Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.

Fertility Choice in a Life Cycle Model with Idiosyncratic Uninsurable Earnings Risk

Kamila Sommer

2014-32

Fertility Choice in a Life Cycle Model with Idiosyncratic Uninsurable Earnings Risk

Kamila Sommer*†
Board of Governors of the Federal Reserve System

April 4, 2014

Abstract

This paper studies the link between rising income uncertainty and household fertility patterns in an Aiyagari-Bewley-Huggett framework augmented to include fertility decisions and infertility risk. Building on Becker and Tomes (1976), I model fertility decisions as sequential, irreversible choices over the number of children, accompanied by parental choices of time and money invested toward improving children's quality. The calibrated model is used to quantify the contribution of earnings uncertainty to the changes in the key fertility indicators between steady states. I show that realistic increases in uninsurable earnings risk lead to a postponement in births by young households, and are associated with a decline in the total number of births. The linkage between earnings risk and fertility patterns highlights the important role that labor market conditions can play in determining both short-term cyclical fluctuations in fertility (such as those in the recent U.S. data) and longer-term demographic trends (such as persistently depressed fertility rates in Southern Europe where youth unemployment rates are high and unemployment spell are very persistent).

JEL: E21, D91, J30

Key Words: Fertility Choice, Life Cycle, Heterogenous Agents, Uninsurable Idiosyncratic Income Risk.

*I would especially like to thank Jonathan Heathcote for his insightful comments and ongoing support. This project has benefited greatly from my discussions with Robert Cumby, Alessandra Fogli, Garance Genicot, Mark Huggett, Larry Jones, Greg Kaplan, Ellen McGrattan, Martin Sommer, Paul Sullivan and Victor Rios-Rull. I also thank Mathias Doepke, Raquel Fernandez, John Knowles, Christos Koulovatianos, Maria Luengo-Prado, Alice Schoonbroodt, Kjetil Storesletten, and Michele Tertilt for their valuable feedback and suggestions in the early stages of this project. A significant part of this research was completed during my stay at the Federal Reserve Bank of Minneapolis. All errors are mine. The views expressed herein are those of the author and not necessarily those of the Federal Reserve Board or the Federal Reserve System.

†Contact: kv28@georgetown.edu.

1 Introduction

Over the last four decades, the average total fertility rate (TFR) in OECD countries has fallen dramatically: from 2.9 in the 1960s to 2.0 in 1975, and then further down to 1.6 in 2000. The decline in fertility has been accompanied by a delay in childbearing: the average age at first birth in OECD countries has increased from 24.0 in 1970 to 27.0 in 2000.[1]

A number of candidate explanations have been put forward to account for declining fertility rates – I briefly survey this literature below. In this paper, I focus on a relatively unexplored mechanism: the link between delaying and reducing fertility on the one hand, and rising labor income risk on the other. In the next section, I will present evidence from micro data consistent with such a link. At the same time, one can motivate exploring a potential link on theoretical grounds. First, thinking of the decision to have a child as an investment in a lumpy durable good (Becker, 1960), recent work in the literatures on irreversible investment and "consumption commitments" (i.e., big-ticket goods with sizable adjustment costs) suggests that link between risk and fertility. For example, Chetty and Szeidl (2007) or Postlewaite, Samuelson, and Silverman (2008) show that consumption commitments can amplify risk aversion with respect to earnings shocks. If earnings shocks become larger, agents may therefore be less willing to commit to children. Fisher and Gervais (2011) show that, in the presence of large transactions costs, young households postpone homeownership when risk is high, preferring to initially rent and save more before buying a home. Using the same logic, if children are a durable good of irreversible nature that requires investment of parental resources, households could postpone (or abandon) childbearing when risk is high, initially preferring to work and save before starting a family.

At the macro level, falling fertility rates have been observed during periods when labor market risk was high. One striking example is provided by the experience of the Central and Eastern European transition economies during the early 1990s, when large increases in unemployment and earnings volatility brought about by the dissolution of centralized wage- and production- setting were associated with a dramatic decline in TFRs: the average TFR in the Central and Eastern European (CEE) region kept at a stable level of about 2.2 between

[1]The following countries were excluded from the calculation of the OECD average due to limited data availability: Australia, New Zealand, Mexico, Korea, and Turkey.

1970 and 1990, but fell dramatically to about 1.2 by the year 2000.[2] Large changes in the TFR and mean age at first birth have also been observed in Western Europe and the United States during periods of economic instability. For instance, the fertility rate fell dramatically during the Great Depression, a period characterized by record-high unemployment rates and high levels of earnings uncertainty. In contrast, the pickup of fertility in the post-war 1940s coincided with a booming U.S. economy and a compression of earnings inequality (see Kopczuk, Saez, and Song (2009) or Goldin and Margo (1992)). Since 2008, fertility has been declining amidst global financial turmoil and rapidly increasing unemployment, in part due to the postponement of births by younger age-cohorts. In particular, the birth rate for women aged 20-24 years declined by 7 percent in 2009, which represents the largest year-to-year decline in this rate since 1973.[3] Overall, the estimated number of births over woman's lifetime (also known as the total fertility rate) declined from 2.1 to 1.9 births per woman between 2007 – the recent peak – and 2011.[4]

Consistent with these observed macro trends, a number of empirical studies support the hypothesis that fertility responds negatively to labor market risk. Using a variety of OECD panel data, Adsera (2004, 2005, 2006) finds that high unemployment and unstable contracts, common in Southern Europe, depress fertility and decrease the probability of a transition to higher-order births.[5] Using the 1992-2002 data from the German Socio-Economic Panel (GSOEP), Bhaumik and Nugent (2006) estimate a strong negative relationship between unemployment uncertainty and childbearing decisions for East German women. Mira and Ahn

[2]The reported CEE average includes the following countries: Bulgaria, Czech Republic, Estonia, Hungary, Latvia, Lithuania, Poland, Romania, Russian Federation, Slovak Republic, and Slovenia.

[3]The birth rate for women aged 25-29 years declined in 2009 as well by 4 percent, as did the rates for women in their thirties. For details, see National Vital Statistics Report Volume 59, Number 3, available at www.cdc.gov/nchs/data/nvsr/nvsr59/nvsr59_03.pdf.

[4]National Vital Statistics Report Volume 61, Number 5, available at www.cdc.gov/nchs/data/nvsr/nvsr61/nvsr61_05.pdf.

[5]Using a panel of 23 OECD nations, Adsera (2004) finds that high unemployment and unstable contracts depress fertility, particularly of young women. Adsera (2005) uses the 1994-2000 waves of the European Household Panel Survey (EHPS) to reconstruct fertility histories of around 48,000 women in 13 West European countries. The estimated Cox proportional-hazard model predicts that in countries where male and female unemployment rates are similar and joblessness is short-lived, fertility rates are around 1.81 (a level close to the replacement rate). However, when unemployment rates are particularly high for women and unemployment is highly persistent, the estimated fertility is only around 1.28. Adsera (2006) uses the 1985 and 1999 Spanish Fertility Surveys to study differences between desired and actual fertility and finds that women in that mid-twenties facing high unemployment rates tend to reduce their fertility below desired levels.

(2002) show that the fertility rate responds negatively to unemployment over the business cycle in 15 out of 21 OECD countries.

In Section 2, I study fertility choices of households with different earnings risk profiles. Combining the estimates of occupational risk by Saks and Shore (2005) with the Census data, I find that the realized fertility of women with husbands in low-risk occupations (e.g., teachers or heath care professionals) is, on average, higher than fertility of women with husbands in high-risk occupations (e.g., sales or arts and sciences). This finding is robust across different age groups. In Figure 1, I explore variation in household fertility over the business cycle, and confirm that the U.S. fertility rate is pro-cyclical.[6] Since the household labor market risk is known to rise during recessions (Storesletten, Telmer, and Yaron, 2004), one interpretation of pro-cyclical fertility is that households postpone births when earnings uncertainty is high.

Starting with Section 3, this paper offers the first quantitative theoretical exploration of the link between earnings risk and fertility patterns. I develop a life cycle model of fertility choice in an Aiyagari-Bewley-Huggett style economy with incomplete markets and idiosyncratic labor market risk. I study unitary households where parents make joint decisions about consumption, savings, family size, and the allocation of resources invested into childrearing. Households face idiosyncratic wage shocks which can be partially self-insured by accumulating precautionary asset holding. Building on Becker (1960) and Becker and Tomes (1976), I model fertility decisions as sequential, irreversible choices over the number of children, accompanied by parental choices of time and money spent on improving children's quality. The decision to have another child can only be made during the first part of the life cycle when parents are fertile. The duration of this fertile period is, however, unknown to parents, who face idiosyncratic permanent infertility shocks. Infertility risk, while low early in the life cycle, increases exponentially with the age of the household. To my best knowledge, this paper is the first quantitative theoretical study where expenditures invested

[6] In order to study the co-movement between fertility and the business cycle, I use the general fertility rate since 1975. The general fertility rate is defined as the number of live births per 1,000 women aged 15-44 in a given year. Figure 1 plots a measure of economic cycle against the fertility rate detrended using the HP filter. The corresponding impulse response function points to a statistically significant relationship between output shocks and fertility changes during 1975-2008, with a lead time of about 1-3 years. The effect of an output shock on the fertility rate dies off after about 5-6 years.

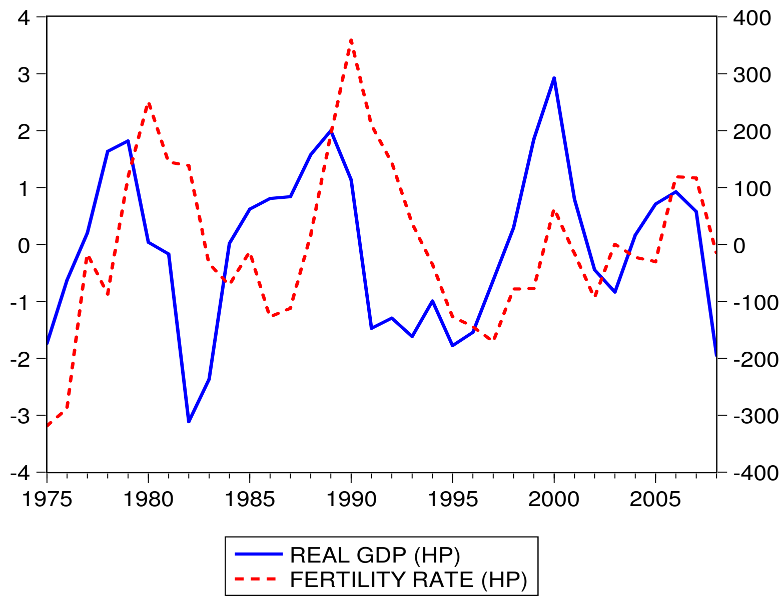

Figure 1: U.S. fertility rate: 1975 - 2008

Notes: Evolution of the general fertility rate over the past century. Source: U.S. Vital Statistics.

into childrearing (and children's quality) are determined endogenously within the model, together with number of children and timing of births. This paper is also the first to explore the role of infertility risk in explaining recent trends in household fertility.

The calibrated model is used to quantify the contribution of earnings uncertainty to the changes in the key fertility indicators between steady states. I show that realistic increases in uninsurable earnings risk lead to a postponement in births by young households, and are associated with a decline in the total number of births. The linkage between earnings risk and fertility patterns highlights the important role that labor market conditions can play in determining both short-term cyclical fluctuations in fertility (such as those in the recent U.S. data) and longer-term demographic trends (such as persistently depressed fertility rates in Southern Europe where youth unemployment rates are high and unemployment spell are very persistent).

The key mechanism generating the postponement of births and the fertility decline in the model is that children are discrete, irreversible choices, and that childrearing requires at least a minimum amount of investment per child. When markets are incomplete and households have limited access to credit, young parents with positive wealth may respond to a fall in household wages by temporarily dis-saving, increasing labor supply (and thus reducing the hours spent on childrearing), or reducing the market expenditures devoted to childrearing. Since parents prefer to smooth consumption, households initially choose to postpone childbearing when labor market risk is high, and work and save more instead. While parents may initially consider their decisions to delay childbearing as temporary, infertility risk means that delayed fertility translates into reduced total fertility. The longer the delay of first and higher-order births, the larger the reduction in fertility. Absent fertility risk, increased labor market risk has a smaller but still negative effect on total fertility.

