Risky Income, Risky Families: Marriage and Divorce in a Volatile Labor Market*

Cezar Santos, David Weiss †

University of Mannheim & University of Pennsylvania

Abstract

There has been a striking increase in American idiosyncratic labor income volatility since 1970, with little attention paid to the effects on families. The formation and dissolution of the typical American family has changed substantially, however, with a notable decline/delay of marriage and, since 1980, declining divorce rates. Furthermore, the elderly are divorcing more. This paper demonstrates a quantitatively important link between income volatility and the changing family. Marriage typically involves children, a large, persistent cost, which causes people to dislike risk; volatility therefore causes less marriage. This effect dominates the increased insurance value of marriage that arises because shocks to income are imperfectly correlated between spouses. Once a couples has married, however, the rising insurance value of marriage also leads to a decline in divorce. On the other hand, the elderly are either retired or near retirement and have grown children, and thus are less susceptible to the effects of volatility. Elderly divorce rises as younger people delay divorce. The model qualitatively matches observed family changes over time, and quantitatively accounts for up to a third of the data.

Keywords: marriage, divorce, fertility, consumption commitments, earnings volatility, search models of marriage, simulated method of moments.

*We would like to thank Jeremy Greenwood, Iourii Manovskii, Dirk Krueger, Rachel McCulloch, Doug Hanley, Fatih Karahan, Kurt Mitman, Serdar Ozkan, and seminar participants at the Penn Macro Seminar for their comments. All remaining errors are our own.

†E-mails: santos@uni-mannheim.de and daweiss@econ.upenn.edu

1 Introduction

"Keep the eyes wide open before marriage and half shut afterwards."

BENJAMIN FRANKLIN

The United States labor market has changed dramatically since the early 1970s, with a stark rise in labor income volatility.[1] While many economists study the effects of these changes on consumption, savings, inequality, and career choices, few have analyzed their effects on families.[2] However, the American family has experienced substantial changes over the same period. It is well known among both academics and the general public that rates of marriage have been on the decline since 1970.[3] While the "Divorce Revolution" of the 1970s has been widely noted by academics and the general public, researchers have documented a substantial subsequent decline in divorce in the years since. Since 1980 the divorce rate in the U.S. has fallen by about a quarter. These trends are evident in Figure 1. The trend in overall divorce masks significant heterogeneity: While divorce rates are falling overall (as seen in the figure), they have been rising for the elderly, more than doubling since 1980. In this paper, we propose and quantitatively evaluate mechanisms, related to mutual spousal insurance and children, through which rising income volatility can explain these three trends in marriage and divorce.

[1] For example, Gottschalk and Moffitt (1994) discuss the growing instability in wages and Katz and Autor (1999) study the changes in wage structure and overall earnings inequality.

[2] See, for instance, Kambourov and Manovskii (2009) for the effects of increasing occupational mobility on inequality, and Krueger and Perri (2006) for the effects of income volatility on consumption inequality. Santos and Weiss (2012) and Sommers (2011) are the only papers we are aware of that relate earnings volatility to families.

[3] For an excellent review of the academic literature pertaining to marriage and divorce, see Stevenson and Wolfers (2007).

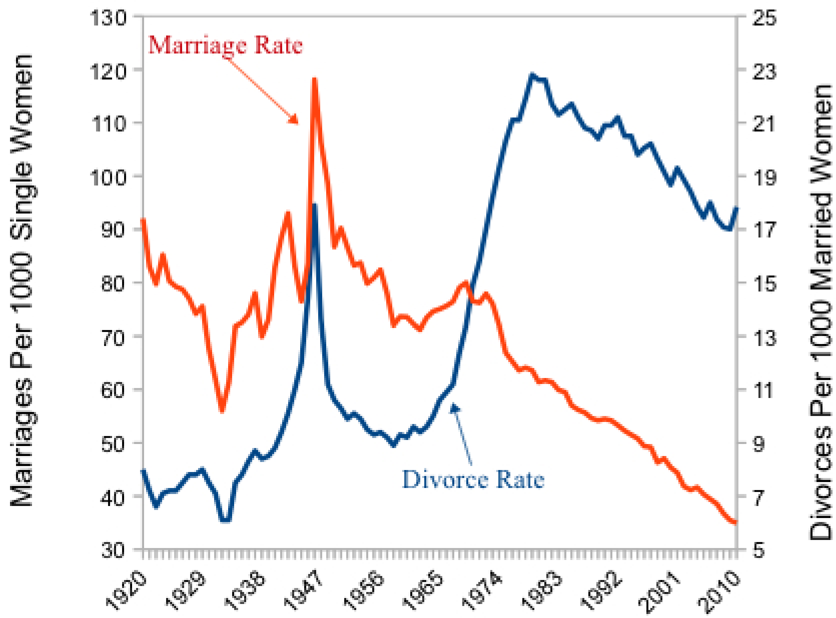

Figure 1: Marriage and Divorce Over Time

Volatility has two opposing effects on marriage. The first comes from the fact that a significant reason for marriage is children. Children represent a consumption commitment, which makes married couples behave as if they are more risk averse than singles.[4] This is because children account for a large fraction of household expenditures, and remain part of the family even after an adverse income shock. When volatility rises, the relative value of being married to being single drops, causing a decline in marriage rates. On the other hand, marriage also allows for diversification of income risk as earnings fluctuations between spouses need not be perfectly correlated. Therefore, higher income volatility makes marriage more desirable due to insurance. It is the interaction between these two aspects of

[4]Note that we use the term "consumption commitments" a little loosely. Generally the literature, such as Chetty and Szeidl (2007), uses the term to describe consumption goods that are costly to adjust. This term applies well to mortgages and housing. Children might not represent a consumption commitment in the general sense of the word– their consumption is perhaps just as flexible as their parents' consumption. However the fact they exist leads to an additional set of goods the household optimally chooses to buy, which in turn generates the same effect as a consumption commitment. Throughout this paper, we abuse the term as such, though we do find evidence that a portion of expenditures on children is relatively fixed, and thus a commitment.

marriage, consumption commitments and mutual spousal insurance, that allows volatility to generate the observed changes in marriage, divorce, and elderly divorce.

Due to the conflicting channels through which volatility affects marriage ex-ante, the net effects are ambiguous. In the quantitative exercise we find that increased volatility leads to less marriage.[5] That is, the importance of consumption commitments within marriage outweigh the gains from the added spousal insurance.

We next turn to examining the mechanism for generating a decline in divorce. Increased volatility leads to less divorce as people value spousal insurance more, causing them to remain married even after realizing unfavorable shocks. Hess (2004) highlights mutual income insurance as an influence on people's decisions to marry and divorce. This mechanism is comparable to the one in Krueger and Perri (2006), in which people are more willing to remain in insurance contracts when risk increases. We examine marriage as a special case of such contracts, which in this paper generates the main channel for the decrease in divorce. The fact that children do not disappear upon divorce amplifies this mechanism.[6]

While divorce rates have been dropping across education and racial lines, they have not dropped for all age groups. The elderly in America have experienced an increase in divorce rates. Sociologists refer to this as the "Gray Divorce Revolution" (Brown and Lin, 2012). From 1980 to 2007, divorce rates for women (men) over age 60 increased from 1.75 (2.47) to 7.32 (6.71) per 1000 married.[7] This fact is indeed predicted by the mechanisms presented in this paper. Retired people face little or no labor market volatility. Those near retirement face a shorter time horizon over which they might expect a spouse to insure them. Furthermore, elderly couples are more likely to have their mortgages paid off, and their children tend to be grown up, implying they have fewer consumption commitments. In short, older people have less use for marriage as insurance. If younger people are divorcing less in order to insure one another, then there are more marginal marriages in existence

[5]In Santos and Weiss (2012), we quantitatively evaluate the effects of increased earnings volatility vis-à-vis other mechanisms in the literature to explain the increase in the age of first marriage that the US has experienced since 1970. We find that increased volatility can account for over a third of the delay. In this paper, rather than focusing on the age of first marriage only, we explore how risk affects both family formation *and* dissolution *over the life cycle*.

[6]Additionally, there is evidence in the literature that divorce rates are sensitive to risk. Hellerstein & Morrill (2011) use state-level data on divorce and unemployment rates to argue that divorce is significantly procyclical. While not, strictly speaking, necessarily a sign of risk, house prices also predict divorce. Farnham et al (2011) show that when house prices fall, divorce rates fall, and use this to explain some of the fall in divorce during the Great Recession.

[7]Data from the American Community Survey and Clarke (1995), following the methodology used in Brown and Lin (2012).

when people become old. The bifurcation of divorce rates in the data is therefore consistent with divorce decisions being driven by insurance motives.

One particularly useful way of looking at the empirical changes in divorce rates is to follow Stevenson and Wolfers (2011) in creating cumulative divorce probabilities by marital tenure, focusing on first marriages. Figure 2 shows the probability, by marriage cohort, that a marriage has ended in divorce by a certain anniversary.[8] Using this graph, we see that marriages that have begun since the 1980s show lower probabilities of divorce.[9] For instance, the marriages that began in the 1990s have divorce probability rates similar to those of marriages that began in the 1960s, before the "Divorce Revolution" of the 1970s.[10]

It is worth noting that the facts we are trying to explain might seem strange at first glance: If marriage rates are going down, this would suggest a decrease in the gains to marriage, which would tend to imply an *increase* in divorce. Indeed many papers on the decrease in marriage employ various mechanisms by which the value of being married relative to unmarried decreases.[11] However, as can be seen in Figure 1, marriage and divorce rates usually move together, with the notable exception of the period of the "Divorce Revolution" of the 1970s. Furthermore, the fact that elderly divorce has risen calls for a culprit that affects older people differently than younger people. We consider the fact that the mechanisms proposed in this paper predict that marriage and overall divorce rates move in the same direction, while elderly divorce moves in the opposite direction, to be very compelling arguments for the story presented in this paper.

While there have been many quantitative papers written about the decline in marriage,[12] the only other paper we are aware of that attempts to quantitatively account for the decline in divorce over the last few decades is Rotz (2011). She claims that the older age of marriage

[8]Qualitatively, our results and theirs are the same. We differ in that we use 3 waves of the Survey of Income and Program Participation (SIPP), and restrict attention to whites.

[9]This data is robust to breakdowns by education group. See Figures 7 and 8 in the Appendix B for the breakdown by different education groups.

[10]One potential explanation as to why divorce rates dropped after 1980 is that there was a bulk of marginal marriages that didn't dissolve until changes during the 1970s, such as the introduction of unilateral divorce and women's liberation. This hypothesis says that divorce rates were unusually high in the 1970s due to these transitory reasons, and dropped afterwards. Given that divorce rates declined over the entire 1980-2005 period, and that much of the decrease started in the mid 1990s, this hypothesis cannot fully explain the observed changes.

[11]For example, Greenwood & Guner (2009) argue that technological improvements in home goods, such as refrigerators and washing machines, have decreased the value of spouses specializing in market vs home goods, leading to less marriage and more divorce. Rios-Rull & Regalia (2001) argue that the narrowing of the gender wage gap is responsible for the increase in single households.

[12]See Wolfers and Stevenson (2007) for a review of the marriage and divorce literature.

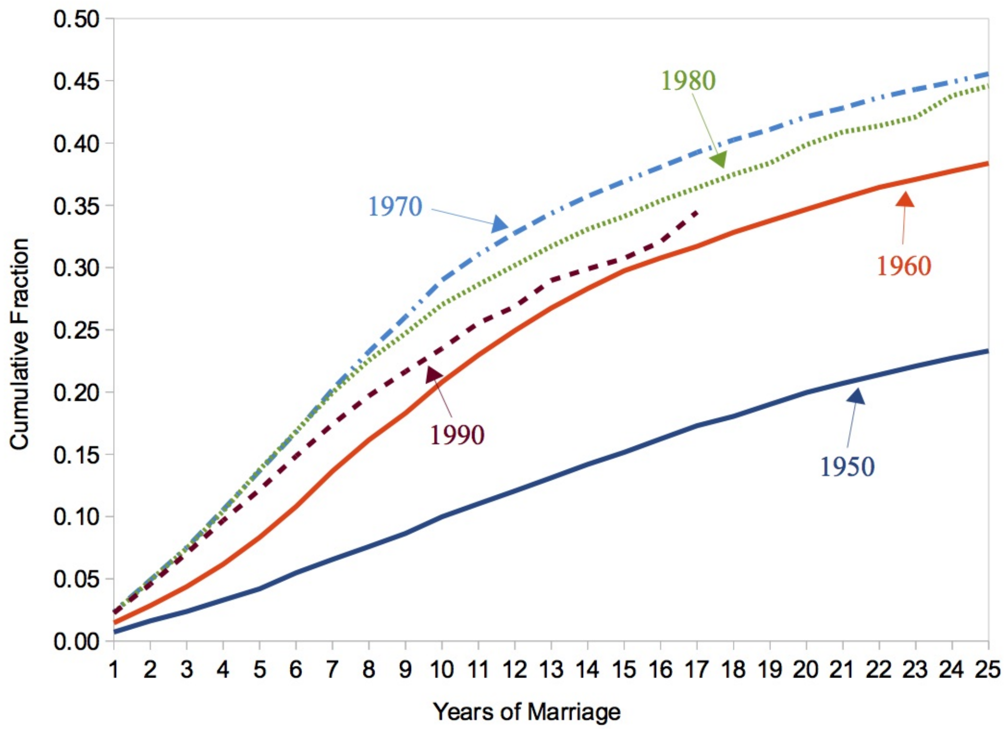

Figure 2: Cumulative Divorce Probabilities, White Americans, First Marriage. Authors Calculation from the SIPP 2001-2008, following Stevenson & Wolfers (2011)

among women has caused a decrease in divorce, as older women make better matrimonial choices and have worse outside options once they actually do marry, incentivizing them to remain married.. Using regression analysis, she argues that controlling for the age of the bride at marriage helps statistically explain the decline in divorce. While compelling, her hypothesis cannot be a complete explanation. This is because, as she notes, the age of the entrants at their first marriage cannot statistically explain the decline in divorce among all demographic groups, such as non-whites or the college educated. Furthermore, this explanation is difficult to reconcile with the rise in the rate of elderly divorce. Increased labor market risk can explain the decline in overall divorce and the increase of divorce among the elderly, as well as the delay in marriage, which may in turn further decrease divorce along the lines of Rotz's argument.

