



**THE PINHAS SAPIR CENTER FOR DEVELOPMENT
TEL AVIV UNIVERSITY**

**"Why Not Settle Down Already"
A Quantitative Analysis of the Delay in Marriage"**

David Weiss¹, Cezar Santos²

Discussion Paper No. 3-15

January 2015

We would like to thank The Pinhas Sapir Center for Development at Tel Aviv University for their financial support.

We are very grateful to Jeremy Greenwood for his advice. The editor Jesús Fernández-Villaverde and three anonymous referees provided invaluable comments. We would also like to thank Nezh Guner, Moshe Hazan, Greg Kaplan, Fatih Karahan, John Knowles, Dirk Krueger, Iouri Manovskii, Serdar Ozkan, Michèle Tertilt, and seminar participants at Penn, U of Montreal, PUC-Rio, EPGE-FGV/RJ, EESP-FGV/SP, the U of São Paulo, ITAM, Mannheim, Konstanz, the 2011 Midwest Macro, the 2011 Minerva Summer Institute at The Hebrew University, the SED-Ghent, the 2011 Cologne Macro Workshop, and the 2011 Asian Meeting of the Econometric Society, the 2012 Royal Economic Society Meeting (Cambridge), St. Louis Fed, and the 2013 North American Meeting of the Econometric Society for their comments. The usual disclaimer applies

¹ David Weiss - The Eitan Berglas School of Economics, Tel Aviv University. Email: davidweiss@post.tau.ac.il

² Cezar Santos - Getulio Vargas Foundation. Email: cezar.santos@fgv.br

Abstract

One of the most striking changes in American society in the last forty years has been the decline and delay in marriage. The fraction of young men and women who have never been married increased significantly between 1970 and 2000. Idiosyncratic labor income volatility also increased over the same period. This paper establishes a quantitatively important link between these two facts. Specifically, if marriage involves consumption commitments, then a rise in income volatility results in a delay in marriage. Marriage, however, also allows for diversification of income risk since earnings fluctuations between spouses need not be perfectly correlated. We assess the hypothesis that rising income volatility contributed to the delay in marriage vis-à-vis other explanations in the literature, using an estimated equilibrium search model of the marriage market. We find that the increase in volatility accounts for about 20% of the observed delay in marriage. Thus, we find that the effects of consumption commitments due to increased income volatility outweigh the effects of the insurance gains provided by spouses.

Keywords: delay in marriage, income volatility, gender wage gap, household technology, search models of marriage, simulated method of moments.

JEL Codes: E24, D13, J12

1 Introduction

One of the most striking changes in American society over the last 40 years has been the decline and delay in first-time marriage. The fraction of young men and women who have never been married increased significantly between 1970 and 2000. This trend has captured the attention of both academic researchers and the general public.¹ The question here, in the vernacular, is: Why not settle down already? The answer we propose relies on the increased labor income volatility observed in this period. In order to quantitatively assess this hypothesis, we build and estimate a structural equilibrium search model of the marriage market.

Figure 1 shows the fraction of never-married American white males, by age, for both 1970 and 2000. This graph illustrates how the onset of marriage has been delayed. The numbers are striking. In 1970, only 26% of 25-year-old white males had never been married. By 2000, this number had more than doubled to 57%. At age 35, only 8% of white males were single in 1970, whereas this number increased to 21% in 2000.² These numbers clearly illustrate the decline and delay in marriage observed in this period.³

The economics literature has documented a rise in idiosyncratic labor income volatility over the same period. Gottschalk and Moffitt (1994), Heathcote, Perri, and Violante (2010), among others, find an increase in the variance of persistent and transitory shocks to income between the late 1960s and 2000. Various effects of this changing labor market have drawn the attention of a wide body of literature.⁴ However, to the best of our knowledge, no quantitative work has been done relating changes in income volatility with changing marriage decisions of young adults.

Figure 2 shows the increase in the median age of marriage for males and the increase in labor income volatility as measured by the standard deviation of persistent income shocks. It is interesting to note that both series exhibit a very similar increase between the late 1960s and 2000. In fact, the correlation between the two series is 0.96. Some empirical

¹For an excellent review of the academic literature, see Stevenson and Wolfers (2007).

²Decline and delay of marriage are two similar but distinct concepts. The decline in marriage is captured by the fraction of people who never marry. There has been an increase in this fraction in the data. The delay in marriage can be expressed as the fraction of people who are married by each age. As it will be clear later, the model is consistent with both.

³The graph for white women looks very similar. For data on cohabitation and by education groups, see appendix B. Detailed explanations about the data sources are contained in Appendix A.

⁴For an excellent overview of this literature with a specific focus on welfare, see Heathcote, Storesletten, and Violante (2011).

Abstract

One of the most striking changes in American society in the last forty years has been the decline and delay in marriage. The fraction of young men and women who have never been married increased significantly between 1970 and 2000. Idiosyncratic labor income volatility also increased over the same period. This paper establishes a quantitatively important link between these two facts. Specifically, if marriage involves consumption commitments, then a rise in income volatility results in a delay in marriage. Marriage, however, also allows for diversification of income risk since earnings fluctuations between spouses need not be perfectly correlated. We assess the hypothesis that rising income volatility contributed to the delay in marriage vis-à-vis other explanations in the literature, using an estimated equilibrium search model of the marriage market. We find that the increase in volatility accounts for about 20% of the observed delay in marriage. Thus, we find that the effects of consumption commitments due to increased income volatility outweigh the effects of the insurance gains provided by spouses.

Keywords: delay in marriage, income volatility, gender wage gap, household technology, search models of marriage, simulated method of moments.

JEL Codes: E24, D13, J12




Figure 1: Fraction of White Males Never Married, by Age

papers have also provided suggestive evidence of the impact of certain aspects of labor market volatility on marriage.⁵

The contribution of this paper is to quantitatively establish the effect of rising labor income volatility on the delay in marriage. We do this by exploring three channels through which income volatility can affect marriage timing. The first and novel effect that we explore in this paper arises from the presence of consumption commitments within marriage. Consumption commitments emerge when households consume goods for which adjustments are costly. These consumption commitments aggravate the effects of income fluctuations: households must cover these commitments, or face costly adjustments, following a bad income realization. To do so, they might need to cut their discretionary consumption substantially, causing a large utility loss. In this paper, we provide evidence that married individuals, compared to their single counterparts, have more consumption commitments; in particular, more married households have children. Therefore, a rise in the volatility of income results in a delay in marriage as these commitments become less desirable. That

⁵For example, using U.S. data, Oppenheimer, Kalmijn, and Lim (1997) argue that difficulties in starting careers in a period of higher volatility have delayed marriage. Ahn and Mira (2001) show that employment risk has caused delay in marriage in Spain. Southall and Gilbert (1996) study the impact of economic distress in 19th century United Kingdom and find that periods with more uncertainty are related with fewer marriages overall as well as higher variability in marriage rates for workers in more volatile occupations.




Figure 2: St. Deviation of Persistent Income Shocks and Median Age of Marriage for Males

is, singles might find it preferable to wait until one receives a favorable income shock, or search longer for a “better” spouse, before settling down with a family.

This paper also includes two other channels through which income volatility will affect marriage. One effect is that of spousal insurance: Marriage allows for diversification of income risk as earnings fluctuations between spouses need not be perfectly correlated. Therefore, higher income volatility may make marriage more desirable due to insurance. This mechanism is highlighted by Hess (2004). Another effect emerges if higher income volatility induces higher income inequality. With higher inequality, the marriage market will 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. Gould and Paserman (2003) find empirical support for this channel. All three of these effects, consumption commitments, spousal insurance, and search incentives, are incorporated in our study. Since these channels work in opposite directions, how rising income volatility will affect the timing of marriage ultimately becomes a question about the net impact of these three effects, which is addressed in our quantitative analysis.

In order to quantitatively assess the impact of increased labor income volatility on marriage decisions, we include two additional relevant changes to the U.S. labor market over this time period: the increased labor force participation of married women and the narrowing of the gender wage gap. Both changes are important determinants of the amount

of insurance spouses provide and thus in the decision to get married. Whether a wife is working or not and how high her earnings are determine how much her income can replace her husband's if he receives a bad shock in the labor market.⁶ In order to generate increased female labor force participation, we follow Greenwood, Seshadri, and Yorukoglu (2005), who make the case that less expensive household goods, such as washing machines and refrigerators, led to the increase in female labor force participation. Regalia and Rios-Rull (2001) argue that the decrease in the gender wage gap is itself important for the delay in marriage. They argue that when women become richer they can afford to be pickier with the mate they choose.⁷ Moreover, Greenwood and Guner (2009) argue that cheaper household goods made the cost of running a household lower. This caused the traditional household setup of the husband specializing in market work and the wife specializing in home production to become obsolete. The result, Greenwood and Guner argue, was a decrease in the gains from trade associated with marriage, and thus a decline in marriage. Since we include both of these channels, we can quantitatively assess their importance vis-à-vis increased income volatility.