Given the rich modeling framework that integrates quality-quantity trade-offs as well as infertility risk, the model offers a number of additional insights related to the response of household behavior to increases in labor market- and infertility risks. For example, it is shown that in economies with higher levels of earning risk, the average quality per child suffers in spite of the fact that households have on average fewer children. This is because parents must work more and save more to insure against the adverse economic environment and, hence, invest fewer resources into their children's quality. Additionally, the calibrated model is used to assess the role of infertility risk in determining household fertility patterns. The model predicts that having perfect control over their fertile window appears to allow households to more optimally time births without running a risk of ending up with suboptimal family sizes: for a given level of earnings uncertainty, both total fertility and mean age at births are always higher in a model with no fertility risk. The model predictions with regard to infertility risk are particularly interesting when viewed in the context of the ongoing advances in reproductive science. Broadly speaking, the results suggest that wider access to reproductive technologies can help to offset some of the adverse effect of increased earnings risk on completed fertility.

A vast body of studies in microeconomics, labor economics, and macroeconomics have explored channels that likely contribute to fertility changes. I will briefly review several recent studies of fertility choice in the quantitative theoretical tradition that are most closely related to this paper. Da Rocha and Fuster (2006) show that high unemployment risk induces women to postpone and space births, which in turn reduces the fertility rate. In their model with job search and human capital accumulation, young women facing low job-finding rates may space births in order to avoid costly career interruptions following the child's birth. Moreover, similarly to this paper, children are costly in terms of goods, so young women may postpone births in order to smooth consumption. Other papers try to connect three trends: increasing female education, increasing female labor market participation, and declining fertility. Conesa (2000) suggests that changes in the timing of fertility decisions resulting from increasing female access to higher education can partially account for the recent fertility decline in advanced economies. In a related study, Caucutt, Guner, and Knowles (2002) argue that better education can explain less than one-third of the increase in mean age at birth, and that the delayed fertility is driven by changes in the marriage markets and increasing returns to female labor market participation. Education and female participation are beyond the scope of the present paper, but it would be interesting to explore the feedback from fertility to these trends. In particular, if women choose to delay fertility in response to labor market risk, they have more time available for education and work.

In terms of the longer-term demographic trend, Greenwood and Seshadri (2002a,b), or Doepke (2004) explain the decline as a result of a production shift from low-skill, labor-intensive agriculture to high-skill manufacturing. Greenwood, Seshadri, and Vandenbroucke (2005) argue that the secular decline is due to the relentless rise in real wages that increased the opportunity cost of children, while the baby boom of the 1950s and 1960s can be explained by an technological progress that lowered the cost of having children. Knowles (2007) argues that improved opportunities for contraception can partly account for increasing female labor participation and lower fertility rates, a theory consistent with the empirical evidence in Goldin and Katz (2002). Boldrin, De Nardi, Jones, and Madrid (2005) argue an increase in government provided old-age pension can help explain the both the long-run demographic decline as well as account for the differences in fertility rates across countries.[7]

The paper is organized as follows. Section 2 presents empirical evidence on the relationship of idiosyncratic labor risk and household fertility. In Section 3, I develop a life cycle model with heterogeneous households, fertility choice, and idiosyncratic earnings risk.

[7] These papers focus only on explaining the secular decline in the fertility rates, but ignore the corresponding changes in timing of births.

Section 4 describes the model's calibration. In Section 5, I discuss the predictions of the benchmark model. Section 6 quantifies the extent to which increases in earnings risk affect household fertility decisions, and discusses the model predictions in the context of the U.S. time-series data. Section 7 studies the role of infertility risk in the model, while Section 8 concludes the paper.

2 Occupational risk and fertility in cross-sectional data

In this section, I examine whether women in high earnings-risk households have different fertility patterns than women in households subject to low levels of earnings risk. In order to answer this question, I use the riskiness of husband's occupation as a proxy for the earnings risk faced by households.[8] Using the PSID income data for male household heads, Saks and Shore (2005) find that teachers, health-care professionals, and engineers face the lowest levels of earnings uncertainty, while men with occupations in math and sciences, sales, and arts and entertainment typically experience high levels of earnings risk.

To construct the data on fertility of households from these occupational groups, I use the 5 percent sample of the 1990 Decennial Census, concentrating on married couples where the husband is not self-employed.[9] Following the estimation strategy of Saks and Shore (2005), I control for the educational attainment of husbands.[10] After the selection criteria are applied, the sample consists of roughly 100,000 married couples with wives between ages 20 and 43 years.

Figure 2 plots the average number of births from the Census data against the estimates of the occupational earnings risk by Saks and Shore (2005) for various age-groups. I find that that the realized fertility of women with husbands in low-risk occupations is, on average, higher than fertility of women with husbands in high-risk occupations. This finding is robust

[8] Occupation is considered a career choice that is connected with a significant accumulation of human capital. Since changes in occupation typically involve large losses of the accumulated human capital, the perceived riskiness of the occupation represents a good proxy for the perceived riskiness of lifetime income.

[9] Since self-employed individuals have been shown to face higher earnings risk than individuals working for wage or salary across occupations (Carroll and Samwick (1997), Carroll and Samwick (1998), or Saks and Shore (2005)), I focus on individuals who are not self-employed. The fertility patterns of the self-employed exhibit similar patterns.

[10] Saks and Shore (2005) estimate the occupational risk for male heads with at least a college degree. For details on sample selection in this paper, see Appendix A.

Figure 2: Fertility and earnings risk of husband's occupation

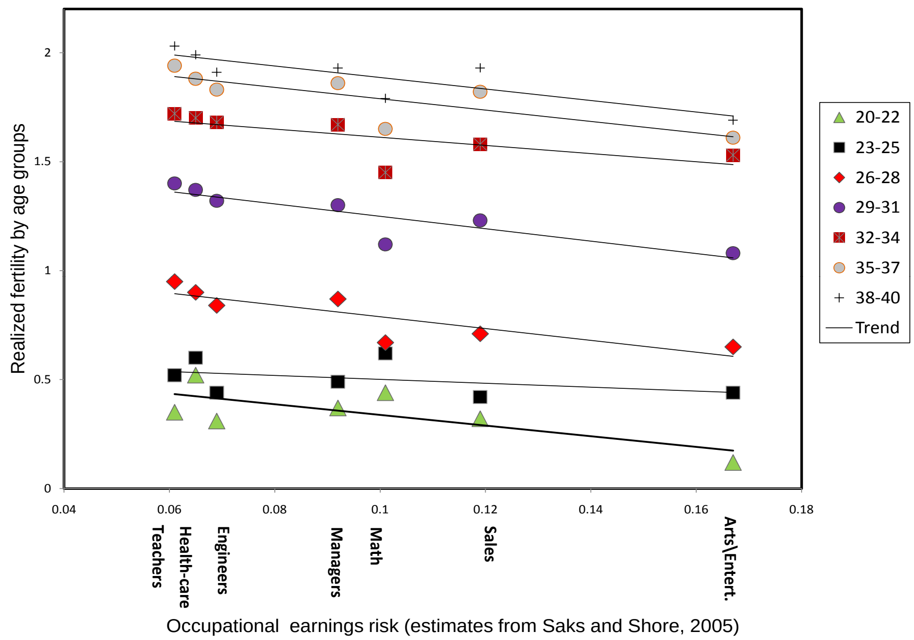

Notes: Age-specific fertility of wives by the riskiness of their husbands' occupation. Occupational earnings risk estimates: Saks and Shore (2005).

across different age groups.[11]

To examine the negative correlation between fertility and earnings risk in a more formal framework, I estimate a simple OLS model of completed fertility using the cross-sectional data set from Census. The dependent variable is the number of births for any given couple. The regressors include occupational dummies and other basic household characteristics such as wife's age, and wife's and husband's income. In Table 1, the occupational dummies are ordered by the earnings risk: the higher the riskiness of the occupation, the lower the rank in the table. The dummy for teachers (i.e., the lowest-risk occupation) is omitted from the regression, and the estimated dummies therefore correspond to the difference in fertility between a given occupational group and fertility of families where husband (teacher) has a low earnings risk.

[11]While household formation and marriage are beyond the scope of this paper, it would be interesting to explore whether women who desire large family sizes select husbands in low risk occupations. For a model of household sorting and self-selection in household formation, see, for example, Fernandez, Guner, and Knowles (2005).

As can be seen in the table, the correlation between the number of births and the riskiness of husband's occupation remains, on average, negative, since the estimated dummies tend to be lower (more negative) for high-risk occupations. Interestingly, the effect of husband's income level on the number of births is positive, indicating that the demand for children rises with household income. On the other hand, the negative relationship between wife's income and the number of births is consistent with the "price of time" theory which posits that higher-earning women have smaller families due to the higher opportunity cost of raising children. The coefficients on all variables are statistically significant at the one percent level, except for the occupational dummy for health care professionals which is significant at the 10 percent level.

Table 1: Regression analysis

Variable	**Coefficient**	**(Std. Err.)**
age of wife	0.390	(0.006)
age of wife squared	-0.005	(0.000)
health care	-0.036	(0.022)
engineer	-0.148	(0.015)
manager	-0.129	(0.013)
math	-0.204	(0.061)
sales	-0.158	(0.014)
arts and entertainment	-0.305	(0.028)
husband's total income	8.06e-07	(0.000)
wife's total income	-1.77e-05	(0.000)
Intercept	-5.987	(0.107)
N	103271	
R^2	0.254	
$F_{(10,103260)}$	3519.162	

3 The benchmark model

To study the impact of changes in labor market risk on household fertility decisions in a formal framework, I build a model based on an Aiyagari-Bewley-Huggett style economy with incomplete markets and uninsurable earnings risk. Building on Becker and Tomes' (1976) idea of children as a durable good whose "quality" can be improved by parental investment of time and money, I next add a dynamic sequential fertility choice, and endogenize the allocation of resources devoted toward childrearing. I consider unitary households where parents make joint decisions about consumption, saving, their family size, and the inputs invested toward improving their children's quality.[12]

The model mirrors the following assumptions. Young households, which start their life cycle childless and with zero asset holdings, have limited access to credit and face idiosyncratic earnings shocks which can be partially self-insured by accumulating precautionary asset holdings. Parents enjoy having children and care for their children's quality which is secured through parental inputs of time and market goods. Children are discrete and irreversible choices that are born in increments of one (e.g., no twins are allowed). The decision to have another child can be made only during the first part of the life cycle when parents are fertile. The exact timing of the last fertile period is, however, unknown to parents who face infertility shocks which render them permanently infertile.

3.1 The demography and endowments

The model economy is inhabited by a continuum of the same-age husband-wife households with identical preferences. The model period is one year. Households start their life together at age 16, and live until age 80 with certainty. During the first 50 periods of life, the household wage process is determined according to an idiosyncratic stochastic process:

$$\ln w_t = \ln w_0 + h(t) + \epsilon_t + \nu_t, \tag{1}$$

where $h(t)$ governs the average age-profile of wages, $\nu_t \sim N(0, \sigma_\nu^2)$ is a transitory shock to income received every period, and ϵ_t is a persistent shock, also received each period, which follows a first-order autoregressive process:

$$\epsilon_t = \rho\epsilon_{t-1} + \psi_t \text{ with } \psi_t \sim IID(0, \sigma_\epsilon^2) \text{ and } \epsilon_1 = 0. \tag{2}$$

[12] In future research, I plan to relax the assumption of unitary households and model male and female wage processes separately.

At age 65, households retire and receive a lump-sum pension:

$$w_t = \bar{w} \text{ for } t > 65, \tag{3}$$

in the form of a transfer from the government.

3.2 Preferences

In the spirit of Becker and Tomes (1976), I assume that each household has a utility function of the form:

$$U = U(c_t, n_t, q_t), \tag{4}$$

where c_t stands for the parental consumption of a nondurable market good, n_t is the number of children at home, and q_t is the quality of each child. In this model, households are not altruistic toward their offspring, leaving no bequest to their children. The expected discounted lifetime utility can then be written as:

$$E_0 \sum_{t=16}^{80} \beta^{t-16} U(c_t, n_t, q_t),$$

with the discount factor $\beta \in (0, 1)$.

