we examine the subset of women who did not work in either 1935 or when subject to the 1955 divorce risk, household savings falls substantially more than for couples in which the wife worked given the 1955 divorce risk.

⁴⁷Note that in each case the change in savings is obtained using the new Pareto weight with both the 1935 divorce profile and the 1955 divorce profile. The Pareto weight in each experiment, thus, does not change across cohorts; it only changes across experiment.

benchmark case of $\chi = 0.3$).

7.2 The Role of Concavity

As noted in the literature review, Eckstein and Lifshitz (2011) concluded that the increase in divorce risk played only a small quantitative role in increasing married women’s LFP. While there are several differences across the two models, the role played by a concave utility function in our environment as opposed to their linear utility specification is a key difference to investigate. We turn to an analysis of the role of concavity in generating the results.

To examine concavity’s role, we assume preferences are almost linear by setting $\gamma = 0.001$ and proceed to *recalibrate* the model.⁴⁸ Figure 8 shows the result of this recalibration for the LFP of married and divorced women, by education. As can be seen in this figure, the calibrated model does very well in reproducing the life-cycle path of LFP for all women (see table A.5). Tables 9 and 10 present the new values of the internally calibrated parameters. Note that although the parameters for the disutility from labor are all larger than in the benchmark calibration, their consumption equivalence is remarkably similar as can be seen by comparing Table 4 with Table 9.

Next, we perform the same exercise as in the previous section: we examine the role of the higher 1955 divorce rate by keeping the rest of the environment at its 1935 benchmark values (as before, we also change the degree of assortativeness of marriage to match that in 1955). The result is shown in Figure 9.⁴⁹

Contrasting the effect of 1955 divorce in the linear model with our original concave model, the higher divorce rate increased female LFP in this case as well. There are two interesting differences, however, to point out with respect to the case of concavity. First, the increase in LFP is significantly smaller, especially when married women are younger. Over the first 3 periods of married women’s working lives the increase (over both education categories) averages 3 percentage points per period. In our original benchmark ($\gamma = 1.5$), it

⁴⁸ An almost linear specification was preferred as a fully linear one could potentially exhibit behavior very different from the model with curvature simply due to computational “kinks”.

⁴⁹ Note that the increase in LFP is driven entirely by the increase in divorce rate – the increase in assortativeness on its own has no effect as can be seen separately the figure.

averaged 21 percentage points per period over the same ages. Second, the response is more or less uniform over women’s lifetime as opposed to being concentrated exclusively over the first few periods of life in the concave case. This is due to the increase in women’s wages resulting from greater labor market experience. In the absence of curvature in the utility function, the wage increase induces significantly more women to work in the later periods of their working life. In the concave model, the main driver of women’s increased LFP is the significant decrease in utility due to the potential fall in future consumption if divorced; the additional experience is secondary. With linearity, the potential fall in consumption does not yield the same disutility and experience plays a more important role over the entire life span.

Overall, divorce alone accounts for around 20% of the increase in married women’s LFP over their working lifetime in the linear case relative to over 40% of this increase with concavity. We conclude that concavity plays an important quantitative role in generating a larger impact from the changed divorce probabilities.

8 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 which is used to obtain a solution to the married household’s allocation problem. We chose $\chi = 0.3$, implying that the wife had lower bargaining power than her husband. Another plausible choice would have been $\chi = 0.5$, where husband and wife receive equal weight in the household allocation problem. With this in mind, we recalibrate the benchmark 1935 model using $\chi = 0.5$.

While most wage parameters remain similar (albeit 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 somewhat more muted. As before, the response in LFP

to increased divorce remains concentrated on the 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 percentage points. 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 structure (see end of section 6.2), although, once again, the effects are smaller. These changes now account for around 72% 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 that the increase in women’s wage increased the value of $\chi = 0.41$ in the 1990s. Since we chose $\chi = 0.3$ for both cohorts, this possibly underestimates the bargaining power of women in the 1955 cohort.

We investigate the consequence of increasing $\chi = 0.5$ for the 1955 cohort, keeping the 1935 cohort fixed at $\chi = 0.3$. Note that 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. As compared to the specification with the Pareto weight of $\chi = 0.3$, the combined changes in wage and family structure result in average lifetime LFP of married college women dropping by 12 percentage points (from 71.3% to 60.2%) and their high school counterparts’ participation is 17 percentage points lower (it drops from 63.7% to 46.7%). The LFP 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.

