

Fertility Risk in the Life-Cycle

Sekyu Choi
Universitat Autònoma de Barcelona,
Barcelona GSE and MOVE*

March 2015

Abstract

In this paper I study fertility decisions of women in the U.S. economy, with special emphasis on the timing of births and abortions over the life-cycle. Given the constant policy debate regarding abortion availability and recent empirical evidence of its positive impact on women outcomes, understanding the fertility process in a broad sense should help guide the discussion. In this paper I present a life-cycle model of consumption-savings and fertility decisions in an environment with uninsurable income shocks and imperfect fertility control. My model presents a unified framework in which both opportunity costs of child rearing and technological restrictions (in the form of contraception effectiveness) have roles to understand lifetime fertility choices.

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

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

*Email: sekyu.choi@uab.cat. This is part of my Phd dissertation at the University of Pennsylvania and previously circulated under the title "Life-cycle Fertility: Means vs. Motives vs. Opportunities". I'd like to acknowledge the support and encouragement of José-Víctor Ríos-Rull throughout this project. This paper benefited from discussions with Jeremy Greenwood, Dirk Krueger, Jesús Fernández-Villaverde, Virginia Sánchez-Marcos, Iouri Manovski, Ken Wolpin, Hal Cole, John Knowles, Nezh Guner, Juan Carlos Conesa, two anonymous referees and seminar participants at several institutions. I gratefully acknowledge financial support from the Spanish Ministry of Economy and Competitiveness through grant ECO2012-32392 and through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV-2011-0075). All errors are mine.

1 Introduction

The policy debate regarding abortion legality and availability resurfaces frequently in the US. This is fuelled in part by the fact that one in four pregnancies ends up in an abortion, a stable statistic over the last decade.¹ In terms of access, there are still signs of restrictions to access: [Upadhyay, Weitz, Jones, Barar, and Foster \(2013\)](#) estimate that a significant number of women get turned down at abortion clinics or are unable to procure an abortion due to strict term limits or lack of clinics in certain regions. As for its impact, there is empirical evidence showing that abortion availability is related to better labor market and health outcomes for women, as noted by [Angrist and Evans \(2000\)](#) and [Coleman \(2011\)](#) respectively, as well as to improvements in eventual child investment decisions by mothers, a result found in [Gruber, Levine, and Staiger \(1999\)](#).

Given the significance of failed contraception and the effects of abortion on women’s outcomes and welfare, in this paper I introduce a model of stochastic fertility in order to understand fertility decisions in a broad sense. The model is characterized by an incomplete markets framework with uncertain income and exogenous marital status transitions (similar to [Hong and Ríos-Rull, 2012](#)), where females cannot control perfectly the timing of births, thus creating incentives to abort some pregnancies. The model is rich enough to predict heterogeneity of births and abortions across educational groups and is useful to quantify the sources of such heterogeneity. Also, given that the model is embedded in a standard incomplete-markets framework (where agents have access to savings), the model is able to produce predictions on the interplay between household assets and fertility decisions over the life-cycle.

The model I present builds on [Mincer \(1963\)](#), [Becker \(1960\)](#) and [Becker \(1965\)](#), who put forward an optimal ”allocation of time” theory to rationalize the negative income-fertility correlation observed in the data: given that child rearing requires time away from the market, individuals with higher skills (thus, higher value of their market time) choose optimally to have fewer children.² In this paper I expand this theory to allow for the possibility of fertility risk, in the sense of agents not being able to fulfill their fertility plans perfectly.

This negative income-fertility correlation can be observed from the left panel of figure 1, where I tabulate information on births from the 1995 National Survey of Family Growth.³ Compared to their college counterparts, individuals in the high school group have higher and earlier birth rates. Since they also have lower earnings (see figure 13 in the appendix) this represents a negative correlation between fertility and labor income.

The second panel of figure 1 on the other hand, hints at the need to expand the ”allocation

¹See for example [Henshaw \(1998\)](#), [Finer and Henshaw \(2006\)](#) and [Kocharkov \(2012\)](#).

²See [Galor and Weil \(1996\)](#) and [Greenwood, Seshadri, and Vandenbroucke \(2005\)](#) for an updated view. See also [Jones, Schoonbroodt, and Tertilt \(2008\)](#) for a discussion in dynastic models, and [Erosa, Fuster, and Restuccia \(2010\)](#) for one in a life-cycle context.

³for a detailed description of the data, see the Appendix.

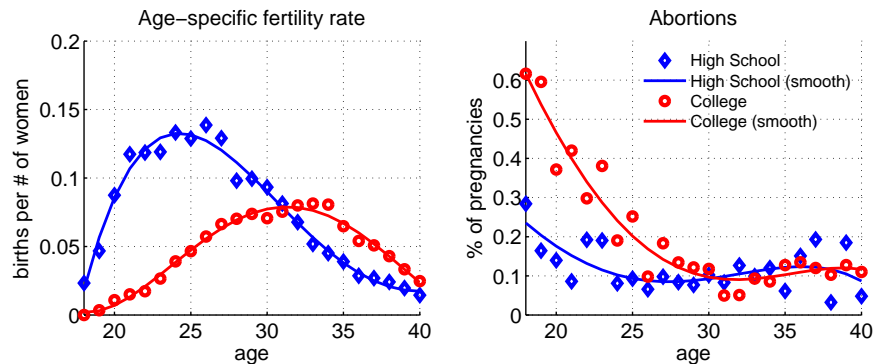


Figure 1: Age-specific fertility rates and abortions percentages, by education attainment of the respondent. Source: 1995 wave of the National Survey of Family Growth. "High School" is the group of all those with a high school diploma or less education; "College" is all the rest. "Smooth" refers to a 4-th order polynomial on age.

of time" story. It documents the prevalence and timing of abortions across educational groups. The figure shows that the rate between number of abortions and pregnancies is higher for the college group earlier in the life-cycle and that the age profile for both groups is decreasing in age, converging at around age 25. Although the allocation of time theory might explain the fact that the college group exhibits a higher abortion rate, the declining trend for both educational groups and its fast convergence hints at some complementary mechanism at play: if the High School group has less opportunity costs from child bearing, why do they chose to abort (relatively) earlier in their life-cycle as their college counterparts do?

Here, I use the findings in [Rosenzweig and Schultz \(1989\)](#) who show that more educated individuals are more efficient using birth control methods. Thus, my model combines both "allocation of time" and "differential fertility risk" mechanisms to account *jointly* for the facts. Assuming the same preferences for children across individuals, I calibrate the model using data from the 1995 National Survey of Family Growth (NSFG), the Panel Study of Income Dynamics (PSID) and the Current Population Survey (CPS) using a simulated method of moments approach.

The model predicts that the "differential fertility risk" mechanism is key to simultaneously account for the differences in births along similarities in abortion trends across educational groups. Thus, this result suggests that the most direct way to affect abortions and fertility decisions is through policies aimed at awareness efforts (e.g., sex-education at early ages) rather than at income subsidies (e.g., child-care subsidies or paid maternal leave). The model also implies an important role of asset accumulation in determining the rates of births and abortions and correctly predicts the relationship between assets and births, given longitudinal evidence from the PSID: the level of assets has a positive and significant effect in predicting births (after controlling for education and marital status of the mother).

Admittedly, my approach suffers from several limitations which are used to simplify the analysis: I assume exogenous educational attainment and marital transitions, which might pose problems given that fertility decisions might be closely intertwined with both education and marriage/divorce choices; I also pose a generic decision choice model, but do not identify the sources of fertility risk per se. Both of these issues pose interesting topics for future research.