We build an equilibrium search model of the marriage market in which the economy is populated by overlapping generations of individuals that optimally choose when to get married and have children. Each person's labor income is risky and households can save in a riskless bond. Married couples face economies of scale in consumption, but also must support their children, if they have any. These child-related expenses are what we consider to be consumption commitments. Parents choose the level of their children's consumption, which is subject to adjustment costs. Married females decide whether or not to work in the market. The model is estimated using the Simulated Method of Moments. We target several moments regarding marriage, fertility, labor force, and consumption choices that are derived from different micro data sets.

⁶This extensive margin labor force participation decision by married women also accounts for the "added-worker effect", which is also an important margin for insurance. For a discussion of the added-worker effect, see, for example, Lundberg (1985).

⁷Another interesting question investigated in Regalia and Rios-Rull (2001) is the rise in single motherhood. While evaluating the causes of the change in single motherhood is interesting, it is beyond the scope of the current exercise. We would like to point out, however, that while rising risk may well make women less willing to have a child out of wedlock, the economic mechanisms introduced in this paper would also imply that men are less willing to marry these women. If women have a stronger preference for children than men do, or are stuck with accidental pregnancies, it is possible that risk causes men's desire to marry to fall more than women's propensity to have children.

Our results show that rising labor income volatility accounts for 20% of the observed delay in marriage. Thus, we find that the effects of consumption commitments and changes to the option value of searching for a spouse due to rising income volatility outweigh the effects of the gains on spousal insurance. Regarding the other channels, we also find that the decrease in the price of home inputs explains 14% of this decline, while the effects of the narrowing of the gender wage gap are negligible. In sum, rising income volatility has substantially contributed to the delay in marriage.

In our model, the effect of increased labor income volatility on the timing of marriage is partially influenced by the presence of consumption commitments. Therefore, this paper contributes to the consumption commitment literature along the lines of Chetty and Szeidl (2007) and Postlewaite, Samuelson, and Silverman (2008). Chetty and Szeidl discuss how risk-averse agents can become even more risk averse in the presence of consumption commitments. Postlewaite, Samuelson, and Silverman study how risk-neutral individuals can behave as if they have preferences about risk when they face commitments. Another interesting paper, Sommer (2014), discusses the role of consumption commitments and rising income volatility. In her paper, she argues that rising volatility leads to a delay in fertility. Our papers differ in the modeling of the spousal insurance within marriage in this paper, in the equilibrium approach to the marriage market used here, and in the quantitative methodologies employed. In a related paper (Santos and Weiss, 2012), we expand upon this notion to discuss how income risk affects not only marriage, but divorce as well. Adding divorce comes at substantial computational cost, and complicates the model due to specifications of how children are dealt with upon divorce. There, we simplify several dimensions of the model described here in order to take divorce into account. For example, we do not compare income risk with other mechanisms in the literature. Greenwood, Guner, Kocharkov and Santos (2012) look at related changes in the American family, namely marriage and divorce patterns in the U.S. over the last 50 years. However, they do not discuss earnings risk or the effects of consumption commitments.

The remainder of this paper is organized as follows: Section 2 presents evidence on the relationship between marital status, risk, and consumption commitments. Section 3 presents the model, and Section 4 discusses the important channels working in the model. Section 5 discusses the estimation procedure. Section 6 discusses the results, and Section 7 concludes.

2 Consumption Commitments, Risk, and Marital Status

The focus of this paper is the relationship between the timing of marriage and income volatility. In this section, we discuss some empirical evidence on the relationship between labor market risk and marriage as well as the presence of consumption commitment within marriage. We start with the latter and provide evidence on a particular form of commitment: children.




Figure 3: Presence of Children in the Household by Marital Status (White Males)

As Figure 3 shows, a strong majority of white married men have children in their household, while the opposite is true for singles. These data show a strong link between marriage and fertility, a notoriously expensive and persistent form of consumption commitment.

Next, Figure 4 provides evidence of the relationship between volatility of income and the timing of marriage for U.S. states between 1970 and 2000. Since we cannot estimate the variance of income shocks at the state level, we use measures of inequality of residual income as a proxy for income volatility.⁸ The figure shows that, in states that experienced a larger increase in our proxy for income volatility, young adults marry at later ages. The relationship between the two variables is positive and statistically significant.⁹

⁸The data we use to estimate the variance of income shocks comes from the PSID, which is not representative at the state level. In order to get a measure of residual income, we control for educational attainment and a polynomial in age using a Mincerian regression for males. The difference between this measure and the volatility of income shocks could potentially come from the presence of individual fixed effects.

⁹We also computed the correlation between other measures for residual inequality and the timing of marriage. For residual inequality, other than the Gini coefficient, we also used: the variance of logarithms and the ratios of percentiles (75/25, 90/10, 90/50). For the timing of marriage, we used both the median and the mean age of marriage as well as the fraction of never-married 25-year-old men. We always found a




Figure 4: Changes in the Timing of Marriage and in the Variance of Residual Income

3 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 can either be single or married. Every agent is endowed with a unit of time every period.

3.1 Production

There are two goods in the economy: a market good, Y , and a home good, n . For the consumption good 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.

significantly positive correlation between all combinations of measures of residual inequality and the timing of marriage.

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

$$y = w\phi_g f_g(a), \quad (2)$$

where w is an idiosyncratic shock and the deterministic trend is composed of ϕ_g , a gender wage gap, and $f_g(a)$, a gender specific deterministic age income profile. We will now discuss each of these terms.

The shock w consists of a persistent shock with innovations ε and persistence δ . These parameters are constant for everyone in the model economy. However, we allow for shocks to be correlated between spouses. For example, if one spouse loses a job and needs to take a new one in a different city, then the other spouse will need to find a new, potentially worse job. The parameter ρ controls this correlation. This allows us to get the appropriate level of spousal insurance in the model. This insurance is a counter mechanism to income volatility causing a delay in marriage, so getting the correct level is important. Thus, we assume that this process takes the following form for singles:

$$\begin{aligned} \ln w &= \delta \ln w_{-1} + \varepsilon \\ \varepsilon &\sim N(0, \sigma_{\varepsilon,t}^2). \end{aligned} \quad (3)$$

For married individuals, the process specifies shocks for each of the two spouses (an arrow above each shock denotes that this is a vector). Thus, the income process for married households takes the following form:

$$\begin{aligned} \ln \vec{w} &= \delta \ln \vec{w}_{-1} + \vec{\varepsilon} \\ \vec{\varepsilon} &\sim N\left(0, \begin{bmatrix} \sigma_{\varepsilon,t}^2 & \rho \\ \rho & \sigma_{\varepsilon,t}^2 \end{bmatrix}\right). \end{aligned} \quad (4)$$

Note that the variances for all shocks are indexed by the time subscript $t \in \{1970, 2000\}$. An increase in volatility is measured by changing σ_{ε}^2 , which control the variance of the shocks.¹⁰

¹⁰In 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.

As noted above, the amount of efficiency units available to an agent also varies with his/her age a according to the function $f_g(a)$. 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 $\phi < 1$ if $g = 2$ (females).

We turn to the home sector now. The home good, n , is produced by a constant elasticity of substitution production function between home inputs, d , and time, h :

$$n = \left[\theta d^\xi + (1 - \theta) h^\xi \right]^{1/\xi}, \quad (5)$$

where θ is the relative weight on home inputs, and ξ is the parameter that controls the elasticity of substitution between home inputs and time.

3.2 Preferences

Preferences of households are additively separable and exhibit constant relative risk aversion (CRRA) over both consumption goods and home goods. We begin with singles. Their period t utility function reads:

$$u^s(c_t, n_t) = \frac{c_t^{1-\eta}}{1-\eta} + \alpha \frac{n_t^{1-\zeta}}{1-\zeta}, \quad (6)$$

where η is the CRRA parameter on the consumption of market goods, ζ is the CRRA parameter on home goods, and α is the relative weight of home goods.

