Life-Cycle Fertility: Means vs. Motives vs. Opportunities

(Job Market Paper)

Sekyu Choi
University of Pennsylvania *

November 2008

Abstract

More educated individuals have fewer children and have them later in life than their less educated counterparts. To understand these facts, I embed a standard fertility theory into a realistic life-cycle, consumption-savings framework. I then use this model to revisit the question of why there is a negative skill-fertility relationship and to assess the extent to which fertility theories can accommodate the cross-sectional and life-cycle variation of births in the data.

I show that fertility risk is a key feature needed in order to account for the level and the timing of births in the life-cycle. By fertility risk I mean both: (i) early in life, women with different educational attainments have different levels of success when carrying out childbearing plans and (ii) biological constraints affect fertility decisions later in life. I estimate the extent of these risks using individual data on pregnancies and abortions for the U.S. during the mid 1990s.

These findings lead me to conclude that fertility theories used in macroeconomics to understand the time-series dimension of the data (e.g., the demographic transition) don't impose enough structure to account for cross-sectional variation nor life-cycle fertility facts. Hence, the additional discipline represented by life-cycle considerations, restricts the basic ingredients for a successful economic theory of fertility.

Keywords: Stochastic fertility, Life-cycle model, Heterogeneous agents, Birth control, Educational attainment

JEL Classification: D31,D91,E21,J13,J31

---

*Email: sechoi@econ.upenn.edu. I owe a debt of gratitude to José-Víctor Ríos-Rull, who has supported me in many ways. This paper benefited from discussions with Jeremy Greenwood, Dirk Krueger, Jesús Fernández-Villaverde, Iourii Manovski, Ken Wolpin, Hal Cole, John Knowles, Stefania Albanesi, Alexander Bick, Cristina Fuentes-Albero, Edith Liu, Clement Joubert, Jon Pogach and seminar participants at the macro lunch and the money/macro seminar at Upenn. All errors are mine.

1 Introduction

In this paper I study life-cycle fertility patterns and the cross sectional distribution of births in the U.S. during the mid 1990s. Using data from the National Survey of Family Growth for the year 1995 (NSFG95), I document two features of recent fertility trends: more educated females (those with at least some college education) have less children and they start childbearing later in life than their less educated counterparts (females with at most a high school diploma).

My approach to understand the facts is as follows. I embed a standard economic model of fertility (the allocation of mother's time variety[1]) into a rich life-cycle, consumption-savings framework[2] where fertility is an endogenous choice subject to idiosyncratic shocks (undesired outcomes). Using this model, I revisit the question of why there exist a negative skill-fertility relationship and ask whether standard fertility theories, which use the notion that individuals with higher education and wages have less children because of the higher opportunity cost of child rearing, can accommodate the cross-sectional and life-cycle variation in the data.

Results from the quantitative exercise show that differential fertility risk is key when trying to account for differences in the timing and the number of births across educational groups. Thus, relying on substitution effects alone, as is done in standard fertility theories, is not sufficient to produce a negative skill-fertility relationship when we try to match life-cycle facts: under imperfect capital markets, non-labor income (savings) is endogenous and higher for more skilled/lucky earners. This creates an income effect that (theoretically) could overpower the opportunity cost of having children. Using my model, I can assess the importance of this margin and show how fertility risk plays a role in generating different outcomes by group.

Throughout this paper, fertility risk has a dual meaning: earlier in life, it represents the fact that there exist failure when using contraceptive technologies and pregnancies may occur sooner than expected or when they were not wanted at all; later in life, fertility risk acquires a different connotation, since females who postpone childbearing find themselves dealing with biological constraints to conceive. I measure fertility risk by estimating my model using individual data from the NSFG95 on pregnancies, abortions and intentions on conceptions (whether pregnancies were planned or not).

My analysis also shows that predictions from fertility theories change significantly under two

[1]This theory was first developed by Mincer (1963) and Becker (1965) and used recently by Greenwood, Seshadri, and Vandenbroucke (2005) and Galor and Weil (1996), among others

[2]The basis of this structure is in Hong and Ríos-Rull (2007) and in Hong (2006). The main difference with those papers is that they don't have endogenous fertility. On the other hand, I depart from general equilibrium due to computational burden considerations

different capital market assumptions: a model calibrated with and without individuals having the ability to save provides very different results in terms of matching the facts. This should pose a cautionary signal to models that try to understand life-cycle fertility in the absence of savings.

My approach borrows insights from the empirical microeconomic literature that studies life-cycle fertility[3] using structural and dynamic models of fertility choice. From that literature, my paper relates to Wolpin (1984) and Hotz and Miller (1993) who acknowledge the importance of the stochastic nature of fertility. Wolpin analyzes how child mortality risk shapes fertility choices using Malaysian data; Hotz and Miller estimate birth control method choices by females in a life-cycle framework. However, my approach differs starkly in terms of assumptions regarding capital markets and preference heterogeneity (this is true for the whole literature and not just the specific papers mentioned above): I assume imperfect capital markets in the sense that agents can save but not borrow against their future earnings; also, I impose the same preferences for all agents, downplaying the role of unobserved heterogeneity in utility.

This paper relates the most to Rosenzweig and Schultz (1989) and Conesa (2000). The first paper provides evidence that more educated individuals are more efficient using different birth control methods, which is the main mechanism through which I obtain a negative skill-fertility relationship. However, Rosenzweig and Schultz (1989) restrict their attention to all-white couples in intact first marriages, which might produce sample bias in their regression estimates. Also, they analyze a time period (late 1960s and early 1970s) when policies regarding birth control were different to the ones in the period I analyze: the pill was still not massively adopted by single females and abortion was not readily available to everyone. On the other hand, Conesa (2000) studies fertility and educational attainment in the U.S. and develops a general equilibrium overlapping generations model in which agents choose whether to conceive period by period and how much to consume and save. The main difference between my paper and Conesa (2000) is that he assumes fertility risk when couples are seeking a birth (in the form of a constant probability of getting pregnant if one chooses to) but perfect control when they don't want a pregnancy.

The structure of the paper is as follows: In the next section I describe my data sources and the main stylized facts I want to explain. In section 3, I pose a simple static model in which I show where standard theories of fertility might fail when moved to a life-cycle setting. In section 4, I describe my model. Sections 5 and 6 describe the functional forms used in the model and the specific estimation method to obtain model parameters. I show the estimation results and some quantitative experiments in section 7. The final section concludes.

[3]See Hotz, Klerman, and Willis (1997) for a survey

2 The Facts

I use information from the National Survey of Family Growth (NSFG) to put forward a set of facts on U.S. fertility. The NSFG is compiled by the National Center for Health Statistics (NCHS) and gathers information on family life, fertility, use of birth control and other health related questions. I use the survey for the year 1995, which comprises around ten thousand women between the ages of 15 and 44.

For every survey participant, the NSFG collects retrospective information on usage of birth control methods, on a monthly basis for up to 5 years. Participants also answer questions on wantedness and timing of births and pregnancy outcomes for all pregnancies conceived during that 5 year period. The survey also contains information on educational attainment, marital status and other background information.

I present age specific fertility rates in figure 1 and age specific abortion rates in figure 2.

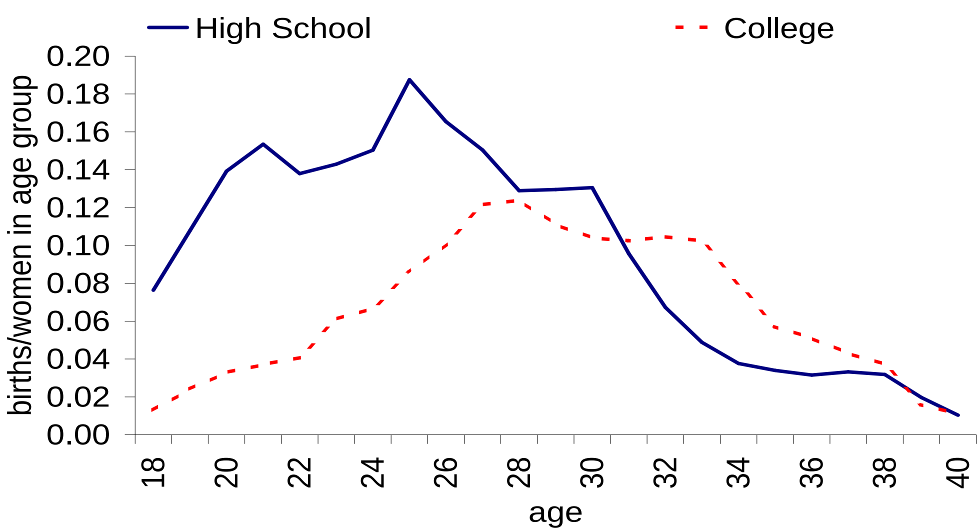

Figure 1: Age-specific fertility rates by education of the mother (NSFG 1995)

Both graphs present information on pregnancies occurring between 1994 and 1995, for each particular education-age group, i.e., I'm focusing on the cross-sectional dimension of the data[4] Age specific fertility rates are defined by the ratio between the number of pregnancies in the specific education-age group and the total number of women in that group. Abortion rates are the number of abortions divided by the total number of women in each group. I divide groups

[4]Hence, I'm assuming NO cohort effects in fertility rates. A more complete discussion is in the Appendix.

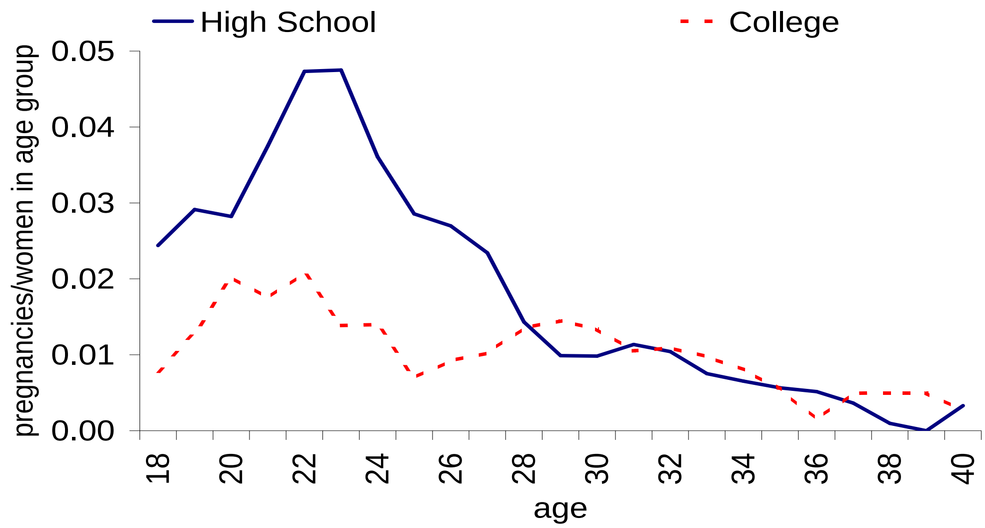

Figure 2: Age-specific abortion rates by education of the mother (NSFG 1995)

according to educational attainment as follows: High School (those without any post-secondary education) and College (those with at least some post-secondary education). In both figures, I show smoothed statistics (moving averages of 3 years).

