Journal of Monetary Economics 83 (2016) 27–38

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Fertility choice in a life cycle model with idiosyncratic uninsurable earnings risk

Kamila Sommer

Federal Reserve Board, 1801 K Street NW, Washington, DC 20006, United States

ARTICLE INFO

Article history:
Received 18 April 2013
Received in revised form
18 August 2016
Accepted 18 August 2016
Available online 27 August 2016

Keywords:
Fertility choice
Life cycle
Uninsurable idiosyncratic income risk

ABSTRACT

Motivated by large shifts in uninsurable earnings risk over time, this paper studies the link between delaying and reducing fertility on the one hand, and earnings and fertility risks on the other. When children are modeled as consumption commitments, increases in earnings risk are associated with a reduction in family sizes and patterns of delayed childbearing. Since household ability to bear children declines with age, the postponement of birth associated with the increased earnings risk drives down the number of birth per family further. An access to in vitro fertilization (IVF) is shown to have only a limited offsetting effect.

Published by Elsevier B.V.

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. Motivated in part by the negative empirical correlation between fertility and income, many economic studies linked changes in income levels to changes in fertility, but generally abstracted from modeling how income risk (among other factors) might affect it as well. (For an extensive review of the literature, see Jones et al., 2008.) Against this backdrop, medical literature—while agnostic about the economic mechanisms behind the observed changes in fertility patterns—suggests the interaction of the delay in childbearing with age-dependent infertility risk as one reason why fertility might be falling. In particular, while age-specific infertility rates have not changed substantially since the 1970s, the number of women with reported fertility problems has risen appreciably, largely because women are attempting to have children at older, less fecund ages (Chandra and Stephen, 1998). Working in the opposite direction, the introduction of new infertility treatment options—in vitro fertilization (IVF) in particular—has likely mitigated some of these effects. This paper builds on the existing economic and medical literatures by studying how income risk interacts with infertility risk in affecting household fertility.

The analysis is motivated by large shifts in uninsurable earnings risk since the 1970s (see, for example, Meghir and Pistaferri, 2004 or Heathcote et al., 2014). Thinking of children as durable goods of irreversible nature that require investment of parental resources (Becker, 1960), research on consumption commitments suggests that—at least on a theoretical level—households could postpone childbearing when earnings risk is high, initially preferring to work and save

E-mail address: kamila.sommer@frb.gov

[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.

http://dx.doi.org/10.1016/j.jmoneco.2016.08.002
0304-3932/Published by Elsevier B.V.

before starting a family.[2] Since household ability to bear children declines with household age, the postponement of births could in turn lead to a (perhaps unintended) reduction in number of births per family.

The next section presents empirical evidence from micro data suggesting the link between delaying and reducing fertility on the one hand, and rising labor income risk on the other. At the macro level, falling fertility rates have been observed during periods when labor market risk was high. 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 et al., 2010 or Goldin and Margo, 1992). Recently, during the Great Recession, fertility declined precipitously amidst global financial turmoil and rapidly increasing unemployment, in part due to the postponement of births by younger age-cohorts.[3]

Starting with Section 3, this paper offers the first quantitative theoretical exploration of the link between earnings risk and fertility. In my Aiyagari–Bewley–Huggett framework augmented with fertility choice, unitary households face idiosyncratic wage shocks and make joint decisions about consumption, savings, family size, timing of births, and the allocation of resources (time and market goods) spent on improving children's quality. Fertility decisions are modeled as sequential, irreversible choices over the number of children. 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 the best of my knowledge, this paper is the first quantitative theoretical study where expenditures invested 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.

Using the exogenous estimates of labor market risk for the 1990s from Meghir and Pistaferri (2004), the model is calibrated based on cross-sectional fertility and income patterns of a U.S. cohort of households who made their fertility decisions in the 1980s and 1990s—a period associated with higher levels of idiosyncratic earnings uncertainty. Next, using Meghir and Pistaferri's (2004) risk estimates for the 1970s, the model is used to quantify the contribution of earnings uncertainty to the changes in the key U.S. fertility indicators between the two steady states.

It is shown that realistic increases in persistent labor market risk observed between the 1970s and 1990s could explain about half of the total decline in the number of births over the period, while accounting for a sizable fraction of the observed delay in childbearing. The key mechanism generating the postponement of births and the fertility decline 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 child-rearing), 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.