For marrieds, we assume a unitary model, i.e., that spouses make decisions jointly when choosing the household's level of consumption goods c , children's consumption c_k (if they have any), and home goods n .¹¹ The amount of household consumption that is enjoyed by the spouses in a married household is determined by the economies of scale in consumption — ψ is the parameter that controls these economies of scale. The period t utility function for each individual married agent then reads:

¹¹Knowles (2013) uses data on men's hours at home and in the market to analyze whether bargaining in marriage is important for macro-economists. His analysis is particularly relevant, since he studies household choices in men's labor supply, a margin we abstract from. Empirically, men's labor market hours changed considerably less than women's. He concludes that "Bargaining in the model has little to add to the analysis of events whose impacts on the value of single life are similar for both sexes." In this paper, we study the impact of income risk. The effects are similar for single men as single women, meeting Knowles' criterion for excluding bargaining. Accordingly, we use a unitary model of household decision making and take men's labor choices exogenously.

$$u^m(c, n, k, c_k, c_{k,-1}, \iota) = \frac{\left(\frac{c}{1+\psi}\right)^{1-\eta}}{1-\eta} + \alpha \frac{\left(\frac{n}{1+\psi}\right)^{1-\zeta}}{1-\zeta} + \mathcal{I}_{\iota>1} \left[\alpha_k(k) \frac{\left(\frac{c_k}{k}\right)^{1-\eta}}{1-\eta} - \mathcal{I}_{\iota=2} (\chi(c_k - c_{k,-1}))^2 \right], \quad (7)$$

where k is the number of children and ι represents the family's fertility status. If the family does not current have a child, $\iota = 0$. If there is a newborn child, such that there was no $c_{k,-1}$, and thus no potential for adjustment costs, $\iota = 1$. When the child is at least one year old, $\iota = 2$. $\mathcal{I}_\iota$ is an indicator function for the variable ι . When there are children ($k > 0$), the family gets flow utility out of each child's consumption, c_k/k , which has relative weight in the utility function of $\alpha_k(k)$.¹² Furthermore, capturing the idea of a consumption commitment, there is a quadratic utility cost for changing the child's consumption from period to period, starting when the child is in her second year. This cost has weight χ .

The expected discounted value of lifetime utility:

$$U(\{c_{t=1}^{t=T}\}, \{n_{t=1}^{t=T}\}, \{c_{k,t=1}^{t=T}\}) = E_{t=1} \left[\sum_{t=1}^{t=T} \mathcal{I}_{s,t} u_t^s(c_t, n_t) + (1 - \mathcal{I}_{s,t}) u_t^m(c_t, n_t, k, c_{k,t}, c_{k,t-1}, \iota_t) \right], \quad (8)$$

where $\mathcal{I}_{s,t}$ is an indicator function that the agent is single in period t .

In addition to the utility derived from the consumption of goods, when individuals first get married, they also enjoy an additive marital bliss utility denoted by γ . This is a stochastic shock drawn from the distribution $\Gamma(\gamma)$. We assume that $\gamma \sim N(\mu_\gamma, \sigma_\gamma^2)$. This utility shock is received only once at the start of married life.¹³ This represents the (stochastic) lifetime discounted utility of being married that arises due to non-economic reasons.

Finally, fertility is endogenous. People may choose whether or not to have a single child each period, up to a total of K children in their lifetime. The decision to have a child is irreversible. When people are married, and have fewer than K children, they draw a fertility bliss shock every period $\lambda \sim N(\mu_{\lambda,k}, \sigma_\lambda^2)$. This shock captures how much people

¹²The main exercise will set limit couples to having at most 1 child, while the robustness exercises will explore flexibility in this dimension.

¹³Since there is no divorce in the model and γ is additively separable, the assumption that the marital bliss shock is completely front-loaded at the time of marriage is without loss of generality. It also makes the computation of the model easier, given that γ will thus not be a state variable.

love their children in addition to any utility they derive from their child's consumption c_k .¹⁴ Note that the parameters of the shock process depend on the number of children in the household. This feature of the model allows for greater flexibility in the flow utility of the marginal child. The presence of the fertility shock allows the model to separate between the decision to get married and fertility choices.

3.3 Budget Sets

All singles divide their time between market and home production at an exogenous rate, such that they work τ_g^s amount of their time on the market, which is allowed to depend on their gender g . Thus, their budget constraint will be given by

$$c + pd + b' = \phi_g w f_g(a) \tau_g^s + (1 + r)b \quad (9)$$

where p is the price of home inputs, ϕ_g is the gender wage gap, w is the idiosyncratic productivity shock, $f_g(a)$ is an age dependent productivity level, b is the individual's current level of assets chosen in the previous period, and b' is the savings chosen today.¹⁵ $(1 + r)$ is the gross interest rate.

When married, spouses pool their resources. Furthermore, when the couple has children, there are consumption commitments. This is modeled as an endogenous expenditure that married agents pay every period, denoted by c_k . Married women have the option of whether to work in the market or work only at home — l^f is the indicator function that women choose to work in the market. Denote by w_1 and w_2 the husband's and wife's wage offers, respectively. The time spent working for the husband (wife) is τ_1^m (τ_2^m). Hence, a couple's budget constraint reads

$$c + pd + c_k + b' = w_1 f_1(a) \tau_1^m + l^f \phi w_2 f_2(a) \tau_2^m + (1 + r)b. \quad (10)$$

¹⁴For the same reason as γ , fertility bliss λ is not a state variable.

¹⁵Agents can only save, not borrow. We note, however, that financing is most effective at smoothing transitory income shocks, rather than permanent ones. In a previous working paper version, we had both types of income shocks. We found that people could effectively smooth out transitory income shocks with assets, and thus these shocks had little impact. We removed them from this version of the model, which can be thought of as allowing them to be freely and completely insured. The shocks that remain are relatively permanent, with an autoregressive coefficient of 0.98, as can be seen in the calibration section. People can still smooth these shocks with self insurance.

3.4 Timing and Marriage

The timing of a period is as follows:

- At the beginning of the period, agents observe the realization of shocks to their wage offers.
- Single agents randomly meet another single agent of the same (model) age and opposite gender and decide whether or not to get married. Marriage is an absorbing state, i.e., there is no divorce.¹⁶
- Married couples, including newlyweds, enter the fertility phase. They draw the bliss shock λ , and decide whether or not to have a child (up to K children).
- Married agents choose whether or not the wife works.¹⁷ All agents optimally divide their income between consumption goods, children's consumption, home inputs, and savings. Consumption takes place and the period ends.

3.5 Decision Making

How do households make their decisions in the model? Single agents decide how to divide their income between the consumption of market, non-market goods, and their asset holdings. 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 market and home-produced goods, and must decide whether the wife should work or not. Moreover, if married, they decide whether to have a child. We will now describe each household's problem recursively.

Let us start with couples. The state vector for married households consists of a wage shock for the husband w , a wage offer shock for the wife w^* , the current assets level b , an expenditure level $c_{k,-1}$ representing how much the couple spent on their child (if they had one) in the previous period, their current fertility status ι , and their age a .

¹⁶This is a simplifying assumption, to make modeling marriage and keeping track of singles distributions easier. Since we are trying to explain timing of first marriages only, the issue is whether or not there are a lot of young divorcés for never-married people to consider marrying. Empirically, there are not. In 2000, the percentage of young adults (under age 30) who had been divorced/separated was roughly 5% (IPUMS-Census). This figure is slightly lower for 1970. Since there are so few of these people to worry about in the data, we exclude them from the model.

¹⁷That is, the extensive, not intensive, margin of female labor force participation.

The first subperiod is the decision making process for fertility. In the beginning of a period, married couples (including newlyweds) receive a draw of a fertility bliss from the distribution $\lambda \sim N(\mu_\lambda, \sigma_\lambda^2)$. They choose to have a child, and thus start providing it with consumption immediately, if the value of being married with an extra child, along with the bliss shock, is larger than the value of not doing so. That is,

$$V^m(w, w^*, b, k+1, c_{k,-1}, 1, a) + \lambda > V^m(w, w^*, b, k, c_{k,-1}, 0, a). \quad (11)$$

Again, 0 in the state for ι represents having no children, while 1 represents having just had a child, and thus there is no adjustment cost to face. For these people, the c_k state is redundant. The policy function for fertility is thus

$$J_k(w, w^*, b, k, \lambda, a) = \begin{cases} 1, & \text{(extra child)} \\ 0, & \text{(no extra child)} \end{cases}$$

The value function for entering the fertility phase is

$$F(w, w^*, b, k, c_{k,-1}, \iota, a) = \int \{J_k(w, w^*, b, k, a, \lambda)[V^m(w, w^*, b, k, c_{k,-1}, 1, a) + \lambda] + (1 - J_k(w, w^*, b, k, a, \lambda))V^m(w, w^*, b, k, c_{k,-1}, 0, a)\} d\Lambda(\lambda).$$

Then the married value function can be written as follows:

$$\begin{aligned} V^m(w, w^*, b, k, c_{k,-1}, \iota, a) &= \max_{l^f \in \{0,1\}, b' \geq 0, c \geq 0, c_k \geq 0, d \geq 0} u^m(c, n, k, c_k, c_{k-1}, \iota) \\ &\quad + \beta E_{w', w^{*'}} F(w', w^{*'}, b', k, c_k, \iota, a+1) \\ &\text{s.t.} \\ c + pd + c_k + b' &= wf(a)\tau_1^m + l^f \phi w^* f_2(a)\tau_2^m + (1+r)b \\ n &= \left[\theta d^\xi + (1-\theta) \left(2 - \tau_1^m - l^f \tau_2^m \right)^\xi \right]^{1/\xi}. \end{aligned} \quad (12)$$

A married household chooses whether or not the wife works this period, l^f , consumption c , consumption of the children c_k , savings b' , and home inputs d . Define the policy functions for the married problem as follows: $l^f = P_l^m(w, w^*, b, k, c_{k,-1}, \iota, a)$ for the woman's labor force decision, $d = P_d^m(w, w^*, b, k, c_{k,-1}, \iota, a)$ for choice of home inputs,

$c = P_c^m(w, w^*, b, k, c_{k,-1}, \iota, a)$ for the consumption decision, $c_k = P_{c_k}^m(w, w^*, b, k, c_{k,-1}, \iota, a)$ for the children's consumption, and $b' = P_b^m(w, w^*, b, k, c_{k,-1}, \iota, a)$ for the savings decision. The continuation value is given by the expected value of being married during the next period, and going into the fertility phase, F , described above. The expectation is taken with respect to the income shocks for both spouses.