The following is a list of stylized facts from the data:

1. The education fertility gap: the high school group has a total fertility rate (TFR)[5] of 2.2, while for the college group, it is 1.5 .

2. Timing of births: high school females start having children earlier than their college counterparts. According to figure 1, the age with the highest fertility rate is 25 for the High school group and 28 for the college one.[6]

3. Failure in fertility plans: the number of aborted pregnancies is higher for the high school group. The abortion rate for high school educated females is approximately 18 per 1000 women, while the number for college educated females is 11

All but the last fact have been well documented in recent economic literature. Since income of more educated individuals is higher, fact 1 above can be restated as the well known negative

[5]Total fertility rates are the sum of of the age specific fertility rates and represent a cross-sectional measure of aggregate fertility

[6]A similar statistic is average age of first birth (22.7 and 26.2 for high school and college respectively)

income-fertility relationship.[7] The differential timing of births is documented and studied by Caucutt, Gunner, and Knowles (2002), who argue that returns to experience as well as marriage markets play an important role in explaining delay in childbirth.

The last fact shows that failure rates are more acute for the high school group. In terms of accuracy of this data, Fu, Darroch, Henshaw, and Kolb (1998) claim that the introduction of computer assisted interviews in the NSFG for the year 1995 helped in reducing underreporting of abortions and unplanned pregnancies. Nevertheless, their study shows (by comparing implied abortion rates from the NSFG to data from abortion providers in the U.S.) that non reported abortion cases are still present and are higher for lower income groups (the high school group). Hence, differences in failure rates by educational groups are likely to be more pronounced if I were to have access to all data.

3 Example: A Static Model

The example below is useful to understand the basic features of standard economic theories of fertility. I also show how a static model of fertility can be modified to account for stochastic fertility outcomes and a dynamic setting (period by period choices). This example then shows where the standard time allocation theory might fail when faced with life-cycle considerations and why fertility risk is a natural solution.

Suppose that individuals derive utility from consumption c and the number of children k in the household. I assume separability in the utility from both elements and $\log$-preferences. Agents have one unit of time which can be sold in the market at rate w and also have access to some non-labor income a. If there are children in the household, agents must spend a fraction $b(k) \in (0,1)$ of their time taking care of them. This function is increasing in k. I ignore good-costs of children to keep the analysis simple. I model fertility choices in a two stage setting. In the first stage agents choose whether to increase the size of their household. During the second stage, agents choose optimal consumption.

In the second stage, agents solve the following problem, given the stock of children k (chosen during the previous stage)

[7] This observation goes back to Becker (1960). Jones and Tertilt (2008) study Census data and find that this negative relationship is robust across time and different definitions of income

$$
\begin{aligned}
V(a,w,k) &= \max_{c} \ \log(c) + \gamma \log(k) \\
s.t. & \\
c + wb(k) &= a + w \\
\Rightarrow V(a,w,k) &= \log\left[a + w(1 - b(k))\right] + \gamma \log(k)
\end{aligned}
$$

Besides being increasing, I assume that for any $k_1 > k_0$, $b(k)$ (time cost of children) satisfies the following

1. $V(0, w, k_0) > V(0, w, k_1)$
2. $\frac{1-b(k_0)}{1-b(k_1)} > \frac{a+w(1-b(k_0))}{a+w(1-b(k_1))}$

Assumption 1 states that if non-labor income is zero, the status quo in terms of family size is always preferred. The second assumption is a restriction on the way $b(k)$ affects the budget constraint of the household in terms of resources and time. Both assumptions are restrictive but provide unambiguous results in the examples below. Once we depart from these assumptions, however, answers must come from a quantitative exercise.

Deterministic Fertility Choice: In the first stage, and given a startup number of kids k_0, the fertility problem is simply

$$v^f = \max\{V(a, w, k_0), V(a, w, k_1)\}$$

with $k_1 > k_0$. The two lemmas below show that the optimal policy function for kids is a step function, that jumps from k_0 (low fertility) to k_1 (high fertility) depending on both wages and non-labor income.

Lemma 3.1 *There exists a unique* $w^*(a, k_0)$ *such that* $V(a, w^*, k_0) = V(a, w^*, k_1)$

Proof In the Appendix. ■

Lemma 3.2 *There exists a unique* $a^*(w, k_0)$ *such that* $V(a^*, w, k_0) = V(a^*, w, k_1)$

Proof In the Appendix. ■

The first lemma says that below some threshold wage w^*, the optimal choice is to have high fertility k_1. This is the standard negative income-fertility result. On the other hand, lemma 3.2 shows an opposing, "nesting effect": above some threshold a^* of non-labor income, individuals would choose higher fertility. Note that in a life-cycle setting, non-labor income can be thought of as savings from previous periods. Hence, the final income-fertility relationship cannot be derived as straightforward as in the static case. This is true in general, when non-labor income and labor earnings are positively correlated.

I will use a similar structure for the full quantitative exercise below. However, this basic framework is not suited to account for heterogeneity in fertility across individuals with the same wage or level of non-labor income. Below I introduce stochastic fertility and imperfect control and show how this extension provides a natural framework to understand the facts.

Stochastic Fertility Choice: Now, assume that individuals must exert contraceptive effort x in order to influence the probability of no-conception $\pi : \mathbb{R} \to (0, 1)$, which is an increasing and concave function. They also face some utility cost $c(x)$ of exerting effort, which is always positive, increasing and convex. Then, the problem during the fertility stage is

$$v_s^f = \max_x \ \ \pi(x)V(a, w, k_0) + [1 - \pi(x)] - c(x)$$

Using the first order condition from this problem as well as assumptions 1 and 2 from before, it can be shown that $\partial x/\partial w > 0$ and $\partial x/\partial a < 0$. Hence, if we define the expected fertility outcome given optimal effort x^* as

$$k_s^*(a, w) = \pi(x^*)k_0 + [1 - \pi(x^*)]\, k_1$$

we get that

$$\begin{aligned} \frac{\partial k_s^*}{\partial w} &< 0 \\ \frac{\partial k_s^*}{\partial a} &> 0 \end{aligned}$$

The optimal policy functions for fertility choice and their relation to both wages and non-labor income are depicted in figures 3 and 4 respectively. In both figures, the deterministic case is shown as a step policy function while the stochastic case is a smooth one.

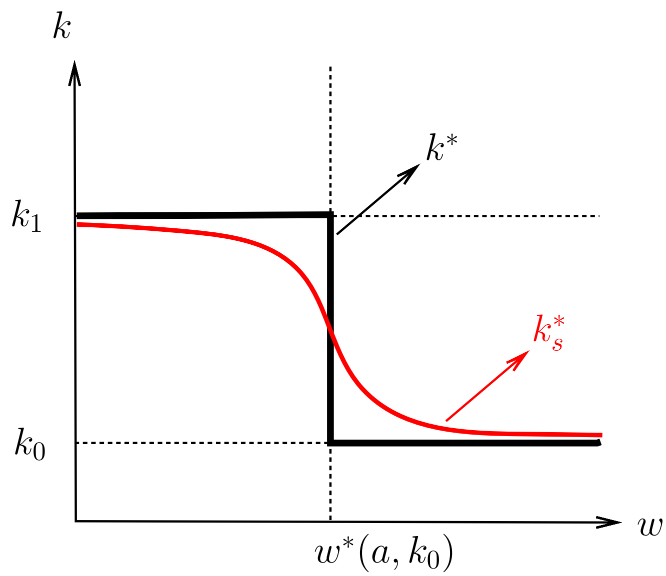

Figure 3: Optimal fertility choices with respect to wages

If we consider an economy populated by a continuum of individuals facing the same problem, stochastic fertility and imperfect control produce a non-degenerate fertility rate (unlike the deterministic case, where the fertility rate is a fixed number): it is an endogenous distribution that depends on incentives and the shape of both $\pi(x)$ and $c(x)$.

As I showed in the previous section, the sign of the wage-fertility relationship is ambiguous if we let labor and non-labor income to be positively correlated (as it is the case within educational groups), since they act as two forces influencing fertility in opposite directions. In the stochastic setting, these forces act in the same way on the optimal contraceptive effort, thus the level of contraceptive failures by skill group cannot be assessed either.

One way of rationalizing the educational fertility gap then, would be to implement a model with heterogeneous preferences for children.[8] However, heterogeneous preferences can not explain higher levels of error in fertility plans by educational group (abortions and unplanned pregnancies). The alternative I propose is to allow for differential effectiveness of birth control effort on fertility outcomes. This approach has been proposed before by Rosenzweig and Schultz (1989) and in the setting below, it also helps in matching the facts on abortion.

[8] See for example, Jones, Schoonbroodt, and Tertilt (2008)

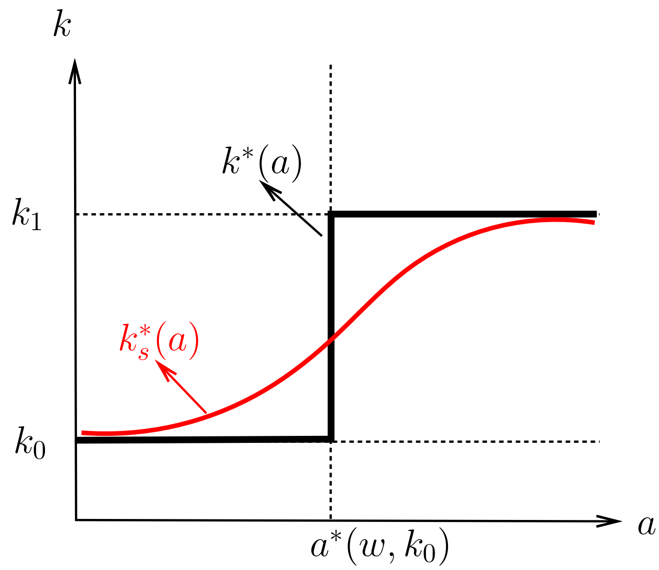

Figure 4: Optimal fertility choices with respect to non-labor income

4 A Quantitative Model

The model environment is an economy populated by agents of different gender (males and females) and education level (high and low) who live finite lives and face three types of exogenous and idiosyncratic shocks: to their life (survival shocks), to their household type (marital transition shocks) and to their earnings (shocks to the value of their market rewards). All agents derive utility from consumption and from the presence of children in the household. Agents supply labor inelastically to the market before retirement and every period they decide how much to consume and save for the future; they cannot borrow.