Additional robustness checks included changing the proportion of a man’s income which must be paid in child support and the 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.⁵⁰

9 Concluding Remarks

This paper developed a dynamic quantitative life-cycle model to evaluate how changes in family structure and economic environment contributed to changes in women’s LFP. The model was calibrated to match key statistics of the 1935 cohort. We then proceeded to change key features of the environment in order to mimic the one faced by the 1955 cohort. The model was successful in predicting increased LFP for women, accounting for over 85% of the LFP gap of married women between the ages of 25-59 (though it overpredicted the LFP of divorced women).

We found that both changes in family (i.e., assortativeness, marriage, divorce, and remarriage patterns, fertility, and child-care costs) and wage structure (skill premium, gender wage gap, and returns to women’s labor market experience) are important in explaining the increase in women’s work. Each group of factors in isolation is able to account for a large proportion (about 55%) of the observed changes in married women’s LFP over their working lives. Furthermore, the increased probability of divorce faced by the 1955 cohort is on its own a key driver of the increase in married women’s work. It alone accounts for around 46% percent of the increase in married women LFP over their working lives (age 25-59).

Our analysis of the mechanisms through which a higher divorce probability affected married women’s propensity to work yielded some surprising results. In particular, the analysis showed that the increase in married women’s LFP was not driven by the desire to obtain greater labor market experience and that higher LFP did not translate into a higher asset level for married couples. Instead, our analysis indicates that married women’s higher LFP is the solution to the conflicting preferences of husbands and wives towards how to adjust marital consumption to higher divorce risk. When confronted with higher divorce risk, married women would prefer to increase savings in order to transfer more assets to the

⁵⁰All results from the robustness checks are available upon request.

divorced state, thus helping them smooth consumption. Men, on the other hand, would prefer to increase consumption in the married state as this is what allows them to smooth consumption. A “compromise” is reached – the solution to a Pareto weighted utility maximization problem – that involves married women working more and thus allowing married savings to decline by less than what it would otherwise.

This paper points to marital instability as one of the key drivers in the increase in married women’s LFP. It takes a significant step by incorporating interactions between family structure and work choices in a computationally tractable dynamic life-cycle setting with endogenous consumption, saving, and accumulation of labor market experience. Its findings contrast with others in the literature (Eckstein and Lifshitz (2011)) and hence clarify the importance of risk aversion and household bargaining in obtaining quantitatively significant results for the role of marital instability in affecting married women’s work.

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.
- 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.
- Cubeddu, L. and J.-V. Ríos-Rull (2003). Families as shocks. *Journal of the European Economic Association* 1(2-3), 671–682.
- 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* 79(6), 1675–1726.
- Fernández, R. (2012). Cultural change as learning: The evolution of female labor force participation over a century. *American Economic Review*, forthcoming.
- Fernández, R., A. Fogli, and C. Olivetti (2004). Mothers and sons: Preference formation and female labor force dynamics. *Quarterly Journal of Economics* 119(4).
- Fernández, R. and J. Wong (2011). The disappearing gender gap: The impact of divorce, wages, and preferences on education choices and women’s work. NBER Working Paper 17508.
- 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* 29, 773–818.
- 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.
- Gray, J. S. (1998). Divorce-law changes, household bargaining, and married women’s labor supply. *American Economic Review* 88(3).
- Greenwood, J., N. Guner, G. Kocharkov, and C. Santos (2012). Technology and the changing family: A unified model of marriage, divorce, educational attainment and married female labor-force participation. NBER Working Papers 17735.
- 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.
- Iyigun, M. and R. P. Walsh (2007). Building the family nest: Premarital investments, marriage markets, and spousal allocations. *Review of Economic Studies* 74(2), 507–535.
- Jacquemet, N. and J.-M. Robin (2011). Marriage with labor supply. working paper.
- Johnson, W. R. and J. Skinner (1986). Labor supply and marital separation. *American Economic Review* 76(3), 455–469.
- Jones, L., R. Manuelli, and E. McGrattan (2003). Why are married women working so much? working paper.