Now, we move on to singles. The value function for singles with wage shock w , asset holdings b , gender g , and age a , after the marriage market, is as follows:

$$\begin{aligned}
V^s(w, b, g, a) &= \max_{b' \geq 0, c \geq 0, d \geq 0} u^s(c, n) + \beta E_{w'} B(w', b', g, a + 1) \\
&\text{s.t.} \\
c + pd + b' &= \phi_g w(x) f_g(a) \tau_g^s + (1 + r)b \\
n &= \left[\theta d^\xi + (1 - \theta)(1 - \tau_g^s)^\xi \right]^{1/\xi}.
\end{aligned} \tag{13}$$

Single households choose consumption c , savings b' , and home inputs d . Define the following policy functions associated with the single agent's problem: $d = P_d^s(w, b, g, a)$ for choice of home inputs, $c = P_c^s(w, b, g, a)$ for the consumption decision, and $b' = P_b^s(w, b, g, a)$ for the savings decision. The continuation value for singles is the expectation of the value function $B(\cdot)$, which represents the value for a single before going through the marriage market (or the “bachelor” phase); and the expectation is taken with respect to the income shocks next period. We will elaborate on the value function $B(\cdot)$ slightly later in this section.

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 γ from the distribution $\Gamma(\gamma)$. 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. Again, as they just married and have no kids, $\iota = 0$. The state $c_{k,-1}$ is redundant and left valueless. A marriage occurs if and only if both agents agree to marriage. Formally, a marriage occurs if and only if

$$\underbrace{V^m(w, w^*, b + b^*, 0, c_{k,-1}, 0, a) + \gamma}_{\text{male's decision}} > V^s(w, b, 1, a) \text{ and } \underbrace{V^m(w, w^*, b + b^*, 0, c_{k,-1}, 0, a) + \gamma}_{\text{female's decision}} > V^s(w^*, b, 2, a). \tag{14}$$

Let the indicator function $J(w, w^*, b, b^*, \gamma, a)$ take a value of 1 if both people agree to the match and a value of 0 otherwise. Thus,

$$J(w, w^*, b, b^*, \gamma, a) = \begin{cases} 1, & \text{if (14) holds,} \\ 0, & \text{otherwise.} \end{cases} \quad (15)$$

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

$$\begin{aligned} B(w, b, g, a) = & \int \int \{J(w, w^*, b, b^*, \gamma, a) [V^m(w, w^*, b + b^*, 0, -1, a) + \gamma] \\ & + (1 - J(w, w^*, b, b^*, \gamma, a)) V^s(w, b, g, a)\} d\widehat{\mathbf{S}}(w^*, b^*, g^*, a) d\Gamma(\gamma), \end{aligned} \quad (16)$$

where $\widehat{\mathbf{S}}(w^*, b^*, g^*, a)$ is the probability distribution of meeting a potential mate from the other gender (g^*) and age a . This will be elaborated on later.

3.6 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 decisions, this distribution will be an equilibrium object. The non-normalized stationary distribution for singles aged $a > 1$ is given by

$$\begin{aligned} \mathbf{S}(w', b', g, a + 1) = & \iiint (1 - J(w, w^*, b, b^*, \gamma, a)) \mathcal{I}(P_b^s(w, b, g, a) \leq b') \times \\ & \times \mathbf{S}(w, b, g, a) d\mathbf{S}(w^*, b^*, g^*, a) d\mathbf{W}^s(w', w) d\Gamma(\gamma), \end{aligned} \quad (17)$$

where g^* represents the opposite gender and $\mathbf{W}$ represents the wage shock process for singles defined above. Singles aged $a = 1$ are distributed over wages according to the invariant distribution of $\mathbf{W}^s$. $\widehat{\mathbf{S}}(w, b, g, a)$ denotes the normalized distribution for singles that determines the probability that single agents will meet in the marriage market, and is defined by

$$\widehat{\mathbf{S}}(w, b, g, a) = \frac{\mathbf{S}(w, b, g, a)}{\int d\mathbf{S}(w, b, g, a)}.$$

We can now formally define the equilibrium for this economy:

Definition 1 *A stationary equilibrium is a set of value functions for singles, couples, bachelors, and fertile couples, $V^s(w, b, g, a)$, $V^m(w, w^*, b, k, c_{k,-1}, \iota, a)$, $B(w, b, g, a)$ and $F(w, w^*, b, k, c_{k,-1}, \iota, a)$; policy functions for single households $P_c^s(w, b, g, a)$, $P_d^s(w, b, g, a)$, and $P_b^s(w, b, g, a)$; policy functions for married households $P_c^m(w, w^*, b, k, c_{k,-1}, \iota, a)$, $P_d^m(w, w^*, b, k, c_{k,-1}, \iota, a)$, $P_l^m(w, w^*, b, k, c_{k,-1}, \iota, a)$, $P_{c_k}^m(w, w^*, b, k, c_{k,-1}, \iota, a)$, and $P_b^m(w, w^*, b, k, c_{k,-1}, \iota, a)$; a matching rule for singles $J(w, w^*, b, b^*, \gamma, a)$; a fertility rule $J_k(w, w^*, b, k, \lambda, a)$, and a stationary distribution for singles $\mathbf{S}(w, b, g, a)$ such that:*

1. *The value function $V^s(w, b, g, a)$ and the policy functions $P_c^s(w, b, g, a)$, $P_d^s(w, b, g, a)$, and $P_b^s(w, b, g, a)$ solve the single's problem (13), given the value function for bachelors $B(w, b, g, a)$ and the distribution for singles $\mathbf{S}(w, b, g, a)$.*
2. *The value function $V^m(w, w^*, b, k, c_{k,-1}, \iota, a)$ and the policy functions $P_c^m(w, w^*, b, k, c_{k,-1}, \iota, a)$, $P_d^m(w, w^*, b, k, c_{k,-1}, \iota, a)$, $P_l^m(w, w^*, b, k, c_{k,-1}, \iota, a)$, $P_{c_k}^m(w, w^*, b, k, c_{k,-1}, \iota, a)$, and $P_b^m(w, w^*, b, k, c_{k,-1}, \iota, a)$ solve the couple's problem (12).*
3. *The value function $B(w, b, g, a)$ solves the bachelor's problem (16), given the value functions for singles and couples, $V^s(w, b, g, a)$ and $V^m(w, w^*, b, k, c_{k,-1}, \iota, a)$, and the matching rule $J(w, w^*, b, b^*, \gamma, a)$.*
4. *The matching rule $J(w, w^*, b, b^*, \gamma, a)$ is determined according to (15), taking as given the value functions $V^s(w, b, g, a)$ and $V^m(w, w^*, b, k, c_{k,-1}, \iota, a)$.*
5. *The fertility rule $J_k(w, w^*, b, k, \lambda, a)$ is determined according to (11), with associated value function $F(w, w^*, b, k, c_{k,-1}, \iota, a)$.*
6. *The stationary distribution $\mathbf{S}(w, b, g, a)$ solves (17), taking as given the matching rule $J(w, w^*, b, b^*, \gamma, a)$ and the policy function $P_b^s(w, b, g, a)$.*

4 Mechanisms

Our purpose is to quantitatively explain the delay in entrance into marriage between 1970 and 2000. There are three exogenous forces that change over time in the model: income volatility, the price of home inputs (which represents technological progress in the home sector), and the gender wage gap. In this section, we discuss the effects of each of these forces in turn.

4.1 Income Volatility

This is the chief hypothesis we propose: The rise in labor income volatility (increasing σ_ϵ^2) has multiple effects. Let's first discuss the role played by the presence of consumption commitments within married households.

