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

Cezar Santos, David Weiss[†]

University of Pennsylvania

Preliminary and incomplete

Abstract

One of the most striking changes in American society since 1970 has been the decline and delay of marriage. The fraction of young men and women who have never been married has increased steadily and significantly. Labor market volatility has increased over the same period. This paper proposes a new hypothesis that links these two facts. Specifically, if marriage involves consumption commitments, then a rise in income volatility could result in a delay of marriage. We assess this new hypothesis vis-à-vis others in the literature, using an estimated structural model. We find that the increase in volatility accounts for about 40% of the observed change in marriage. The decrease in the price of home inputs explains about a third of the data. The effects of the narrowing of the gender wage gap are small.

Keywords: decline and delay of marriage, income volatility, gender wage gap, home production, technological progress in the household, female labor force participation, young adults, search models of marriage, structural estimation, simulated method of moments.

*September 2011. We are very grateful to Jeremy Greenwood, Iourii Manovskii, and Dirk Krueger for their advice. We would also like to thank Fatih Karahan, Serdar Ozkan, Greg Kaplan, and seminar participants at the Penn Macro Lunch, the 2011 Midwest Macro Conference, the 2011 Minerva Summer Institute at the Hebrew University, the SED Meetings Ghent, the 2011 Cologne Macro Workshop and the 2011 Asian Meeting of the Econometric Society for their comments. The usual disclaimer applies.

[†]E-mails: santosca@econ.upenn.edu and daweiss@econ.upenn.edu

1 Introduction

One of the most striking changes in American society over the last 40 years has been the decline and delay of marriage. The fraction of young men and women who have never been married has steadily and significantly increased since before 1970. This trend has captured the attention of both academic researchers and the general public[1]. In the popular press, the current young generation is being referred to as "boomerang kids", due to their increased propensity to move back home in their 20s, rather than to get married and start their own households. Frequently, these young adults are being described as "failing to launch"[2]. The question here is: Why not settle down already?

To illustrate the trend, Figure 1 shows the increase, by age, of the percentage of young white males who have never been married[3]. This graph illustrates how marriage has been delayed. The numbers are striking. In 1970, only 26% of 25 year old white males were never married. By 2000, this number had more than doubled to 57%[4].

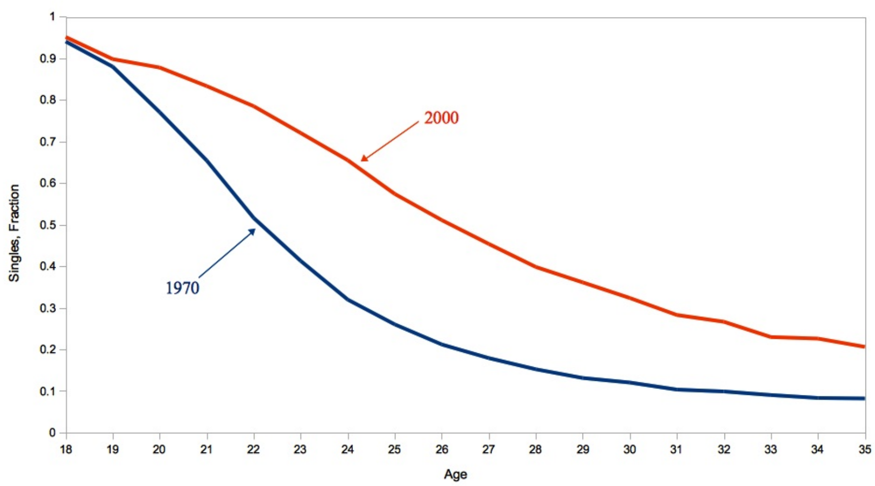

Figure 1: Percentage of White Males Never Married, by Age. Source: US Census

The economics literature has also noted a rise in labor market volatility over this period. That is, many studies have documented increasing variances of shocks in income processes,

[1]For an excellent review of the academic literature, see Stevenson and Wolfers (2007).
[2]New York Times magazine, August 18th 2010
[3]Our definition of married includes cohabiting couples that have children.
[4]The graph for white women looks very similar. For data on cohabitation and by education groups, see appendix A.

such as Gottschalk and Moffitt (1994), Katz and Autor (1999), and Heathcote, Perri and Violante (2010). The effects of the changing labor market have been widely discussed[5]. However, to the best of our knowledge, no work has been done relating changes in labor market volatility with changing marriage decisions of young adults.

The contribution of this paper is to establish and quantitatively assess the link between rising income volatility and the delay of marriage. Specifically, suppose marriage involves consumption commitments such as children, mortgages, etc. Then a rise in the volatility of income would result in a delay in marriage as these commitments become less desirable. That is, it might be preferable to wait until one receives a favorable income shock, or search longer for a better spouse, before settling down with a family. However, there are other channels by which changing income volatility affects marriage decisions. One is that marriage provides insurance between spouses. The idea here is that marriage allows for diversification of risk since income shocks between spouses need not be perfectly correlated. Therefore, when income volatility increases, marriage might become more desirable. This mechanism is highlighted by Hess (2004). Additionally, Gould and Paserman (2003) argue that rising inequality would lead to a delay in marriage as there could be further gains from search. All three of these channels are incorporated in our study.

In order to carefully assess the impact of increasing labor market volatility on marriage decisions, two additional changes to the US labor market over this time period must be included. The first is the declining gender wage gap. This is important as it speaks to how well a woman's income can replace her husband's income when he receives a bad shock in the labor market. Additionally, Regalia, Rios-Rull and Short (2008) argue that the decrease in the gender wage gap is itself a major source of the delay in marriage. They argue that when women become richer they can afford to be pickier with the mate they choose. The second significant change is the increase in the labor force participation rate of married women. This trend is another important factor influencing spousal insurance. Greenwood, Seshadri and Yorukoglu (2005) make the case that less expensive household goods, such as washing machines and refrigerators, led to the increase in labor force participation. Additionally, Greenwood and Guner (2009) argued that since cheaper household goods made the cost of running a household lower, there were decreases in the gains from marriage. Men no longer

[5]For example, Kambourov and Manovskii (2009) argue that increasing occupational mobility, which is related to income volatility, may be important to understand the increase in income inequality in the US. Krueger and Perri (2005) discuss the relationship between rising income inequality and consumption inequality.

need women to specialize in the production of home goods, and women have the time to go out into the labor force themselves. The results are a decrease in marriage and an increase in the labor force participation of married women. For our purposes, including a narrowing gender wage gap and declining prices of home inputs will help quantify the effects of rising labor market volatility. That is, higher wages and more participation for married women lead to more spousal insurance, which is fundamental for our quantitative analysis.