This paper is related to recent literature using dynamic models to analyze the interplay between fertility and different economic outcomes. For example, [Alvarez \(1999\)](#) uses a Barro-Becker model of dynasties to study persistence of wealth across generations. [Caucutt, Guner, and Knowles \(2002\)](#) study how the optimal timing of fertility depends on the wage distribution. [Greenwood, Guner, and Knowles \(2003\)](#) analyze the interaction between fertility, the marriage market and the income distribution. [Soares and Falcao \(2008\)](#) propose a framework with endogenous fertility, labor supply and child investment decisions to analyze how gains in life expectancy changed the role of women and their incentives for having big families. My model is most closely related to [Conesa \(2000\)](#) and [Sommer \(2009\)](#): they both study life-cycle fertility decisions in environments with uninsurable labor income risk. I extend this framework by considering a richer demographic structure and the presence of fertility risk, hence, my environment provides predictions for abortions.

My approach borrows insights from the empirical microeconomic literature that studies life-cycle fertility⁴ using structural and dynamic models of fertility choice. From that literature, my paper relates the most to [Wolpin \(1984\)](#) and [Hotz and Miller \(1993\)](#) who acknowledge the importance of the stochastic nature of fertility. The earlier, analyzes how child mortality risk shapes fertility choices using Malaysian data; [Hotz and Miller \(1993\)](#) on the other hand, estimate birth control method choices by females in a life-cycle framework. My approach differs from theirs in that I acknowledge costly abortions and thus, the possibility of error in implementing fertility plans. Moreover, I assume imperfect capital markets in the sense that agents can save but not borrow against their future earnings while the above mentioned literature assumes away the existence of capital markets. In terms of environment, I also introduce exogenous marital transitions which help in explaining aggregate births and abortion rates for the entire female population. In terms of preferences, I impose the same utility function for all agents, downplaying the role of unobserved heterogeneity in preferences to account for the data.

2 A Quantitative Model

The economy is populated by agents of different gender (males and females) and education level (high school and college). I abstract from the education decisions in order to simplify the analysis. Agents 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

⁴See [Hotz, Klerman, and Willis \(1997\)](#) for a survey

to their efficiency units of labor). 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. This structure is similar to models in [Hong and Ríos-Rull \(2012\)](#) and [Hong \(2008\)](#).

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 a fraction of their time at home rearing their children and this time cannot be substituted away (i.e., there is no child care). Males are not affected by this time requirement.

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

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

asterisks represent values for spouses (when applicable). Age is indexed by $i = \{i_0, \dots, I\}$, a is the amount of real assets in the household, $k = \{1, 2, \dots, K\}$ represents the number of children living in the household, $e \in \{\underline{e}, \bar{e}\}$ represents the education type of the agent (low, high), $m = \{1, 2, 3\}$ is the type of household (1 = single, 2 = married, 3 = widowed/divorced), and ϵ is the value of the multiplicative shock to labor earnings. For ease of exposition, in some sections of the paper I use the following partition of the state space $\tilde{z} = \{e, e^*, m, \epsilon, \epsilon^*, i^*\}$ so that $z = \{i, a, k\} \times \tilde{z}$.

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 (\bar{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 $\delta_{i,e}$ and δ_{i^*,e^*}^* for females and males respectively.

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 (or status quo). 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. The exertion of this effort comes at a utility cost $C(x)$.

With complementary probability $(1 - \pi)$, a pregnancy occurs. Agents have the opportunity of getting an abortion at a utility cost κ . If the agent desires to keep the pregnancy, the household increases its size by one. There is no child mortality risk in the model nor multiple births. Note that the effort space is the real line, so trying to get a pregnancy can be identified with exerting a negative level of effort $x < 0$ which minimizes the probability of status quo (no pregnancy).

I make the assumption that children are attached to females. I don’t keep track of the age nor the sex of children in the household due to the computational burden of doing so. 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.⁵ 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.⁶

Markets. Agents 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 at the market rate r .

Labor endowments. Agents are endowed with state dependent efficiency profiles, $\varepsilon_{i,m,e}$ for females and $\varepsilon_{i^*,m^*,e^*}$ 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

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

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 function $b(m, k) \in (0, 1)$, which depend on marital status m and number of kids in the household k . Thus, labor income of females/mothers is given by

⁵This hazard rate is independent for each child in the household (if $k > 1$).

⁶This timing assumption simplifies the calculation of expectations over future states.

$$b(m, k)w \in \varepsilon_{i, m, e}$$

Since I don't keep track of ages of children in the household, $b(m, k)$ is not age dependent. In this formulation, $b(m, k)$ represents a *motherhood penalty* in the sense of a reduced wage bill for female agents with kids.

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 valued by the household decision maker. 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; it also creates consumption habits from marriage, since the marginal utility from consumption will differ between single and married households.⁷

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 adds simplicity to the model. Finally, 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 woman in this model: (1) work-fertile stage, (2) work - infertile stage and (3) Retirement. In the first stage, women make both fertility and consumption-savings decisions. The following Bellman equation represents the problem of agents during this stage, before making fertility decisions:

⁷See [Hong and Ríos-Rull \(2012\)](#).

$$\begin{aligned}
V(i, a, k, \tilde{z}) &= \max_{c, y} u(c|z) + \gamma g(k) + \delta_{i, e} \beta E [v_f(i + 1, a', k', \tilde{z}')|z] \\
&st : \\
c + y &= (1 + r)a + b(m, k)w\epsilon_i \epsilon_{i, m, e} \quad \text{if } m = \{1, 3\} \text{ or } m = 2, i^* \geq i_r \\
c + y &= (1 + r)a + b(m, k)w\epsilon_i \epsilon_{i, 2, e} + w\epsilon_{i^*}^* \epsilon_{i^*, 2, e}^* \quad \text{if } m = 2, i^* < i_r \\
a' &= \Phi(y, z'|z)
\end{aligned} \tag{2}$$

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.⁸ To update the number of children present in the household, I apply a binomial distribution with parameter s_k , i.e.,

$$P(k' = k_0 | k) = \left(\frac{k!}{k!(k - k_0)!} \right) s_k^{k - k_0} (1 - s_k)^{k_0}$$

Given optimal policies in the problem above, females make contraceptive effort choices. This is represented by value function v_f below:

$$\begin{aligned}
v_f(i, a, k, \tilde{z}) &= \max_x \pi(x|i, m, e) V(i, a, k, \tilde{z}) \\
&+ [1 - \pi(x|i, m, e)] \max \left\{ \begin{array}{l} V(i, a, k + 1, \tilde{z}), \\ V(i, a, k, \tilde{z}) - \kappa \end{array} \right\} \\
&- \frac{1}{2} x^2
\end{aligned} \tag{4}$$

This value function is a convex combination of the continuation values $V(\cdot)$ 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 . Note that even though there are discrete

⁸The particular form of Φ is given by:

$$\Phi(y, z'|z) = \begin{cases} y & if & (m' = 2 | m = 2) \\ y & if & (m' = 1, 3 | m = 1, 3) \\ y & if & (m' = 3 | m = 2) \text{ (widowhood)} \\ 0.5y & if & (m' = 3 | m = 2) \text{ (divorce)} \\ y + a^* & if & (m' = 2 | m = 1) \end{cases} \tag{3}$$

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 split and divided equally among agents, hence $a' = 0.5y$. 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.

outcomes following this optimization problem (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.⁹

This setup allows 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 and restricted to be symmetric around zero. This general specification allows me to capture biological constraints 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. The specifics of the dynamic problem depend on whether individuals are or not retired. Before retirement ($i \leq i_r$), the problem of the agent is:

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

After retirement, the problem reduces to

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

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 savings. If agents are married to working age individuals, they have access to $w \epsilon_{i^*}^* \varepsilon_{i^*, 2, e}^*$ (the extra labor income).