During the first part of their life-cycle, female agents are fertile (can conceive children) and decide on contraceptive efforts period by period. This effort influences imperfectly the probability of conception. Unwanted pregnancies can be aborted; both contraceptive effort and abortions come at a utility cost. After a birth, female agents must spend some time at home rearing their children. Male agents are not affected by this requirement.

State space. Let z be the state space that defines an agent in this economy. Throughout the discussion, I focus on the female's point of view:

$$z = \{e, e^*, i, k, m, \epsilon, \epsilon^*, i^*, a\} \tag{1}$$

Asterisks represent values for spouses (when applicable). Age is indexed by $i = \{i_0, ..., I\}$, $k = \{1, 2, ..., K\}$ represents the number of children living in the household (not the same as parity), $m = \{1, 2, 3\}$ is the type of household (1 = single, 2 = married, 3 = widowed/divorced[9]), $e \in \{\underline{e}, \overline{e}\}$ represents the education type of the agent (low, high), ϵ is the value of the multiplicative shock to labor earnings and a is the amount of real assets in the household.

For ease of exposition, in some sections of the paper I use the following partition of the sate space $\widetilde{z} = \{m, e, e^*, \epsilon, \epsilon^*, i^*\}$ so that $z = \{i, a, k\} \times \widetilde{z}$.

The Life-cycle proper. All agents start life at age i_0 (first year of adulthood) being one of two educational types: low ($\underline{e}$) or high ($\overline{e}$). This type doesn't change and can be considered as a decision taken before the events in the model. Agents can also start life as married or single and with or without children.

The maximum lifespan for all agents is of I years. Survival from age i to $i+1$ is subject to state dependent mortality risk, i.e., the probability of surviving an additional year depends on the gender and the educational type of the agent. I denote this probability as $\delta_{i,e}$. The probability for males is $\delta^*_{i^*,e^*}$

With regard to labor markets, agents work until they reach age i_r. The retirement age is common for males and females. Female agents also make fertility decisions from i_0 to i_f, the last fertile age. This cut-off for the fertile period is common and known to all female agents.

Fertility and children. During their fertile years, females choose effort to determine the probability of a pregnancy. I denote this effort as $x \in \mathbb{R}$, which translates into a probability $\pi(x|i, m, e) \in (0, 1)$ of *no* conception. This stochastic production function of no pregnancies depends on the age of the female agent (to capture biological constraints on women's reproductive systems), her marital status (since conception opportunities might differ if a mate is present or not) and her education. Evidence of this last point is in Rosenzweig and Schultz (1989), who estimate differential effectiveness rates of contraceptive use by educational attainment. The exertion of this effort comes at a utility cost of $C(x)$.

With complementary probability $(1 - \pi)$, a pregnancy occurs. If the pregnancy falls into the category of "unplanned/unwanted" (i.e., a positive amount of contraceptive effort was exerted), agents have the opportunity of getting an abortion at a utility cost κ_e. This cost depends on the educational level of the agent. If the pregnancy is intended (i.e., $x < 0$) the agent keeps the

[9]Features of widowed vs. divorced households are unified in a single state, since their distinctions in the data are not significant

child and the household increases its size by one.[10]

I make the assumption that children are attached to female agents. I don't keep track of the age nor the sex of children in the household for reasons of computational burden. Instead, households face a constant hazard rate for the permanence of children in the household. I denote this hazard by s_k, which means that on average, children spend $1/s_k$ periods attached to their mothers. Finally, no children can stay in the household after retirement of the mother.

Marital states. The transition through different marital status is stochastic and exogenous. The probability of going from m to m' (conditional on both spouses being alive, in case of agents being married) is given by $\Gamma_{i,e}(m'|m)$. I assume that mortality shocks hit the household before marital transition shocks.[11]

Markets. Agents can sell their time to a spot market for labor, receiving a fixed price of w. They can also save positive amounts of resources, i.e., they can rent assets for the market rate r.

Labor endowments. Agents are endowed with state dependent efficiency profiles, $\varepsilon_{i,m,e}$ for females and $\varepsilon^*_{i^*,m^*}$ for males. They also face idiosyncratic and persistent multiplicative income shocks (ϵ and ϵ^*). The processes generating these shocks are also state dependent. Hence, for males of age i^*, marital status m^* and education level e^*, labor income is given by

$$w\epsilon^*\varepsilon^*_{i^*,m^*,e^*}$$

Note that w is the market rental rate for efficiency units of labor. On the other hand, if children are present in the household, females need to devote some time taking care of them. These time requirements are reflected in $b(m,k) \in (0,1)$, so that labor income of females/mothers is given by

$$w\epsilon\varepsilon_{i,m,e}(1-b(m,k))$$

Since I don't keep track of ages of children in the household, $b(m,k)$ is not time dependent. This simplifying assumption is in contrast of evidence that children require more time and money as they grow old.[12]

[10] There is no child mortality risk nor multiple births

[11] This is important to calculate expectations over future states

[12] For example, see Attanasio, Low, and Sánchez-Marcos (2008)

Preferences. Agents in the economy derive utility from per period consumption and the number of kids in the household. Hence, children are treated as durable goods in terms of utility and their characteristics (such as age and sex) are not qualities that enter agents utility function. In this paper I restrict attention to preferences that are separable in consumption and number of children of the form

$$u(c|z) + \gamma g(k)$$

Preferences for consumption depend on the characteristics of the household (z), namely, the number of members living under the same roof. This is to capture economies of scale in consumption and the idea that marriage might create consumption habits.[13]

Since the focus of this paper is on females and fertility, utility of married households is taken to be that of the female member. This could be the result of using unitary theories of the household or theories that allow for intra-household bargaining and the female having all the bargaining power. This assumption is restrictive, but necessary to keep this a feasible exercise.

Agents in this economy don't have the ability/desire of leaving bequests upon death and don't receive utility from their children once they leave the household.

The Dynamic problem when fertile. There are three distinct stages in the life-cycle of a female agent: (1) work-fertile stage, (2) work - infertile stage and (3) Retirement. Figure 5 presents the timing of events during the first stage, when females make both fertility and consumption/savings choices.

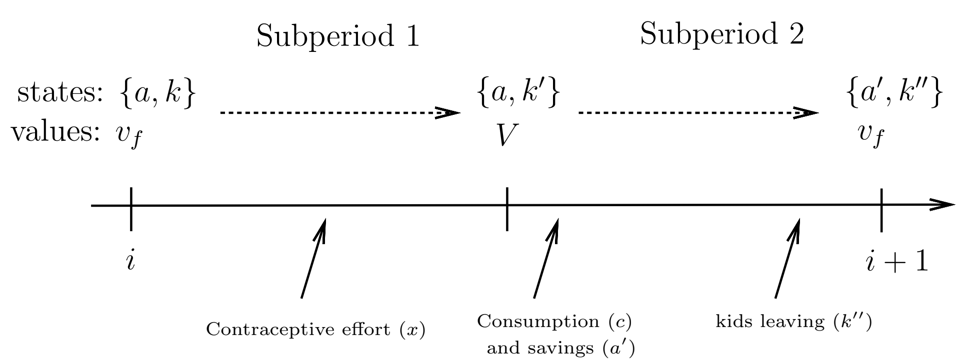

Figure 5: Time ine of Events during fertile years

[13]See Hong and Ríos-Rull (2007).

As seen in the figure, agents in this stage transit between subperiod 1, where they make fertility decisions and subperiod 2, where they choose consumption and savings for the future. Before transiting to subperiod 1 again, households face an updating in their stock of children (due to kids leaving their mothers).

The following bellman equation represents the problem of agents during sub-period 2 (once they have made contraceptive effort choices):

$$
\begin{aligned}
V(i,a,k,\widetilde{z}) &= \max_{c,y} u(c|z) + \gamma g(k) + \delta_{i,e}\beta E\left[v_f(i+1,a',k',\widetilde{z}')|z\right] && (2)\\
st &: \\
c+y &= (1+r)a + w\epsilon_i\varepsilon_{i,m,e}(1-b(m,k)) && \text{if } m=\{1,3\}\\
c+y &= (1+r)a + w\epsilon_i\varepsilon_{i,2,e}(1-b(2,k)) + w\epsilon^*_{i^*}\varepsilon^*_{i^*,2,e^*} && \text{if } m=2, i^*<i_r\\
a' &= \Phi(y,z'|z)
\end{aligned}
$$

Where m represents current marital status ($m = 1,2,3$ stands for single, married and widowed/divorced respectively). The budget constraint accounts for different states, since married agents receive extra income from their spouses' labor, but only if the spouse is not retired ($i^* < i_r$). The Φ operator translates the amount of savings into next period assets given marital transitions and future states.[14]

Given optimal policies in subperiod 2, females make contraceptive effort choices in subperiod 1. The problem faced by them is:

[14] The particular form of Φ is given by:

$$
\Phi(y,z'|z) = \begin{cases} y & if \quad (m'=2|m=2) \\ y & if \quad (m'=1,3|m=1,3) \\ y & if \quad (m'=3|m=2) \text{ (widowhood)} \\ \chi y & if \quad (m'=3|m=2) \text{ (divorce)} \\ y+a^* & if \quad (m'=2|m=1) \end{cases} \qquad (3)
$$

where (m',m) refers to a transition from m to m' next period. For example, when going from $m=2$ (married) to $m=3$ (through divorce), assets next period are a fraction χ of what is saved today, where $\chi \in (0,1)$ reflects the partition of assets after a divorce. Note that when going from $m=1$ (single) to $m=2$ (married), assets next period are given by current savings plus what the prospective spouse brings to the household. This last variable (a^*) is a random variable that depends on the distribution of single agents of the opposite sex in the economy.

$$
\begin{aligned}
v_f(i,a,k,\widetilde{z}) \;&=\; \max_x \;\; \pi(x|i,m,e)V(i,a,k,\widetilde{z}) \\
&\quad + [1-\pi(x|i,m,e)] \max \left\{ \begin{array}{c} V(i,a,k+1,\widetilde{z}), \\ V(i,a,k,\widetilde{z}) - \kappa_e \end{array} \right\} \\
&\quad - \; C(x)
\end{aligned} \tag{4}
$$

The value function at this stage is a convex combination of the continuation values with and without a new pregnancy. In the case of pregnancy (which occurs with probability $(1-\pi(\cdot))$), agents have the chance of having an abortion at utility cost κ_e. Note that even though there are discrete outcomes following this stage (number of children in the household), the effort function convexifies the problem maintaining smoothness of the value function, which proves useful for solving (3) using standard continuous methods.[15]

My approach to model fertility choices differs from those who try to understand choices of specific birth control methods by women.[16] The setup above doesn't distinguish between different contraceptive methods nor their efficacy, but is general and its implementation straightforward.