In order to carefully assess the impact of increasing labor market volatility on marriage decisions, two additional changes to the US labor market over this time period must be included. The first is the declining gender wage gap. This is important as it speaks to how well a woman's income can replace her husband's income when he receives a bad shock in the labor market. The second significant change is the increase in the labor force participation rate of married women. This trend is clearly an important factor influencing spousal insurance. Greenwood, Seshadri and Yorukoglu (2005) make the case that less expensive household goods, such as washing machines and refrigerators, led to the increase in labor force participation. We include progress in household technology in order to generate increased participation. For our purposes, including a narrowing gender wage gap and declining prices of home inputs will help quantify the effects of rising labor market volatility. That is, higher wages and more participation for married women lead to more spousal insurance, which is fundamental for our quantitative analysis.

The forces mentioned in the previous paragraph themselves are relevant to marriage decisions. Regalia, Rios-Rull and Short (2008) argue that the decrease in the gender wage gap is itself a major source of the delay in marriage. They argue that when women become richer they can afford to be pickier with the mate they choose. Moreover, Greenwood and Guner (2009) argue that since cheaper household goods made the cost of running a household lower, there were decreases in the gains from marriage. Men no longer need women to specialize in the production of home goods, and women have the time to go out into the labor force themselves. The result is a decrease in marriage. Since we include both of these channels, we can quantitatively assess their importance.

We build and estimate a structural model, containing marriage as a consumption commitment, as well as a form of consumption insurance. Our results show that rising income volatility explains about 40% of the decline in marriage. This shows that the effects of consumption commitment and added gains to search due to rising labor market volatility dramatically outweigh the effects of the gains to spousal insurance. The decrease in the price of home inputs explains about a third of this decline. The narrowing of the gender wage gap has small effects. In sum, out of the three mechanisms studied, rising labor market volatility is the most important cause of the delay of marriage.

Marriage, in our model, serves as a form of consumption commitment. Therefore, this paper is a contribution to the consumption commitment literature along the lines of Chetty and Szeidl (2007) and Postlewaite, Samuelson and Silverman (2008). Chetty and Szeidl (2007) discuss how risk averse agents can become even more risk averse in the presence of consumption commitments. Postlewaite, Samuelson and Silverman (2008) study how risk neutral individuals can behave as if they have preferences over risk when they face commitments. When labor market volatility increases, we argue that the effect is to delay marriage decisions. Our mechanism assumes that the difference between marriage and other forms of social arrangements, such as cohabitation, is consumption commitments. One paper that discusses the role of consumption commitments within a family economics context and rising income volatility is Sommer (2009). In her paper, she argues that rising volatility leads to a delay in fertility (children being a form of consumption commitments). Given her emphasis on fertility, she doesn't model the marriage market as we do here. Moreover, she doesn't take into account the change in the labor force participation rate of married women or the gender wage gap as we do. These are important ingredients for our study given that they affect the ability of married households to cover the costs of having children, and, more significantly, they influence the importance of spousal insurance.

This paper is organized as follows: Section 2 presents evidence on consumption commitments by marital status. 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

The main premise of this paper is that marriage comes with consumption commitments. We provide evidence on two forms of commitments: children and housing[6].

First we turn to children. As Figure 2 shows, a strong majority of married people have children, while the opposite is true for singles. For example, according to Panel (a) in the figure, in 2000, 59% of 25 year old married males had at least one child in their household, whereas only 9% of their single counterparts did - those numbers were 68%

[6]All data in this Section refers to white American males and, unless otherwise noted, come from the US Census.

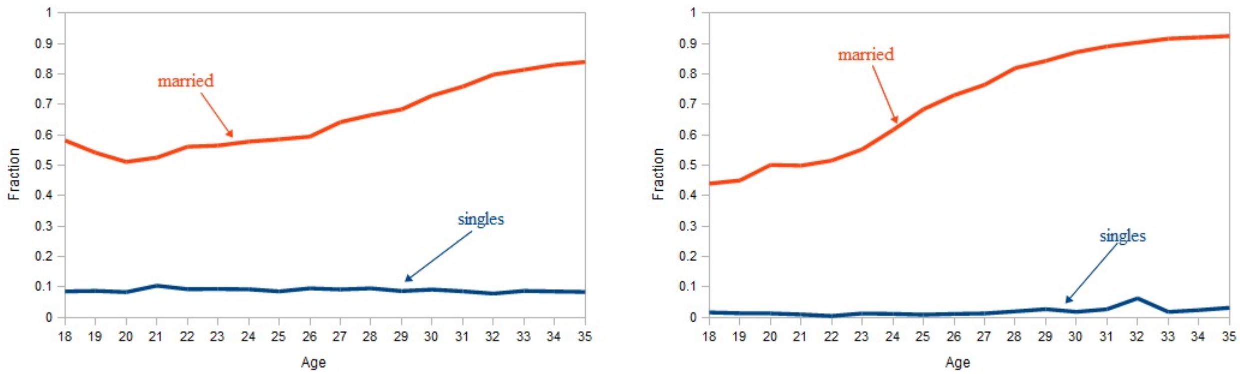

(a) 2000 (b) 1970

Figure 2: Presence of Children in the Household by Marital Status

and 1% respectively in 1970; see Panel (b). What Figure 2 misses is the fact that some married people are newlyweds who simply haven't had time to start a family yet. Using the Panel Study of Income Dynamics (PSID), we can follow married couples to see if they have children shortly after marriage. These numbers are reported in Table 1. These data show a strong link between marriage and fertility, a notoriously expensive and persistent form of consumption commitment.

Table 1: Marriage and Fertility - White Males

	Married with children*	Single without children
1970	84%	98%
2000	81%	90%

* = Children now or within the next two years.

As further evidence of the relationship between marriage and consumption commitments, we turn to housing. Chetty and Szeidl (2007) categorize home ownership as a form of consumption commitment, while renting is not [7]. As Figure 3 shows, married males are much more likely to own their own homes than singles. At age 25, in 2000, 51% of married males are home owners whereas only 26% of singles are. The difference increases to 71% versus 38% by age 30. In 1970, the gap is even wider.

[7]Additionally, Karahan and Rhee (2011) argue that home ownership results in "lock-in" effects, that discourage geographic responses to changes in labor market conditions. The effects of this "lock-in" are similar to consumption commitments.