Consumption commitments emerge when households consume goods for which adjustments are costly. In our model, these consumption commitments are embodied in the choice variable c_k and parameter κ that come with the endogenous decision to have children. These consumption commitments aggravate the effects of income fluctuations by effectively causing an increase in risk aversion among married agents relative to single agents. This effect comes in two ways. The first is that the mere presence of children has people optimally spending money on them. In turn, parents spend less on their own consumption, moving to a steeper, and thus more risk averse, area of their utility function. The second is due to the adjustment costs κ . This parameter amplifies the cost to couples of adjusting consumption in response to shocks. Due to these factors, a rise in the volatility of income results in a delay in marriage as these commitments become less desirable. That is, singles might find it preferable to wait until one receives a favorable income shock, or search longer for a “better” spouse, before settling down with a family. By delaying marriage, individuals expect to earn higher income in the future (given the growth in wages over the life cycle) and accumulate more assets which will help them cover the consumption commitments associated with married life.

There are other channels through which income volatility will affect marriage. One 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 (γ), 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 agents may search longer for a better spouse.

Another 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. 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.

All three effects discussed here are incorporated in our study and, since they work in opposite directions, how rising income volatility will affect the timing of marriage ultimately becomes a question about the net impact of these three effects, which is addressed in our quantitative analysis.

4.2 Price of Home Inputs

Another exogenous change present in the model are improvements in the technology of the home sector, modeled here as a decrease in the price of the inputs used in home production. Greenwood and Guner (2009) explain in detail the mechanism by which such a decrease in the price of inputs for home production (such as washing machines) would be likely to cause a decrease in marriage. The idea is simple: If marriage allows men and women to specialize according to their comparative advantages of market production and home production, respectively, then a decrease in the price of goods used as inputs for home production would tend to decrease the gains from specialization. As the prices of home inputs decrease, females have an incentive to work in the market given the substitutability of time and home inputs in the production function of home goods. If the marginal utility of home goods declines faster than that of market goods, married households will spend less on home inputs compared to less well-off single households.¹⁸ This will be especially true for younger and poorer individuals. Thus young single households will benefit more from improvements in the technology of home production and, as the gains from marriage decrease, people will postpone marriage.

4.3 Gender Wage Gap

The final mechanism explored is the narrowing of the gender wage gap. This is one of the channels explored by Regalia and Rios-Rull (2001). Again, we will highlight the various channels through which a change in the gender wage gap affects marriage decisions. With a smaller gender gap in income, women make relatively more than they did before, as compared to men. This causes two opposing effects on marriage.

The first effect appears in the changes for a female when she is single. With a lower gender wage gap, women are richer than before. They can now afford a better standard of

¹⁸This will be the case in our quantitative analysis, since the estimation procedure yields $\eta < \zeta$.

living while they are still single and gives them a better option outside of marriage. With this more attractive outside option, women can afford to be pickier with the mate they choose and thus they search longer. This will cause a delay in marriage.

The second effect of a lower gender wage gap is related to married life. As women are richer, they are able to provide more resources to a married household. This will make them economically “more attractive” to men. *Ceteris paribus*, men will be more likely to marry in order to enjoy the extra income provided by their now-richer wives. This effect will then cause more marriages to take place.

The net effect of the gender wage gap changing over time is thus ambiguous. We quantitatively analyze these channels to determine the net effect of the gender wage gap.

5 Matching the Model to the Data

The model period is 1 year. Given the age gap of approximately 2 years between the age of marriage for a male and a female (which remained approximately constant through the period analyzed), the same model age actually corresponds to this two-year gap in the data, i.e., age 1 in the model corresponds to age 18 (16) for males (females) in the data.

5.1 Computation

In order to numerically solve the model, we use backwards induction on the value functions. The model is solved for males from ages 18 to 35 (16 to 33 for females).¹⁹ At this final age, we need a terminal condition. This terminal condition is determined by solving a slightly modified version of the model for an extra 30 years: After age 35 (33 for females), the marriage market is shut down and couples cannot have children anymore, but the problems are otherwise the same as the ones described above. Agents live until age 65 (63 for females), after which they die with certainty.²⁰

We solve two steady states for the model; one that represents the world in 1970 and the other in 2000. Most parameters are kept constant for both steady states. The only parameters that change are those that govern the variance of income shocks, the gender wage gap, the price of household inputs, and the mean of the marital bliss shock distribution μ_γ and fertility bliss shock $\mu_{\lambda,k}$. The reason for changing these means across time periods

¹⁹We restrict the ages at which individuals can marry due to computational considerations.

²⁰Equivalently, they can be thought of as having retired with no more financial responsibility towards their children.

will be elaborated on later. A more detailed discussion of how the parameters in the model are calibrated/estimated will now follow.

5.2 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 us start with preference parameters. The time discount factor β is set to 0.98, which is the inverse of the gross interest rate, discussed below, and is also similar to what is used in the literature. The coefficient of relative risk aversion (CRRA) for market goods is set to 2.0, which is also standard in the macroeconomic literature.²¹ 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 $\psi = 0.7$.

The parameters for the production function of non-market goods were estimated by McGrattan, Rogerson, and Wright (1997) using business cycle frequency data. Their numbers are used by Greenwood and Guner (2009) in a model of the marriage market. We also use their numbers in this paper.

The maximum number of children K is set to 1 for the main exercise, with a robustness analysis where $K = 2$ in Section 6.1. This allows the model to simultaneously have an intensive margin for households to adjust their fertility choices, while maintaining computational tractability.

We can also set some parameters related to the labor market. The correlation of spousal persistent shocks ρ is set to 0.25, the number estimated by Hyslop (2001) using data from the PSID. The fraction of time spent working in the market is computed using data from the U.S. Census. We compute the number of hours worked in a week and divide by 112, the number of non-sleeping hours in a week. These numbers are allowed to vary by marital status and gender, as displayed in Table 1. The life-cycle profile $f_g(a)$ that controls the average level of efficiency units supplied at every age for each gender is computed by fitting a cubic polynomial over the mean income at each different age in the U.S. Census.²² We choose a cubic polynomial because it provides a very good fit to the non-parametric data. Finally, we estimate the income process parameters $\Upsilon = (\delta, \sigma_{\epsilon,1970}^2, \sigma_{\epsilon,2000}^2)$ from the PSID.

²¹We perform a number of robustness analyses on these parameter values and others in Section 6.1.

²²The results are very similar if we use data from the PSID. We use the larger sample from the U.S. Census to get tighter estimates.

Table 1: Parameters Set Using a Priori Information

Parameter	Description	Value	Source
<u>Preferences</u>			
β	Time discount factor	0.98	Standard/Interest Rate
λ	CRRA —consumption	2.0	Standard
ψ	Economies of scale	0.7	OECD equiv. scale
<u>Technology</u>			
θ	Weight on home inputs in production	0.206	McGrattan et al (1997)
ξ	CES home production	0.189	McGrattan et al (1997)
K	Maximum number of children	1	—
<u>Income</u>			
ρ	Correlation of spousal pers. shocks	0.25	Hyslop (2001)
τ_1^s	% of time at work (single males)	0.37	U.S. Census
τ_2^s	% of time at work (single females)	0.35	U.S.Census
τ_1^m	% of time at work (married males)	0.40	U.S.Census
τ_2^m	% of time at work (married females)	0.32	U.S.Census
$f_g(a)$	Age profile of income	—	U.S.Census
Υ	Income Shock Process	—	PSID
<u>Prices</u>			
—	Decline in the price of home inputs	6%/year	Greenwood & Guner (2009)
r	Interest rate	2%	Kaplan & Violante (2014)

The estimation procedure is detailed in Appendix C. The estimated parameter values are reported in Table 2.

Since this is a partial equilibrium model with respect to capital and home goods markets, we have to make some assumptions about prices. We set the interest rate to $r = 0.02$, following Kaplan and Violante (2014).²³ For the decline in the price of home inputs, we use 6%, the number estimated by Greenwood and Guner (2009). This number falls in the middle of other available estimates: the Gordon (1990) quality-adjusted price index for home appliances fell at 10% a year in the postwar period; on the other hand, the price of

²³Kaplan and Violante (2014) find that the average real net of tax return on total household wealth is 1.67%, and on illiquid wealth to be 2.29%. We choose 2% as a midpoint. Their data come from the Survey of Consumer Finances. We perform a robustness analysis on the discount factor β in Section 6.1.

kitchen and other household appliances from the National Income and Product Accounts (NIPA) declined at about 1.5% a year since 1950.

Table 2: Parameters for the Income Process

Parameter	Description	Param.
δ	Autoregressive Coefficient	0.9803
$\sigma_{\varepsilon,1970}^2$	Shock Variance	0.0113
$\sigma_{\varepsilon,2000}^2$	Shock Variance	0.0342

5.3 Estimation

The remaining parameters are estimated by 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 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 $\hat{\Omega}$ is given by the square root of the diagonal of the following matrix:

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

where Σ is the variance-covariance matrix of data moments, $g_1(\hat{\Omega}) = \partial g(\hat{\Omega}) / \partial \Omega$, 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 13 parameters so that we have the following vector of parameters to be estimated: $\Omega = (\alpha, \zeta, p, \kappa, \phi_{1970}, \phi_{2000}, \mu_{\gamma,1970}, \sigma_{\gamma}, \mu_{\gamma,2000}, \mu_{\lambda,1970}, \sigma_{\lambda}, \mu_{\lambda,2000}, \alpha_k)$.