Moreover, I allow the probability of no conception to be flexible enough so that overall fertility is not only due to failed birth control but also as the result of conscious efforts of females to start a family. Specifically, this means that the domain of π is the entire real line (contraceptive effort can be negative, in order to maximize the probability of conception) and the cost function is always positive, increasing away from zero. This general specification allows me to capture biological constraint on human fertility, which play a role in determining the optimal timing of births later in life.

The dynamic problem after fertile years. Once agents are past the fertile stage (cannot produce more children), they keep choosing optimal paths for consumption and savings until death. This stage in the life-cycle can also be divided into two: before and after retirement.

Before retirement ($i \leq i_r$), the problem of the agent is:

[15] Details of the numerical solution procedure are in the Appendix.

[16] See for example Hotz and Miller (1993) and Rosenzweig and Schultz (1985).

$$
\begin{aligned}
V(i,a,k,\widetilde{z}) &= \max_{c,y} u(c|z) + \gamma g(k) + \delta_{i,e}\beta E\left[V(i+1,a',k',\widetilde{z}')|z\right] \qquad (5)\\
st &: \\
c+y &= (1+r)a + w\epsilon_i\varepsilon_{i,m,e}(1-b(m,k)) \qquad \text{if } m=\{1,3\}\\
c+y &= (1+r)a + w\epsilon_i\varepsilon_{i,2,e}(1-b(2,k)) + w\epsilon^*_{i^*}\varepsilon^*_{i^*,2,e^*} \qquad \text{if } m=2, i^*<i_r\\
a' &= \Phi(y,z'|z)
\end{aligned}
$$

The main difference between this Bellman equation and the one in (6) is that the continuation value is the same function V and the stock of children can only decrease from period to period.

After retirement, the problem reduces to

$$
\begin{aligned}
V(i,a,0,\widetilde{z}) &= \max_{c,y} u(c|z) + \gamma g(k) + \delta_{i,e}\beta E\left[V(i+1,a',0,\widetilde{z}')|z\right] \qquad (6)\\
st &: \\
c+y &= (1+r)a \qquad \text{if } m=\{1,3\}\\
c+y &= (1+r)a + w\epsilon^*_{i^*}\varepsilon^*_{i^*,2,e^*} \qquad \text{if } m=2, i^*<i_r\\
a' &= \Phi(y,z'|z)
\end{aligned}
$$

at this stage no children are present in the household ($k=0\,\forall i\geq i_r$) and the only resources available for non-married agents are past savings. On the other hand, if agents are married to working age individuals, they enjoy the extra labor income $w\epsilon^*_{i^*}\varepsilon^*_{i^*,2,e^*}$.

5 Taking the Model to the Data

The solution of this model is a set of policy functions $x^{opt}(z|\Theta), y^{opt}(z|\Theta)$ for contraceptive effort and savings respectively, given the current state z and other parameters, Θ (including prices). As it's usual, analytical expressions for the optimal policies are unfeasible, so I approximate them using numerical solutions to an empirical model with the following quantitative features.

Demographics and life-cycle. All agents start life at age 18 and cannot live longer than 95 years. Retirement is at 65 and the last fertile age is 40. A model period is one year when

$i \in \{18, ..., 40\}$, 5 years when $i \in \{40+1, ..., 65\}$ and 10 years when $i \in \{65+1, ..., 95\}$.[17]. Age specific mortality rates are taken from the National Center for Health Statistics and adjusted for educational attainment, as in Hong (2006).[18]

I divide educational or skill types into those with at most a high school diploma or GED, and those with some post secondary education (college, community college, vocational school, etc.). To calculate the proportion of these types, I use the Current Population Survey (CPS) between 1990 and 1995. The proportion of high school individuals is around 40%. The majority of agents start life as single and childless, but I allow some of them to be married and have children. The proportion of married 18 year old females in the CPS is around 93% and females with kids is around 9%. When performing simulations of the model, I distribute women uniformly according to these statistics to determine their initial state.

Since non-married females can always find a (new) partner in the model, I need information on who they'd marry. Also from the CPS, I compute the proportion of couples by age and educational attainment of the partners, the age distribution of male partners for married females and the relative asset position of both non-married males and non-married females.[19] Given this information, I construct education-specific grids with probabilities of marrying someone of characteristics given by $\{e^*, i^*, a^*\}$ (education, age and assets of prospective husbands). Since I'm not computing equilibrium, this procedure doesn't check for internal consistency of measures of agents (as in Hong and Ríos-Rull (2007), where all these probabilities are endogenous objects).

Transitions between marital states come from the Panel Study of Income Dynamics (PSID) for the years 1990-1995. I follow all heads of household older than 18 years old (inclusive) and compute annual age and education specific transition probabilities between three states: single, married and divorced/widowed.[20] Given variable specification in the PSID, married couples include cohabitating couples.

Preferences. I use an additively separable specification for instantaneous, per period utility: $u(c|z) + \gamma g(k)$. The marginal utility from consumption depends on the size of a household:

[17] I do this to reduce the state space of the model. For details, see the Appendix.

[18] Given the mapping from model periods to actual years, all age specific variables used in the computation are recalculated, depending on the stage of model life-cycle. Details in the Appendix.

[19] My proxy for individual assets is the sum of interest, dividend and rent income as defined in the March supplements of the CPS

[20] Some transitions are theoretically zero (for example, from single to widowed/divorced or from married to single.) and others are complement to each other.

$$u(c|z) \equiv \widetilde{u}\left(\frac{c}{1+\mathbf{1}_{\{m=2\}}\phi_m+\mathbf{1}_{\{k>0\}}k\phi_k}\right) \tag{7}$$

Where $\mathbf{1}_{\{cond\}}$ is the indicator function that takes a value of one when "$cond$" is true and zero otherwise; ϕ_m and ϕ_k are equivalence scales which discount consumption in married households and households with children respectively. If $\phi_m, \phi_k < 1$, economies of scale in consumption exist in the household: expenditures to maintain the level of per capita utility constant, grow proportionally less than the number of household members.

The specific functional forms for $\widetilde{u}$ and g are given by

$$\widetilde{u}(\mathbf{c}) = \frac{\mathbf{c}^{1-\eta_c}-1}{1-\eta_c} \qquad\qquad g(k) = \frac{(1+k)^{1-\eta_k}-1}{1-\eta_k} \tag{8}$$

Fertility. I use the following function for π (the probability of NO conception, given effort x):

$$\pi(x|i,m,e) = \begin{cases} \pi^{+}(x|i,m,e) & if \quad x>0 \\ \pi^{-}(x|i,m,e) & if \quad x<0 \end{cases} \tag{9}$$

where

$$\pi^{+}(x|i,m,e) = \frac{\exp\{x\}}{\exp\{x\}+\varphi^{+}_{i,m,e}\exp\{-x\}}$$

and

$$\pi^{-}(x|i,m,e) = \frac{\exp\{x\}}{\exp\{x\}+\varphi^{-}_{i}\exp\{-x\}}$$

In general, π is a modified logistic function with φ as a shift parameter. Regalia and Ríos-Rull (2001) use a similar framework to study fertility choice in an equilibrium model. Note that the higher $\varphi^{+}_{i,m,e}$, the higher the probability of a pregnancy when effort (x) is positive (females trying to avoid fertility), which means that I can parameterize higher difficulty in controlling fertility by increasing $\varphi^{+}_{i,m,e}$. However, if women are trying to get pregnant (negative x), parameterizing π through the same $\varphi^{+}_{i,m,e}$ would not be realistic, since it would mean that ability in using contraceptive methods is negatively correlated with the ability of procuring a conception when it is desired. Hence, I use a different shifter, φ^{-}_{i} for this case. Note that contraceptive ability $\varphi^{+}_{i,m,e}$ depends on age, marital status and education and its parameterization is given by

$$\varphi^+_{i,m,e} = \widetilde{\varphi}^+_{i,m} + \mathbf{1}_{\{e=\underline{e}\}}\overline{\varphi}_i$$

While ability in using contraceptive technology might depend on marital status and education, I assume that the technology of procuring a pregnancy depends mostly on biological constraints. This is represented by φ^-_i depending only on age.

I parameterize these age profiles using polynomial approximations on age: given the order for the polynomial (p, the same for all profiles) the number of parameters to be determined is then $p+1$ (the number of polynomial coefficients) times 5 profiles: for singles ($\widetilde{\varphi}^+_{i,1}$), for married ($\widetilde{\varphi}^+_{i,2}$), for divorced/widowed ($\widetilde{\varphi}^+_{i,3}$), for the extra risk faced by the high school group ($\overline{\varphi}_i$) and for the biological fertility profile φ^-_i. For example, the profile for married individuals is given by

$$\widetilde{\varphi}^+_{i,2} = \alpha^2_0 + \alpha^2_1 i + \alpha^2_2 i^2 + \ldots + \alpha^2_p i^p$$

while the profile for the excess fertility risk faced by the high school group is

$$\overline{\varphi}_i = \overline{\alpha}_0 + \overline{\alpha}_1 i + \overline{\alpha} i^2 + \ldots + \overline{\alpha}_p i^p$$

I reduce the number of parameters by assuming that the fertility control technology for singles is the same than for divorced/widowed agents ($\widetilde{\varphi}^+_{i,1} = \widetilde{\varphi}^+_{i,3}$).

On the other hand, I parameterize the utility cost of exerting contraceptive effort as

$$C(x) = \frac{x^2}{2}\xi$$

This cost function is symmetric around zero, so I use it for both sides of the fertility problem: when females want to prevent or are seeking a pregnancy. This is not restrictive, given the asymmetric structure of π.

Earnings and Labor Supply. Endowments of labor efficiency profiles come from the CPS (years 1990-1995). I calculate annual labor earnings for the two educational groups (high school and college), by age and marital status. As in Hong and Ríos-Rull (2007) and Hong (2006), I use annual earnings since they capture differences in the intensive margin of earnings by sex and marital status better than hourly earnings. To account for inflation, I adjust nominal values by the GDP deflator for the year 2000.

I restrict attention to childless females throughout the sample period. For males, I don't make that distinction, since the change in income due to the presence of own children in the household is not significant.