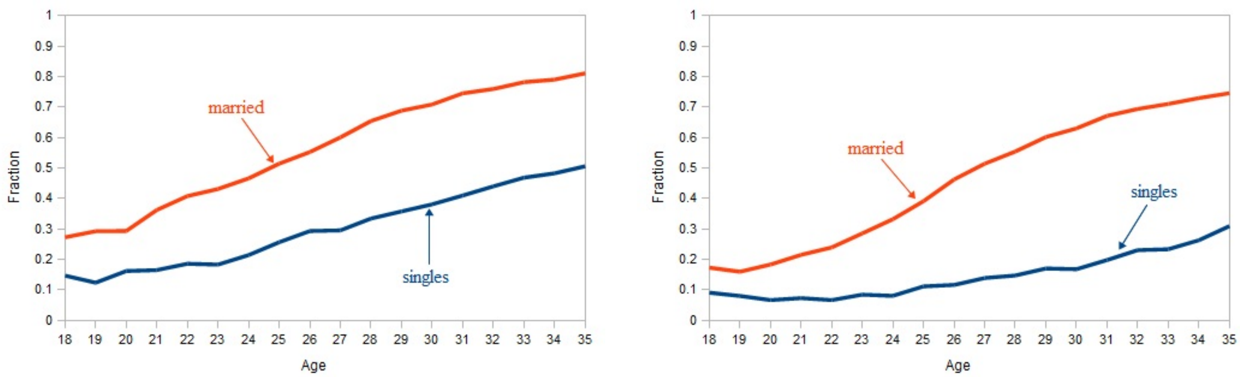

(a) 2000 (b) 1970

Figure 3: Home ownership by Marital Status

3 The Model

There are 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 wages in the model are equal to efficiency units of labor supplied.

The amount of efficiency units of labor, I, each agent supplies follows a stochastic process around a deterministic trend.

$$I = w\phi_g f_g(a), \tag{2}$$

where w is an idiosyncratic shock, while the deterministic trend is composed of ϕ_g, a gender wage gap, and $f_g(a)$, a gender specific deterministic age income profile.

The shock w consists of a persistent shock z (with innovations η) and a transitory shock

ϵ. An increase in volatility is measured by changing σ_η^2 (σ_ϵ^2), which control the variance of the persistent (transitory) shocks. The shock process is specific to the agent's marital status; both the variance and the persistence 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. Additionally, we allow for persistent 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. 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):

$$\begin{aligned} \ln w_s &= z_s + \epsilon_s \\ z_s &= \delta_s z_{s,-1} + \eta_s \\ \eta_s &\sim N(0, \sigma_{\eta,s}^2) \\ \varepsilon_s &\sim N(0, \sigma_{\varepsilon,s}^2). \end{aligned} \tag{3}$$

For couples (denoted by the subscript m), the process takes the following form:

$$\begin{aligned} \ln w_m &= z_m + \epsilon_s \\ z_m &= \delta_m z_{m,-1} + \eta_m \\ \eta_m &\sim N\left(0, \begin{bmatrix} \sigma_{\eta,m}^2 & \rho \\ \rho & \sigma_{\eta,m}^2 \end{bmatrix}\right) \\ \varepsilon_m &\sim N\left(0, \begin{bmatrix} \sigma_{\epsilon,m,t}^2 & 0 \\ 0 & \sigma_{\epsilon,m}^2 \end{bmatrix}\right). \end{aligned} \tag{4}$$

The parameter ρ controls the correlation of spousal shocks. This allows us to get the appropriate level of spousal insurance in the model. This insurance is a counter mechanism to labor market volatility causing a delay in marriage, so getting the level right is important.

As noted, the amount of efficiency units available to an agent also varies with his/her age a according to the function $f_g(a)$. 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).

The home good, n, is produced by a constant elasticity of substitution production function between home inputs, d, and time, h:

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

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

3.2 Preferences

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

$$u^s(c,n) = \frac{c^{1-\lambda}}{1-\lambda} + \alpha\frac{n^{1-\zeta}}{1-\zeta}, \tag{6}$$

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

For married agents, we assume a unitary model, i.e. spouses make decisions jointly. Agents pool their resources, and enjoy economies of scale in consumption - ψ is the parameter controlling economies of scale in marriage. The utility function for married agents then reads:

$$u^m(c,n) = \frac{\left(\frac{c}{1+\psi}\right)^{1-\lambda}}{1-\lambda} + \alpha\frac{\left(\frac{n}{1+\psi}\right)^{1-\zeta}}{1-\zeta}. \tag{7}$$

The agent's total utility is equal to the expected discount value of lifetime utility.

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

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

When individuals first get married, they also enjoy a marital bliss utility increase γ. This is a stochastic shock drawn from the distribution $\Gamma(\gamma)$. We assume that $\gamma \sim N(\mu_\gamma, \sigma_\gamma)$. This utility shock is received only once at the start of married life. This represents the (stochastic) lifetime discounted utility of being married that arises due to non-economic reasons.

Finally, a wife incurs a utility cost of κ_w (κ_h) if she moves into (out of) the labor force. This cost is a reduced form way of capturing the various costs married women encounter when changing their labor force status. It allows the model to generate movements into and out of the labor force in accordance with the data, ensuring appropriate spousal insurance. With a slight abuse of notation, define the function $\kappa(l)$ to represent the utility cost from moving into and out of the labor force such that $\kappa(0) = \kappa_w$ ($\kappa(1) = \kappa_h$) if the wife is moving into (out of) the labor force.

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, which is allowed to depend on their gender. Thus, their budget constraint will be given by

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

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. $(1 + r)$ is the gross interest rate.

When married, spouses pool their resources, and enjoy economies of scale in consumption. Furthermore, there are consumption commitments. This is modeled as a lump sum cost that married agents have to 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 and w^* 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

$$c + pd + c_k + b' = wf_1(a)\tau_1^m + l^f \phi w^* f_2(a)\tau_2^m + (1 + r)b. \tag{10}$$

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.

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 to get married. Note that marriage is an absorbing state, i.e. there is no divorce[8].
- Married agents choose whether or not the wife works[9]. All agents optimally divide their income between consumption goods, home inputs and savings. Consumption takes place.

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 and non-market goods. They also have to decide whether or not to get married with a potential mate. Married agents have a similar consumption decision regarding market and home-produced goods and they have to decide whether the wife should work or not. We will now describe each household's problem recursively.

Let's start with couples. The state vector for married households consists of a wage for the husband w, a wage offer for the wife w^*, the current assets level b, an indicator function l representing whether or not the wife worked last period, and their age a. Then

[8] This is a simplifying assumption, to make modeling marriage and keeping track of singles distributions easier. Since we are trying to explain timing of first marriage, the only issue is whether or not there are a lot of young divorces for single 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 not too many of these people to worry about in the data, we exclude them from the model.

[9] That is, the extensive, not intensive, margin of female labor force participation.

the married value function can be written as follows

$$
\begin{aligned}
V^m(w, w^*, b, l, a) &= \max_{l^f \in \{0,1\}, b' \geq 0, c \geq 0, d \geq 0} u^m(c, n) - \mathcal{I}(l^f \neq l)\kappa(l) + \beta E_{w', w^{*\prime}} V^m(w', w^{*\prime}, b', l^f, a+1) \\
&\text{s.t.} \\
c + pd + c_k + b' &= w f_1(a)\tau_1^m + l^f \phi w^* f_2(a)\tau_2^m + (1+r)b \\
n &= \left[\theta d^\xi + (1-\theta)\left(2 - \tau_1^m - l^f \tau_2^m\right)^\xi\right]^{1/\xi}
\end{aligned}
\tag{11}
$$

where $\mathcal{I}(x)$ is an indicator function that takes the value of 1 if x is true and 0 otherwise. Define the policy functions for the married problem as follows: $l^f = P_l^m(w, w^*, b, l, a)$ for the woman's labor force decision, $d = P_d^m(w, w^*, b, l, a)$ for choice of home inputs, $c = P_c^m(w, w^*, b, l, a)$ for the consumption decision, and $b' = P_b^m(w, w^*, b, l, a)$ for the savings decision.