This paper also relates to a very active literature exploring the interactions between family formation and dissolution and other economic choices people make. For example, Mazzocco et. al. (2007) document how labor supply changes with respect to marriage and divorce. Marriage allows for specialization based on the comparative advantage between men and women in market and home production, while divorce marks the end of this arrangment. Marriage therefore decreases female (market) labor supply, while divorce undoes this effect. Guvenen and Rendall (2012) explore how growth in divorce increases women's incentives to educate themselves as a form of self-insurance. They show how this results in a further increase in divorce, and a reduction in marriage, as women become more self sufficient. This demonstrates a clear mechanism for how income risk and family choices can affect women's education and career decisions. Tavares (2011) examines how the US tax code, which uses households rather than individuals as the unit of taxation, affects female education choices and labor force decisions, using the marriage market as an amplification mechanism. This aspect of the tax code also affects perceived risk by people in a progressive taxation environment, leaving room for future work. Finally, Chakraborty et. al. (2012) looks at how differences in divorce rates across countries can account for differences in female labor force participation.

In order to assess quantitatively our hypothesis, we build an equilibrium search model of the marriage market, complete with a life cycle component. People make choices with respect to consumption, savings, fertility, marriage, divorce, and married women's labor force participation. In order to accurately measure spousal insurance, we account for both married women's labor force participation rates and the gender wage gap. Each person's labor income is risky, and households can save in a riskless bond market. Married couples benefit from economies of scale in consumption but implicitly face consumption commitments that persist even after marriage ends, such as children. Household decisions are determined through Nash bargaining between the spouses.

The model is estimated using the Simulated Method of Moments. We target several moments regarding marriage, divorce, fertility, labor force, and consumption choices that are derived from different micro data sets. All of the model moments are calibrated to data from 2005. We then perform a counterfactual simulation using 1970 labor income volatility parameters in order to assess what the current marriage and divorce statistics would be had volatility not increased.

Based on our counterfactual analysis, we find that increased earnings volatility is indeed capable of qualitatively generating all of the changes in marriage and divorce discussed above, namely: a decline in marriage rates, declining overall divorce rates, and an increase in the divorce rate among the elderly. Quantitatively, the model is capable of explaining 35% of the decline in divorce, 17% of the decline in marriage, and 9% of the rise in elderly divorce.

The remainder of the paper is organized as follows. We begin by describing the model in Section 2. We then discuss the mechanisms at work in the model in Section 3. Calibration and estimation procedures are documented in Section 4. Section 5 presents the results of the model. We conclude in Section 6.

2 The Model

The economy is populated by overlapping generations of men and women. There is a unit measure of each gender, g, and age, a. Agents age stochastically between four age groups that represent different stages of the life cycle: the 20s, 30s, 40s/50s, and 60s. Agents can either be single or married, and either have children or not. They begin life as a single at age $t = 1$, with no assets, no children, no child support obligations (or receipts), and an initial income realization drawn from the invariant distribution of the stochastic process for income shocks.

2.1 Production

There is one good produced in the economy, a market good denoted Y. There is a linear production function, with labor as the only input:

$$Y = AL, \tag{1}$$

where A is a technology parameter normalized to 1, and L is aggregate market labor supply. This implies that the wage in the model is equal to the efficiency units of labor supplied.

The amount of efficiency units of labor, y, supplied by each agent follows a stochastic process around a deterministic trend:

$$y = w\phi_g f(t), \tag{2}$$

where t is the individual's age, w is an idiosyncratic shock, and the deterministic trend is composed of $\phi_g(t)$, a gender wage gap, and $f(t)$, a deterministic age income profile. We will now discuss each of these terms.

The shock w consists of a persistent shock, with innovations η. The shock process is specific to the agent's marital status; both the variance and the autocorrelation of the shocks may be different between the two groups. This allows for the fact that married and single agents may behave differently, especially in the presence of consumption commitments. For example, perhaps people who are married are less likely to want to switch careers, since such moves typically involve a short run cost of lower wages during retraining. For this reason, we assume that divorced individuals with children follow the same income process as married people.[13] Additionally, we allow for persistent shocks to be correlated between spouses. For example, if one spouse loses a job and needs to move to take a new one in a different city, then the other spouse will also need to move and find a new, potentially worse job. Since we are not modeling behavior in the labor market explicitly, we must account for differences in labor market outcomes by estimating different income processes by marital status. Thus, we assume that this process takes the following form for singles (denoted by the subscript s):

$$\begin{aligned} \ln w_s &= \delta_s \ln w_{s,-1} + \eta_s \\ & \eta_s \sim N(0, \sigma^2_{\eta,s,\tau}). \end{aligned} \tag{3}$$

For married individuals (denoted by the subscript m), the process specifies shocks for each of the two spouses (an arrow above each shock denotes that this is a vector). The parameter ρ controls the correlation of spousal shocks. This is important as it allows the

[13] Data limitations prevent a third category of income processes.

model to attain the appropriate level of spousal insurance. Thus, the income process for married households takes the following form:

$$\begin{aligned} \ln \vec{w}_m &= \ln \vec{w}_{m,-1} + \vec{\eta}_m \\ \vec{\eta}_m &\sim N\left(0, \begin{bmatrix} \sigma^2_{\eta,m,\tau} & \rho \\ \rho & \sigma^2_{\eta,m,\tau} \end{bmatrix}\right). \end{aligned} \tag{4}$$

Note that the variances for all shocks are indexed by the time period subscript $\tau \in \{1980, 2005\}$. An increase in volatility is measured by changing σ^2_η, which controls the variance of the shocks.[14]

As noted above, the amount of efficiency units available to an agent also varies with his/her age t according to the function $f(t)$. This is intended to capture the average life cycle increase in earnings observed in the data.

Females supply a fraction ϕ compared to males; this accounts for the gender wage gap. Define the function ϕ_g that takes the value of 1 if $g = 1$ (males) or $\phi < 1$ for all t if $g = 2$ (females).

2.2 Preferences

Preferences of households are additively separable and exhibit constant relative risk aversion (CRRA) over both adults' consumption and children's consumption. The period utility function from consumption reads:

$$u(c, c_k) = \begin{cases} \frac{c^{1-\zeta}}{1-\zeta} + \alpha \frac{(c_k - \underline{c_k})^{1-\zeta}}{1-\zeta}, & k = 1 \\ \frac{c^{1-\zeta}}{1-\zeta} & , k = 0 \end{cases}, \tag{5}$$

where c is the adult's consumption and c_k is the child's consumption. $k \in \{0, 1\}$ is an indicator that the household has children. λ is the CRRA parameter controlling risk aversion.

The agent's total utility is equal to the expected discounted value of his or her lifetime utility:

[14] In the numerical analysis below, these continuous income processes are discretized using the method described in Kopecky and Suen (2010). Using their method is crucial, as the income processes exhibit high persistence.

$$U\left(\{c_{t=1}^{t=T}\},\{c_{k,t=1}^{t=T}\}\right)=E_{t=1}\left[\sum_{t=1}^{t=T}\beta^{t-1}u(c_t,c_{k,t})\right], \tag{6}$$

where β is a common discount factor.

When married, people enjoy a utility benefit of marital bliss, b. At the time of marriage, potential spouses draw the shock b from the distribution $\Upsilon(b)\sim N(\mu_b,\sigma_b^2)$. During marriage, the marital bliss evolves over time from state to state according to $\Xi(b,b')$, an AR(1) process with autocorrelation δ_b and innovations $\epsilon_b\sim N(0,\sigma_{bb}^2)$. Notice that the variance of these innovations is different than the variance of the initial draw. When married people have a child present in the household, and the other spouse is the parent of the child, then they receive an extra flow utility of $\bar{b_k}$. Notice that this flow utility is deterministic, and does not evolve. Additionally, married people draw a transitory marital bliss shock every period, $\lambda\sim N(0,\sigma_\lambda^2)$.

Moreover, if the wife works, both spouses suffer a cost of ψ. This cost is ex ante heterogeneous among women, and represents the cost to the household of the woman not maintaining the home. If there are children present, the family suffers an additional homogenous cost of the wife working of ψ_k.

2.3 Budget Sets

The budget constraint for single men is given by

$$c+a'=wf(t)+(1+r)a-\tau, \tag{7}$$

where w is the idiosyncratic productivity shock, $f(t)$ is the age-dependent productivity level, a is the individual's current level of assets chosen in the previous period, and a' is the savings chosen today. $(1+r)$ is the gross interest rate. τ is child support transfers to former wives (which is 0 if the man does not have any children).[15]

The budget constraint for single women is given by

$$c+c_k+a'=\phi(t)wf(t)+(1+r)a+\tau, \tag{8}$$

[15] There is no alimony in this model, as is the case in the literature. Voena (2012) documents that only 10% of divorced women in the Longitudinal Survey of Young and Mature Women receive alimony. Those that do receive relatively little, just 15% of household income. Given the limited nature of alimony, we feel that this simplifying assumption is reasonable.

where $c_k = 0$ and $\tau = 0$ if she has no children.

When married, spouses pool their resources. Adult consumption is a public good, as is child consumption. Married women have the option of whether to work in the market or work only at home, and l^f is the indicator function that women choose to work in the market. The husband's and wife's wage offers are denoted by w_1 and w_2. ξ represents economies of scale in household consumption. Hence, a couple's budget constraint reads

$$\frac{c}{\xi} + c_k + a' = w_1 f(t) + l^f \phi(t) w_2 f(t) + (1+r)a + \tau. \tag{9}$$

Here τ represents commitments from previous marriages.

Additionally, there is a consumption floor. If a household (either single or married) cannot afford to consume above the floor, there is assumed to be an exogenous transfer from an unmodeled government.

2.4 Period Timing

The timing of a period is as follows:

1. At the beginning of the period, everyone updates their wage offer. Married people update their bliss shock.

2. Single individuals randomly meet another single person of the same age group and opposite gender in the marriage market. They draw a marital bliss shock $b \backsim N(\mu_b, \sigma_b^2)$, and if they don't have children from a previous marriage, their first marriage period fertility bliss shock (described in 4). They then decide whether or not to marry. During marriage, bliss evolves according to an AR(1) : $b = \delta_b b_{-1} + \epsilon_b$, $\epsilon_b \sim N(0, \sigma_{bb}^2)$. We also assume that singles see the first period draw of λ. These assumptions ensure that they have all relevant information at the time of marriage.

3. If the individual remains single in the period, he/she pays or receives any child support payments (if applicable) and makes a consumption/savings decision.

4. Married people make allocation decisions. If they are under age 40 and do not already have children, they draw a fertility bliss shock $\gamma_f \backsim N(0, \sigma_{\gamma_f}^2)$.[16] They engage in Nash

[16]While this fertility bliss shock explicitly is a period-specific utility of choosing to have a child, high values can also be thought of as the failure of contraceptive methods, and low values can be thought of as attempts at reproduction failing, in addition to any idiosyncratic reasons why people may or may not choose to have children at a given point of time.

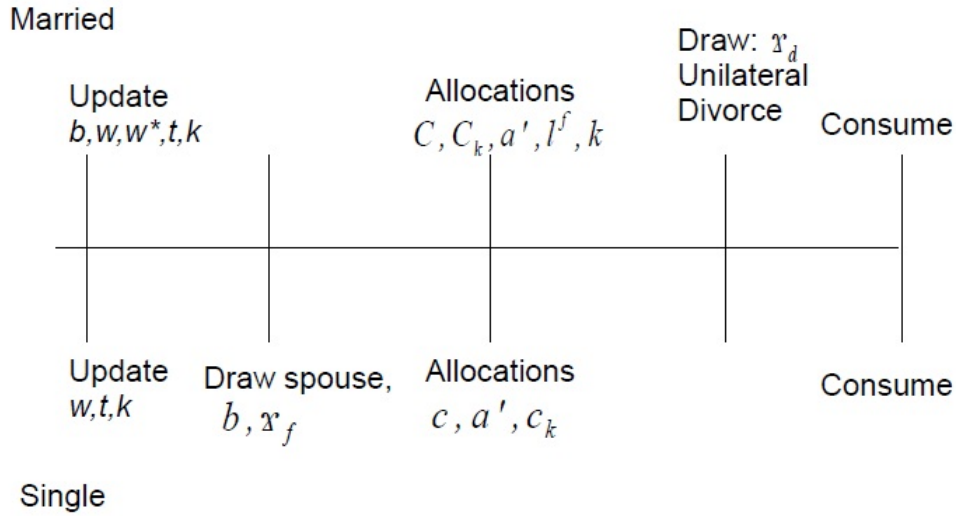

Figure 3: Timeline of the Model

bargaining over whether to have a child, consumption of the adults, consumption of children, female labor force participation, and savings.[17] They commit to this allocation should they remain married.