⁹Details of the numerical solution procedure are in the Appendix.

3 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. A model period is one year. 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. Age specific mortality rates are taken from the National Center for Health Statistics and adjusted for educational attainment, as in [Hong \(2008\)](#).

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 never married 18 year old females in the CPS is around 93% and the proportion females with kids is around 9%. When performing simulations of the model, I distribute women 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.¹⁰ 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 \(2012\)](#), 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. Given variable specification in the PSID, married couples include cohabitating couples.

Preferences. I use an additively separable specification for the per period utility derived from consumption and number of children in the household, $u(c|z) + \gamma g(k)$. The utility from consump-

¹⁰My proxy for individual assets is the sum of interest, dividend and rent income as defined in the March supplements of the CPS.

tion is of the constant relative risk aversion (CRRA) type, which also depends on the size and composition of the household, as in [Heathcote, Storesletten, and Violante \(2012\)](#)¹¹

$$u(c|z) \equiv \omega(z) \left(\frac{\mathbf{c}(z)^{1-\eta_c} - 1}{1 - \eta_c} \right) \quad (7)$$

where $\omega(z)$ equals the number of adults in the household and $\mathbf{c}(z) = c/(1 + 0.7 \times \mathbf{1}_{\{m=2\}} + 0.5 \times k)$ represents public consumption inside the household, deflated by a standard OECD equivalence scale ($\mathbf{1}_{\{m=2\}}$ is an indicator function which equals one if there are two adults in the household). In this setup, direct utility and economies of scale related to consumption inside the household depend on its size and composition. On the other hand, the direct utility over children (g) is given by

$$g(k) = \frac{(1+k)^{1-\eta_k} - 1}{1 - \eta_k} \quad (8)$$

Fertility. I assume that the maximum number of children a women can have in her lifetime is five ($K = 5$). I use the following function for π , the probability of NO conception (or status quo), given effort x :

$$\pi(x|i, m, e) = \frac{1}{1 + \varphi_{i,m,e} \exp\{-x\}} \quad (9)$$

π is a modified logistic function with $\varphi_{i,m,e}$ as a shift parameter. 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 this parameter. I further parameterize $\varphi_{i,m,e}$ as

$$\varphi_{i,m,e} = \mathbf{1}_{\{m=1,3\}} \tilde{\varphi}_{i,s} + \mathbf{1}_{\{m=2\}} \tilde{\varphi}_{i,m} + \mathbf{1}_{\{e=\underline{e}\}} \bar{\varphi}_i$$

where $\tilde{\varphi}_{i,s}$ is the shifter for non-married females, $\tilde{\varphi}_{i,m}$ is the shifter for married ones and $\bar{\varphi}_i$ is a shifter for females in the low education group. These parameters are age specific, in order to account for biological restrictions on conceptions. I assume that these parameters decay linearly to zero from age 30 onwards, so to parameterize $\varphi_{i,m,e}$, I only need 3 parameters.

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 \(2012\)](#) and [Hong \(2008\)](#), I use annual earnings since they capture differences in the intensive margin of earnings by sex and

¹¹See [Bick and Choi \(2013\)](#) for a discussion on the implications of different specifications of $\omega(z)$ for predictions of life-cycle consumption.

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.

I attribute the time cost of child-rearing $b(m, k)$ to annual labor income differentials of females in fertile age (18 to 40 years old) by number of children. This is not exactly 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 a law firm, etc.) other than through hours worked. Since the presence of children is persistent in the household, these "time costs" reflect both contemporaneous as well as dynamic effects (loss of occupation specific human capital, for example). The income differential values are in table 1.

Children	Not Married	Married
0	100.0%	100.0%
1	94.1%	73.5%
2	83.1%	62.5%
3	59.0%	47.4%
4	38.7%	36.7%
5+	18.8%	27.2%

Table 1: Ratio between annual labor income of females aged 18 to 40, with k children and different marital states, with respect to females in the same marital status but with NO children. Source: 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, since the ratio of annual labor income to that of a childless female is decreasing. The cost increases faster in the number of kids for married women than for non-married ones, denoting possible specialization in housework/household production for women with children. Since women in the model chose the timing of fertility, and given fixed income profiles for different educational groups, these statistics are endogenous in the model, thus, I calibrate the values of $b(m, k)$ to match the facts in table 1.

For the 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 $e = \{\underline{e}, \bar{e}\}$) from [Hong \(2008\)](#), who uses the PSID between 1986-1992 to compute maximum likelihood estimates. These parameters are in table 2. As is common, I discretize both continuous processes using the method proposed by [Tauchen \(1986\)](#).

	High School (\underline{e})	College (\bar{e})
ρ_e male	0.9023	0.9117
ρ_e female	0.8602	0.8750
σ_e male	0.2208	0.2037
σ_e female	0.2984	0.2760

Table 2: Parameterization of income shocks, by gender and educational level. Source: [Hong \(2008\)](#).

4 Calibration

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 4.66%, the average of the 1-year Treasury Bill Rate (monthly auction averages).¹² I let the discount factor β to be $1/(1+r)$ and set $\eta_c = 1.5$, a standard value in the literature.

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 initial 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 (v) calculate the error of the iteration (the sum of square values of the difference between every data and model moment) (vi) 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 by education: 46 moments (23 ages \times 2 education levels)
- Age profile of abortion rates by education: 46 moments
- Income ratio between mothers and non-mothers by marital status: 10 moments

In total, there are 102 moments to match. On the other hand, the model has 16 parameters to be determined jointly (which makes this an overidentified system):

- curvature in the utility of children: η_k
- multiplicative parameter in utility of children: γ
- utility cost of an abortion: κ
- contraceptive ability parameters: $\tilde{\varphi}_{i,0}, \tilde{\varphi}_{i,s}, \tilde{\varphi}_{i,m}$

¹²Series id TB1YA, on the St. Louis Fed Economic Data webpage.

- contraceptive ability shifter for low skill/education group: $\bar{\varphi}_i$
- time cost of children: $b(m, k)$

The calibrated parameters are displayed in table 3 where I show those related to individual preferences and to fertility, and in table 4 where I show the parameterization of $b(m, k)$, the time cost of children for females.

Parameter	Description	Value
η_k	Curvature in preference for Children	2.02
γ	Multiplicative constant, pref. for Children	5.02
κ	Utility cost of an abortion	40.84
$\tilde{\varphi}_{i_0,s}$	Shifter in fertility control tech. (Not-Married)	0.014
$\tilde{\varphi}_{i_0,m}$	Shifter in fertility control tech. (Married)	0.100
$\bar{\varphi}_{i_0}$	Shifter in fertility control tech. (\underline{e})	0.035

Table 3: Model Parameters: preferences and fertility process

Children (k)	Not-Married ($m \neq 2$)	Married ($m = 2$)
1	0.99	0.90
2	0.75	0.75
3	0.55	0.59
4	0.37	0.47
5	0.17	0.35

Table 4: Model Parameters: time costs of children for different marital states, $b(m, k)$.

The model implications of certain parameter values are straightforward. Increasing the different contraceptive ability parameters leads to more pregnancies, which given a cost of abortions and preference for children, leads to both higher number of births and abortions. On the other hand, increasing γ , the multiplicative parameter in the preference for children, increases desired fertility, without an increase in abortions (and thus, a fall in the rate of abortions), while movements in the cost of abortions affect individuals in terms of their exerted contraceptive effort and the number of abortions.