I attribute the time cost of child-rearing $b(m,k)$ to annual labor income differentials of females in fertile age (18 to 40) by number of children. This is different than accounting for hours worked by number of children in the household; it stands alternatively for different ways in which a child might change earnings ability of the mother (e.g., getting a job with more flexible schedule but lower pay, getting a job with lower pay but closer to home, not getting tenured at an academic job or not being made partner at the firm, etc.) other than through hours worked per week. The computed values are in table 1.

Children	$b(m=\{1,3\},k)$	$b(m=2,k)$
0	0%	0%
1	5.9%	26.5%
2	16.9%	37.5%
3	41.0%	52.6%
4	61.3%	63.3%
5+	81.2%	72.8%

Table 1: Time cost of Children (in terms of full time work), CPS 1990-1995.

As seen from the table, time cost of children (or time away from the best paid market alternative) is increasing in the number of children present in the household. Note also that the cost increases faster in the number of kids for married women than for single ones.

For earnings shocks, I use an AR(1) specification

$$\epsilon_e' = \rho_e \epsilon_e + \mu_e' \tag{10}$$

where $\mu_e \sim N(0, \sigma_e)$. These shocks are gender and education specific. I take values of ρ_e, σ_e (for $e = \{\underline{e}, \overline{e}\}$) from Hong (2006), who uses the PSID between 1986-1992 to compute maximum likelihood estimates. As is common, I discretize both continuous processes using the method proposed by Tauchen (1986).

6 Estimation

Given the partial equilibrium nature of the exercise, I set several model parameters exogenously. First, the rental price of efficiency units of labor w is normalized to 1. I set the interest rate to equal the average of the 1-year Treasury Bill Rate (monthly auction averages).[21] I let the discount factor β to be $1/(1+r)$. For equivalence scales, I use $\phi_m = 0.7$ and $\phi_k = 0.5$ (i.e., the OECD values).

The rest of the model parameters are determined jointly, by minimizing the square difference between data and model moments. The procedure is standard in the literature: (i) select which data targets to use (ii) guess values for model parameters (iii) solve the model and calculate optimal policies (iv) simulate life-cycles for a large number of individuals and compute model equivalents to the data targets (vi) calculate the error of the iteration (the sum of square values of the difference between every data and model moment) (vii) if the error is less than a pre-specified tolerance, exit; if not, update parameters according to some predefined rule and repeat from step (iii) until convergence. This is a simplified simulated method of moments estimation procedure, where the weighting matrix for moments is the identity matrix.

The list of moments is as follows:

- Age profile of pregnancy rates for non-married females by education[22]: 46 moments = 23 ages $\times$ 2 education levels
- Age profile of pregnancy rates for married females by education: 46 moments = 23 ages $\times$ 2 education levels
- Age profile of abortion rates by education: 46 moments = 23 ages $\times$ 2 education levels
- Age profile of unplanned pregnancy rates by education: 46 moments = 23 ages $\times$ 2 education levels

In total, there are 184 moments to match. On the other hand, the number of model parameters depends on the choice of order for the polynomials that define age profiles for fertility parameters. I chose $p = 6$, so each age profile is defined by 7 coefficients. Thus, the model has 34 parameters to be determined jointly:

- curvature in the utility of consumption η_c (1)
- curvature in the utility of children η_k (1)
- multiplicative parameter in utility of children γ (1)

[21]Series id TB1YA, on the St. Louis Fed Economic Data webpage.

[22]Note that I merge the statistics for both single and widowed/divorced females.

- utility cost of an abortion κ_e (2)
- utility cost of contraceptive effort ξ (1)
- contraceptive ability parameter $\widetilde{\varphi}_{i,m}$ (14 parameters = 7 coefficients $\times$ 2 marital states)
- conception ability parameter ϕ_i^- (7)
- contraceptive ability shifter for low skill/education group $\overline{\varphi}_i$ (7)

Solution to the model is by backwards recursion. In the last period of life there is no continuation value (I assume no bequests motives nor life insurance) hence optimal policies and value functions can be calculated recursively from the next to last period. Details of the procedure are in the Appendix.

7 Results and Experiments

The estimated parameters are in table 2. The full list of coefficients for the polynomial functions used in parameterizing conception probabilities is in the Appendix. In figure 6, I present instead the implied age profiles given by the estimated polynomials.

Parameter	value
η_c	1.79
η_k	1.45
γ	0.77
κ_{HC}	3.95
$\kappa_{College}$	1.42
ξ	0.95

Table 2: Model Parameters

The estimated curvature in the utility of consumption (η_c) equals 1.78, which is in line with the rest of the literature (the usual number lays between 1.5 and 2). Preferences are close to being homothetic: the value of η_k (1.45) is close to the one for η_c. Overall, the value of both parameters indicate that consumption and children enter as complements in the utility function, so females enjoy the presence of children more when consumption levels are higher. Given increasing wages during the earlier stages of the life-cycle, this means that females would like to postpone childbearing as much as possible.

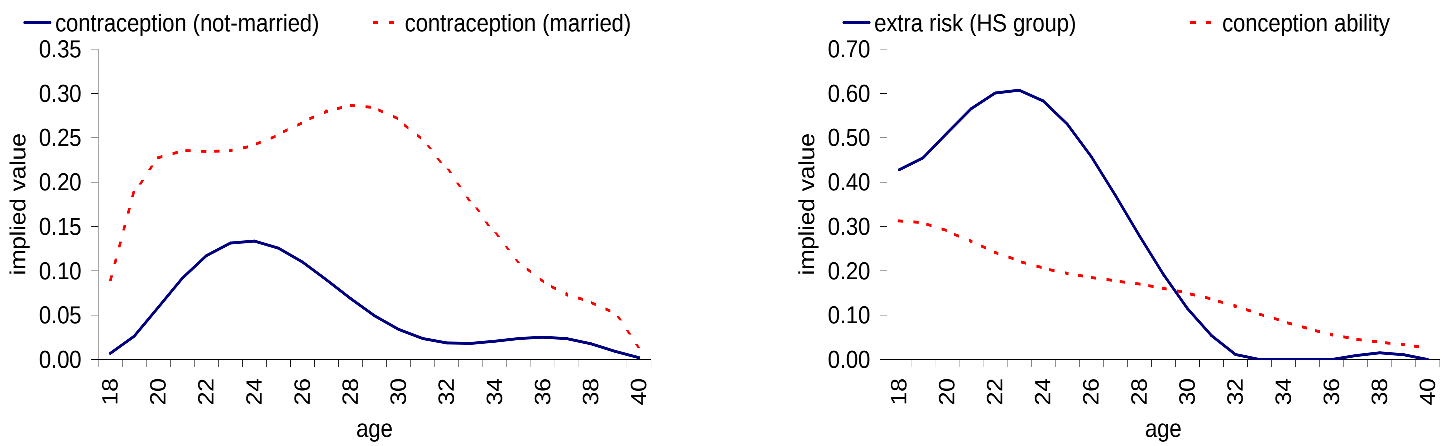

(a) $\widetilde{\varphi}_{i,\{1,3\}}$ and $\widetilde{\varphi}_{i,2}$

(b) $\overline{\varphi}_i$ and ϕ_i^-

Figure 6: Estimated parameters for contraceptive technology and conception ability

The utility cost of an abortion is around two and a half times higher for high school individuals than for college individuals (3.95 vs. 1.42 respectively).

Figure 6 shows the age-profiles that describe fertility technology and restrictions during the life-cycle. All profiles are decreasing in age which is a reflection of decreasing chances of conception late in the fertile stage of life. Contraceptive parameters for married individuals are higher than for single ones, which means that birth control is easier when there is no steady partner of the opposite sex in the household. With respect to the risk faced by the high school group, results imply that females in that group have 17% more chances (on average in their lifetime) of having an unwanted pregnancy than their college counterparts. The calculation of this percentage comes from using the age profile $\overline{\varphi}_i$: the approximated probability of an unwanted pregnancy (when exerting a very low amount of effort) for every age is given by $\overline{\varphi}_i/1+\overline{\varphi}_i$ (The average for all ages is 17%).

Figures 7 to 9 show the goodness of fit of the model.

Overall, the model does a good job in replicating the stylized facts with respect to the number and timing of births across educational groups. Both simulated abortion and fertility rates follow closely their data counterparts; on the other hand, the rate of unwanted pregnancies is overpredicted for the high school group but the overall qualitative features of the data are