Finally, Section 5.2 tests how effective a broad-based adoption of IVF technology might be in mitigating the effects of increased labor market risk on household fertility. In particular, each household that realizes an infertility shock can choose whether or not to undergo up to two cycles of IVF treatment, calibrated to the most recent success rates. IVF allows households to more optimally time births by reducing age-dependent infertility risk. However, the adoption of IVF technology only partially offsets the effects of increased labor market risk on the total number of births, in part because IVF success rates are relatively low for women in their late thirties and early forties when—in a high earnings risk regime—households ideally desire to have higher-order births.

A vast body of studies in microeconomics, labor economics, and macroeconomics have explored channels that likely contribute to fertility changes. Most closely related to this paper, Santos and Weiss (2016) show that when marriage entails consumption commitments, then a rise in earnings risk can explain a sizable portion of the delay in marriage between the 1970s and the 2000s. In another related work, Da Rocha and Fuster (2006) show that, in a model with job search and female human capital accumulation, high unemployment risk induces women to postpone and space births, which in turn reduces the fertility rate. Other papers try to connect three trends: increasing female education, increasing female labor market participation, and declining fertility. For example, Conesa (2002) 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 contrast, Caucutt et al. (2002) argue that better education can explain less than one-third of the increase in

[2] For example, Chetty and Szeidl (2007) or Postlewaite et al. (2008) show that consumption commitments (i.e., big-ticket goods with sizable adjustment costs) 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.

[3] Between 2007 and 2011, the birth rate for women between ages 20–24 declined to the lowest level ever recorded in the U.S. while the birth rate for women between ages 25–29 reached the lowest level since 1976. 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 (Hamilton et al., 2012). Indeed, as shown in the online appendix, the U.S. fertility rate over the last 40 years has been pro-cyclical. Since the household labor market risk is known to rise during recessions (Storesletten et al., 2004), one interpretation of pro-cyclical fertility is that households postpone births when earnings uncertainty is high.

Table 1
OLS regression of number of births on earnings risk.

Variable	Coefficient	(Std. Err.)
Age of wife	0.390***	(0.00638)
Age of wife squared	−0.005***	(0.00009)
Occup. earnings risk	−1.597***	(0.15790)
Husband's total income	0.001***	(0.00009)
Wife's total income	−0.018***	(0.00017)
Intercept	−5.984***	(0.10687)
N	103271	
R^2	0.253	
$F_{(5,103265)}$	7010.8	

Note: The table shows the estimates for the OLS regression discussed in Section 2, wherein the number of births for any married couples is regressed on estimates of occupational uninsurable earnings risk associated with the husband's occupation, and other basic household characteristics, such as wife's age, and wife's and husband's income. The estimates of occupational earnings risk associated with each husband's occupation are drawn from Saks and Shore (2005). *, **, and *** denote that the estimated coefficient is statistically significant at 10, 5, and 1 percent levels, respectively.

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.

2. Earnings risk and fertility in data

Over the last four decades, U.S. household fertility patterns have changed significantly. The mean ages at 1st and 2nd births increased by 3.5 and 3.6 years between 1970 and 2000 (from 21.4 to 24.9 and 24.1 to 27.7), respectively, with the steepest increase from 1970 to 1990 (Mathews and Hamilton, 2009). At the same time, women who made fertility decisions in the 1960s and 1970s had, on average, 2.5 children by age 45, compared to 1.9 birth of women who made such decisions in the 1980s and 1990s.[4] The changes in fertility trends coincided with large shifts in earnings and fertility risks. In particular, a large body of literature documented sizable shifts in microeconomic earnings uncertainty since the early 1970s, with the lion's share of these increases being attributed to increases in uninsurable idiosyncratic earnings risk.[5] Moreover, since infertility risk increases exponentially with the age of the household (Trussell and Wilson, 1985; Wallace and Kelsey, 2010), the number of women with reported infertility issues has risen appreciably since the 1970s, as more women are attempting to give births at older, less fecund ages (Chandra and Stephen, 1998).

