Summer 2017

Three essays on income dynamics and demographic economics

Lev Lvovskiy
University of Iowa

THREE ESSAYS ON INCOME DYNAMICS AND DEMOGRAPHIC ECONOMICS

by

Lev Lvovskiy

A thesis submitted in partial fulfillment of the
requirements for the Doctor of Philosophy
degree in Economics
in the Graduate College of
The University of Iowa

August 2017

Thesis Supervisor: Professor Alice Schoonbroodt

Graduate College
The University of Iowa
Iowa City, Iowa

CERTIFICATE OF APPROVAL

PH.D. THESIS

This is to certify that the Ph.D. thesis of

Lev Lvovskiy

has been approved by the Examining Committee for the thesis requirement for the Doctor of Philosophy degree in Economics at the August 2017 graduation.

Thesis Committee: ____________________________

Alice Schoonbroodt, Thesis Supervisor

Martin Gervais

Nicolas Ziebarth

Yu-Fai (Michael) Choi

David Frisvold

ACKNOWLEDGEMENTS

I would like to thank my advisor Alice Schoonbroodt as well as the other members of my committee: Martin Gervais, Nicolas Ziebarth, Michael (Yu-Fai) Choi and David Frisvold for their valuable feedback and comments on my research.

I also owe a special thanks to my friends Lawrence Warren, Michael Andrews, Chander Kochar and Yuanyang Liu for their support and guidance over the years and for the exceptionally interesting and gratifying discussions of economics. In addition I want to thank many other fellow economics graduate students without whom this dissertation would not be possible.

I am very grateful to my parents Elena and Yakov for providing me with the opportunity to pursue the Ph.D. I am heavily indebted to Dr. Alissa Klots for the initial encouragement and continuous support throughout the course of my Ph.D. Finally, I would like to thank Renea Jay, Adriana Gama, Dr. Alexandra Nica and Mikhail Kulak who made my experience writing this dissertation more enjoyable.

ABSTRACT

This dissertation consists of three chapters. The first chapter addresses the roles of changes in assisted reproductive technologies, returns to female experience and abortion rates in explaining the historical trend of child adoption. The second chapter assesses the effects of increased income inequality and decreased income mobility on timing of births and marriages and on the single motherhood rates. The third chapter establishes the importance of accounting for marital state in the models of indirect income uncertainty inference.

Chapter 1 aims to explain the μ-shaped historical trend of child adoption in the US by emphasizing the role of the changes in the demand side of the market for child adoption. I argue that changes on the demand side such as increasing returns to female human capital and innovations in Assisted Reproductive Technology (ART) have played a major role in shaping the historical adoption trend along with the changes in the supply side, namely, increase in the abortion rates. I present a life-cycle model, in which an agent makes a fertility-timing decision based on the returns to her human capital and age-specific probability of conception. Under the assumption that adoption is an alternative to childbearing, i.e. an agent chooses to adopt after she fails to conceive, the presented model uses historical trends of returns to human capital and success rate of ART to explain changes in adoption trends. According to the model, increasing returns to female human capital were responsible for the delay in childbearing and therefore the increase in the demand for adoption until

the 1970s. After 1970, the legalization of abortion decreased the supply of orphans, while innovations in ART decreased the demand by allowing women to have biological children at later ages. Around 1980, the effect of increasing returns to human capital overturned the one of advances in ART, which resulted in a slow recovery of the adoption trend.

Chapter 2 studies the dramatic transformation that the typical American family has undergone since the 1950s. Marriage and fertility have been delayed, while single-motherhood rates have increased. The link between these facts emanates from the greater delay in marriage than that in first births. As "the Gap" between the age at first birth and the age at first marriage becomes negative for some women, out-of-wedlock first births increase. In my analyses I focus on the increase in income inequality and the decrease in income mobility — observed across two National Longitudinal Survey of Youth (NLSY) cohorts of women — to account for the above facts using an equilibrium two-sided search framework in which agents make marriage and fertility choices over the life-cycle. Marriage is a commitment device for consumption-sharing, providing spouses with partial insurance against idiosyncratic earnings risk. Agents derive utility from children, but children also involve a risky commitment to future monetary and time costs. According to my model, two observed trends in the income process produce these changes in the respective timings of marriage and fertility. First, the increase in income inequality produces incentives to delay marriage. Since single women tend to face higher income risk than do married women, all else being equal, a decline in marriages when young implies delayed births, which are perceived to be risky. Second, the decrease in income mobility also delays marriage as the insurance value of marriage decreases but accelerates fertility because it becomes less risky to have a child. The model qualitatively matches the observed changes in family formation and quantitatively accounts for a significant portion of the observed changes in marriage and fertility timing between the two NLSY cohorts.

In Chapter 3 I aim to add to the indirect income uncertainty inference literature. The currently existing models used to infer earnings uncertainty from consumption decisions of individuals either use married couples as a unit of analysis or treat married individuals as singles. Income pooling and less than perfect correlation of earnings in marital unions provide spouses with marital income insurance. Not accounting for the marital insurance biases the uncertainty estimation results. In this chapter I demonstrate some properties of the marital insurance bias in a stylized analytical model. In order to access the potential magnitude of the marital bias I build a structural model which accounts for marital insurance. I then compare the estimation results of the model which accounts for marriage with the results of one that does not after using them on the simulated data set. In addition I introduce a non-parametric income process in the structural model used for the indirect uncertainty inference. The main advantage of the resulting model is that, unlike the typical models in this area, it can be used on short-term panel data.

PUBLIC ABSTRACT

This dissertation consists of three chapters. The first chapter addresses the roles of changes in assisted reproductive technologies, returns to female experience and abortion rates in explaining the historical trend of child adoption. The second chapter assesses the effects of increased income inequality and decreased income mobility on timing of births and marriages and on the single motherhood rates. The third chapter establishes the importance of accounting for marital state in the models of indirect income uncertainty inference.

Chapter 1 aims to explain the historical trend of child adoption in the US by emphasizing the role of the changes in the demand side of the market for child adoption. I argue that changes on the demand side such as increasing returns to female human capital and innovations in Assisted Reproductive Technology (ART) have played a major role in shaping the historical adoption trend along with the changes in the supply side, namely, increase in the abortion rates. I present a model, in which an agent makes a fertility-timing decision based on the returns to her human capital and age-specific probability of conception. Under the assumption that adoption is an alternative to childbearing, i.e. an agent chooses to adopt after she fails to conceive, the presented model uses historical trends of returns to human capital and success rate of ART to explain changes in adoption trends. According to the model, increasing returns to female human capital were responsible for the delay in childbearing and therefore the increase in the demand for adoption until the 1970s. After 1970, the

legalization of abortion decreased the supply of orphans, while innovations in ART decreased the demand by allowing women to have biological children at later ages. Around 1980, the effect of increasing returns to human capital overturned the one of advances in ART, which resulted in a slow recovery of the adoption trend.

Chapter 2 studies the dramatic transformation that the typical American family has undergone since the 1950s. Marriage and fertility have been delayed, while single-motherhood rates have increased. The link between these facts emanates from the greater delay in marriage than that in first births. As “the Gap” between the age at first birth and the age at first marriage becomes negative for some women, out-of-wedlock first births increase. In my analyses I focus on the increase in income inequality and the decrease in income mobility — observed across two National Longitudinal Survey of Youth (NLSY) cohorts of women — to account for the above facts using an equilibrium two-sided search framework in which agents make marriage and fertility choices over the life-cycle. Marriage is a commitment device for consumption-sharing, providing spouses with partial insurance against idiosyncratic earnings risk. Agents derive utility from children, but children also involve a risky commitment to future monetary and time costs. According to my model, two observed trends in the income process produce these changes in the respective timings of marriage and fertility. First, the increase in income inequality produces incentives to delay marriage. Since single women tend to face higher income risk than do married women, all else being equal, a decline in marriages when young implies delayed births, which are perceived to be risky. Second, the decrease in income mobility also delays marriage as the insurance value of marriage decreases but accelerates fertility because it becomes less risky to have a child. The model qualitatively matches the observed changes in family formation and quantitatively accounts for a significant portion of the observed changes in marriage and fertility timing between the two NLSY cohorts.

In Chapter 3 I aim to add to the indirect income uncertainty inference literature. The currently existing models used to infer earnings uncertainty from consumption decisions of individuals either use married couples as a unit of analysis or treat married individuals as singles. Income pooling and less than perfect correlation of earnings in marital unions provide spouses with marital income insurance. Not accounting for the marital insurance biases the uncertainty estimation results. In this paper I demonstrate some properties of the marital insurance bias in a stylized analytical model. In order to access the potential magnitude of the marital bias I build a structural model which accounts for marital insurance. I then compare the estimation results of the model which accounts for marriage with the results of one that does not after using them on the simulated data set. In addition I introduce a non-parametric income process in the structural model used for the indirect uncertainty inference. The main advantage of the resulting model is that, unlike the typical models in this area, it can be used on short-term panel data.

TABLE OF CONTENTS

LIST OF TABLES

Table

LIST OF FIGURES

Figure

CHAPTER 1
ART VS ABORTION: EXPLAINING TRENDS IN CHILD ADOPTION

1.1 Introduction

The adoption of a child is an important socio-economic phenomenon which has a significant impact on the welfare of both adopted children and parents. Although the share of adoptions is small relative to the overall U.S. birthrate (about 3-3.5% in recent years[1]), since most adoptive parents are part of the top educational and income layers of society (the part of the population having the least number of children per capita), adoption constitutes up to 25% of live births of this social group.[2]

Despite the socioeconomic importance of adoption, few economists have studied it. The main areas of adoption research are the market for adoptions and its policy implications,[3] and using data on adoption to study nature versus nurture contributions.[4] The most notable work in collecting and analyzing adoption data was done by Moriguchi (2012).

This study aims to advance the understanding of trends in child adoption by studying the supply and demand sides of this market with a special emphasis on the latter. To achieve this goal, I first develop an analytical model where the

[1]Moriguchi (2012) used CWIG 2011a data to estimate that in 2000-2008 adoption rate was from 32 to 36 per 1000 births

[2]Rough estimation can be made by combining numbers from NSAP 2007 and Caucutt, Guner and Knowles (2002).

[3]Landes and Posner (1978), Medoff(1993), Hansen and Hansen (2006)

[4]Case et al. (2000), Sacerdote(2002), Plug and Vijverberg (2003)

agent chooses the timing of a birth and quantity of children. The agent faces the following trade-off with regard to the birth timing: the benefit of delaying fertility is accumulating human capital through on-the-job training to receive higher wages. At the same time, the disadvantage of fertility delay is the decreasing probability of conception. The failure to conceive a child in this model leads to adopting one. In the second step of this project, I simulate historical trends of adoption using the solution of the agent's problem at different returns to human capital and success rate of Assisted Reproductive Technologies (ART).

The main finding of this paper is that the major changes in patterns of adoption rate in the U.S. which took place from 1950-2010 can be explained by the demand side factors such as the increase in returns to human capital, combined with a breakthrough in ART which took place around 1970, rather than the supply side factors, such as the legalization of abortion.

1.2 Stylized Historical Facts

Nowadays, the U.S. government has a variety of special programs related to child adoption. Celebrities adopt children from all over the world. This, along with many other conditions, makes adoption not only more affordable and socially acceptable, but even trendy. Nevertheless, recent absolute and relative adoption numbers have been significantly below their historical peak in 1970.[5] Figure 1 represents the adoption rate per thousand births based on data found in Moriguchi (2012). To de-

[5] See Moriguchi (2012)

scribe adoption trends, I specify two periods: prior to 1970 and after. During the first period, the adoption rate was growing exponentially, whereas in the second period the trend is parabolic. The main goal of this paper is to provide a possible explanation for those changes between the periods.

Figure 1.1: Child Adoption Rate per 1000 Births

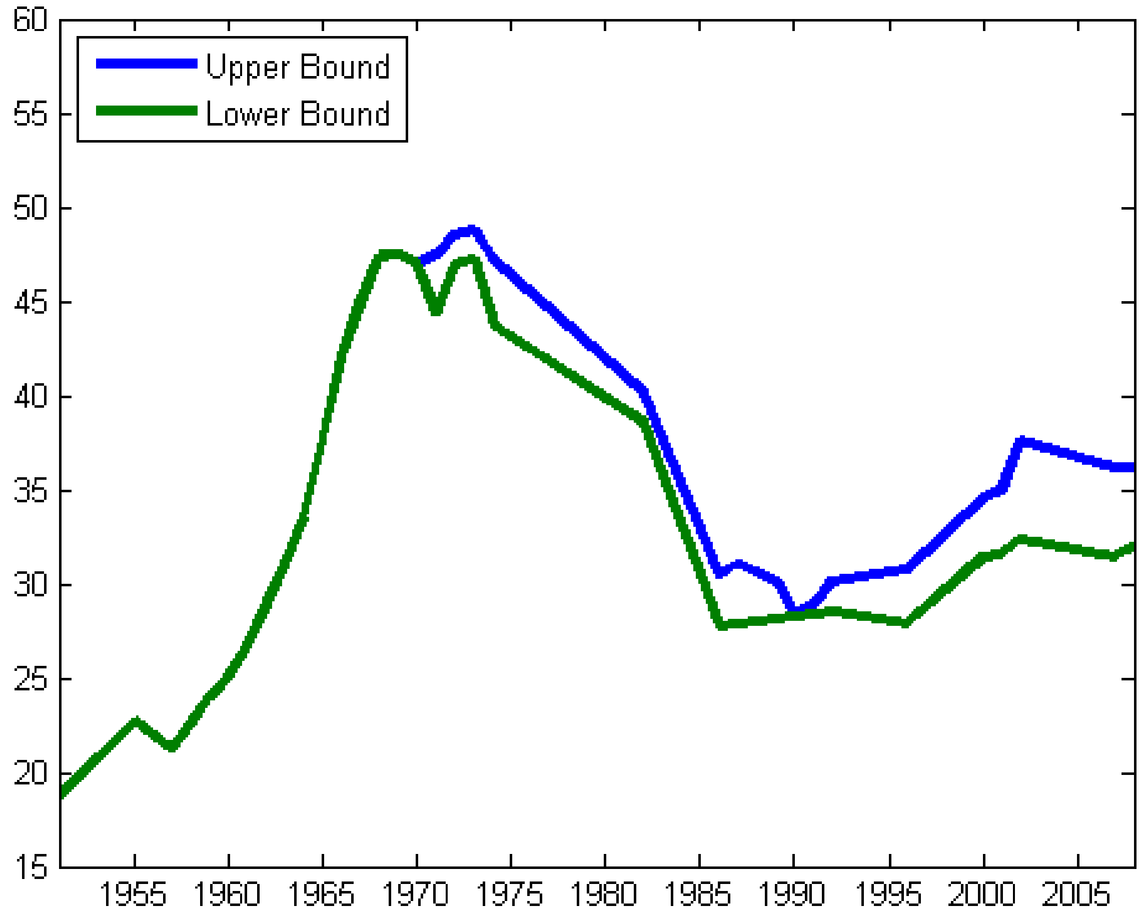

Source: Moriguchi (2012)

The exponential growth of the adoption rate observed in 1950-1970 can be explained by increasing returns to female work experience. A number of studies show that both female and male returns to work experience have significantly increased over the last 70 years, with females experiencing a steeper increase. To the best of my knowledge, unfortunately, research on the earlier period (1950-1970) produces estimations only for male returns to experience.[6] More recent papers studying female returns to experience found that during the last 40 years returns for women have been growing faster than for men.[7] Thus, it seems plausible that female returns to experience had been growing before 1970 as well, but were not an object of economists' interest. In addition to this intensive margin, female participation in the labor force was permanently growing in the post-WWII period constituting an extensive margin.

But how can we explain the behavior of the adoption rate in the second period of the data range? The most popular explanation of this phenomenon originates from the supply side. Moriguchi (2012) argues that the primary cause of the fall of the adoption rate is the rise of abortion rate. Indeed, the shape of the abortion rate in the U.S. during the second half of the Twentieth Century may support this explanation. But people respond to incentives: if women who previously were delaying fertility due to increasing earnings and the possibility to substitute childbearing with adoption started to lose this option, it should have motivated them to have children earlier, when the probability of conceiving is higher. Rather, as in Figure 2, we see

[6] Jacob Mincer's (1974) paper is considered a cornerstone work in this area, another modern revisitation of which was made by Heckman, Lochner and Todd in 2003.

[7] Many authors studied this question using PSID data. O'Neill and Polachek (1993) found that during 1970-1980 female returns to experience increased more than male, Blau and Kahn (1997) got similar results and also found that that in 1989 returns to women experience where 25% higher than for men. In a more recent paper, Olivetti (2006), after correction for selection bias, found that during 1970-1990 period marginal returns to experience for women grew by 25%

that the median age of women at first birth increased faster after 1970.

Figure 1.2: Mean Age at Birth

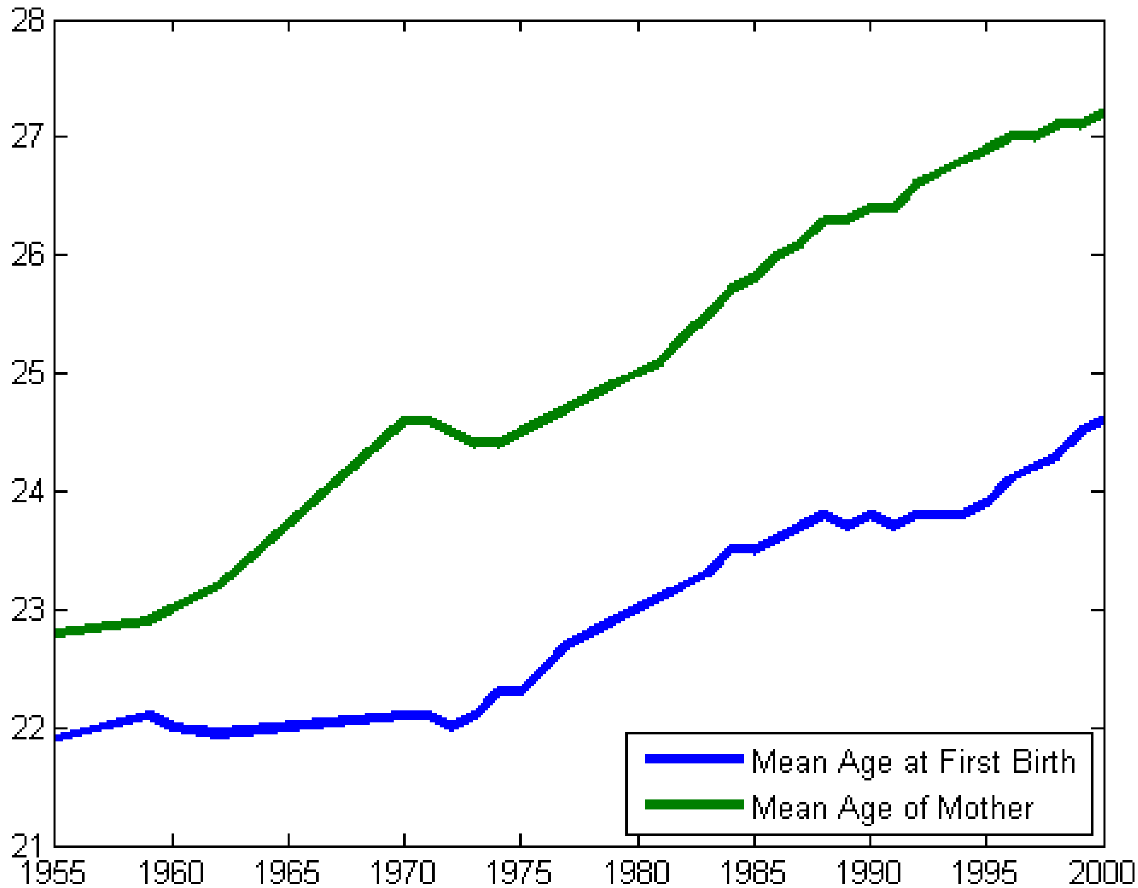

Source: National Vital Statistics Reports, Vol. 55, No. 1, 09/29/2006

I argue that both the rapid decrease of the adoption rate and faster increase of the median age at first birth in the early 1970's were, at least partially, due to a breakthrough in the methods of ART which took place several years before 1970 and became widespread in the early 1970's. The first commercial sperm bank in the U.S. opened in New York in 1970, with several others opening around the country in the following years.[8] In 1973, the Commissioners on Uniform State Laws approved the Uniform Parentage Act, standardizing legal issues of artificially inseminated births. This kind of ART, while not curing female infertility, increases the probability of conception and mitigates all problems related to male infecundity, and many other methods including IVF were developed in the following years. Moriguchi (2012) uses SART and CDC data[9] to show that the ART success rate increased from 6% in 1985 to over 30% in 2009; these numbers are conservative since artificial insemination is not included.

In sum, the general idea of this project is to combine the facts about returns to experience and change in ART described above to explain the big change in the adoption pattern which happened in 1970-1973.

1.3 The Model

I use a life-cycle model similar to those in Razin (1980), Happel et al. (1984), Cigno and Ermish (1989), Blackburn, Bloom and Neumark (1992) and others. In these models, the fertility timing choice is related to on-the-job investment in human capital (or "learning-by-doing") and returns to accumulated human capital (returns to experience). In particular, I use a modification of the age-specific human capital accumulation function from Olivetti (2006). A distinctive feature of this specification is that the rate of accumulation of human capital is higher for young agents and lower

[8]Barney (2005)

[9]SART stands for Society for Assisted Reproductive Technology, and CDC for Center for Disease Control and Prevention

for older ones. The key distinction of my model from previous research is the use of an age-related probability of conception based on one estimated in Van Noord-Zaadstra et al. (1991) instead of a constant fertility window bound. This approach allows me to introduce an indirect choice of an agent between adoption and childbearing. In other words, the agent is not making an explicit decision whether to adopt or have a biological child, but instead influences the probabilities of these outcomes by choosing the time of a conception attempt. A detailed description of the model follows.

1.3.1 Model Setup

The agent in my model is a woman who lives from time 0 — the beginning of her career — to time R — an exogenous retirement date. The agent maximizes her expected lifetime value of consumption $V(\cdot)$ and utility from having children $U(\cdot)$ by making two ex-ante decisions. She chooses T — the time when she will make an attempt to conceive and/or adopt and $\overline{K}$ — the number of attempts that she is going to make. All biological (k_b) and adopted children (k_a) are assumed to appear simultaneously in the agent's life. Note that $\overline{K}$ is the maximum number of children $K = k_b + k_a$ that the agent will have if all of her conception and/or adoption attempts are successful.

Given the assumptions, agent solves:

$$\max_{T,\overline{K}} \quad \mathbb{E}_K V(T,K) + \mathbb{E}_{k_b,k_a} U(k_b, k_a, T) \tag{1.1}$$

1.3.1.1 Family Formation

When the agent chooses to make an attempt to conceive a child at period T, she succeeds with probability $\pi_b(T)$ (probability of being fertile at a given age, $\frac{\partial \pi_b(T)}{\partial (T)} < 0$). If she fails to conceive, the agent adopts a child with probability π_a or finishes the current attempt child-free otherwise. If an agent appeared to be infertile at the previous attempt, during the next one she will only try to adopt a child. For example if the maximum number of attempts $\overline{K} = 1$, probabilities that she ends up with a biological, adopted child or child free are $p_{1,0}(1) = \pi_b(T)$, $p_{0,1}(1) = (1 - \pi_b(T))\pi_a$ and $p_{0,0}(1) = (1 - \pi_b(T))(1 - \pi_a)$ respectively. Probabilities that the agent will have one or zero children in this case are $P_1(1) = p_{1,0}(1) + p_{0,1}(1)$ and $P_0(1) = p_{0,0}$. In such a way, by choosing $T, \overline{K}$, the agent implicitly chooses probabilities of having a certain number of children $P_K(T, \overline{K})$ and distribution over composition of adopted and biological children $p_{k_b,k_a}(T, \overline{K})$. Then the expected utility agent derives from parenting is:

$$\mathbb{E}_{k_b,k_a} U(k_b, k_a, T) = \sum_{\{k_b,k_a\}:\ k_b+k_a \leq \overline{K}} p_{k_b,k_a}(T, \overline{K}) U(k_b, k_a, T) \tag{1.2}$$

I allow the utility from a child to be a function of length of the time when a child is present in the agent's life and the type of the child. In particular, I assume that

$$U(k_b, k_a, T) = \lambda u(k_b, R - T) + u(k_a, R - T) \tag{1.3}$$

where $\lambda > 1$. The assumption that the agent draws higher utility from her biological children produces incentives to attempt conception earlier in life even when the probability of adoption in case of failure is $\pi_a = 1$.

Given the probabilities of the various family formation outcomes, adoption rate (AR) can be computed as a ratio of the expected number of adopted children to the expected number of all children born to a representative agent. For example if the optimal number of attempts to become a mother is $\overline{K} = 1$, the adoption rate is:

$$AR(T,1) = \frac{p_{0,1}(1)}{p_{0,1}(1) + p_{1,0}(1)} = \frac{(1-\pi_b)\pi_a}{(1-\pi_b)\pi_a + \pi_b} \tag{1.4}$$

In sum, the cost of fertility delay (increase in T) originates from the decreasing probability of having biological children, the decreasing probability of having children and the decrease in the utility the agent draws from being a parent as the agent enjoys being a parent for less time.[10]

1.3.1.2 Life-Time Earnings

In each period, the agent receives an exogenous market rental rate (normalized to 1) per unit of her current level of human capital (experience). Following Olivetti (2006), the agent accumulates human capital through a learning-by-doing process (i.e., current stock of human capital depends on its past value and the number of hours worked in the previous period). The agent is endowed with H hours per period.

[10]Note that the assumption that utility derived from parenting decreases with the age at birth is just a shortcut to the usual notion of receiving that utility during lower number of periods $u(k, R-T) \sim \sum_T^R \tilde{u}(k)$.