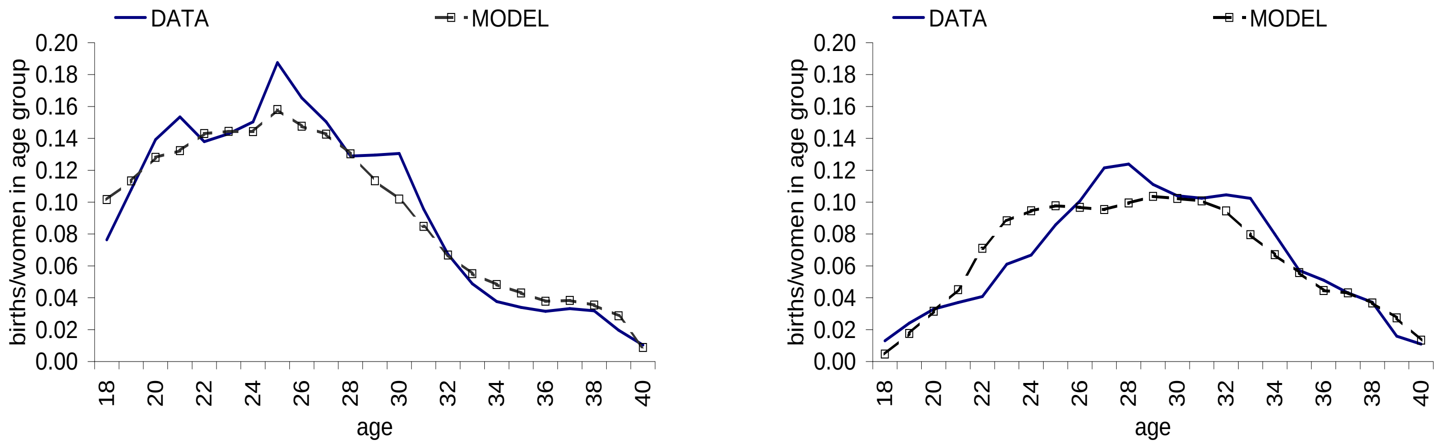

(a) High School

(b) College

Figure 7: Age specific fertility rates: Data and Model

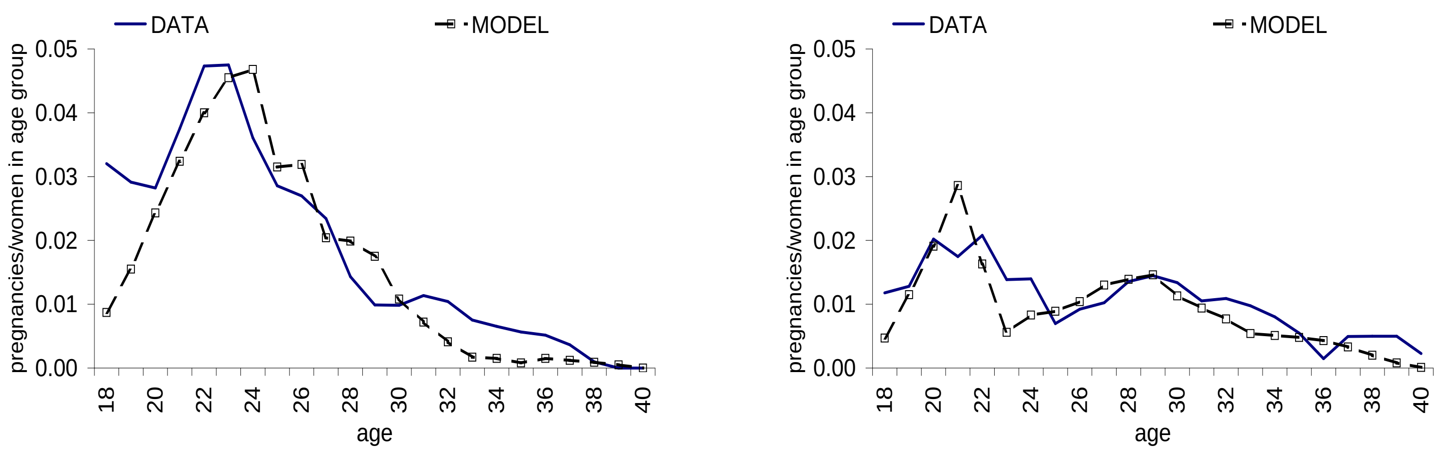

(a) High School

(b) College

Figure 8: Abortion rates: Data and Model

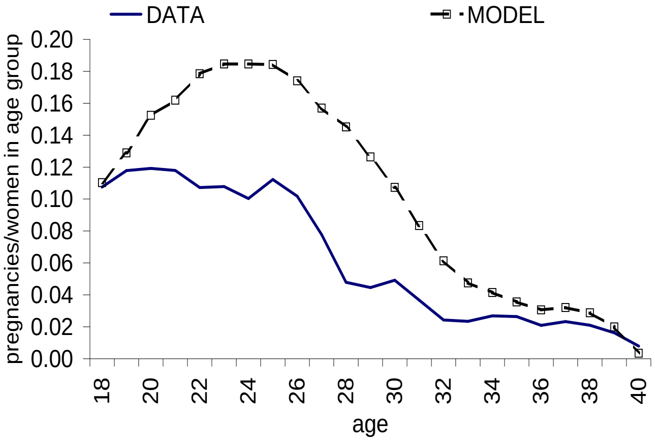

(a) High School

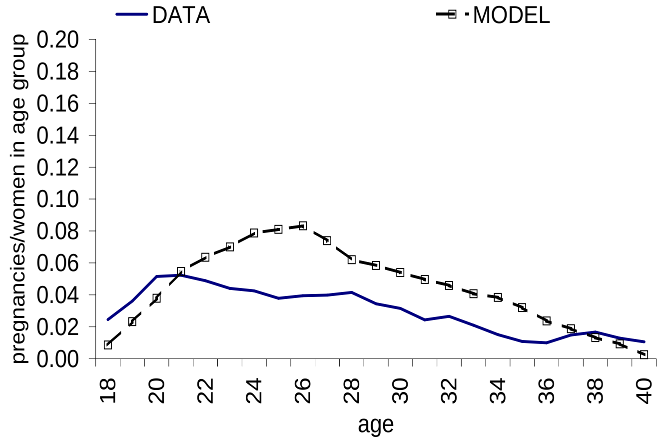

(b) College

Figure 9: Abortion rates: Data and Model

preserved (differences among educational groups).

Quantitative Experiments: Below I present experiments that show the importance of the two main ingredients of my theory, namely differential fertility risk across educational groups and self-insurance (the ability to save for the future).

Figure 10 shows the comparison between the baseline model and the case when the model is calibrated to match the facts without extra risk for the high school group ($\overline{\varphi}_i = 0, \forall i$). As seen from the figure, the latter model is unable to match the higher fertility rates of high school educated individuals, while the rates for the college group are slightly overpredicted. In terms of timing of births, this alternative model correctly accounts for the delay in childbearing by females in the college group as opposed to those in the high school group; however, this difference is less pronounced than in the baseline model.

In the model where the two educational groups face the same contraceptive technology, the ratio between their TFRs is 0.94 (1.56 and 1.67 for high school and college educated individuals respectively) while the ratio in the data (and the baseline) is closer to 0.7. Hence, differential fertility risk earlier in life accounts for most of the differences in the number of births across educational groups during the life-cycle, leaving a small role for differences in wages. This result

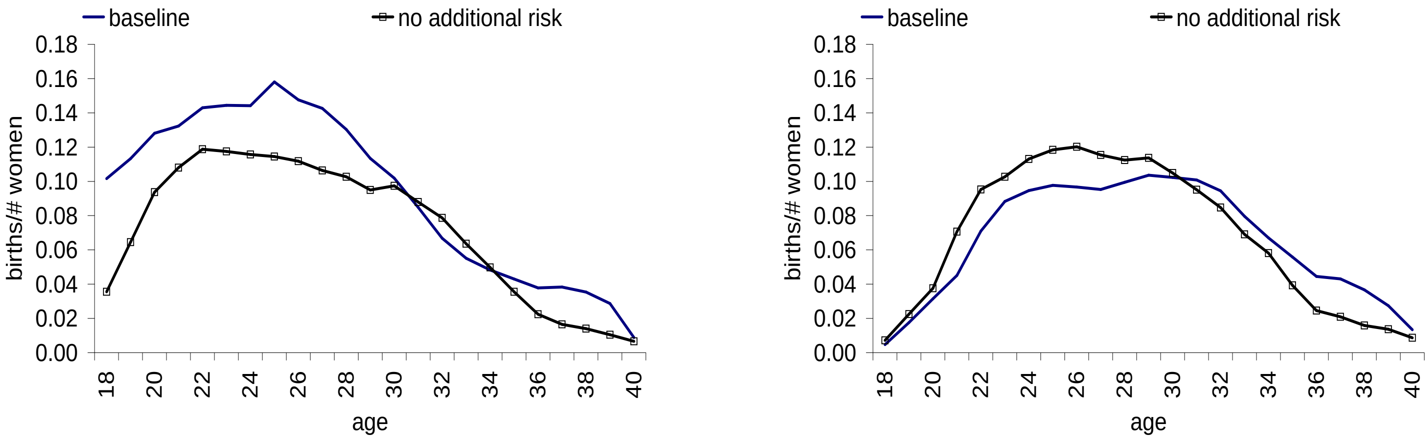

(a) High School

(b) College

Figure 10: Age-specific fertility rates, different models

hints that differences in wage profiles help mostly in predicting the different timing of births: the flatter profiles of life-cycle wages of high school educated females makes them choose early childbirth (at the margin) given the complementarity between consumption and children.[23] On the other hand, college individuals face rapidly growing wage profiles, which induce them to delay fertility.

In the next simulation, I take the baseline model, shut down the ability to save and recalibrate the economy. Since I assume that after retirement agents don't receive any income other than past savings, I set the last period in the no-savings simulation to be i_r and discard the retirement stage (setting the value functions after age i_r equal to zero for any point in the state space).

The predicted fertility profiles of the no-savings case are in figure 11. Results from simulating this version of the model (no savings but with differential fertility risks) shows that the choice of assumption regarding capital markets changes the predictions of fertility models in non-trivial directions. It also helps in understand the intertemporal margins faced by individuals in the life-cycle.

Given the inability to save, all income has to be consumed at the end of each period. For

[23]After re-estimating the parameters, $\eta_c = 1.17$ and $\eta_k = 1.11$, so consumption and children are still complements, but the extent of this complementarity is smaller than in the baseline

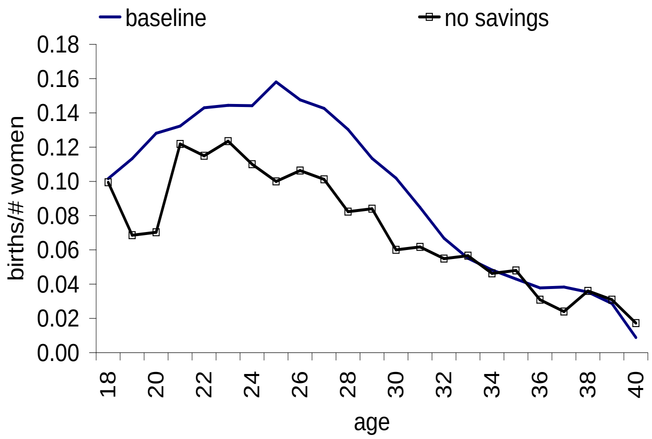

(a) High School

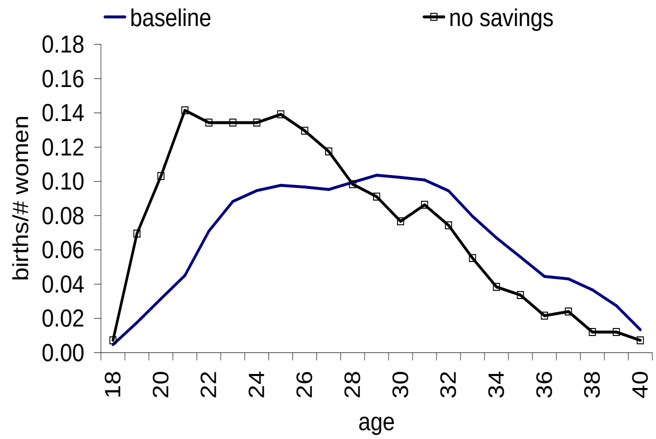

(b) College

Figure 11: Age-specific fertility rates, different models

individuals in both educational groups, consumption in this setting is, on average, higher than what they would have chosen had they been able to save. For individuals in the high school group, this leads to marginal lower fertility rates than in the baseline. The college group, however, faces increasing wage profiles, thus increasing utility. Given complementarity between consumption and children in the utility function, this 'forced' higher consumption for the group leads them to conceive more children to increase utility when wages and consumption are high. Since fertility choice is dynamic and females are restricted to one child per period, the timing of first child birth is shifted towards younger ages, which is a major counterfactual prediction of this variation of the model.