The value function for singles of gender g after the marriage market is as follows:

$$
\begin{aligned}
V^s(w, b, g, a) &= \max_{b' \geq 0, c \geq 0, d \geq 0} u^s(c, n) + \beta E_{w'} B(w', b', g, a+1) \\
&\text{s.t.} \\
c + pd + b' &= w f_g(a)\tau_g^s + (1+r)b \\
n &= \left[\theta d^\xi + (1-\theta)(1-\tau_g^s)^\xi\right]^{1/\xi}.
\end{aligned}
\tag{12}
$$

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.

Now, we can turn our analysis to the marriage phase. Every single person draws a potential partner of the opposite gender randomly. 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. A marriage occurs if and only if both agents agree to marriage. Formally, a marriage occurs if and only if

$$
\underbrace{V^m(w, w^*, b+b^*, 1, a) + \gamma > V^s(w, b, 1, a)}_{\text{male's decision}} \text{ and } \underbrace{V^m(w, w^*, b+b^*, 1, a) + \gamma > V^s(w^*, b, 2, a)}_{\text{female's decision}}. \tag{13}
$$

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

$$J(w, w^*, b, b^*, \gamma, a) = \begin{cases} 1, & \text{if (13) holds,} \\ 0, & \text{otherwise.} \end{cases} \tag{14}$$

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

$$\begin{aligned} B(w, b, g, a) &= \int\int \{J(w, w^*, b, b^*, \gamma, a)\,[V^m(w, w^*, b + b^*, 1, a) + \gamma] \\ &\quad + (1 - J(w, w^*, b, b^*, \gamma, a))\, V^s(w, b, g, a)\} d\widehat{\mathbf{S}}(w^*, b^*, g^*, a) d\Gamma(\gamma), \end{aligned} \tag{15}$$

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

$$\begin{aligned} \mathbf{S}(w', b', g, a + 1) = \iiiint & (1 - J(w, w^*, b, b^*, \gamma, a))\, \mathcal{I}(P_b^s(w, b, g, a) \leq b') \times \\ & \times \mathbf{S}(w, b, g, a) d\mathbf{S}(w^*, b^*, g^*, a) d\mathbf{W}^s(w', w) d\Gamma(\gamma) \end{aligned} \tag{16}$$

where g^* represents the opposite gender and $\mathbf{W}^s$ represents the wage 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 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 and* *bachelors,* $V^s(w,b,g,a)$*,* $V^m(w,w^*,b,l,a)$*, and* $B(w,b,g,a)$*, policy functions* $P_c^s(w,b,g,a)$*,* $P_d^s(w,b,g,a)$*,* $P_c^m(w,w^*,b,l,a)$*,* $P_d^m(w,w^*,b,l,a)$*,* $P_l^m(w,w^*,b,l,a)$*, a matching rule for singles* $J(w,w^*,b,b^*,\gamma,a)$*, 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* $P_c^s(w,b,g,a)$*,* $P_d^s(w,b,g,a)$*,* $P_b^s(w,b,g,a)$ *solve the single's problem (12), 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* $V^m(w,w^*,b,l,a)$ *and the policy functions* $P_c^m(w,w^*,b,l,a)$*,* $P_d^m(w,w^*,b,l,a)$*,* $P_l^m(w,w^*,b,l,a)$ *solve the couple's problem (11).*

3. *The value function* $B(w,b,g,a)$ *solves the bachelor's problem (15), given the value functions for singles and couples,* $V^s(w,b,g,a)$ *and* $V^m(w,w^*,b,l,a)$*, 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 (14), taking as given the value functions* $V^s(w,b,g,a)$ *and* $V^m(w,w^*,b,l,a)$*.*

5. *The stationary distribution* $\mathbf{S}(w,b,g,a)$ *solves (16), taking as given the matching rule* $J(w,w^*,b,b^*,\gamma,a)$*.*

4 Mechanisms

Our purpose is to quantitatively explain the delay in marriage between 1970 and 2000. In this section we discuss the mechanisms in the model.

4.1 Labor Market Volatility

This is the chief hypothesis we propose. Rising labor market volatility, as defined by increasing variances to both persistent and transitory shocks (increasing $\sigma_{\epsilon,m}^2$, $\sigma_{\eta,m}^2$, $\sigma_{\epsilon,s}^2$, and $\sigma_{\eta,s}^2$), has multiple effects.

1. Consumption commitments, as embodied in the parameter c_k, cause an increase in risk aversion among married agents relative to single agents. When volatility increases, this decreases the value of being married relative to being single, causing *less* marriage.

2. "Waiting for Mr. Right" effect of increased benefits from search. That is, when volatility increases, the cross-section distribution of income becomes more disperse.

As a result, the benefit from searching for a spouse with higher income increases. This would cause *less* marriage.

3. When labor market volatility increases, mutual insurance between spouses becomes more important. This would tend to cause *more* marriage.

If effects 1 and 2 are quantitatively more significant than effect 3, then increasing volatility will lead to less marriage.

4.2 Price of Home Inputs

Greenwood and Guner (2009) explain in detail the mechanism by which a decrease in the price of inputs for home production (such as washing machines) would tend 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. This mechanism is in our model.

However we should note the presence of an additional channel by which the change in the price of home inputs affects the marriage decision. Marriage in our model is in some ways a normal good. There is a consumption cost to marriage (c_k), and a utility gain γ. While there are other pros and cons to getting married in the model, to the extent that μ_γ determines an agent's choice of whether or not to marry, the agents are more likely to marry when they are wealthier and can more easily cover the fixed costs. If marriage is a normal good, then a decrease in the price of home inputs will have an income effect leading to more marriage.

4.3 Gender Wage Gap

The final mechanism explored is narrowing the gender wage gap. Again, we will highlight the various channels through which a change in the gender wage gap affects marriage decisions.

1. The (economic) quality of women increases due to their higher incomes, leading men to be more willing to marry. This should *increase* marriage.

2. Women are happier to stay single and search longer, since they are wealthier. That is, their outside option from marriage becomes more attractive. This should *decrease* marriage.

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 (that 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

We solve using backwards induction on value functions. The model is solved for males from ages 18 to 35 (16 to 33 for females). After age 35 (33 for females), the marriage market is shut down. They live until age 65 (63 for females).

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 gap, the price of household inputs, and the mean of the marital bliss shock distribution μ_γ. The reason for changing μ_γ 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 2.

The time discount factor β is set to 0.96 and the coefficient of relative risk aversion (CRRA) for market goods is set to 2.0. Both values are standard in the macroeconomic literature.