The quality of children is determined by parents through their inputs of time, l_t, and goods, x_t, spent on childrearing. Similarly as in Becker and Tomes (1976), I assume that the quality of each child within a family is the same and takes on the form:

$$q_t = f(x_t, l_t, n_t), \tag{6}$$

where x_t and l_t are the total amount of goods and the fraction of time invested toward childrearing ($l_t \in [0, 1]$), respectively. Although households do not value leisure, their labor supply is determined endogenously within the model as a fraction of the total time that is not spent on childrearing. Lastly, while household spending on children is discretionary, a minimum level of investment in children's quality is required for families with children so that $q_t \geq \underline{q}$ if $n_t > 0$.

3.3 Process for dependent children

In order to build a model that can match the life cycle household patterns of parental spending on children, I assume that parents enjoy and make expenditures only on children who are young and live at home. Ideally, one might like to think of such children as children younger than a certain age. However, the recursive structure of this model makes keeping track of children's ages difficult, as it requires integrating a history of the past fertility decisions into the state space of the problem (for details, see Hotz and Miller (1988)).[13] For the purposes of this paper, I assume that parents have two types of children: children who are young and still live at home (n_t), and children who have become financially independent and have left home. The law of motion of the children ever born to the household (n_t^b) is deterministic and follows the process:

$$n_{t+1}^b = n_t^b + K_t \text{ where } K_t = \{0, 1\}, \tag{7}$$

with $K_t = 1$ when a household has an additional child next period and $K_t = 0$ otherwise. The number of dependent children which still live at home is assumed to be distributed binomially:

$$n_{t+1} \sim Bi(n_t + K_t, p) \text{ with } n_{16}^b = n_{16} = 0, \tag{8}$$

where p is a time-invariant probability that a child becomes independent and leaves home.

3.4 Infertility risk

In order to build a further element of realism into the model, households face a binary idiosyncratic infertility shock $f_t = \{I, F\}$ which arrives at the beginning of every period. Only parents that are fertile in a given period (i.e., $f_t = F$) can choose to have another child, while parents once hit by the infertility shock remain infertile forever (i.e., if $f_t = I$, then

[13] The history of the past fertility decisions would be summarized by a vector of zeros and ones, depending on whether the household had a child or not at each period of the fertile horizon. When the number of fertile periods is large, keeping track of children's ages becomes computationally intractable.

$f_{t+j} = I \ \forall j \geq 0$). The time-variant probability p_t^I that a household becomes permanently infertile is assumed to rise exponentially with the age of parents, t, and to become one at age 45 after which no households are able to conceive a child:[14]

$$p_t^I = \begin{cases} \omega_1 \exp^{\omega_2 t} & \text{if } t < 45 \\ 1 & \text{if } t \geq 45. \end{cases} \tag{9}$$

3.5 Dynamic program of fertile parents

Parents who have not lost their ability to bear children (i.e., $f_t = F$) solve the recursive problem:

$$V_t(a_t, n_t, w_t, f_t = F) = \max_{c_t, a_{t+1}, x_t, l_t, K_t = \{0,1\}} u(c_t, n_t, q_t) + \beta E_t V_{t+1}(a_{t+1}, n_{t+1}, w_{t+1}, f_{t+1} = \{I, F\})$$

subject to

$$A_{t+1} = \begin{cases} (1+r)(A_t + (1-l_t)w_t - c_t - x_t) & \text{if } t \leq R\,; \\ (1+r)(A_t - c_t + \bar{w}) & \text{if } R < t \leq T\,, \end{cases} \tag{10}$$

$$q_t = f(x_t, l_t, n_t) \text{ with } q_t \geq \underline{q} \text{ if } n_t > 0, \tag{11}$$

$$n_{t+1}^b = n_t^b + K_t \text{ with } n_{16}^b = 0, \tag{12}$$

$$n_{t+1} \sim Bi(n_t + K_t, p) \text{ with } n_{16} = 0, \tag{13}$$

$$A_{t+1} \geq 0, \tag{14}$$

by choosing the parental consumption of the nondurable market good, c_t, savings, A_{t+1}, and the time, l_t, and market goods, x_t, inputs into the production function for the children's quality, $f(x_t, l_t, n_t)$, with $\underline{q}$ imposing a lower bound on children's quality. In addition to these continuous choices, households also make a discrete decision whether to have a child next period ($K_t = 1$) or not ($K_t = 0$). Households also face uncertainty about their ability to bear children next period: $f_{t+1} = \{I, F\}$ follows the process from Section 7. Equation (12) determines the law of motion for the stock of children ever born, n_t^b, to the household, while

[14]The fertile horizon of the household is based on the fertility cycle of the wife and reflects the fact that most women do not give birth after the age of 45.

equation (13) summarizes the law of motion for children at home, n_t. In the baseline model, no borrowing is allowed (equation 14). $(1+r)$ is the gross rate of return on a single asset in the economy; and w_t represents household wages and follows the process described in the equation (1).

3.6 Dynamic program of infertile parents

Parents who have lost their ability to bear children (i.e., $f_t = I$) can no longer increase their family size and, therefore, solve the problem:

$$V_t(a_t, n_t, w_t, f_t = I) = \max_{c_t, a_{t+1}, x_t, l_t} u(c_t, n_t, q_t) + \beta E_t V_{t+1}(a_{t+1}, n_{t+1}, w_{t+1}, f_{t+1} = I),$$

by choosing the optimal allocations of consumption, savings, and resources devoted to childrearing, subject to the set of constraints and transition equations (10), (11), and (14), and subject to the law of motion:

$$n_{t+1} \sim Bi(n_t, p) \text{ with } n_{16} = 0. \tag{15}$$

4 Calibration

The calibration strategy involves fixing some parameter values exogenously, and estimating the remaining parameters using the method of simulated moments based on cross-sectional patterns of fertility, income, consumption, and saving. Table 2 summarizes parameters which were drawn from other studies or were calculated directly from the data. Table 3 contains eight estimated parameters based on moments described in Table 4 that are constructed using the data from the 1979 National Longitudinal Survey of Youth (NLSY79), the 2004 American Time Use Survey (ATUS), and the 2000 waves of the Consumption Expenditure Survey (CEX). Appendix A.1 provides details on the sample selection and the calculation of moment conditions from these data sets.

4.1 Infertility risk and and earnings process

Trussell and Wilson (1985) provide point estimates for the fraction of couples who are permanently infertile by the woman's age. The authors' point estimates, fitted by an exponential function in t, represent the benchmark cumulative distribution function (c.d.f.) of the permanent infertility risk (Figure 3). The c.d.f. of the permanent infertility risk is in turn used to calculate the sequence of the time-variant probabilities, p_t^I, in equation (9) which are derived so that the fraction of permanently infertile households of any given age in the model matches exactly the corresponding fraction in the data.[15] Wallace and Kelsey (2010) provide further evidence that women's ability to conceive declines dramatically with age. Using data for women in the U.K., U.S. and Europe to estimate ovarian reserves, the authors find that for 95 percent of women by the age of 30 years only 12 percent of their eggs are present and by the age of 40 years only 3 percent remains (see their Figure 5; shown here in the Appendix A.2). While assisted reproductive technology has much improved since the first successful in-vitro fertilization (IVF) transfer in the early 1980s, its costs run high. For example, an average IVF cycle costs between \$10,000 and \$20,000. At the same time, the probability that the embryo transfer results in a live birth is 38 percent for women under age 35 and only 22 percent for women between ages 38 and 40, according to Mayo Clinic's website.[16]

Three parameters are needed to parametrize the stochastic components of the idiosyncratic earnings process in equation (1): the serial correlation coefficient, ρ, the standard deviation of the innovation term, σ_ϵ, for the persistent shock, and a standard deviation of the innovation, σ_ν, for the transitory shock.

Various authors have estimated the stochastic process for logged labor earnings using the PSID data. Controlling for household observable characteristics (such as education and age), Card (1991), Hubbard, Skinner, and Zeldes (1995), and Storesletten, Telmer, and Yaron (1998) estimate a ρ in the range from 0.88 to 0.96, and a σ_ϵ in the range between 0.12 and 0.25. Assuming the presence of a unit root, Meghir and Pistaferri (2004) find that σ_ϵ

[15]In the data, about 97 percent of all couples are infertile at age 45. In the model, the cumulative probability that a household is permanently infertile at age 45 is set to 1.

[16]http://fertilitysuccessrates.com/ivfclinics/Rochester/Mayo-Clinic-Assisted-Reproductive-Technologies/report204.html

increased from about 0.15 in the 1970s to 0.21 in the 1980s.[17] Meanwhile, the estimates for σ_ν range between 0.15 and 0.24.

For the purposes of this paper, ρ and σ_ν are set to the middle of the spectrum of the available estimates, i.e., 0.95 and 0.17, respectively. Since the model is calibrated to match fertility choices of the NLSY79 cohort of agents who mostly made their fertility decisions in the 1980s and 1990s, my choice for σ_ϵ of 0.21 lies at the upper end of the available estimates, as work by Meghir and Pistaferri (2004) suggests that households in the 1980s and the 1990s faced on average a higher level of persistent labor earnings uncertainty than the earlier cohorts.

To avoid numerical integration, earnings process (1) is implemented as a discrete approximation to the otherwise continuous earnings process. The autoregressive process is approximated with a seven-state Markov chain with innovations being i.i.d. and transition probabilities chosen following Tauchen (1986). For the transitory shocks, I use an i.i.d. two-state Markov chain.

The average age-profile for wages, $h(t)$, is calculated from the 2004 CPS by dividing the family labor income, defined as a sum of yearly earnings of both spouses in husband-wife families, by the sum of total hours worked by the couple (Figure 4). The average age of the couple is taken to represent the age of the household. The profile is smoothed using a cubic polynomial in age.

Retired households receive a pension transfer $\bar{w}$ which is proportional to the household earnings in the last working period, with a replacement rate b. Using the Health Retirement Survey data and the Social Security Administration records, Munnell and Soto (2006) report that the 1999-2002 median replacement rate for newly retired workers was about 42 percent of worker average indexed earnings (higher for earnings-poor individuals and lower for earnings-rich individuals, due to the progressiveness of the system). On a household basis, the Social Security benefits provide an average replacement rate of 44 percent; 58 percent for a couple with a non-working spouse and 41 for couples where both spouses work. For the purposes

[17] Using PSID data, Meghir and Pistaferri (2004) provide historical estimates for the variance of persistent shock for the period 1969-1991. Since the variance is known to fluctuate year by year, I compute the mean variance for the 1980s and 1990s by taking an arithmetic average across the authors' variances for the periods 1970-1979 and 1980-1989.

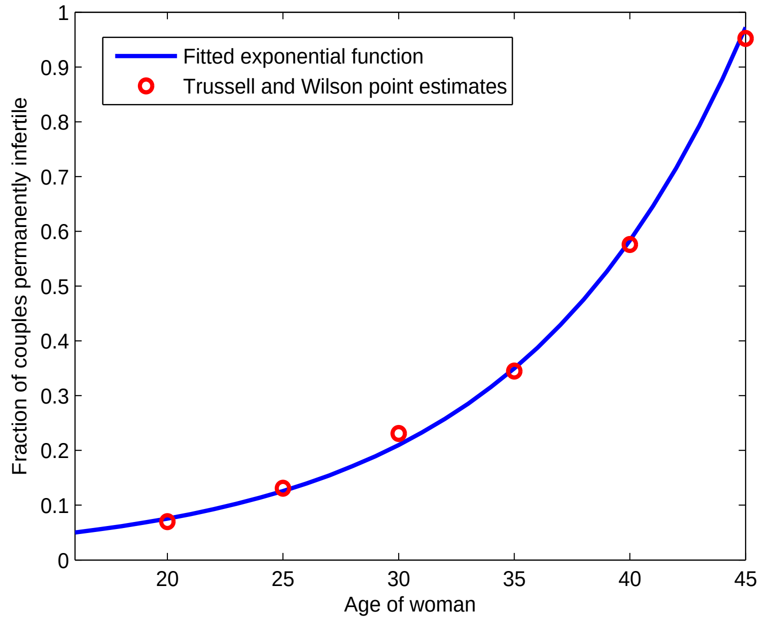

Figure 3: Fraction of couples permanently infertile by age of wife

of this paper, the replacement rate b is set to 0.4.