While the agent is childless, she devotes all her time to work since she does not value leisure. Once she becomes a mother, each period she spends an exogenously given fraction of her time τ per child. Since the choice of the time allocation is degenerate in my model, I can rewrite human capital accumulation function in a non-recursive form $\Theta \equiv \Theta(t, 0)$ before the agent attempts to become a mother and $\Theta \equiv \Theta(t, T, K)$ once she starts to take care of K children.

The only source of uncertainty about the agent's future earnings is the realization of number of children K so the expected value of a life-time consumption is:

$$\mathbb{E}_K V(T, K) = \sum_{t=1}^{T} \beta^t v\left(\Theta(t, 0)\right) + \sum_{K=0}^{\overline{K}} P_K(T, \overline{K}) \left[\sum_{t=T+1}^{R} \beta^t v\left(\Theta(t, T, K)\right)\right] \tag{1.5}$$

where $v(\cdot)$ — per period utility of consumption.

There are two benefits of delaying motherhood in terms of the value of life-time consumption. First, delaying motherhood results in fewer periods of not working full time. Another benefit of later motherhood arises from the decreasing accumulation rate of human capital. The intuition behind this property of the income process borrowed from Olivetti (2006), is that the amount of human capital accumulated from an additional hour of work declines as an individual ages. In other words, the human capital accumulation rate à la Olivetti is age-specific as opposed to the commonly used tenure-specific human capital accumulation à la Mincer. Figure 1.3 demonstrates the difference between the two processes for an individual who has been absent from the labor force for ten years. Analytically, the difference between the life

Figure 1.3: Experience Accumulation

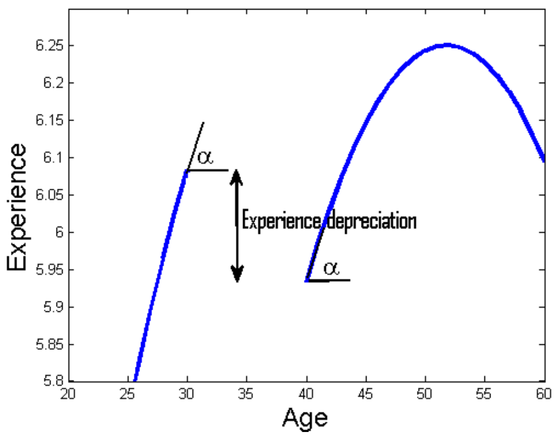

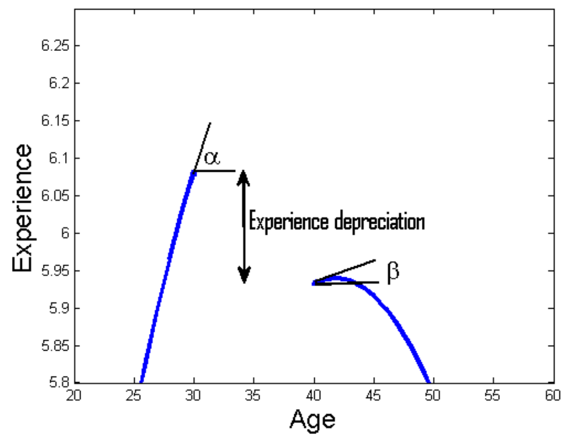

time earnings depicted on the left vs right panels of Figure 1.3 is:

$$\begin{aligned} & \int_0^T \Theta(t)dt + \int_{T+10}^R -\delta + \Theta(t-10)dt \\ VS \quad & \int_0^T \Theta(t)dt + \int_{T+10}^R -\delta + \Theta(t) + \Theta(T) - \Theta(T+10)dt \end{aligned}$$

1.4 Simulation of the Historical Trend

The main purpose of this paper is to explain the changes in the adoption rate through the changes in the human capital accumulation function (Θ), innovations in ART (π_b) and changes in the abortion rate. The latter is represented by changes in the probability of adoption π_a. In order to assess the plausibility of such explanation and the quantitative importance of those changes, I use my model to simulate the 1950-2010 U.S. adoption trend. Whenever possible, I use the parameter values estimated in the other papers; for periods when numbers are unavailable, I use extrapolations of known trends.

1.4.1 Parameters

Since I model the behavior of individuals in a life-cycle perspective, the simulation includes predictions of the adoption rates of successive generations of representative agents. Each generation is characterized by the specific parameters of the Θ, π_a and $\pi_b(g)$ functions.

1.4.1.1 Human Capital Accumulation Θ

As mentioned in Section 2, to the best of my knowledge, there are no published estimates for female returns to labor experience prior to 1970. To obtain these numbers, I assume that trends in female experience returns observed after 1970 were present before as well. A possible source of bias from this assumption is that if the increase in returns to female experience had been growing slower before 1970, my simulation will produce a higher increase in the adoption rate during the pre-1970 period. Olivetti (2006), after fixing selection bias problems, used PSID data to estimate an age-specific human capital accumulation function where the rate of accumulation depends on both the age of the agent and hours worked in the previous period. The degenerate choice of labor hours in my model allows me to interpolate Olivetti's recursive human capital accumulation specification in a non-recursive, fertility timing and number of children –specific functions of a form:

$$\Theta(t,T,N) = e^{\eta_0^{T,N} + \eta_1^{T,N} t + \eta_2^{T,N} t^2} \tag{1.6}$$

for the 1970 and 1990 cohorts,[11] where $\delta = 0.2277$—depreciation rate of human capital. Estimations for the 1970s are: $\eta_0 = -0.0499, \eta_1 = 0.0429, \eta_2 = -0.0015$; for 1990s cohort: $\eta_0 = -0.0886, \eta_1 = 0.0948, \eta_2 = -0.0017$. For the simulation, I use linear extrapolations of these coefficients.

1.4.1.2 Probability of Being Fertile π_b

Van Noord-Zaadstra et al. (1991), in their study of women in artificial insemination programs, estimated a critical age probability of successful pregnancy function; i.e. their function is constant until age 30 and starts to decrease exponentially after this critical age. To obtain a differentiable probability function, I re-estimated it with a new functional form:

$$\pi_b(t) = \frac{1}{\alpha_1 + \alpha_2 e^{\alpha_3 t}} \tag{1.7}$$

The parameters of this function are: $\alpha_1 = 1.0887$, $\alpha_2 = 0.0046$, and "natural" (before 1970) value of $\alpha_3 = 0.2821$. There are two ways to think about the post-1970 evolution of α_3; the first is to assume that from that point the success rate of ART started to grow linearly though 2008, where the 32% of ART success is equivalent to $\alpha_3 = 0.1813$. The second possible trend is some breakthrough in 1970 and a gradual increase after that. The latter version is supported by historical evidence. Barney (2005), for example, points out that donor insemination also solved the problem of "social" infertility giving the opportunity of maternity to unmarried women, widows

[11] The 1970 cohort consists of women who started their career (where 20 years old) in 1970

and lesbians. It also provided the same opportunity to couples suffering from male infertility, which, according to various studies,[12] accounts for 40% to 50% of all cases of couple's infertility. Of course, not all couples can agree on such an option (in this case spouses may have different levels of altruism towards a child), and we also do not know the size of "social" infertility. Therefore, I assume the size of the breakthrough to be a 4% increase in the probability of conception. I simulate the latter version of ART success rate evaluation.

1.4.1.3 Availability of Adoption Opportunity π_a

The popular statement that the adoption market is capacity constrained is not fully correct — the number of children in foster care system was never close to zero during the time period of interest. Rather, during some periods, children meeting specific preferences of adoptive parents are not available.[13] In such a case, the reader may think about π_a as some measure reflecting the tightness of the adoption market. Also, for the purposes of my project, I am rather interested in changes in $\hat{\pi}$ than in its absolute level. I construct the values of this parameter as the following:

$$\pi_a = 1 - N_I - N_L + N_{IA} \tag{1.8}$$

where N_I and N_L are illegal and legal abortions per thousand live births respectively, and N_{IA} is the number of intercountry adoptions per thousand live births.

[12] Sharma et al. (1999), Poongothai, Gopenath and Manonayaki (2009) etc.

[13] For example, most adoptive parents look for a healthy white girl without cases of cancer or alcoholic addiction in her biological family history.

I use the number of legal abortions compiled in Johnston (2011). The number of illegal abortions is estimated from number of deaths caused by this procedure in Tietze (1975). Finally, the number of intercountry adoptions is taken from Moriguchi (2012).

1.4.1.4 Extensive Margin

Another factor that has changed significantly over time and is related to the adoption rate is female labor force participation. Toossi (2002) found that female labor force participation experienced a practically linear growth between 1950 and 1990. To be able to account for extensive margin I, therefore, take labor force participation rate in 1950 to be equal to 32.58% with an annual growth of 0.61%.

1.4.1.5 Relative Value of Children λ

I choose values of altruism. Intuitively, λ should be greater than or equal to one, since the majority of people prefer to have biological children rather than adopted. There could be several reasons for such preferences, one measurable difference between these two types of children is their expected health. Various studies show that adopted children have higher health risks; for example, the 2007 National Study of Adoptive Parents (NSAP) shows that while the percentage of all children with special health care needs is 19%, the same number for adopted children is 39% which suggests $\lambda = 2$. However, disparity between the adopted and biological children is likely to be the highest exactly in terms of health, which means that $\lambda = 2$ is rather an upper bound. For the simulation I assume the midpoint $\lambda = 1.5$.

1.4.1.6 Utility Functions and Parameters

Finally, I assume utility functions over consumption and children of the following forms:

$$v(c) = \frac{c^{1-\gamma}}{1-\gamma} \tag{1.9}$$

$$U(k_b, k_a, T) = \mu(R-T)(\lambda k_b + k_a)$$

Parameters γ and $\mu =$ are chosen such that the model matches the adoption rate in 1950.

1.4.2 Simulation

Figure 1.4 demonstrates that the changes in returns to experience, ART and abortion rate used in the developed stylized model are able to match the shape and the general magnitude of the historical trend. The simulated trend allows to precisely track the pre-1970 increase in historic rates. It somewhat overshoots the peak level but further quite closely reflects the decrease. Following a slight overshoot in the 1980s, it levels out, although at a lower magnitude than the data. The kink in the simulated trend right before the 1970 is due to the discrete maximum birth parity choice assumed in the model and is produced by agents switching from $\overline{K} = 3$ to $\overline{K} = 2$ as the time cost of children increases.

Figure 1.4: Simulation of the Adoption Trend

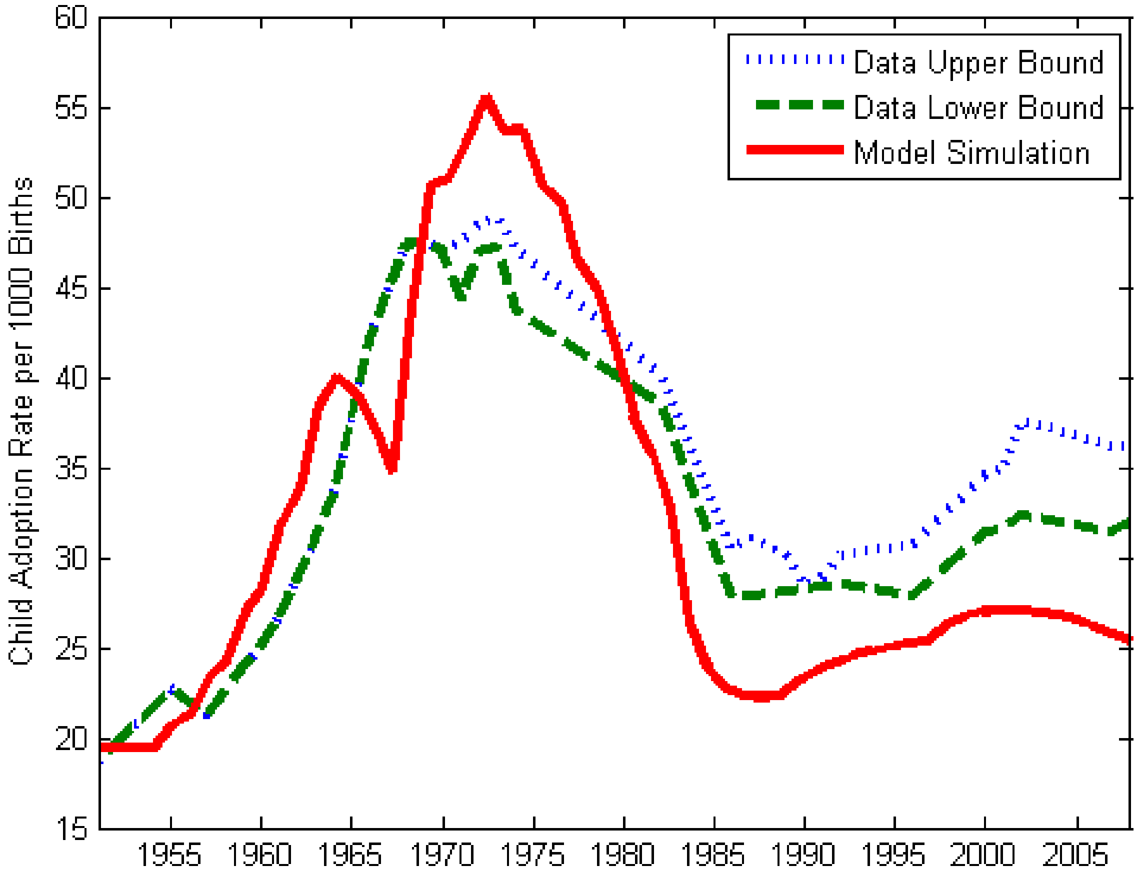

1.4.3 Counter-Factual Experiments

Since all the exogenous changes in the model interact with each other, in order to grasp the idea of their relative importance, I run several counter-factual experiments.

Figure 1.5 shows the model's prediction if the only exogenous change was increasing returns to experience. Since such a change only increases the benefit of fertility delay, the adoption rate continues to grow throughout the entire period and reaches 16% by 2008.

In the second experiment, I combine changes of increasing returns to experience

Figure 1.5: Only Returns to Experience Change

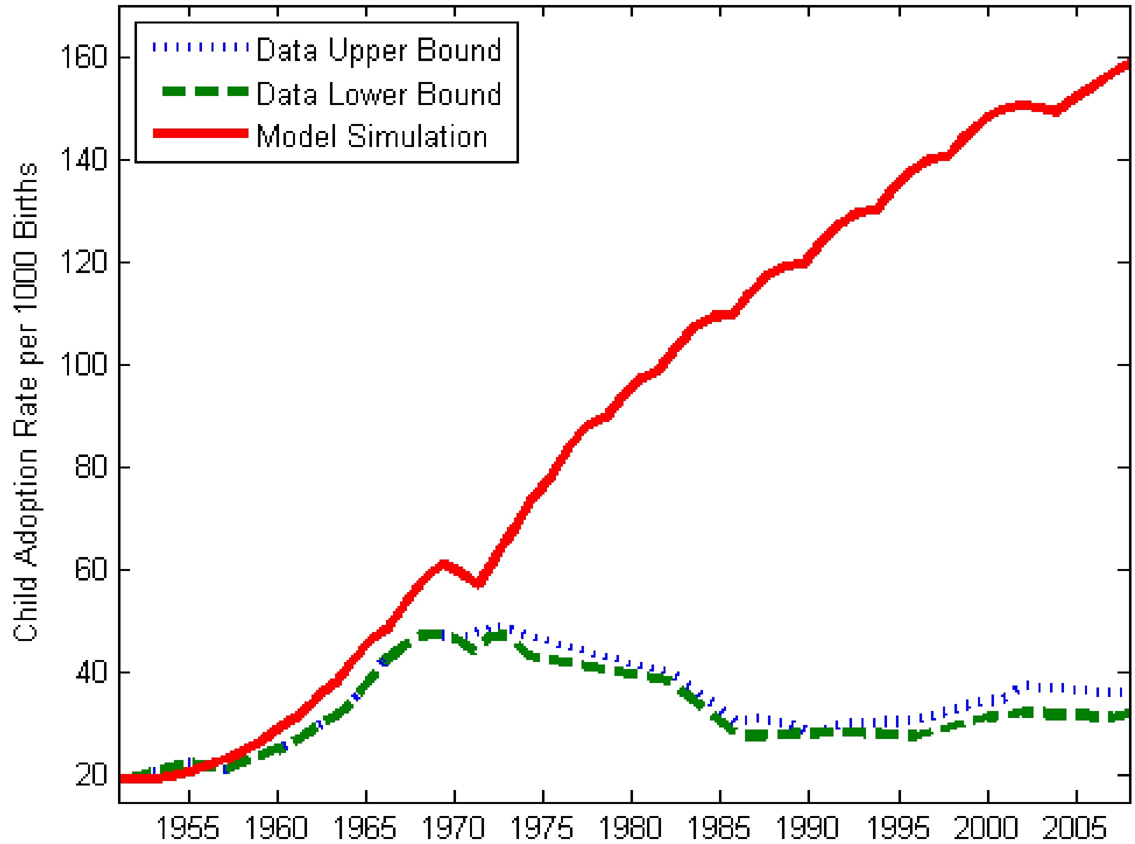

with the effect of the increase in abortion rate (represented by π_a). As one can see in Figure 1.6, when it becomes harder to adopt a child, the adoption rate first declines slightly (1970s) but later the increase in the demand side overturns its effect.

Finally, Figure 1.7 presents the simulation of the model when the only sources of exogenous change are returns to experience and innovations in ART. Out of the three counter-factual experiments, this one most closely resembles the data. The simulation captures the shape of the historical trend while consistently over-predicting it since the supply of children for adoption (probability of adoption) does not decrease in this case. Still, this suggests the major role of the demand side changes in shaping

Figure 1.6: Returns to Experience and Abortion Rate Changes

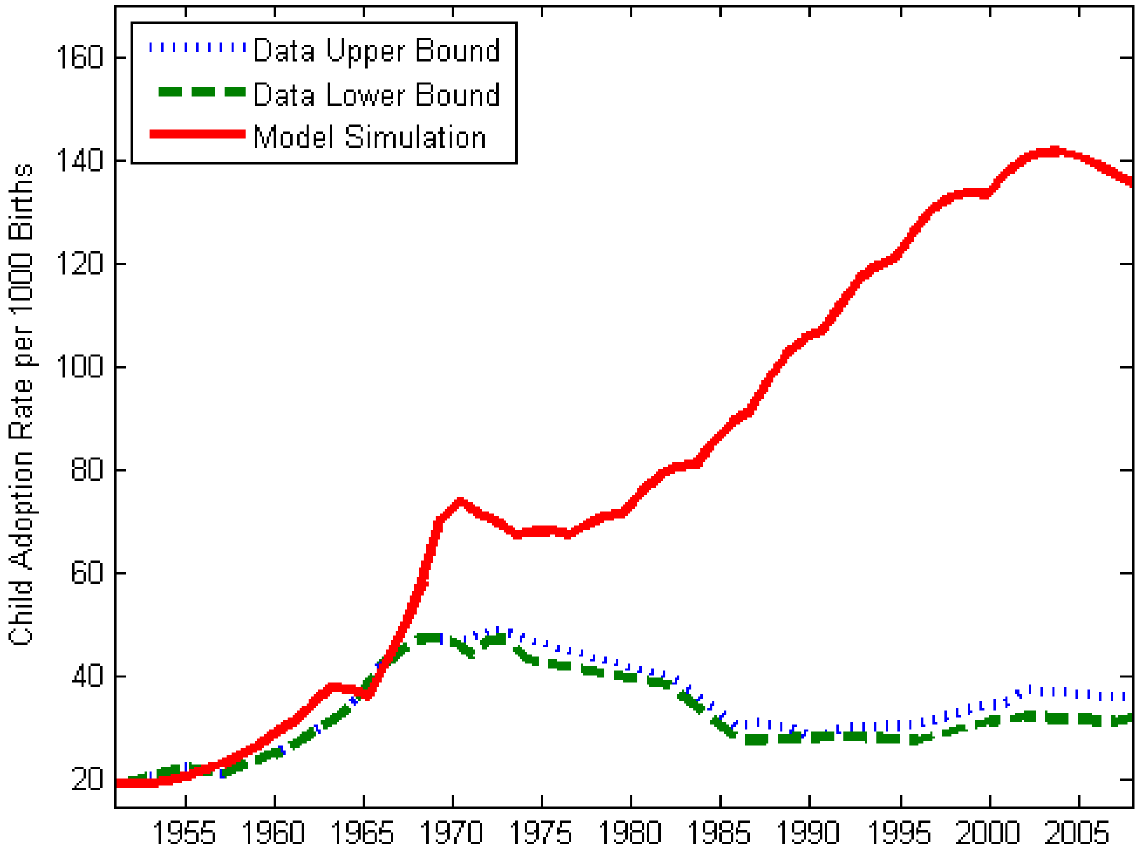

the historical trend.

Figure 1.7: Returns to Experience and ART Changes

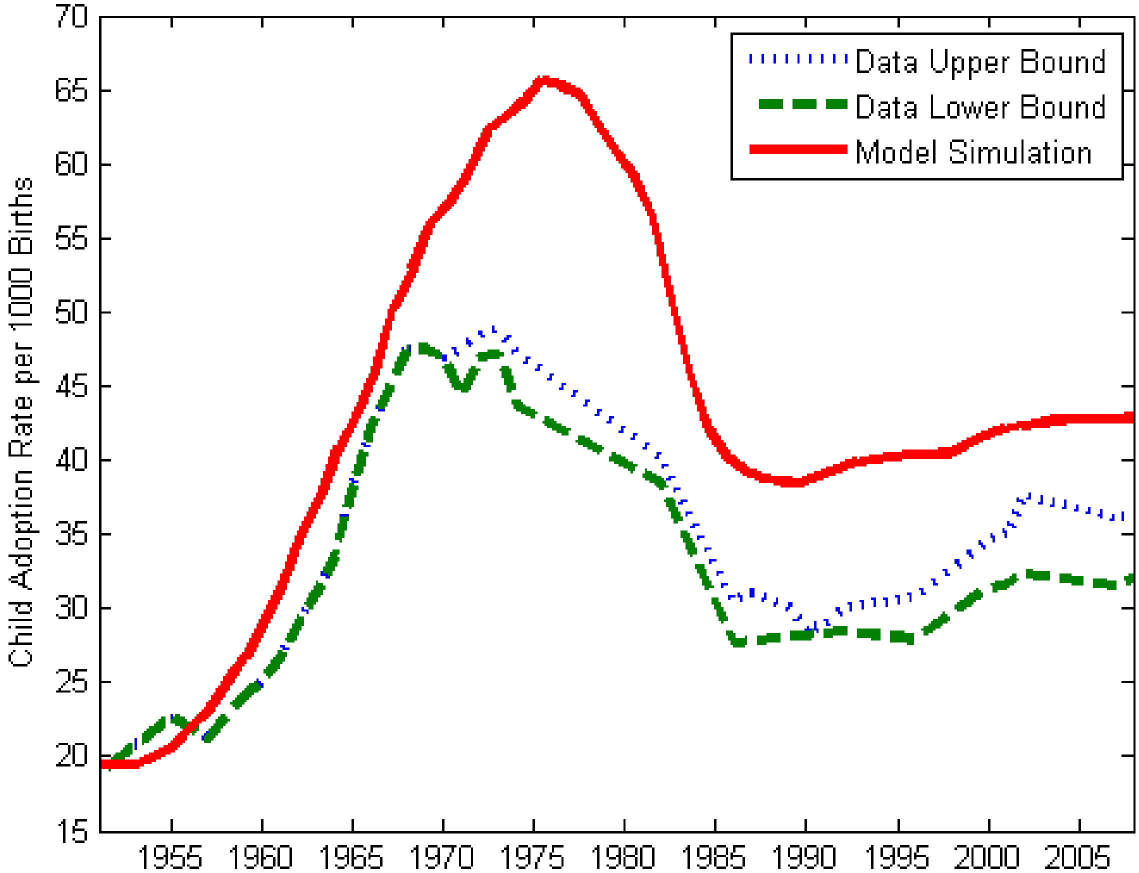

1.5 Conclusion

In this study, I examined the roles of changes in demand and supply sides of the child adoption market in producing the historical child adoption trend. First, I developed a life-cycle model of a woman's optimal fertility timing conditional on returns to her work experience, the current state of ART development and availability of adoption. Intuition behind the process is that, if the effect of the innovations in reproductive technologies is higher than the age-related decline in fecundity due to an increasing optimal fertility timing, then it produces a decline in adoption while still increasing the age of mothers. I treated adoption as an alternative means of family formation for women who are unable to conceive. I used the model to simulate the 1950-2010-s adoption trend. I found that increasing returns to female work experience and the development of ART during the latter half of the Twentieth Century can explain the shape of adoption trends in the U.S.. The alternative hypothesis that the adoption market in the U.S. was mainly driven by the "supply" side — i.e., the fall of adoption rates after 1970 is due to the increase of the abortion rate — was also considered in a counter-factual experiment but failed to either produce the observed shape of the historical trend or come close in magnitudes to it. The "demand side" hypothesis presented in this paper is also consistent with the observed higher rate of increase of the median age of women at birth as well. The logic of this explanation was shown analytically and quantitatively

CHAPTER 2
MIND THE GAP: WHAT EXPLAINS CHANGES IN THE RELATIVE TIMING OF MARRIAGE AND FERTILITY?

2.1 Introduction

Over the past half-century, the American family has seen dramatic demographic shifts. Marriage and fertility have been delayed while single-motherhood rates have been increasing. The three trends are bonded together through the decrease in "the Gap," i.e., the difference between the timing of first birth and first marriage, which emanates from the greater delay in marriage than that in fertility. As the Gap between the timing of birth and marriage decreases and even becomes negative, the share of women giving birth out-of-wedlock increases. This paper aims to account for the four closely-related trends using observed changes in income inequality and mobility. To do this, I use an equilibrium two-sided marriage search framework where agents make marriage and fertility choices over the life-cycle. In my model, marriage provides partial insurance against idiosyncratic earnings risk, while children represent a monetary and time commitment. In the model, I highlight two key mechanisms: the first links an increase in income inequality first and foremost with the delay in marriage timing, while the second mainly establishes a positive relationship between income mobility and fertility timing. In a nutshell, increasing income inequality produces incentives to delay marriage, which in turn makes women postpone fertility since single women, lacking partial insurance from marriage, face a higher level of earnings risk. Moreover, the decrease in income mobility also delays marriage as the insurance motive for marriage weakens but accelerates fertility since it becomes less risky to have a child. The model qualitatively matches the observed changes in family formation and quantitatively accounts for 42% and 40% of the observed changes in marriage and fertility timing, respectively, between the National Longitudinal Survey of Youth cohorts of women.