- Kneip, T. and G. Bauer (2007). Effects of different divorce probabilities on female labor force participation and fertility. Mannheim Centre for European Social Research Working Paper 102.
- 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).
- Love, D. (2010). The effects of marital status and children on savings and portfolio choice. *Review of Financial Studies* 23(1), 385–432.
- Lundberg, S. and E. Rose (1999). The determinants of specialization within marriage. Unpublished paper. University of Washington.
- Mazzocco, M., C. Ruiz, and S. Yamaguchi (2007). Labor supply, wealth dynamics, and marriage decisions. mimeo.
- 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.
- Papps, K. L. (2006). The effects of divorce risk on the labour supply of married couples. IZA Discussion Paper No. 2395.
- 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.
- Rios-Rull, J.-V. and F. Regalia (2001). What accounts for the increase in the number of single households? working paper.
- Ríos-Rull, J. V. and V. Sanchez-Marcos (2002). College attainment of women. *Review of Economic Dynamics* 5(4).
- Stevenson, B. (2008). Divorce law and women's labor supply. *Journal of Empirical Legal Studies* 5(4).
- Storesletten, K., C. I. Telmer, and A. Yaron (2004). Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy* 112(3), 695–717.
- van der Klaauw, W. (1996). Female labour supply and marital status decisions: A life-cycle model. *Review of Economic studies* 63(2).
- Voena, A. (2012). Yours, mine and ours: Do divorce laws affect the intertemporal behavior of married couples? working paper.
- Weiss, Y. (1997). The formation and dissolution of families: Why marry? who marries whom? and what happens upon marriage and divorce? In R. Rosenzweig and O. Stark (Eds.), *Handbook of Population Economics*. Amsterdam Elsevier Science.

Figures and Tables

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





Figure 2: LFP for married and divorced women by education (left) and skill premium and the gender wage ratio (right) for the 1935 cohort, model vs data.



Note: \bullet = data. ($\square, \circ, \triangle$) = model predictions

Figure 3: Model predictions for married women with 1955 assortativeness and 1955 divorce rates




Figure 4: Model predictions for married women with various 1955 family structure changes




Figure 5: Model predictions for divorced women with 1955 family structure changes

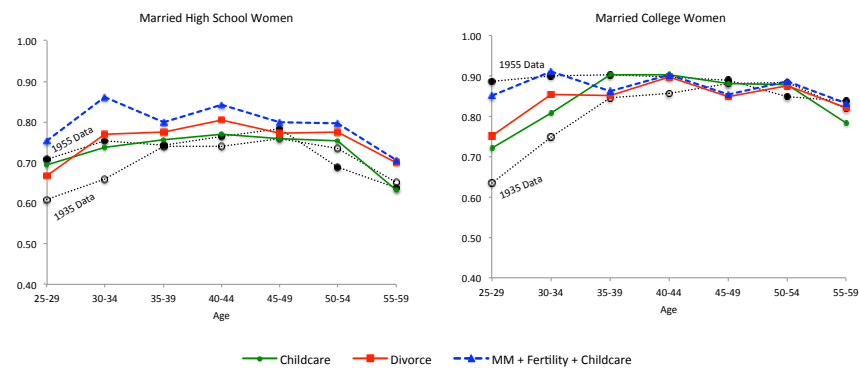


Figure 6: Model predictions for married women with 1955 wage and family changes




Figure 7: Model predictions for divorced women with 1955 wage and family changes





Figure 8: LFP for married and divorced women by education and age for the 1935 cohort (left) and skill premium and the gender wage ratio for the 1935 cohort by age (right), linear model vs data.



Note: ● = data. (□, ○, △) = model predictions

Figure 9: Model predictions for married women with 1955 assortativeness and 1955 divorce rates, linear model




Table 1: 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}	0.1651
	γ_{lm2}	-0.00308
Regression log wage on age and age ² , College men	γ_{hm1}	0.2276
	γ_{hm2}	-0.00285
Returns to experience, women	γ_{ef1}	0.1041
	γ_{ef2}	-0.003
Persistence of wage residuals, by educ.	ρ_l	0.90985
	ρ_h	0.96920
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 A.1
	$\{\tau_{hm,t}\}_{t=2,t^R}$	
Probability of divorce, remarriage and first marriage	$\{d_{eg,t}, r_{eg,t}, m_{eg,t}\}_{t=1,\dots,4}$	see Tables A.1, 6