While the existing economic literature has largely abstracted from connecting the increases in these risks with the observed changes in household fertility, a suggestive correlation between delaying and reducing fertility on the one hand and rising earnings risk on the other can be found in the micro data. To document this correlation, in the first step, I use the riskiness of husband's occupation from Saks and Shore (2005) as a proxy for the earnings risk faced by households.[6] Specifically, 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. Their estimates are merged with the 5 percent sample of the 1990 Decennial Census, concentrating on married couples where the husband is not self-employed.[7] After the selection criteria are applied, the sample consists of roughly 100,000 married couples with wives between ages 20 and 43 years.

In the second step, I estimate a simple OLS model of completed fertility using the cross-sectional data set from the Census (Table 1). The dependent variable is the number of births for any given couple. The regressors include estimates of earnings risk associated with the husband's occupation, and other basic household characteristics such as wife's age, and wife's and husband's income (in thousands of dollars). Notably, the estimated coefficient on riskiness of husband's occupation is negative and statistically significant at a one percent level, confirming the negative correlation between fertility

[4] The average number of births for women who made childbearing choices in the 1960s and 1970s comes from Jones et al. (2008) 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.

[5] Next to studies mentioned in the Introduction, see also Levy and Murnane (1992), Gottschalk (1997), or DeBacker et al. (2013).

[6] 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.

[7] Self-employed individuals have been shown to face higher earnings risk than individuals working for wage or salary across occupations; see, for example, Saks and Shore (2005). Following the estimation strategy of Saks and Shore (2005) who estimate the occupational risk for male heads with at least a college degree, educational attainment of husbands is also controlled for. For details on sample selection in this paper, see the online appendix.

and earnings risk.[8] The signs of the remaining coefficients are aligned with economic theory. Namely, the effect of husband's income level on the number of births is positive, indicating that the demand for children rises with household income. 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 rest of the paper builds a structural model of household fertility behavior in the presence of earnings and fertility risks that attempts to explain the negative correlation between earnings risk and fertility observed in the micro data, and to quantify the contribution of the increase in earnings risk and its interaction with infertility risk to the changes in the timing and number of births over time.

3. The benchmark model

The model is based on the following assumptions. Young unitary 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 (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. 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 18, and live until age 80 with certainty. During the working state of life (through age 65), the household wage process is determined according to an idiosyncratic stochastic process $\ln w_t = \ln w_0 + h(t) + \epsilon_t + \nu_t$, where $h(t)$ governs the average age-profile of wages, and $\nu_t \sim N(0, \sigma_\nu^2)$ is a transitory shock to income received every period. The persistent shock, ϵ_t, also received each period, follows a first-order auto-regressive process $\epsilon_t = \rho\epsilon_{t-1} + \psi_t$, where $\psi_t \sim IID(0, \sigma_\epsilon^2)$ and $\epsilon_1 = 0$. After retirement ($t > 65$), households receive a pension transfer $w_t = \overline{w}$ from the government.

3.2. Preferences

In the spirit of Becker and Tomes (1976), each household has a per-period utility function $U = U(c_t, n_t, q_t)$, 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.[9] 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, $q_t = f(x_t, l_t, n_t)$, is a function of the total amount of goods ($x_t \geq 0$) and the fraction of time invested toward childrearing ($l_t \in [0, 1]$), respectively.[10] 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 $n_t > 0$. Households discount future at the rate $\beta \in (0, 1)$.

3.3. Process for dependent children

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 $n_{t+1}^b = n_t^b + K_t$ where $K_t = \{0, 1\}$, 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, with $n_{t+1} \sim Bi(n_t + K_t, p)$, with $n_{18}^b = n_{18} = 0$ and p being a time-invariant probability that a child becomes independent and leaves home. Finally, it is assumed that parents enjoy and make expenditures only on children who are young and live at home.[11]

[8] The online appendix contains an alternative regression specification that uses dummy indicators for husband's occupation, rather than the continuous measure of the riskiness of husband's education from Saks and Shore (2005). All results go through. The appendix also contains a graphical representation of the relationship in the data.

[9] In this model, households are not altruistic toward their offspring, leaving no bequest to their children.

[10] 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.

[11] 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).

Table 2
Parameters and moments.