5.3.1 Estimated Parameters

In addition to the parameters discussed in the previous section, we still need to estimate 13 extra parameters. 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.

Let us first start with parameters that influence the production and consumption of home goods: the weight of home goods in the utility function α , the CRRA for home goods ζ , and the initial level for the price of home inputs in 1970 p .²⁴ The data moment we use to identify the parameter p is the fraction of income spent on household operations in 2000. According to the U.S. National Income and Product Accounts (NIPA), this number is approximately 10.5%. Greenwood and Guner (2009) also include food as an example of their measure of home goods; according to NIPA, this would lead to approximately 40% of consumption share. We target an intermediate number: Household Operations, Utilities, and Personal Care. In 2000, this number was 23% of household consumption according to the Consumer Expenditure Survey (CEX).²⁵ Since home goods are produced using time and, in our model, married females choose whether to work in the market or not, we use the labor force participation rate (LFPR) of married females as data targets to identify the parameters that control the utility of home goods (α and ζ). We target LFPR in both 1970 and 2000 since this can give us information on the elasticity of labor supplied by married females. According to the U.S. Census, the LFPR for married females was 0.42 in 1970 and 0.72 in 2000.

In our model, married females are able to move into and out of the labor force freely. Consider the difference in incentives to change labor force status for women with and without kids. Women with kids face commitments, and are thus more willing to change their labor force status in order to adjust their consumption less. Thus, the movements in and out of the labor force help to identify κ . In the data, we measure these movements using PSID data. Since the PSID data is a panel data set, we can follow married females over time and observe how often they move. The data targets we use are the fractions of wives that moved into and out of the labor force in 2000. We use 2000 since there is more of a disconnect between being married and having children than in 1970, allowing for better

²⁴For the price of home inputs in 2000, we decrease the price p by 6% per year, the number reported by Greenwood and Guner (2009) —see Table 1.

²⁵We use the same sample as Aguiar and Hurst (2013).

identification. The percentage of wives (with children) that moved into the labor force in that year was 21.3% (21.8%), the percentage that moved out was 4.7% (7%).²⁶

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 dummy using Census data from both 1970 and 2000. We run this regression using observed wages for individuals that both work and report positive income. The coefficient on the gender dummy is our data target for the gender wage gap. The value of the estimates are 0.67 for 1970 and 0.75 for 2000.

We now turn to the parameters that govern the marital bliss shocks in 1970: $\mu_{\gamma,1970}$ and σ_{γ} . These parameters govern the average level and dispersion of match qualities in the economy. They control both the number and timing of marriages. Imagine that the variance of the Γ distribution was 0, for instance. In that case, a potential couple would not have to worry about all the different potential relationships that are also available in the economy, as they are all the same. Then $\mu_{\gamma,1970}$ would only control the level of marriages that take place in equilibrium. With a more dispersed distribution, which is controlled by the parameter σ_{γ} , potential mates might prefer to wait for a better draw. In order to identify these two parameters, we target the overall age profile of single males in 1970, which informs on both the overall level of marriages and their timing. The same logic applies to the fertility parameters, $\mu_{\lambda,k,1970}$, $\mu_{\lambda,k,2000}$ and σ_k . They control both the number and timing of births. We therefore target the profile of the fraction of married couples with at least one child, in both 1970 and 2000.

As mentioned above, we also allow the mean of the distribution for match qualities, μ_{γ} , and fertility bliss, μ_{λ} to change across steady states. This is done in order to guarantee that the model will be able to match the entire change in the timing and level of marriages and births that took place between 1970 and 2000. The goal is to study the effects of rising volatility, reduced price of home inputs, and the narrowing gender wage gap on marriage. However, it is unlikely that these three forces can explain the entire decline. Thus, we can think of this change in μ_{γ} and μ_{λ} as explaining the residual change of the delay in marriage and fertility.²⁷

²⁶In Section 6.1, we recalibrate the model imposing $\kappa = 0$ as a robustness exercise.

²⁷For example, this residual can be thought of as containing other explanations for the delay in marriage, like changes in social norms, improvements in contraception technology, etc. See Stevenson and Wolfers (2007) for a discussion of different explanations.

At first glance, it may seem unnecessary to include a residual: Why not simply see how much the channels in the model can account for? The problem with this is that in order to use moments from both steady states, such as the gender wage gap and women’s labor force participation, we need to get the right levels of single and married people, with and without children, in the model in both time periods. We therefore need to include this residual to ensure that the model explains all the data. Moreover, the inclusion of these residuals allows us to decompose the strength of various channels without forcing them to match the observed trend in the data. That is, nothing in the estimation strategy imposes that our proposed mechanisms explain any of the observed decline in marriage.

In our model, married couples with children choose amount of consumption for their child, c_k , every period. A key parameter for this choice is α_k , which controls the relative value of a child’s consumption to her parent’s consumption. We then choose to target the average fraction of household expenditures attributable to children in households that have both a husband and a wife and *one* child. Considering that marriage often results in more than one child, we consider this to be a reasonable lower bound on the level of expenditures faced by married parents. Betson et al. (2001) estimate the fraction of a household’s consumption expenditure that is attributable to one child using data from the 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 child. 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. 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%. This is the number we use as our target.

5.4 Model Fit

In this section, we discuss the fit of the model, in regard to both the moments used in the estimation and non-targeted statistics. We estimate a total of 13 parameters by targeting 30 data moments. The data moments are the 10 statistics included in Table 4 along with 20 selected points in the profiles of singles by age and fraction of married people with children

Table 3: Estimated Parameters

Parameter	Description	Value	SE
α	Utility Weight on Home Goods	4.24	0.27
ζ	CRRA Parameter on Home Goods	4.07	0.07
p	Price of home inputs, 2000	0.57	0.06
κ	Adj. cost on child consumption	112.46	49.14
ϕ_{1970}	Gender Gap - 1970	0.50	0.002
ϕ_{2000}	Gender Gap - 2000	0.60	0.003
$\mu_{\gamma,1970}$	Mean marital bliss shock - 1970	143.96	58.92
σ_{γ}	St. Deviation of marital bliss shock	83.73	17.78
$\mu_{\gamma,2000}$	Mean marital bliss shock - 2000	-100.23	12.22
$\mu_{\lambda,1970}$	Mean fertility bliss shock - 1970	557.48	21.95
σ_{λ}	St. Deviation of fertility bliss shock	10.10	1.08
$\mu_{\lambda,2000}$	Mean fertility bliss shock - 2000	431.45	26.39
α_k	Utility Weight on Child's consumption	1.45	0.07

by age, for both 1970 and 2000. These profiles are graphed in Figures 5 and 7 respectively. The estimated parameter values are reported in Table 3.

We can observe the narrowing of the gender wage gap, represented by an increase in the relative income of women (an increase in the value of ϕ over time). The CRRA parameter for home goods ζ is estimated to be larger than the CRRA parameter on market goods. This means that the marginal utility of home-produced goods decreases faster than the marginal utility of market goods. As discussed in Section 4, this means that younger, poorer single households benefit more from the decline in the price of home inputs. Finally, we can observe that the parameter that controls the average level of marital bliss shocks, μ_{γ} , decreases over time. This means that there is indeed a residual delay in marriage left unexplained by the forces explicitly modeled in this paper. In Section 6, we will quantify the quantitative power of each of these forces.

Figure 5 compares the fraction of single males at each age in the model and in the data. The model generates a good fit both in terms of the level of marriages that take place and also their timing. In the estimation, we only target the life cycle profile of single males, not females. Figure 6 plots the fraction of single females at each age both in the model and in




Figure 5: Model Fit —Fraction of Single Males by Age

the data. The similarity of the model and data in this dimension validates the simplifying assumption that men marry women 2 years younger than themselves.

Additionally, Figure 7 shows the model fit of the fraction of married couples with children, both in 1970 and 2000. The fact that the model successfully replicates these profiles is important as the connection between marriage and consumption commitments in the model is children.

Table 4 compares the statistics generated by the model with the other data targets. First, the model is able to generate an increase in the labor force participation rate of the same magnitude as the one observed in the data.²⁸ This is done with a combination of the parameters that control the utility of home goods and the exogenous forces over time in both the price of home inputs and the gender wage gap. The movements of married females into and out of the labor force are matched. The observed gender wage gap, measured only on observed wages, is also matched for both years. The model also generates the same fraction of expenditures on home inputs as the fraction of expenditure of household

²⁸ Another way of looking at the labor force participation numbers is to compare weekly hours worked by married women in the model versus the data. Knowles (2013) provides these data: married women worked on average 13.76 hours in 1967-1975 and 23.47 in 1997-2001. The model delivers 15.05 in 1970 and 25.8 in 2000.