A key parameter for the entire exercise is the curvature in the utility of children η_k , which controls the distribution of the number of children across households. However, η_k also controls the risk aversion of individuals with respect to fertility risks. Thus, movements in this parameter have also a level effect in total fertility rates. With this caveat in mind, figure 2 is informative on how η_k affects the calibration.

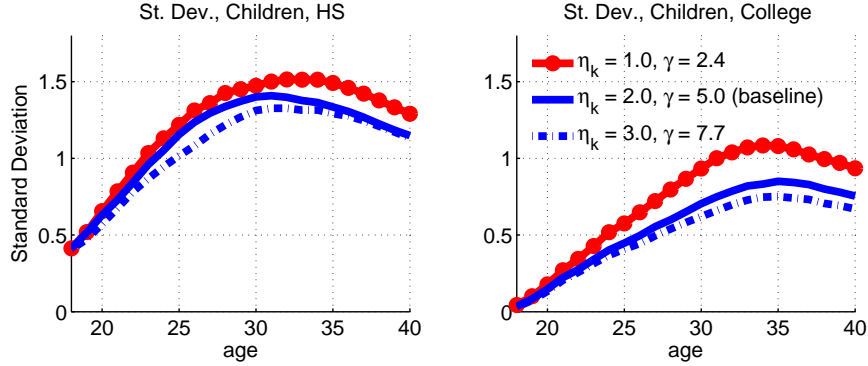


Figure 2: Standard deviation in number of Children inside the household, given different curvature (η_k) parameters in the preference for Children. Level parameters (γ) are adjusted in each case to produce the same average number of Children as in the baseline.

The figure shows the standard deviation in the number of Children across households at each age by educational group. For the three cases of η_k , the level parameter γ is set such that the average number of Children across models is the same as in the baseline. Figure 2 shows clearly that the curvature parameter η_k is negatively related to the amount of dispersion in the number of Children, with its effect being more pronounced for the College group.

The calibration of $b(m, k)$ conveys the time cost of child rearing in the model, which coupled with the endogenous timing of fertility, give rise to the model equivalent of the ratios presented in table 1. As discussed in the introduction, the differential fertility across education groups is determined in part by how they are affected by these time costs, as well as the extra fertility risk bore by the High School group.

5 Results

Figure 3 and table 5 show the model performance in terms of matching the targeted moments. As seen from the figure (two superior panels), the model matches well the differences in the profile of births across education groups, with a slight underprediction for college fertility. On the other hand, the two bottom panels of figure 3 show that the model predicts accurately both the level of abortions per pregnancy and the decreasing nature of these profiles over the life-cycle. The model also is able to match closely the targets presented in table 5, which represent the time cost of rearing children in terms of lost labor market opportunities.

The model also fits the qualitative differences in birth and abortion profiles across education and marital status groups: predicted birth rates are highest for married in the HS group, while abortion rates are highest for singles of both groups and degenerate at zero for married individuals,

again for both groups. This is shown in figures 11 and 12 (in the appendix). Since the calibration does not target statistics by marital status, this is a fair goodness of fit test for the model. The inability to match specific profiles at this level of disaggregation implies the existence of subtle mechanisms not present in the model. However, all these results imply that the demographic composition is important to match aggregate birth and abortion statistics by educational group, since the differences in marriage and divorce rates are imposed exogenously to agents in the model.

Results from the calibration procedure show that the contraceptive technology shifters (φ parameters in table 3) follow closely birth rates across different demographic groups: the parameter for married females $\tilde{\varphi}_{i_0,m}$ is higher than the one for non-married females $\tilde{\varphi}_{i_0,s}$ and the shifter for the low educational group is positive, with a value of $\bar{\varphi}_{i_0} = 0.035$. On the other hand, we see from table 4 that the gradient in function $b(m, k)$ is similar across marital states for all number of children, while the endogenous outcome (income ratios in table 5) shows a faster decline in annual earnings for married mothers, showing that $b(m, k)$ influences the timing of births over the life-cycle and across marital states in a non-trivial way.

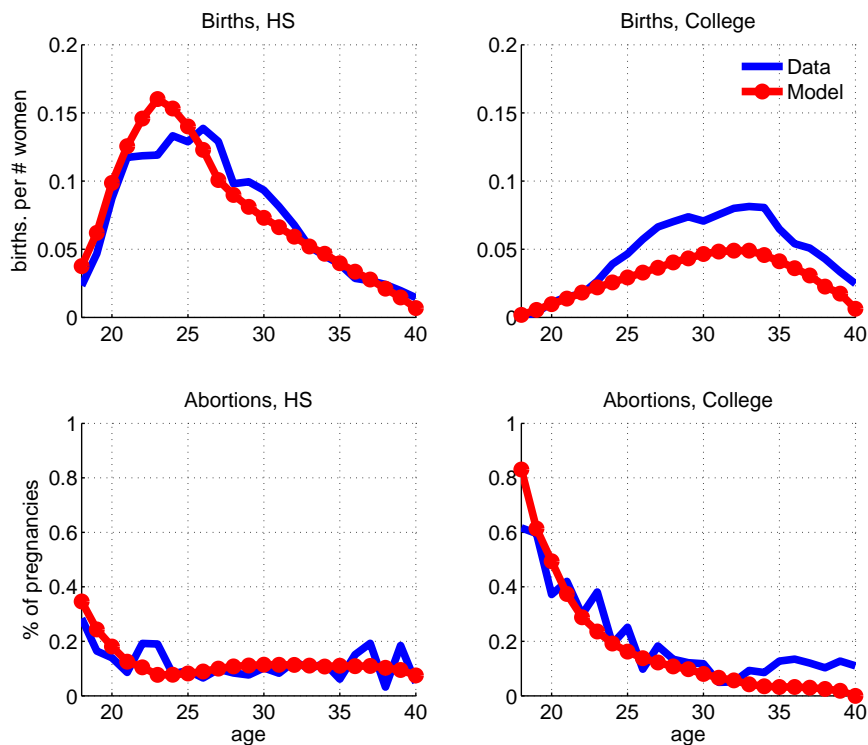


Figure 3: Births per woman and abortion rates, Model versus Data (Baseline)

Figure 4 shows consumption and asset accumulation across educational groups. Assets for both

Children	Not Married		Married	
	model	data	model	data
1	93.0%	94.1%	72.5%	73.5%
2	86.7%	83.1%	61.2%	62.5%
3	58.0%	59.0%	45.8%	47.4%
4	37.9%	38.7%	37.7%	36.7%
5	19.6%	18.8%	26.0%	27.2%

Table 5: Model versus data on ratio between annual labor income of females aged 18 to 40, with k children and different marital states, with respect to childless females. Data Source: CPS 1990-1995.

groups are constrained (at zero) early in the life-cycle, showing that most individuals in the model are constrained in terms of borrowing during their most fertile years. In terms of consumption, the average profile for both groups shows a humped shape, induced by standard forces (borrowing constraints early in life) and the presence of children.¹³

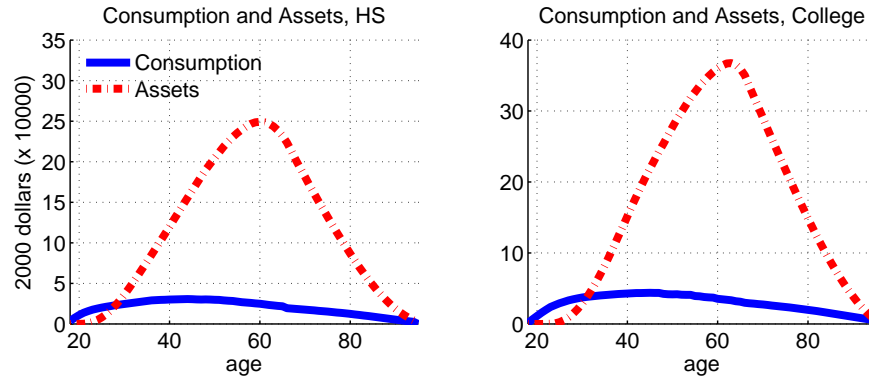


Figure 4: Consumption and assets in the model, HS and College.