4.2 Preferences

Following the literature on fertility choice (see, for example, Becker, Murphy, and Tamura (1990), Ranjan (1999), de la Croix and Doepke (2003), or Jones, Schoonbroodt, and Tertilt (2008)), the preferences are modeled as additively separable between consumption and fertility choices (i.e., the number of children and the children's quality):

$$U(c,n,q) = \frac{c^{1-\gamma}}{1-\gamma} + \zeta \frac{(nq)^{1-\kappa}}{1-\kappa}, \tag{16}$$

with $\gamma > 0$ and $\kappa > 0$. The constant relative risk aversion preferences over consumption are standard, and are characterized by the risk aversion coefficient, γ, which determines the household desire to smooth consumption across time and states. The existing estimates of γ typically lie in a range between 1 and 3. To model household preferences over the number of children and their quality, I adopt a generalized version of the preference specification

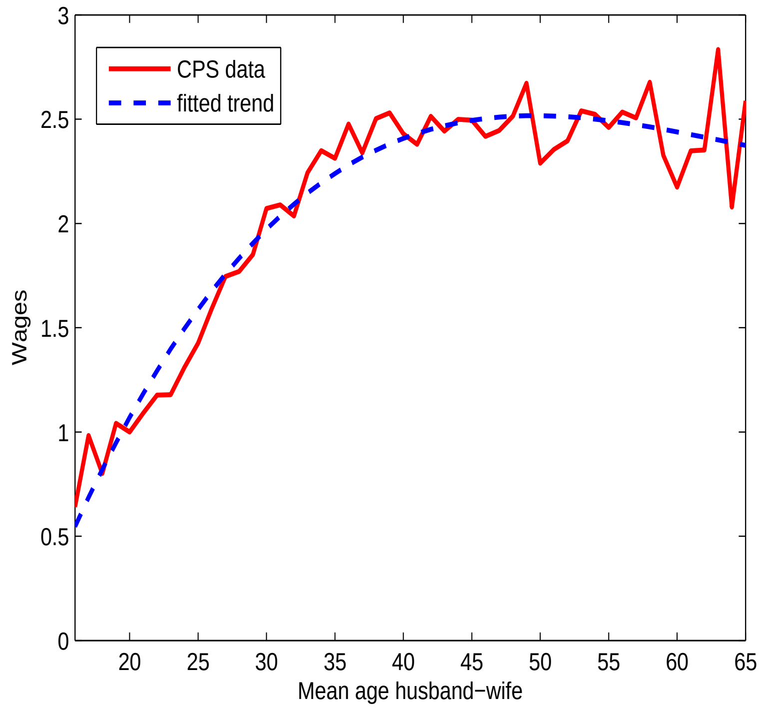

Figure 4: Mean wages for husband-wife households by mean age of couple

in de la Croix and Doepke (2003).

In order to characterize the household preferences described in equation (16), four parameters are needed: the three which determine the utility function (γ, κ, ζ), plus the discount factor β. In order to minimize the number of calibration targets, the annual gross interest rate $(1+r)$ is set equal to 1.04 so that $\beta = \frac{1}{1+r}$. I set the relative risk aversion, γ, to 1.5. The remaining two preference parameters ζ and κ are calibrated.

4.3 Production function for children's quality

The production function for the children's quality takes on the constant elasticity of substitution (CES) form

$$q_t = [\mu(\frac{x_t}{n_t^{\psi_1}})^\theta + (1-\mu)(\frac{l_t}{n_t^{\psi_2}})^\theta]^{1/\theta}, \tag{17}$$

where $\mu \in [0,1]$ is the production share, and $\frac{1}{1-\theta}$ with $\theta \in (-\infty, 1]$ represents the elasticity of substitution between time (l_t) and goods (x_t) devoted to childrearing, while parameters

ψ_1 and ψ_2 represent the household economies of scale in the time and market expenditures spent on childrearing. The CES production function is popular in applied research due to its flexibility regarding the degree of substitution between production inputs.[18] Since very little is known about the degree of substitutability of time and market expenditures in children's production, no a priori assumption is made about the value of θ. Instead, both CES production parameters μ and θ as well as the parameters ψ_1 and ψ_2 are calibrated. The lower bound on children's quality, $\underline{q}$, from section 3.2 is calibrated as well.

4.4 Process for dependent children

In order to determine process (8), a value for the time-invariant probability p that a child leaves home is needed. Since a child can separate from the household in any period, p is calibrated so that the number of children living with mature-age parents at home matches the number of children living at home in the data.[19]

Table 2: Exogenous parameters

Parameter	Value
Gross interest rate $(1+r)$	1.04
Discount factor β	$\frac{1}{1+r}$
Risk aversion coefficient γ	1.5
Age-profile of wages $h(t)$	computed from 2004 CPS
Persistence coefficient ρ	0.95
Std. of persistent shock σ_ϵ	0.21
Std. of transitory shock σ_ν	0.17
Replacement rate b	0.40

4.5 Moment conditions for the simulated method of moments

Based on the previous discussion, eight structural parameters must be calibrated: the scale and curvature preference parameters, ζ and κ; the production share, μ; the elasticity of

[18] When $\theta = 1$, production inputs are perfect substitutes. Conversely, when $\theta = -\infty$, the inputs are perfect complements. $\theta = 0$ gives a Cobb-Douglas production function.

[19] Besides data on the number of children born to respondents, NLSY79 also collects information about the number of children living at home.

Table 3: Calibrated parameters

Parameter	Value
Preference curvature κ	0.14
Preference scale ψ	3.34
Production share μ	0.33
Elasticity of substitution in production $\frac{1}{1-\theta}$	$\frac{1}{1-0.73}$
Lower bound on children's consumption $\underline{q}$	0.34
Household economies to money input to production ψ_1	0.66
Household economies to time input to production ψ_2	0.54
Probability that a child stays at home $(1-p)$	0.98

Table 4: Moments targeted in the estimation

Calibration target	Data	Model	Data Source
Completed fertility of households	1.90	1.90	NLSY79
Mean age at 1st birth	25.5	25.5	NLSY79
Mean number of children at home for households age 35	1.43	1.43	NLSY79
Expenditures on childrearing to earnings	0.40	0.40	Lino (2008)
Elasticity of market expenditures w.r.t. number of children	0.34	0.43	CEX
Elasticity of childrearing time w.r.t. number of children	0.25	0.25	ATUS
Correlation between earnings and fertility at age 20	-0.20	-0.20	NLSY79
Correlation between earnings and fertility at age 45	-0.02	-0.02	NLSY79

substitution between time and market goods in the production function, θ; the lower bound on the children's quality, $\underline{q}$; the parameters of the economies of scale, ψ_1 and ψ_2 in the production function of children's quality; and the probability that a child leaves home at any given period, p. Let $\Theta = (\zeta, \kappa, \mu, \theta, \underline{q}, \psi_1, \psi_2, p)$ define the vector of structural parameters to calibrate. The parameter values Θ are determined so that the resulting statistics in the model economy $G_j(\Theta)$ are determined by the eight specified targets G_j for $j = 1, \ldots, 8$ measured in the U.S. cross-section.

The data for the eight targets come from three different sources: NLSY79, ATUS, and CEX. The NLSY data is used to estimate the average number of children ever born to a household (1.9), the mean age of a parent at first birth (25.5), and the average number of children at home for parents age 35 (1.4), as well as the correlation coefficient between number of births and labor earnings for parents of age 20 (-0.20) and of age 45 (-0.02). The average childrearing expenditures to labor income ratio (0.40) for households with children is drawn from Lino (2008) who, using the Consumer Expenditure Survey (CEX), estimates that an average dual-earner household with two children between ages 0 and 17 spends roughly 40 percent of the household earnings on direct expenses connected with childrearing (e.g., food, housing, education, transportation, babysitting, and daycare).

The ATUS and the CEX are used to help determine the economies of scale to market goods (ψ_1) and time (ψ_2) inputs in childrearing. Previous studies (Doepke, Hazan, and Maoz (2007)) used the slope coefficient from the regression of logged time on a constant and a logged number of children in the household to represent the economies of scale to time input. However, setting ψ_2 to the slope coefficient can be misleading if households trade off quality per child for bigger family sizes (e.g., $\frac{\partial q_t}{\partial n_t} < 0$) since, in such a case, the estimated slope coefficient overstates the true economies of scale. To estimate the parameters ψ_1 and ψ_2 directly, the method of indirect inference is applied (for an overview, see Smith (2008)). First, in order to pin down the household economies to the expenditure input ψ_1, I run an auxiliary regression,

$$\ln x_t = \alpha_0 + \alpha_1 \ln n_t, \tag{18}$$

using the 2000 CEX data, with x_t representing the total children-specific expenses and n_t determining the number of own children in the household. The slope coefficient α_1, estimated at 0.34, represents the elasticity of childrearing expenditures with respect to the number of children at home. Similarly, in order to estimate the economies of scale to the time input ψ_2, the 2004 ATUS data are used to run an OLS regression

$$\ln l_t = \gamma_0 + \gamma_1 \ln n_t, \tag{19}$$

with l_t representing the total time spent by respondents on childrearing and n_t capturing the number of own children at home. The estimated elasticity of time spend on childrearing with respect to the number of children is 0.25. Equations (18) and (19) provide the last two moment conditions for the method of simulated moments, with the elasticities α_1 and γ_1

yielding calibration targets.[20]

5 Fit of the benchmark model

Figure 5 shows the baseline life cycle profiles. The generated profiles of earnings, savings, and consumption in Panel A follow expected patterns.[21] Panel B compares the age-specific cumulative births generated by the model with the corresponding NLSY79 estimates. The simulated birth profile matches the data well for households between ages 30 and 45, although the average number of births by the younger agents differs slightly from the data, in part due to the low levels of household heterogeneity early in the life cycle. Panel C shows the simulated age-profile for the number of children at home. The model, calibrated to match the number of children living at home for households aged 35, does a very nice job matching the data for households younger than 40 years of age. However, the model tends to overstate the number of children at home for older households.[22]

Allocation of Resources toward Childrearing: In addition to the standard profiles, I next explore the predicted patterns of household expenditures on children. The high elasticity of substitution between time and market expenditures in childrearing – estimated at about 3.7 (see Table 3) – has implications for the allocation of resources devoted toward childrearing across wage groups and along the life cycle. First, in the model, low-wage households have a low opportunity cost of spending time at home and, as such, specialize in home production of children's quality. Since the opportunity cost of staying at home and caring for children increases with household wages, high-wage households prefer to substitute time at home for market expenditures.[23] Second, in a model with deterministic wage growth

[20]Indeed, the model predicts that the actual values for the economies of scale parameters ψ_1 and ψ_2 lie well above the corresponding elasticities α_1 and γ_1. This implies that parents face trade-offs between the number and quality of children, and reduce the quality of each child in order to increase their family size.

[21]The household earnings approximately triples between ages 25 and the peak at age 58. At the same time, the savings profile generated by the model peaks in the retirement year at about 3.7 times of the mean earnings. The evidence from the Survey of Consumer Finances (SCF) 2007 suggests that the age-profile of households peaks roughly at age 55 and triples between age 22 and the peak, while the savings profile peaks at age 65 at about 3.5 times of mean earnings.

[22]In an alternative calibration exercise, the number of children at home at the household age of 45 was targeted to pin down the parameter p. Such calibrated model tends to understate the number of children at home for households between ages 30 and 45, but improves the match for households aged 45+.