Figure 6: Fraction of Single Females by Age

operations observed in the data. Finally, the fraction of expenditures that are measured as consumption commitments is also matched.

5.4.1 Non-Targeted Statistics

The model also provides some predictions for statistics that were not targeted in the estimation procedure outlined above. The ability of the model to match such a non-targeted statistic serves as a validation of the model.

When examining the effects of *risk* on marriage choices, a natural counterpart to examine the effects of *outcomes*. Specifically, we can empirically investigate how innovations to an individual's income affect his marital choices, and compare the data to the model. Accordingly, we run a linear probability model regression in which the dependent variable is whether or not a single male gets married, conditional on innovations in income in both the model and the data.²⁹ We also add a cubic polynomial in age as a control and, for

²⁹We also ran a similar regression with the level of income (and not differences) as the explanatory variable. However, we must note that, by running the regression in levels in the data, we are not filtering out any fixed effects (which are controlled for in the differences specification). The model nonetheless generates very similar estimates to the ones obtained with actual data.




Figure 7: Fraction of Married Couples with Children, Model and Data

actual data, add dummies for educational attainment.³⁰ The results are reported in Table 5. In the data, the coefficient on income differences is not significantly different from zero. The model counterpart is also very close to zero and is contained in the 95% confidence interval of the estimates in the data. Note also that the R^2 for both regressions is small, indicating that innovations in income do not explain much of the variation in the decision to get married; what seems to be important then is the amount of volatility households face and not the innovation immediately preceding marriage. Overall, the fact that the model generates very similar estimates to the ones obtained with actual data is reassuring.

The model is able to match several important features of the data, targeted or not. Crucially, the model generates the same pattern of selection out of singlehood and into marriage that is observed in the data. We can now use the model to understand the contributions of several channels to the observed delay in marriage.

³⁰For the actual data, we use a sample of white men from the PSID, since we need a panel data set for this exercise given that we must follow an individual over multiple periods of time to determine income innovations and whether he will get married.

Table 4: Model Fit —Targeted Moments

Statistic	Model	Data
Female LFP - 1970	.42	.42
Female LFP - 2000	.72	.72
Observed Gender Gap - 1970	.67	.66
Observed Gender Gap - 2000	.77	.76
% of wives moving into LF (2000)	21.4	23.8
% of wives moving out of LF (2000)	5.9	6.3
% of wives moving into LF - w/ children (2000)	21.3	19.8
% of wives moving out of LF - w/ children (2000)	7.2	6.9
Fraction of household expenditures on home inputs in 2000	.23	.23
Childrens Consumption: % of couple's expenditures, 2000	.30	.30

Table 5: Linear Probability Model —Marriage and Innovations in Income

	Coefficient	St. Error	95% CI	R^2
Data	0.001	0.005	[-.009, .011]	0.0072
Model	0.005	—	—	0.0648

Dependent Variable: Marital status dummy (married or single)

6 Results

In this section we decompose the effects of various mechanisms on the delay in marriage. To do this, we perform a series of counterfactuals that aim to isolate the effect of each particular channel. Each counterfactual works as follows: From the 1970 steady state, we change all parameters to the 2000 values, *except for the parameter of interest*. For example, when we study income volatility, we change the gender wage gap, the price of home inputs, and the residual component (μ_γ), and see how much is left to be explained by volatility. The counterfactual question is “What would have happened to the timing of marriage had income volatility not increased?” We then look at how much each mechanism affects the change from the model benchmark in 1970 to the model benchmark in 2000.

The results for the counterfactuals, except for the gender wage gap, are plotted in the different panels of Figure 8. Each figure plots the fraction of single males at each age for the benchmark years of 1970 and 2000, as well as the fraction that is computed under the counterfactual assumptions. We exclude the gender wage gap exercise since the effects are too minimal to be discernible in a figure.




Figure 8: Fraction of Single Males by Age —Effects of Different Channels on Marriage

The effect of rising income volatility on marriage can be inspected in panel (a). It is clear that shutting down any increase in income volatility causes more marriages to take place since we observe a lower fraction of singles in the counterfactual. As a way of quantifying the effect of increased income volatility, we average over the ages to examine how much of the overall decline in marriage between 1970 and 2000 is left to be explained once income volatility is kept at the 1970 level. The result is 20% of the decline in marriage is attributable to changes in income risk. Furthermore, this shows that the effects of consumption commitments and added gains to search due to rising income volatility dramatically outweigh the effects of the gains to spousal insurance.

Panel (b) plots the effects of the counterfactual for the technological progress in the home sector. Again, we average across the ages. The results show that declining prices for home inputs are also an important factor: they account for 14% of the decline in marriage. Meanwhile, the narrowing of the gender wage gap leads only to slightly less marriage in the economy; a change of only 1%. This is because, on one hand, when women earn more money, they find it easier to remain single; on the other hand, they become more attractive to men. Quantitatively, it turns out that these effects mostly cancel each other out, resulting in a small overall effect of the narrowing gender wage gap.

In sum, the results show that two channels (increasing income volatility and declining home input prices) have strong quantitative effects that lead to delays in marriage, while a third channel (the narrowing gender wage gap) does not. The effect of declining home inputs is significant, but going forward we will focus on the income volatility channel, as this paper seeks to establish a quantitative relationship between that channel and marriage delay. About one-fifth of the observed change between 1970 and 2000 can be attributed to higher income volatility. As previously shown, one of the factors of such a result is the role that consumption commitments play within marriage. In the next section, we show that income risk is still important for marriage decisions using a wide variety of robustness exercises.

6.1 Robustness

In section 4, we outline the mechanisms by which income risk may cause a delay in marriage through consumption commitments. The two channels outlined are the adjustment costs associated with changing children's consumption, and the fact that having children, and the resultant choice to spend money on them, moves people to a steeper part of their utility

curve, making them implicitly more risk averse. In this section, we perform a number of robustness exercises to check the result. In all of the analyses, we recalibrate the model, and then report an average effect of changing labor market volatility in precisely the same way outlined above. All results from this section can be found in Table 6.

A natural check of the results is to examine the effects of income risk assuming that couples do not face adjustment costs regarding their expenditures on children. This is done by imposing $\kappa = 0$. As expected, the result is weaker. It does not go away entirely, still allowing for volatility to account for 12% of the change. This is due to the fact that people still spend money on their children leaving them relatively more risk averse.

The next experiment asks how the results would change if couples were permitted to have up to two children rather than just one. In order to perform this exercise, we allow in the calibration for the mean utility bliss of the second child to differ from the first child and target the fraction of families that have at least two kids by age 33. As reported in Table 6, increased volatility accounts for 15% of the delay in marriage. The reason for this slightly lower number is that now couples have one extra margin of adjustment, namely the number of children they have over their life cycle.

Since the main mechanism here works through the risk taking of households, another natural check is to change the household's coefficient of relative risk aversion η . For this exercise, we use $\eta = 1.5$, a value also widely used in the literature. Naturally, with less risk aversion the effects of consumption commitments are lessened, but higher volatility is still responsible for 17% of changes in marriage over the period.

As previously discussed, the effect of volatility is through the risk that couples may, in the future, receive an adverse shock. The effects of volatility, therefore, are intricately linked to the discount factor β . Hence, our next experiment increases the rate at which households discount the future, i.e. we use a lower discount factor and set $\beta = 0.96$, also a standard value. In this scenario, increased volatility accounts for 16% of the rise in marriage between 1970 and 2000.

One important factor in the model is home production. Specifically, how substitutable are women with market goods? This level of substitutability matters as it affects their ability to smooth household income shocks: if the husband receives a bad shock but the wife cannot find a good substitute for her home production, then volatility will matter more as she cannot smooth out the adverse shock. We explore this possibility by allowing the substitutability between a married woman and home goods to be half as substitutable as in

Table 6: Robustness

Robustness Exercise	Δ in Marriage due to Volatility
No adjustment costs	12%
Two children	15%
Lower CRRA, $\eta = 1.5$	17%
Lower discount factor, $\beta = .96$	16%
Lower substitutability in home production function, $(1/2)\xi$	14%

the benchmark analysis. The results suggest that this parameter matters, however the fact that volatility has a quantitatively significant effect survives, as can be seen from Table 6.

7 Conclusions

There have been drastic changes in American society over the last 40 years. In particular, young adults have been delaying marriage. We contribute toward answering the most natural question: Why?

We propose a new hypothesis: increasing income volatility has led to a delay in marriage. The idea behind this hypothesis is simple. If marriage involves consumption commitments, such as children, then an increase in income volatility makes marriage less desirable. Young singles will thus delay marriage until a later point when they will ostensibly have greater incomes or accrued assets to offset these commitments. Despite the implicit insurance between spouses, this channel is quantitatively important.