5. Married people receive a transitory bliss shock and decide, unilaterally, whether to divorce or execute the allocation decision. People cannot divorce in the period that they get married or the period they choose to have a child.[18]

6. Child support transfers for divorcées are determined given the husband's income.

7. Consumption takes place.

The timeline described above is illustrated in Figure 3.

[17]When it comes to labor supply, people choose the extensive, not intensive, margin of married female labor force participation. All other labor supplies are fixed.

[18]This assumption is a little unusual and merits discussion. This shock helps convexify the choice space. From the point of view of the sub-time period when married couples make their allocation decisions, they are picking a probability that there will be a bad enough realization of this shock, causing divorce. The added smoothness aids in value function iteration. For an example of a similar technique used in this literature, see Regalia & Rios-Rull (2001).

2.5 Decision Making

How do households make their decisions in the model? Singles make a consumption/savings choice, along with a child's consumption choice if a single woman has a child from a previous marriage. They also have to decide whether or not to get married to a potential mate. Married agents have a similar consumption decision regarding savings and the consumption of adults and children, and must also decide whether the wife should work or not, whether or not to have a child if they currently do not have one (and are fecund), and/or whether or not to divorce.[19] We will now describe each household's problem recursively.

2.5.1 Value Functions and Policy Functions

Let's start with singles. The state vector for single households, (w, a, τ, k, ψ, t), consists of a wage shock w, their current asset holdings a, their child support obligations (or receipts) τ, their age t, and for women, an indicator function k, representing whether or not there is a child at home, and their idiosyncratic costs of working should they marry, ψ. Then the single value function for men and women without children can be written as follows:

$$V_g^s(w, a, \tau, 0, \psi, t) = \max_{c,a'} u(c, 0) + \beta E_{w',t'} B_g(w', a', \tau, k, \psi, t') \tag{10}$$
$$\text{s. t.}$$
$$c + b' = \phi_g w f_g(t) + (1 + r)b$$

where g represents gender. If a single woman has children from a previous marriage:

$$V_2^s(w, a, \tau, 1, \psi, t) = \max_{c,c_k,a'} u(c, c_k) - \psi_k + \beta E_{w',,t'} B_2(w', a', \tau, k, \psi, t') \tag{11}$$
$$\text{s. t.}$$
$$c + c_k + b' = \phi_g w f(t) + (1 + r)b + \tau$$

[19] We are implicitly assuming that people can only have children with one partner, as they cannot have children if they already have a child, and children only grow up once the people are no longer fecund. This assumption is made in order to simplify the model, but it is not a bad assumption in terms of the data. Using the National Survey of Family Growth, we find that only 12% of white fathers have had children with multiple partners.

Single households choose consumption c, and savings a', and if a single woman has children, she chooses their consumption c_k. Define the following policy functions associated with the single agent's problem: $c = P^s_{c,g}(w, a, \tau, k, \psi, t)$ for the consumption decision (by gender), $a' = P^s_{a,g}(w, a, \tau, k, \psi, t)$ for the savings decision, and $c_k = P^s_{c_k,g}(w, a, \tau, k, \psi, t)$ for choice of children's consumption. The continuation value for singles is the expectation of the value function $B_g(\cdot)$, which represents the value for a single before going through the marriage market (or the "bachelor" phase). The expectation is taken with respect to the income shocks next period and the possibility that the agent might grow older. We will elaborate on the value function $B_g(\cdot)$ slightly later in this section.

Let us discuss the decisions faced by married couples. The state vector, $(w, w^*, a, \tau, k, b, \psi, \gamma_f, t)$, consists of the husband's wage shock w, the wife's wage shock w^*, the household's asset holdings a, child support obligations from previous marriages τ, whether or not the couple has kids k, marital bliss level b, disutility from the wife working ψ, the bliss shock for having a child this same period γ_f, and age t. Thus,

$$V^m_g\left(w, w^*, a, \tau, k, b, \psi, \gamma_f, t\right) = u(\tilde{c}, \tilde{c_k}) + b + k\bar{b_k} - \psi\widetilde{l^f} + \gamma_f\widetilde{k} + \beta E_{w', w^{*\prime}, , b, \gamma_f, t'}\left[V^m_g\left(w', w^{*\prime}, \tilde{a}, \tau', k', b', \psi, \gamma'_f, t'\right)\right], \tag{12}$$

where $\widetilde{l^f}$ is the indicator function that the woman is working, and $\widetilde{k}$ is the indicator function that the family just had a child in this period. Note that if $k > 0$ (the couple already has children) or $\tau \neq 0$ (the husband had children in a previous relationship), then $\widetilde{k}$ must be 0, since couples can only have children once. People receive flow utility b from being married, which evolves stochastically as previously described. If they have children, and are married to the other parent of the children, they receive an extra flow utility of $\bar{b_k}$, which is deterministic.

Notice that there is no maximization done in the equation above. That is because the consumption, savings, labor supply, and fertility choices are the outcomes of Nash bargaining between the spouses. Given the possibility of divorce, spouses may well disagree on both the savings and fertility choices. This setup allows them to naturally bargain over the intertemporal decisions. The continuation value is given by the expected value of being married during the next period, where the expectation is taken with respect to the income shocks for both spouses, the new marital bliss shock b', the bliss shock for fertility next period γ_f, and whether or not they age t'.

To determine allocations, married people solve the following problem:

$$\left(\widetilde{c}, \widetilde{c_k}, \widetilde{a'}, \widetilde{l^f}, \widetilde{k}\right) = \arg \max_{c, c_k, a', l^f, k} \left(V_1^m\left(w, w^*, a, \tau, k, b, \psi, \gamma_f, t\right) - V_1^d\left(w, w^*, a, \tau, k, \psi, t\right)\right)^{\eta} \times$$
$$\times \left(V_2^m\left(w^*, w, a, \tau, k, \psi, \gamma_f, t\right) - V_2^d\left(w^*, w, a, \tau, k, \psi, t\right)\right)^{(1-\eta)}. \tag{13}$$

Here V_g^d represents the value of divorcing for each gender, discussed below. Policy functions for the married problem are defined as follows: $l^f = P_l^m\left(w, w^*, a, \tau, k, b, \psi, t\right)$ for the woman's labor force decision; $c = P_c^m\left(w, w^*, a, \tau, k, b, \psi, \gamma_f, t\right)$ for the consumption decision; $c_k = P_{c_k}^m\left(w, w^*, a, \tau, k, b, \psi, \gamma_f, t\right)$ for the consumption decision for children; $k = P_k^m\left(w, w^*, a, \tau, k, b, \psi, \gamma_f, t\right)$ for the decision to have a child; and $a' = P_a^m\left(w, w^*, a, \tau, k, b, \psi, \gamma_f, t\right)$ for the savings decision.

2.5.2 Divorce Value Functions

When people divorce, they split their assets evenly.[20] If there are children present, the man is required to pay a certain fraction of his income to the woman in future child support payments. These transfers are given by $\tau^n = \chi w f_1(t)$, for some fraction of income χ. The divorced woman will then maximize, period by period, on how much to spend on her children. The man is assumed to still derive utility out of his child's consumption in the future. To simplify the state space, we make the assumption that he believes that whatever the woman decides to spend on the children in the period of divorce is how much the children will consume in the future. While not choices, we denote the transfer "policy" as $\tau = P_{\tau,g}^m\left(w, w^*, a, \tau, k, b, \psi, \gamma_f, t\right)$.[21]

The value functions for men are then given by:

[20] Mazzocco et. al. (2007) analyze divorce settlements from the National Longitudinal Study of the High School Class of 1972 (NLS-72). They find that the average percentage of household wealth allocated to the wife upon divorce is 49.6%. The assumption of splitting assets evenly is therefore quite reasonable.

[21] Note that a divorced woman is not required to spend the money she receives in child support on the child. This allows the model to capture the flavor of Weiss & Willis (1985), which studies the inefficiencies of child support upon divorce.

$$V_1^d(w, w^*, a, \tau, k, \psi, t) = V_1^s(w, a/2, \tau^n, \psi, t) + \frac{1}{1-\beta(1-\eta)} \alpha u(\hat{c_k}(w^*, a/2, \tau^n, \psi, t)) - \omega_d(k). \tag{14}$$

For women:

$$V_2^d(w^*, w, a, \tau, k, \psi, t) = V_2^s(w^*, a/2, \tau^n, \psi, t), \tag{15}$$

where $\hat{c_k}(w^*, a/2, \tau^n, \psi, t)$ is the amount the woman chooses to spend on her child in the period of divorce. η is the rate at which children grow up and move out, which depends on the aging process in the model. The term multiplied by $\frac{1}{1-\beta(1-\eta)}$ is the sum of the expected discounted flow utilities that the father will get from the children's future consumption. We frontload this term to the divorce period so that it does not need to be carried as a state variable. Finally, $\omega_d(k)$ is the utility cost of divorce, which is allowed to vary based on the presence of children.

2.5.3 Divorce Decision

As mentioned previously, couples make a commitment to follow a certain allocation. They then draw a transitory bliss shock. Divorce is then unilateral. That is, a divorce will happen if and only if for at least one g disolving the marriage is preferable. λ is the transitory marital bliss shock discussed above. That is, people divorce if for at least one g:

$$V_g^d(w, w^*, a, \tau, k, \psi, t) > V_g^m(w, w^*, a, \tau, k, b, \psi, \gamma_f, t) + \lambda \tag{16}$$

The divorce policy function is given by $P_d^m(w, w^*, a, \tau, k, b, \psi, \gamma_f, t, \lambda)$. Integrating over all λ gives the probability a couple divorces:

$$\pi_d(w, w^*, a, \tau, k, b, \psi, \gamma_f, t) = \int P_d^m(w, w^*, a, \tau, k, b, \psi, \gamma_f, t, \lambda)\, d\Lambda(\lambda).$$

2.5.4 Marriage Decisions

We can now turn our analysis to the marriage phase. In the beginning of the period, every single person randomly draws a potential partner of the opposite gender from the distribution of available singles of that particular age. Each potential couple draws a marital bliss shock b from the distribution $\Upsilon(b)$. Each potential spouse will agree to marriage if and only if the continuation value in married life plus the marital bliss shock is larger than the continuation value as a single. A marriage occurs if and only if both agents agree to marriage. Formally, a marriage occurs if and only if

$$\underbrace{V_1^m(w, w^*, a + a^*, \tau, k, b, \psi^*, \gamma_f, t) + \lambda > V_1^s(w, a, \tau, k, t)}_{\text{male's decision}} \tag{17}$$

and

$$\underbrace{V_2^m(w, w^*, a + a^*, \tau, k, b, \psi^*, \gamma_f, t) + \lambda > V_2^s(w^*, a^*, \tau^*, k^*, \psi^*, t)}_{\text{female's decision}}.$$

Denote $z = (w, a, \tau, k, t)$ and $z^* = (w^*, a^*, \tau^*, k^*, \psi^*, t)$; that is, they are vectors of the state space of singles meeting in the marriage market. Let the indicator function $J(z, z^*, b, \gamma_f, \lambda, t)$ take a value of 1 if both people agree to the match and a value of 0 otherwise. Thus,

$$J(z, z^*, b, \gamma_f, \lambda, t) = \begin{cases} 1, & \text{if (17) holds,} \\ 0, & \text{otherwise.} \end{cases} \tag{18}$$

We can now write the value function before the marriage market (the "bachelor" phase):

$$\begin{aligned} B_g(w, a, \tau, k, \psi, t) \quad = \quad & \iiiint \{J(z, z^*, b, \gamma_f, \lambda, t) \left[V_g^m(w, w^*, a + a^*, \tau, k, b, \psi^*, \gamma_f, t)\right] \\ & + (1 - J(z, z^*, b, \gamma_f, \lambda, t)) V_g^s(w, a, \tau, k, \psi, t)\} \times \\ & \times d\widehat{\mathbf{S}_{\mathbf{g}^*}}(w^*, a^*, \tau^*, k^*, \psi^*, t) d\Upsilon(b) d\Gamma(\gamma_f) d\Lambda(\lambda), \end{aligned} \tag{19}$$

where $\widehat{\mathbf{S}_{\mathbf{g}^*}}(w^*, a^*, \tau^*, k^*, \psi^*, t)$ is the probability distribution of meeting a potential mate from the other gender (g^*) and age t. This will be elaborated on later. $\Upsilon(b)$ is the invariant distribution of marital bliss shocks, $\Gamma(\gamma_f)$ is the distribution of fertility bliss shocks, and $\Lambda(\lambda)$ is the distribution of transitory marital bliss shocks.