A key dimension of the model, which leads to differences in the profiles of births and abortions across education groups, is the one dealing with the opportunity cost of children vis a vis the fertility risk faced by each group. Figure 5 shows the life-cycle profiles of annual earnings for females, by number of children in the household.

As noted earlier, the ratio between average annual earnings between mothers and childless females (by marital status) is calibrated to match CPS data. Thus, the differences observed in the figure match reality and show that the *motherhood* penalty for the High School group, measured as the vertical distance between earning profiles, is significantly less than for the college group. Thus,

¹³A detailed analysis of consumption profiles and demographic structure can be found in [Bick and Choi \(2013\)](#).

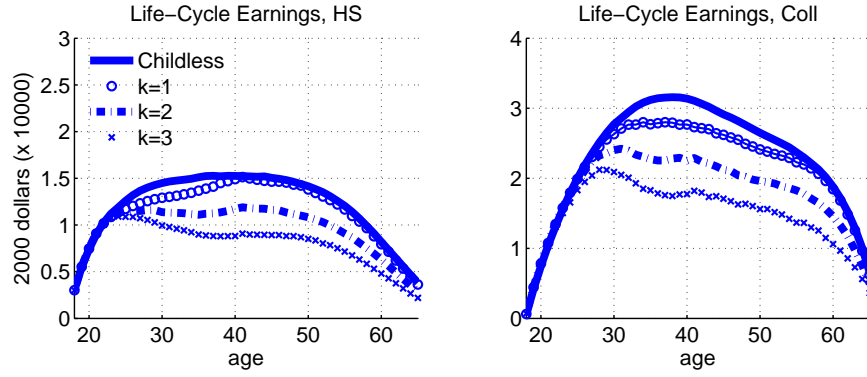


Figure 5: Annual female earnings in the model, HS and College. Earnings are shown by number of children currently in the household.

just in terms of the opportunity cost of having children, it is reasonable to think that the college group prefers less fertility and leans more towards abortions earlier in life, which is the observed pattern for the college group in figure 3. However, the main challenge that the opportunity cost of children mechanism faces is in explaining the similar profile of abortions for the High School group, who need to sacrifice less of their income to rear for children. To match both a higher number of births with a decreasing profile of abortions for this education group, the calibration results in a positive value of the extra fertility risk parameter for High School females.

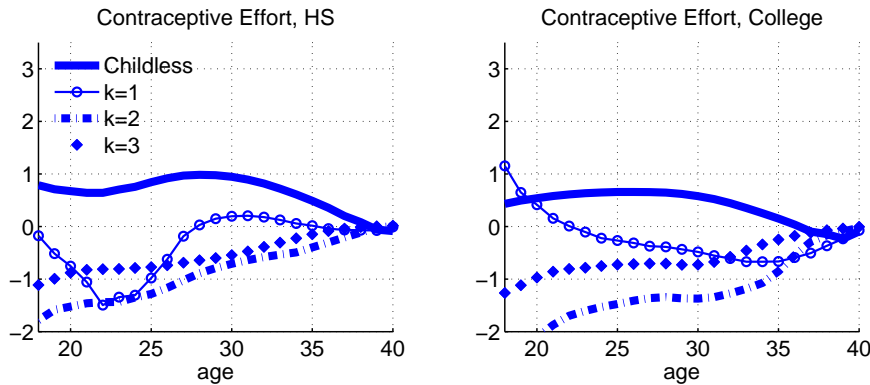


Figure 6: Average contraceptive effort by number of children currently in the household, HS and College.

Figure 6 shows exerted contraceptive effort over the life-cycle. The figure shows that on average individuals in both groups exert positive effort when they are childless (positive values for Childless individuals in the figure), but once they have children, they start exerting negative effort, to increase

their family size. This negative contraceptive effort declines with the number of children already in the household, which reflects the effect of the concavity in the utility for children. Note also that the effort is higher for the High School group when compared to the college one, especially for childless individuals throughout the life-cycle. Consistent with the imposed restriction of no fertility beyond age 40, all contraceptive/conception efforts converge to zero at that age.

5.1 Exploring the Roots of Differential Fertility Patterns.

Below I present two exercises which showcase how different factors affect fertility decisions in the model. To simplify exposition, I show figures for the High School group only (the group which experiences the extra fertility risk in parameter $\bar{\varphi}_{i_0}$). Figure 7 shows birth and abortion profiles for the baseline model and two different cases: one where the income process and marriage market of the High School group are equalized to that of the College group and one where only the marriage market is equalized. Both experiments make females in the HS group enjoy higher levels of income, either because of higher average income profiles/shocks or marriage to higher earners.




Figure 7: Birth and abortion profiles comparison for the High School Group, for Baseline model and a model where College characteristics are given to the HS group (HS=Coll).

The figure shows that the effect of these changes on the baseline predictions for births and abortions are small. Both increasing earnings and improving marriage prospects for women in the HS education group increase slightly their birth profiles, specially earlier in the life-cycle, while the effect on abortions is slightly negative.

In comparison, consider figure 8 where I show outcomes of the baseline model, a model with no extra fertility risk for the HS group and a model where the cost of abortions is zero. As seen from the figure, equalizing the fertility risk across education groups ("No HS Risk" line) decreases significantly births over the life-cycle, while shifting the shape of the abortion profile towards a higher percentage of abortions earlier in the life-cycle: there are less early pregnancies and births

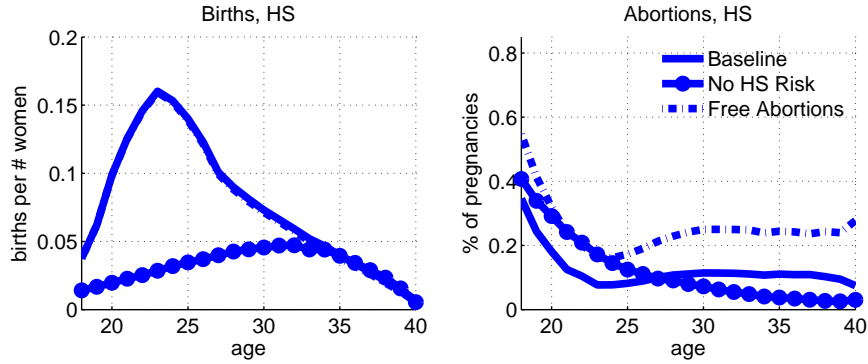


Figure 8: Birth and abortion profiles comparison for the High School Group, for Baseline model, a model with no differential fertility risk ("No HS Risk") and a model where the utility cost of abortions is set to zero ("Free Abortions").

(due to $\bar{\varphi}_{i_0}$ being zero) and the ratio of abortions to pregnancies increases before age 30 and decreases after that age when compared to the baseline. The figure shows that once this extra risk is removed, the profiles of births and abortions become very close to the ones for the college group.

On the other hand, the "Free abortions" economy exhibits a birth profile for the HS group which overlaps with the baseline, while the fraction of abortions increases for all ages: agents in this economy are simply substituting the relatively costly (ex-ante) birth control with costless (ex-post) control, i.e., abortions.