Exogenous parameters			
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		
Estimated parameters			
Preference curvature κ	0.14		
Preference scale ψ	3.50		
Production share μ	0.35		
Elasticity of substitution in production $\frac{1}{1-\theta}$	$\frac{1}{1-0.70}$		
Lower bound on children's consumption $\underline{q}$	0.34		
Household economies to money input to production ψ_1	0.91		
Household economies to time input to production ψ_2	0.54		
Probability that a child stays at home $(1-p)$	0.98		
Targeted moments	Model	Data	Source
Average Number of Births at age 45	1.90	1.90	NLSY79
Average Number of Births at age 25	0.80	0.80	NLSY79
Mean number of children at home for households at 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	0.41	0.41	CEX
Elasticity of childrearing time w.r.t. number of children	0.25	0.23	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

Note: The calibration strategy involves fixing some parameter values exogenously, and estimating the remaining parameters using the method of simulated moments. The table summarizes all the parameters used to calibrate the model, as well as the data moments targeted in the estimation. The parameters and the data moments are discussed in detail in Section 4.

3.4. *Infertility risk*

Households face a binary idiosyncratic age-dependent infertility shock $f_t = \{I, F\}$ which arrives at the beginning of every period with a probability p_t^I. 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 $f_{t+j} = I \ \forall j \geq 0$).

3.5. *Dynamic program of fertile parents*

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

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

subject to

$$A_{t+1} = \begin{cases} (1+r)(A_t + (1-l_t)w_t - c_t - x_t) & \text{if } t \leq R; \\ (1+r)(A_t - c_t + \overline{w}) & \text{if } R < t \leq T, \end{cases} \tag{2}$$

$$q_t = f(x_t, l_t, n_t) \quad \text{with } q_t \geq \underline{q} \text{ if } n_t > 0, \tag{3}$$

$$n_{t+1} \sim Bi(n_t + K_t, p) \quad \text{with } n_{18} = 0, \tag{4}$$

by choosing the parental consumption ($c_t > 0$), savings ($A_{t+1} \geq 0$), and the time ($l_t \geq 0$) and money ($x_t \geq 0$) inputs into the production of the children's quality, q. Households also make a discrete decision whether to have a child next period ($K_t = 1$) or not ($K_t = 0$), and face uncertainty about their fertility status ($f_{t+1} = \{I, F\}$) next period. Eq. (4) summarizes the law of motion for children at home, n_t.

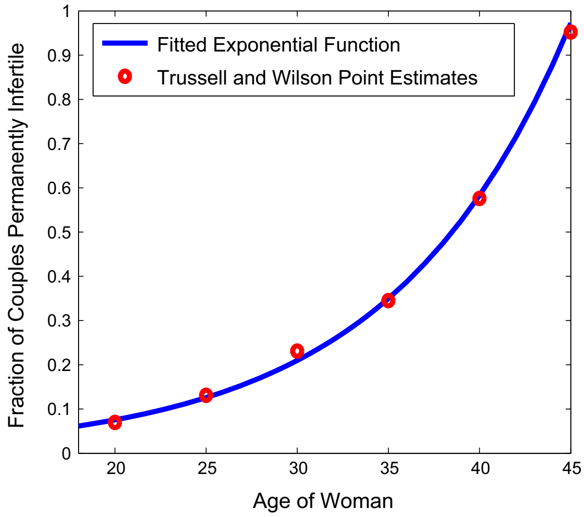

Fig. 1. Fraction of couples permanently infertile by age of wife. *Note:* The figure shows the benchmark cumulative distribution function of the permanent infertility risk (sourced from Trussell and Wilson, 1985) from which the age-dependent probabilities, p_t^I, associated with age-specific infertility shocks are derived. The point estimates (red dots) from Trussell and Wilson (1985) are fitted by an exponential function (blue solid line). The figure is discussed in Section 4.1. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

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:

$$V_t(a_t, n_t, w_t, f_t = I) = \max_{c_t, a_{t+1}, x_t, l_t} u(c_t, n_t, q_t) + \beta E_t V_{t+1}(a_{t+1}, n_{t+1}, w_{t+1}, f_{t+1} = I), \tag{5}$$

subject to the constraints (2) and (3), and to the law of motion $n_{t+1} \sim Bi(n_t, p)$, an analogue of Eq. (4) above.[12]

4. Calibration

The calibration strategy involves fixing some parameter values exogenously, and estimating the remaining parameters using the method of simulated moments.[13] All parameters are summarized in Table 2.

4.1. Infertility risk 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 (Fig. 1), from which the age-dependent probabilities, p_t^I, associated with a permanent infertility shock are derived. The probabilities, p_t^I, 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. 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.