A substantial amount of research devoted to marriage and fertility delay – as well as to the increase in single motherhood – tends to treat the trends separately.[1] As a result, those theories aiming to explain changes in only one of them often fail to match qualitative facts related to the others. Although a number of papers acknowledge a fundamental interdependence between the marriage and fertility timing choices, this paper is the first to explicitly link them to out-of-wedlock fertility through the (overlooked by the literature) decrease in the Gap phenomenon, to the best of my knowledge. Of course, the decrease in the Gap is related to non-marital *first* births first and foremost, but most of the recent increase in overall out-of-wedlock fertility can be attributed to this very category.[2]

One way to depict the facts about marriage and fertility timing described above is shown in Figure 2.1. Here, I use the two cohorts from the National Longitudinal Survey of Youth (NLSY). The first consists of women born between 1956 and 1964 (labeled "60s") and the second of women born between 1980 and 1984 (labeled "80s"). As can be seen in Figure 2.1, the fraction of ever-married women by age (red

[1]The few exceptions are Caucutt, Guner & Knowles (2002); Regalia, Rios-Rull & Short (2011); Santos, Weiss (2016)

[2]Wu, Bumpass & Musick (2001)

lines) decreased by more than the share of mothers (black lines) did across the two cohorts. These shifts can be summarized by the change in the conditional medians. In particular, the median age at first birth was delayed from age 22.5 to age 23.6, while the median age at first marriage was delayed from age 21.2 to age 24.3. Hence, the Gap as described by the difference between the medians decreased by almost two years. As the Gap becomes negative and, hence, first births occur before marriage for some, it tends to translate into an increase in the share of women who gave birth to their first child out-of-wedlock (blue lines). In Section 3, I show that these facts hold more broadly: for example, the Gap decrease is robust to redefining it in terms of mean ages and to the exclusion of all "shotgun marriage" observations. I also show that it is applicable to all major socio-economic groups of women and qualitatively present in Current Population Survey data (CPS).

Studying the three trends together with an emphasis on the Gap gives a wider perspective compared to studying marriage, fertility and single motherhood separately. As documented in Section 3, the decrease in the Gap is relevant to all major socio-economic groups of US women. Single motherhood is often attributed to poor, low-educated and black women. The difference is that for more affluent women, the Gap was initially much more substantial than that of their less advantaged counterparts. One implication of the universality of the decrease in the Gap is that although we do not currently observe high levels of out-of-wedlock births among educated women, the share of educated women who give birth before marriage has more than tripled. Hence, if the Gap continues its trend, single motherhood will no longer be a sub-group phenomenon.

Figure 2.1: Decrease in the Gap from a Life-Cycle Perspective

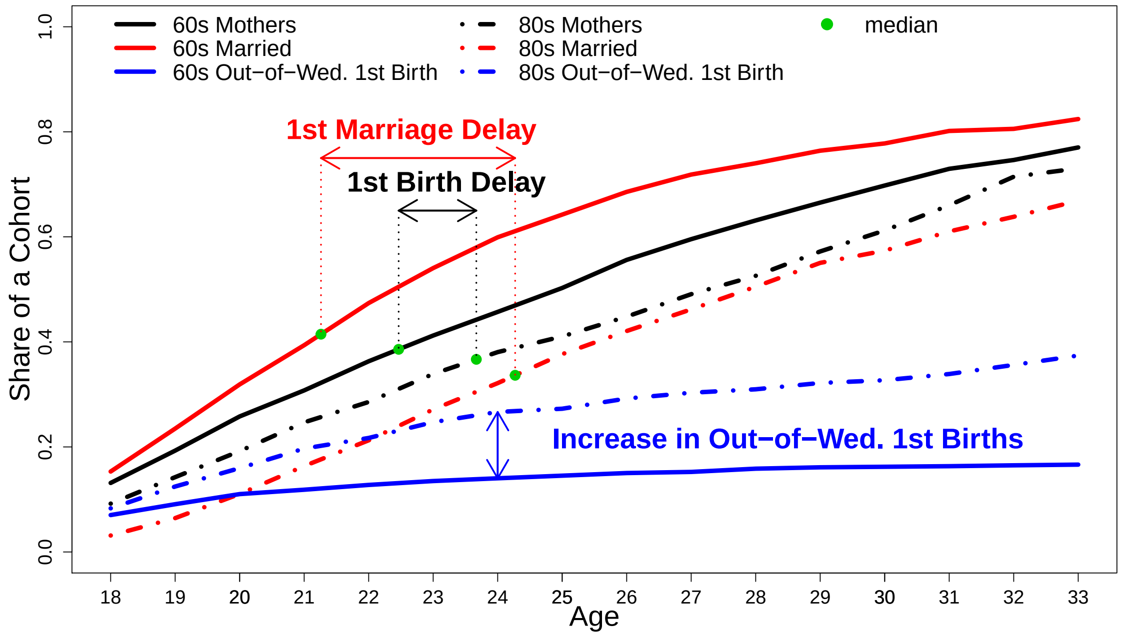

Note: Throughout the paper, National Longitudinal Survey of Youth (NLSY) is my primary data source. Since the data for the later cohort (NLSY97) are only available through the age of 33, what I refer to as median age is, in fact, the median age conditional on making fertility/marriage choices prior to age 33.

The two trends in the income process that this paper argues are important drivers of changes in marriage and fertility timing are the increase in income inequality and the decrease in income mobility over the last 50 years. To estimate these changes, I use a non-pararametric income process similar to that in De Nardi, Fella & Paz-Pardo (2016), which, among other benefits, allows me to clearly distinguish changes in income inequality from those in income mobility. The basic structure of the process is that at every age, agents receive one of N age-specific wages, which represent quantiles – i.e., the share of agents receiving this particular wage is exactly $1/N$. In the beginning of the next period, wages evolve according to an age-specific transition matrix. In such a setting, income inequality, defined as the variance of log of age-specific wages, is disentangled from income mobility as represented by transition probabilities. I find that inequality mainly increased among men, especially at higher ages. This is similar to the findings in Heathcote, Perri, & Violante (2010) and other recent papers using various data sets and methods.[3] With regard to income mobility, I find that its decrease was especially pronounced among women. This decrease in income mobility and female mobility in particular is also confirmed by Orzag & Director (CBO 2007), who used Social Security Administration data (SSA), and a number of other studies (see the discussion in Section 2).

The nature of marriage and fertility is closely related to the notions of risk, insurance and commitment – the classic wedding vow starts with the words "for better or for worse, for richer, for poorer..." Moreover, we traditionally regard families with children, and single mothers especially, as the most vulnerable (i.e. most exposed to risk) members of society. It is probably for this same reason that marriage and fertility have traditionally been viewed as inherently related demographic choices. After all, an additional exposure to risk due to the future expenditure commitments associated with child-rearing was compensated by marriage, a social institution providing informal insurance. Until relatively recently, the conventional wisdom supported by

[3]Katz & Murphy (1991); Heathcote, Perri, & Violante (2010); Debacker et.al. (2013)

available empirical evidence[4] suggested that the ongoing increase in income inequality has also implied an increase in income risk. Indeed, as the variance of log earnings increases, all else being equal, so does the conditional variance of log earnings growth; therefore, an increase in income inequality produces an increase in income risk. Since an increase in volatility makes single motherhood riskier and marriage insurance more desirable, if both moments of the income process were rising, then explaining the decrease in the Gap and the associated increase in out-of-wedlock fertility would be problematic. This puzzle has motivated researchers to seek alternative mechanisms unrelated to the risk-insurance link between marriage and fertility; several of them are discussed in Section 3. The results of this search were often unfruitful, as stated by Ellwood & Jenks (2004): "Indeed, it is only a slight exaggeration to say that quantitative social scientists main contribution to our understanding of this change has been to show that nothing caused single-parent families to become more common."

In the current paper, I show that, what is missing in the conventional logic is the fact that income risk (volatility) is a function of both income inequality – how far apart the potential conditional earnings realizations are – and income mobility – how likely the transitions to those realizations are. Therefore, a secular decrease in income mobility can more than offset the effect of income inequality on income volatility. For example, if the income of an individual can grow by either \$100 or \$200 with equal probabilities in the next period, the income risk that she faces would increase with the variance of the outcomes (say, if the new values were \$100 and

[4]See the discussion of the PSID in Section 2.

\$210). If, however, in addition her income mobility were to decrease (say, if the probability of a \$100 outcome became 99%), the resulting income risk may actually decline. Moreover, according to recent findings in the Social Security Administration data, this is exactly what happened. Given that, in my proposed model, an increase in income inequality produces a delay in both marriage and fertility timing. Meanwhile, decreasing income mobility produces a decrease in the Gap and an increase in the out-of-wedlock fertility by making childbearing less risky while the insurance provided in marriage less desirable. That is, it accelerates fertility while enhancing the delay in marriage.

In order to relate changes in the timing of first birth and marriage to changes in income inequality and mobility, I build a model similar to that of Aiyagari, Greenwood & Guner (2000), who consider a two-sided marriage search equilibrium framework. Agents live for a predetermined number of periods and make choices with respect to consumption and marriage, while couples and single females also decide on fertility.

In my model, marriage represents a long-term commitment to consumption-sharing. Agents, therefore, value marriage because this commitment partially insures them against idiosyncratic income risk. On the other hand, the requirement to split consumption incentivizes agents to decline marriage if the prospective partner is relatively poor. In addition, agents derive utility from children, but being a parent means committing oneself to future monetary and time costs associated with child-rearing. In order to decrease the complexity of the model, I assume that marriage is an absorbing state and marriages happen only within a cohort. I also assume that there are no savings and no borrowing. Because all agents are 21 to 31 years old, the no-savings assumption is not crucial since this is a period of a rapid earnings growth. The no-borrowing assumption is important for the agent's behavior, but, on the other hand, it is not too unnatural: the real-life ability of young individuals to borrow against their uncertain future income is very limited as well. Married agents enjoy consumption economies of scale and the choices of a married couple are the collective outcome of a Kalai-Smorodinsky bargaining process with exogenously fixed equal bargaining powers.

In the model, I highlight two key mechanisms: Mechanism 1 links an increase in income inequality with the delay in marriage timing, while Mechanism 2 establishes a positive relationship between earnings mobility and fertility timing. According to Mechanism 1, when earnings inequality increases, agents at the top of the earnings distribution become more selective in choosing whom they want to split their consumption with, and thus delay marriage. Agents who are still viewed as marriageable by the top-earners, will delay marriage because they become more likely to meet a single top-earner in the future and value of being married to him also becomes higher. In a two-sided environment, this causes a chain reaction as other types of agents also delay marriage because the best of their potential spouse-types do the same. Hence, marriage is delayed throughout the earnings distribution. The intuition behind Mechanism 2 is that high earnings mobility means (1) increased chances of low future earnings, where the monetary costs of a child outweigh the benefits, and (2) increased chances of high earnings, where the opportunity cost of time spent on child-rearing outweighs the benefits of (additional) children. Hence, an increase in income mobility makes agents delay fertility until later in life, when income mobility becomes lower.

According to my model, the observed changes in the income process are combined with the two mechanisms as follows. First, the increase in income inequality delays marriage in accordance with Mechanism 1. Since single women tend to face higher income risk than do married women, all else being equal, a decrease in marriages when young immediately implies delayed births according to Mechanism 2. Second, the decrease in income mobility also delays marriage as the insurance value of marriage decreases but accelerates fertility since it becomes less risky to have a child; that is, Mechanism 2 undoes part of the initial delay in birth.

Taken together, the increase in income inequality delays both marriage and fertility and the decrease in income mobility delays marriage further. It accelerates fertility, however, producing the decrease in the Gap between the two and therefore increases the incidence of single-motherhood.

Finally, I perform several quantitative experiments. First, having estimated the income process described above for the two cohorts of agents from the National Longitudinal Survey of Youth (NLSY79 & NLSY97), I calibrate the model to match age-specific shares of (1) ever-married women, (2) women who gave birth and (3) never-married mothers from age 21 to 31 of the initial cohort of women (NLSY79). I then simulate the change in the aforementioned age-specific shares by changing the income process to that estimated for the later cohort (NLSY97). I find that changes in the income process can account for 42% and 40% of the changes in the timing of marriage and fertility, respectively, between the two cohorts across ages. To decompose these responses into the effects of changes in income inequality and mobility, I simulate the model while changing only one feature of the income process at a time. I find that the change in income inequality accounts for 35% of the delay in marriage and 335% of the observed changes in fertility behavior. The latter effect is large because the increase in income inequality delays fertility through both delay in marriage and increasing income risk. The decrease in income mobility accounts for 14% of changes in marriage behavior between the two NLSY cohorts of women and for -124% of the observed changes in fertility.

Hence, income inequality and mobility jointly delay marriage. Their interaction suggests that for a given increase in income inequality, the effect of income mobility on the timing of marriage is weakened (42% $<$ 35% + 14%). On the other hand, while the increase in income inequality delays fertility markedly, the decrease in income mobility tends to accelerate it appreciably. Unlike the case of marriage, given the increase in income inequality, the effect of the decrease in mobility on the timing of birth is actually more potent. The intuition for the interaction effect on marriage timing is related to the fact that when income inequality is high, there are more types of agents whose incomes are so far apart that they would not agree to marry each other. This additional decrease in the value of marriage brought by decline in mobility would not be able to thwart those marriages. The intuition for the interaction effect on fertility, however, relates to the share of singles, which differs across low versus high income inequality environments. I analyze these interactions in more detail in Section 6.

The rest of the paper is organized as follows: In Section 2, I review the related literature; Section 3 presents empirical evidence related to the decrease in the Gap and tests the ability of several other hypotheses to explain it; Section 4 presents the model; Section 5 presents intuition for the two key mechanisms; in Section 6, I calibrate the model and run several counterfactual experiments; Section 7 concludes.

2.2 Related Literature

In the first part of this section, I overview the papers partially overlapping with my own in terms of the object of study, highlighting the differences and similarities in our approaches. In the second part, I review the literature that serves as a foundation of the model I develop.

2.2.1 Marriage, Fertility and Single Motherhood: Separately & Together

Marriage, fertility and out-of-wedlock childbearing are inherently related, and should therefore be studied together because theories aiming to explain changes in only one of them naturally cannot be used to study their interactions. For example, the papers by Happel, Hill & Low (1984); Blackburn, Bloom & Neumark (1992); Kohler & Kohler (2002); Kreyenfeld (2010); Adsera (2004, 2005, 2006); Sommer (2014) and Vanderbroucke (2014) etc. can be used to explain delay and decline in fertility but cannot be applied to the Gap since they treat marriage exogenously. On the other side of the spectrum, Keeley (1974); Loughran (2002); Gould & Paserman (2003); Coughlin & Drewianka (2011) etc. focus predominantly on marriage timing, and, therefore, are incapable of explaining the decrease in the relative timing as well. Finally, scholars of out-of-wedlock fertility such as Wilson (1987); Ellwood & Jencks (2004); Charles (2010); Lundberg & Pollak (2014, 2015, 2016) often consider income/education or race-specific mechanisms since most out-of-wedlock births occur among women of a low socio-economic status. Although such mechanisms can contribute to the share of single mothers among certain population groups, they are unlikely explanations for the decrease in the Gap, which, as I document, exists among all races and educational groups of women.

A number of theories produce simultaneous delay in marriage and fertility by elaborating on Becker's (1974) limits to the household specialization hypothesis. According to Becker, the value of marriage comes from household specialization of labor, which typically assumes that the husband has a comparative advantage in the labor market, while the wife is more efficient in home production (including childbearing). Typically, in such theories some exogenous innovation delays fertility by decreasing the incentives of individuals to have children. Since having children is seen as part of the value of being married, the undesirability of children also delays marriage. Examples of such innovations are an increase in income volatility as in Santos & Weiss (2012, 2016), an increase in the general income level and decrease in the gender-wage gap as in Regalia et al. (2008) and the invention of orally-administered birth control medication as in Goldin & Katz (2002). Although such mechanisms are indeed able to produce delays in both marriage and fertility, they naturally fail to also account for the decrease in the Gap and increase in single motherhood, since the decreased desire to have children was the original reason for marriage being less attractive and, hence, delayed. Because out-of-wedlock fertility is disproportionately concentrated among low-educated and black women, researchers often propose race or socio-economic class-specific mechanisms, such as hypothesis of black incarceration, unmarriageable men, etc. I show that the decrease in the Gap is relevant to all major socio-economic groups of women, and therefore the group-specific mechanisms may contribute to but cannot explain the observed trends for the whole population of women.

Closely related to my work, Santos & Weiss (2012, 2016) used PSID as their primary data source and were therefore using an increasing trend in volatility as a driving force for their mechanism (see the following discussion of the PSID). The authors proposed an explanation of the marriage and fertility delay in the spirit of Becker's (1974) specialization hypothesis. Given that an increase in income mobility disincentivizes childbearing, Santos & Weiss argued that it will also delay marriage by limiting the degree of potential specialization. But even if income mobility indeed increased, such a theory would still fail to explain the increase in single motherhood. If it is too risky to have a child while being married, it should be even riskier to have a child in the absence of marital insurance.

Another relevant paper, Regalia, Rios-Rull & Short (RRS) (2011), explains the increase in single-motherhood and decrease in marriage with the closing of the gender-wage gap. The intuition behind their explanation is yet another rendition of Becker's limits to specialization hypothesis – richer females can now afford to spend more time searching for a better partner. As argued before, this type of mechanism usually cannot produce an increase in single motherhood since the factors delaying marriage are exactly those demotivating fertility – in the case of RRS, higher female wages also should make the time cost of child-rearing binding. RRS overcome this problem by making special assumptions forcing some fraction of women to have children independently of their desire to be a mother. The current paper is complementary to RSS in that the gender-wage gap margin is also present here, but my income process also incorporates other important margins such as richer income inequality and income mobility. These appear in RRS in a rather reduced form.

This paper is also related to Caucutt, Guner & Knowles (2002), who showed how a delay in marriage provides incentives for fertility delay. However, studying relative changes in the timing of marriage and fertility, as well as focusing on the effects of income inequality and mobility in a unified framework of marriage and fertility choices, was beyond the scope of their research.

2.2.2 Mechanisms, Trends and Modeling

The validation of the crucial assumption that marriage provides partial insurance, among others, can be found in Hess (2004), who showed that income insurance arising through the marriage commitment plays an important role in marriage formation as well as in marriage longevity.

The hypothesis that there is a link between income inequality and age at first marriage, both of which have been increasing for the last 50 years,[5] was initially proposed by Keeley (1974). This result was derived in a one-sided search model (a version of a reservation wage model). Keeley's intuition was that increased earnings inequality increases the mean value of potential spouses whose wages are distributed above the reservation value. Therefore, under some conditions, returns to search increase as well as search duration. The fact that increasing inequality is associated with delay in marriage has been empirically confirmed by a number of papers.[6] However, Keeley's original intuition cannot be applied in the two-sided search case. For example, when men can reject marriage proposals, increased inequality of the upper tail of the male distribution does not necessary increase the gains to search for all females: rich males who are becoming even richer may now reject marrying some poorer women to avoid sharing their consumption with them.

To the best of my knowledge, this paper is the first to provide intuition about the way in which income inequality affects marriage timing in a two-sided search framework. Although increasing income inequality does not always theoretically delay marriage in the setting I use, marriage is delayed for parameter values similar to U.S. data (see the discussion of mechanisms in Section 5).

The decreasing trend in income mobility, which is responsible for the decrease

---

[5]For marriage delay trends, see a review by Stevenson & Wolfers (2007) and a documentation of the trend by Goldstein & Kenney (2001); Arroyo et. al. (2013).

[6]Oppenheimer, Kalmijn & Lim (1997); Loughran (2002); Gould & Paserman (2003); Watson & Mclanahan (2011); Coughlin & Drewianka (2011); Kearney & Levine (2012).

in the Gap and provides incentives for out-of-wedlock fertility in my model, was documented and confirmed only relatively recently, when economists gained access to the SSA data.[7] Before the SSA became available, the largest longitudinal data set at the researcher's disposal was the Panel Survey of Income Dynamics, and studies based on it often produced increasing trends in income volatility and mobility.[8] Neveu (2015) resolves this contradiction. He finds that the unbalanced nature of PSID explains much of the disparity between the PSID and SSA results.

The second key mechanism in my model, which links income mobility with fertility delay, dates back to Ranjan (1999) and has been widely accepted in the literature since. The relevance of this mechanism was shown in several papers, including Wong (2011) and Sommer (2014), both of whom used the variation in income risk associated with different occupations while treating marriage as exogenous. Kohler & Kohler (2002); Kreyenfeld (2005); Adser (2004); Bhaumik, S. K., and J. B. Nugent (2006); Vandenbroucke (2012); Goldstein et al. (2013) etc. found the same effect of "negative uncertainty," which is an exposure to the risk of decline in future earnings

[7]Orzag & Director (CBO 2007) find that the probability of an increase or drop in the real wage decreased for all ages. Other studies mainly examined income volatility. Because we know that income inequality has increased and income volatility has decreased, this necessarily leads to a decrease in income mobility. Sabelhaus & Song (2010) show a parallel downward shift in the life-cycle trend of income volatility. Guvenen et al. (2014) show a decrease in cross sectional income volatility. Guvenen (2016) shows that the cross sectional dispersion of income growth has decreased for both genders since 1980 at the least. Finally, Neveu (2015) confirms this fact using PSID, and documents that income volatility was declining for both genders in the mid-1980s, after which it continued to decrease only among females and stabilized among males.

[8]See Moffitt & Gottschalk (1994, 1995, 2012); Katz & Dickens (1994); Heathcote & Perri (2010); Shin & Solon (2011).

on fertility timing.

The model I use is similar in spirit to the one in Aiyagari Greenwood & Guner (2000) and a growing number of papers following it.[9] From a technical point of view, an important distinction between models of this type and mine is that I introduce a non-parametric income process similar to that of De Nardi, Fella & Paz-Pardo (2016). This non-parametric structure decreases the size of the state-space and computational intensity, while accommodating more periods of the life-cycle and a richer income process.[10] Also, having a finer income grid allows me to achieve a good calibration fit without adding utility shocks such as "the marital bliss", "the value of love" or stochastic childbearing, which is common in this type of models. The second advantage of this income process is that it clearly distinguishes between income inequality and income mobility. Finally, I add to the modeling literature by addressing a concern about the uniqueness of the equilibrium. Most models of this type are non-stationary in a life-cycle sense: i.e., distributions of agents in the marriage market evolve endogenously. The outside option of marriage today depends on the distribution of single agents tomorrow, while the latter itself depends on the marriage decisions (i.e., on the outside options) of agents today. Thus, it is generally

[9]Caucutt, Guner & Knowles (2002); Chade & Ventura (2002); Greenwood, Guner,& Knowles (2003); Erosa, Fuster, & Restuccia (2002,2010); Fernandez, Guner, & Knowles (2005); Da-Rocha & Fuster (2006); Regalia, Rios-Rull & Short (2011); Guner, Kaygusuz & Ventura (2008); Guner, & Knowles (2009); Knowles (2012).

[10]For example, Caucutt Guner and Knowles (2002) model a three period life-cycle with seven, nine and eleven grid points for income in each period, respectively. Regalia, Rios-Rull & Short (2011) model four period life-cycle, and have only two wages for each gender at every age.

possible to have multiple self-fulfilling equilibria in this type of setting. Although I do not provide general conditions for uniqueness, I propose a dominated strategies elimination algorithm that can numerically check whether an equilibrium is unique given parameter values, example of the algorithm can be found in the Appendix.

2.3 Empirical Investigation

In this section, I document several facts related to the decrease in the Gap, and test the robustness of this phenomenon to several peculiar explanations existing in the literature. While some of them may play a role in explaining demographic trends, I show that they cannot individually explain the full extent of the decrease in the Gap phenomenon occurring across all demographic groups.

Table 2.1 shows the delay in median age at first marriage (Δ A1M), the delay in median age at first birth (Δ A1B), and the Gap between the two median ages among women born in the 1960s (Gap 60s) and among the later cohort (Gap 80s). It also notes the percentage of out-of-wedlock first births in each cohort. The first fact that one can see in Table 2.1 is that all major socio-economic groups of women experienced the decrease in relative timing, which led to increase in the rates of first out-of-wedlock motherhood. Nevertheless, white women with a college degree or above, did not undergo the crossover of the median ages at first birth and first marriage as a group, since their initial Gap was larger than that of the other groups of women. Nevertheless, I document a greater than three-fold increase in the share of first births to unmarried women in this subgroup of the population.

Table 2.1: Domain of the Decrease in the Gap

	Δ A1M	Δ A1B	Gap 60s	Gap 80s	% 1st bir. to single 60s	% 1st bir. to single 80s
All women	3.01	1.15	1.20	-0.65	21	54
White, $\geq$ College	2.56	1.61	3.46	2.51	4	14
Black, $\geq$ College	4.56	3.31	-1.00	-2.26	48	74
White, $<$ College	2.95	0.55	1.86	-0.55	15	56
Black, $<$ College	3.46	1.45	-2.46	-4.46	71	88

Source: National Longitudinal Survey of Youth (NLSY79 & NLSY97)

Note: Δ A1M and Δ A1B denote delays in the median ages at first marriage and birth respectively.

Still, the overwhelming majority of out-of-wedlock births today are to low-educated and black women, or, as Ellwood & Jencks (2004) note, “the increase in non-marital childbearing, [...], mainly affected non-white women and white women without college degrees.” Given such an uneven spread of out-of-wedlock fertility, it is no surprise that many scientists looked for and proposed a number of income- and race-specific explanations of the phenomenon. A few examples of such explanations are a hypothesis of “non-marriageable” men proposed by Wilson (1987) and a hypothesis related to the black incarceration rate.