The parameters for the production function of non-market goods are taken from McGrattan, Rogerson and Wright (1997) and are also used by Greenwood and Guner (2009). McGrattan et al (1997) estimate these parameters using business cycle frequency data.

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 is computed using the US 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 2.

For the parameter ψ that controls the degree of economies of scale in a household, we use the OECD equivalence scale. For the decline in the price of home inputs, we set it to 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 kitchen and other household appliances from the National Income and Product Accounts declined at about 1.5% a year since 1950.

The life-cycle profile of income for each gender is computed by fitting a cubic polynomial over the mean income at each different age in the US Census.

For the consumption floor, we use the number provided by Kaplan (2010). In his model, also of young adults, Kaplan fixes unemployment benefits to be \$500/month and a consumption floor of \$100/month. We take \$600/month (\$7200/year) to be our consumption floor.

In this model, assets have the flavor of liquid assets, as precautionary savings is the main motive for asset accumulation. We pick $r = 0$ as many forms of liquid assets have zero returns (such as checking accounts), and may even be negative after adjusting for inflation.

5.3 Structural Estimation

The remaining parameters are estimated by the Simulated Method of Moments. We first need a set of data moments. 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 the 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:

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

In our case, we need to estimate 9 parameters (in addition to 12 parameters that control the income processes) so that we have the following vector of parameters to be estimated: $\Omega = (\alpha, \zeta, \kappa_h, \kappa_w, \mu_{\gamma,1970}, \mu_{\gamma,2000}, \sigma_\gamma, p, c_k, \Upsilon)$, where Υ is a vector that contains the labor

Table 2: Parameters set using a priori information

Parameter	Description	Value	Source
β	Time Discount Factor	.96	Standard
λ	CRRA Consumption	2.0	Standard
θ	Weight on home inputs in production	.206	McGrattan et al (1997)
ξ	CES home production	.189	McGrattan et al (1997)
ρ	Correlation of Spousal Persistent shocks	.25	Hyslop (2001)
τ_1^s	% of time spent at work (single males)	.37	US Census
τ_2^s	% of time spent at work (single females)	.35	US Census
τ_1^m	% of time spent at work (married males)	.40	US Census
τ_2^m	% of time spent at work (married females)	.32	US Census
ψ	Economies of Scale	.7	OECD equiv. scale
–	Decline in the price of home inputs	6%/year	Greenwood and Guner (2009)
$f_g(a)$	Age Profile of Income	–	US Census
–	Consumption Floor	\$7200/Year	Kaplan (2010)
r	Interest Rate	0%	Return on Liquid Assets

market parameters[10].

5.3.1 Labor Market Parameters

For the data on income processes, we use data on white men from the Panel Study of Income Dynamics (PSID)[11] for the years 1968-1997. We first run a Mincer-style regression for every year in the sample controlling for education and age, generate the residuals, and then estimate the parameters from (3) and (4) using the 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. The results[12] of this estimation procedure are reported in the Data column in Table 3. Even though this procedure is popular in the literature, estimates by

[10] That is, $\Upsilon = \left(\delta_{s,1970}, \delta_{s,2000}, \sigma^2_{\eta,s,1970}, \sigma^2_{\eta,s,2000}, \sigma^2_{\epsilon,s,1970}, \sigma^2_{\epsilon,s,2000}, \delta_{m,1970}, \delta_{m,2000}, \sigma^2_{\eta,m,1970}, \sigma^2_{\eta,m,2000}, \sigma^2_{\epsilon,m,1970}, \sigma^2_{\epsilon,m,2000}\right)$

[11] For details on sample selection and estimation procedure, see appendix B.

[12] 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, we opted for the simpler model described above.

marital status are not common. However, the variances of the persistent shocks that we estimate are in line with the numbers reported by Heathcote, Storesletten and Violante (2010) and Meghir and Pistaferri (2004), for example. The variances of the transitory shocks are a little higher than their estimates.

To get the gender wage gap, we run a Mincerian regression controlling for age, education, and a gender dummy using Census data[13] from both 1970 and 2000. We run this regression on observed wages. The coefficient on the gender dummy is our data target for the gender wage gap. The value of the estimates are $\phi_{1970} = 0.67$ and $\phi_{2000} = 0.75$.

Simply using these estimates in the model generates sample selection problems. Specifically, for the income process, there is selection as to who is married/single. If singles wait for good persistent 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, we run the regression on observed wages. Clearly, there is selection as to which women are working and which are not. To solve these problems, we take an indirect inference approach[14]. That is:

1. Guess parameter values for married and single agents, and a gender wage gap .
2. Solve and simulate the model.
3. Run the same estimators on the simulated data as on actual data.
4. Check if model estimates match data estimates.

According to Table 3, the model is able to generate estimates for the auxiliary GMM model that are very close to the estimates obtained with actual data from the PSID.

5.3.2 Other Estimated Parameters

The first step is to choose the data moments (again, in addition to the labor market parameters discussed in the previous section). We try to choose 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.

[13] Again, see appendix for details.
[14] For a detailed description of this technique, see Gourieroux and Monfort (1996).

Table 3: Parameters for the Income Process

Parameter	Description	Data	Model	Input	SE
Married					
$\delta_{m,1970}$	Autoregressive Coefficient	0.9959	0.9811	0.9959	N/A
$\delta_{m,2000}$	Autoregressive Coefficient	0.9959	0.9783	0.9959	N/A
$\sigma^2_{\eta,m,1970}$	Persistent Shock Variance	0.0063	0.0062	0.0050	N/A
$\sigma^2_{\eta,m,2000}$	Persistent Shock Variance	0.0214	0.0238	0.0210	N/A
$\sigma^2_{\varepsilon,m,1970}$	Transitory Shock Variance	0.0987	0.1029	0.1111	N/A
$\sigma^2_{\varepsilon,m,2000}$	Transitory Shock Variance	0.1155	0.1133	0.1250	N/A
Singles					
$\delta_{s,1970}$	Autoregressive Coefficient	0.9344	0.9386	0.8600	N/A
$\delta_{s,2000}$	Autoregressive Coefficient	0.9344	0.9301	0.8950	N/A
$\sigma^2_{\eta,s,1970}$	Persistent Shock Variance	0.0069	0.0069	0.0110	N/A
$\sigma^2_{\eta,s,2000}$	Persistent Shock Variance	0.0233	0.0245	0.0320	N/A
$\sigma^2_{\varepsilon,s,1970}$	Transitory Shock Variance	0.1309	0.1269	0.1400	N/A
$\sigma^2_{\varepsilon,s,2000}$	Transitory Shock Variance	0.1532	0.1503	0.1600	N/A

Let's 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 price of home inputs in 2000 p [15]. One of the targets we use to identify these parameters is the fraction of income spent on household operation 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, Personal Care and Entertainment. In 2000, this number was 23% of household consumption according to the Consumer Expenditure Survey (CEX). 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 these parameters. We target LFPR in both 1970 and 2000 since this can give us information on the elasticity of labor supplied by married females. The LFPR for married females was 0.42

[15] For the price of home inputs in 1970, we increase the price p by 6% per year, the number reported by Greenwood and Guner (2009) - see Table 2.

in 1970 and 0.72 in 2000 according to the US Census.