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 et al. (1995), and Storesletten et al. (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 increased from about 0.15 in the 1970s to 0.21 in the 1980s (see the online appendix for details). 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

[12] In order to implement the inequality constraint $q_t \geq \underline{q}$, all households who violate the condition receive infinitely large negative utility so that households optimize away from such an outcome. Still, some particularly unlucky household theoretically might not be able to afford to consume above the floor on children's quality, requiring an exogenous transfer from an unmodeled government (see, for example, Santos and Weiss, 2016). However, in my experiments, no such cases are detected.

[13] The online appendix provides details on the sample selection and the calculation of moment conditions from these data sets.

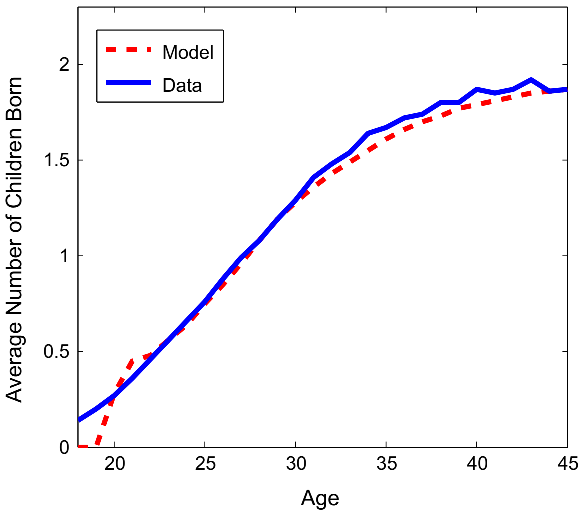

Fig. 2. Age-specific cumulative average number of births. *Note:* The figure shows the simulated age-profile of cumulative births (red dashed line) and compares it against the profile derived from the NLSY79 data (blue solid line). The figure is discussed in Section 4.6. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

Table 3
Effect of earnings and infertility risk on the number and timing of births.

	Model version	*Earnings risk*								
		Number of Births			Age at 1st Birth			Age at 2nd Birth		
		(A) No risk	(B) 1970s	(C) 1990s	(A) No risk	(B) 1970s	(C) 1990s	(A) No risk	(B) 1970s	(C) 1990s
Infertility risk	(1) Baseline	2.4	2.2	1.9	20.0	22.7	24.0	21.0	24.8	29.8
	(2) IVF	2.5	2.3	2.0	20.2	24.1	26.1	21.2	25.1	30.0
	(3) No risk	2.9	2.6	2.4	20.3	25.3	28.5	21.3	27.1	34.1

Note: The table summarizes the simulated effects of earnings and infertility risks on household childbearing patterns, discussed in detail in Section 5. Row (1) shows how number of births and the mean ages at first and second births vary in the baseline model with different levels of uninsurable earnings risk: no earnings risk (column A), the 1970s level (column B), and the 1990s level (column C). Row (2) shows how household fertility changes across these three earnings risk regimes once IVF is introduced. Row (3) captures the effect of earnings risk on household fertility when all infertility risk is eliminated from the model.

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.

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.[14] The retirement transfer, $w_t = \overline{w}$, is proportional to the household earnings in the last working period, with a replacement rate of 0.4.[15]

4.2. Preferences

Following the literature on fertility choice, the preferences are modeled as additively separable between consumption and fertility choices (i.e., the number of children and the children's quality): $U(c,n,q) = \frac{c^{1-\gamma}}{1-\gamma} + \zeta\frac{(nq)^{1-\kappa}}{1-\kappa}$, with $\gamma > 0$ and $\kappa > 0$. The constant relative risk aversion preferences over consumption are standard. To model household preferences over the number of children and their quality, a generalized version of the preference specification in De la Croix and Doepke (2003) is adopted. To parametrize these preferences, four parameters are needed: (γ, κ, ζ), plus the discount factor β. I thus set γ to a standard value of 1.5, and let the annual gross interest rate $(1+r) = 1.04$ so that $\beta = \frac{1}{1+r}$. The remaining two preference parameters ζ and κ are calibrated.

[14] 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.
[15] Using the Health Retirement Survey data and the Social Security Administration records, Munnell and Soto (2006) report that, on a household basis, the Social Security benefits provide an average replacement rate of 44 percent.