These figures confirm the ideas discussed above and in the introduction. To simultaneously account for births and abortion profiles across education groups, one needs to expand the theory to include the notion of differential fertility risks. The discussion above also shows that this margin is quantitatively important, relative to other margins in the model (income and marriage market differentials).

5.2 Assets and Fertility.

Since the fertility theory presented here is embedded in a life-cycle, incomplete markets model, I can assess the importance of asset accumulation for fertility decisions. In this section, I show how model predictions line up with empirical evidence from the US economy. The data requirements to assess the relationship between household assets and fertility decisions are strong: I need longitudinal information on assets and household composition around (or close to) the time period which I used to calibrate the model. The only data-set fulfilling these requirements is the Panel Study of Income Dynamics (PSID) which contains information on household wealth between the years 1999 and 2009 along with detailed information on household characteristics, especially income and demographics.

I calculate wealth from the PSID by adding the value of different assets owned by the household and subtracting all household debt. Home equity is also included. As a first pass, in figure 9 I show assets to income ratios, for both the PSID data and the model. The figure shows that the model lines up relatively well with the data, with less accuracy later in life (beyond age 55).

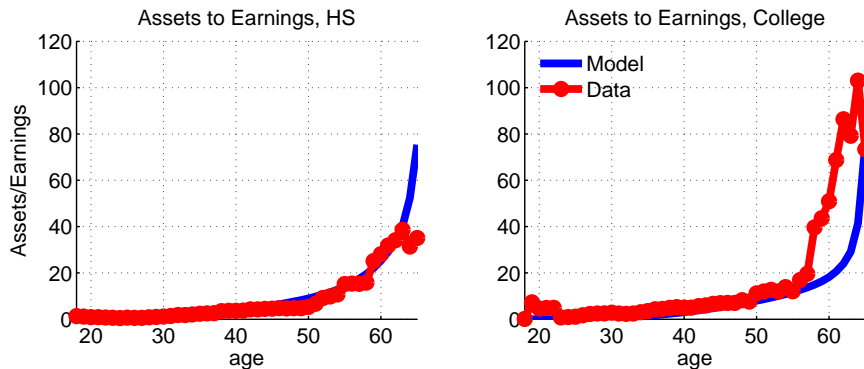


Figure 9: Assets to income ratios (household level). Data is from the PSID (1999-2009) family and wealth supplements.

Using the longitudinal dimension of the PSID, I impute births by comparing the number of own children inside the household across survey years and checking whether the age of the youngest own child is less or equal than 2 (the survey is performed every two years). I perform the analysis separating between married couples and single households (whose head is a female).

	Data		Model	
	coefficient	std. error	coefficient	std. error
Lagged earnings	0.1298	0.0860	0.1349***	0.0143
Lagged earnings (male)	0.1044	0.0925	0.0790***	0.0084
Lagged Assets	0.0062***	0.0024	0.0082***	0.0031
Obs.	2843		92599	
R-squared	0.0054		0.0002	

Table 6: Coefficients and standard errors from a linear, fixed effects-panel data regression on BIRTHS. Only married households considered. Data is from the PSID (1999-2009) main file and wealth supplements. Earnings and assets are in hundreds of thousands of current US dollars (2000). Regressions control for age, age squared, race, number of children already in the household and education dummies (for both members of the couple); 1,2 and 3 stars represent significance at the 10, 5 and 1% respectively.

I estimate fixed effect, linear regressions between a dummy variable indicating whether a birth has occurred in the household and individual earnings and household assets. I include the latter variables lagged one period (two actual years) in order to analyze whether labor earnings/assets

have predictive power on the probability of a birth. The earnings variable in the PSID is defined as total annual earnings, which matches my definition of earnings in the model. Wealth is defined as all wealth owned by the household (including housing) net of current debt. Both earnings and assets are in hundreds of thousand of 2000 dollars in the regressions.

I estimate the regression for both the PSID data and simulated data from my baseline model (fifty thousand simulated life-cycles). In the PSID regressions I control for age, education and race of each member of the household, as well as number of previous children. When using the PSID data, I also restrict the sample to include households where both male and females have positive earnings.

Table 6 shows the results of this exercise for the married sample. As seen in the table, lagged assets have significant predictive power on current births, both in the data and the model. Note that both male and female lagged earnings in the PSID regression affect births positively, but are statistically non-significant. The model correctly predicts the sign of earnings, but attributes significance, as opposed to the data.

	Data		Model	
	coefficient	std. error	coefficient	std. error
Lagged earnings	0.0498*	0.0269	0.0020	0.0013
Lagged Assets	0.0001	0.0003	0.0018***	0.0005
Obs.	3155		117401	
R-squared	0.0838		0.0002	

Table 7: Coefficients and standard errors from a linear, fixed effects-panel data regression on BIRTHS. Only single households considered. Data is from the PSID (1999-2009) main file and wealth supplements. Earnings and assets are in hundreds of thousands of current US dollars (2000). Regressions control for age, age squared, race, number of children already in the household and education dummies (for both members of the couple); 1,2 and 3 stars represent significance at the 10, 5 and 1% respectively.

Table 7 shows the results for single females: again, earnings and assets have positive effects on births, but the PSID regression shows that the effect is only statistically significant for earnings. For the model, the reverse is true in that lagged assets are the only ones which are significant.

Overall, the model and the data seem to agree on the effect of monetary variables predicting births across married and single households. Note that the calibration of the model did not include these regression coefficients, so the fact that model predictions are in the ballpark of what the PSID shows, is a sign of the goodness of fit of the model and its predictive power. Taken as a whole, this exercise shows that the baseline model provides a good first step in acknowledging and quantifying the role of assets for fertility decisions.

Assessing the effect of assets and earnings on abortions empirically is not possible given the available data. The National Survey of Family Growth, from where abortion statistics are compiled,

does not record asset information and has very limited information on both labor force participation and labor earnings of the mother at the time of conception. Thus, a similar regression analysis as the one performed for births is not possible.

In this regard, one can use the model to tease out the role of assets in determining abortion dynamics over the life-cycle. The next figure shows how the predictions of the baseline model depend on the presence of asset accumulation.