There are a few implicit assumptions about this marriage decision that need to be made explicit. First is that we assume people combine assets, which is a standard assumption in this literature. The second is that they see the first period of their marriage's fertility bliss draw when deciding to marry. Since one of the main reasons people marry is to have children, it seems natural that they know how much they want to have children upon deciding to marry. Finally, we also assume that the couple's transfer payments are derived from the man. That is, upon remarriage, women stop receiving transfers from their previous husbands, but men do not lose their obligations. This assumption is made for tractability: If we did not do it this way, we would have to keep track of the spouses' previous obligations individually in case of divorce. Legally, this is not accurate. Men are required to continue to pay child support to their ex-wives even if they remarry. However, voluntary extra payments for the children may well cease, and women may be less demanding of their former spouses given they have a new husband. The fact that men cut back on supporting their children when their ex-wives remarry is one of the main mechanisms studied in Chiappori and Weiss (2006). This assumption adds tractability to the model, and is reasonable, subject to the difficulty in studying voluntary supplementary child support payments and how payments change upon remarriage.

2.6 Equilibrium

Before we formally define the equilibrium for this economy, we must first elaborate on the distribution of single agents, since this distribution appears in the dynamic programming problem for bachelors. Note that, because of the endogenous marriage and divorce decisions, this distribution will be an equilibrium object and will also depend on the (equilibrium) distribution of married agents, since there are flows from married life into singlehood due to the presence of divorce.

Let $\mathbf{S_g}(w, a, \tau, k, \psi, t)$ and $\mathbf{M_g}(w, w^*, a, \tau, k, b, \psi, \gamma_f, t)$ denote the distributions of single and married agents, respectively. These distributions will evolve from period to period according to the policy functions and stochastic processes; the full equations can be found in Appendix C. Here, it is sufficient to note that there will be an updating operator $\mathbf{T}$ that relates the next period's distributions with the current period's distributions:

$$\begin{pmatrix} \mathbf{S_g}(w', a', \tau', k', \psi, t+1) \\ \mathbf{M_g}(w', w^{*\prime}, a', \tau, k', b', \psi', \gamma_f', t+1) \end{pmatrix} = \mathbf{T} \begin{pmatrix} \mathbf{S_g}(w, a, \tau, k, \psi, t) \\ \mathbf{M_g}(w, w^*, a, \tau, k, b, \psi, \gamma_f, t) \end{pmatrix}$$

In a stationary equilibrium, the distributions must be a fixed point of the updating operator $\mathbf{T}$.

We can now formally define the economy:

Definition 1 *A stationary equilibrium is a set of value functions for singles, couples, divorced individuals, and bachelors,* $V_g^s(w, a, \tau, k, \psi, t)$*,* $V_g^m(w, w^*, a, \tau, k, b, \psi, \gamma_f, t)$*,* $V_g^d(w, w^*, a, \tau, k, \psi, t)$ *,* *and* $B_g(w,a,\tau,k,\psi,t)$*; policy functions for single households* $P^s_{c,g}(w,a,\tau,k,\psi,t)$*,* $P^s_{c_k}(w,a,\tau,k,\psi,t)$*, and* $P^s_{a,g}(w,a,\tau,k,\psi,t)$*; policy functions for married households* $P^m_c(w,w^*,a,\tau,k,b,\psi,\gamma_f,t)$*,* $P^m_{c_k}(w,w^*,a,\tau,k,b,\psi,\gamma_f,t)$*,* $P^m_l(w,w^*,a,\tau,k,b,\psi,\gamma_f,t)$*,* $P^m_k(w,w^*,a,\tau,k,b,\psi,\gamma_f,t)$*, and* $P^m_a(w,w^*,a,\tau,k,b,\psi,\gamma_f,t)$*; a matching rule for singles* $J(z,z^*,b,\gamma_f,\lambda,t)$*; and a stationary distribution for singles* $\mathbf{S_g}(w,a,\tau,k,\psi,t)$ *and married people* $\mathbf{M}(w,w^*,a,\tau,k,b,\psi,\gamma_f,t)$*, such that:*

1. *The value function* $V^s_g(w,a,\tau,k,\psi,t)$ *and the policy functions* $P^s_{c,g}(w,a,\tau,k,\psi,t)$*,* $P^s_{c_k}(w,a,\tau,k,\psi,t)$*, and* $P^s_{a,g}(w,a,\tau,k,\psi,t)$ *solve the single's problem (10), given the value function for bachelors* $B_g(w,a,\tau,k,\psi,t)$ *and the distribution for singles* $\mathbf{S_g}(w,a,\tau,k,\psi,t)$*.*

2. *The value function* $V^m_g(w,w^*,a,\tau,k,b,\psi,\gamma_f,t)$ *and the policy functions* $P^m_c(w,w^*,a,\tau,k,b,\psi,\gamma_f,t)$*,* $P^m_{c_k}(w,w^*,a,\tau,k,b,\psi,\gamma_f,t)$*,* $P^m_l(w,w^*,a,\tau,k,b,\psi,\gamma_f,t)$*,* $P^m_k(w,w^*,a,\tau,k,b,\psi,\gamma_f,t)$*, and* $P^m_a(w,w^*,a,\tau,k,b,\psi,\gamma_f,t)$ *solve the couple's problems (12 and 13).*

3. *The value function* $B_g(w,a,\tau,k,\psi,t)$ *solves the bachelor's problem (19), given the value functions for singles and couples,* $V^s_g(w,a,\tau,k,\psi,t)$ *and* $V^m_g(w,w^*,a,\tau,k,b,\psi,\gamma_f,t)$*, and the matching rule* $J(x,x^*,b,b^*,\gamma,\lambda,a)$*.*

4. *The matching rule* $J(x,x^*,b,b^*,\gamma,\lambda,a)$ *is determined according to (18), taking as given the value functions* $V^s_g(w,a,\tau,k,\psi,t)$ *and* $V^m_g(w,w^*,a,\tau,k,b,\psi,\gamma_f,t)$*.*

5. *The divorce policy function* $P^m_d(w,w^*,a,\tau,k,b,\psi,\gamma_f,t,\gamma_d)$ *is determined according to (16), taking as given the value functions* $V^d_g(w,w^*,a,\tau,k,\psi,t)$*, and* $V^m_g(w,w^*,a,\tau,k,b,\psi,\gamma_f,t)$*.*

6. *The stationary distribution* $\mathbf{S_g}(w,a,\tau,k,\psi,t)$ *solves (20), taking as given the matching rule* $J(z,z^*,b,\gamma_f,\lambda,t)$*, the policy function* $P^s_{a,g}(w,a,\tau,k,\psi,t)$*, the divorce policy function* $P^m_d(w,w^*,a,\tau,k,b,\psi,\gamma_f,t,\gamma_d)$*, and the transfer policy rule* $P^m_\tau(w,w^*,a,\tau,k,b,\psi,\gamma_f,t)$*, while simultaneously* $\mathbf{M}(w,w^*,a,\tau,k,b,\psi,\gamma_f,t)$ *solves (21), taking as given the matching rule* $J(z,z^*,b,\gamma_f,\lambda,t)$*, the policy function* $P^m_a(w,w^*,a,\tau,k,b,\psi,\gamma_f,t)$*, and the divorce policy function* $P^m_d(w,w^*,a,\tau,k,b,\psi,\gamma_f,t,\lambda)$ *as expressed by* $\pi_d(w,w^*,a,\tau,k,b,\psi,\gamma_f,t,\lambda)$*.*

3 Mechanisms

We examine the effects of labor income volatility on family formation and dissolution. The data speaks of three trends: A decline in marriage rates, a decline in divorce overall, and an increase in divorce among the elderly. This section outlines the mechanisms through which income volatility affects all three of these phenomena.

3.1 Marriage

Earnings volatility influences marriage through three channels. The first is through the consumption commitments associated with marriage. One of the main reasons people get married is to have children. Children represent a consumption commitment, in the sense that once they exist they must consume, and having children is an irreversible decision. In this model, the amount of consumption devoted to a child is variable, but the mere fact that children exist makes their parents more risk averse. This is due to the fact that people spend substantial amounts of money on their children, which moves their own consumption to a steeper part of the utility function. When volatility increases, having children thus becomes thus less desirable. Since children are one of the main reasons to get married, marriage also becomes less desirable.

Another effect arises if higher income volatility induces higher income inequality. If workers are subject to more volatile persistent shocks, we should expect to see a more dispersed wage distribution in the population. That means that the marriage market will also be populated by a more dispersed distribution of potential mates. Hence, the option value of searching for a spouse increases as single individuals search longer for "better" matches. Conditional on a value for the non-economic reasons for marriage (love and children), if all potential mates are similar, then there is no reason to keep searching. However, if the distribution of potential mates is very dispersed, then people may search longer for a better spouse. This mechanism is highlighted, for example, by Gould and Paserman (2003).

The final effect comes from the availability of spousal insurance: Marriage allows for diversification of income risk since earnings fluctuations between spouses need not be perfectly correlated (though they are somewhat correlated both in the data and model). Hess (2004) analyzes the importance of insurance in marriage. For example, if a husband receives a bad income realization, the wife's income could help the household to smooth consumption. This possibility is not available for singles. Therefore, higher income volatility may make marriage more desirable due to this insurance aspect. The amount of insurance in this model is dependent on the level of divorce individuals expect to occur. If people expect divorce in the future, then spouses do not provide much insurance. Furthermore, the gender wage gap and depth of female labor force participation speak to the amount of a husband's income that a wife can replace. This model is rich enough to account for all of these facets of spousal insurance in a quantitative fashion.

3.2 Divorce

Marriage implicitly acts as a risk-sharing contract. When people are married, they combine resources for their mutual use. In these types of insurance contracts it is the person who receives the favorable shock that wants to leave the agreement. That is, if a woman receives a positive wage shock, and the man receives a negative wage shock, the woman is more likely to want to divorce the man. Krueger and Perri (2006) show that an increase in volatility makes people more willing to remain in insurance contracts, as the person who would normally want to leave now values the insurance more. Applied to this context, earnings volatility makes remaining married (with spousal insurance) more attractive, leading to less divorce.

The effect is stronger than perceived at first glance. Once people are married with consumption commitments, the value of spousal insurance rises. The idea here is that commitments, such as kids, create a lock-in. Upon divorce, the consumption commitment created by children does not go away.. The presence of children in people's lives make them more risk averse, as discussed previously. So, when people have children, they value insurance from their spouses much more than they did when single. This amplifies the mechanism, reducing rates of divorce–especially among the young.

3.2.1 Elderly Divorce

Once people retire, they experience no labor income risk. When their children grow up, they face fewer consumption commitments. The mechanism outlined above for how divorce is affected by volatility therefore does not apply to the elderly.

Thus, one might think that elderly divorce rates should not be affected by when volatility changes. However, volatility causes some people who might otherwise have divorced to remain married for insurance reasons. When they age, and that insurance is no longer relevant, they divorce. Simply put, volatility causes people to delay the decision to divorce, leading to both a decline in divorce rates for the young and a rise for the elderly.

4 Matching the Model to the Data

The model period is one year. There are four age groups in the model, representing the 20s, 30s, 40s/50s, and 60s, respectively. In order to set parameters for this model we follow a combination of setting parameters a priori, and estimating the model on micro data. In this section, we first discuss the parameters set a priori, and then the estimated parameters along with the identification strategy.

4.1 Parameters Calibrated a Priori

Some parameters are standard in the literature or have direct counterparts in the data. These parameters are listed in Table 1, and we briefly comment on them now.

Let's start with preference parameters. The time discount factor β is set to a standard value of 0.97. The coefficient of relative risk aversion (CRRA) is set to 2.0, which is also standard in the macroeconomic literature. For the parameter ψ that controls the degree of economies of scale in a household, we use the OECD equivalence scale. According to this scale, a second adult in the household only needs 70% of the consumption of the first adult in order to maintain the same standard of living. So we set $\psi = 0.7$.

Table 1: Parameters Set Using a Priori Information

Parameter	**Description**	**Value**	**Source**
Preferences			
β	Time discount factor	0.97	Standard
ζ	CRRA —consumption	2.0	Standard
ψ	Economies of scale	0.7	OECD equiv. scale
Income			
ρ	Correlation of spousal pers. shocks	0.25	Hyslop (2001)
$f(t)$	Age profile of income	–	U.S. Census
χ	Fraction of male income in child support	.25	Wisconsin guideline
Demographics			
π_a	Probability of aging	–	1/number of years in group
Prices			
–	Consumption floor	$2,640	Kaplan (2010)

A few parameters that control the amount of efficiency units of labor supplied by individuals can also be set here. The correlation of spousal persistent shocks, ρ, is set to 0.25, the number estimated by Hyslop (2001) using data from the Panel Study of Income Dynamics (PSID). The life cycle profile, $f(t)$, that controls the average level of efficiency units supplied at every age group is computed by calculating the mean income at each different age in the American Community Survey (ACS) relative to those in their 20s.[22] Child support payments are determined to be 25% of the man's income upon divorce assuming 2 children per family and using the Wisconsin guideline.. The Wisconsin guidelines are a set of rules for determining child support obligations, which do so by setting a fraction of income to be paid, where that fraction is dependent on the number of children. These guidelines are commonly used in the literature, for example by Takayama and Tanaka (2012).[23] Since there is stochastic aging in this model, we also calibrate the aging probabilities. These are given by the inverse of the number of years contained in an age group; i.e., since the 20s, 30s, and 60s age groups last for 10 years each, the probability of aging is 0.1. For the 40s/50s age group, it is 0.05.