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: $q_t = \left[\mu\left(\frac{x_t}{n_t^{\psi_1}}\right)^\theta + (1-\mu)\left(\frac{l_t}{n_t^{\psi_2}}\right)^\theta\right]^{1/\theta}$, where $\mu \in [0,1]$ is the production share, and $\frac{1}{1-\theta}$ represents the elasticity of substitution between time (l_t) and goods (x_t) devoted to childrearing. Parameters ψ_1 and ψ_2 represent the household economies of scale in the time and market expenditures spent on childrearing. CES parameters μ and θ along with the parameters ψ_1 and ψ_2 are all estimated. 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 the process (4), 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.

4.5. Estimation

Based on the previous discussion, eight structural parameters must be calibrated to match targets computed from the data: $(\zeta, \kappa, \mu, \theta, \underline{q}, \psi_1, \psi_2, p)$.[16] The data targets are computed from the NLSY79, ATUS, and CEX, and are summarized in Table 2.

The value of the curvature parameter on the production of children's quality, κ, relative to the curvature parameter on own non-durable consumption, γ, affects the age profile of fertility relative to that of non-durable consumption, thereby impacting the timing of births over the life cycle. This motivates the use of the average number of births at age 25—the mean age at first birth computed from the NLSY79—as the targeted moment (0.8). The preference scale parameter, ψ, affects the amount of utility received from the production of children's quality relative to parental non-durable consumption, and is used to match the average number of children ever born to a household (1.9). The probability with which a child leaves home at a given period affects the household's age-profile of the number of children living at home; the mean number of children at home at age 35 (1.4) is thus used as a target. The elasticity of substitution of time and money in children's production, $\frac{1}{1-\theta}$, and the production share, μ, jointly affect the age-profile of correlation between number of births and earnings; as such these correlations at ages 25 and 45 are used as calibration targets (-0.20 and -0.02, respectively).[17] Finally, the lower bound on children's quality, $\underline{q}$, affects the amount of resources used toward childrearing relative to earnings and motives the choice to use the ratio of households childrearing expenditures to earnings as a calibration target (0.4).[18]

Finally, to pin down the household economies to childrearing ψ_1 and ψ_2, I run auxiliary regressions $\ln x_t = \alpha_0 + \alpha_1 \ln n_t$ and $\ln l_t = \gamma_0 + \gamma_1 \ln n_t$, where x_t and l_t represent the total amount of money and time spent on own children (n_t) by the CEX and ATUS families, respectively. The slope coefficients α_1 and γ_l—estimated at 0.41 and 0.23, respectively—represent the elasticity of money and time inputs into childrearing with respect to the number of children at home, and provide the last two moment conditions for the method of simulated moments.

4.6. Predictions of the benchmark model

The simulated age-profile of cumulative births, shown in Fig. 2, matches its NLSY79 counterpart well for households between ages 22 and 45, although the average number of births for very young households differs slightly from the data, in part due to unmodeled teenage pregnancies early in the life cycle.[19]

The high elasticity of substitution between time and market expenditures in childrearing—estimated at about 3.3 (see Table 2)—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. Second, in a model with deterministic wage growth over the life cycle (as in this paper), 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.

[16] 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.

[17] In this paper, a relatively high degree of substitutability between time and money inputs into the production of children's quality is needed to match the age-profile of correlations between earnings and fertility.

[18] Using the CEX data, Lino (2008) 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).

[19] In the model, the first child can be born at age 19.

5. Results

This section discusses the effects of earnings and fertility risks on household childbearing patterns, and also discusses how these risks interact in affecting household fertility. The main results are summarized in Table 3. The first row in Table 3, discussed in Section 5.1, shows how number of births and the mean ages at first and second birth vary in the baseline model with different levels of uninsurable earnings risk: no earnings risk (columns A), the 1970s level (columns B), and the 1990s level (columns C). The second row, discussed in Section 5.2, shows how household fertility changes across various earnings risk regimes once IVF is introduced, and is used to study how effective IVF technology might be in offsetting effects of higher earnings risk on household fertility. Finally, the third row, discussed in Section 5.3, captures the effect of earnings risk on household fertility in the extreme case when all infertility risk is eliminated from the model, and is used to quantify the relative importance and the interaction of earnings and fertility risks in the model.