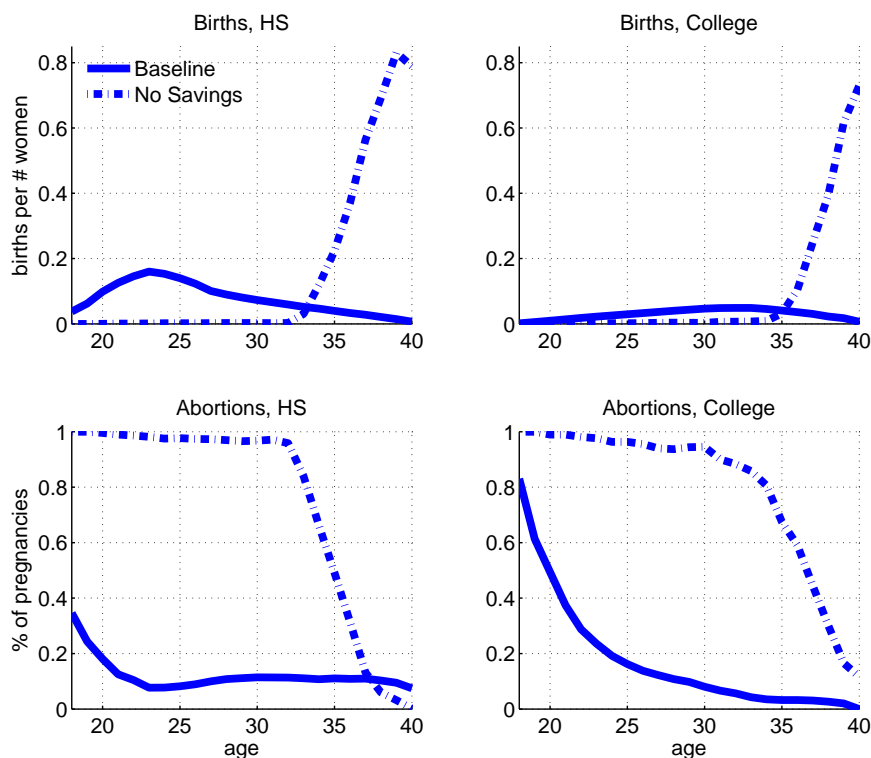


Figure 10: Births and abortion rates, Baseline vs. Model without savings

Figure 10 shows a comparison of births and abortions for both the baseline model, and an alternative model where individuals have no access to savings, so they must consume all their labor earnings every period. In the latter economy, we can observe that the average number of births is higher for the high school group, while abortions rates are higher for the college group, a prediction consistent with the baseline and the previous discussion. However, the most striking feature of this model is the timing of fertility decisions: both the high school and college groups decide to postpone fertility decisions as late as possible, resulting in very high abortion rates for all females earlier in the life-cycle. This result is due to the fact that in the absence of savings to smooth consumption across ages and different family sizes (given economies of scale), agents use children and fertility

decisions to try to equalize the value of $u(c|z)$ (utility from consumption) across different states of the world (income shocks and family sizes).

This exercise shows that the presence of assets in the model affects optimal fertility and abortion decisions in a dramatic way. Although the specific direction in which assets do so depends on several different factors, such as the curvature of the utility function, marriage markets, the timing of biological restrictions to fertility (age limit of 40 to conceive), etc., the model poses a very significant role of assets for these decisions.

6 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, number of births and abortions 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, but they can save. In this setting, I include endogenous decisions about fertility, but where the outcome of those decisions cannot be controlled perfectly (i.e., fertility risk). The exercise shows that to comprehend all dimension of fertility decisions, one must take a broad approach, considering both economic (opportunity cost of time) and non-economic (differential fertility risk) mechanisms.

Calibration of the model shows that in order to simultaneously account for the observed heterogeneity in births and abortion profiles across educational groups, one needs to extend the model to allow differential fertility risk across these groups. The model is useful to account for the sources behind fertility and abortion decisions and matches empirical facts regarding household assets and births. Thus, it extends important earlier work on contraceptive decisions by women in the sense that it allows for stochastic fertility outcomes and is a first step towards understanding the effects of market incompleteness on fertility decisions and household welfare.

References

- ALVAREZ, F. (1999): “Social Mobility: The Barro-Becker children meet the Latiner-Loury Dynasties,” 2(1), 65–103.
- ANGRIST, J. D., AND W. N. EVANS (2000): “Schooling and labor market consequences of the 1970 state abortion reforms,” *Research in Labor Economics*, 18, 75–113.
- 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,” *The Economic Journal*, 75, 493–517.
- BICK, A., AND S. CHOI (2013): “Revisiting the effect of household size on consumption over the life-cycle,” *Journal of Economic Dynamics and Control*, 37(12), 2998 – 3011.
- CAUCUTT, E., N. GUNER, AND J. KNOWLES (2002): “Why do Women Wait? Matching, Wage Inequality and Incentives for Fertility Delay,” *Review of Economic Dynamics*, 5(4), 815–855.
- COLEMAN, P. K. (2011): “Abortion and mental health: quantitative synthesis and analysis of research published 1995 to 2009,” *The British Journal of Psychiatry*, 199, 180–186.
- CONESA, J. C. (2000): “Educational Attainment and Timing of Fertility Decisions,” Working Paper 0010 CREB, UNiversidad de Barcelona.
- EROSA, A., L. FUSTER, AND D. RESTUCCIA (2010): “A Quantitative Theory of the Gender Gap in Wages,” Mimeo.
- FINER, L. B., AND S. K. HENSHAW (2006): “Disparities in rates of unintended pregnancy in the United States, 1994 and 2001,” *Perspectives on sexual and reproductive health*, 38(2), 90–96.
- 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/June 1998.
- GALOR, O., AND D. WEIL (1996): “The Gender Gap, Fertility and Growth,” *American Economic Review*, 86(3), 374–87.
- GREENWOOD, J., N. GUNER, AND J. KNOWLES (2003): “More on Marriage, Fertility and the Distribution of Income,” *International Economic Review*, 44(3), 827–862.
- GREENWOOD, J., A. SESHADRI, AND G. VANDENBROUCKE (2005): “The Baby Boom and Baby Bust,” *American Economic Review*, 95(1), 183–207.
- GRUBER, J., P. LEVINE, AND D. STAIGER (1999): “Abortion Legalization And Child Living Circumstances: Who Is The “Marginal Child”?,” *The Quarterly Journal of Economics*, 114(1), 263–291.
- HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2012): “Consumption and Labor Supply with Partial Insurance: An Analytical Framework,” Mimeo.

- HEER, B., AND A. MAUNER (2004): *DGE Models: Computational Methods and Applications*. Springer-Verlag.
- HENSHAW, S. (1998): “Unintended Pregnancy in the United States,” *Family Planning Perspectives*, 30.
- HONG, J. H. (2008): “Life Insurance and the Value of Spouses: Labor Supply vs. Household Production,” Mimeo.
- HONG, J. H., AND J.-V. RÍOS-RULL (2012): “Life Insurance and Household Consumption,” *American Economic Review*, 102(7), 3701–3730.
- 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.
- JUDD, K. L. (1998): *Numerical Methods in Economics*. MIT Press.
- KOCHARKOV, G. (2012): “Abortions and Inequality,” Mimeo, University of Konstanz.
- 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.
- ROSENZWEIG, M. R., AND T. P. SCHULTZ (1989): “Schooling, Information and Nonmarket Productivity: Contraceptive Use and its Effectiveness,” *International Economic Review*, 30(2), 457–477.
- SOARES, R. R., AND B. L. S. FALCAO (2008): “The Demographic Transition and the Sexual Division of Labor,” *Journal of Political Economy*, 116(6), 1058–1104.
- SOMMER, K. (2009): “Fertility Choice in a Life Cycle Model with Idiosyncratic Uninsurable Earnings Risk,” Mimeo, Georgetown University.
- TAUCHEN, G. (1986): “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions,” 20, 177–181.
- UPADHYAY, U. D., T. A. WEITZ, R. K. JONES, R. E. BARAR, AND D. G. FOSTER (2013): “Denial of Abortion Because of Provider Gestational Age Limits in the United States,” Forthcoming, *American Journal of Public Health*.
- WOLPIN, K. (1984): “An Estimable Dynamic Stochastic Model of Fertility and Child Mortality,” *Journal of Political Economy*, 92(5), 852–74.