Since this is a partial equilibrium model with respect to capital, we have to make some assumptions about prices. We set the interest rate to $r = 0.01$, a standard value for the risk-free interest rate. For the consumption floor, we use data provided by Kaplan (2010), who also studied income risk during this phase of life. Based on his calculations, the median monthly benefit for his National Longitudinal Survey of Youth (NLSY) sample is \$220/month, and covers the time period in question. We take this number (which amounts to \$2,640/year) to be our consumption floor.

4.2 Estimation

The remaining parameters are estimated using the Simulated Method of Moments. We first need a set of data moments that will inform on the parameters of the model. For a given set of parameter values, the model will generate statistics that can be compared to the data targets. The parameter values are then chosen to minimize some weighted distance between the model statistics and the data targets. Let Ω be the vector of parameters to be

[22] The results are very similar if we use data from the PSID. We use the larger sample from the ACS to get tighter estimates. Additionally, for the 20s age group, we use the observations for individuals aged 25-29 to abstract from issues regarding college.

[23] Since the 1980s, 15 states have adopted the Wisconsin guidelines. For more on the effectiveness of the child support system, see Garfinkel et. al. (1998).

estimated, and $g(\Omega)$ the difference between model moments and data moments at parameter Ω. We use a diagonal weighting matrix, W. The estimation procedure solves the following problem:

$$\min_{\Omega} g(\Omega)' W g(\Omega).$$

The vector of the standard errors for the estimator $\widehat{\Omega}$ is given by the square root of the diagonal of the following matrix:

$$V(\widehat{\Omega}) = \frac{1}{n} \left[g_1(\widehat{\Omega})' W g_1(\widehat{\Omega}) \right]^{-1} g_1(\widehat{\Omega})' W \Sigma W g_1(\widehat{\Omega}) \left[g_1(\widehat{\Omega})' W g_1(\widehat{\Omega}) \right]^{-1},$$

where Σ is the variance-covariance matrix of data moments, $g_1(\widehat{\Omega}) = \partial g(\widehat{\Omega})/\partial \Omega$, and n is the number of observations. The data moments derive from multiple data sets. The moments are independent across data sets. Therefore, Σ is a block diagonal matrix, with each block corresponding to a different data set. Each block is weighted by the number of observations in the block relative to the total number of observations.

In our case, we need to estimate thirteen parameters on top of the labor market parameters so we have the following vector of parameters to be estimated: $\Omega = \big(\alpha, \underline{c_k}, \psi, \psi_k, \omega(k), \bar{b_k}, \mu_b, \sigma_b, \delta_b, \sigma_{bb}, \mu_{\gamma_f},$ where Θ is a vector containing the estimated labor market parameters that are discussed in the next subsection.

4.2.1 Labor Market Parameters

For the data on income processes, we use data on white men from the Panel Study of Income Dynamics (PSID) for the years 1968–2009.[24] We first run a Mincerian regression for every year in the sample, controlling for education and a cubic polynomial in age. We then obtain our measure for residual income by generating the residuals of this regression. Using this measure for residual income, we estimate the parameters from (3) and (4) using Generalized Method of Moments (GMM). Note that we separately estimate the parameters for the process for married and single individuals since individuals from the two different groups might behave differently in the labor market.[25] The results of this estimation procedure

[24] For details on sample selection and estimation procedure, see Appendix D.

[25] The stock of divorced individuals is too small to estimate a third process in the PSID. We estimate the income processes both including divorced with married and single men, and do not find a substantial difference in our parameter estimates.

are reported in the Data column in Table 2 below.[26] Although this procedure is popular in the literature, estimates by marital status are not common. This difference aside, the variances of the shocks that we estimate are in line with the numbers reported by Heathcote, Storesletten, and Violante (2010) and Meghir and Pistaferri (2004).

To get a measure for the gender wage gap in the data, we run a Mincerian regression using log wages as a dependent variable and controlling for age, education, and a gender/age dummies using ACS data from 2005. We run this regression using observed wages for individuals that both work and report positive income. The coefficients on the gender dummies, age by age, is our data target for the gender wage gap ϕ. The value of this estimate is 0.74.

Using these estimates alone for the variance of the shocks to income processes and for the gender wage gap in the model generates sample selection problems. Specifically, for the income process, there is selection involved in who is married and who is not. If singles wait for good income shocks before getting married, then we would expect to truncate the top of the distribution of shocks into married people. This would make the observed shock process for singles not volatile enough. Additionally, for the gender wage gap, the estimate is obtained from a regression on observed wages. Clearly, there is selection involved in which women are working and which are not. To solve these problems, we take an indirect inference approach.[27] That is, our estimation procedure will make use of the model in order to estimate the parameters that control the income processes and the gender wage gap by adopting the following steps:

1. Guess parameter values for the income process for both married and single agents, as well as for the gender wage gap.

2. Solve and simulate the model in order to generate artificial data from the model.

3. Run the same GMM estimator on the simulated data as on actual data.

4. Check if the GMM estimates from the model match the data estimates.

We must emphasize that this estimation is performed in conjunction with the other parameters described in the next section.

[26] We also estimated the parameters for an age-specific income process in the spirit of Karahan and Ozkan (2010). Since the results were similar to the ones obtained here and we obtained tighter estimates for this simpler model, we opted for the simpler model described above.

[27] For a detailed description of this technique, see Gourieroux and Monfort (1996).

Table 2: Parameters for the Income Process

Parameter	Description	Data	Model	Param.
<u>**Married**</u>				
δ_m	Autoregressive coefficient	0.976	0.976	
$\sigma^2_{\eta,m}$	Variance of shock	0.016	0.015	
<u>**Singles**</u>				
δ_s	Autoregressive coefficient	0.929	0.922	
$\sigma^2_{\eta,s}$	Variance of shock	0.027	0.024	

4.2.2 Other Estimated Parameters

In addition to the labor market parameters discussed in the previous section, we still need to estimate thirteen additional parameters. As mentioned above, however, all parameters are estimated simultaneously.

In order to identify these parameters from the data, we try to choose data targets that will inform on the parameters we are estimating. Since we are jointly estimating all parameters, what follows is a heuristic argument as to how different data moments inform on model parameters.

The parameter ψ is the utility cost to a couple of having the woman work, representing the opportunity cost of what she could have been doing at home. ψ_k represents the additional cost of women working when they have kids. The statistics that most closely relate to these parameters are therefore the labor force participation rates (LFPR) of married women with and without kids. According to the American Community Survey (ACS), in 2005 the LFPR of married females ages 20-59 without children was 0.73, and 0.67 for those with children.

We now turn to parameters dealing with marriage and divorce. Marital bliss parameters include $\mu_b, \sigma_b, \delta_b$, and σ_{bb}, representing the mean bliss shock, the dispersion of the initial draw of this bliss shock, the persistence of the bliss shock during marriage, and the standard deviation of the innovations to the bliss shock, respectively. Furthermore, there is a parameter ω that is the utility cost of divorcing when the couple does not have children. Finally, there is a transitory bliss shock to smooth out the divorce decision, with mean 0, and standard deviation σ_λ. These are six parameters. To identify these paremeters, we target both marriage and divorce rates in the 20s, 30s, and 40s/50s (six targets). Marriage rates speak to the utility benefit of being married, along with the dispersion of these utilities. When bliss (μ_b) is high, marriage rates are high. When dispersion (σ_b) is high, people search longer for a better marital bliss, reducing marriage rates, but also affecting their profile over the life cycle. We further target the entire profile of the fraction of marriages that end in divorce by marital tenure (four targets). The shape of this profile helps identify how the value of being married changes over time, which is particularly useful for identifying the persistence of marital bliss (δ_b), along with the variance of innovations (σ_{bb}) and transitory shocks (σ_λ). The cost of divorce for people without kids (ω) is identified by the ratio of divorce rates of people aged 20-39 with kids to the divorce rates of those without kids. We calculate this ratio in the combined 2001, 2004, and 2008 Survey of Income and Program Participation (SIPP) to be 0.68.

Marriage and divorce rates, in the data, are much lower for the elderly than for the rest of the population. While we have discussed the identification of marriage and divorce parameters for the young and middle-aged, the model has difficulty replicating these rates for the elderly. There may be some changes that happen in old age that are not captured well in the model. For example, elderly people may be less willing to accomodate the changes in life that come with marriage and divorce. They may also have a harder time convincing their adult children to accept new family members. For these reasons, we estimate a separate divorce cost, ω_o, in order to allow the model to capture these life cycle changes. We do so by targeting divorce rates for people in their 60s.

Fertility is a separate decision from marriage in this model. The fertility parameters are $\bar{b}_k$ (the deterministic period utility flow of having a child present during marriage), and σ_{γ_f} (the standard deviation of the mean 0 fertility bliss shock). We target a notion of completed fertility, which is the fraction of women who have had a child by age 40, and a notion of the fertility rate, which is the rate at which childless married couples under 40 have children. The idea is that the period utility of having a child should speak to the total level of fertility, while the variance should give some indication of the timing. When variance is high, people are more likely to delay childbearing until later years while they wait for a good fertility shock, leading to a low fertility rate. When variance is low, they have children immediately. By targeting our notions of both completed fertility and fertility rates, we are able to separately identify these parameters.

A crucial parameter in this model is α. It governs how much people care about their children's consumption relative to their own. If α is very low, then children do not represent a significant consumption commitment. If α is high, then they represent a large commitment. The natural target in the data for this parameter is how much people spend on their children. Betson et al. (2001) estimate the fraction of a household's consumption expenditure that is attributable to one, two, or three children using data from the Consumption Expenditure Survey (CEX). This is not a straightforward calculation since it is not immediately clear how to divide the expenditures of certain goods (like shelter or utilities, for example) between the parents and the children. That is, the focus of the problem is to determine how parents reallocate consumption within the household in order to make room for the child's consumption. The idea Betson et al. use is to determine what the child's consumption is by comparing the welfare of childless couples and couples with one child, two, or three children. The authors then estimate Engel curves based on food expenditures in order to keep the standard of living constant. Following this methodology, the authors estimate the average fraction of consumption expenditures spent on one child to be 30.1%; on two children, 43.9%; and on three children, 52%. Since fertility has been roughly constant at two children per woman over this time period, we use the figure for two children as our target.[28] Note that this fraction does vary with the income of the household, which is not captured in the model. However, much of the heterogeneity that we observe in the data is not present in the model (for example, differences in education and individual fixed effects). Moreover, we are more interested in the type of risk an individual of a certain type faces throughout his or her lifetime, and not specifically in the cross-sectional variation observed in the data.

The final parameter is the subsistence consumption for children, $\underline{c_k}$. When $\underline{c_k}$ is large, and income grows over the life cycle, people choose to have children later in life, as the value of this cost becomes smaller relative to income. To identify this parameter, we therefore target the relative fertility rates between people in their 20s and 30s.[29] The ratio between this rate for those in their 20s and those in their 30s is the target. Using ACS data, we find this value to be 0.39.

In total, we use nineteen targets for the thirteen parameters discussed in this subsection, leaving the model over-identified. Note, however, that this is in addition to simultaneously estimating five labor market parameters on the five targets discussed in 4.2.1.[30]

[28] From Carter et. al. (2006), the Total Fertility Rate for white women in 1970 was 2.3 and 2.0 in 1998.

[29] We adjust the fertility rates to reflect the fact that some people never have a child, our notion of completed fertility discussed above. The rate we use is $\frac{NP}{NP+CL-NC}$, where NP is the number of new parents, CL is the number of childless adults, and NK is the number of people who never have children. This measure is calculated in the exact same way in both the data and the model, where the number of

Table 3: Estimated Parameters

Parameter	Description	Value
α	Utility weight on children	1.62
$\underline{c_k}$	Subsistence level for children	0.73
μ_γ	Mean marital bliss shock	-15.3
σ_b	Dispersion of marital bliss draw	10.57
δ_γ	Persistence of marital bliss shock	0.90
σ_λ	Transitory divorce shock	25.57
σ_{bb}	St. dev of mar. bliss innovation	1.76
ω_d	Utility cost of divorce	100
$\bar{b_k}$	Flow utility of having children	15.15
σ_f	St. deviation of fertility bliss shock	109.40
$\omega_{d,o}$	Divorce cost (old)	10.45
ψ_h	High cost of LFP	0.24
ψ_k	Extra cost of LFP with children	1.60

4.3 Model Fit (PRELIMINARY AND INCOMPLETE)

In this section, we discuss the fit of the model in regard to the moments used in the estimation. We estimate a total of 19 parameters by targeting 24 data moments. The estimated parameter values are reported in Table 2 (which contains the labor market parameters) and Table 3 (which contains the remaining parameters).

Table 4 compares the statistics generated by the model with the other data targets. Overall, the model does a good job matching these additional moments. Regarding marriage rates, for example, the model generates the decreasing likelihood of getting married as the individual gets older; i.e., marriage rates go down with age. Divorce rates also exhibit the same pattern, both in the model and in the data. Moreover, the data shows that people with children are less likely to divorce; the model is also consistent with this fact. Additionally, women without children participate more in the market, and the model also generates this pattern. Finally, the indirect inference approach yields income processes for the model that are very similar to their data counterparts.

people who never have children is the fraction who are childless at age 40.