[23]The model generates a correlation between household wages and time spent (expenditures) on childrea-

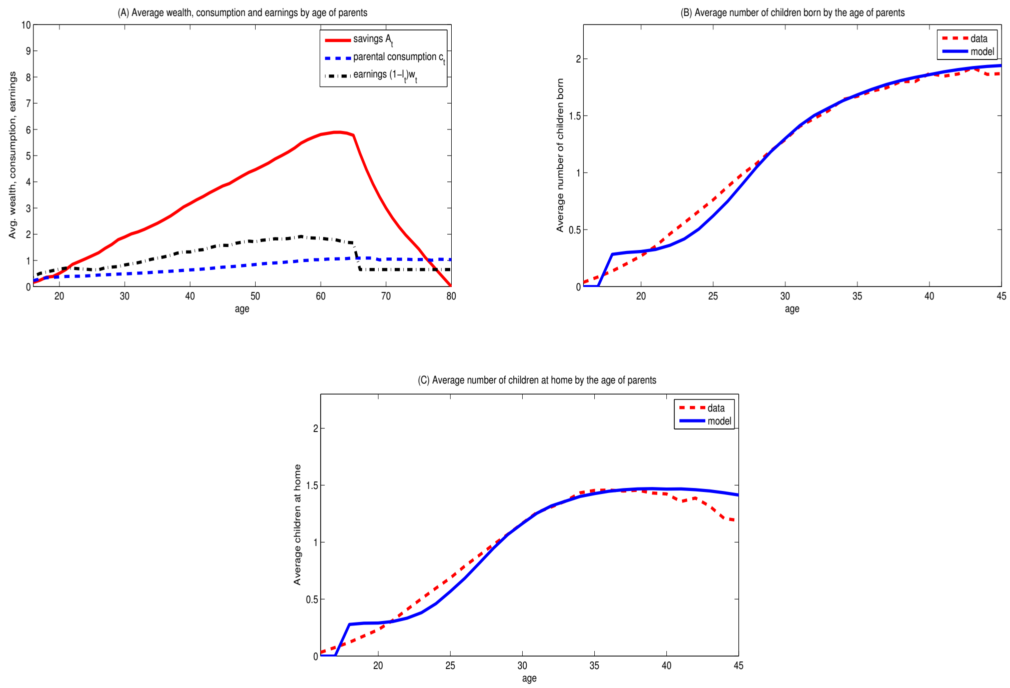

Figure 5: Benchmark profiles

Notes: Age-profiles of parental consumption, saving, earnings (Panel A). Age-specific cumulative average number of births (Panel B). Age-profile of average number of children at home (Panel C). Data source: NSLY79.

over the life cycle (as in here), young working families – who have a low opportunity cost of time relative to older workers – choose to invest time (rather than money) into children's production (Figure (7)).

To document the empirical variation in allocation of resources devoted to childrearing across income groups, Figure 6 captures the average expenditures on a child under age 6 by dual-earner families by household income group in the 2000 CEX data.[24] The CEX estimates confirm that household spending on the directly measured child-specific items increases with household income – a finding consistent with Lino (2008).[25] However, the exponential rise

ring of -0.23. (0.71), respectively.

[24] I focus on households who report positive spending on childcare services (such as baby-sitting or daycare), as reported zero spending on such services by dual-earners points to an alternative childcare arrangement (e.g., childcare is provided by an unpaid family member).

[25] CEX collects only limited data on expenditures directly attributable to children or childrearing. In particular, CEX collect information on children's clothing (for boys, girls and infants), toys, playground equipment, babysitting and daycare. The remainder of consumption expenditures is collected on household

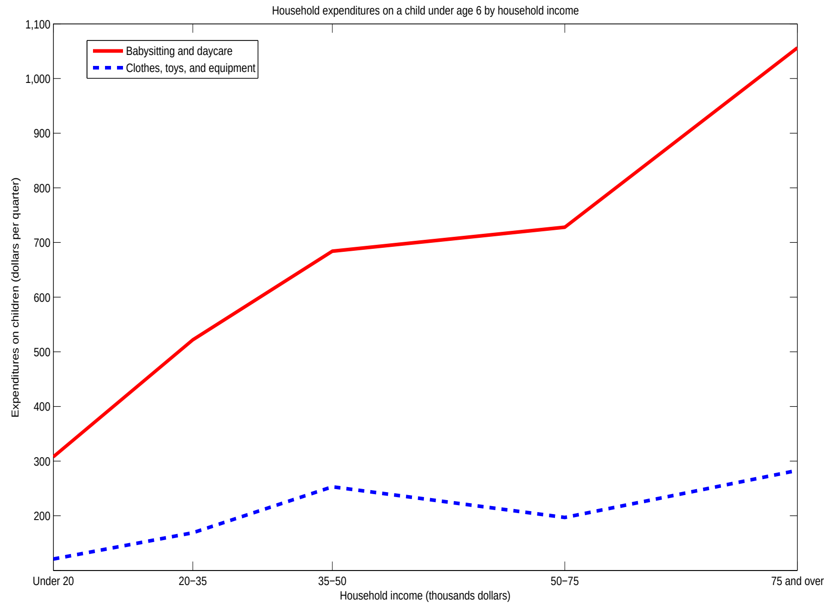

Figure 6: Expenditures on a child under age 6 by dual-earner families (CEX)

in spending on childcare services relative to low-income households indicates that high-income households indeed substitute some of the time devoted to childrearing with market expenditures.[26]

Relationship between Income and Fertility: A large number of studies has explored the relationship between income and fertility. Using the analogy with durable goods, Becker (1965) argues that the number of children and income should be positively correlated, but perhaps weakly so because childrearing is "a time-intensive activity that uses many hours which could be used to work" (page 510). While a full consensus on the relationship between income and fertility has not yet been reached (see, for example, Heckman and Walker (1990) for review), most studies agree the life cycle fertility is likely to respond to changes

basis. Using the CEX data, Lino (2008) uses per-capita method to impute the total expenditures on children by dual-earner families. The author finds that all components of total household expenditures on children rise with household income (see his Table 1).

[26]In particular, households earning less than $20,000 per year on average spend $120 per quarter on children's clothes, toys and equipment, while a household making over $75,000 dollars per year spends roughly 1.8 times more. The differences in household spending on babysitting and daycare by income class are, however, more pronounced. Households with yearly incomes of less than $20,000 spend roughly $300 per quarter on babysitting and daycare, while households earning over $75,000 per year spend 3.5-times more on such services.

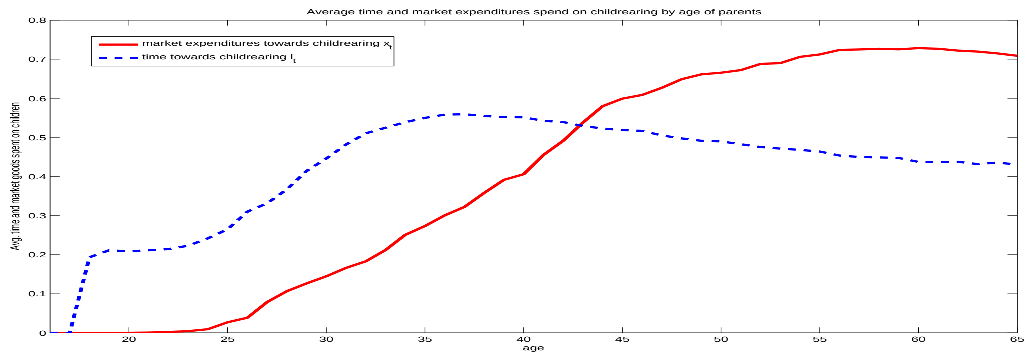

Figure 7: Time and money spent on children in the benchmark model

in household wages in non-linear fashion, in part because of the offsetting income and substitution effects.[27] In the model, the correlation between household market wages and the number of births is slightly negative (around -0.1) for households between ages 18 and 21 (Panel A in Figure 8). This means that low-wage households have, on average, their first child sooner than high-wage households who face a higher opportunity cost of childrearing than their low-wage peers. However, since children are a normal good, the income effect associated with higher wage dominates in a long-run, as the correlation coefficient flips in sign at age 22. Since young households on a whole prefer to use time (rather than money) to care for their offspring (as discussed above), the reduction in hours worked by households with children in turn lowers total household earnings, so that the correlation between births and labor income is lower than the correlation between births and wages per unit of labor (Panel B in Figure 8).

6 Relationship between earnings risk and fertility

In this section, I use the calibrated model to evaluate the impact on changes in uninsurable labor market risk on household fertility between steady states. Figure 9 summarizes the predicted relationship between the standard deviation of persistent earnings shock, σ_ϵ,

[27]When an important cost of children is parents' own time, higher wage families face higher price of children relative in terms of other consumption, but as long as children are a normal good, the income effect associated with the higher wage implies that the demand for children should increase. Which effect dominates then determines the sign of the total effect of the wage change on the demand for children.

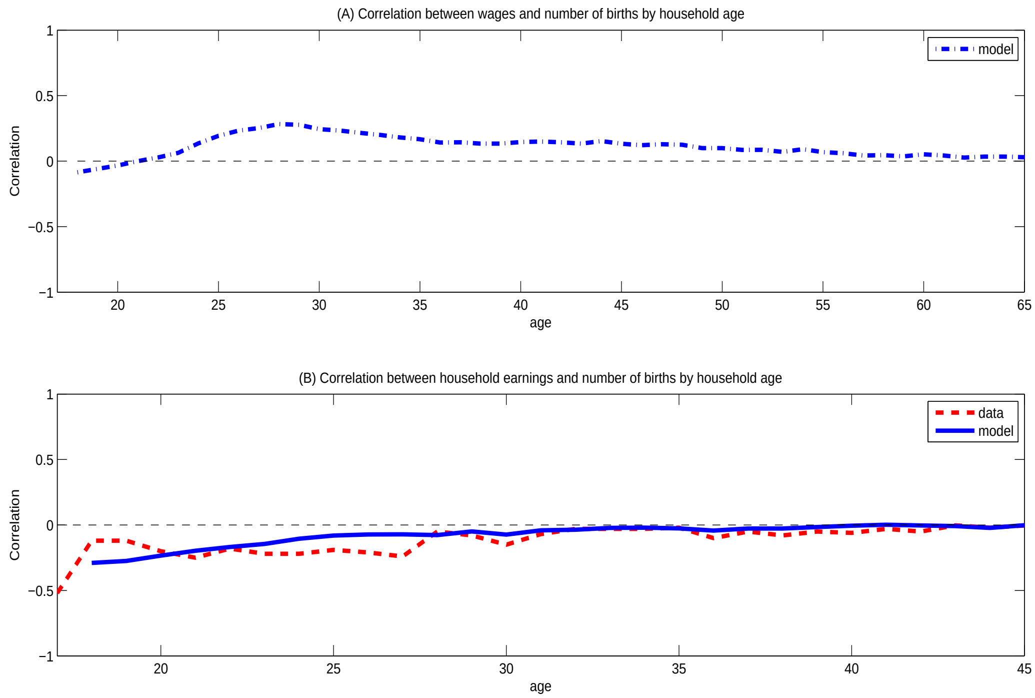

Figure 8: Relationship between fertility, earnings, and wages

Notes: Age-specific correlation between household earnings and births for husband-wife couples (Panel A). Age-specific between market wage and births (Panel B). Data source: NLSY79.

household fertility, and the timing of first and second births. The persistent component of household earnings is commonly assumed to represent uninsurable idiosyncratic earnings risk. As can be seen from the figure, increased earnings risk leads to a postponement of the first birth, increased gap between births, and reduced number of births per household. As long as the standard deviation of earnings remains relatively low (e.g., σ between 0.01 and 0.12), the predicted changes in household fertility are small: the number of births per family declines from roughly 2.6 to 2.4, and the mean age at first birth rises by roughly two years (from 20 to 22 years). At the same time, the average time-gap between the first and second births remains constant.