5.1. Earnings risk and fertility

Row 1 of Table 3 show how increases in uninsurable labor market risk of the magnitudes observed in the data affect the number and timing of births. When the labor market increases from levels observed from the 1970s to the levels observed in the 1990s, the number of births by age 45 falls from 2.2 to 1.9, and the mean age at first birth rises from 22.7 to 24. The delay in childbearing is even more pronounced for higher-order births: the mean age at the second birth increases fully by 5 years, from 24.8 to 29.8. To put the simulated steady-state results into context of the U.S. time-series data (previously discussed in Section 2), the mean ages at 1st and 2nd birth increased by 3.5 and 3.6 years between 1970 and 2000 (from 21.4 to 24.9 and 24.1 to 27.7), while the average number of births at age 45 declined from 2.5 for women who made fertility decisions in the 1960s and 1970s to 1.9 for women who made such decisions in the 1980s and 1990s. Viewed in isolation, increases in the uninsurable earnings risk that are in line with the U.S. experience could explain about one half of the decline in fertility and about a third of the increase in the mean age at first birth, while matching broadly the changes in the timing of the second birth.

Turning to the mechanism, children are a durable good of an irreversible nature and childrearing is costly, as at least a minimum amount of time and money invested into each offspring is needed to maintain the average quality per child above $\underline{q}$. Hence, while having children provides households with utility, it also limits their ability to insure against adverse wage shocks through increased saving or labor supply. As such, 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. The longer the delay of first and higher-order births, the larger the reduction in fertility.

The next sections discuss the effect of infertility risk and changes in its treatment on household fertility.

5.2. Infertility risk: an IVF application

This section explores how effective this technology might be in attenuating the decline in family sizes associated with the delay of births generated by the heightened labor market risk. By way of background, the first "tube-baby" was born in the United States in 1981. However, the use of IVF technology became more commonplace only in the mid-1990s, and grew in popularity especially during in the second half of the 2000s.[20] Table 4 shows the evolution of age-specific IVF success rates over time. The odds of a birth out of an IVF cycle improved markedly since the late 1990s, but mostly for younger women who are generally more fertile. In particular, between 1997 and 2012 (the latest data point available), the success rates increased fully 11 percentage points (from 30 percent to 41 percent) for women less than age 35, but only about 4 percentage points for women ages 35–40. For women age 40+, the success rates were about unchanged over that period, likely reflecting that even the technological improvements in IVF treatment cannot completely undo the effects of aging on household ability to conceive.

To quantify the effect of IVF technology in mitigating the effects of aging on family size, this section introduces IVF technology into the model.[21] In particular, once an infertility shock $f_t = I$ is realized for the first time (i.e., $f_{t-1} = F$), each household is allowed to choose whether to undergo IVF treatment or not. In accordance with data in Table 4, it is assumed that two embryos are transferred per IVF cycle. Households that choose to undergo IVF treatment thus face three possible birth outcomes: no birth, a singleton birth, or a twin birth.

[20] Although the use of IVF is still relatively rare compared to the potential demand (likely largely due to its high cost), the number of administered IVF cycles more than doubled over the past 10 years, and today roughly one percent of all infants born in the U.S. every year are conceived through IVF. Source: Center for Disease Control and Prevention, http://www.cdc.gov/art/reports/index.html.

[21] As in the data, IVF technology is not available in the baseline model which is calibrated to the fertility profiles of the NLSY79 women who made their fertility choices mostly in the 1980s and 1990s.

Table 4
IVF success rates and number of embryos transferred per IVF cycle by woman's age.

IVF success rates (%) Source	Age	< 35	35–37	38–40	> 40	41–42	> 42
CDC	1997	29.4	24.4	16.8	8.3	–	–
	1999	31.0	25.2	18.6	9.5	–	–
	2001	33.5	27.3	18.6	10.4	–	–
	2003	36.5	30.3	20.4	10.6	–	–
SART	2003	37.5	30.4	20.2	–	11.2	4.5
	2004	36.6	29.3	19.5	–	10.6	3.9
	2005	37.1	29.2	19.7	–	10.6	3.5
	2006	38.8	30.6	20.6	–	10.9	4.3
	2007	39.9	30.5	21	–	11.7	4.6
	2008	41.3	31.1	22.2	–	12.3	4.1
	2009	41.4	31.7	22.3	–	12.6	4.2
	2010	41.7	31.9	22.1	–	12.5	4.1
	2011	40.1	31.9	21.6	–	12.2	4.2
	2012	40.7	31.3	22.2	–	11.8	3.9
Number of embryos transferred per IVF cycle Source	Age	< 35	35–37	38–40	> 40	41–42	> 42
SART	2012	1.9	2	2.4	–	2.9	2.9