Wilson (1987) argued that the marriage rates among low-educated women decreased because low-educated men – victims of globalization and a skill-biased technical change – simply became too poor to marry them. This hypothesis was formally tested by Wood (1995), who found only a small effect (3%) of change in the quality of the male marriage pool on low-educated female marriage rates.

A different argument in a similar spirit was proposed by Lundberg & Pollak (2014, 2015, 2016). The authors argue that since low-educated individuals invest less in their children, they do not value the commitment provided by marriage as much as do highly-educated people. While this argument may be true, one can alternatively argue that low-income individuals may value other properties of marriage, such as a lower per-parent monetary cost of a child, more than their affluent counterparts do.

Yet another example of a group-specific explanation is the hypothesis that the increase in single motherhood rates among black women was caused by disproportionate incarceration rates among black men. There is no doubt that this effect is a contributor to the high rate of first out-of-wedlock births among low-educated black women. At the same time, as argued above, a race-specific mechanism is an unlikely cause of changes found among all groups of women.

Another possible explanation of the observed demographic changes is a so-called cultural drift. One example of this type is a decrease in shotgun marriages.[11] This idea was articulated by Akerlof, Yellen and & Katz (1996), who argued that such

[11] In common usage, shotgun marriages are understood as taking place between conception and childbirth. Because the quality of the data is often insufficient to allow for the identification of the precise timing of birth and marriage, shotgun marriage is usually defined as happening within a year of childbirth for research purposes.

a cultural change can explain most of the change in out-of-wedlock fertility. Carlson, Mclanahan & England (2004), however, use a richer data set and do not find any effect of the decrease in the incidence of shotgun marriages on out-of-wedlock fertility. I address the issue of shotgun marriages in Table 2.2. The first row presents the same statistics as in Table 2.1, after deleting all observations in which the couple is married within a year of childbearing. The second row presents an experiment in which all women with shotgun marriages were rewritten in the sample as "never married." The two numerical experiments suggest that although the change in shotgun marriages may contribute to the increase in single motherhood, it cannot fully explain the decrease in the Gap.

Another explanation related to a cultural drift is that people simply value marriage to a lesser degree than in previous years. Of course, it is difficult to address the robustness and magnitude of the effect of a change in preferences. Even the argument of whether preference changes occurred or not can be fruitless. Tucker (2000) provides an overview of studies collecting data on self-reported attitudes towards marriage and other demographic choices. Most such studies do not find any significant differences between the family values of people from different demographic groups or over time.

Table 2.2: Shotgun Marriages

	Δ A1M	Δ A1B	Gap 60s	Gap 80s	% 1st bir. to single 60s	% 1st bir. to single 80s
delete shotgun observations	2.51	-0.71	2.30	-0.90	28	65
rewrite shotgun as "never married"	2.51	1.15	0.50	-0.85	52	73

A separate branch of the – related – out-of-wedlock literature considers single motherhood as a static phenomenon. For example, Willis (1999) and Choo & Siow (2006) both develop static equilibrium models in which some women optimally choose to become single mothers. Indeed, one can imagine a scenario whereby most women do not change their marriage or fertility behavior, but a new category of "static" single mothers appears in society. In such a scenario, we would observe a decrease in the Gap, but it would be meaningless to study it. The true change would be the appearance of this new population category, rather than changes in relative timing. Table 2.3 presents three experiments addressing the theory of "static" single motherhood. The first row presents the results when only women who are ever-married and mothers by the age of 33 are included in the sample. Because all women in this reduced sample make both fertility and marriage choices in their life-cycle, in the second row, I show the median Gap between the two events rather than the Gap between medians. Finally, the third row of Table 2.3 presents statistics for the case in which all women who are single mothers at age 33 are excluded from the sample. Table 2.3 provides evidence that the effect of "static" single mothers, although a contributing factor, cannot explain the decrease in the Gap phenomenon.

Table 2.3: "Static" Single Motherhood

Status at age 33	Δ A1M	Δ A1B	Gap 60s	Gap 80s	% 1st bir. to single 60s	% 1st bir. to single 80s
Married & Mother	2.06	0.35	1.76	0.05	14	33
Married & Mother (Δ median gap)			0.64	0.05		
Not a Single Mother	2.96	2.56	1.45	1.05	13	33

2.4 The Model

2.4.1 Basic Environment

The economy is populated by agents who make choices for a predetermined number of periods T, at the end of which they receive a continuation value, which is equivalent to living an additional R periods without further available choices. Agents differ in gender $g \in \{(m)ale, (f)emale\}$, wages, marital status and number of children. Only married couples and single women can have children. Agents draw utility from consumption and number of children; men do not distinguish stepchildren from their own.

At the beginning of every period, agents observe their new wage realizations and single agents proceed to the marriage market. Here, they have a chance of being randomly matched with a single agent of the opposite gender. If two agents are matched, they decide whether they want to marry each other, continuing as a married couple if they both agree. Unmatched and unmarried agents proceed as single. After the marriage market phase, all agents make consumption choices. Single females and married couples also decide whether they want to have an additional child or not: $k_t \in \{0, 1\}$.

Once two agents marry, they permanently exit the marriage market to fulfill the no-divorce model assumption. Distributions of single agents are not normalized: i.e., if some share of the population is already married, some single agents would be matched with nobody.

Each child imposes a fixed monetary cost η_m, which does not differ by family type. In addition, there is a time cost of raising a child η_τ, which is bigger for married couples – this assumption reflects time spent with the child by both parents concurrently. Since I consider decisions made when agents are between 21 and 31 years old, children are assumed to never leave the household.

The total measure of male agents is normalized to 1 and equals that of females. Let w_t^f and w_t^m denote the period t wages of females and males, respectively. Wages can take one of the N *age* $\times$ *gender*-specific values in every period:

$$w_t^g \in W_t^g \equiv \{w_{t,1}^g, ..., w_{t,N}^g\}$$

Empirically, these wages are estimated as mean wages of the N *age* $\times$ *gender*-

specific quantiles, and the measure of men and women with a given wage at any period is therefore $1/N$.

All wages evolve according to a *age* $\times$ *gender*-specific transition matrix:

$$\Pi_t^g \equiv \begin{array}{|c|c|c|c|} \hline & w_{t+1,1}^g & \cdots & w_{t+1,2}^g \\ \hline w_{t,1}^g & \pi_{t,1,1}^g & \cdots & \pi_{t,1,N}^g \\ \hline \vdots & \vdots & \ddots & \vdots \\ \hline w_{t,N}^g & \pi_{t,N,1}^g & \cdots & \pi_{t,N,N}^g \\ \hline \end{array},$$

where $\pi_{t,i,i'}^g = Pr(w_{t+1}^g = w_{t+1,i'}^g | w_t^g = w_{t,i}^g)$ is the probability that an agent of gender g will have a wage $w_{t+1,i'}^g$ at age $t+1$, conditional on having a wage $w_{t,i}^g$ at age t.

Let $\mu_{t,i} \in [0, 1/N]$ be the measure of men of wage-type i who are single in the beginning of period t (i.e. the probability of matching with this type of male in the marriage market), and let the vector of all matching probabilities for single males be $\mathcal{M}_t \equiv \{\mu_{t,i}\}_{i=1}^N$. The type of a single female is specified by her wage and the number of children that she has: K_{t-1}; let a type-specific measure and a vector of matching probabilities for single females be denoted as $\phi_{t,i}(K_{t-1}) \in [0, 1/N]$ and $\Phi_t \equiv \{\{\phi_{t,i}(K_{t-1})\}_{i=1}^N\}_{k=0}^{t-1} \equiv \{\phi_{t,j}\}_{j=1}^{N\times t}$, respectively.

2.4.2 The Equilibrium

In order to define the equilibrium for this economy, I first list the problems that different types of agents face at every age. Since the model is non-stationary, in order to define the value functions, I start with the last period of agent's lives T.

2.4.2.1 Period T

Let the value of male single life be:

$$M_T(w_T^m) = \max_c U(c) + \beta M^{CV}(w_T^m) \tag{2.1}$$

subject to

$$c \leq w_T^m,$$

where $M^{CV}(w_T^m)$ is the continuation value of living another R periods as a single male as a function of his wage in the last period. Let $C_T^m(w_T^m)$ be the optimal consumption choice in the problem above.

The value of female single life:

$$F_T(w_T^f, K_{T-1}) = \max_{c, k_T \in \{0,1\}} U(c) + V(K_T) + \beta F^{CV}(w_T^f, K_T), \tag{2.2}$$

subject to

$$c + \eta_m K_T \leq (1 - \eta_\tau K_T) w_T^f$$

$$K_T = K_{T-1} + k_T,$$

where η_m, η_τ are the monetary and time costs, respectively, per child and $F^{CV}(w_T^f, K_T)$ is a continuation value of living another R periods as a single female as a function of her wage in the last active period and her completed fertility. Let the optimal consumption and fertility choices of a single female at period T be denoted as $C_T^f(w_T^f, K_{T-1})$ and $k_T^f(w_T^f, K_{T-1})$.

I define the value of the life of each spouse in a married couple as:

$$MC_T(w_T^f, K_{T-1}, w_T^m) = \max_{c, k_T \in \{0,1\}} U\left(\frac{c}{1+\gamma}\right) + V(K_T) + \beta MC^{CV}(w_T^f, K_T, w_T^m), \tag{2.3}$$

subject to

$$c + \eta_m K_T \leq (1 - \alpha \eta_\tau^{MC} K_T) w_T^f + (1 - (1-\alpha)\eta_\tau^{MC} K_T) w_T^m,$$

$$K_T = K_{T-1} + k_T.$$

Because the bargaining powers are endogenously set to be equal, the above problem is equivalent to the Kalai-Smorodinsky bargaining outcome. Married agents pool their incomes together and choose aggregate consumption, after which each of them enjoys a fraction $\frac{1}{1+\gamma}$ of it. $\gamma \in [0, 1]$ represents family consumption economies of scale. A child is a public good, so each spouse derives utility from being a parent. Finally, the way spouses split their parenting duties is fixed exogenously, and α is the fraction of parenting time contributed by the mother. In addition, in order to account for the fact that parents sometimes engage in parenting activities simultaneously, the total time cost of a child for a married couple is higher than that of a single mother:

$\eta_\tau^{MC} > \eta_\tau$.

The solution of the married couple problem, then, is represented by consumption and fertility decision rules: $C_T^{mc}(w_T^f, K_{T-1}, w_T^m)$ and $k_T^{mc}(w_T^f, K_{T-1}, w_T^m)$.

Given these values of being single and married, during the marriage market phase at the beginning of period T, two agents will marry if and only if

$$I_T(w_T^f, K_{T-1}, w_T^m) = \begin{cases} 1, & MC_T(w_T^f, K_{T-1}, w_T^m) \geq F_T(w_T^f, K_{T-1}) \quad \text{and} \\ & MC_T(w_T^f, K_{T-1}, w_T^m) \geq M_T(w_T^m) \\ 0, & otherwise \end{cases} \tag{2.4}$$

Now I can define the expected value of being a single male before the marriage market phase in period T begins:

$$\begin{aligned} EM_T(w_T^m, \Phi_T) =& \Sigma_j \phi_{T,j} I(w_{T,j}^f, K_{T-1,j}, w_T^m) MC_T(w_{T,j}^f, K_{T-1,j}, w_T^m) + \\ & \Sigma_j \phi_{T,j} \left(1 - I(w_{T,j}^f, K_{T-1,j}, w_T^m)\right) M_T(w_T^m) + \\ & \left(1 - \Sigma_j \phi_{T,j}\right) M_T(w_T^m), \end{aligned} \tag{2.5}$$

The first line of Equation 2.5 represents the expected utility of meeting and marrying a female agent; the second line stands for the cases when agents would be matched but marriage does not occur, and the last line represents cases when an agent will not be matched with a female agent because some agents are already married.

Similarly, the expected value of being a single female before the marriage market phase in period T begins is:

$$\begin{aligned} EF_T(w_T^f, K_{T-1}, \mathcal{M}_T) =& \Sigma_i \mu_{T,i} I(w_T^f, K_{T-1}, w_{T,i}^m) MC_T(w_T^f, K_{T-1}, w_{T,i}^m) + \\ & \Sigma_i \mu_{T,i} \left(1 - I(w_T^f, K_{T-1}, w_{T,i}^m)\right) F_T(w_T^f, K_{T-1}) + \\ & \left(1 - \Sigma_i \mu_{T,i}\right) F_T(w_T^f, K_{T-1}). \end{aligned} \tag{2.6}$$

2.4.2.2 Period $t \leq T-1$

All period-indexed agent's problems are identical, except for that of period T. The value of being male at $t \in \{1, .., T-1\}$ is

$$M_t(w_t^m, \Phi_{t+1}) = \max_c U(c) + \beta \mathbb{E}_{w_{t+1}^m} \left[EM_{t+1}(w_{t+1}^m, \Phi_{t+1}) | w_t^m \right], \tag{2.7}$$

subject to

$$c \leq w_t^m.$$

Similarly to the problem in period T, the decision rule for consumption as a single male is $C_t^m(w_t^m, \Phi_{t+1})$.

The value of single life for a female is

$$F_t(w_t^f, K_{t-1}, \mathcal{M}_{t+1}) = \max_{c, k_t \in \{0,1\}} U(c) + V(K_t) + \beta \mathbb{E}_{w_{t+1}^f} \left[EF_{t+1}(w_{t+1}^f, K_t, \mathcal{M}_{t+1}) | w_t^f \right], \tag{2.8}$$

subject to

$$\begin{aligned} & c + \eta_m K_t \leq (1 - \eta_\tau K_T) w_t^f \\ & K_t = K_{t-1} + k_t. \end{aligned}$$

The optimal decision rules associated with this problem are: $C_t^f(w_t^f, K_{t-1}, \mathcal{M}_{t+1})$ and $k_t^f(w_t^f, K_{t-1}, \mathcal{M}_{t+1})$ for consumption and fertility, respectively.

The value of the life of each spouse in a married couple is

$$MC_t(w_t^f, K_{t-1}, w_t^m) = \max_{c, k_t \in \{0,1\}} U\left(\frac{c}{1+\gamma}\right) + V(K_t) + \beta \mathbb{E}_{w_{t+1}^f, w_{t+1}^m}\left[MC_{t+1}(w_{t+1}^f, K_t, w_{t+1}^m) | w_t^f, w_t^m\right], \tag{2.9}$$

subject to

$$c + \eta_m K_t \leq (1 - \alpha \eta_\tau^{MC} K_t) w_t^f + (1 - (1-\alpha)\eta_\tau^{MC} K_t) w_t^m$$

$$K_t = K_{t-1} + k_t.$$

The optimal decision rules for a married couple associated with this problem are: $C_t^{mc}(w_t^f, K_{t-1}, w_t^m)$ and $k_t^{mc}(w_t^f, K_{t-1}, w_t^m)$.

For period $t \leq T-1$, two agents agree to marry if and only if

$$I_t(w_t^f, K_{t-1}, w_t^m, \Phi_{t+1}, \mathcal{M}_{t+1}) = \begin{cases} 1, & MC_t(w_t^f, K_{t-1}, w_t^m) \geq F_t(w_t^f, K_{t-1}, \mathcal{M}_{t+1}) \\ & \text{and} \\ & MC_t(w_t^f, K_{t-1}, w_t^m) \geq M_t(w_t^m, \Phi_{t+1}) \\ 0, & otherwise \end{cases} \tag{2.10}$$

Finally, the expected value of being a single male before the marriage market phase in period $t \leq T-1$ begins is

$$\begin{aligned} EM_t(w_t^m, \Phi_t) =& \Sigma_j \phi_{t,j} I(w_{t,j}^f, K_{t-1,j}, w_t^m, \Phi_{t+1}, \mathcal{M}_{t+1}) MC_t(w_{t,j}^f, K_{t-1,j}, w_t^m) + \\ & \Sigma_j \phi_{t,j} \left(1 - I(w_{t,j}^f, K_{t-1,j}, w_t^m, \Phi_{t+1}, \mathcal{M}_{t+1})\right) M_t(w_t^m, \Phi_{t+1}) + \\ & (1 - \Sigma_j \phi_{t,j}) \, M_t(w_t^m, \Phi_{t+1}); \end{aligned} \tag{2.11}$$

the expected value of being a single female before the marriage market phase in period $t \leq T-1$ begins is

$$\begin{aligned} EF_t(w_t^f, K_{t-1}, \mathcal{M}_t) =& \Sigma_i \mu_{t,i} I(w_t^f, K_{t-1}, w_{t,i}^m) MC_t(w_t^f, K_{t-1}, w_{t,i}^m) + \\ & \Sigma_i \mu_{t,i} \left(1 - I(w_t^f, K_{T-1}, w_{t,i}^m)\right) F_T(w_t^f, K_{t-1}, \mathcal{M}_{t+1}) + \\ & (1 - \Sigma_i \mu_{t,i}) \, F_t(w_t^f, K_{t-1}, \mathcal{M}_{t+1}). \end{aligned} \tag{2.12}$$

2.4.2.3 Distribution Updating

Once a male of type i and a female of type j are matched and agree to marry, they permanently exit the distribution of singles. The measure of type i single males after the marriage phase is then given by

$$\hat{\mu}_{t,i} = \mu_{t,i} - \Sigma_j I_t(w_{t,j}^f, K_{t-1,j}, w_{t,i}^m, \Phi_{t+1}, \mathcal{M}_{t+1}) \times \mu_{t,i} \times \phi_{t,j}. \tag{2.13}$$

After that, all measures of single males evolve according to the wage transition matrix.

Similarly, the measure of type j single females after the marriage phase is given by

$$\hat{\phi}_{t,j} = \phi_{t,j} - \Sigma_i I_t(w^f_{t,j}, K_{t-1,j}, w^m_{t,i}, \Phi_{t+1}, \mathcal{M}_{t+1}) \times \mu_{t,i} \times \phi_{t,j}. \tag{2.14}$$

After that, the matching probabilities of single females are updated according to their respective fertility decisions, evolving according to the wage transition matrix at the end of the period.

2.4.2.4 Definition

An equilibrium is a collection of value functions and consumption, fertility and marriage decision rules together with vectors of matching probabilities of single agents at the beginning of every period. These are such that all decision rules are optimal, taking the decisions of other agents and vectors of matching probabilities of single agents as given, and such that the value functions and measures of single agents are generated by those decision rules.

Definition 1. An equilibrium is a set of fertility decision rules by single women and married couples $k^f_t(w^f_t, K_{t-1}, \mathcal{M}_{t+1})$, $k^{mc}_t(w^f_t, K_{t-1}, w^m_t)$, a set of marriage decision rules, $I_t(w^f_t, K_{t-1}, w^m_t, \Phi_{t+1}, \mathcal{M}_{t+1})$ and a set of matching probabilities, $\Phi_t, \mathcal{M}_t$ of all types of agents at every age $\forall \{w^g_{t,i}\}_{i=1..N, g=\{f,m\}}, \quad \forall K^f_{t-1} = 0..t-1, \quad \forall K^{mc}_{t-1} = 0..t-1, \quad \forall t = 1..T$, such that:

1. The consumption, fertility and marriage rules are optimal taking the matching probabilities as given.
2. The matching probabilities are consistent with the marriage decision rules and fertility decision rules of single females.

2.4.3 Computation

From the Definition 1, it follows that, in equilibrium, matching probabilities $\{\Phi_t, \mathcal{M}_t\}_{t=1}^T$ are the fixed points of the mappings implied by the marriage acceptance rules $I_t(w_t^f, K_{t-1}, w_t^m, \Phi_{t+1}, \mathcal{M}_{t+1})$ and fertility decision rules of single females $k_t^f(w_t^f, K_{t-1}, \mathcal{M}_{t+1})$. This equilibrium can be found using the following iteration on the matching probabilities:

1. Make an initial guess of $\{\Phi_t, \mathcal{M}_t\}_{t=1}^T$
2. Given the current guess, compute value functions and decision rules of agents at every period by backwards induction.
3. Given marriage acceptance rules $I_t(w_t^f, K_{t-1}, w_t^m, \Phi_{t+1}, \mathcal{M}_{t+1})$ and fertility decision rules of single females $k_t^f(w_t^f, K_{t-1}, \mathcal{M}_{t+1})$, update the matching probabilities $\{\Phi_t, \mathcal{M}_t\}_{t=1}^T$
4. If the matching probabilities have changed during Step 3., Return to Step 2., treating $\{\Phi_t, \mathcal{M}_t\}_{t=1}^T$ as a new guess.

2.5 Mechanisms

The goal of this section is to provide intuition for the two key mechanisms, through which the increase in income inequality and decrease in income mobility produce the four facts characterizing the changes in timings of marriage and fertility: namely, delay in marriage and fertility, a decrease in the Gap in relative timing and an increase in out-of-wedlock fertility.

2.5.1 Mechanism 1: Income Inequality & Marriage Delay

The key assumptions producing a delay in marriage as a response to increasing (male) inequality in my model are: (1) marriage is a commitment (I assume no divorce, but it is not a necessary assumption); (2) there are limits to bargaining inside of the family (I assume equal bargaining weights, which is a simplifying, but not necessary assumption); (3) endogenous evolution of the distribution of agents; (4) finite horizon. Assumptions (1) and (2) motivate agents to marry only within a certain income range around their own wage-type. When income inequality among males increases (for example, the top male type becomes richer), the lowest wage-type of women that the top-income male type agrees to marry (reservation type) will increase as well. The value of search for high-income females rises, since the average quality of their marriage pool increases. Moreover, since the top men are pickier, more of them will be single in later periods – in this way, Assumption (3) reinforces the marriage delay incentives. Given that, high-income females become pickier as well and increase their reservation type for early periods. If the high-type women's own income has not changed, they will return to their old reservation types in later periods (because Assumption (4) provides for a finite-horizon model). Now, returns to delay for middle-income men increase as well, since they will be able to marry high-type females later. In turn, middle-income men raising their reservation types early in life will motivate low-income women to delay marriage, and so on. With this type of chain reaction to the increase in inequality, even changes at the very top of the income distribution would trickle down through it and alter marriage reservation types across the board. A numerical example of the above intuition can be found in the Appendix.

2.5.2 Mechanism 2: Income Mobility & Fertility Timing

This mechanism generates a positive association between income mobility (and income risk in general) and the timing of fertility and is responsible for the decrease in the Gap and increase in out-of-wedlock fertility. The intuition is straightforward and can be described as follows. Since raising a child costs money, there is a threshold income $\underline{w}$, below which an agent would optimally choose to stay childless. Given the concomitant time cost, there is a threshold wage $\overline{w}$, earning above which will make an agent choose a childless life as well because the opportunity cost of a child will be too high. In the absence of income mobility, once an agent's wage is between the two thresholds, she will choose to have a child, staying childless otherwise. In the presence of income mobility, however, even if the current wage of an agent satisfies $w \in [\underline{w}, \overline{w}]$, an agent may choose to defer childbearing if it is sufficiently likely that the future wage will be outside of this interval.

2.6 Matching the Model to the Data

This section aims to establish the quantitative importance of my model. In order to do so, I calibrate the model to the marriage and fertility behavior of the earlier cohort of women (NLSY79); I then perform a simulation of the demographic choices made over the life-cycle by women from the later cohort (NLSY97) and perform several other counterfactual experiments in order to assess the magnitude of changes that can be attributed to the increase in income inequality and decrease in income mobility.

2.6.1 Estimation of the Income Process

I estimate a ten-quantile ($N = 10$) income process from the NLSY79 and NLSY97 data sets. Due to data limitations, I consider only individuals aged 21 to 31. One period in the model corresponds to one year in the data.

The estimation procedure consists of the following steps:

1. Split all income observations into *age* $\times$ *gender* groups.

2. Split all the incomes inside of every *age* $\times$ *gender* group into N quantiles and compute the means of these quantiles. The size of the resulting income grid for each cohort of individuals is $\{m, f\} \times N \times T = 2 \times 10 \times 10$.

3. For every two consecutive periods (ages), compute a transition matrix (quantile mobility matrix) according to $Pr(w^g_{t+1} = w^g_{t+1,i'} | w^g_t = w^g_{t,i})$, which is equal to the fraction of individuals whose income was in quantile i at age t and moved to the *age* $\times$ *gender*-specific income quantile i' the next year. Overall, there are $T - 1$ quantile mobility matrices of size $N \times N = 10 \times 10$.

2.6.2 Calibration

In addition to the income process, I need nine further parameters in order to compute the model. I take four of them from the literature, choosing the rest by matching the model with the data.

2.6.2.1 Parameters Set from Existing Estimates

Because agents do not value leisure in my model, the time cost of child-rearing should reflect the labor market opportunity cost. Schoonbroodt (2016) provides estimates of parenting time during typical working hours. Under the assumption of a 40-hour work week, her results suggest that the time cost of a child is 18.5% of working time, i.e. $\eta_\tau = 0.185$. Folbre et al. (2005) provide an estimate for the overlap in parenting timing for the two-parent families, namely, $\eta_\tau^{MC} = 1.26\eta_\tau$. I follow the computation of Schoonbroodt (2016) for relative shares of paternal and maternal parenting time, setting the mother's share to $\alpha = 0.7$.

I use the standard discounting rate $\beta = 0.98$. Finally, the OECD average consumption economies of scale estimate is 0.7. I do not fix γ at this level, but require that it satisfy $\gamma \in [0.65, 0.75]$.

Finally, I assume that the continuation value is the expected utility of an agent over the next fourteen periods ($R = 14$), assuming that the transition probabilities and income values are the same as in period T. I choose fourteen periods because I do not track fertility histories in the model, but after fourteen periods most previously-born children are expected to become eighteen years old and leave their parents. Choosing a different number of periods would influence some of the estimated parameters, but would not significantly alter the overall fit of the calibration.

2.6.2.2 Parameters Estimated from the Calibration

I assume utility functions of the following form:

$$U(c) = \frac{c^{1-\sigma_c}}{1-\sigma_c},$$
$$V(K) = \psi \frac{K^{1-\sigma_k}}{1-\sigma_k}.$$

Overall, I estimate five parameters by calibration: three from the utility function (σ_c, σ_k, ψ); fixed monetary cost per child (η_m); economies of scale (γ), forced to be in the interval $\gamma \in [0.65, 0.75]$.