Table 2: 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	$\{\tau_{lf,t}\}_{t=1,t^R}$	see Figure A.1
	$\{\tau_{hf,t}\}_{t=1,t^R}$	
Wage depreciation from not working	δ	0.0839
Moments	Data	1935 Bench
LFP of married women, by age and education (14)	see Table 3	see Table 3
LFP of divorced women, by age and education (14)	see Table 3	see Table 3
Skill Premium by age, women (7)	see Table 3	see Table 3
Ratio of male to female wages by age (7)	see Table 3	see Table 3
Skill Premium (lifetime), men	1.43	1.44

Table 3: Calibration targets, model vs. data

	Married LFP				Wages	
	College		HS		Gender Wage Ratio	
Ages	Data	Model	Data	Model	Data	Model
25-29	34.59	33.41	26.40	27.51	0.75	0.70
30-34	32.74	34.96	34.22	33.58	0.65	0.62
35-39	47.10	47.92	44.81	45.29	0.60	0.58
40-44	62.34	61.94	54.72	55.60	0.55	0.58
45-49	69.97	68.87	57.30	58.04	0.55	0.60
50-54	68.85	69.81	56.69	57.53	0.59	0.58
55-59	62.50	62.55	48.22	47.59	0.63	0.63
	Divorced LFP				Wages, Women	
	College		HS		Skill Premium	
Ages	Data	Model	Data	Model	Data	Model
25-29	63.37	63.46	60.90	60.11	1.40	1.45
30-34	75.00	75.43	66.05	65.68	1.44	1.45
35-39	84.65	84.34	73.95	72.95	1.39	1.43
40-44	85.67	86.66	73.87	74.79	1.34	1.39
45-49	88.10	87.58	76.00	75.95	1.39	1.45
50-54	88.41	87.69	73.38	74.11	1.48	1.44
55-59	81.96	80.58	65.13	63.93	1.53	1.48
	Skill Premium, men:				1.43	1.44

Note: The total distance between moments and targets is computed as a weighted average of the difference between each moment and target. The moments and targets for the gender wage ratio and skill premium are both multiplied by 100. The weights are such that the 28 moments for LFP receive half of the total weight and the 15 moments for wages receive the other half of the weight. Within each group, the moments are equally weighted.

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: 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 6: Per period divorce and remarriage probabilities

		d_1	d_2	d_3	d_4	r
1935 Cohort						
Men	HS	5.87	8.60	8.49	6.44	37.23
	College	6.45	6.03	4.67	5.67	42.48
Women	HS	5.39	4.93	5.34	5.38	38.84
	College	3.44	3.84	4.62	5.46	38.49
1955 Cohort						
Men	HS	18.04	11.60	7.85	4.70	35.21
	College	17.42	9.17	5.65	5.06	40.18
Women	HS	14.48	9.90	5.36	3.61	36.11
	College	13.99	9.90	6.17	2.93	35.38

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. The divorce numbers were derived using 2004 SIPP data. The remarriage 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 7: Effects from Divorce: LFP

	Experience (LFP levels)			Bargaining (LFP differences)		
	Benchmark	Divorce	(2)-(1)	Original	$\chi = 0.5$	$\chi = 0.7$
ages	(1)	(2)	(3)	(4)	(5)	(6)
A. High School						
25-29	19.72	51.88	32.16	32.38	36.05	38.87
30-34	29.53	52.13	22.60	22.77	26.78	30.02
35-39	39.82	48.45	8.63	8.62	13.25	17.17
40-44	51.49	54.17	2.69	2.69	5.81	7.31
45-49	54.95	56.54	1.58	1.58	2.65	1.54
50-54	55.20	57.06	1.86	1.87	2.71	1.81
55-59	47.51	48.59	1.08	1.08	1.87	-0.32
B. College						
25-29	23.60	59.13	35.52	37.76	42.85	46.77
30-34	29.30	55.68	26.38	27.56	32.01	35.59
35-39	40.80	53.76	12.96	13.33	18.99	23.79
40-44	60.36	58.99	-1.37	-1.38	3.16	6.80
45-49	65.70	65.98	0.27	0.25	2.84	5.00
50-54	70.12	69.55	-0.58	-0.58	1.50	4.24
55-59	60.97	63.06	2.09	2.09	4.22	6.40

Notes: Column 1 shows the LFP values for HS and college women in the benchmark model with exogenous returns to experience, while column 2 shows the analogous values for the experiment with 1955 divorce probabilities. Their difference is shown in column 3. Columns 4-6 shows the corresponding difference for (4) the model with endogenous returns to experience, (5) the model with $\chi = 0.5$ and (6) $\chi = 0.7$. The benchmark for all cases corresponds to the baseline 1935 setup with the 1955 assortativeness.