7 Appendix

7.1 Extra Figures and Tables

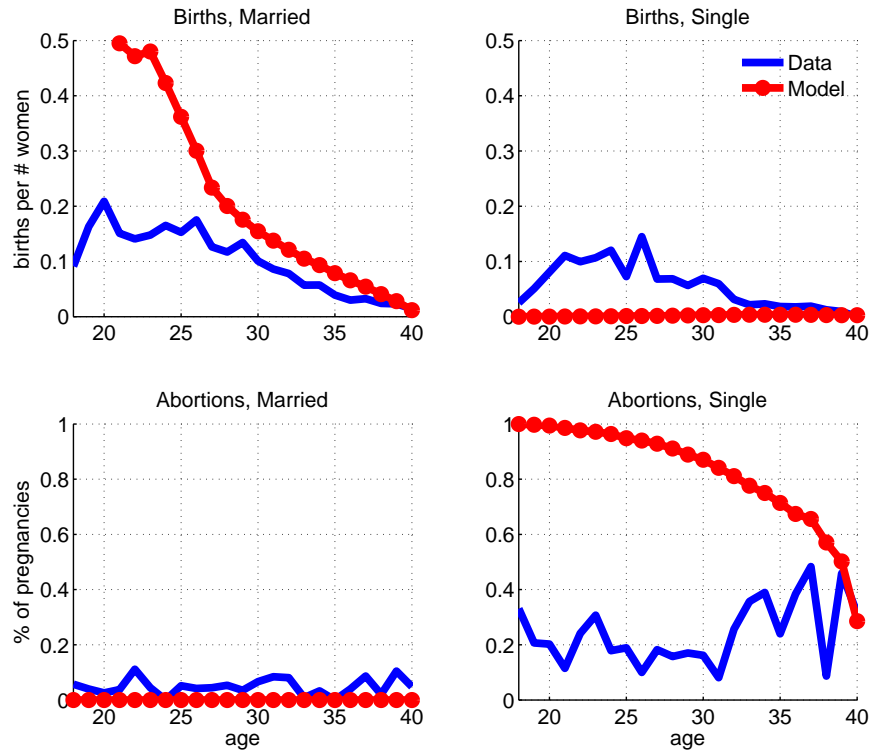


Figure 11: Births per woman and abortion rates, Model versus Data: High School (e)

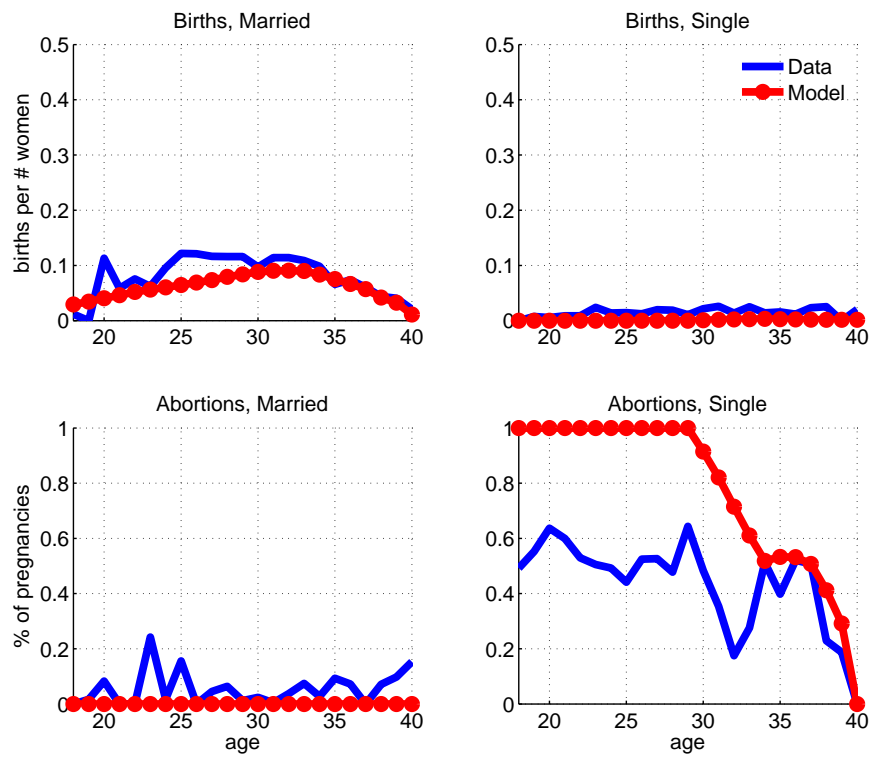


Figure 12: Births per woman and abortion rates, Model versus Data: College (\bar{e})

7.2 Data

NSFG: The National Survey of Family Growth 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 sexual activity and usage of birth control methods on a monthly basis for up to 5 years.¹⁴ Participants also answer questions on timing of births and pregnancy outcomes for all pregnancies conceived during that 5 year window. In terms of accuracy of this data, [Fu, Darroch, Henshaw, and Kolb \(1998\)](#) show that the introduction of computer assisted interviews in the NSFG for the year 1995 was accompanied by a reduction of 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 survey also contains information on educational attainment, marital status and other background information of the respondents.

I present age-specific fertility rates and percentage of aborted pregnancies in figure 1, separated by educational level: all those with a high school diploma or less ("High School") versus respondents with at least one year of college education ("College"). Age-specific fertility rates are computed as the number of live births by women of certain age and education group, divided by the total number of women in that group. Similarly, the abortion statistic reflects the number of pregnancies ending in an abortion, divided by the total number pregnancies in the group.¹⁵ In both figures I restrict the analysis to respondents ages 18 to 40. This choice follows two reasons: first, before age 18 and after age 40, there is not much information regarding both births and abortions (especially after age 40). Second, I want to abstract from fertility decisions of minors (those less than 18) since factors like education, cohabitation and parental influence might play important roles, which I abstract from in the quantitative model from the next section.

CPS: Figure 13 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.

To characterize the labor market, I also use gender and education specific idiosyncratic labor shocks. These shocks come from estimates from [Hong \(2008\)](#), 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

¹⁴Some participants have shorter retrospective information, if their active sex life has been shorter than 5 years.

¹⁵The figures are created by weighting each observation with the provided weight in the survey. The information at each age also compiles the information for each individual's five year retrospective information: for example, a 40 year old respondent in 1995 provides retrospective information from 1991 to 1995, which I average and use for the 40 year old group. Statistics using only the last 12 months of information for each respondent are very similar.

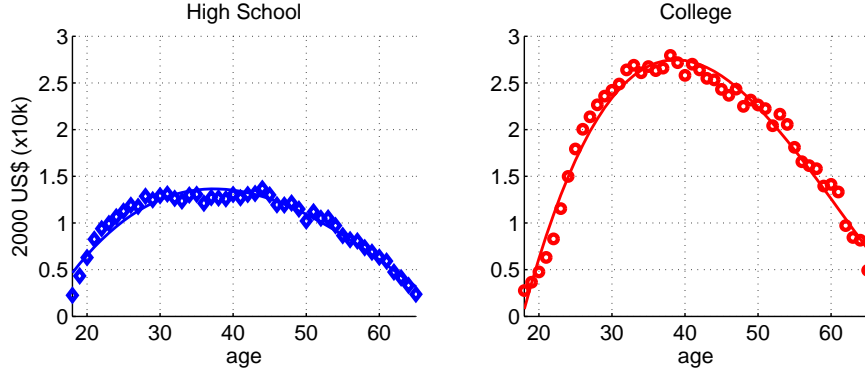


Figure 13: Average annual earnings for women, March Supplements of the Current Population Survey (1990-1995). The figures show the unconditional averages (i.e., not controlling for observables nor on being a worker or not) for *childless* women, and a 4-th order polynomial on age. "High School" is the group of all those with a high school diploma or less education; "College" is all the rest.

educated males (irrespective of presence of children in the household), in order to measure positive assortative matching in the marriage market.

I compute yearly survival probabilities by educational group using the information in [Hong \(2008\)](#). I interpolate his 5 year values and smooth the resulting series with a second order polynomial.

PSID: 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.

7.3 Further details on computation and the model

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 $P(i^* = i + 2) = 0.19$, 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 lessened.

For assets, I calculate from CPS data the average annual non-labor income (dividends, interests and rents) for both non-married males and females. Single males have on average 20% higher non-labor income than non-married 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.

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 6 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.