2.6.2.3 Calibration Targets

I calibrate the model by matching 33 targets based on the marriage and fertility behavior of the initial NLSY cohort. I target age-specific shares of ever-married women (eleven), age-specific shares of mothers (eleven) and age-specific shares of women who gave their first birth out-of-wedlock (eleven).

2.6.2.4 Fitting the Initial Cohort

Figure 2.2 compares the age-specific shares of mothers, married women and never-married mothers produced by the model with those estimated from the data. Values of the calibrated parameters are shown in Table 2.4. Overall, the model is able to closely fit the data, especially given the absence of any exogenous utility shocks in the model.

Figure 2.2: Fit of the 1960s Cohort

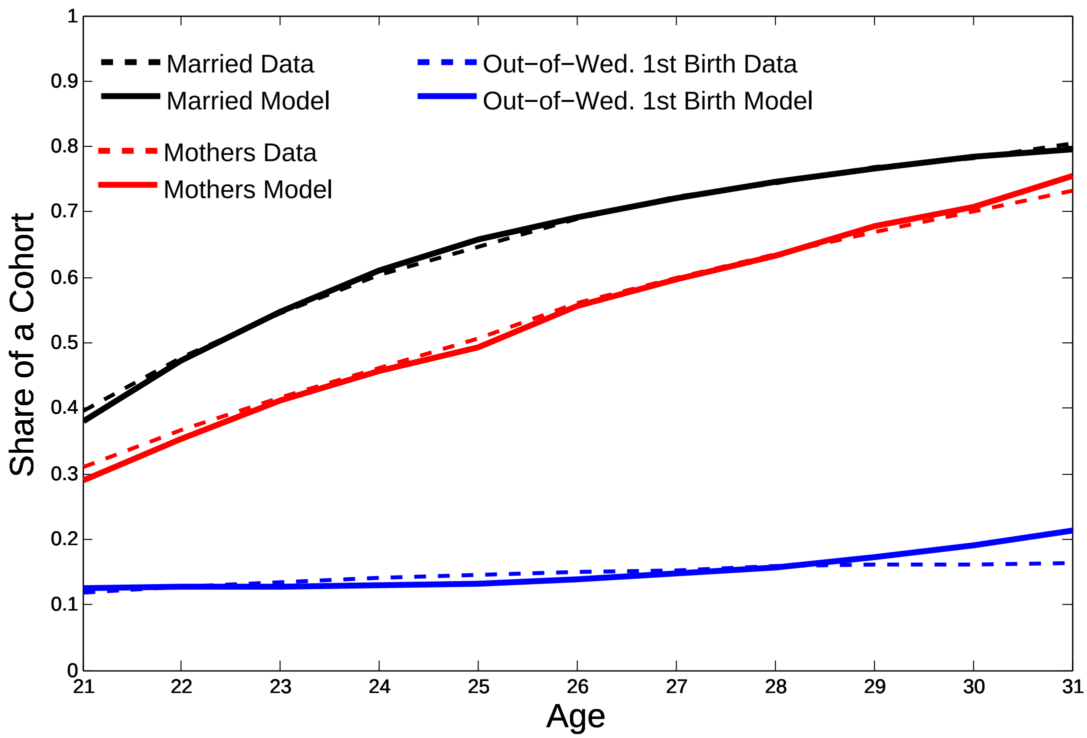

2.6.3 Simulation and Counterfactual Experiments

2.6.3.1 Simulation of the 1980s Cohort

In order to assess the overall effect of changes in the income process on marriage and fertility timing, I simulate the demographic behavior of the later NLSY cohort by changing the income process to one estimated from NLSY97.

Table 2.4: Parameters Estimated by Fitting the 1980s Cohort

σ_c	σ_k	ψ	η_m	γ
risk aversion	r. a. in K	C/N	monetary cost	ec. scale
0.27	0.98	4.34	4,890 (in 2012$)	0.73

Figure 2.3: Change in Marriage Behavior: Model vs Data

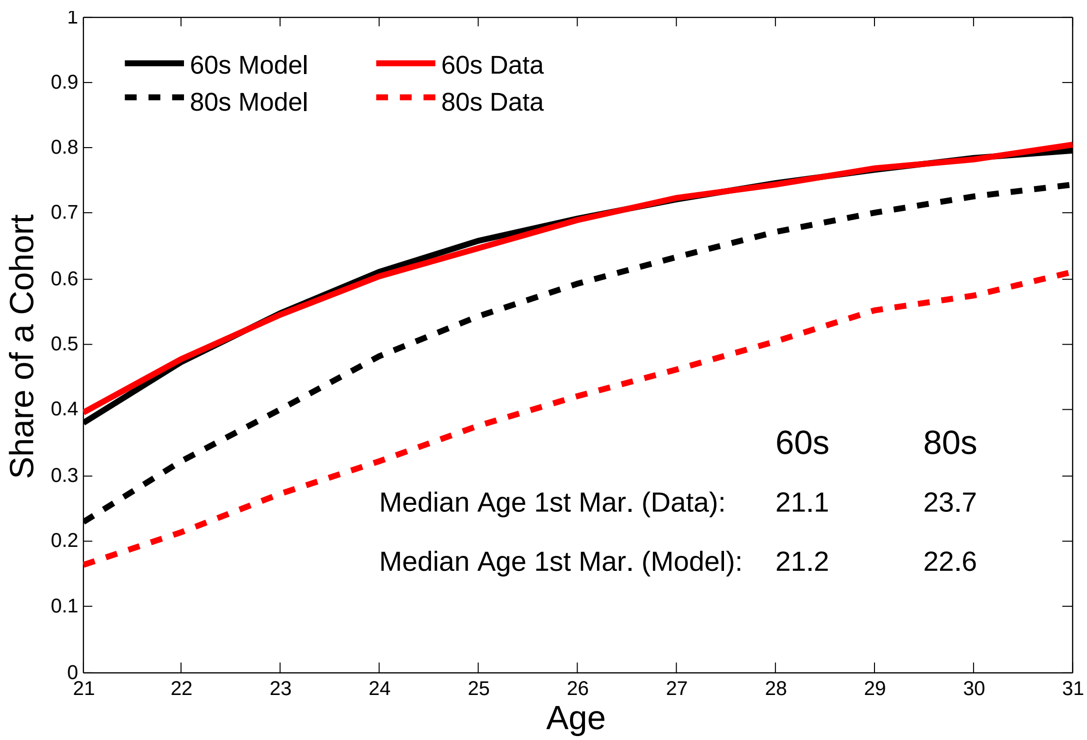

Figures 2.3, 2.4 and 2.5 compare the changes produced by the model with those observed in the data in age-specific shares of married women, mothers and never-married mothers. Changes in the income process account for 42% and 40% of the observed changes in age-specific shares of married women and mothers, respec-

tively. These percentages are calculated as a ratio of the area between the black lines (decrease produced by the model) and the area between the red lines (decrease in the age-specific shares observed in the data); if the produced direction of change is opposite to the observed in the data, then the ratio is negative. The produced simulation matches the qualitative fact of greater delay in marriage relative to that in fertility, which in turn produces an increase in single motherhood.

Figure 2.4: Change in Fertility Behavior: Model vs Data

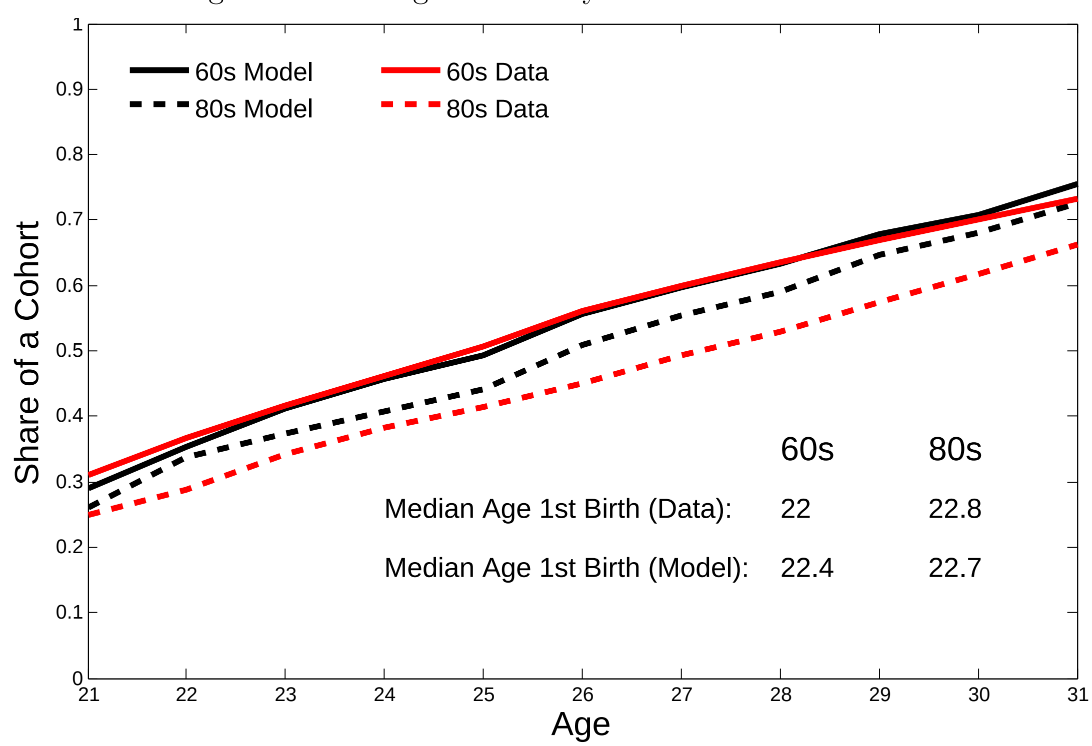

Figure 2.5: Change in Single Motherhood: Model vs Data

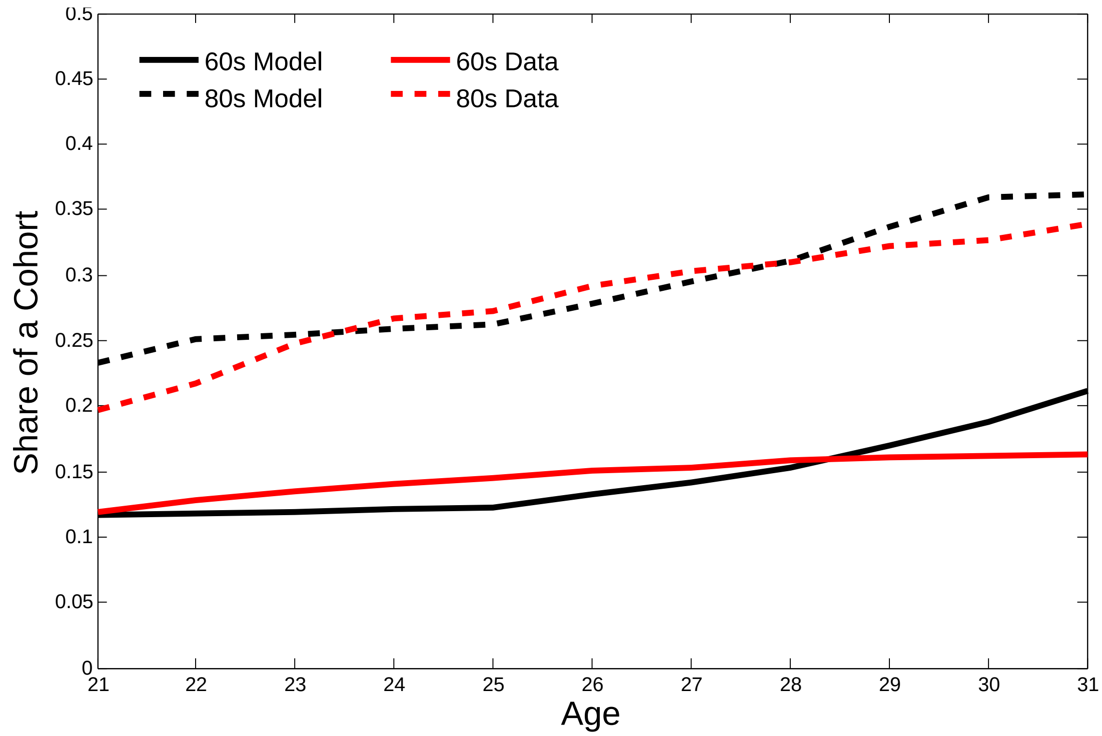

2.6.3.2 Experiment 1. Income Inequality

The modeling structure of my income process allows me to simulate secular effects of income inequality and income mobility. Such experiments can also serve as validity checks by matching the theoretical predictions of my model.

According to the logic of my model, a secular increase in income inequality should delay marriage while delaying fertility to an even greater extent. The reason for the latter is that income inequality delays fertility both directly – through the increase in income risk – and indirectly through the decrease in marriage rates. Figure 2.6 shows the prediction of the model if only wage distributions W_t^g are changed. For such a case, simulation accounts for 35% of the observed change in marriage and

335% of the observed changes in fertility. Since having a child becomes too risky when married, let alone when single, the predicted share of never-married mothers decreases and actually becomes 0 for the first several years of the life-cycle. Note that by updating the wage distributions I do not only change the spread of wages but also their relative levels. One concern could be that the observed in Figure 2.6 changes are due to change in the gender-wage gap rather than change in inequality. I address this concern in Experiment 3.

Figure 2.6: Experiment 1. Only Inequality has Changed

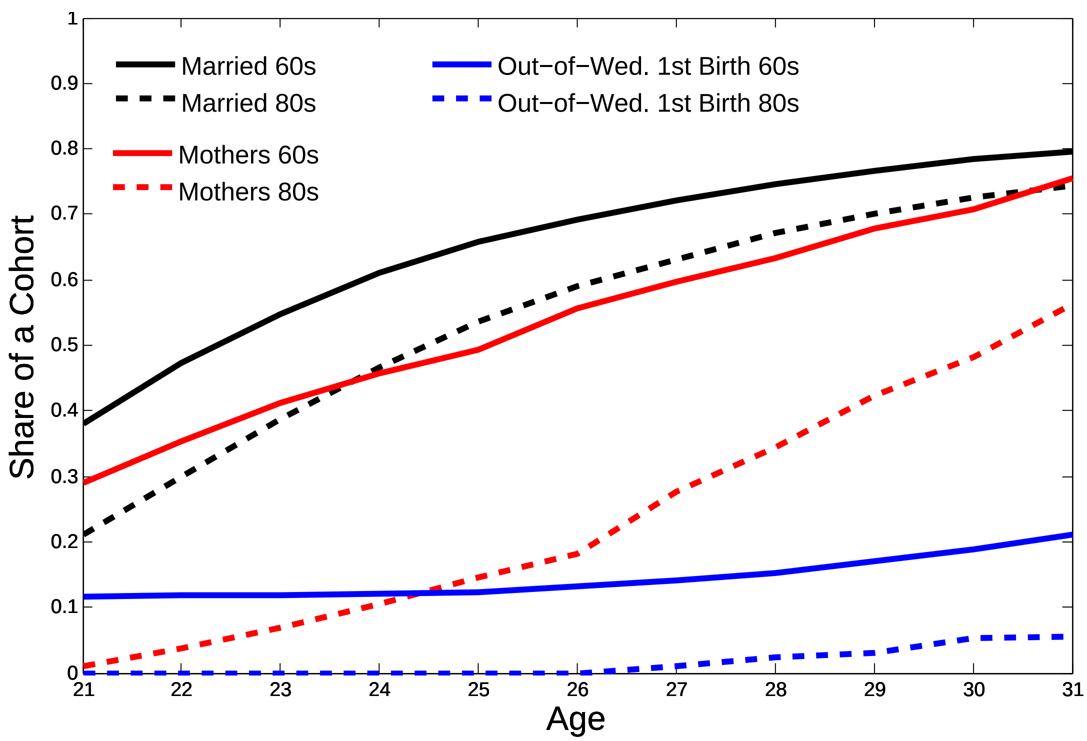

2.6.3.3 Experiment 2. Income Mobility

A decrease in income mobility is expected to accelerate fertility because having a child becomes less risky; at the same time, it produces a delay in marriage, since agents would value the associated insurance less. Figure 2.7 presents a counterfactual experiment in which I only update the transition matrices Π_t^g. In this case, the model accounts for 14% and -124% of the observed changes in marriage and fertility behavior, respectively.

Hence, the results of the two experiments confirm the theoretical predictions about the effects of income inequality and income mobility on the respective timings of marriage and fertility.

In addition, the two experiments suggest a negative interaction between an increase in inequality and a decrease in mobility with respect to their effects on marriage timing; the interaction is positive with respect to their effects on timing of fertility. Taken separately, the two changes in income process produce a larger delay in marriage than they will when they change together (35% + 14% > 40%). The intuition behind this is that when inequality is low, more types of agents would agree to marry each other in principle – that is, if the outside option is to remain single forever. In this case, a decrease in the desirability of marital insurance can affect the marital strategies of many types of agents. When income inequality is high, different types are less compatible with each other even in principle (even if the male agent were the last man on earth). Here, a decrease in income risk produced by changing mobility would not influence some of the marital choices, since marriage is already rejected in those cases.

Contrary to the negative interaction between the two changes in the income process in terms of their effects on marriage, higher income inequality makes any decrease in income mobility more potent (335% - 124% > 41%). When income inequality is high, there are more single women, who are more vulnerable to risk and hence more responsive to changes in income risk produced by a decrease in income mobility.

Figure 2.7: Experiment 2. Only Mobility has Changed

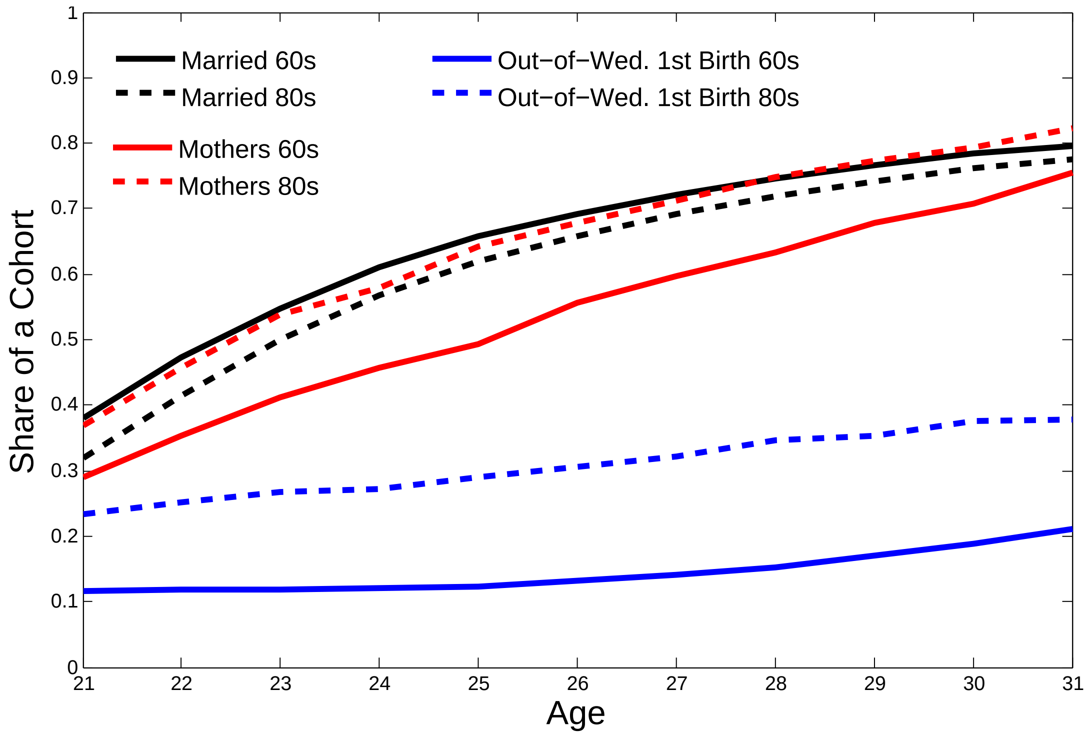

2.6.3.4 Experiment 3. Gender-Wage Gap

Finally, in order to make the model comparable with Regalia, Rios-Rull & Short (RRS) (2011), I want to assess the impact of changes in the gender-wage gap (GWG) on the marriage and fertility timing of individuals. In their model, RRS use a stylized income process which does not account for heterogeneity, so when the authors mark up a female wage it indeed only produces decrease in the GWG. This is not the case with the income process that I consider in this paper. If I multiply all female wages by a coefficient, it results in both, females becoming richer and the female distribution of wages moving relative to the male distribution of wages. Figure 2.8 shows the effects of scaling up all female wages in the 1960s cohort in such a way that the ratio of the median wages is equal to the one for the 1980s cohort. RRS theorized that an increase in relative female well-being would increase her reservation value and, therefore, she would spend more time searching for a partner. This is contrary to what one can see in Figure 2.8. The main reason behind that is that when the relative distributions of wages of male and female agents change, it could produce an increase in the overlap of marriageable types. Consider for example three female and male types of agents with the initial wages $w^f \in \{10, 45, 90\}$ and $w^m \in \{20, 60, 100\}$, respectively. For simplicity, I also assume that two agents would agree to marry once their wages differ by less than a factor of 2. So in the initial setting, pairs (w^f_1, w^m_2), (w^f_1, w^m_3), (w^f_2, w^m_1), (w^f_2, w^m_3) would not agree to marry each other. Once the GWG decreases and female wages are multiplied by, say, 2, a pair of agents (w^f_2, w^m_3) would agree to marry each other and marriage rates would rise. This is evidently the case in experiment 3 where the agents marry faster than before. On the other hand, since the price of female time increases, they would have fewer children, and especially out-of-wedlock where husbands can not alleviate part of the time cost.

Figure 2.8: Experiment 3. Only Gander-Wage Gap has Changed

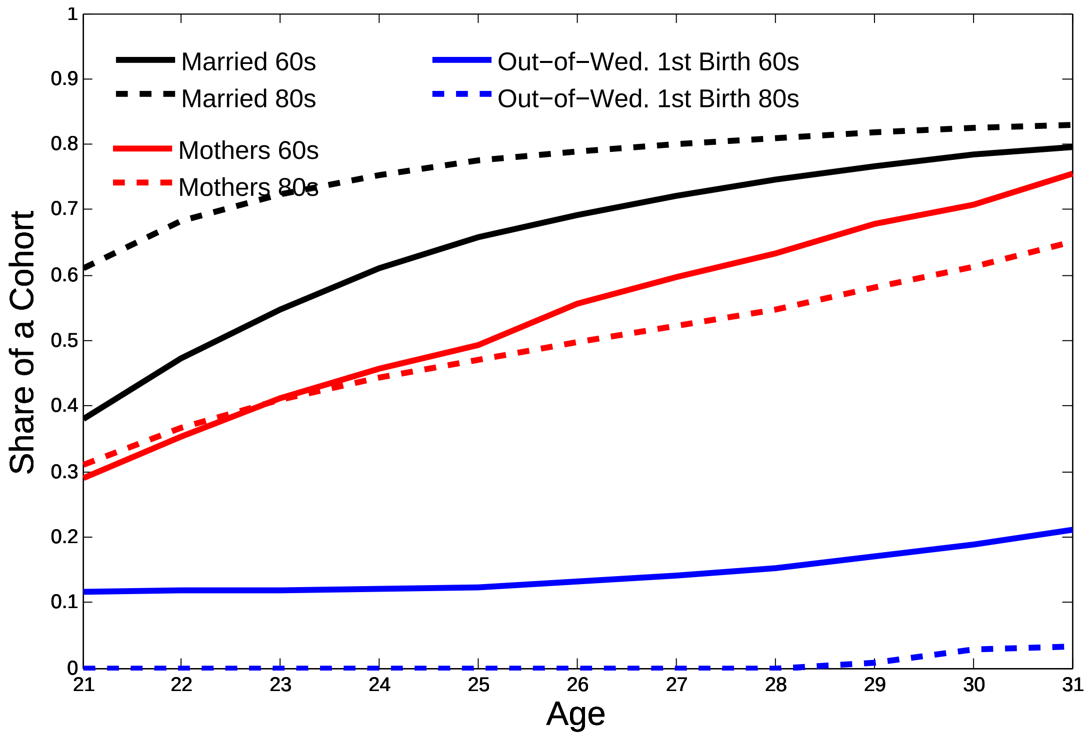

2.7 Conclusion

Large demographic changes taking place over the last half-century have attracted much attention in both the academic literature and popular media. In this paper, I focus on four trends characterizing the evolution of the respective timings of first births and first marriages. Although there are substantial literatures devoted to delay in marriage, postponement of childbearing and increases in single motherhood, most researchers treat these phenomena separately. I point out that the three trends are bonded together by the process of a decrease in the Gap between the timings of fertility and marriage. This phenomenon emanates from the ongoing trend of delay in childbearing and an even larger delay in marriage. As a result of this, the probability that a woman will have her first child prior to marriage increases. As argued by Ellwood & Jencks (2004), moreover, most of the recent increase in out-of-wedlock fertility is actually driven by an increase in the share of out-of-wedlock first births.

I show how studying these related trends separately can be misleading: the mechanisms proposed to explain one of the changes often fail to produce the qualitative facts related to the others. For example, a number of theories produce a simultaneous delay in marriage and fertility by elaborating on Becker's limits to household specialization hypothesis. In such theories, some exogenous innovation typically delays fertility by decreasing the willingness of individuals to have children. Because having children is seen as part of the value of being married, the undesirability of children also delays marriage. Examples of such innovations are an increase in income volatility as in Santos & Weiss (2012, 2016), an increase in the general income level in conjunction with a decrease in the gender-wage gap as in Regalia et al. (2008) and the invention of orally-administered birth control medication as in Goldin & Katz (2002). Although such mechanisms are, indeed, able to produce delays in both marriage and fertility, they naturally fail to account for the decrease in the Gap and increase in single motherhood. After all, by the logic of these mechanisms, the unwillingness to have children was the original reason for delaying marriage. Since out-of-wedlock fertility is disproportionately concentrated among low-educated and black women, researchers often propose race or socio-economic class-specific mechanisms, such as the hypotheses of black incarceration, unmarriageable men, etc. I show that the decrease in the Gap is relevant to all major socio-economic groups of women, and, therefore, group-specific mechanisms may contribute to but cannot explain the observed trends for the female population in general.