The changes in household fertility behavior are more pronounced for medium to high levels of earnings risk. Increasing the earnings risk from 0.12 to 0.21 would reduce the number of births per household from 2.4 to 1.9. This decrease in the number of births is significant not only because of its magnitude (the number of births would fall by over 20

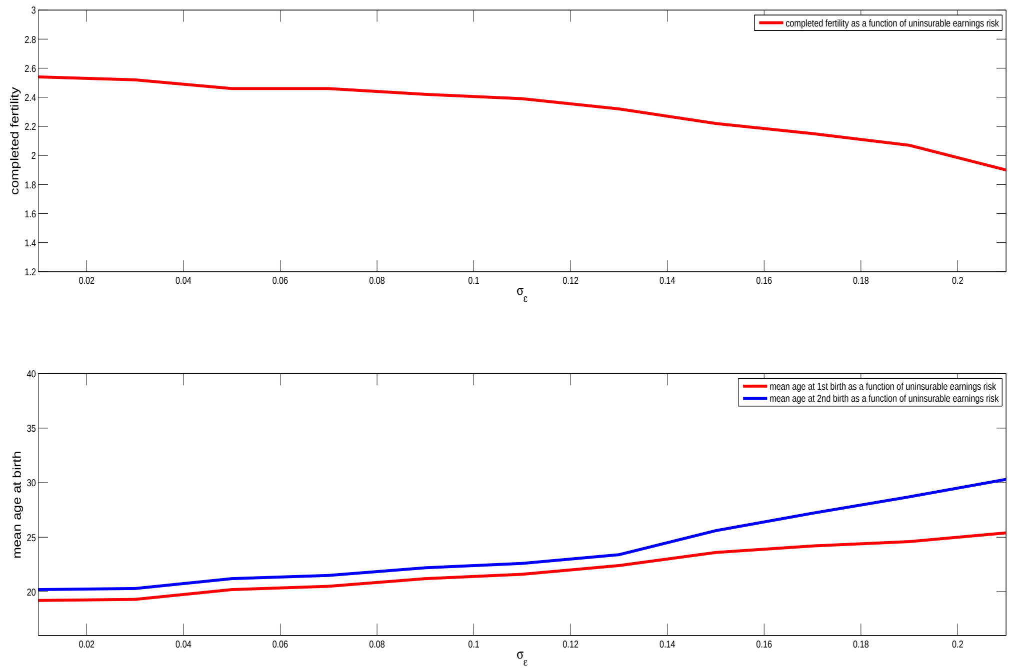

Figure 9: Impact of changes in persistent earnings risk on completed fertility and timing of first and second births

percent), but also because completed fertility would drop below the replacement rate of 2.1 (i.e., the birth level needed to prevent a demographic decline).

Moreover, when the earnings risk increases from 0.12 to 0.21, the sizable decline in the number of births is accompanied by increases in the age at first and second births. The age at first birth rises by additional 3.2 years, while the gap between the births widens dramatically, as households significantly postpone the birth of their second child. The age at second birth jumps from 23 to 30 years, and the number of households with higher-order births falls. Interestingly, the predicted postponement of births in response to increased earnings uncertainty is consistent with empirical findings in Amialchuk (2008) who – using panel data constructed from the 1968-1993 waves of the PSID – documents a negative effect of a persistent shock to husband's earnings on the fertility patterns of married couples. In particular, Amialchuk (2008) finds that women with husbands displaced from a job (due to, for example, a factory closing or a lay-off) postpone births, with the postponement being

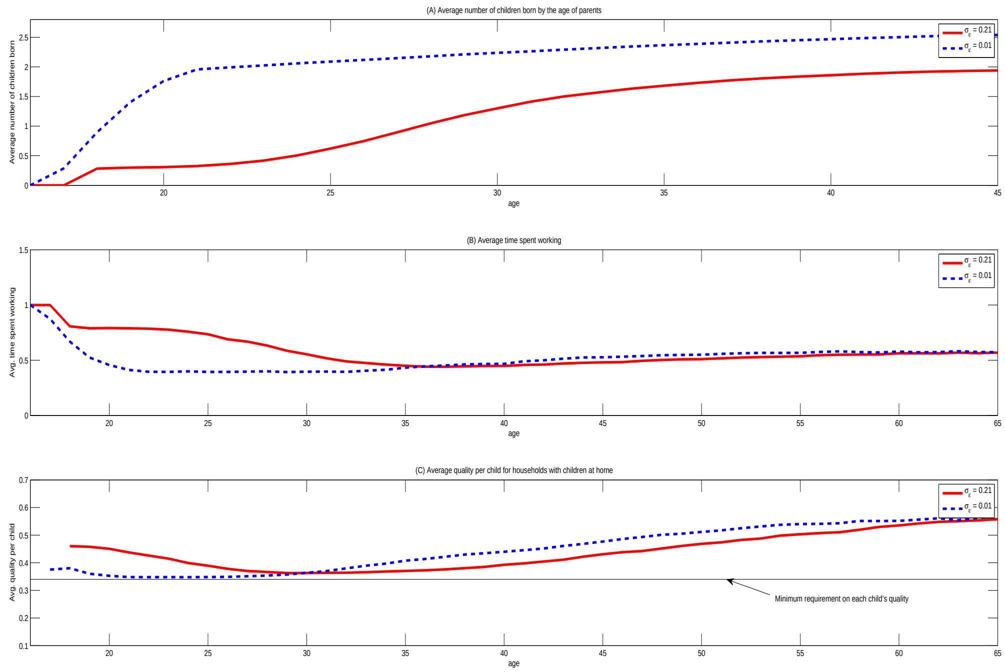

Figure 10: Age-specific birth rates (Panel A), time spent working (Panel B), average quality per child for households with children (Panel C) under low and medium levels of earnings risk

particularly robust for the second and subsequent births.

To put the simulated results of the steady-state model into context of the U.S. time-series data, the mean age at first birth increased by 3.5 years between 1970 and 2000, with the steepest increase from 1975 to 1985. At the same time, women who made childrearing decisions in the 1960s and 1970s had on average 2.5 children compared to 1.9 births for women who made fertility choices in the 1980s and 1990s.[28] The changes in U.S. fertility trends coincided with large shifts in microeconomic labor market risk (see Gottschalk (1997), Levy and Murnane (1992), or Heathcote and Violante (2004) for review). A large share of

[28] Rather than using the TFR which provides a misleading estimate of the actual rate of childbearing when the age at childbearing is changing (Bongaarts (1999)), I use estimates of completed cohort fertility. The average number of births for women who made childbearing choices in the 1960s and 1970s comes from Jones and Tertilt (2006) who use the Decennial Census data between 1900 and 1990 to construct the average number of births by women's birth-cohorts. The 2000 wave of the Decennial Census no longer collects information on the number of children ever born. NLSY79 is thus used to compute births of women who made their fertility choices in the later years.

the increased risk has been attributed to the persistent component of household earnings which has risen from roughly 0.15 in the 1970s to 0.21 in the 1990s (Meghir and Pistaferri, 2004).[29] The model predicts that an increase in persistent risk of this magnitude would lead the fertility rate to decline from 2.2 to 1.9, while the mean age at first and second births would increase by roughly 2 and 5 years, respectively. Viewed in an isolation, a permanent increase in persistent risk that is in line with the U.S. experience could therefore explain about one half of the fertility decline and of the increase in mean age at first birth in the United States between the 1970s and 1990s, while matching broadly the changes in the timing of the second birth.

That said, the estimated effect should be viewed in a broader context of other changes likely driving U.S. fertility. First, increases in household productivity likely offset some of the negative effect on increased labor market risk on household fertility in the long-run U.S. data. (In contrast, depressed income levels of young workers in other countries or during economic downturns likely bolster the negative effect of risk). For example, an increase in risk from 0.15 to 0.21 combined with a 10 percent increase in real wages would reduce fertility from 2.1 to 1.9, while increasing the mean age at first and second births by 1.2 years and 3 years, respectively. Second, changes in the U.S. fertility coincided with significant shifts in educational attainment of women, and changes in marriage and mating. Unfortunately, due to computational constraints, it is currently impossible to incorporate these channels into the model given its already large state space. However, it would be interesting to see whether changes in household risk could link some of these trends. For example, if women choose to delay fertility in response to labor market risk, they have more time available for education and work. Alternatively, if education could be used as a hedge against earnings risk, higher educational attainment and delayed fertility could be tightly linked together insofar as women postpone fertility in order to minimize lifetime earnings uncertainty through increased education. In this sense, increased levels of education could attenuate the negative effect of earnings risk on household family sizes, but at the same time could bolster the delay in timing of births.

[29] Using the PSID data, Meghir and Pistaferri (2004) provide yearly estimates of the standard deviation of the persistent shock to labor earnings for the period 1969-1991, and find that the standard deviation of the persistent shock was rising throughout the 1970s and in the early 1980s.

Turning to the mechanism generating delayed childbearing and lower fertility response to higher risk, it is helpful to recall that in the model children are a durable good of an irreversible nature. Moreover, children are costly: childrearing requires at least a minimum amount of time and money invested into each offspring, so that the average quality per child is always above $\underline{q}$. Thus, childrearing leads to a higher utility but also limits a household's ability to insure against adverse wage shocks through increased saving or labor supply. When labor market risk is high and adverse spells are persistent, parents initially choose to postpone childbearing, and work and save more instead. The delay in childbearing is particularly pronounced for higher-order births, because the amount of resources required for childrearing increases – albeit at a decreasing rate (due to the economies of scale) – with number of children at home. While parents may initially consider their decision to delay childrearing as temporary, the infertility risk tends to reduce the total number of births and the number of households with no or only one child rises. Section 7 discusses the effect of infertility risk on completed fertility.

To illustrate this mechanism, Panel A in Figure 10 shows the age-specific birth rates under different levels of risk. When persistent risk is low (i.e., $\sigma_\epsilon = 0.01$), the average household has 2 children by the age of 25. However, when risk rises to the U.S. level (i.e., $\sigma_\epsilon = 0.21$), the birth activity of households is more spread out, and an average household has only one child by the age of 30. Turning to labor supply, Panel B demonstrates that increases in labor market risk have a large effect on the labor supply of young households. In particular, households between ages 16 and 30 work 30 percent less when risk is low than in an environment with U.S. earnings volatility, as households have children sooner and spend time at home caring for them. Surprisingly, the labor supply of households over 35 is relatively inelastic with regard to the increases in the persistent component of labor market risk, in part because middle-aged households substitute time with market goods in childrearing (see the discussion in Section 5). Moreover, increased risk leads to higher aggregate saving in the economy. The average predicted assets-to-wage (A_t/w_t) ratio in a low risk environment ($\sigma_\epsilon = 0.01$) is roughly 25 percent lower than in the economy with U.S. level of risk ($\sigma_\epsilon = 0.21$): 1.8 versus 2.4.[30]

[30]Broadly speaking, related mechanisms have been explored in other studies. First, thinking about the

The increased risk has also an effect on the average quality per child in the economy. In a low risk economy, young parents – who have over two children by the age of 25 – reduce the quality of each child in exchange for larger family sizes. As a result, the minimum quality requirement on average binds for parents in their twenties in a low risk economy, who in general have 2+ children. Over time, as parents' wage profile increases (due to the life cycle effects) and some children leave home, the average quality per child rises again. In contrast, in an economy with U.S. level of risk, the minimum quality requirement binds mostly for households in their late twenties and early thirties, when an average household has the first and second child. In this case, children's quality suffers not because of large family sizes, but because parents work more and save more to insure against adverse economic environment.

7 Effect of Infertility Risk

Figure 11 compares predictions of the baseline model against a model with no fertility risk. When fertility risk is absent, the total number of births per family is higher for all values of earnings risk relative to the baseline model, but – perhaps surprisingly – so are the mean ages at first and second birth. For example, an average household gives the first birth at age 29 in the model with no fertility risk at the baseline level of earnings risk ($\sigma_\epsilon = 0.21$) – nearly four years later than in the benchmark model with fertility risk. This pronounced delay in childbearing has, however, no adverse effect on completed fertility; an average household has 2.6 children in a model with no fertility risk, compared with 1.9 in the baseline model. In

link between labor market risk and household savings, many studies simulated the role of income uncertainty in individual savings decisions, and found that precautionary saving can explain a sizable portion of U.S. household savings (see, for example, Skinner (1985), Zeldes (1989), Caballero (1991), Deaton (1991), Carroll (1992, 1991), Carroll and Samwick (1997, 1998), Hubbard, Skinner, and Zeldes (1994), Aiyagari (1994)). Attanasio, Low, and Sanchez-Marcos (2005) and Ortigueira and Siassi (2010) explore the role of labor supply as an insurance against idiosyncratic earnings risk within the family. The studies find that intra-household risk sharing (through labor supply adjustment) has the largest impact among households with limited financial wealth. The authors find that households with sizable financial wealth use mainly savings to smooth consumption across unemployment spells, while poorer households rely on spousal labor supply during adverse income spells. While intra-household insurance is beyond the scope of this paper, the results in the listed studies support this study's predictions that young households (with limited assets) are likely to respond to increased labor market risk by increasing labor supply, while older households respond primarily through increased saving. The general alignment of the model predictions regarding the response of household savings and labor supply to increases in idiosyncratic uninsurable earnings risk with previous studies is reassuring.