8 Conclusion

In this paper I study life-cycle fertility in the U.S., focusing on birth profile differences across educational groups (high school and college). To understand the facts on timing and number of births during the life-cycle, I develop a structural model where agents transit through different marital states, face idiosyncratic survival and earnings risk and capital markets are incomplete (individuals cannot borrow against their future earnings). In this setting, I embed a standard fertility model (the "time allocation of mothers" variety) and add the assumption of imperfect control of individuals over fertility outcomes. From the analysis, I conclude that differential fertility risk (in the form of ability to control fertility plans) across education groups is the main determinant of differences in timing and levels of fertility, while differences in marriage/labor markets play minor roles. This shows that standard fertility theories, which rely solely on substitution effects to produce negative skill-fertility relationships, cannot account for life-cycle nor cross sectional facts.

References

ATTANASIO, O. P., H. LOW, AND V. SÁNCHEZ-MARCOS (2008): "Explaining Changes in Female Labor Supply in a Life-Cycle Model," *American Economic Review*, 98(4), 1517–42.

BECKER, G. S. (1960): "An Economic Analysis of Fertility," *Demographic and Economic Change in Developed Countries*, 11.

——— (1965): "A Theory of the Allocation of Time," *Economic Journal*, 75, 493–517.

CAUCUTT, E., N. GUNNER, AND J. KNOWLES (2002): "Why do Women Wait? Matching, Wage Inequality and Incentives for Fertility Delay," *Review of Economic Dynamics*, 5(4), 815–855.

CONESA, J. C. (2000): "Educational Attainment and Timing of Fertility Decisions," Working Paper 0010 CREB, UNiversidad de Barcelona.

FU, H., J. DARROCH, S. HENSHAW, AND E. KOLB (1998): "Measuring the extent of abortion underreporting in the 1995 National Survey of Family Growth," *Family Planning Perspectives*, May/Jun 1998.

GALOR, O., AND D. WEIL (1996): "The Gender Gap, Fertility and Growth," *American Economic Review*, 86(3), 374–87.

GRAY, G. A., AND T. G. KOLDA (2006): "Algorithm 856: APPSPACK 4.0: Asynchronous Parallel Pattern Search for Derivative-Free Optimization," *ACM Transactions on Mathematical Software*, 32(3), 485–507.

GREENWOOD, J., A. SESHADRI, AND G. VANDENBROUCKE (2005): "The Baby Boom and Baby Bust," *American Economic Review*, 95(1), 183–207.

HEER, B., AND A. MAUNER (2004): *DGE Models: Computational Methods and Applications*. Springer-Verlag.

HONG, J. H. (2006): "Life Insurance and the Value of Spouses: Labor Supply vs. Household Production," Mimeo.

HONG, J. H., AND J.-V. RÍOS-RULL (2007): "Social Security, Life Insurance and Annuities for Families," *Journal of Monetary Economics*, 54, 118–140.

HOTZ, J. V., J. A. KLERMAN, AND R. J. WILLIS (1997): "The Economics of Fertility in Developed Countries," in *Handbook of Population and Family Economics*, ed. by M. R. Rosenzweig, and O. Stark, pp. 275–347. Elsevier.

HOTZ, J. V., AND R. MILLER (1993): "Conditional Choice Probabilities and the Estimation of Dynamic Models," *Review of Economic Studies*, 60(3), 497–529.

JONES, L. E., A. SCHOONBROODT, AND M. TERTILT (2008): "Fertility Theories: Can They Explain the Negative Fertility-Income Relationship?," NBER working paper No. 14266.

JONES, L. E., AND M. TERTILT (2008): "An Economic History of Fertility in the U.S.:1826-1960," in *Frontiers of Family Economics*, ed. by P. Rupert. Emerald Press (forthcoming).

JUDD, K. L. (1998): *Numerical Methods in Economics*. MIT Press.

KEANE, M. P., AND K. I. WOLPIN (2006): "The Role of Labor and Marriage Markets, Preference Heterogeneity and the Welfare System in the Life cycle Decisions of Black, Hispanic and White Women," PIER working paper 06-004.

KOLDA, T. G. (2005): "Revisiting asynchronous parallel pattern search for nonlinear optimization," *SIAM Journal on Optimization*, 16(2), 563–586.

MINCER, J. (1963): "Market Prices, Opportunity Costs and Income Effects," in *Measurement in Economics: Studies in Mathematical Economics in Honor of Yehuda Grunfeld.*, ed. by C. C. et. al. Stanford University Press.

REGALIA, F., AND J.-V. RÍOS-RULL (2001): "What Accounts for the Increase in Single Households and for the Properties of Fertility?," Mimeo, University of Pennsylvania. First version, 1998.

ROSENZWEIG, M. R., AND T. P. SCHULTZ (1985): "The Demand for and Supply of Births: Fertility and its Life Cycle Consequences," *American Economic Review*, 75(5), 992–1015.

——— (1989): "Schooling, Information and Nonmarket Productivity: Contraceptive Use and its Effectiveness," *International Economic Review*, 30(2), 457–477.

TAUCHEN, G. (1986): "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions," 20, 177–181.

WOLPIN, K. (1984): "An Estimable Dynamic Stochastic Model of Fertility and Child Mortality," *Journal of Political Economy*, 92(5), 852–74.

9 Appendix

9.1 Proofs

Proof of Lemma 3.1. Define the value of not increasing family size

$$\Delta V(a, w, k_0) \equiv V(a, w, k_0) - V(a, w, k_1)$$

existence of w^* comes from continuity of the $\log$ function and the use of the intermediate value theorem. First,

$$\begin{aligned}\lim_{w \to 0} \Delta V &= \log \left\{ \frac{a + w(1 - b(k_0))}{a + w(1 - b(k_1))} \right\} + \gamma \log\{\frac{k_0}{k_1}\} \\ &= 0 + \gamma \log\{\frac{k_0}{k_1}\} \\ &< 0\end{aligned}$$

On the other hand, $\lim_{w \to \infty} \Delta V > 0$, by assumption 1. Hence, there must exist at least one wage such that $\Delta V = 0$. For uniqueness, we require $\frac{\partial \Delta V}{\partial w} \geq 0$, which comes from using assumption 2 ■

Proof of Lemma 3.2. This proof is analogous to the previous one. First, note that $\lim_{a \to 0} \Delta V = \log(k_0/k_1) < 0$. On the other hand, $\lim_{a \to \infty} \Delta V > 0$ by assumption 1, so applying the same logic as above, a^* exists. For uniqueness, we have that

$$\frac{\partial \Delta V}{\partial a} = \frac{1}{a + w(1 - b(k_0))} - \frac{1}{a + w(1 - b(k_1))}$$

which is strictly negative, because $b(k)$ is increasing. ■

9.2 Data

Figure 12 shows the profiles for labor endowments, computed from march supplements of the Current Population Survey (years 1990 to 1995). In the figure I show annual earnings for females, between ages 18 to 65, corrected for inflation using the GDP deflator for the year 2000. These profiles are smoothed using a 5th order polynomial.

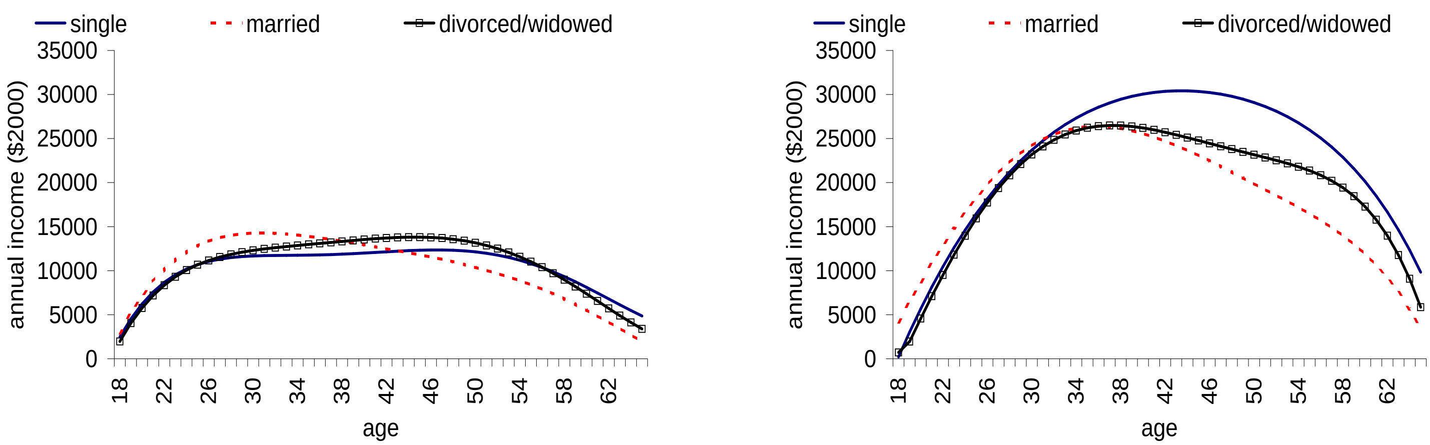

(a) High School

(b) College

Figure 12: Labor Endowments by educational group

To characterize the labor market, I also use gender and education specific idiosyncratic labor shocks. These shocks come from estimates from Hong (2006), who uses labor earnings data from the PSID to calculate the unobserved component of annual labor earnings. I use a standard discretization of the continuous AR(1) described in the paper. I choose to discretize the four processes (2 education groups and 2 genders) by a 3 state markov system. The standard in the literature is to use at least 5 states, but computational burden prevents me from using a more detailed shock structure. However, results in the paper don't rely in the dimensionality of these shocks.

Also from the CPS, I calculate the proportion of females (by education) married to college educated males (irrespective of presence of children in the household), in order to measure positive assortative matching in the marriage market. As seen in figure 13, marriage indeed shows the positive assortative matching property.

I compute yearly survival probabilities by educational group using the information in Hong (2006). I interpolate his 5 year values and smooth the resulting series with a second order polynomial. The resulting probabilities for female individuals are in figure 14.