In our model, married females would be able to move into and out of the labor force freely if it were not for the parameters κ_h and κ_w. In the absence of these parameters, this fact might lead to counterfactually high levels of movements into and out of the labor force, in turn leading to too much consumption insurance between husband and wife. So, we choose these two parameters so that the model generates reasonable movements. In the data, we measure these movements using PSID data. Since this is a panel data set, we can follow married females over time and we can observe how often they move. The data targets we use are the fraction of wives that move into and out of the labor force in 1970. The percentage into the labor force is 4%, the percentage out is 7%.

In our model, couples have to pay a fixed consumption cost c_k every period; this is supposed to represent all the consumption commitments that married agents have to incur. Most couples throughout the period of the analysis had children; most of them had more than one[16]. 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. We use the Engel estimate from Betson et al (2001) of 30.1%. This is an estimate in the middle of the ranges of possible values that they report. Considering that marriage often results in more than one child, mortgages, and other commitments, we consider this to be a be a reasonable lower bound on the consumption commitments faced by married couples.

We then turn to the parameters that govern the marital bliss shocks in 1970: $\mu_{\gamma,1970}$ and σ_γ. These parameters govern the strength and variety 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. Then, a potential couple wouldn't have to worry about all the different potential relationships are also available in the economy - they are all the same. Then, $\mu_{\gamma,1970}$ would control the level of marriages. With a more dispersed distribution, potential mates might prefer to wait for a better draw. In order to identify these two parameters, we target the overall profile of single males in 1970.

The final parameter is the mean of the marital bliss in 2000 $\mu_{\gamma,2000}$. For this parameter, we target the profile of single males in 2000. This forces the model to match the profile of marriages in both steady states. At first glance, it seems unnecessary to do this: why not simply see how much model forces can account for? However, in order to recover the labor market parameters through indirect inference, we need to get the right amount of

[16]For data on the relationship between marital status and consumption commitments, see Section 2

single/married people in the model in both time periods. We need to include this parameter to ensure that the model explains all the data. It can be considered a residual.

The estimated parameter values are reported in Table 4 and the model fit is discussed in Section 5.4.

Table 4: Estimated Parameters

Parameter	Description	Value	SE
α	Utility Weight on Home Goods	.99	N/A
ζ	CRRA Parameter on Home Goods	3.79	N/A
κ_h	Cost of wife leaving the labor force	.58	N/A
κ_w	Cost of wife entering the labor force	1.17	N/A
$\mu_{\gamma,1970}$	Mean marital bliss shock , 1970	144.85	N/A
σ_γ	St. Deviation of marital bliss shock	47.37	N/A
$\mu_{\gamma,2000}$	Mean marital bliss shock , 2000	45.00	N/A
p	Price of Durables, 2000	4.38	N/A
c_k	Marital Consumption Commitments	.44	N/A
ϕ_{1970}	Gender Wage Gap, 1970	.63	N/A
ϕ_{2000}	Gender Wage Gap, 2000	.79	N/A

5.4 Model Fit

Let us first discuss how well the model is able to reproduce the data statistics targeted in the estimation procedure. We target the age profile of percentage of singles in 1970 and 2000. We target female labor force participation rate in both periods, female movements in and out of the labor market in 1970, percentage of household's expenditures spent on home inputs in 2000, and consumption commitments as a fraction of couple's expenditures in 2000. We start with marriage. As Figure 4 shows, the model is able to hit the profile of singles both in 1970 and 2000 quite well.

Table 5 compares the model statistics with the other data targets. Overall, the model does a good job matching these additional moments.

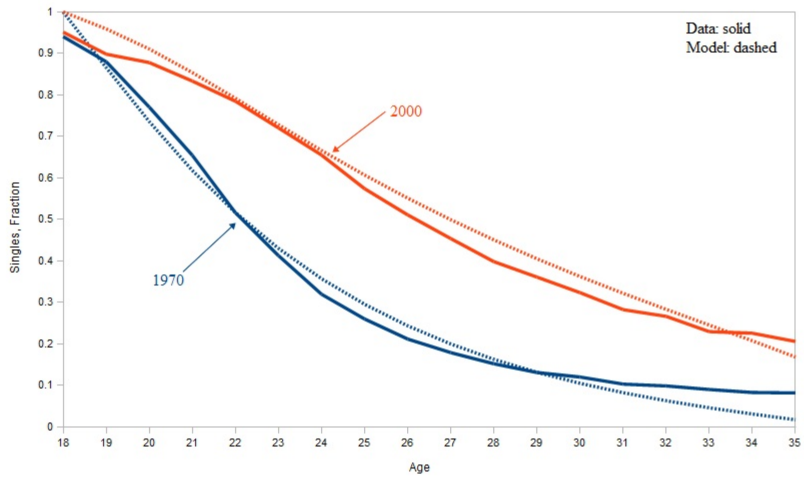

Figure 4: Model fit: Single Males Profiles

6 Results

In this section we decompose the effects of various mechanisms on the delay in marriage. To do this, we change all parameters to the 2000 values, *except for the parameter of interest.* For example, when we study labor market 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 had labor market volatility not increased?" We look at how much each mechanism affects the change from the model benchmark in 1970 to the model benchmark in 2000. The results are in Figure 5.

As a way of quantifying the various effects illustrated in Figure 5, we look at how much of the overall decline in marriage at age 25 is explained by each of the mechanisms[17]. Increasing labor market volatility is responsible for 42% of the change. This shows that the effects of consumption commitment and added gains to search due to rising labor market volatility dramatically outweigh the effects of the gains to spousal insurance. Labor market risk is a major determinant of marriage decisions.

Declining prices of home inputs account for 33% of the decline in marriage. This is clearly an important factor as well, as highlighted by Greenwood and Guner (2009). The

[17]These results are not sensitive to the choice of this specific age.

Table 5: Model Fit

Statistic	Model	Data
Female LFP - 1970	.43	.42
Female LFP - 2000	.78	.72
Observed Gender Gap - 1970	.67	.67
Observed Gender Gap - 2000	.75	.75
% of wives moving into LF in 1970	.05	.04
% of wives moving out of LF in 1970	.07	.07
Fraction of household expenditures on home inputs in 2000	.24	.23
Consumption commitments: % of couple's expenditures, 2000	.29	.30

narrowing of the gender wage gap explains only 7% of the difference between the benchmark profiles. The weak effect of the narrowing gender wage gap is the result of the two opposing forces mentioned above. 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. It turns out, quantitatively, that these effects mostly cancel out.