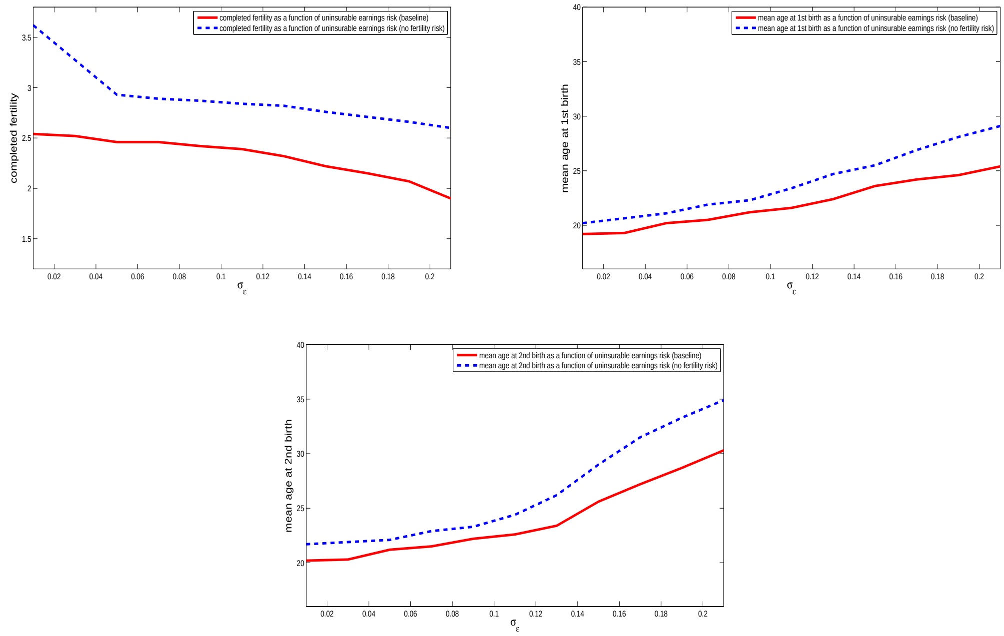

Figure 11: Impact of infertility risk on the response of household fertility patterns to increases in uninsurable earnings risk

general, having perfect control over their fertile window appears to allow households to more optimally time births without running a risk of ending up with suboptimal family sizes.

The ability of households to significantly delay childbearing without compromising their total family sizes has also implications for the response of household fertility patterns to increased labor market risk. In the model with no fertility risk, households optimally respond to increases in earnings risk by further delays in childbearing but catch up on the desired number of births in their late thirties and early forties. In particular, when labor market risk increases from 0.15 to 0.21 in the model with no fertility risk, the mean age at first and second birth increases by 3.6 and 5.9 years to 29 and 35 years of age, respectively, while the total number of births per household falls ever so slightly from 2.8 to 2.6.

The model predictions with regard to infertility risk are particularly interesting when viewed in the context of the ongoing advances in the reproductive science. In particular, the model predicts that wider access to reproductive technologies can help to offset some of the adverse effect of increased earnings risk (and – broadly speaking – other factors contributing to delayed childbearing) on completed fertility.

8 Conclusions

This paper studied the relationship between household fertility choices and microeconomic uninsurable earnings uncertainty using a life cycle model of fertility in which unitary households face idiosyncratic uninsurable labor market and infertility risks, and make joint decisions about consumption, savings, the family size, and the allocation of parental time and market goods invested toward improving their children's quality. Children were modeled as durable goods of irreversible nature that are born sequentially, and require at least a minimum amount of parental investment per child. The analysis was motivated by the fact that although the link between labor market risk and household fertility has been suggested in both time-series and cross-sectional data, it has not been rigorously explored in a quantitative micro-founded framework. To this end, the main contribution of this paper lies in offering a quantitative theoretical exploration of the link between uninsurable idiosyncratic earnings risk and fertility patterns.

Using the carefully calibrated model and comparing model predictions across steady states with varying levels of uninsurable idiosyncratic earnings risk, I showed that higher earnings uncertainty is associated with lowered fertility rate and increased mean ages at births. This is because when risk is high, households postpone childbearing, initially preferring to work more and to accumulate more precautionary savings. The birth postponement in turn interacts with the risk of infertility which rises exponentially with the age of the mother, reducing the total number of births per family. The documented linkage between earnings risk and fertility patterns highlights the important role that labor market conditions can play in determining both short-term cyclical fluctuations in fertility (such as those in the recent U.S. data) and longer-term demographic trends. In my view, the findings presented in this paper that point to the prominent role that labor market conditions can play in driving changes in population growth are of first order importance when viewed in the context of the population aging in many European countries where labor market risk – primarily for young households – is high and persistent and fertility rates are very low.

That said, the results presented in this paper are hopeful. First, while several recent studies document sizable increases in earnings uncertainty in a cross-section of advanced economies, they also point to the role of labor market institutions, and tax and transfer systems which can offset the rise in household labor market risk (see, for example, Domeij and Floden (2009) or Jappelli and Pistaferri (2009)). While the current analysis was focused on studying fertility in a U.S. style economy, I see studying of the differences in fertility across countries with varying levels of earnings uncertainty and institutional features as an important avenue for future work. Moreover, the presented results point to the key role of infertility risk in determining household fertility patterns. Broadly speaking, the results presented in Section 7 suggest that wider access to the reproductive technology can help to offset the negative effect of earnings risk on completed fertility.

Earnings risk is certainly not the only determinant of fertility choices. Education, career, changes in marriage and mating, and increasing contraceptive use have been reported as important factors affecting household fertility choices. Unfortunately, due to computational constraints, it is currently impossible to incorporate these channels into the model in order to study an interaction of all these factors in a unified framework. However, given possible inter-linkages and feedback effects amongst these channels, I view this as an important avenue for future work.

References

Adsera, A. (2004): "Changing Fertility Rates in Developed Countries: The Impact of Labor Market Institutions," *Journal of Population Economics*, 17(1), 17–43.

——— (2005): "Vanishing Children: From High Unemployment to Low Fertility in Developed Countries," *American Economic Review*, 95(2), 189–193.

——— (2006): "An Economic Analysis of the Gap Between Desired and Actual Fertility: The Case of Spain," *Review of Economics of the Household*, 4(1), 75–95.

Aiyagari, S. R. (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, 109(3), 659–84.

Amialchuk, A. (2008): ""The Effect of Husband's Job Displacement on Fertility"," *mimeo*.

Attanasio, O., H. Low, and V. Sanchez-Marcos (2005): "Female Labor Supply As Insurance Against Idiosyncratic Risk," *Journal of the European Economic Association*, 3(2-3), 755–764.

Becker, G. S. (1960): "An Economic Analysis of Fertility," in *Demographic and Economic Change in Developed Countries*, ed. by G. S. Becker. Princeton, NJ: Princeton University Press.

——— (1965): "A Theory of the Allocation of Time," *The Economic Journal*, 75(299), 493–517.

Becker, G. S., K. M. Murphy, and R. Tamura (1990): "Human Capital, Fertility, and Economic Growth," *Journal of Political Economy*, 98(5), S12–37.

Becker, G. S., and N. Tomes (1976): "Child Endowments and the Quantity and Quality of Children," *Journal of Political Economy*, 84(4), S143–162.

Bhaumik, S. K., and J. B. Nugent (2006): "Does Economic Uncertainty Affect the Decision to Bear Children? Evidence from East and West Germany," CEDI Discussion Paper Series 06-06, Centre for Economic Development and Institutions(CEDI), Brunel University.

Boldrin, M., M. De Nardi, L. Jones, and C. Madrid (2005): "Fertility and Social Security," *NBER working paper.*

Bongaarts, J. (1999): "Fertility Decline in the Developed World: Where Will It End?," *American Economic Review*, 89(2), 256–260.

Caballero, R. J. (1991): "Earnings Uncertainty and Aggregate Wealth Accumulation," *American Economic Review*, 81(4), 859–71.

Card, D. (1991): "Intertemporal Labor Supply: An Assessment," NBER Working Paper 3602, National Bureau of Economic Research.

Carroll, C. D. (1991): "Buffer Stock Saving and the Permanent Income Hypothesis," Discussion paper.

——— (1992): "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 23(1992-2), 61–156.

Carroll, C. D., and A. A. Samwick (1997): "The Nature of Precautionary Wealth," *Journal of Monetary Economics*, 40(1), 41–71.

——— (1998): "How Important Is Precautionary Saving?," *The Review of Economics and Statistics*, 80(3), 410–419.

Caucutt, E., N. Guner, and J. Knowles (2002): "Why Do Women Wait? Matching, Wage Inequality, and the Incentives for Fertility Delay," *Review of Economic Dynamics*, 5(4), 815–855.

Chetty, R., and A. Szeidl (2007): "Consumption Commitments and Risk Preferences*," *The Quarterly Journal of Economics*, 122(2), 831–877.

Conesa, J. C. (2000): "Educational Attainment and Timing of Fertility Decisions," *Mimeo.*

Da Rocha, J. M., and L. Fuster (2006): "Why Are Fertility Rates And Female Employment Ratios Positively Correlated Across O.E.C.D. Countries?," *International Economic Review*, 47(4), 1187–1222.

de la Croix, D., and M. Doepke (2003): “Inequality and Growth: Why Differential Fertility Matters,” *American Economic Review*, 93(4), 1091–1113.

Deaton, A. (1991): “Saving and Liquidity Constraints,” *Econometrica*, 59(5), 1221–48.

Doepke, M. (2004): “Accounting for fertility decline during the transition to growth,” *Journal of Economic Growth*, 9(3), 347–383.

Doepke, M., M. Hazan, and Y. Maoz (2007): “The Baby Boom and World War II: A Macroeconomic Analysis,” NBER Working Paper 13707, National Bureau of Economic Research.

Domeij, D., and M. Floden (2009): “Inequality Trends in Sweden 1978-2004,” Discussion paper, Mimeo, Stockholm School of Economics.

Fernandez, R., N. Guner, and J. Knowles (2005): “Love and Money: A Theoretical and Empirical Analysis of Household Sorting and Inequality,” *The Quarterly Journal of Economics*, 120(1), 273–344.

Fisher, J., and M. Gervais (2011): “Why Has Home Ownership Fallen Among the Young?,” *International Economic Review*, 52(3), 883–912.

Goldin, C., and L. Katz (2002): “The Power of the Pill: Oral Contraceptives and Women’s Career and Marriage Decisions,” *Journal of Political Economy*, 110(4), 730–770.

Goldin, C., and R. Margo (1992): “The Great Compression: The Wage Structure in the United States at Mid-Century,” *The Quarterly Journal of Economics*, 107(1), 1–34.

Gottschalk, P. (1997): “Inequality, Income Growth, and Mobility: The Basic Facts,” *Journal of Economic Perspectives*, 11(2), 21–40.

Greenwood, J., and A. Seshadri (2002a): “Technological Progress and Economic Transformation,” Economie d’avant garde research reports, Economie d’Avant Garde.

Greenwood, J., and A. Seshadri (2002b): “The US demographic transition,” *American Economic Review*, pp. 153–159.

Greenwood, J., A. Seshadri, and G. Vandenbroucke (2005): "The baby boom and baby bust," *American Economic Review*, 95(1), 183–207.

Heathcote, J., and G. Violante (2004): "The Macroeconomic Implications of Rising Wage Inequality in the United States," Discussion paper.

Heckman, J., and J. Walker (1990): "The Relationship Between Wages and Income and the Timing and Spacing of Births: Evidence from Swedish Longitudinal Data," *Econometrica*, pp. 1411–1441.

Hotz, V. J., and R. A. Miller (1988): "An Empirical Analysis of Life Cycle Fertility and Female Labor Supply," *Econometrica*, 56(1), 91–118.

Hubbard, R. G., J. Skinner, and S. P. Zeldes (1994): "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving," *Carnegie-Rochester Conference Series on Public Policy*, 40(1), 59–125.