Table 8: Effects from divorce: assets

	Benchmark (asset levels)			Differences	
	Original	Divorce	(2)-(1)	$\chi = 0.5$	$\chi = 0.7$
ages	(1)	(2)	(3)	(4)	(5)
A. High School					
25-29	-4.88	-5.01	-0.13	0.54	1.56
30-34	-7.61	-8.12	-0.51	0.56	2.23
35-39	-4.37	-4.73	-0.37	0.31	1.39
40-44	-0.28	-0.65	-0.37	-0.23	0.09
45-49	4.40	5.39	1.00	0.80	0.89
50-54	12.57	16.55	3.98	2.70	1.19
55-59	22.37	27.25	4.88	3.11	1.66
B. College					
25-29	-6.87	-7.39	-0.53	1.60	1.94
30-34	-11.25	-12.40	-1.15	2.57	3.18
35-39	-7.33	-8.12	-0.79	1.67	2.07
40-44	-0.98	-1.20	-0.22	0.20	0.28
45-49	7.10	12.55	5.45	1.71	0.98
50-54	19.22	27.53	8.31	2.41	1.49
55-59	31.93	41.24	9.31	3.39	2.81

Notes: Column 1 shows the values for household asset for HS and college women in the benchmark model, while column 2 shows the analogous values for the experiment with 1955 divorce probabilities. Their difference is shown in column 3. Columns 4-5 shows the corresponding difference for (4) the model with $\chi = 0.5$ and (5) $\chi = 0.7$. The benchmark for all cases corresponds to the baseline 1935 setup with the 1955 assortativeness.

Table 9: Linear model: disutility of labor parameters, $\psi_{em}(k)$

		Married	Married with Child	Divorced	Divorced with Child
Parameter values	High School	0.1454	0.3468	0.1311	0.1762
	College	0.0986	0.1729	0.0587	0.1400
Consumption equivalence	High School	12.3%	34.1%	14.9%	16.1%
	College	6.1%	13.5%	7.7%	9.0%

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 10: Linear model: parameters calibrated internally

Parameter		1935 Bench
Disutility of labor for women	$\psi_e^s(k_t)$	see Table 9
Childcare costs for children aged 0-4 and 5-9	κ_{young}	4.0596
	κ_{old}	0.5753
Initial intercept log wages for men, by education	$\tau_{lm,1}$	0.2860
	$\tau_{hm,1}$	2.7458
Time varying log wage intercepts, women	$\{\tau_{lf,t}\}_{t=1,t^R}$	see Tab A.4
	$\{\tau_{hf,t}\}_{t=1,t^R}$	
Wage depreciation from not working	δ	0.0899
Moments	Data	1935 Bench
LFP of married women, by age and education (14)	see Tab A.5	see Tab A.5
LFP of divorced women, by age and education (14)	see Tab A.5	see Tab A.5
Skill Premium by age, women (7)	see Tab A.5	see Tab A.5
Ratio of male to female wages by age (7)	see Tab A.5	see Tab A.5
Skill Premium (lifetime), men	1.43	1.43

FOR ONLINE PUBLICATION

A Appendix: Figures and Tables

Figure A.1: Implied log wage intercepts for men and women, by education - 1935 cohort


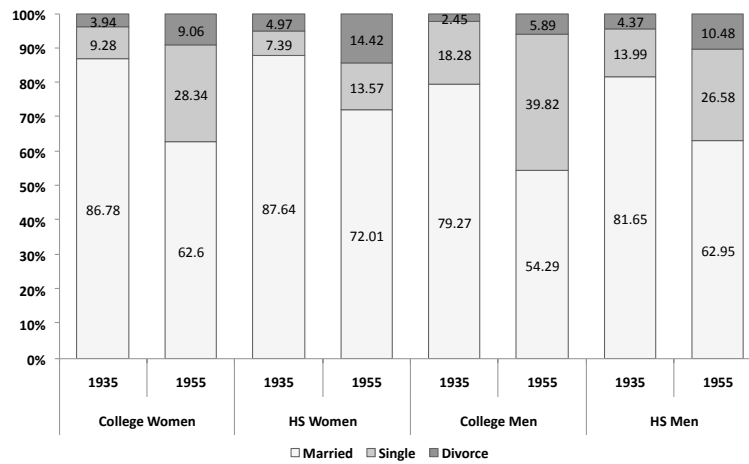


Figure A.2: Proportion of each gender-education group in each marital state during the ages of 25-29, by cohort.