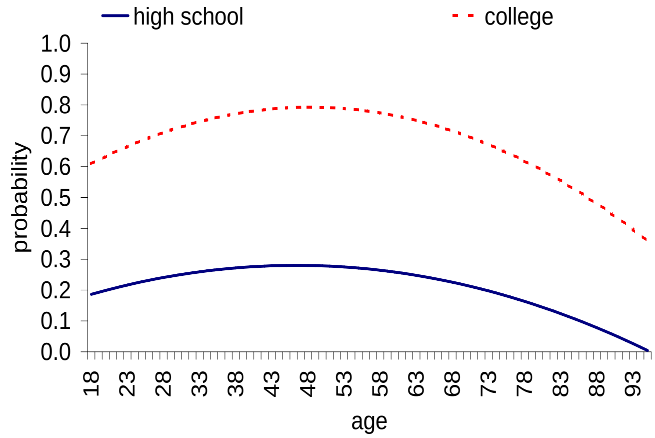

Figure 13: Probability of being married to college male, by education of female

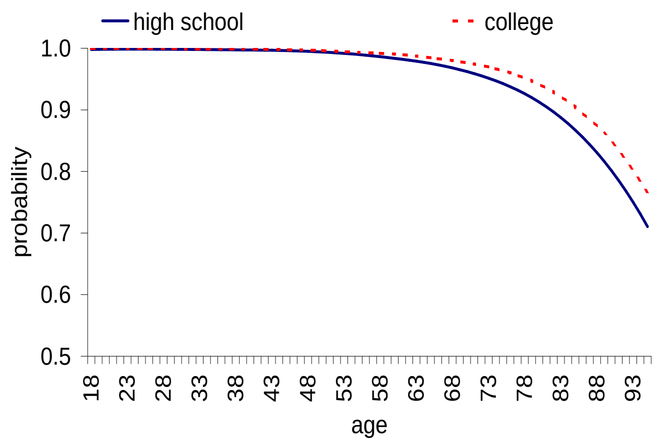

Figure 14: Survival probability by education

To calculate transitions through marital states, I use the Panel Study of Income Dynamics (PSID) for the years 1990 through 1995. I use heads of household and wives (as defined in the PSID) to compute the following probabilities, by education and age: probability of remaining single, the probability of remaining married and the probability of getting married conditional on being divorced/widowed. Given these three probabilities, I can span all transitions (e.g., some probabilities are zero by definition and others are just complements). I extrapolate these probabilities when necessary since the PSID doesn't have many observations for young/old heads of household. Given the short span of my chosen sample, individuals contribute at most 5 observations/years, making these probabilities a cross-section description of marital transitions during the mid 1990s in the U.S. Figure 15 shows these transitions.

I assume simple age and asset distribution of prospective male partners. For ages I consider only 3 possible alternatives: same age, one year older and two years older ($i^* = \{i, i+1, i+2\}$), each occurring with probabilities $P(i^* = i) = 0.4$, $P(i^* = i+1) = 0.41$ and $P(i^* = i+2) = 0.19$,

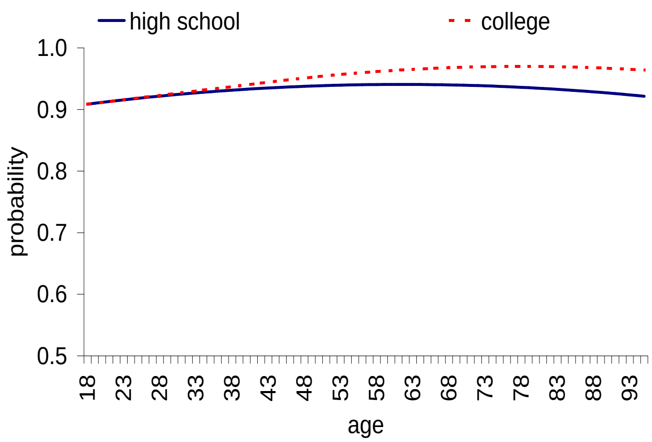

(a) single to single

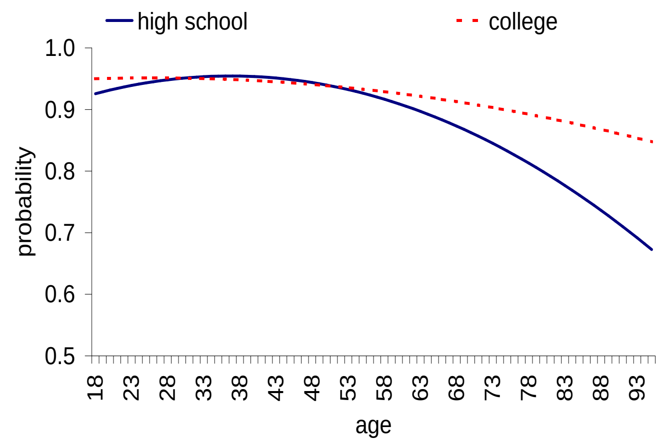

(b) married to married

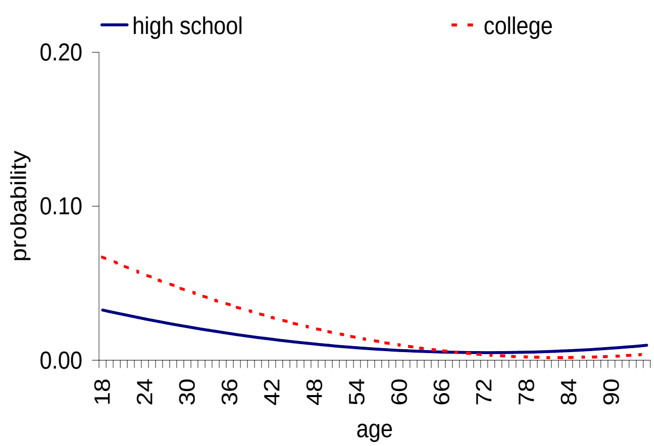

(c) div/wid to married

Figure 15: Transition probabilities for marital states

which come from CPS data. Age of partners is important since they determine the extra income for the household in terms of partner's labor earnings and the probability of death (hence, transitioning to widowhood status). Since the profiles for both characteristics are smoothed, the tradeoff between accuracy and simplicity of the solution by assuming such a narrow age distribution is lessen.

For assets, I calculate from CPS data the average annual non-labor income (dividends, interests and rents) for both single males and females. Single males have on average 20% higher non-labor income than single females. Hence, I create a simple three point distribution for assets of prospective partners $a^* = \{1.1a, 1.2a, 1.3a\}$, centered around the fact that on average $a^*/a = 1.2$. This simple distribution is uniform (equal probabilities for each point). Changing this distribution doesn't alter any of the qualitative results from the exercise.

A note on Total Fertility Rates and cohort effects: throughout the paper, I assume no cohort effects in fertility rates. Although fertility has experienced significant changes during the 20th century, fertility rates for the cohorts considered in my analysis are quite stable. Figure 16 shows age specific fertility rates computed from the internet release of Vital Statistics of the United States for the year 1995 (tables 1-7).

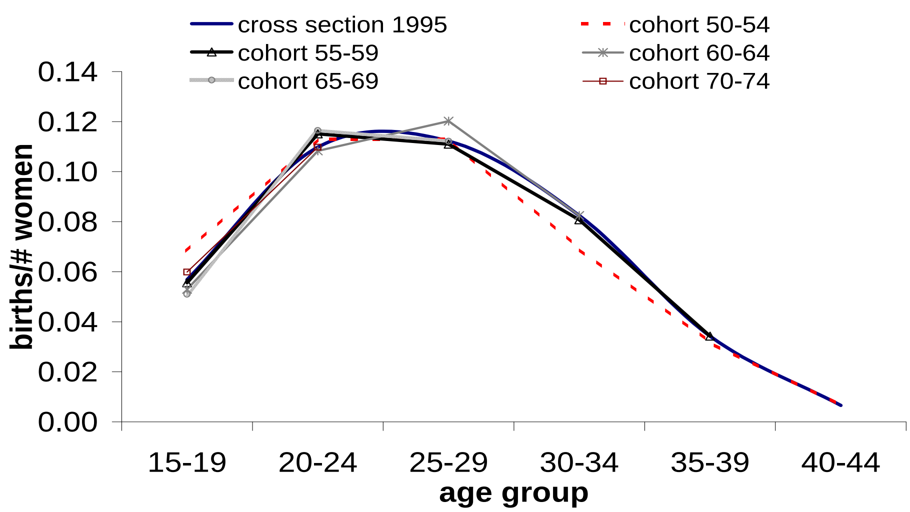

Figure 16: Age Specific Fertility Rates by birth cohort

The figure shows both total fertility rates for the cross section in 1995 and for cohorts (denoted by year of birth) across multiple survey years. The differences between the cross-sectional profile and actual cohort profiles is minimal. This comes from the fact that I am considering a small window in the life-cycle of cohorts that are close together (at most 20 years between births).

9.3 Computation and Estimation details

To solve the model, I use a Chebyshev regression (as described in Judd (1998) and Heer and Mauner (2004)) to approximate the optimal policies for savings and contraceptive effort and the value function along the asset space (the only continuous state variable in the model). My approximation is described by 7 collocation points and the use of a Chebyshev polynomial of degree 5. Increasing both the number of collocation points and/or the order of the polynomial doesn't improve the quality of the approximation significantly.

A note on the non-standard mapping between model periods and years: I adopt this procedure to save on computational time. A similar feature is present in Keane and Wolpin (2006). In my model, $i_f = 40$ (last fertile age), $i_r = 45$ (stands for a retirement age of 65 years) and $I = 48$ (represents the age of 95, the last period of life).

I assume that an individual aged $i \in \{i_f + 1, ..., i_r\}$ experiences one model period as the average of 5 real years; when $i \in \{i_r + 1, ..., I\}$, the experience is that of 10 averaged years. The external data used (and described in the previous section) is treated accordingly depending on the age of the individual: earnings, survival and transition probabilities, etc., are averaged in groups of 5 or 10 years accordingly.

Some notes on the estimation: The parameterization of age profiles for contraceptive and conception ability are in table 3

	$\widetilde{\varphi}^+_{i,m=\{1,3\}}$ (not-married)	$\widetilde{\varphi}^+_{i,2}$ (married)	$\overline{\varphi}_i$ (high-school)	φ^-_i (conception)
constant	2.05E-02	-1.17E-01	4.74E-01	2.89E-01
i	-3.86E-02	2.81E-01	-1.02E-01	4.34E-02
i^2	2.99E-02	-8.57E-02	6.68E-02	-2.30E-02
i^3	-5.34E-03	1.25E-02	-1.19E-02	3.58E-03
i^4	3.98E-04	-9.12E-04	8.47E-04	-2.59E-04
i^5	-1.35E-05	3.18E-05	-2.69E-05	8.85E-06
i^6	1.71E-07	-4.24E-07	3.17E-07	-1.15E-07

Table 3: polynomial coefficients for parameterization of fertility profiles

To accelerate the estimation algorithm, I use a Beowulf cluster with 20 processors. I parallelize at the parameter level, using the APPSPACK software available free on the web. See Gray and Kolda (2006) and Kolda (2005) for details.