——— (1995): "Precautionary Saving and Social Insurance," *Journal of Political Economy*, 103(2), 360–399.

Jappelli, T., and L. Pistaferri (2009): "Does Consumption Inequality Track Income Inequality in Italy?," Csef working papers, Centre for Studies in Economics and Finance (CSEF), University of Naples, Italy.

Jones, L. E., A. Schoonbroodt, and M. Tertilt (2008): "Fertility Theories: Can They Explain the Negative Fertility-Income Relationship?," Working Paper 14266, National Bureau of Economic Research.

Jones, L. E., and M. Tertilt (2006): "An Economic History of Fertility in the U.S.: 1826-1960," Discussion paper.

Knowles, J. (2007): "High-powered Jobs: Can Contraception Technology Explain Trends in Womens Occupational Choice," Discussion paper, Mimeo, University of Pennsylvania.

Kopczuk, W., E. Saez, and J. Song (2009): "Earnings Inequality and Mobility in the United States: Evidence from Social Security Data since 1937," *forthcoming Quarterly Journal of Economics*.

Levy, F., and R. J. Murnane (1992): "U.S. Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations," *Journal of Economic Literature*, 30(3), 1333–81.

Lino, M. (2008): "Expenditures on Children by Families: 2007 Annual Report," Discussion paper, U.S. Department of Agriculture, Center for Nutrition Policy and Promotion.

Meghir, C., and L. Pistaferri (2004): "Income Variance Dynamics and Heterogeneity," *Econometrica*, 72(1), 1–32.

Mira, P., and N. Ahn (2002): "A Note on the Changing Relationship Between Fertility and Female Employment in Developed Countries," *Journal of Population Economics*, 15, 667–682.

Munnell, A. H., and M. Soto (2006): "What Replacement Rates Do Households Actually Experience In Retirement?," Working paper, Center for Retirement Research, Boston College.

Ortigueira, S., and N. Siassi (2010): "How Important Is Intra-household Risk Sharing for Savings and Labor Supply?," *EUI Working Papers*, (ECO 2010/36).

Postlewaite, A., L. Samuelson, and D. Silverman (2008): "Consumption Commitments and Employment Contracts," *Review of Economic Studies*, 75(2), 559–578.

Ranjan, P. (1999): "Fertility Behaviour under Income Uncertainty," *European Journal of Population*, 15(1), 25–43.

Saks, R., and S. Shore (2005): "Risk and Career Choice," *Advances in Economic Analysis & Policy*, 5(1), 1414–1414.

Skinner, J. (1985): "Variable Lifespan and the Intertemporal Elasticity of Consumption," *The Review of Economics and Statistics*, 67(4), 616–23.

Smith, T. (2008): "Indirect Inference," in *The New Palgrave Dictionary of Economics*, ed. by S. Durlauf, and L. Blume. London: Palgrave Macmillan.

Storesletten, K., C. I. Telmer, and A. Yaron (1998): "The Risk Sharing Implications of Alternative Social Security Arrangements," GSIA Working Papers 252, Carnegie Mellon University, Tepper School of Business.

——— (2004): "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy*, 112(3), 695–717.

Tauchen, G. (1986): "Finite State Markov-chain Approximations to Univariate and Vector Autoregressions," *Economics Letters*, 20(2), 177–181.

Trussell, J., and C. Wilson (1985): "Sterility in a Population with Natural Fertility," *Population Studies*, 32(2), 269–286.

Wallace, W. H. B., and T. W. Kelsey (2010): "Human Ovarian Reserve from Conception to the Menopause," *PLoS ONE*, 5(1), http://www.plosone.org/article/info%3Adoi%2F10.1371%2Fjournal.pone.0008772pone-0008772-g005.

Zeldes, S. P. (1989): "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *The Quarterly Journal of Economics*, 104(2), 275–98.

A Appendix

A.1 Data sources

A.1.1 NLSY79

The National Longitudinal Survey of Young 1979 (NLSY79) is a nationally representative sample of 12,686 young men and women ages 14 to 22 when they were first interviewed in 1979. These individuals were interviewed annually through 1994 and are currently interviewed on a biennial basis. This study uses the data from 1982 through 2004 to construct fertility moments used either to estimate the model, or used elsewhere in the paper. The information on the number of children ever born, the age at first birth, the number of children at home by respondents, and the wages and salaries for both the respondent and his/her spouse were used to compute the statistics reported in the paper. To compute the average number of births and the mean age at first birth, the 2004 data were used. To construct the life cycle profiles in Figures 5 and 8, the panel dimension of the data set was employed. To construct the age-profile of correlation between household earnings and births, I define the total household earnings as the sum of the nonnegative reported salaries and wages for a respondent and his/her spouse or unmarried partner. If a respondent does not have a spouse or a partner, the labor income of the spouse is replaced with zero. All earnings variables are deflated by the CPI-U. Sampling weights were employed to create a nationally representative sample of households, and respondents with the missing data for variables used to construct the data moments were dropped from the analysis.

A.1.2 CEX

The Consumer Expenditure Survey (CEX) is a quarterly survey of household expenditures. Household expenditures are aggregated mostly at a household level, with only few expenses being directly attributable to children (i.e., expenses on children's clothing, toys, equipment, daycare, and babysitting are directly recorded for children). To estimate the OLS regression in equation (18), only the observations of child-specific expenditure components for the last quarter for husband and wife consumer units with own children younger than 18 years of age are used is used to run the auxiliary regression for indirect inference, specified in equation (18), as well as to construct expenditure profiles in Figure 6.

A.1.3 ATUS

The American Time Use Survey (ATUS) is a survey on time use in the United States. ATUS respondents (i.e., one household member age 15 or over) are interviewed only one time about how they spent their time on the previous day, where they were, and whom they were with. Questions related to caring for or helping household children are asked in the survey. The survey defines the time spent on caring or helping household children as "time spent doing activities to care for or help any child (under age 18), regardless of relationship to the survey respondent or the physical or mental health status of the person being helped." In this study, I use the 2004 data on "primary childcare activities." These activities include time spent providing physical care; playing with children; reading with children; assistance with homework; attending children's events; taking care of children's healthcare needs; and dropping off, picking up, and waiting for children. Passive childcare done as a primary activity (such as "keeping an eye on my son while he swam in the pool") also is included.[31] Only respondents with (i) a spouse or unmarried partner and (ii) own children in the household have been included in the analysis. These reported data on minutes spent on primary childcare activities in the previous day have been used to estimate the auxiliary regression in equation (19), and to construct the profiles of time spent with children used elsewhere in the paper.

[31] In the survey, a child's presence during the activity is not enough in itself to classify the activity as childcare. For example, "watching television with my child" is coded as a leisure activity, not childcare.

A.1.4 Census

I use the 5 percent sample of the Decennial Census from 1990, available publicly at the IPUMS-USA website. I concentrate on married couples. Matching respondents with their spouses yields a total of about 2.5 million husband-wife observations. Since educational attainment is known to affect the timing of births, I consider only households in which the husband attained a bachelor's degree. Similarly, since fertility patterns are known to vary between urban and rural areas, I consider only households which resided at urban, non-farm regions at the time of the interview. These sub-sampling criteria decrease the sample's heterogeneity without restricting the sample size excessively. In order to explore the effect of the husband's occupation on household fertility, I use the Census 1950 three-digit occupational codes to assign husbands to one of the seven occupational groups from Saks and Shore (2005). I assign husbands the following occupations: teachers, heath-care professionals (e.g., medical and dental technicians, nurses, optometrists, and pharmacists), engineers, managers, math and sciences (e.g., mathematicians, physicists, and other natural scientists), sales workers, and arts and entertainments (e.g., artists and arts teachers, actors, dancers and dance teachers, musicians and music teachers). Husbands whose occupation lies outside those studied in Saks and Shore (2005) have been omitted from the analysis. Since self-employed individuals have been shown to face higher earnings risk than individuals working for wage or salary across occupations (see, for example, Carroll and Samwick (1997) or Saks and Shore (2005)), I focus on individuals who are not self-employed.

A.2 Further Evidence on Infertility Risk

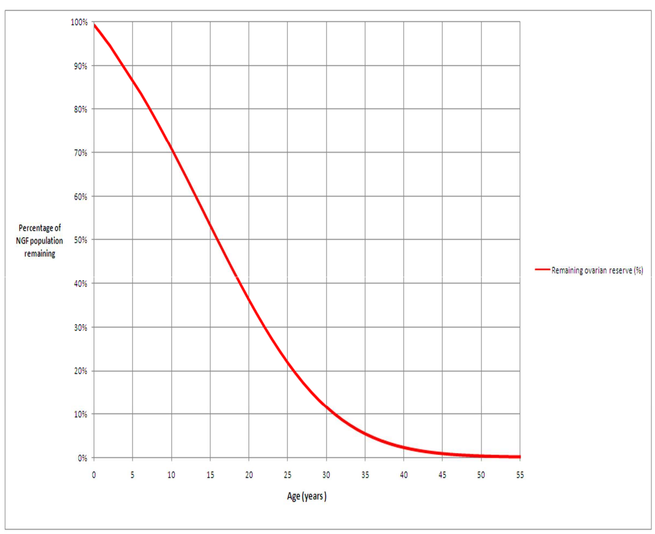

Figure 12: Ovarian reserve (NGF) by the age of a women (Wallace and Kelsey (2010))

A.3 Numerical solution and algorithm

The household problem is solved numerically by backward recursion from the terminal period. At each state, I solve for the value function and the optimal policy rules, given the current state variables and the solution to the value function in the next period. In the model, households face sterility shocks which render them permanently infertile. A household with fertile members and a state vector $(a_t, n_t, \epsilon_t, \nu_t, f_t = F, t)$ solves the problem

$$V(a_t, n_t, \epsilon_t, \nu_t, F, t) = \max \begin{cases} \max_{a_{t+1}, x_t, l_t} u(c_t, n_t, q_t) + \beta E_t V(a_{t+1}, n_{t+1}, \epsilon_{t+1}, \nu_{t+1}, f_{t+1}, t+1) \\ \max_{a_{t+1}, x_t, l_t} u(c_t, n_t, q_t) + \beta E_t V(a_{t+1}, n_{t+1}, \epsilon_{t+1}, \nu_{t+1}, f_{t+1}, t+1) \end{cases}$$

subject to the constraints and transition equations specified in equations (10) to (14), with $f_{t+1} = \{I, F\}$.

A household with infertile members and state vectors $(a_t, n_t, \epsilon_t, \nu_t, f_t = I, t)$ solves the recursive problem of the form

$$V(a_t, n_t, \epsilon_t, \nu_t, I, t) = \max_{a_{t+1}, x_t, l_t} u(c_t, n_t, q_t) + \beta E_t V(a_{t+1}, n_{t+1}, \epsilon_t, \nu_t, I, t+1),$$

subject to the constraints and transition equations in equations (10), (11), (14), and (15).

The complications for the solution of the household problem arise from the presence of a discrete choice. The discrete fertility choice implies that the value function will not necessarily be concave or differentiable at any stage of the life cycle. Therefore, I employ finite dynamic programming methods and only approximate the solution to the household problem by solving the household problem on a grid.

The algorithm used to solve the household problem is as follows. First, I guess the values for the parameters to be estimated. Given the guesses (and given the remaining parameters summarized in Table 2), I use finite dynamic programming to solve for optimal decision rules for savings $a(a_t, n_t, \epsilon_t, \nu_t, f_t, t)$, number of children at home $n(a_t, n_t, \epsilon_t, \nu_t, f_t, t)$, and time $l(a_t, n_t, \epsilon_t, \nu_t, f_t, t)$ and market expenditures $x(a_t, n_t, \epsilon_t, \nu_t, f_t, t)$ devoted to childrearing. Next, I simulate the shock histories for 10,000 households. Using the simulated histories and the optimal decision rules, I compute the targeted moments for the model economy. Last, I use the method of simulated moments to pin down the values for estimated parameters which produce moments summarized in Table (3). Since the differentiability of the objective function in the estimated parameters is not guaranteed, I use a minimization procedure that does not rely on the existence of the gradient (simplex). Once the estimated parameters are identified, I resolve the household problem and save the optimal decision rules.