To summarize, out of the three modeled mechanisms, rising labor market volatility is the most important cause of the delay of marriage.

6.1 Decomposition of Volatility Shocks

The income process used in this paper consists of two shocks, namely a persistent and a transitory shock. In this subsection, we decompose the overall effect of rising income volatility into the contributions of each particular type of shock.

Figure 6 graphs the relative importance of each type of shock to explain the delay of marriage. The increasing volatility of persistent shocks accounts for virtually all of the overall effect. This is intuitive given that these shocks are harder to insure against either through spousal insurance or by accumulating assets. Notice that the effects of the growing volatility of transitory shocks decrease over the life-cycle. This reason is that, as people accumulate assets, they are more able to self insure against this type of shocks.

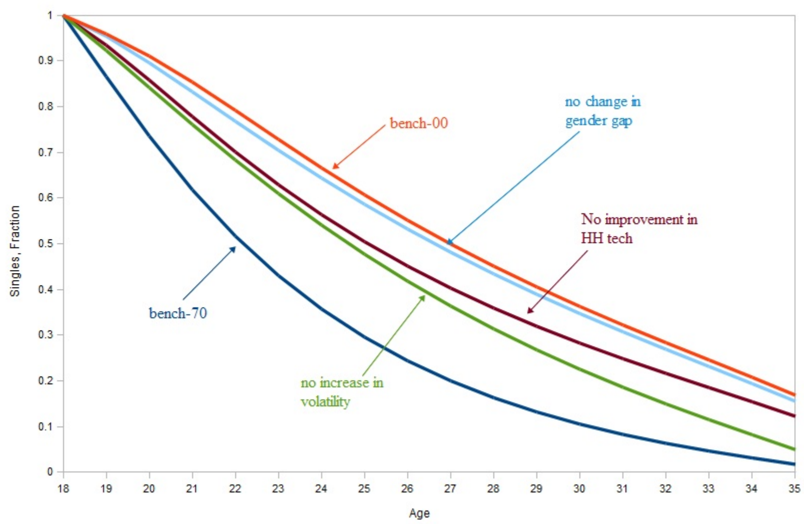

Figure 5: Single Males Profile - Different Channels

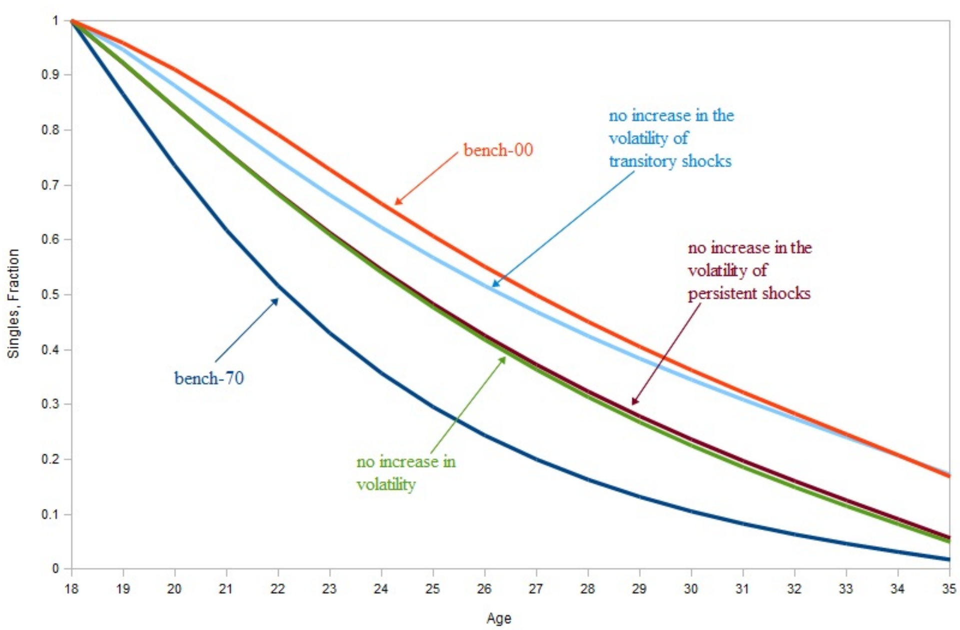

Figure 6: Single Males Profile - Persistent and Transitory Shocks

7 Conclusions

There have been drastic changes in American society over the last 40 years. Young adults have been delaying marriage, in a manner that is frequently described as "failing to launch". We contribute towards answering the most natural question: Why?

We propose a new hypothesis: increasing income volatility has led to a delay in marriage. The idea is simple: if marriage involves consumption commitments, such as kids, then an increase in labor market volatility makes marriage less desirable. 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 structurally estimate a model with increasing income volatility, a narrowing gender wage gap, and decreasing prices of home inputs. We find that rising income volatility explains about 40% of the decline in marriage. The decrease in the price of home inputs explains about a third of this decline. The narrowing of the gender wage gap has small effects.

Note that the median age at first marriage over the 20th century was U-shaped, as Figure 7 shows. While data limitations hurt our ability to evaluate key statistics for the earlier part of the century, leaving our hypothesis untestable, it is possible that the labor market was volatile during the beginning of the century, causing delayed marriage. With a more complete data set, perhaps increasing labor market volatility could explain the full time series.

The framework developed here could also be used to address different questions. For example, in the presence of consumption commitments, individuals may sort themselves into jobs or occupations with more or less risk depending on their marital status. Another possibility is the analysis of the impact on family formation of different government policies that affect the labor market. We leave these possibilities for future research.

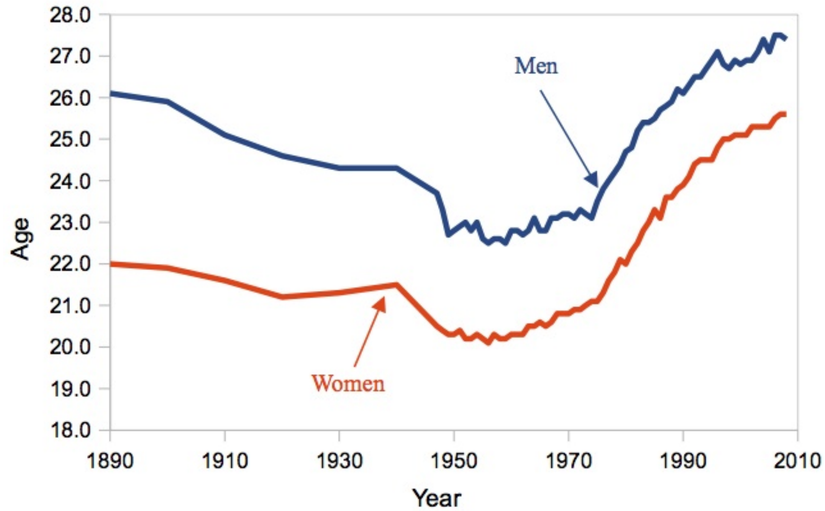

Figure 7: Median age at first Marriage. Sources: US Census Bureau

References

[1] Attanasio, Orazio, Hamish Low and Virginia Sanchez-Marcos. 2008. “Explaining Changes in Female Labour Supply in a Life-cycle Model”, *American Economic Review*, 98 (4), 1517-52.