[30]The five labor market parameters are the persistence and variance of shocks to income processes for both single and married people, along with the gender wage gap.

Table 4: Model Fit —Targeted Moments

Moment	**Data**	**Model**
Fraction HH expenditures on children	0.45	0.48
Ratio 20/30 adjusted fertility rates	0.39	0.46
Marriage rates: 20s/30s	8.29%	16.56%
Marriage rates: 40s/50s	3.36%	2.17%
Marriage rates: 60s	0.90%	0.71%
Divorce rates: 20s	4.08%	4.33%
Divorce rates: 30s	3.21%	2.84%
Divorce rates: 40s/50s	1.94%	1.90%
Divorce rates: 60s	0.73%	0.27%
Rel. divorce rates with/without children	0.68	0.54
% Women who have children by 40	0.77	0.82
Fertility rate w/o children	0.28	0.35
Married women's LFP w/o children	0.73	0.75
Married women's LFP w/ children	0.67	0.59
Single persistence	0.929	0.922
Single variance	0.027	0.024
Married persistence	0.976	0.976
Married variance	0.016	0.015

Furthermore, we target the profile of divorce rates by marital tenure. Figure 4 shows the fit between the model and the data. Overall, the model does quite well matching the general shape of this profile.

5 Results (PRELIMINARY AND INCOMPLETE)

In this section we analyze how important labor income volatility is for marriage and divorce decisions and how such changes might be heterogenous over an individual's life cycle. To do this, we perform a counterfactual simulation aimed at isolating only the effect of earnings volatility on these decisions. The counterfactual asks the following question: How would the world have looked in 2005 had income volatility not increased from its 1970 level? More precisely, we re-solve the model changing only the variances of the idiosyncratic income

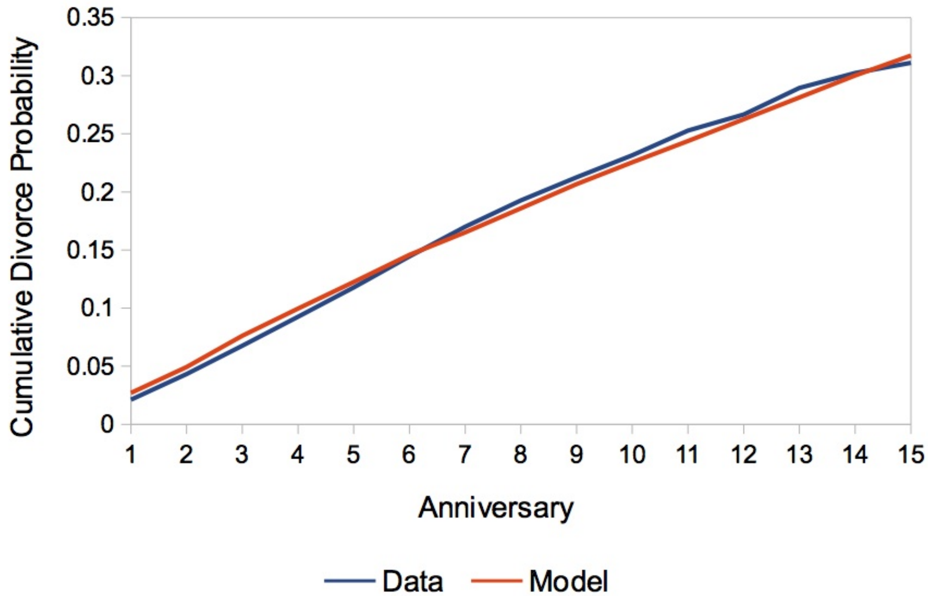

Figure 4: Cumulative Divorce Probability by Marital Tenure

shocks processes. We then compare how much closer to the 1970 data this new equilibrium is than the 2005 benchmark. The exercise thus isolates only the effects stemming from labor market volatility.

The results for the overall cumulative divorce probability can be inspected in Figure 5. First note that the data exhibits an interesting pattern across tenure length: In 1970, divorces were less likely to occur in the early years of marriage, and the opposite was true after the sixth anniversary. The counterfactual equilibrium generates exactly this same pattern. This can be explained by the fact that couples in the earlier years of their marriage are less likely to have children. These people are more inclined to divorce a spouse who receives a bad income shock and try their luck again in the marriage market. When volatility is high, such behavior is more prevalent. This mechanism can thus generate the qualitative pattern that we observe in the data.

To get a sense of how quantitatively important this mechanism is in explaining divorce, we report the percentage of the difference between the 1970 and 2005 data that the model can account for in Figure 6. For each anniversary, the model can explain a substantial fraction of the change and, on average, this mechanism accounts for 35% of the changes

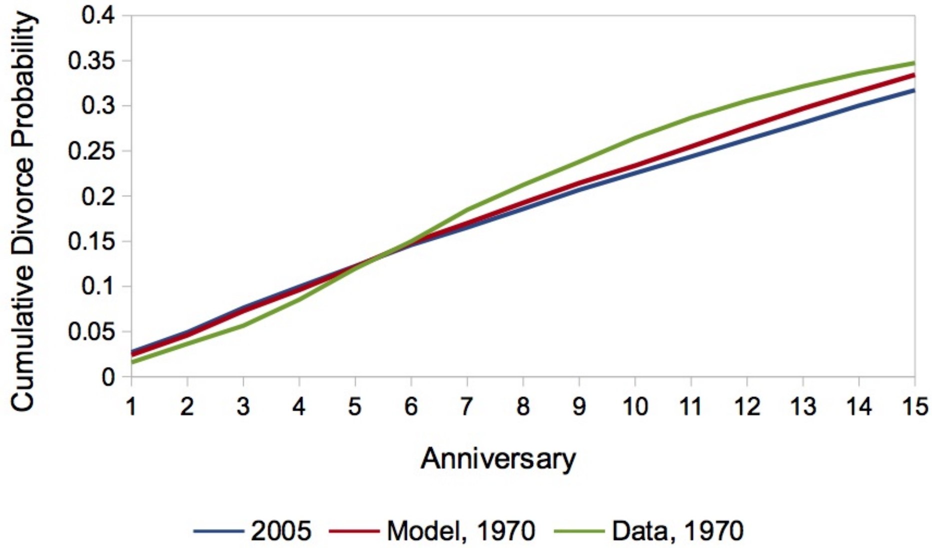

Figure 5: Cumulative Divorce Probability, Benchmark versus Counterfactual

in divorce between these two dates. The reason in the model for this observed decrease in divorce rates as earnings volatility increases is the lock-in effect caused by the presence of consumption commitments such as children. When volatility is higher, individuals prefer to stay married and enjoy the extra insurance and the economies of scale, since getting divorced means facing the commitments alone, even if the individual income changes.[31]

We noted above that divorce rates exhibit a different trend for older individuals. In particular, while divorce rates are decreasing overall, they are increasing for those above 60 years old (the "Gray Divorce Revolution"). We can examine what the model predicts for this age group in Table 5. Interestingly, even though the model predicts an overall decrease in divorce, it also predicts an *increase* in the divorce rates for those above 60 years old. The model predicts 9% of the observed changes between 1970 and 2005 data. The mere fact that it can generate this asymmetric behavior across age groups is noteworthy. The reason for this increase in divorce rates for the elderly comes from the fact that they don't face most of the lock-into-marriage effect of the younger generations: They are less likely

[31]Note that, exactly at the fifth anniversary, the model prediction appears incorrect. This is because the fifth anniversary is the cross-over point between the two lines, i.e., the divorce tenure profiles between 1970 and 2005 are essentially the same at that point. Hence, at that particular anniversary, there is essentially no difference in the data to be explained. This data point is therefore irrelevant.

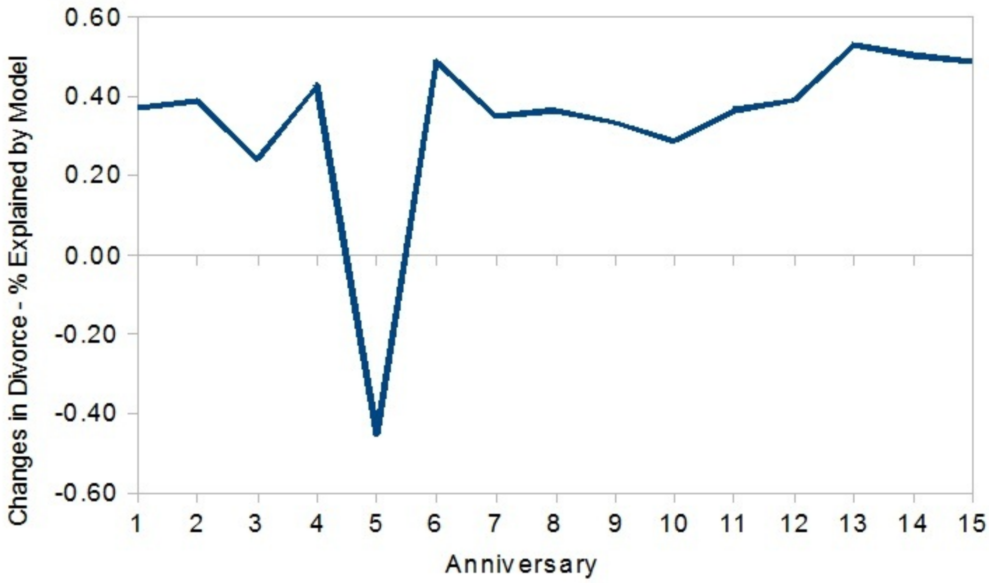

Figure 6: Fraction of the Changes in Divorce between 1970 and 2005 that can be Accounted For with Increased Earnings Volatility

to have children in their household (their children are older and have already moved out) and they have less need for spousal insurance (either their income is relatively stable due to retirement or they have a short time span in their working lives to worry about such fluctuations).

Finally, we can also inspect the model's predictions concerning marriage. Predicted changes in marriage rates for the model are also reported in Table 5. We observe a decrease in the marriage rates between 1970 and 2005 data; the model is able to explain 17% of those changes. What this means is that the added gains from spousal insurance are quanti-

Table 5: Marriage and Divorce Rates - Counterfactual

Moment of Interest	**% Change Data (1970-2005)**	**% Change Model(1970-2005)**	**Model/Data Change**
Marriage Rates	↓47%	↓8%	17%
Divorce Rates (60s)	↑ 300%	↑ 16%	9%

tatively less important than the adverse risk effects caused by the presence of consumption commitments within marriage, even though these commitments are *choices* made by the agents.

6 Conclusions

There have been stark changes in the way people form and dissolve families in the US over the last four decades. In general, people are less likely to both get married or divorced. But for older generations, the story is somewhat different: they are *more* likely to get a divorce. What could explain these positive comovements on marriage and divorce and the differential trend in divorce across age groups? In this paper, we evaluate the importance of rising earnings volatility to explain all of these changes.

When labor income is more volatile, agents are less likely to get married due to the consumption commitments (such as children) associated with marriage, even though these commitments might arise from some of their choices. The commitments also create a lock-into-marriage effect that causes a drop in divorce. For the older generations, these commitments are less important (their children have already grown up and moved out). Thus, the elderly are more prone to divorce.

In order to assess the quantitative importance of this mechanism, we build a life cycle equilibrium model of the marriage market in which marriage, divorce, and fertility are all endogenous decisions. We take the model to the data by calibrating its parameters to hit certain data targets for 2005. By performing a counterfactual analysis, we find that the model is able to explain 35% of the observed decline in divorce rates, 9% of the *rise* in divorce rates among the elderly, and 17% of the fall in marriage rates.

The ideas presented in this paper can be very important for a variety of economic issues. For example, these changes in family formation can affect asset accumulation in response to higher risk. On the other hand, by changing family composition, earnings volatility might also have an impact on the levels of parental investment on children. We leave this for future research.

References

[1] Betson, David, Carol Huffine, David Price, and Jane Venohr. 2001. *A Review of California's Statewide Uniform Child Support Guideline*. Denver, Colorado: Policy Studies Inc.

[2] Brown, Susan, and I-Fen Lin. 2012. "The Gray Divorce Revolution: Rising Divorce among Middle-aged and Older Adults, 1990-2009." Working Paper Series WP-12-04. National Center for Family & Marriage Research at Bowling Green State University.

[3] Burdett, Ken and Melvyn G. Coles. 1997. "Marriage and Class," *Quarterly Journal of Economics*, 112(1), 141-68.

[4] Carter, Susan B. et al., eds. 2006. *Historical Statistics of the United States, Earliest Times to the Present.* Millennial Edition. New York: Cambridge University Press.

[5] Chakraborty, Indraneel, Hans Holter, and Serhiy Stepanchuk. "Marriage Stability, Taxation and Aggregate Labor Supply in the U.S. vs. Europe." 2012. Mimeo.

[6] Chiappori, Pierre-Andre and Yoram Weiss. 2007. "Divorce, Remarriage and Child Support," *Journal of Labor Economics*, 25(1), 37-74.

[7] Chetty, Raj, and Adam Szeidl. 2007. "Consumption Commitments and Risk Preferences," *Quarterly Journal of Economics*, 122, no. 2: 831-877.

[8] Clarke, S.C. (1995). "Advance Report of Final Divorce Statistics, 1989 and 1990," *Monthly Vital Statistics Report*, Vol 43, No. 9. Hyattsville, MD: National Center for Health Statistics.