I then propose an explanation using two exogenous changes in the income process in order to explain all of the qualitative facts related to the respective marriage and fertility timings. First, an increase in income inequality delays both marriage and fertility; second, a decrease in income mobility decreases income risk and further delays marriage while accelerating fertility, which produces a decrease in the Gap and an increase in single motherhood.

In order to determine the quantitative importance of my mechanisms, I build an equilibrium model of marriage and fertility, matching it with data on women born around 1960 (NLSY79). Although the benchmark model is relatively simple, it is able to produce a good fit of the data. After calibrating the model to match the behavior of the 1960 cohort, I update the income process estimates and find that, according to my model, they account for 42% and 40% of the changes in marriage and fertility timing, respectively, between the two NLSY cohorts.

A better understanding of the driving forces behind the demographic changes of interest is important because of their impact on the lives and well-being of individuals. For example, the empirical literature finds that being born to an unmarried mother brings certain disadvantages, including lower human capital accumulation and adverse health outcomes.[12] Marriage also plays an important role in the general

[12]McLanahan & Sandefur (2009) document various adverse effects of being raised in a single-parent family on human capital formation and schooling outcomes of children. They also show that cohabitation is not a substitute for marriage in this respect. Waldfogel et.

standard of living of individuals, savings behavior and even the propensity to buy a house.[13]

The purpose of the benchmark model was to demonstrate the effects of income inequality and income mobility on the respective timings of marriage and fertility. Therefore, I emphasized simplicity. In order to test some policy applications of my model, I intend to first enrich the benchmark model with several related margins, such as investment into child quality and endogenous labor supply.

Once those elements are integrated in the model, it can be used to study the self-induction and propagation of income redistribution policies. Most redistribution policies affect both income inequality and income mobility, which would in turn affect both population growth rates and investment in in turn human capital of children. Such policies could therefore play an important role in determining the income inequality of future generations. For example, lower human capital accumulation by children raised out-of-wedlock tends to be translated into an increase in income inequality once those children enter the labor market. Now let's consider a redistributive policy, such as the progressive labor income tax: on the one hand, it decreases income inequality and thus promotes marriage; on the other hand, it also

al. (2010) provide an extensive review of the literature documenting the effect of family structure on the health outcomes of children.

[13]Due to the existence of household consumption economies of scale, two individuals living together can afford a higher level of consumption than they could living separately. This effect of marriage can be partially replicated by cohabitation, but it differs from marriage in that it does not provide commitment. For example, a cohabiting couple is less likely to take a mortgage, which requires a long-term commitment. Some recent publications include Browning, Chiappori & Lewbel (2013) – economies of scale, Knoll, Tamborini & Whitman (2012) – savings behavior and Fisher & Gervais (2011) – home ownership.

decreases income mobility, which decreases the Gap and increases single motherhood. Thus, the effect of the redistributive policy on income inequality for the future generation is ambiguous and can be assessed quantitatively with the use of an extended version of my model.

I also want to use this version of the model to quantify the (often unintended) effects of different public policies on marriage and fertility timings. While some policies, like progressive income taxation, affect both inequality and mobility, others have an impact on one or the other first and foremost. For example, anti-discrimination policies and worker retraining programs would mainly contribute to an increase in income mobility, while trade unions would presumably have an opposite effect. Policies regulating wealth accumulation and inheritance, on the other hand, would primarily influence inequality. The intuition provided in this paper allows for the evaluation of the effects of these and other polices from a new point of view.

It also appeals to use the presented framework to study inter-generational income mobility and economic growth. For example, De la Croix & Doepke (2003) argue that high income inequality may undermine growth because it leads to under-investment in the education and human capital of children. Hence, redistributive policies may be desirable. I expect to come to the same conclusion, but for a different reason; according to my work, a decrease in income inequality will increase the human capital of the next generation through the higher parental investments of two-parent households.

CHAPTER 3
ACCOUNTING FOR MARRIAGE IN A MODEL OF INDIRECT INCOME UNCERTAINTY INFERENCE

3.1 Introduction

Income uncertainty is a crucial part of the income process due to its major impact on the inter-temporal choices of individuals. There are two general approaches to estimating it. The first one is to directly ask individuals about their perceived income risks, but surveys that include questions about income uncertainty are very scarce and usually do not go far enough in time.[1] An alternative, and relatively more recent approach is an indirect earnings uncertainty inference (IUI).[2] The core idea of this method is that since the inter-temporal decisions of an individual are made in accordance with the level of uncertainty she faces (experienced uncertainty), an econometrician can infer that uncertainty by observing her economic choices.

This paper aims to advance the indirect earnings uncertainty inference (IUI) literature by accounting for marital insurance in a structural model. The common shortcoming of the currently existing IUI models is that they tend to ignore marital state. For example, Guvenen (2007) employs a mixed sample of single and pseudo-single individuals i.e. married, who are assigned a per-adult equivalent of consump-

[1]Guiso, Japelli & Pistaferri (1998) use Bank of Italy's Survey of Household Income and Wealth; Dominitz (1998) and Manski (2004) use Survey of Economic Expectations; Ramos & Schuler (2006) study British Household Panel Survey; See also Hurd (2008), Attanasio & Augsburg (2015). Questions regarding the subjective earnings expectations can also be found in the Health and Retirement Study

[2]Guvenen (2007), Guvenen & Smith (2014)

tion and treated as singles. Guvenen & Smith (2014) use married couples as a unit of their analysis. The notion that the spouses in two-earner families enjoy partial income insurance is well established in the literature.[3] The presence of marital earnings insurance creates a discrepancy between the underlying uncertainty — a fundamental property of the income process and the experienced uncertainty, the one that individuals face given the insurance they have. As a result, not accounting for marriage in a pseudo-single approach produces bias in the estimates, while using married couples as a unit of analysis makes results sensitive to changes in the environment (changes in marriage rates, assortative mating, institutions) while also creating a selection bias.

I begin with establishing the marital insurance bias in a stylized analytical model which resembles Guvenen (2007) and Guvenen & Smith (2014) in a way that it uses a simplified version of the heterogeneous income profiles (HIP) process as its foundation.[4] I show that the resulting marital insurance bias may be both positive or negative depending on the co-movement of spousal earnings. After that, I build a structural IUI model which accounts for the marital state of the individuals. The application of this model to simulated data reveals a significant and systematic bias of the pseudo-single approach.

[3]See Hess (2004), Schneider & Reich (2014), Sopchokchai (2016)

[4]The HIP process used in Guvenen (2007) and Guvenen & Smith (2014) has a rather complex parametric structure which roughly consists of the usual AR(1) process and an individual-specific linear time trend. In order to match the life-cycle properties of consumption inequality this process also requires a learning process so that the agents don't know their true life-cycle trend in the beginning of their careers. The stylized earnings process utilized in section 3.2 does not fully capture the HIP process, however, it does not affect the validity of the conclusions with regard to marital insurance bias.

As a secondary contribution of this project, I introduce a non-parametric income process similar to that in De Nardi, Fella & Paz-Pardo (2016) (DFP) but modified to handle private information about the agent's future earnings. Besides the usual benefits of a non-parametric income process, the major advantage of the one introduced here is that it allows to perform IUI with short panel data.[5] One of the most important problems of the IUI models is to correctly pin down the expectations that an individual has about her future earnings. If we observe a risk-averse individual whose earnings today were \$1000 and she decided to save \$500 it could mean that she expects to earn \$0 in future and is 100% sure about that or that she expects to earn \$1000 in future but is very uncertain. The usual strategy in the models which use HIP income process is to estimate the life-cycle earnings trajectory of an individual from the long panel of data and then assume that her beliefs should be distributed around that trend. The income process introduced in this paper allows to infer individuals' expectations from the population dynamics as I demonstrate in section 3.3.2.2.

3.2 The Analytical Model

In this section, I consider two stylized models — one which accounts for marriage and one that does not. I then compare the indirect inference results obtained from the two models and discuss the results.

[5]As of today, no direct comparison of the HIP process and the one from DFP is made. However, as it is shown in DFP, their non-parametric process is much simpler while is still able to capture the first four data moments with high precision.

3.2.1 Single Agent's Problem

Consider the following setting. All agents live for two periods. Each agent is characterized by her earnings profile which consists of her first period earnings $w_{1,i}$ and second period earnings $\hat{w}_{2,i} = w_{2,i} + \hat{\epsilon}_i$ where $\epsilon_i \sim \mathcal{N}(0, \sigma_i)$. The only decision of the agent is her consumption/savings choice in period 1. For simplicity, I assume that the agent's preferences over consumption are represented by exponential utility. Given that, the single agent's problem is:

$$V_1(w_{1,i}, w_{2,i}, \sigma_i) \equiv \max_{S_i} -e^{C_1} + \beta \mathbb{E}_{\epsilon_i} \left[-e^{C_2} \right] \tag{3.1}$$

s.t.

$$C_1 + S_i \leq w_{1,i}$$

$$C_2 \leq w_{2,i} + (1+r)S_i$$

The indirect inference problem then can be defined as following: the agent knows her $w_{1,i}, w_{2,i}$ and σ_i, while an econometrician only observes $w_{1,i}, S_i$ and knows $w_{2,i}$ and her savings decision S and wants to infer σ_i. The goal of the econometrician is to find the degree of earnings uncertainty which in this model is represented by σ_i.[6]

Substituting budget constraints and integrating (3.1) over ϵ_i gives:

[6]The common problem in such type of models is that the econometrician can only observe $w_{1,i}, S_i$ and $\hat{w}_{2,i} = w_{2,i} + \hat{\epsilon}_i$ — the realization of the agent's second period earnings. "Knowledge" of $w_{2,i}$ — agent's expected future wage requires additional structure and assumptions. I will elaborate on this point in subsection 3.3.2.2.

$$V_1(w_{1,i}, w_{2,i}, \sigma_i) \equiv \max_S -e^{-\theta(w_{1,i}-S)} + \beta e^{-\theta(w_{2,i}+(1+r)S)+\theta^2\sigma_i^2/2}$$

Then the f.o.c. of the single agent's problem are:

$$\sigma_i^2 = \frac{2(2+r)S}{\theta} + \frac{2(w_{2,i} - w_{1,i})}{\theta} - \frac{2log(\beta)}{\theta^2} \tag{3.2}$$

Equation (3.1) contains only variables that by assumption are known to the econometrician, and therefore this equation can be used for indirect inference. I will denote results obtained with the use of equation (3.2) as $\hat{\sigma}_i^2(w_{1,i}, w_{2,i}, S_i)$.

3.2.2 Married Agent's Problem

Now let's turn to the married agent's problem. I assume the same basic income process where each individual earns $w_{1,i}$ and $\hat{w}_{2,i} = w_{2,i} + \hat{\epsilon}_i, \quad \epsilon_i \sim \mathcal{N}(0, \sigma_i)$ at the first and second period, respectively. The correlation between the spousal earnings is ρ. I also assume that spouses pool their incomes and split their consumption according to the Kalai-Smorodinsky bargaining process with exogenously fixed equal bargaining powers. This setting is equivalent to each spouse solving the following problem:

$$VM_1(W_{1,i,j}, W_{2,i,j}, \sigma_i, \sigma_j) \equiv \max_{S_{i,j}} U(C_1/2) + \beta \mathbb{E}_{\epsilon_i}\left[\mathbb{E}_{\epsilon_j}\left[U(C_2/2)\right]\right] \tag{3.3}$$

s.t.

$$C_1 + 2S_{i,j} \leq W_{1,i,j}$$

$$C_2 \leq W_{2,i,j} + (1+r)2S_{i,j}$$

where $W_{t,i,j} = w_{t,i} + w_{t,j}$; C — total family consumption; and $S_{i,j}$ — savings per spouse.

Given the functional forms f.o.c are:

$$\sigma_i^2(1+\rho^2) + 2\rho\sigma_i\sigma_j + \sigma_j^2 = \frac{8(2+r)S_{i,j}}{\theta} + \frac{8(0.5W_{2,i,j} - 0.5W_{1,i,j})}{\theta} - \frac{8log(\beta)}{\theta^2} \quad (3.4)$$

Equation (3.4), as its version for singles equation (3.2), contains only variables that are assumed to be known by the econometrician. However, equation (3.4) is not enough to infer the underlying uncertainties, since it produces the quadratic line $(\sigma_i, \sigma_j(\sigma_i))$. Basically, equation (3.4) means that one either needs additional information or additional assumption in order to identify the exact values of the spouse's uncertainties. I will denote the uncertainties inferred with equation (3.4) as $\tilde{\sigma}_i^2(W_{1,i,j}, W_{2,i,j}, S_{i,j})$.

3.2.3 Indirect Inference Results Comparison

The two most commonly used assumptions that are made in papers that do not include marriage in the structural models are a) taking married couples as a unit of analysis and b) a pseudo-single approach i.e. assigning each spouse a per-adult equivalent of savings and treating them as singles. The first approach, by construction, results in selection bias since one cannot include both married and single individuals in the sample in this case. This approach also only estimates the level of uncertainty for marital unions which does not allow us to learn about individual income risks. In this section, I compare the results of the indirect uncertainty inference made with the model which accounts for marriage with those of the one which uses a pseudo-single approach.

According to a pseudo-single approach, a couple characterized by $w_{1,i}, w_{1,j}, w_{2,i}, w_{2,j}, S_{i,j}$ will be treated as two single individuals each of whom have savings $S_{i,j}$ since it is a per-spouse amount of savings. According to (3.2), the uncertainties that would be inferred are:

$$\hat{\sigma}_i^2(w_{1,i}, w_{2,i}, S_{i,j}) = \frac{2(2+r)S_{i,j}}{\theta} + \frac{2(w_{2,i} - w_{1,i})}{\theta} - \frac{2log(\beta)}{\theta^2} \tag{3.5}$$
$$\hat{\sigma}_j^2(w_{1,j}, w_{2,j}, S_{i,j}) = \frac{2(2+r)S_{i,j}}{\theta} + \frac{2(w_{2,j} - w_{1,j})}{\theta} - \frac{2log(\beta)}{\theta^2}$$

While the true uncertainties according to (3.4) are:

$$\tilde{\sigma}_i^2(1+\rho^2) + 2\rho\tilde{\sigma}_i\tilde{\sigma}_j + \tilde{\sigma}_j^2 = \frac{8(2+r)S_{i,j}}{\theta} + \frac{8(0.5W_{2,i,j} - 0.5W_{1,i,j})}{\theta} - \frac{8log(\beta)}{\theta^2} \tag{3.6}$$

Two facts are evident from the comparison of equations in (3.5) and (3.6). The first is that it is possible that $\hat{\sigma}_i^2, \hat{\sigma}_j^2$ will satisfy (3.6), but most of the time, magnitudes will be different. Secondly, as it has been mentioned before, equation (3.6) requires additional assumptions, while (3.5) does not. This happens because the pseudo-single inference implicitly assumes a point on the (σ_i, σ_j)-plain, while (3.6) gives a quadratic line on that plain.

In order to find the direction of the bias and its magnitude, I need to make an assumption about the relationship of σ_i and σ_j. For the purposes of computational simplicity I assume $\sigma_i = \sigma_j = \sigma$. Then (3.6) will give us:

$$\sigma^2 = \tilde{\sigma}_i^2 = \tilde{\sigma}_j^2 = \frac{8(2+r)S_{i,j}}{[\rho^2+2\rho+2]\theta} + \frac{8(0.5W_{2,i,j} - 0.5W_{1,i,j})}{[\rho^2+2\rho+2]\theta} - \frac{8log(\beta)}{[\rho^2+2\rho+2]\theta^2} \tag{3.7}$$

combining this with (3.5) gives:

$$\sigma^2 = \frac{2(\hat{\sigma}_i + \hat{\sigma}_j)}{[\rho^2+2\rho+2]} \tag{3.8}$$

and the average bias is :

$$\begin{aligned} Average\ Bias &= \frac{\hat{\sigma}_i^2 + \hat{\sigma}_j^2 - 2\sigma^2}{2} = \sigma^2 \frac{[\rho^2+2\rho-2]}{4} \\ \Rightarrow Average\ Bias &\in [-\frac{3}{4}\sigma^2, \frac{1}{4}\sigma^2] \end{aligned} \tag{3.9}$$

Equation (3.9) suggests that depending on the correlation between spouses' incomes ρ, the average bias can be both positive or negative. It is more likely to be negative since even small positive correlation between spousal incomes will still provide them with partial insurance, so that the pseudo-single approach will underestimate the uncertainty that individuals face.

According to D. Hyslop (2001), the average correlation between spousal earnings is about $\rho = 0.35$. Then a ballpark value of the average bias is $-0.29\sigma^2$ — the underestimation of uncertainty by a third of its value.

3.3 Structural Model

In this section I introduce the structural model which accounts for the marital state. I then use this model to quantify the marital bias by comparing the indirect inference estimates produced by this model with those produced by the model which treats married individuals as pseudo-singles. The intention of this paper is to show the importance of accounting for marriage in the IUI models. For this purpose, I apply both models to the simulated data which allows me to compare inferred results with the true level of uncertainty that was used to generate the data sample. Producing the corrected estimates from the data is out of scope of the current project and is left for future work.

The secondary contribution of this paper is modifying the non-parametric income process similar to one in De Nardi, Fella & Paz-Pardo (2016), such that it can handle the unobserved uncertainty and, therefore, be used for the IUI. The major advantage of this income process, relative to the typically assumed in the IUI models heterogeneous income profiles process (HIP), is that the latter requires significantly longer panel data in order to produce IUI estimates. The reason behind that is discussed in subsection 3.3.2.2.

3.3.1 Income Process

I assume a non-parametric income process similar to one in the De Nardi, Fella & Paz-Pardo (2016). The model I propose is different from theirs in that I introduce an unobservable (to an econometrician) private information about agent's future earnings.

The benefit of this income process with regard to the indirect income uncertainty inference is that it allows to use much shorter panels of data than the usual income processes (e.g. RIP or HIP). I will elaborate on this notion later.

3.3.1.1 General Income Process

The income process has the following structure. In every period, an agent can have one of the N period-specific wages $w_t \in W_t \equiv \{w_{t,1}, ..., w_{t,N}\}$. Empirically, these wages are estimated as mean wages of the N age-specific quantiles, and the measure of agents with a given wage at any period is therefore $1/N$.

All wages evolve according to a period-specific transition matrix:

$$\Pi_t \equiv \begin{array}{|c|c|c|c|}
\hline
 & w_{t+1,1} & \cdots & w_{t+1,2} \\
\hline
w_{t,1} & \pi_{t,1,1} & \cdots & \pi_{t,1,N} \\
\hline
\vdots & \vdots & \ddots & \vdots \\
\hline
w_{t,N} & \pi_{t,N,1} & \cdots & \pi_{t,N,N} \\
\hline
\end{array}$$

where $\pi_{t,i,i'} = Pr(w_{t+1} = w_{t+1,i'} | w_t = w_{t,i})$ is the probability that an agent will have a wage $w^g_{t+1,i'}$ in period $t+1$, conditional on having a wage $w^g_{t,i}$ in period t.

I assume that both, possible wage arrays and transition matrices $\{W_t, \Pi_t\}_1^T$ are common knowledge and are also observable by an econometrician.

3.3.1.2 Private Information

In addition to the general income process, in the beginning of every period, each agent receives a private (from an econometrician) signal $l = \hat{l} \in L \equiv \{1, ..., N\}$ which informs her that she is more likely to transit to wage $w_{t+1} = w_{t+1,\hat{l}}$ in the next period. If the quality of that signal is $\lambda \in [0, 1]$, then her probability of receiving wage $w_{t+1,\hat{l}}$ is $\pi_{t,i,\hat{l}}(l) = (1 - \lambda)\pi_{t,i,\hat{l}} + \lambda$. Probability of transiting to any other wage decreases proportionally:

$$\Pi_t(w_{t,i}, \hat{l}) = \{(1 - \lambda_{t,i})\pi_{t,i,1}\ , \ \dots\ , \ (1 - \lambda_{t,i})\pi_{t,i,\hat{l}} + \lambda_{t,i}\ , \ \dots\ , \ (1 - \lambda_{t,i})\pi_{t,i,N}\}$$

In order for the individual transition probabilities across the agents and the signals they receive to be consistent with the general population transition matrices Π_t, two conditions should apply. First, all agents sharing the same current wage $w_{t,i}$ must have the same quality of information λ, while it may vary across the agents who receive different wages in this period and over periods i.e. $\lambda \equiv \lambda(i, t)$. The second condition dictates that the ex-ante probability of receiving signal $l = j$ is $Pr(l = j|w_t = w_{t,i}) = \pi_{t,i,j}$.

Note that neither signal l nor quality of that signal λ is observable by an econometrician.

3.3.2 Income Uncertainty Inference Procedure

I illustrate the IUI procedure for models which use the above income process with a two-period single agent's problem where the quality of information λ is the same across all agents. In order to test the proposed IUI procedure, I apply it to the simulated data.

3.3.2.1 Single Agent's Problem

The value of being a single agent with wage $w_1 = w_{1,i}$ and signal $l = \hat{l}$ in the first period is:

$$\begin{aligned} VS(w_{1,i}, \hat{l}) \equiv & \max_S U(w_{1,i} - S) \\ & + \beta(\lambda + (1-\lambda)\pi_{i,\hat{l}})U(w_{2,\hat{l}} + (1+r)S) \\ & + \beta \sum_{l^* \in L \setminus \{\hat{l}\}} ((1-\lambda)\pi_{i,l^*})U(w_{2,l^*} + (1+r)S) \end{aligned} \tag{3.10}$$

where S denotes savings.

The first order condition of (3.10) can be written as:

$$\lambda = 1 - \frac{\frac{U'(w_{1,i}-S)}{(1+r)\beta} - U'(w_{2,\hat{l}} + (1+r)S)}{\sum_{l \in L} [\pi_{i,l} U'(w_{2,l} + (1+r)S)] - U'(w_{2,\hat{l}} + (1+r)S)} \tag{3.11}$$

Equation (3.11) may become more intuitive if one considers the extreme cases. When $\lambda = 1$ i.e. the agent is certain that her future wage will be $w_{2,\hat{l}}$, savings equate the inter-temporal marginal utilities of consumption. If, at the other extreme, $\lambda = 0$, the private signal is worthless and the term subtracted from numerator and denominator cancels out. Since the agent does not have any private information at this extreme, she chooses savings which equate the marginal utility of consumption in period 1 and the expected marginal utility of consumption in period 2 given the common knowledge transition matrix Π_t.

3.3.2.2 Indirect Inference

Equation (3.11) relates the degree of uncertainty (inverse of the quality of the private signal λ) with savings S. However, equation (3.11) is not enough for the indirect uncertainty inference, since from the econometrician's point of view, there are two unknowns in it — λ and l.

Consider a numerical example in Figure 3.1 which demonstrates the optimal savings of three agents with a current wage $w_{1,2}$ who received signals 1, 2 and 3. As the agent with the lowest signal becomes more certain that her earnings would fall, she saves more. At the same time, agents who received high signals save less as their certainty rises.

Figure 3.1 also demonstrates that equation (3.11) cannot be directly used to infer the degree of uncertainty. For example, if an econometrician observes an agent with no savings $S = 0$, it could be an agent who expected higher a wage ($l = 3$) while being less certain or an agent who was 100 % certain that her earnings would remain on the same level ($l = 2$).

This problem is common in the IUI models. In a model based on the heterogeneous income process, the same indeterminacy requires an econometrician to know the mean of the distribution of future earnings in order to estimate the variance. If the econometrician has a long-enough panel data, she can estimate the individual's long-run earnings trajectory and then assume that the agent's future earnings expectations follow that trajectory. This knowledge is costly in a sense that in order to pin down the life-cycle trajectory, one needs to have long-enough panel data-sets.

Figure 3.1: Optimal Savings Conditional on Private Information

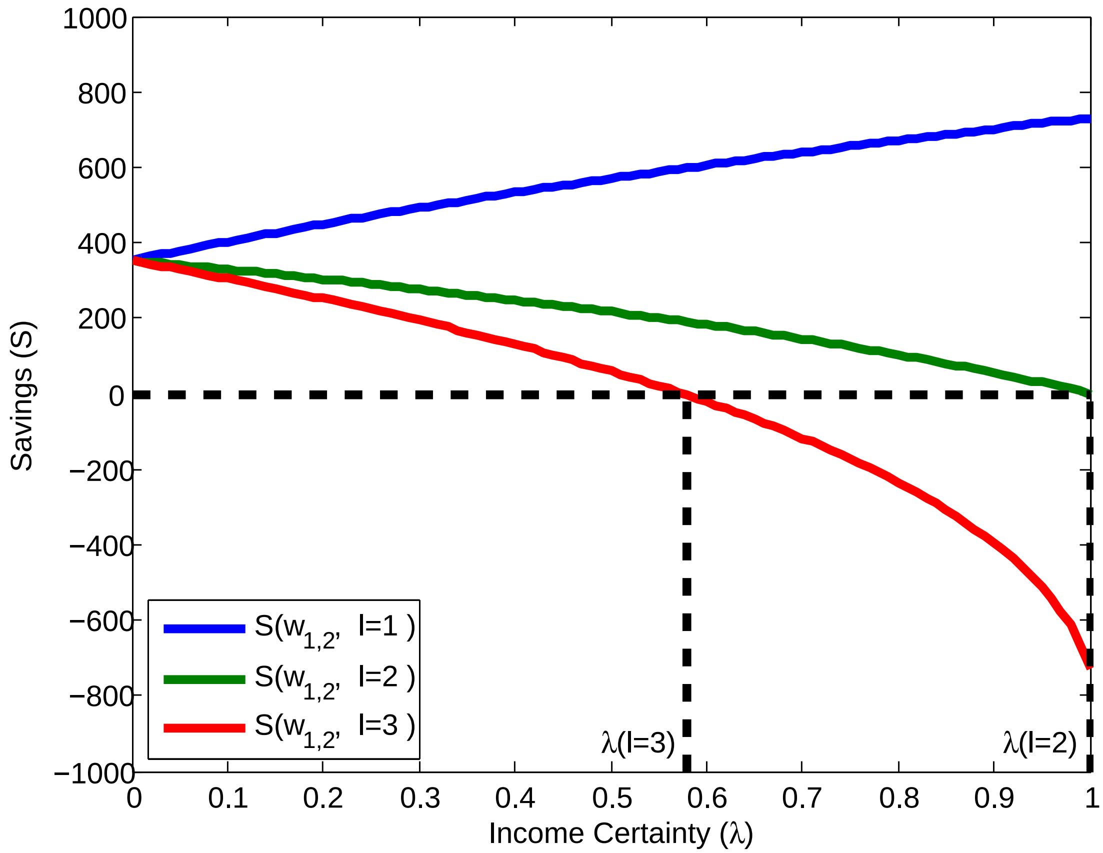

The structure of the income process that I introduce in this project allows the usage of short panel data-sets. Since this income process basically handles the whole distribution of agents, the econometrician can infer income uncertainty by matching the predicted and observed distributions, recall the example on Figure 3.1. Individual observation cannot determine the level of quality of the private information λ. However, if the rest of the observations with $w_1 = w_{1,2}$ concentrate around values of $S = -700, S = 700$, it would suggest that the true level is $\lambda = 1$. If most of the observations are found around the values $S = 200, S = 600$ we conclude that $\lambda = 0.57$. I describe the IUI algorithm more precisely in the next subsection.