Note: Current Population Survey. Sample consists of all white men and women in each birth cohort. "Married" is defined as "married, with spouse present"; "Divorced" is defined as either "divorced" or "separated". 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 A.3: Linear model: implied log wage intercepts for men and women, by education.




Table A.1: Transition probabilities from single to first marriage

		Men		Women	
	age	College	High School	College	High School
1935 Cohort	30	51.24	44.19	18.14	41.01
	35	26.06	18.90	19.52	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 A.2: 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 A.3: Implied log wage intercepts for men and women, by education

age	College Men	HS Men	age	College Women	HS Women
25-29	0.61	0.09	25-29	0.61	0.10
30-34	0.69	0.19	30-34	0.63	0.11
35-39	0.59	0.14	35-39	0.61	0.10
40-44	0.54	0.10	40-44	0.75	0.12
45-49	0.47	0.04	45-49	0.79	0.14
50-54	0.54	0.02	50-54	0.80	0.16
55-59	0.48	-0.06	55-59	0.84	0.18

Table A.4: Implied log wage intercepts for men and women, by education - Linear model

	College Men	HS Men		College Women	HS Women
25-29	0.29	0.20	25-29	4.45	2.73
30-34	0.37	0.30	30-34	4.27	2.46
35-39	0.27	0.25	35-39	4.21	2.07
40-44	0.22	0.21	40-44	3.98	1.86
45-49	0.15	0.15	45-49	3.24	1.57
50-54	0.21	0.13	50-54	3.35	1.60
55-59	0.16	0.05	55-59	3.23	1.55

Table A.5: Calibration targets linear model, model vs. data

	Married LFP				Wages	
	College		HS		Gender Wage Ratio	
Ages	Data	Model	Data	Model	Data	Model
25-29	34.59	36.02	26.40	28.06	0.75	0.72
30-34	32.74	35.53	34.22	35.33	0.65	0.64
35-39	47.10	45.40	44.81	43.37	0.60	0.60
40-44	62.34	64.07	54.72	56.54	0.55	0.59
45-49	69.97	68.14	57.30	56.72	0.55	0.62
50-54	68.85	68.66	56.69	58.61	0.59	0.60
55-59	62.50	60.59	48.22	48.67	0.63	0.65
	Divorced LFP				Wages, Women	
	College		HS		Skill Premium	
Ages	Data	Model	Data	Model	Data	Model
25-29	60.90	59.80	60.90	59.80	1.40	1.42
30-34	66.05	65.00	66.05	65.00	1.44	1.42
35-39	73.95	74.52	73.95	74.52	1.39	1.40
40-44	73.87	76.49	73.87	76.49	1.34	1.36
45-49	76.00	76.28	76.00	76.28	1.39	1.42
50-54	73.38	75.02	73.38	75.02	1.48	1.41
55-59	65.13	65.24	65.13	65.24	1.53	1.44
	Skill Premium, men:				1.43	1.43

Note: The total distance between moments and targets is computed as a weighted average of the difference between each moment and target. The moments and targets for the gender wage ratio and skill premium are both multiplied by 100. The weights are such that the 28 moments for LFP receive half of the total weight and the 15 moments for wages receive the other half of the weight. Within each group, the moments are equally weighted.

B Appendix: Data

B.1 LFP

We use the Current Population Survey from 1962- 2010 to compute labor force participation rates. We construct two synthetic cohorts of women born between 1934-1936 and 1954-1956 and compute their LFP as the proportion of women during each age bracket (e.g. 25-29) who report being in the labor force.