[9] Farnham, Martin, Lucie Schmidt and Purvi Sevak. 2011. "House Prices and Marital Stability," *American Economic Review: Papers & Proceedings*, 101, no. 3: 615-619.

[10] Gottschalk, Peter and Robert Moffitt. 1994. "The Growth of Earnings Instability in the U.S. Labor Market," *Brookings Papers on Economic Activity*, 25, 217-272.

[11] Garfinkel, Irwin, Cynthia Miller, Sara S. McLanahan, and Thomas L. Hanson. 1998. "Deadbeat Dads or Inept States? A Comparison of Child Support Enforcement Systems,"*Evaluation Review.* 22(6): 717-750.

[12] Gourieroux, Christian and Alain Monfort. 1996. *Simulation-Based Econometric Methods*. New York: Oxford University Press.

[13] Gould, Eric and Daniele Paserman. 2003. "Waiting for Mr. Right: Rising Inequality and Declining Marriage Rates," *Journal of Urban Economics*, 53 (2), 257-281.

[14] Greenwood, Jeremy and Nezih Guner. 2009. "Marriage and Divorce since World War II: Analyzing the Role of Technological Progress on the Formation of Households," *NBER Macroeconomics Annual 2008*, 23: 231-276.

[15] Greenwood, Jeremy, Nezih Guner, and John Knowles. 2003. "More on Marriage, Fertility, and the Distribution of Income," *International Economic Review*, 44 (3), 827-862.

[16] Heathcote, Jonathan, Fabrizio Perri, and Gianluca Violante. 2010. "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States 1967-2006," *Review of Economic Dynamics*, 13 (1), 15-51.

[17] Hellerstein, Judith and Melinda Morrill. "Booms, Busts, and Divorce." 2011. Mimeo.

[18] Hess, Gregory. 2004. "Marriage and Consumption Insurance: What's Love Got to Do with It?", *Journal of Political Economy*, 112(2), 290-318.

[19] Hyslop, Dean. 2001. "Rising U.S. Earnings Inequality and Family Labor Supply: The Covariance Structure of Intrafamily Earnings," *American Economic Review*, 91 (4), 755-777.

[20] Kaplan, Greg. 2010. "Moving Back Home: Insurance Against Labor Market Risk," Federal Reserve Bank of Minneapolis Working Paper 667, mimeo.

[21] Kambourov, Gueorgui and Iourii Manovskii. 2008. "Rising Occupational and Industry Mobility in the United States: 1968-1997," *International Economic Review*, 49 (1), 41-79.

[22] Kambourov, Gueorgui and Iourii Manovskii. 2009. "Occupational Mobility and Wage Inequality," *Review of Economic Studies*, 76 (2), 731-759.

[23] Karahan, Fatih and Serdar Ozkan. 2010. "On the Persistence of Income Shocks over the Life Cycle: Evidence and Implications," PIER Working Paper 10-012.

[24] Katz, Lawrence F. and David H. Autor. 1999. "Changes in the Wage Structure and Earnings Inequality," in the *Handbook of Labor Economics*, ed. O. Ashenfelter and D. Card, edition 1, volume 3, chapter 26, 1463-1555.

[25] McKelvey, Richard D. and Thomas R. Palfrey. 1995. "Quantal Response Equilibrium for Normal Form Games,"*Games and Economic Behavior*, 10, 6-38.

[26] Kopecky, Karen A. and Richard M. H. Suen. 2010. "Finite State Markov-chain Approximations to Highly Persistent Processes," *Review of Economic Dynamics*, 13 (3), 701-714.

[27] Krueger, Dirk and Fabrizio Perri. 2006. "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory," *Review of Economic Studies*, 73 (1), 163-193.

[28] Mazzocco, Maurizio, Claudia Ruiz, and Shintaro Yamaguchi. "Labor Supply, Wealth Dynamics, and Marriage Decisions." 2007. Mimeo.

[29] Meghir, Costas and Luigi Pistaferri. 2004. "Income Variance Dynamics and Heterogeneity," *Econometrica*, 72 (1), 1-32.

[30] Neeman, Zvika, Newman, Andrew, and Claudia Olivetti. "Are Career Women Good For Marriage?" 2008. Mimeo.

[31] Pericoli, Felippo and Luigi Ventura. 2011. "Family Dissolution and Precautionary Savings: An Empirical Analysis," MPRA Paper No. 36354.

[32] Postlewaite, Andrew, Larry Samuelson, and Dan Silverman. 2008. "Consumption Commitments and Employment Contracts," *Review of Economic Studies*, 75, 559-578.

[33] Regalia, Ferdinando, Jose-Victor Rios-Rull, and Jacob Short. "What Accounts for the Increase in Single Households?" 2008. Mimeo.

[34] Rotz, Dana. "Why Have Divorce Rates Fallen? The Role of Women's Age at Marriage." 2011. Mimeo.

[35] Santos, Cezar and David Weiss. "Why Not Settle Down Already? A Quantitative Analysis of the Delay in Marriage." 2012. Mimeo.

[36] Sommer, Kamila. 2009. "Fertility Choice in a Life Cycle Model with Idiosyncratic Uninsurable Earnings Risk," Ph.D. Dissertation, Georgetown University.

[37] Stevenson, Betsey and Justin Wolfers. 2007. "Marriage and Divorce: Changes and their Driving Forces," *Journal of Economic Perspectives*, 21(2) 27-52.

[38] Stevenson, Betsey and Justin Wolfers. 2011. "Trends in Marital Stability," in *Research Handbook on the Economics of Family Law*, ed. L. Cohen and J. Wright, Chapter 4, 96-108.

[39] Takayama, Naoki and Satoshi Tanaka. 2012. "Men's Rush to Marriage: Implications of the Child Support Enforcement Policy for Marriage, Fertility, and Long-Term Inequality." 2012. Mimeo.

[40] Voena, Alessandra. "Yours, Mine and Ours: Do Divorce Laws Affect The Intertemporal Behavior Of Married Couples?" 2012. Mimeo.

[41] Weiss, Yoram and Robert Willis. 1985. "Children as Collective Goods and Divorce Settlements," *Journal of Labor Economics*, 3(3), 268-292

A Data Sources

This appendix describes the sources of the data for selected tables and figures in the paper that contain actual data.

Figure 1: The data for marriage and divorce rates comes from Carter et. al. (2006) though 1995. Data from 1996 through 1999 comes from U.S. Census Bureau, Statistical Abstract of the United States: 2004-2005. Data starting in 2000 comes from CDC/NCHS National Vital Statistics System (NVSS). Some statistics are reported in these sources as rates per 1000 population. To convert to rates per 1000 married or single women, we use the Integrated Public Use Microdata Series (IPUMS) Current Population Survey (CPS) for before 2000. After 2000, we use the American Community Survey (ACS). Divorce rates are per 1000 married women. Marriage rates are per 1000 single women, above age 15. Due to data limitations, this series is not restricted by race/education.

Figure 2: The data for the cumulative divorce probabilities by marital tenure comes from three waves of the Survey of Income and Program Participation (SIPP). We combine the 2001, 2004, and 2008 SIPPs, and follow the procedure outlined in Stevenson and Wolfers (2011). We restrict attention to the white population in order to maintain consistency with the rest of the data in this paper. **Figures 7 and 8** perform the same analysis by education group. For marriages from 1990, the figures are somewhat noisy towards the later anniversaries; the reason for this is the relatively smaller sample size.

B Cumulative Divorce by Education

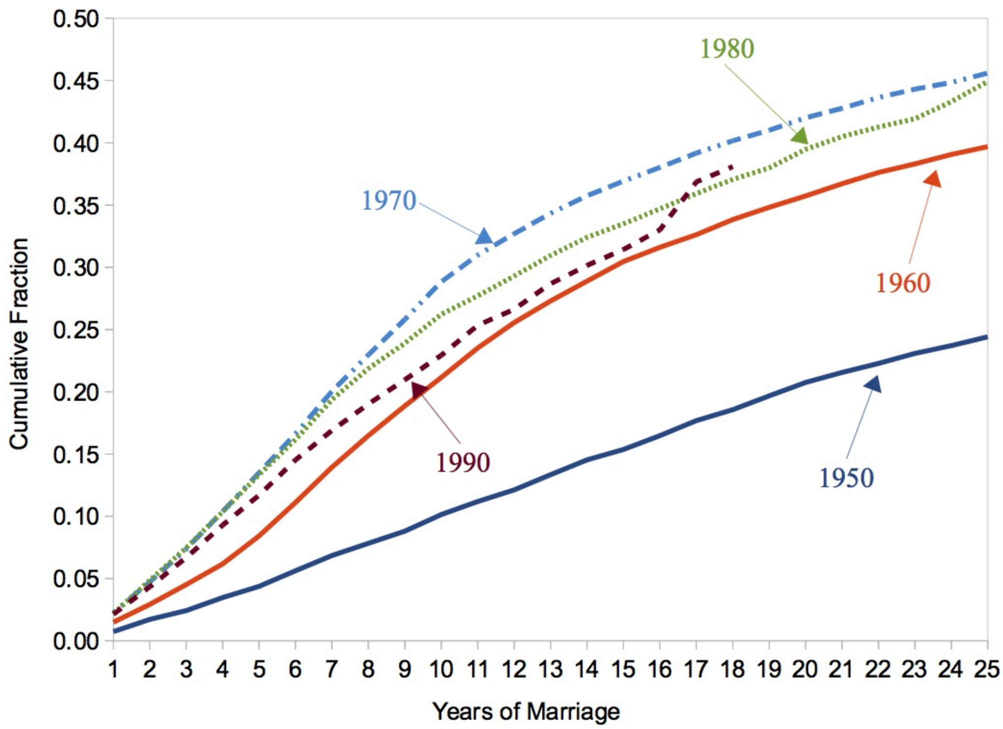

Figure 7: Cumulative Divorce Probabilities, White Americans, High School Educated or More, First Marriage. Authors' Calculation from the SIPP 2001-2008, following Stevenson & Wolfers (2011)

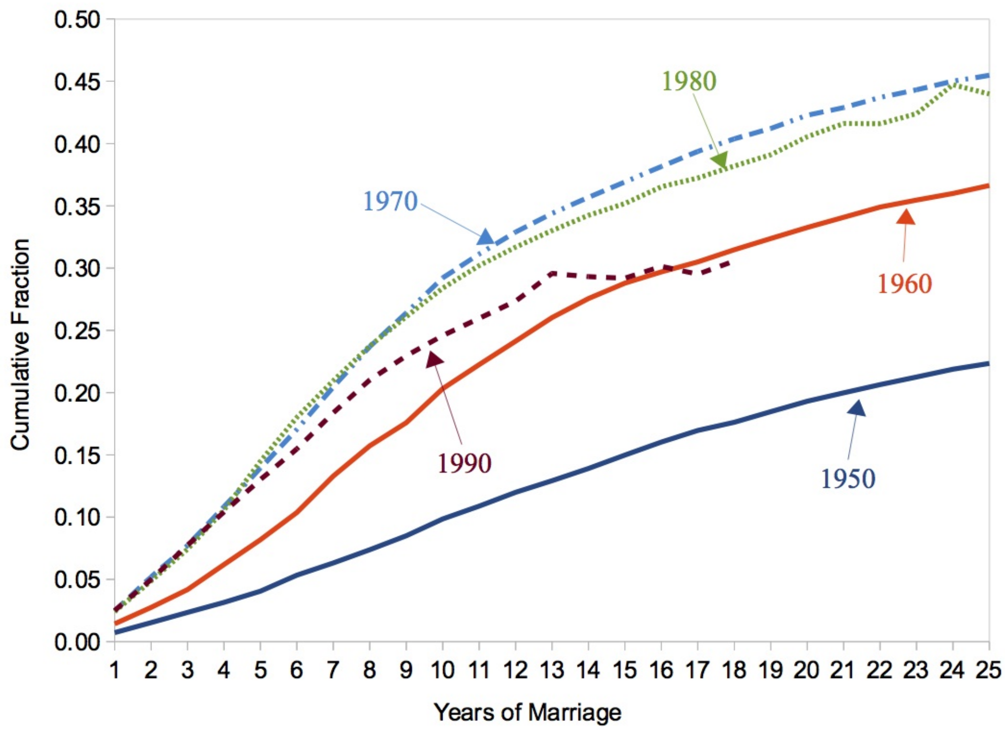

Figure 8: Cumulative Divorce Probabilities, White Americans, High School Educated or Less, First Marriage. Authors' Calculation from the SIPP 2001-2008, following Stevenson & Wolfers (2011).