3.3.2.3 Income Inference Algorithm

While discussing the example in Figure 3.1, I only used the information about an individual's initial wage w_1 and her savings level S. But besides that information, the econometrician also observes the next period earnings w_2. Given that an observation contains both w_1 and w_2 realizations for each agent, the econometrician can compute the discrete probability distribution over the possible values of signal l using transition matrix Π.

Suppose there is an individual with an earnings profile $w_{1,i}, w_{2,j}$. If we guess that $\lambda = \hat{\lambda}$, then the probability that such agent has received signal l is

$$Pr(l = j \in L | w_1 = w_{1,i}, w_2 = w_{2,j}) = \tag{3.12}$$
$$\frac{Pr(l = j | w_1 = w_{1,i}) Pr(w_2 = w_{2,j} | l = j, w_1 = w_{1,i})}{Pr(w_2 = w_{2,j} | w_1 = w_{1,i})}$$

Simplifying (3.12) yields:

$$Pr(l = j) = \hat{\lambda} + (1 - \hat{\lambda})\pi_{i,j} \tag{3.13}$$
$$Pr(l = j^*) = (1 - \hat{\lambda})\pi_{i,j^*} \qquad , \forall j^* \in L \setminus \{j\}$$

These probabilities would also allow us to put weights on the individual observations of savings.

Given wage arrays W_t, transition matrices Π_t, data on the first-period savings of each individual and equations (3.11) and (3.13), one can infer levels of income uncertainty with the following procedure :

1. Guess $\hat{\lambda}$.

2. For each individual, compute $\{Pr(l = 1|i, j), ..., Pr(l = N|i, j)\}$ — the probability distribution over the possible private signals that an agent with $w_1 = w_{1,i}$ and $w_2 = w_{1,j}$ could receive according to (3.13).

3. For each individual and for each of the possible signals, compute savings that we expect such an individual to make according to (3.11).

4. Compute the distance between the observed and expected distributions of savings.[7]

This procedure requires search over values of $\hat{\lambda}$ until the smallest distance in step 4 is found.

Figure 3.2 shows results of applying the above procedure to infer level of income uncertainty λ to the simulated data. In the simulation, I use a 10-quantile income process, 2000 observations and I repeat the simulation 200 times. Overall, the results depicted in Figure 3.2 suggest that the proposed indirect income uncertainty inference procedure coupled with the assumed income process works well.

3.3.3 Accounting for Marital Insurance

In this subsection, I first introduce the married agent's problem. After that, I compare the results of the indirect earnings uncertainty inference obtained with this structural model with those obtained with a single agent's model under the pseudo-

[7]There are several options to compare the observed and predicted results. As a result of the procedure for each individual, we will have a set of predicted savings values as well as probabilities with which those realizations should occur. One way is to compute the expected predicted savings and compare them with the observed ones. Another way would be to choose λ's in order to match the overall distribution. Finally, one could simply choose the single most probable signal on step 2. and compute savings only for it.

Figure 3.2: Inferred Uncertainty (Single Agents)

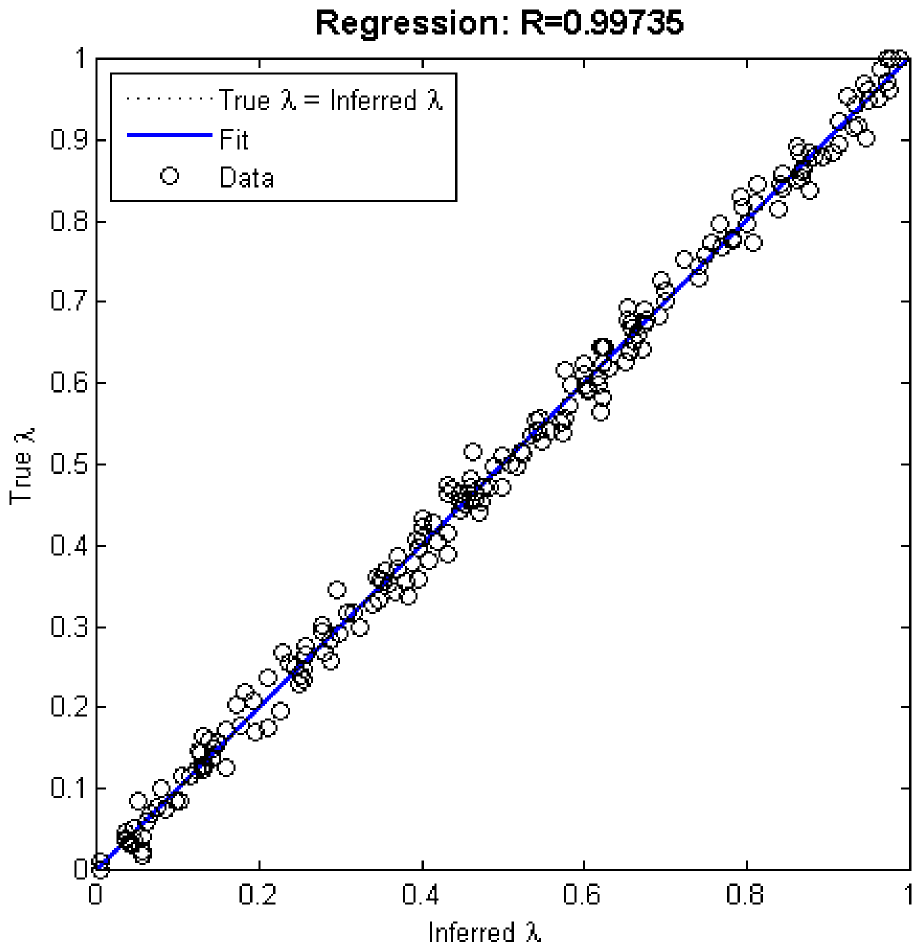

single approach.

3.3.3.1 The Married Agent's Problem

As in the analytical model, I assume that marriage is an exogenous state i.e. there is no decision to marry and no divorce. Spouses share the private information and make the consumption–savings decision according to the Kalai-Smorodinsky bargaining process with exogenously fixed equal bargaining powers. This setting is equivalent to each spouse in a couple with the first period wages $w_{1,f}, w_{1,m}$ and signals l_f, l_m solving the following problem:

$$\begin{aligned} VM(w^f_{1,i}, w^m_{1,j}, \hat{l}_f, \hat{l}_m) \equiv \max_S U & \left(\frac{w^f_{1,i} + w^m_{1,j} - 2S}{2} \right) \\ + \beta\lambda^2 U & \left(\frac{w_{2,\hat{l}_f} + w_{2,\hat{l}_m} + (1+r)2S}{2} \right) \\ + \beta\lambda(1-\lambda) \sum_{l_m \in L} & \left[\pi_{j,l_m} U \left(\frac{w_{2,\hat{l}_f} + w_{2,l_m} + (1+r)2S}{2} \right) \right] \\ + \beta\lambda(1-\lambda) \sum_{l_f \in L} & \left[\pi_{i,l_f} U \left(\frac{w_{2,l_f} + w_{2,\hat{l}_m} + (1+r)2S}{2} \right) \right] \\ + \beta(1-\lambda)^2 \sum_{l_f \in L} & \left[\pi_{i,l_f} \sum_{l_m \in L} \pi_{j,l_m} U \left(\frac{w_{2,l_f} + w_{2,l_m} + (1+r)2S}{2} \right) \right] \end{aligned} \tag{3.14}$$

where S – per spouse savings.

The IUI algorithm is similar to the case of the single agent, except now we need to take care of the joint dynamics of earnings.

Figure 3.3 shows the results of using a structural model which accounts for marital insurance to infer the level of income uncertainty λ from the simulated data of married couples' behavior. In the simulation, I use 10-quantile income process, 2000 observations and I repeat the simulation 200 times. As in the case of the analytical model in section 3.2.2, it is harder to infer uncertainty from the joint savings decisions, as a result of that, the precision of the IUI is lower than in the case of single agent's uncertainty inference.

3.3.4 Marital Bias

Now, once the validity of the IUI with the non-parametric income process is established, I turn to the simulation of marital bias. In order to do that, I generate behavior of 2000 married couples and treat them as a 4000 single individuals with

Figure 3.3: Inferred Uncertainty (Married Agents)

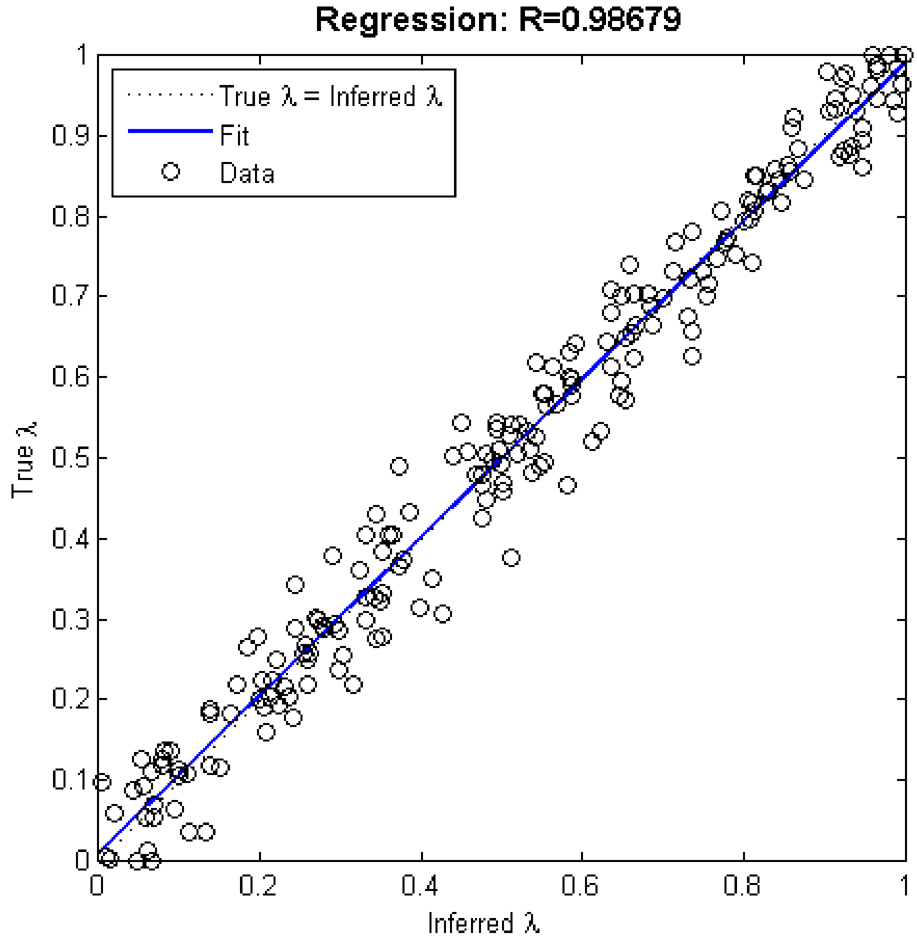

each of the former spouses assumed to have the per-spouse level of savings. After that, I use the single agent's structural model to infer the quality of the private signal λ.

3.3.4.1 Simulation of the Spousal Wage Correlation

Consider a married couple where each spouse is facing the income process defined above, let me denote their current earnings by w_t^m, w_t^f. Since both transition matrices Π_t and wage arrays W_t are fixed a priory, I cannot directly assume the joint dynamics of w_t^m and w_t^f. However, instead of imposing a correlation between wages, I can make an assumption about the correlation of signals that spouses receive every period. In such a way, covariance between the spousal incomes will not be a continuous variable, but rather will be chosen from a set of possible values. For example, under the assumption that $corr(l^m, l^f) = 1$, spouses will always receive the same signal $l^m = l^f$, and correlation of their wages will be the highest, although not perfect. If the process is estimated from data rather than simulated, one would simply compute the special matrices of joint transition probabilities for married couples.

3.3.4.2 Results

It is illustrative to look at the two extreme cases — when spousal earnings correlation is at its maximal and minimal values. Figure 3.4 demonstrates the marital bias in case of perfectly negative correlation of spousal signals (left panel) and in case of the perfect positive correlation of spousal signals (right panel). Recall that in the simple HIP example from section 3.2.2, the negative correlation was producing a systematic underestimation of uncertainty, while the positive correlation was producing a systematic overestimation of uncertainty. As Figure 3.4 shows, this is clearly not the case in this setting.

In the case of negative correlation, the pseudo-single approach results in the underestimation of uncertainty when it is high (overestimation of λ), and in the smaller overestimation when it is low (underestimation of λ). The mechanics of this bias is shown on the left panel of Figure 3.5. When spousal earnings correlation is negative, agents enjoy practically perfect insurance, which results in a moderate savings when the quality of their private information is low and zero savings when it is high. As it is illustrated in Figure 3.5a, when the true certainty is low, so that married agents save $S = 200$ (point A_1), the IUI algorithm decides that certainty is much higher, since it is the closest point at which single agents are expected to save around that value (point B_1). Similar situation happens on the other side of the certainty spectrum. When algorithm observes concentration of savings around the zero level (point A_2) but does not observe any savings at $S = -800, S - 800$ it decides that it is more likely to be point B_2 with the lower certainty level since more types of singles are expected to have similar savings.

When spousal earnings correlation is high, the pseudo-single approach always underestimates uncertainty (overestimates λ) and the magnitude of bias decreases with certainty (Figure 3.4b). The reason behind that is that when the spousal private signals are perfectly correlated and the quality of those signals is high, the per-spouse savings decision of couples converges to the optimal savings of single agents with the same signals (Figure 3.5b). This creates the convergence to the unbiased estimates at high levels of certainty.

When the quality of the private information is low, correlation of those signals declines in importance, and earnings insurance of married couples increases. As a result of that, when the true uncertainty is high (point A_1), the IUI algorithm concludes that it should be lower since it does not acknowledge that those savings were made in the presence of marital insurance.

Since the correlation between spousal earnings is a function of correlation between spousal signals and the quality of those signals λ, at this point I cannot produce a pseudo-single bias estimation line consistent with $\rho = 0.35$ (as estimated by Hyslop (2001)). If, instead, I assume that $\rho = 0.35$ is the correlation between spousal signals, then the marital bias looks somewhat similar to Figure 3.4a i.e. it is negative for $\lambda > 0.55$ and positive for $\lambda < 0.55$ (i.e. estimated λ below (above) the true λ).

Figure 3.4: Pseudo-Single Inference Results

(a) Negative Correlation (b) Positive Correlation

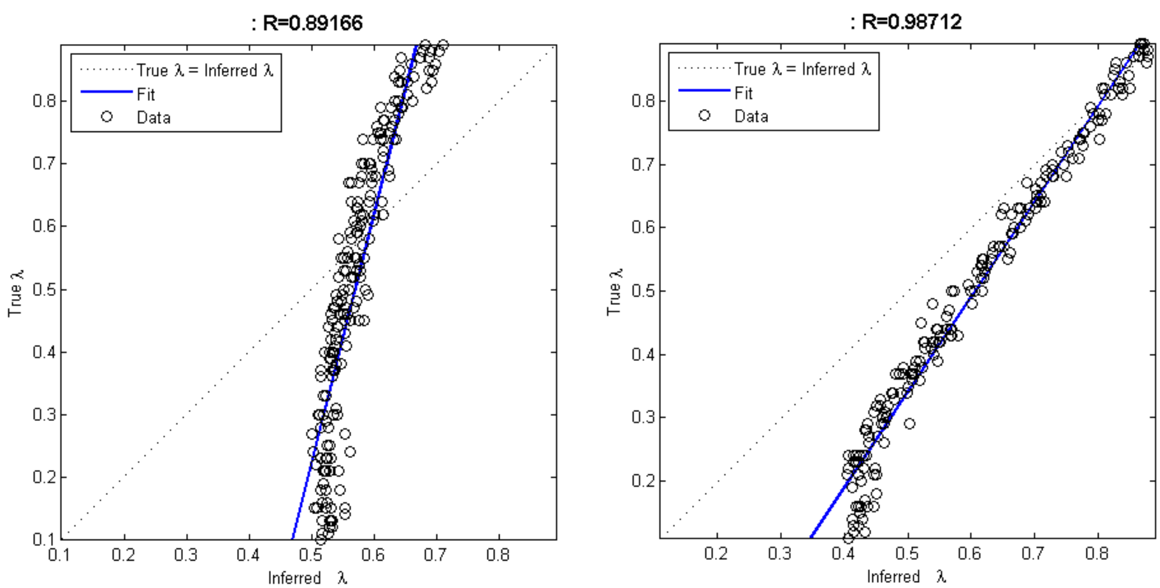

Figure 3.5: Pseudo-Single Inference

(a) Negative Correlation

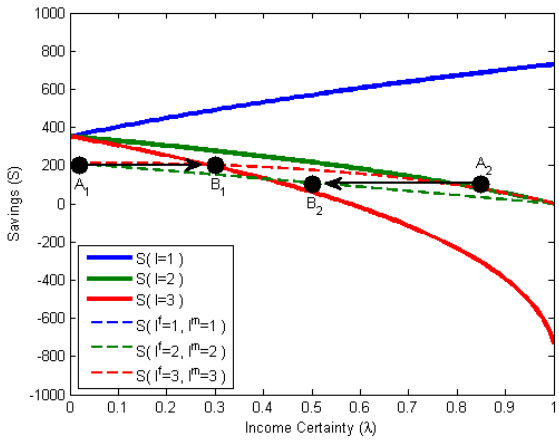

(b) Positive Correlation

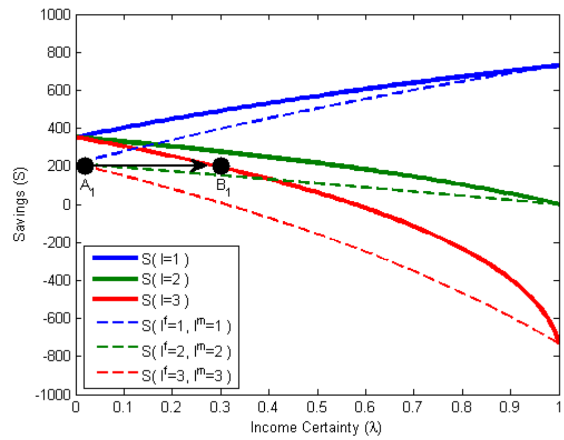

3.4 Conclusion

The goal of this paper was to demonstrate the importance of accounting for the marital state in models of the indirect income uncertainty inference (IUI). Since most of the currently existing IUI models are based on the heterogeneous income profiles earnings process (HIP), I first demonstrate the marital insurance bias in a stylized version of those models. I find that the direction of the bias depends on the co-movement of the spousal earnings and that the magnitude of the average bias is proportional to the size of the underlying uncertainty.

The importance of accounting for the marital state is also demonstrated in a quantitative exercise where the IUI is performed on the simulated data-set. A pseudo-single approach is implemented i.e. each married couple was considered as two single agents with the per-spouse equivalent of savings. Bias produced by not accounting for the marital insurance appears to be large in magnitude and the resulting from such approach estimates fail to resemble the true values used for the simulation.

As a secondary contribution to the IUI literature, I introduce a IUI model based on the non-parametric income process. Besides the usual virtues of a non-parametric process, it also allows to perform the indirect inference with short panel data-sets which is not the case with IUI models based on the HIP process.

Due to the important role earnings uncertainty plays in the inter-temporal choices of individuals, applying the structural model built in this paper to real data is a natural next step and is left for the future work.

APPENDIX A
APPENDIX TO CHAPTER 2

A.1 Example of Mechanism 1

Consider, for example, a simplified version of the full model. Agents make decisions for two periods ($T = 2$), there is no continuation value ($R = 0$) and there are only three wage-types of agents for each gender ($N = 3$):

$w^m \in \{w_1^m, w_2^m, w_3^m\}$ and $w^f \in \{w_1^f, w_2^f, w_3^f\}$.

Also, let there be no income mobility (the transition matrix is represented by an identity matrix). For simplicity, I assume linear utility from consumption ($U(c) = c$). In order to focus explicitly on marriage behavior, assume that agents do not value children ($V(K) = 0$), so that the only incentive to marry is to enjoy consumption economies of scale ($\gamma = 0.5$). Let there be no discounting ($\beta = 1$).

Initially ($t = 1$), all agents are single: the matching probabilities are

$$\phi_1(w^f) = \mu_1(w^m) = \left\{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right\}$$

Since Period 2 is the last period, all marriages will happen according to

$$I_2(w^f, w^m) = \begin{cases} 1, & U(w^f, w^m) \geq U(w^f) \quad \text{and} \\ & U(w^f, w^m) \geq U(w^m) \\ 0, & otherwise \end{cases} \tag{A.1}$$

Because the only choice agents make in this setting is whether to marry or not, the equilibrium can be represented using the marriage-reservation values, which would be $wage \times age$-specific ($RV^g_{t,i}$).

Case 1: Low Male Income Inequality

Let wages be $w^m \in \{2,3,5\}$ and $w^f \in \{2,3,5\}$.

In Period 2, reservation values would be equal to the agents' own wages:

$$RV^g_{2,i} = w^g_i.$$

Then, the second period marriage matrix will be:

$I_2(w^f, w^m) =$

	$w^m_1 = 2$	$w^m_2 = 3$	$w^m_3 = 5$
$w^f_1 = 2$	1	1	0
$w^f_2 = 3$	1	1	1
$w^f_3 = 5$	0	1	1

Equilibrium in this game can be given in terms of the above marriage matrix at Period 2 and a marriage matrix in Period 1:

$I_1(w^f, w^m) =$

	$w_1^m = 2$	$w_2^m = 3$	$w_3^m = 5$
$w_1^f = 2$	1	1	0
$w_2^f = 3$	1	1	1
$w_3^f = 5$	0	1	1

Associated with this equilibrium vectors of matching probabilities at Period 2 are:

$$\Phi_2 = \mathcal{M}_2 \equiv \left\{\frac{1}{9}, 0, \frac{1}{9}\right\}$$

and period 1 reservation values:

$$RV_{1,1}^g = U(w_1^g) + \beta\left[\frac{U(w_1^g, w_1^g)}{9} + \frac{8U(w_1^g)}{9}\right] = 2 + \frac{1}{9}\frac{2+2}{1.5} + \frac{8}{9}2 = 4.07,$$

$$RV_{1,2}^g = 6.29,$$

$$RV_{1,3}^g = 10.18.$$

Case 2: High Male Income Inequality

Now, let male income inequality increase: for example, set the new $w_3^m = 7$. Marriage decisions at $t = 2$ can be represented by:

$$
I_2(w^f, w^m) = \begin{array}{|c|c|c|c|}
\hline
 & w_1^m = 2 & w_2^m = 3 & \mathbf{w_3^m = 7} \\
\hline
w_1^f = 2 & 1 & 1 & 0 \\
\hline
w_2^f = 3 & 1 & 1 & \mathbf{0} \\
\hline
w_3^f = 5 & 0 & 1 & 1 \\
\hline
\end{array},
$$

which means that the top male type has now become too selective, and existing consumption economies of scale would not compensate for the loss in consumption associated with marriage on a second wage-type female.

The first-period marriage matrix associated with the new equilibrium is:

$$
I_1(w^f, w^m) = \begin{array}{|c|c|c|c|}
\hline
 & w_1^m = 2 & w_2^m = 3 & \mathbf{w_3^m = 7} \\
\hline
w_1^f = 2 & 1 & \mathbf{0} & 0 \\
\hline
w_2^f = 3 & 1 & 1 & \mathbf{0} \\
\hline
w_3^f = 5 & 0 & \mathbf{0} & 1 \\
\hline
\end{array}.
$$

The matching probabilities associated with the high inequality equilibrium at Period 2 are:

$$
\Phi_2 \equiv \left\{ \frac{2}{9}, \frac{1}{9}, \frac{2}{9} \right\}, \qquad \mathcal{M}_2 \equiv \left\{ \frac{1}{9}, \frac{2}{9}, \frac{2}{9} \right\}.
$$

It is useful to see how this increase in the wage of one type of male has influenced the equilibrium reservation strategy of all other agents. Change in the reservation values between the low- and high-inequality equilibria also is schematically shown as changes in $I_1(w^f, w^m)$ below. The highest wage-type female agent now receives a higher utility gain from marriage with the top male type. In addition, since w_3^f becomes more selective, there would be higher chances to meet this type of male in Period 2. As such, females of type w_3^f will refuse to marry males of type w_2^m in order to have an option for a better marriage in Period 2. Males of type w_2^m know that in the last period, they would be able to marry w_3^f if matched; in addition, there would be more single w_3^f in Period 2 because fewer of them get marry in Period 1. Given that, males of type w_2^m will also reject marriage proposals from females of type w_1^f. With such a chain reaction to the increase in inequality, even changes at the very top of the distribution would propagate through the income distribution and alter the reservation values of all (and change the marriage acceptance behavior for some) individuals in the economy.