Data is available for the most part of the life-cycles of these women. However, note that the women who are aged 26-29 in 1962 correspond directly to the women from our cohort. However, we do not have any information about these women’s behavior when they are 25. Thus, in order to compute the LFP for the first period of the women from the 1935 cohort, we also add to the sample women who are 25 years old in 1962 (these women would have been born in 1933). The opposite problem happens for the 1955 cohort. The data ends in 2010 when the median woman of that cohort is 55 years old. The labor force participation behavior for those women during their last period of work (aged 55-59) is computed from the LFP behavior of the women who are aged between 55-59 in 2010 (some of which belong to our cohort - those aged 55 and 56 - and some who were born between 1951-1953).

B.2 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.

B.3 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:

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.⁵¹

First, we construct data on hourly wages ($y_{em,t}$) for men of education level e using data on earnings and total hours worked. 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:

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

where \mathbb{D}_{emt} is a set of year dummies.

Given the residuals from the regression above, we estimate the parameters for the

⁵¹The 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.

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.

In order to estimate the time-varying age intercepts for men $\tau_{em,t}$, 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 A.1 in the Online Appendix these are reasonable assumptions. We use \mathbb{D}_{emt} to compute the estimates for the time-varying intercepts.

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 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 $\tau_{em,5} = \tau_{em,4}$.

B.4 Divorce and Remarriage Rates

To compute divorce rates 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 by increasing 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. 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 who divorced in each age bracket.

For remarriage rates, due to sample size restrictions, we proceed in a different way.⁵² 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.⁵³

B.5 Fertility

The timing of fertility shocks is as follows. For the 1935 cohort, single women and college women of all marital status receive a fertility shock in the first period that takes the value of zero or one. Divorced women and married high-school women receive a fertility shock that takes the value of zero or two. The probability of receiving a non-zero value is calibrated so as to match the initial proportions (age 25-29) in the data by marital status and education as reported in Table A.2. 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 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.⁵⁴

⁵²The 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 employ the same procedure.

⁵³We 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.

⁵⁴The 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

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 A.2. 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.⁵⁵ 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.

B.6 Consumption Deflator:

We use an altered McClements scale ($\mathbf{e}(k_t; s)$) in order to deflate household consumption. Table B.6 reproduces the original McClements scale in normalized for one adult.

Table B.6: 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).

B.7 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

if they have college education and 43% probability if they have high school.

⁵⁵Once 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.

average earnings over their lifecycle. More specifically, we take each individual's last observed earnings y_T and compute social security benefits as follows: 90% of y_T up to a first threshold equal to $0.38\bar{y}_T$, where \bar{y}_T is the average observed earnings in the economy during the last period, plus 32% of y_T from this bendpoint to a higher bendpoint equal to $1.59\bar{y}_T$, plus 15% of the remaining y_T exceeding this last bendpoint. 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.

C 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 using value function iteration. 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 two other potential continuous state variables: the persistent components of earnings w_t^{em} and of the wife, w_t^{ef} . Handling more than one continuous state variable is possible but computationally costly. We choose to discretize these two variables, leaving assets as our only continuous state. We use 25 nodes whose locations are age-dependent for each of the earnings components.

During the working stages of the lifecycle, our model combines a discrete decision (whether the woman participates in the market) and a continuous decision (the amount of savings). This combination may lead to non-concavities in the value function. Furthermore, the existence of transitions across marital states also requires some attention. For all periods $t > 4$, since there are no longer any transitions across different marital status, the 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 (working vs not working) intersect only once and thus the woman's participation decision switches at that point from not working to working. Thus conditional on all other states, for all values $a_t < a_t^*$, the woman works and for all values $a_t > a_t^*$ the woman does not work. 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 also 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. So, 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. Thus the algorithm is as follows:

- Given $a_{m,t+1,0}^d$ and all the other state variables, the ex-wife chooses $a_{f,t+1,0}^d$ and P_t .
- The ex-husband takes $a_{f,t+1,0}^d$ chosen by the wife and chooses $a_{m,t+1,1}^d$.
- If $a_{m,t+1,1}^d = a_{m,t+1,0}^d$ then this decision point is done. If not, then the wife will take $a_{m,t+1,1}^d$ as given and solve for a new value of $a_{f,t+1,1}^d$ and P_t .
- This process is iterated until convergence, defined as $|a_{m,t+1,j}^d - a_{m,t+1,j+1}^d| < 1^{-10}$

If the process above converges after D iterations, then the asset levels at time $t + 1$ for the divorced female and male are given by $a_{f,t+1}^d = a_{f,t+1,D}^d$ and $a_{m,t+1}^d = a_{m,t+1,D}^d$