C Stationary Distributions

The non-normalized stationary distribution for singles is given by

$$\begin{aligned}
\mathbf{S_g}(w', a', \tau', k', \psi, t+1) = & \iiiint \pi(t)\,(1 - J(z, z^*, b, \gamma_f, t)) \times \\
& \mathcal{I}(P_a^s(w, a, \tau, k, \psi, t) \leq a')\times \\
& \mathbf{S_g}(w, a, \tau, k, \psi, t) d\widehat{\mathbf{S_{g^*}}}\,(w^*, a^*, \tau^*, k^*, \psi^*, t)\, d\Upsilon(b) d\Gamma(\gamma_f) d\mathbf{W}^s(w', w) \\
& + \iiiint (1 - \pi(t+1))\,(1 - J(z, z^*, b, \gamma_f, t+1)) \times \\
& \mathcal{I}(P_a^s(w, a, \tau, k, \psi, t+1) \leq a')\times \\
& \mathbf{S_g}(w, a, \tau, k, \psi, t+1) d\widehat{\mathbf{S_{g^*}}}\,(w^*, a^*, \tau^*, k^*, \psi^*, t+1)\, d\Upsilon(b) d\Gamma(\gamma_f) d\mathbf{W}^s(w', w) \\
& + \iiiint \pi(t)\pi_d(w, w^*, a, \tau, k, b, \psi, \gamma_f, t)\times \\
& \mathcal{I}(P_a^s(w, a, P_{\tau,g}^m(w, w^*, a, \tau, k, b, \psi, \gamma_f t), k, \psi, t) \leq a')\times \\
& \mathbf{M}(w, w^*, a, \tau, k, b, \psi, \gamma_f, t) d\mathbf{W}^m(w', w) \\
& + \iiiint (1 - \pi(t+1))\pi_d(w, w^*, a, \tau, k, b, \psi, \gamma_f, t+1)\times \\
& \mathcal{I}(P_a^s(w, a, P_{\tau,g}^m(w, w^*, a, \tau, k, b, \psi, \gamma_f t+1), k, \psi, t+1) \leq a')\times \\
& \mathbf{M}(w, w^*, a, \tau, k, b, \psi, \gamma_f, t+1) d\mathbf{W}^m(w', w),
\end{aligned} \tag{20}$$

where g^* represents the opposite gender and $\mathbf{W}^s$ represents the wage shock process for singles defined above. $\mathcal{I}(statement)$ is an indicator function that takes the value of 1 when *statement* is true and 0 otherwise. Singles aged $a = 1$ are distributed over wages according to the invariant distribution of $\mathbf{W}^s$. Notice that there are terms both for singles not marrying (the first terms) and married people divorcing (the later terms). Additionally, $\pi(t)$ represents the probability of a person aging from age t. $\widehat{\mathbf{S_g}}(w, a, \tau, k, \psi, t)$ denotes the normalized distribution for singles that determines the probability that single agents will meet in the marriage market, and is defined by

$$\widehat{\mathbf{S_g}}(w, a, \tau, k, \psi, t) \;\; = \;\; \frac{\mathbf{S_g}(w, a, \tau, k, \psi, t)}{\int d\mathbf{S_g}(w, a, \tau, k, \psi, t)}.$$

Additionally, the distribution of married people is an equilibrium object as there are flows from marriage into singles, affecting the singles distribution. It is given by

$$\begin{aligned}
\mathbf{M}(w', w^{*\prime}, a', \tau, k', b', \psi', \gamma_f', t+1) = & \iint \iiiint \pi(t)\left(J(z, z^*, b, \gamma_f, t)\right) \times \\
& \mathcal{I}(P_a^m(w, w^*, a+a^*, \tau, k^*, b, \psi^*, \gamma_f, t) \leq a') \times \\
& \mathcal{I}(P_k^m(w, w^*, a+a^*, \tau, k^*, b, \psi^*, \gamma_f, t) = k')\mathbf{S_g}(w, a, \tau, k, \psi, t) \times \\
& d\widehat{\mathbf{S}_{\mathbf{g}^*}}(w^*, a^*, \tau^*, k^*, \psi^*, t) d\Upsilon(b) d\Gamma(\gamma_f) \times \\
& d\mathbf{W}^m(w', w, w^*, w'^{\star}) d\Gamma(\gamma_f') d\Xi(b, b') \\
& + \iint \iiiint (1-\pi(t+1))\left(J(z, z^*, b, \gamma_f, t+1)\right) \times \\
& \mathcal{I}(P_a^m(w, w^*, a+a^*, \tau, k^*, b, \psi^*, \gamma_f, t+1) \leq a') \times \\
& \mathcal{I}(P_k^m(w, w^*, a+a^*, \tau, k^*, b, \psi^*, \gamma_f, t+1) = k')\mathbf{S_g}(w, a, \tau, k, \psi, t+1) \times \\
& d\widehat{\mathbf{S}_{\mathbf{g}^*}}(w^*, a^*, \tau^*, k^*, \psi^*, t+1) d\Upsilon(b) d\Gamma(\gamma_f) \times \\
& d\mathbf{W}^m(w', w, w^*, w^{*\prime}) d\Gamma(\gamma_f') d\Xi(b, b') \\
& + \iint \pi(t)(1-\pi_d(w, w^*, a, \tau, k, b, \psi, \gamma_f, t)) \times \\
& \mathcal{I}(P_a^m(w, w^*, a, \tau, k, b, \psi, \gamma_f, t) \leq a') \times \\
& \mathbf{M}(w, w^*, a, \tau, k, b, \psi, \gamma_f, t) \\
& \mathcal{I}(P_k^m(w, w^*, a+a^*, \tau, k^*, b, \psi^*, \gamma_f, t) = k') \times \\
& d\mathbf{W}^m(w', w, w^{\star}, w^{*\prime}) d\Xi(b, b') \\
& + \iint (1-\pi(t+1))(1-\pi_d(w, w^*, a, \tau, k, b, \psi, \gamma_f, t+1)) \\
& \mathcal{I}(P_a^m(w, w^*, a, \tau, k, b, \psi, \gamma_f, t+1) \leq a') \times \\
& \mathcal{I}(P_k^m(w, w^*, a+a^*, \tau, k^*, b, \psi^*, \gamma_f, t+1) = k') \times \\
& \mathbf{M}(w, w^*, a, \tau, k, b, \psi, \gamma_f, t+1) \\
& d\mathbf{W}^m(w', w, w^*, w'^*) d\Xi(b, b'),
\end{aligned} \tag{21}$$

where $\mathbf{W}^m(w', w, w^*, w^{*\prime})$ represents the probability of husband and wife jointly moving from shock (w, w^*) to (w', w'^*). The first term shows how single people marrying and gaining in age are reflected. The second term shows single people marrying and not gaining in age. The third term is married people remaining married and gaining in age. The final term is married people remaining married and not gaining in age.

D Estimation of Income Processes

We use data from the PSID for all waves between 1968 and 2009. As described in the text, we separately estimate the processes described in Section 2.1 for married and single individuals. We use data for male respondents that satisfies the following criteria for at least three years (which need not be consecutive): (i) the individual reported positive earnings and hours; (ii) his age is between 18 and 64; (iii) he worked between 520 and 5100 hours during the year; and (iv) he had an hourly wage above half of the prevailing minimum wage at the time. We also exclude people from the Survey of Economic Opportunity (SEO) sub-sample in 1968. These criteria are fairly standard in the literature. See, for example, Karahan and Ozkan (2010).

First, in order to generate the residual earnings, we run a cross section Mincerian regression for each year, controlling for education and a polynomial in age. Residuals generated from these regressions are used in the estimation procedure. We estimate a slightly modified version of the processes described in Section 2.1 in order to include individual fixed effects (which are not present in the model). We estimate time-varying variances for each shock for each year and HP-filter these time series for the variances. These HP-filtered variances for the shocks are reported in Table 2. The standard errors are computed using a bootstrap procedure. For a formal proof of identification of the parameters, see Karahan and Ozkan (2010).

Starting in 1997 the PSID becomes a biennial survey. Heathcote et. al. (2010) describe how identification of income process parameters is still feasible with this change.

E Computation

We solve two steady states for the model; one that represents the world in 2005 and another that represents a counterfactual world in 2005, namely what things would have looked like had labor income volatility *not* increased. All parameters except those governing the volatility of the labor market are kept constant between the baseline and counterfactual worlds. In order to do so, we need to numerically solve this model.

Computation for this model is done in Fortran 90 using OpenMP protocols. Below we describe the value function iteration process used to solve the model in general, plus a numerical routine we develop in order to solve the married household's bargaining problem. We run into convergence issues when solving the married household's bargaining problem, and formulate a novel approach using the notion of a Quantal Response Equilibrium (see McKelvey and Palfrey, 1995). Both the problem and the solution are discussed in detail below.

E.1 Value Function Iteration

The model is solved by iterating on the value functions and distributions of single/married people until convergence, using the following standard algorithm:

1. Guess value functions and distributions.
2. Given the value functions and distributions, calculate policy functions and the implied value functions.
3. Use the policy function to calculate the implied distributions.
4. Check that the guessed value functions and guessed distributions are within a small tolerance level of the implied value functions and distributions. If so, then end the algorithm. Otherwise, continue.

Notice that this a rather substantial fixed point problem. We need to solve for value functions and distributions for each of these types of agents over all the age groups.

E.2 Solving the Married Household's Bargaining Problem

Solving the Nash bargaining problem (13) is computationally difficult. There are 5 control variables, (c, c_k, a', l^f, k), that need to be jointly solved for, over a bargaining space that has numerically defined, rather than analytic functions. Furthermore, there is no guarantee that a solution within marriage exists– people may prefer divorce. In the case when divorce is preferred, surplus is negative.

Given these difficulties in solving the Nash bargaining problem, we follow Greenwood et. al. (2003) and convert the Nash bargaining problem to an equivalent Pareto problem:

$$\max_{c,c_k,a',l^f,k} \theta_p(u(c) + u(c_k) - l^f\psi + k\gamma_f + \beta EV_1^m(w', w^{*\prime}, a', \tau', k, b', \psi, \gamma_f', t'))+ \tag{22}$$
$$(1-\theta_p)(u(c) + u(c_k) - l^f\psi + k\gamma_f + \beta EV_2^m(w', w^{*\prime}, a', \tau', k', b', \psi, \gamma_f', t'))$$

Notice that the couple agrees on the value of today's utility flows from consumption, fertility, and labor choices, but may disagree on the future utilities. For simplicity, we drop the flow utility from marital bliss from the above equation, as it is simply a shift in utility weights, and does not affect any margins of decision making. The Pareto weight θ_p that guarantees that (22) is equivalent to (13) is given by:

$$\theta_p = \frac{V_2^m - V_2^d}{(V_2^m - V_2^d) + (V_1^m - V_1^d)} \tag{23}$$

The solution to this problem is attained by implementing the following steps:

1. Given the guesses for V^m & V^d, calculate the Pareto weight θ_p.
2. Solve the following for both the cases when the couple does and does not decide to have children (if they do not have children, $c_k = 0$):
   (a) Given choices for savings and labor force participation of the wife (LFP), we know exactly how much the couple will spend on their own consumption and their child's consumption from the problem's first order conditions. Solve thus for the optimal choices of c and c_k. This is the "continuous" choice, as the numbers can be computed exactly given the Pareto weights and cash on hand.
   (b) Given how resources will be allocated at each savings/LFP choice, pick the best savings/LFP decision using grid search. This is the "discrete" choice since there are 2 LFP choices (work or not), and the number of asset points on the grid.
3. Calculate the cutoff strategy in γ_f for the fertility decision if the couple can choose to have children.
4. Calculate the cutoff strategy in the transitory bliss shock. That is, determine how

high a realization of the transitory bliss shock is needed in order to keep the couple married.

As we state above, the couple agrees on the value of consuming/working today, but potentially not on the dynamic choices of asset holdings and fertility. That is, left to their own devices, each spouse might make a different consumption/savings choice. The family's allocation is a compromise between the preferences of the husband and wife, with that compromise determined by the Pareto weight θ_p.

The computational issue that arises is that small changes in θ_p can lead to discrete changes in outcomes, and thus discrete changes in the value functions for each spouse. These changes in value functions induce a change in θ_p, as seen in (23). This causes cycling in the value functions during iteration and prevents convergence of the algorithm. For example, consider a case where the man has relatively high bargaining power. That is, θ_p is high. Then, assume that the allocation the household picks is relatively close to husband's optimal choice. This causes the value function for the husband to be large and that of the wife to be small. In the next iteration of the value function, the Pareto weight for the husband will *decrease*, giving the wife higher bargaining power. The household allocation then becomes more similar to the wife's prefered outcome, leading to a rise in θ_p, restarting the cycle. In short, the household's objective function is twin-peaked in the statespace, and cycles between these two peaks as the Pareto weight changes. While overall household utility converges, as the household is essentially indifferent between the two choices, individual utility within the household varies widely between the iterations.

To solve this issue, we approximate the solution by having households make choices along the lines of a Quantile Response Equilibrium, as introduced by McKelvey and Palfrey (1995). That is, households choose all feasible allocations with positive probabilities, rather than make just one optimal choice. These probabilities are given by:

$$P(i) = \frac{e^{\chi V(i)}}{\sum_i e^{\chi V(i)}} \tag{24}$$

That is, the choice of allocation i is given by $P(i)$. $V(i)$ is the household utility from choice i, as determined by the Pareto problem (22). χ is a numerical parameter. Choices that lead to higher household utility get picked with higher probability. χ determines how weighted the probability distribution is towards choices with higher household utility. When χ is close to 0, all choices are "played" with equal probability. When $\chi \to \infty$, the household only “plays” the choice with highest utility with probability approaching 1.

Using this notion, the household puts large weight on both allocations associated with the twin peaks in its objective function. While the weights on each of these choices may vary slightly during value function iteration, individual utility within the household does not move much, allowing for convergence.