$I_1(w^f, w^m)$=

	w_1^m=2	m_2=3	$\mathbf{w_3^m}$**=7**
w_1^f=2	1, w_1^f: $I_2(w_1^f, w_2^m) = 1$ $\mu_2(w_2^m) \nearrow$	**0**, w_2^m: $I_2(w_3^f, w_2^m)$=1 $\phi_2(w_3^f) \nearrow$	0
w_2^f=3	1, w_2^f: $I_2(w_2^f, w_3^m)$=1 $\mu_2(w_3^f) \nearrow$, $\mu_2(w_2^m) \nearrow$	1	**0**
w_3^f=5	0	**0**, w_3^f: $U(w_3^f, w_3^m) \nearrow$ $\mu_2(w_3^m) \nearrow$	1

A.2 Example of Dominated Strategy Elimination Procedure

Consider the same setting with the same parameter values as introduced in Appendix A.1. The second period marriage matrix will be:

$I_2(w^f, w^m) =$

	$w_1^m = 2$	$w_2^m = 3$	$w_3^m = 5$
$w_1^f = 2$	1	1	0
$w_2^f = 3$	1	1	1
$w_3^f = 5$	0	1	1

A.2.1 Optimistic Beliefs Elimination

We can find an equilibrium of this game by an "optimistic beliefs dominated strategies elimination" procedure. Let initial beliefs about matching probabilities in Period 2 be the most optimistic ones possible, i.e. all the agents of the opposite gender are expected to be single in the marriage market of Period 2:

Iteration 1.

$\tilde{\Phi}_2 = \tilde{\mathcal{M}}_2 \equiv \left\{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right\}$

Then, reservation values in Period 1 would be:

$$RV^g_{1,1} = U(w^g_1) + \beta\left[\frac{U(w^g_1, w^g_1)}{3} + \frac{U(w^g_1, w^g_2)}{3} + \frac{U(w^g_1)}{3}\right] = 2 + \frac{1}{3}\left[\frac{2+2}{1.5} + \frac{2+3}{1.5} + 2\right] = 4.66,$$

$$RV^g_{1,2} = 7.22,$$

$$RV^g_{1,3} = 10.66.$$

And utility values from marriage in Period 1 are:

$$(1+\beta)U(w^g_1, w^g_1) = 2\frac{2+2}{1.5} = 5.33,$$

$$(1+\beta)U(w^g_1, w^g_2) = (1+\beta)U(w^g_2, w^g_1) = 2\frac{2+3}{1.5} = 6.66,$$

$$(1+\beta)U(w^g_2, w^g_2) = 2\frac{3+3}{1.5} = 8,$$

$$(1+\beta)U(w^g_1, w^g_3) = (1+\beta)U(w^g_3, w^g_1) = 2\frac{2+5}{1.5} = 9.33,$$

$$(1+\beta)U(w^g_2, w^g_3) = (1+\beta)U(w^g_3, w^g_2) = 2\frac{3+5}{1.5} = 10.66,$$

$$(1+\beta)U(w^g_3, w^g_3) = 2\frac{5+5}{1.5} = 13.33.$$

Beliefs $\tilde{\Phi}_2 = \tilde{\mathcal{M}}_2 \equiv \left\{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right\}$ are associated with reservation strategies such that no two types would agree to marry each other in Period 1. Some such strategies are obviously strictly dominated – for example, two agents of type 3 would always marry each other in Period 1. In fact, even with such optimistic beliefs, some types of agents would still marry each other if matched in Period 1:

$$I_1(w_3^g, w_3^g) = 1 : \quad (1+\beta)U(w_3^g, w_3^g) \geq RV_{1,3}^g,$$

$$I_1(w_2^g, w_3^g) = I_1(w_3^g, w_2^g) = 1 : \quad (1+\beta)U(w_2^g, w_3^g) \geq RV_{1,3}^g,$$

$$I_1(w_2^g, w_2^g) = 1 : \quad (1+\beta)U(w_2^g, w_2^g) \geq RV_{1,2}^g,$$

$$I_1(w_1^g, w_1^g) = 1 : \quad (1+\beta)U(w_1^g, w_1^g) \geq RV_{1,1}^g.$$

I now update beliefs about the matching probabilities in Period 2 by eliminating beliefs associated with dominated reservation strategies:

$$\tilde{\mu}_{1,1} = \mu_{1,1} - \mu_{1,1} \times \phi_{1,1},$$

$$\tilde{\mu}_{1,2} = \mu_{1,2} - \mu_{1,2} \times \phi_{1,2} - \mu_{1,2} \times \phi_{1,3},$$

$$\tilde{\mu}_{1,3} = \mu_{1,3} - \mu_{1,3} \times \phi_{1,2} - \mu_{1,3} \times \phi_{1,3}.$$

and making a similar update of female matching probabilities.

Iteration 2.

The new beliefs about the matching probabilities in Period 2 are: $\tilde{\Phi}_2 = \tilde{\mathcal{M}}_2 \equiv \left\{\frac{2}{9}, \frac{1}{9}, \frac{1}{9}\right\}$

Then, reservation values in Period 1 would be:

$$RV_{1,1}^{g} = U(w_1^g) + \beta\left[\frac{2U(w_1^g, w_1^g)}{9} + \frac{U(w_1^g, w_2^g)}{9} + \frac{6U(w_1^g)}{9}\right] = 4.29,$$

$$RV_{1,2}^{g} = 6.44,$$

$$RV_{1,3}^{g} = 10.22.$$

Given the new reservation values, the following marriages will occur under these beliefs:

$$I_1(w_1^g, w_2^g) = I_1(w_2^g, w_1^g) = 1 : \quad (1+\beta)U(w_1^g, w_2^g) \geq RV_{1,2}^{g}$$

$$+\,\textbf{all the marriages from iteration 1.}$$

Note, that all the marriages that were taking place in Iteration 1, also take place given the updated beliefs. This monotonicity of reservation values is a crucial for the uniqueness result. This property of the model is due to the assumption that matching probabilities are not normalized after the marriage market takes place.

I update beliefs about the matching probabilities in Period 2 by eliminating beliefs associated with dominated reservation strategies:

$$\tilde{\mu}_{1,1} = \mu_{1,1} - \mu_{1,1} \times \phi_{1,1} - \mu_{1,1} \times \phi_{1,2},$$

$$\tilde{\mu}_{1,2} = \mu_{1,2} - \mu_{1,2} \times \phi_{1,1} - \mu_{1,2} \times \phi_{1,2} - \mu_{1,2} \times \phi_{1,3},$$

$$\tilde{\mu}_{1,3} = \mu_{1,3} - \mu_{1,3} \times \phi_{1,2} - \mu_{1,3} \times \phi_{1,3}.$$

and, similarly, update female matching probabilities.

Iteration 3.

New beliefs are given by:

$\tilde{\Phi}_2 = \tilde{\mathcal{M}}_2 \equiv \left\{\frac{1}{9}, 0, \frac{1}{9}\right\}$

Then, reservation values in Period 1 would be:

$$RV_{1,1}^g = U(w_1^g) + \beta\left[\frac{U(w_1^g, w_1^g)}{9} + \frac{8U(w_1^g)}{9}\right] = 4.07,$$

$$RV_{1,2}^g = 6.29,$$

$$RV_{1,3}^g = 10.18.$$

Given this reservation values, there is no change in marriage behavior between Iteration 2 and Iteration 3. That is we have reached the equilibrium, and the associated marriage matrix in period 1 is:

$I_1(w^f, w^m) =$

	$w_1^m = 2$	$w_2^m = 3$	$w_3^m = 5$
$w_1^f = 2$	1	1	0
$w_2^f = 3$	1	1	1
$w_3^f = 5$	0	1	1

A.2.2 Pessimistic Beliefs Elimination

We can now perform similar iterations, but this time, we will find an equilibrium of this game by a "pessimistic beliefs dominated strategies elimination" procedure. Let initial beliefs about matching probabilities in Period 2 be the most pessimistic ones possible, i.e. all the agents of the opposite gender are expected to be married prior to the marriage market of Period 2:

Iteration 1.

$\tilde{\Phi}_2 = \tilde{\mathcal{M}}_2 \equiv \left\{ 0, 0, 0 \right\}$

Then, reservation values in Period 1 would be:

$$RV_{1,1}^{g} = U(w_1^g) + \beta\left[U(w_1^g)\right] = 2 + 2 = 4,$$

$$RV_{1,2}^{g} = 6,$$

$$RV_{1,3}^{g} = 10.$$

And utility values from marriage in Period 1 are the same as in the optimistic elimination case:

$$(1+\beta)U(w_1^g, w_1^g) = 5.33,$$

$$(1+\beta)U(w_1^g, w_2^g) = 6.66,$$

$$(1+\beta)U(w_2^g, w_2^g) = 8,$$

$$(1+\beta)U(w_1^g, w_3^g) = 9.33,$$

$$(1+\beta)U(w_2^g, w_3^g) = 10.66,$$

$$(1+\beta)U(w_3^g, w_3^g) = 13.33.$$

Beliefs $\tilde{\Phi}_2 = \tilde{\mathcal{M}}_2 \equiv \left\{ 0, 0, 0 \right\}$ are associated with reservation strategies such that any two types would agree to marry in Period 1. Again, some such strategies are obviously strictly dominated – for example, an agent of type 3 would never marry an agent of type 1, because even if she expects to never marry in future, her value of living alone is higher than that of spiting consumption with an agent of type 1:

$$RV^g_{1,3} = 10 > (1+\beta)U(w^g_1, w^g_3) = 9.33.$$

It is easy to check that all the other types would agree to marry each other given this beliefs.

I now update beliefs about the matching probabilities in Period 2 by eliminating beliefs associated with dominated reservation strategies:

$$\tilde{\mu}_{1,1} = \mu_{1,1} - \mu_{1,1} \times \phi_{1,1} - \mu_{1,1} \times \phi_{1,2},$$

$$\tilde{\mu}_{1,2} = \mu_{1,2} - \mu_{1,2} \times \phi_{1,1} - \mu_{1,2} \times \phi_{1,2} - \mu_{1,2} \times \phi_{1,3},$$

$$\tilde{\mu}_{1,3} = \mu_{1,3} - \mu_{1,3} \times \phi_{1,2} - \mu_{1,3} \times \phi_{1,3}.$$

and making a similar update of female matching probabilities.

Iteration 2.

The new beliefs about the matching probabilities in Period 2 are: $\tilde{\Phi}_2 = \tilde{\mathcal{M}}_2 \equiv \left\{\frac{1}{9}, 0, \frac{1}{9}\right\}$

Then, reservation values in Period 1 would be:

$$RV^g_{1,1} = U(w^g_1) + \beta\left[\frac{2U(w^g_1, w^g_1)}{9} + \frac{U(w^g_1, w^g_2)}{9} + \frac{6U(w^g_1)}{9}\right] = 4.07,$$

$$RV^g_{1,2} = 6.29,$$

$$RV^g_{1,3} = 10.18.$$

Given the new reservation values, there is no change in marriage decisions. That is we again have reached the equilibrium. Moreover, this equilibrium is the same as the one attained through the optimistic elimination. Note, that we were sequentially ruling out only the strictly dominated strategies, which is any other beliefs about matching probabilities can not be supported in the equilibrium.

Given any parameter values, I can run this two-sided elimination algorithm, and if the iterating procedures converge to the same equilibrium, this equilibrium is unique.

REFERENCES

S Rao Aiyagari, Jeremy Greenwood, and Nezih Guner. On the state of the union. *Journal of Political Economy*, 108(2):213–244, 2000.

Julia Arroyo, Krista K Payne, Susan L Brown, and Wendy D Manning. Crossover in median age at first marriage and first birth: Thirty years of change. *National Center for Family & Marriage Research. http://ncfmr. bgsu. edu/pdf/family_profiles/file107893. pdf*, 2013.

Orazio Attanasio and Britta Augsburg. Subjective expectations and income processes in rural india. *Universtiy College London. Mimeo*, 2012.

Sandra Barney. Accessing medicalized donor sperm in the us and britain: an historical narrative. *Sexualities*, 8(2):205–220, 2005.

Sumon K Bhaumik and Jeffrey B Nugent. Does economic uncertainty affect the decision to bear children? evidence from east and west germany. 2005.

McKinley L Blackburn, David E Bloom, and David Neumark. Fertility timing, wages, and human capital. *Journal of Population Economics*, 6(1):1–30, 1993.

Francine D Blau and Lawrence M Kahn. Swimming upstream: Trends in the gender wage differential in the 1980s. *Journal of labor Economics*, 15(1, Part 1):1–42, 1997.

Martin Browning, Pierre-Andre Chiappori, and Arthur Lewbel. Estimating consumption economies of scale, adult equivalence scales, and household bargaining power. *The Review of Economic Studies*, page rdt019, 2013.

Anne Case, I-Fen Lin, and Sara McLanahan. How hungry is the selfish gene? *The Economic Journal*, 110(466):781–804, 2000.

Elizabeth M Caucutt, Nezih Guner, and John Knowles. Why do women wait? matching, wage inequality, and the incentives for fertility delay. *Review of Economic Dynamics*, 5(4):815–855, 2002.

Hector Chade and Gustavo Ventura. Taxes and marriage: A two-sided search analysis. *International Economic Review*, 43(3):955–985, 2002.

Eugene Choo and Aloysius Siow. Who marries whom and why. *Journal of political Economy*, 114(1):175–201, 2006.

Alessandro Cigno and John Ermisch. A microeconomic analysis of the timing of births. *European Economic Review*, 33(4):737–760, 1989.

Tristan P Coughlin and Scott D Drewianka. Can rising inequality explain aggregate trends in marriage? evidence from us states, 1977-2005. *The BE Journal of Economic Analysis & Policy*, 11(1), 2011.

José María Da-Rocha, Luisa Fuster, et al. Why are fertility and female participation rates positively correlated across oecd countries. *International Economic Review*, 2002.

Molly W Dahl, Thomas C DeLeire, and Jonathan A Schwabish. Trends in earnings variability over the past 20 years. Congress of the United States, Congressional Budget Office, 2007.

David De La Croix and Matthias Doepke. Inequality and growth: why differential fertility matters. *The American Economic Review*, 93(4):1091–1113, 2003.

Mariacristina De Nardi, Giulio Fella, and Gonzalo Paz Pardo. The implications of richer earnings dynamics for consumption and wealth. Technical report, National Bureau of Economic Research, 2016.

Jason DeBacker, Bradley Heim, Vasia Panousi, Shanthi Ramnath, and Ivan Vidangos. Rising inequality: transitory or persistent? new evidence from a panel of us tax returns. *Brookings Papers on Economic Activity*, 2013(1):67–142, 2013.

Adeline Delavande and Susann Rohwedder. Eliciting subjective probabilities in internet surveys. *Public Opinion Quarterly*, 72(5):866–891, 2008.

Jeff Dominitz. Earnings expectations, revisions, and realizations. *Review of Economics and statistics*, 80(3):374–388, 1998.

David T Ellwood and Christopher Jencks. The spread of single-parent families in the united states since 1960. 2004.

David T Ellwood and Christopher Jencks. The uneven spread of single-parent families: What do we know? where do we look for answers? *Social inequality*, 1:3–77, 2004.

Andrés Erosa, Luisa Fuster, and Diego Restuccia. Fertility decisions and gender differences in labor turnover, employment, and wages. *Review of Economic Dynamics*, 5(4):856–891, 2002.

Andrés Erosa, Luisa Fuster, and Diego Restuccia. A general equilibrium analysis of parental leave policies. *Review of Economic Dynamics*, 13(4):742–758, 2010.

Raquel Fernandez, Nezih Guner, and John Knowles. Love and money: A theoretical and empirical analysis of household sorting and inequality. *The Quarterly Journal of Economics*, 120(1):273–344, 2005.

Jonas DM Fisher and Martin Gervais. Why has home ownership fallen among the young? *International Economic Review*, 52(3):883–912, 2011.

Claudia Goldin and Lawrence F Katz. The power of the pill: Oral contraceptives and womens career and marriage decisions. *Journal of political Economy*, 110(4):730–770, 2002.

Joshua R Goldstein and Catherine T Kenney. Marriage delayed or marriage forgone? new cohort forecasts of first marriage for us women. *American Sociological Review*, pages 506–519, 2001.

Peter Gottschalk, Robert Moffitt, Lawrence F Katz, and William T Dickens. The growth of earnings instability in the us labor market. *Brookings Papers on Economic Activity*, 1994(2):217–272, 1994.

Eric D Gould and M Daniele Paserman. Waiting for mr. right: rising inequality and declining marriage rates. *Journal of Urban Economics*, 53(2):257–281, 2003.

Jeremy Greenwood, Nezih Guner, and John A Knowles. More on marriage, fertility, and the distribution of income. *International Economic Review*, 44(3):827–862, 2003.

Luigi Guiso, Tullio Jappelli, and Luigi Pistaferri. What determines earnings and employment risk? 1998.

Nezih Guner, Remzi Kaygusuz, and Gustavo Ventura. Taxation, aggregates and the household. 2008.

Nezih Guner and John Knowles. Why is the rate of single-parenthood lower in canada than in the us? a dynamic equilibrium analysis of welfare policies. *Canadian Journal of Economics/Revue canadienne d'économique*, 42(1):56–89, 2009.

Fatih Guvenen. Learning your earning: Are labor income shocks really very persistent? *The American economic review*, 97(3):687–712, 2007.

Fatih Guvenen. Income inequality and income risk: Old myths vs. new facts. *Lecture Slides available at https://fguvenendotcom. files. wordpress. com/2014/04/handout_inequality_risk_webversion_may20151. pdf*, 2015.

Fatih Guvenen, Serdar Ozkan, and Jae Song. The nature of countercyclical income risk. *Journal of Political Economy*, 122(3):621–660, 2014.

Fatih Guvenen and Anthony A Smith Jr. Inferring labor income risk and partial insurance from economic choices. *Econometrica*, pages 2085–2129, 2014.

Brady E Hamilton and Stephanie J Ventura. Fertility and abortion rates in the united states, 1960–2002. *International journal of andrology*, 29(1):34–45, 2006.

Mary Eschelbach Hansen and Bradley A Hansen. The economics of adoption of children from foster care. *Child Welfare*, 85(3):559, 2006.

Stephen K Happel, Jane K Hill, and Stuart A Low. An economic analysis of the timing of childbirth. *Population Studies*, 38(2):299–311, 1984.

Jonathan Heathcote, Fabrizio Perri, and Giovanni L Violante. Unequal we stand: An empirical analysis of economic inequality in the united states, 1967–2006. *Review of Economic dynamics*, 13(1):15–51, 2010.

Jonathan Heathcote, Kjetil Storesletten, and Giovanni L Violante. The macroeconomic implications of rising wage inequality in the united states. *Journal of political economy*, 118(4):681–722, 2010.

James J Heckman, Lance J Lochner, and Petra E Todd. Fifty years of mincer earnings regressions. Technical report, National Bureau of Economic Research, 2003.

Gregory D Hess. Marriage and consumption insurance: Whats love got to do with it? *Journal of Political Economy*, 112(2):290–318, 2004.

Michael D Hurd. Subjective probabilities in household surveys. *Annu. Rev. Econ.*, 1(1):543–562, 2009.

Dean R Hyslop. Rising us earnings inequality and family labor supply: The covariance structure of intrafamily earnings. *American Economic Review*, pages 755–777, 2001.

Wm. Robert Johnston. United states abortion rates, 1960-2008. Technical report, 2014.

Lawrence F Katz and Kevin M Murphy. Changes in relative wages, 1963–1987: supply and demand factors. *The quarterly journal of economics*, 107(1):35–78, 1992.

Michael C Keeley. *A Model of Marital Formation: The Determinants of the Optimal Age at First Marriage*. PhD thesis, University of Chicago, 1974.

Melissa AZ Knoll, Christopher R Tamborini, and Kevin Whitman. I do want to save: Marriage and retirement savings in young households. *Journal of Marriage and Family*, 74(1):86–100, 2012.

John A Knowles. Why are married men working so much? an aggregate analysis of intra-household bargaining and labour supply. *The Review of Economic Studies*, 80(3):1055–1085, 2013.

Elisabeth M Landes and Richard A Posner. The economics of the baby shortage. *The Journal of Legal Studies*, 7(2):323–348, 1978.

David S Loughran. The effect of male wage inequality on female age at first marriage. *Review of Economics and Statistics*, 84(2):237–250, 2002.

Shelly Lundberg and Robert A Pollak. Cohabitation and the uneven retreat from marriage in the united states, 1950–2010. In *Human capital in history: The American record*, pages 241–272. University of Chicago Press, 2014.

Shelly Lundberg and Robert A Pollak. The evolving role of marriage: 1950–2010. *The Future of Children*, 25(2):29–50, 2015.

Shelly Lundberg, Robert A Pollak, and Jenna Stearns. Family inequality: Diverging patterns in marriage, cohabitation, and childbearing. *The Journal of Economic Perspectives*, 30(2):79–101, 2016.

Charles F Manski. Measuring expectations. *Econometrica*, 72(5):1329–1376, 2004.

S McLanahan and MJ Sanderfur. Growing up with a single parent: What hurts and what helps, cambridge, ma, harvard university press. 1994.

Marshall H Medoff. An empirical analysis of adoption. *Economic inquiry*, 31(1):59–70, 1993.

Amalia R Miller. The effects of motherhood timing on career path. *Journal of Population Economics*, 24(3):1071–1100, 2011.

Jacob A Mincer. Age and experience profiles of earnings. In *Schooling, experience, and earnings*, pages 64–82. NBER, 1974.

Robert A Moffitt and Peter Gottschalk. Trends in the transitory variance of male earnings methods and evidence. *Journal of Human Resources*, 47(1):204–236, 2012.

Robert A Moffitt, Peter Gottschalk, et al. Trends in the variances of permanent and transitory earnings in the us and their relation to earnings mobility. Technical report, Boston College Department of Economics, 1998.

Chiaki Moriguchi et al. The evolution of child adoption in the united states, 1950-2010: An economic analysis of historical trends. *Economic Review*, 63(3):265–285, 2012.

Andre R Neveu. Earnings volatility trends and the great moderation: a multifactor residual approach. *Atlantic Economic Journal*, 43(2):229–245, 2015.

Claudia Olivetti. Changes in women's hours of market work: The role of returns to experience. *Review of Economic Dynamics*, 9(4):557–587, 2006.

June O'Neill and Solomon Polachek. Why the gender gap in wages narrowed in the 1980s. *Journal of Labor economics*, 11(1, Part 1):205–228, 1993.

Erik Plug and Wim Vijverberg. Schooling, family background, and adoption: Is it nature or is it nurture? *Journal of political Economy*, 111(3):611–641, 2003.

Jens Poongothai, Ts Gopenath, Swaminathan Manonayaki, et al. Genetics of human male infertility. *Singapore Med J*, 50(4):336–347, 2009.

Xavier Ramos and Christian Schluter. Subjective income expectations and income risk. 2006.

Priya Ranjan. Fertility behaviour under income uncertainty. *European Journal of Population/Revue Européenne de Démographie*, 15(1):25–43, 1999.

Ferdinando Regalia, José-Vıctor Rıos-Rull, and J Short. What accounts for the increase in the number of single households. *Unpublished manuscript*, 2011.

John Sabelhaus and Jae Song. The great moderation in micro labor earnings. *Journal of Monetary Economics*, 57(4):391–403, 2010.

Bruce Sacerdote. The nature and nurture of economic outcomes. Technical report, National Bureau of Economic Research, 2000.

Cezar Santos and David Weiss. Risky income, risky families: Marriage and divorce in a volatile labor market. Technical report, mimeo, 2012.

Cezar Santos and David Weiss. why not settle down already? a quantitative analysis of the delay in marriage. *International Economic Review*, 57(2):425–452, 2016.

Daniel Schneider and Adam Reich. Marrying aint hard when you got a union card? labor union membership and first marriage. *Social Problems*, 61(4):625–643, 2014.

Rakesh K Sharma, Fabio F Pasqualotto, David R Nelson, Anthony J Thomas, and Ashok Agarwal. The reactive oxygen speciestotal antioxidant capacity score is a new measure of oxidative stress to predict male infertility. *Human Reproduction*, 14(11):2801–2807, 1999.

Michelle Sheran. The career and family choices of women: A dynamic analysis of labor force participation, schooling, marriage, and fertility decisions. *Review of Economic Dynamics*, 10(3):367–399, 2007.

Donggyun Shin and Gary Solon. Trends in men's earnings volatility: What does the panel study of income dynamics show? *Journal of public Economics*, 95(7):973–982, 2011.

Duangsuda Sopchokchai. *Essays in Health and Income Dynamics*. PhD thesis, University OF Ottawa, 2016.

Betsey Stevenson and Justin Wolfers. Marriage and divorce: Changes and their driving forces. *The Journal of Economic Perspectives*, 21(2):27–52, 2007.

Christopher Tietze. The effect of legalization of abortion on population growth and public health. *Family Planning Perspectives*, pages 123–127, 1975.

Mitra Toossi. A century of change: The us labor force, 1950-2050. *Monthly Lab. Rev.*, 125:15, 2002.

Boukje M van Noord-Zaadstra, Caspar W Looman, Hans Alsbach, J Dik Habbema, Egbert R te Velde, and Jan Karbaat. Delaying childbearing: effect of age on fecundity and outcome of pregnancy. *Bmj*, 302(6789):1361–1365, 1991.

Sharon Vandivere, Karen Malm, and Laura Radel. *Adoption USA: A chartbook based on the 2007 National Survey of Adoptive Parents*. US Department of Health and Human Services, Office of the Assistant Secretary for Planning and Evaluation, 2009.

Jane Waldfogel, Terry-Ann Craigie, and Jeanne Brooks-Gunn. Fragile families and child wellbeing. *The Future of children/Center for the Future of Children, the David and Lucile Packard Foundation*, 20(2):87, 2010.