We quantitatively assess this new hypothesis vis-à-vis others in the literature. In this paper, we estimate a structural search model of the marriage market with increasing income volatility, a narrowing gender wage gap, and decreasing prices of home inputs. We find that rising income volatility explains 20% of the decline in marriage. The decrease in the price of home inputs also explains about 14% of this decline. The narrowing of the gender wage gap has a small effect. Our analysis shows that our findings are robust to specifications of the model, specifically the presence of adjustment costs. In sum, we find that our hypothesis is quantitatively important, and that rising income volatility has a substantial impact on the delay in marriage.

The framework developed here could also be used to address different questions. For example, in the presence of consumption commitments, individuals may sort into jobs or

occupations with greater or lesser risk based on their marital status. Another possibility is an analysis of the impact on household formation of different government policies that affect the labor market. This mechanism can also be important accounting for the rise of singles that live with their parents, as young adults may increasingly value insurance from their parents.³¹ Moreover, this may also be important to explain the rise in cohabitation, which allows for some mutual insurance with lower commitment.³² We leave these possibilities for future research.

References

- [1] Aguiar, Mark and Erik Hurst. 2013. “Deconstructing Life Cycle Expenditure,” *Journal of Political Economy*, 121 (3), 437-492.
- [2] Ahn, Namkee and Pedro Mira. 2001. “Job Bust, Baby Bust? Evidence from Spain,” *Journal of Population Economics*, 14, 505-521.
- [3] 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.
- [4] Chetty, Raj, and Adam Szeidl. 2007. “Consumption Commitments and Risk Preferences,” *Quarterly Journal of Economics*, 122 (2), 831-877.
- [5] Gemici, Ahu and Steve Laufer. 2012. “Marriage and Cohabitation,” mimeo.
- [6] Gordon, Robert J. 1990. *The Measurement of Durable Goods Prices*. Chicago: University of Chicago Press.
- [7] 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.
- [8] Gould, Eric and Daniele Paserman. 2003. “Waiting for Mr. Right: Rising Inequality and Declining Marriage Rates,” *Journal of Urban Economics*, 53 (2), 257-281.

³¹See Kaplan (2012) for a study on the importance of parents as a source of insurance.

³²Gemici and Laufer (2012) study the difference between marriage and cohabitation with regards to levels of insurance and commitment.

- [9] 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.
- [10] Greenwood, Jeremy, Nezih Guner, Georgi Kocharkov and Cezar Santos. 2012. "Technology and the Changing Family: A Unified Model of Marriage, Divorce, Educational Attainment and Married Female Labor Force Participation," NBER Working Paper 17735.
- [11] Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu. 2005. "Engines of Liberation," *Review of Economic Studies*, 72 (1), 109-133.
- [12] Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. 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.
- [13] Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2011. "From Wages to Welfare: Decomposing Gains and Losses from Rising Inequality," mimeo.
- [14] Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2010. "The Macroeconomic Implications of Rising Wage Inequality in the United States," *Journal of Political Economy*, 118 (4), 681-722.
- [15] Hess, Gregory. 2004. "Marriage and Consumption Insurance: What's Love Got to Do with It?," *Journal of Political Economy*, 112 (2), 290-318.
- [16] 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.
- [17] Kaplan, Greg. 2012. "Moving Back Home: Insurance Against Labor Market Risk," *Journal of Political Economy*, 120 (3), 446-512.
- [18] Kaplan, Greg and Giovanni Violante. 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 82 (4), 1199-1239.

- [19] Karahan, Fatih and Serdar Ozkan. 2013. "On the Persistence of Income Shocks over the Life Cycle: Evidence and Implications," *Review of Economic Dynamics*, Vol. 16(3), pp. 452-476 (2013)
- [20] Knowles, John. 2013. "Why are Married Men Working So Much? An Aggregate Analysis of Intra-Household Bargaining and Labour Supply," *Review of Economic Studies*, 80, 1055–1085.
- [21] 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.
- [22] Lundberg, Shelly. 1985. "The Added Worker Effect," *Journal of Labor Economics*, 3 (1), 11-37.
- [23] McGrattan, Ellen, Richard Rogerson and Randall Wright. 1997. "An Equilibrium Model of the Business Cycle with Household Production and Fiscal Policy," *International Economic Review*, 38, 267-290.
- [24] Oppenheimer, Valerie K., Matthijs Kalmijn, and Nelson Lim. 1997. "Men's Career Development and Marriage Timing During a Period of Rising Inequality," *Demography*, 34 (3), 311-330.
- [25] Postlewaite, Andrew, Larry Samuelson, and Dan Silverman. 2008. "Consumption Commitments and Employment Contracts," *Review of Economic Studies*, 75 (2), 559-578.
- [26] Regalia, Ferdinando and Jose-Victor Rios-Rull. 2001. "What Accounts for the Increase in Single Households?," mimeo.
- [27] Santos, Cezar and David Weiss. 2012. "Risky Income, Risky Family: Marriage and Divorce in a Volatile Labor Market," mimeo.
- [28] Sommer, Kamila. 2014. "Fertility Choice in a Life Cycle Model with Idiosyncratic Uninsurable Earnings Risk", *Finance and Economics Discussion Series 2014-32*. Board of Governors of the Federal Reserve System (U.S.).

- [29] Southall, Humphrey and David Gilbert. 1996. "A Good Time to Wed?: Marriage and Economic Distress in England and Wales, 1839-1914," *Economic History Review*, 49 (1), 35-57.
- [30] Stevenson, Betsey and Justin Wolfers. 2007. "Marriage and Divorce: Changes and their Driving Forces," *Journal of Economic Perspectives*, 21 (2), 27-52.

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 never-married white males comes from the Integrated Public Use Microdata Series (IPUMS) Census for both 1970 and 2000.

Figure 2: The data for the median age of individuals at their first marriage comes from the U.S. Census Bureau (Table MS-2). The standard deviation of persistent shocks are the values estimated by Heathcote, Storesletten, and Violante (2010). Their series for the variances is smoothed using the HP-filter. The non-filtered data is very noisy but is also positively correlated with the series for age at first marriage (0.40). The values for the variances reported in their paper are very similar to the ones obtained here in Section 5.

Figure 3: The fraction of households with children is computed using data on families headed by a white male from the IPUMS-Census data for both 1970 and 2000.

Figure 7: The data for married couples with children comes from the IPUMS-Census data for both 1970 and 2000. We compute the fraction of married couples with at least one child in the household, by age.

Figures 9 and 10: The data for never-married white males comes from the IPUMS-Census for 1970 and 2000. A male is college-educated if he has at least 16 years of education.

Figure 11: The data for white males comes from the IPUMS-Census for both 1970 and 2000. The lines labeled “1970” and “2000” are the same as in Figure 1. Singles that are not cohabitating are represented in the line “2000 - no cohabitation”. Note that, for 2000, we do not consider individuals that are not married but *cohabit* and have children as singles. This explains the small difference between the lines “2000” and “2000 - legal”.

B Other Data on the Delay in Marriage

In this section, we present data on the delay in marriage for different groups of the population for both 1970 and 2000. Figures 9 and 10 plot the fraction of white males that are single conditional on their educational attainment, i.e., whether they have a college degree or not. It is clear that marriage has been delayed by individuals of both education groups. Moreover, Greenwood, Guner, Kocharkov and Santos (2012) show that, even in an economy with a much higher number of college-educated women, marriage does not decrease substantially (see their Table 5).




Figure 9: Percentage of College-Educated White Males Never Married, by Age

One form of living arrangement that we abstract from in this paper is cohabitation. Young adults could have been opting to cohabitate instead of getting married in 2000. Figure 11 shows that this is not the case. Even though there is a fraction of the population that currently cohabitates, an increase in the fraction of singles among young adults is clearly visible in the figure. Second, if we added a cohabitation possibility, the model would presumably give a rise in cohabitation in addition to a decline in marriage, as people would be able to enjoy some of the economies of scale and a degree of mutual insurance implicit in cohabitation, without the commitment involved in marriage. The difference in commitment between cohabitation and marriage is the subject of Gemici and Laufer (2012), and beyond the scope of this exercise. It is for this reason that we abstract from




Figure 10: Percentage of Non-College-Educated White Males Never Married, by Age

cohabitation entirely. Note that our definition of married individual differs a little from the legal definition reported in the Census. In particular, we treat individuals that *cohabitate and have children* as effectively being married. This causes the small adjustment observed in Figure 11.




Figure 11: Living Arrangements of Young Adults

C Estimation of Income Processes

We use data from the PSID for all waves between 1968 and 2009. After 1997, the PSID becomes a biennial survey. We follow Heathcote, Storesletten, and Violante (2010) in estimating processes for the biennial PSID.³³ 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.

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 3.1 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 (2013).

³³See their online appendix section 2.3.