[2] 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.

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

[4] Gordon, Robert J. 1990. *The Measurement of Durable Goods Prices.* Chicago: University of Chicago Press.

[5] 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.

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

[7] Gould, Eric and Daniele Paserman. 2003. "Waiting for Mr. Right: Rising Inequality and Declining Marriage Rates", Journal of Urban Economics, 53, 257-28.

[8] 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.

[9] Greenwood, Jeremy, Ananth Seshadri and Mehmet Yorukoglu. 2005. "Engines of Liberation", *Review of Economic Studies*, 72 (1), 109-133.

[10] 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.

[11] Heathcote, Jonathan, Kjetil Storesletten and Gianluca Violante. 2010. "The Macroeconomic Implications of Rising Wage Inequality in the United States", *Journal of Political Economy*, 118 (4), 681-722.

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

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

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

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

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

[17] Karahan, Fatih and Suryun Rhee. 2011. "Housing and the Labor Market: the Role of Migration on Aggregate Unemployment", University of Pennsylvania, mimeo.

[18] Katz, Lawrence F. and David H. Autor. 1999. "Changes in the wage structure and earnings inequality" In *Handbook of Labor Economics*, ed. O. Ashenfelter and D. Card, edition 1, volume 3, chapter 26, 1463-1555.

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

[20] McGrattan, Ellen, Richard Rogerson and Randall Wright. 1997. "An Equilibrium Model of the Business Cycle", *International Economic Review*, 38, 267-90.

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

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

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

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

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

A Other Data on the Delay of Marriage

In this section, we present data on the delay of marriage for different groups of the population. Figures 8 and 9 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's clear that marriage has been delayed by individuals of both education groups.

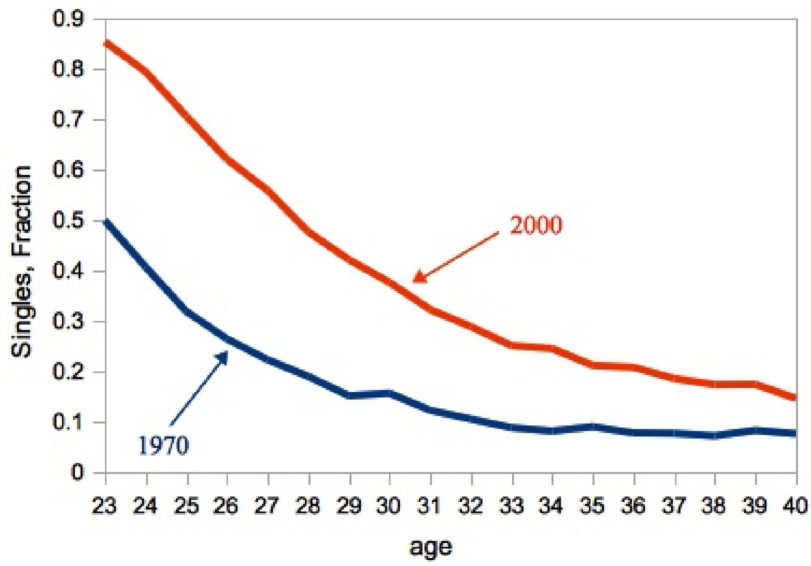

Figure 8: Percentage of College Educated White Males Never Married, by Age. Source: US Census.

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 10 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. 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 10.

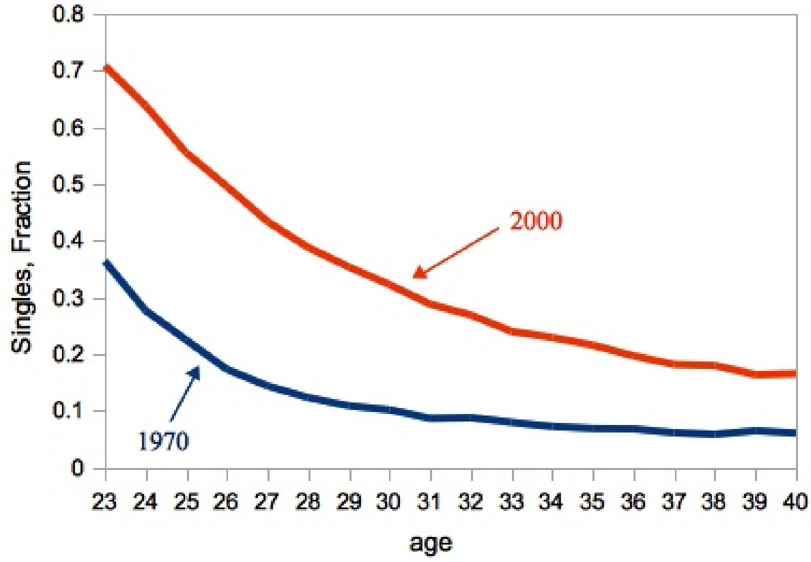

Figure 9: Percentage of Non-College Educated White Males Never Married, by Age. Source: US Census.

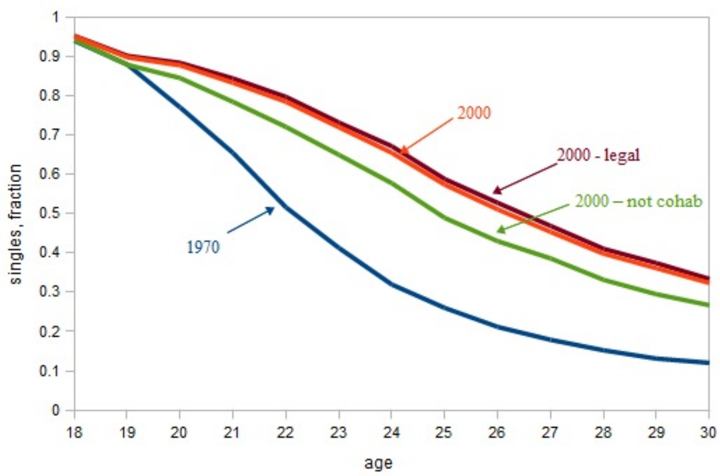

Figure 10: Living Arrangements of Young Adults. Source: US Census

B Estimation of Income Processes

We use data from the PSID for all waves between 1968 and 1997. As described in the text, we separately estimate the processes described in Subsection 3.1 for married and single individuals. We use data for male respondents that satisfy the following criteria for at least three not necessarily consecutive years: (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, (iv) 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 Subsection 3.1 in order to include individual fixed effects (that 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 the ones reported in Table 3. The standard errors are computed using a bootstrap procedure. For a formal proof of identification of the parameters, see Karahan and Ozkan (2010).