Note: The table shows the evolution of IVF success rates by woman's age over time, and the average number of embryos transferred per IVF cycle by women's age as of 2012. IVF success rates are defined as a probability of a live birth out of an IVF cycle. The data are sourced from the Center for Disease Control and Prevention (CDC) and the Society for Assisted Reproductive Technologies (SART).

To calibrate the probabilities associated with each outcome, I use age-specific IVF success rates for 2012 from Table 4, and set the probability of a twin birth conditional on an IVF success to 29 percent (McLernon, 2010).[22] Each household is allowed to conduct a maximum of two IVF cycles in the period in which the household receives the permanent infertility shock for the first time: namely, the success rates are calculated so that if the first IVF cycle fails (i.e., no birth is realized), a household repeats the cycle one more time.

Turning to the results, a comparison of rows 1 and 2 in Table 3 suggests that the introduction of IVF technology increases the realized number of birth per family relative to the baseline model, but is not able to fully compensate for the decline in births associated with the increased labor market risk. In particular, when IVF technology is made available to women who made their fertility choices during the 1990s, the number of births by age 45 rises to 2.0 (relative to 1.9 in the baseline model with no IVF technology), but is still lower than the number of births (2.2) for women who faced the level of earnings risk associated with the 1970s.

The ability of IVF technology to offset the effects of increased labor market risk is limited for several reasons. First, when earnings risk is at the 1990s level, women optimally postpone the third and subsequent births into their late thirties and early forties, a stage of the life cycle when infertility risk is high and IVF success rates are relatively low.[23] In contrast, when earnings risk at the 1970s level, women have generally children in their twenties, a stage of the life cycle when infertility risk is low even without IVF. Second, the introduction of IVF technology is associated with a further delay in births, as households rely on the technology as a way to overcome a potential infertility shock. For example, the mean age at first birth increases from 24 to 26.1 in the economy with the 1990s levels of earnings risk, thereby increasing the probability that an infertility shock is in fact realized. To see this, compare row 1/columns C with row 2/columns C in Table 3.

Finally, in the context of this model, the estimated effect likely represents a lower bound on how effective IVF might be in offsetting the effects of increased labor market risk on fertility. In the model, the IVF option enables households to have an additional singleton or twin birth when the infertility shock is realized for the first time. However, households are not allowed to repeat the treatment in later ages. Mechanically, as the number of admissible IVF cycles increases, the probability of a success at any given age converges to unity, and one embryo per IVF cycle is transferred, the model with the IVF extension (row 1 in Table 3) folds into the model with no infertility risk (row 3). That said, the currently high costs of IVF treatment likely reduce the full potential effect of the option to receive and to repeat treatments relative to the model where the treatment is free.

[22] Note that while a probability of a twin birth conditional on IVF success is about 29 percent, the unconditional probability of a twin birth out of an IVF cycle is much lower: about 12 percent for women below age 35. This probability further falls with age to about 2 percent for women ages 42 and above, due to the declining probability of an IVF success.

[23] A relatively frequent occurrence of a third birth (or above) is generally needed to push number of births per family above 2.0.

5.3. Interaction of earnings and fertility risks

Finally, I quantify the interaction between income and infertility risk. In particular, the calibrated model is used to answer the following questions: If there is one risk and not the other, which is more important quantitatively? And how do these risks interact within the model?

A comparison of rows 1 and 3 in Table 3 suggests that both risks have significant effects on household fertility. In an economy with the 1990s level of earnings risk and no infertility risk, households give on average 2.4 births by age 45 (row 3/columns C). This is the same average as in the economy with the baseline infertility risk but no earnings risk (row 1/columns A). However, the two risks interact and amplify each other. In an economy where neither risk is present, the average number of births (2.9) is much larger than that (1.9) in an economy with these risks at their baseline levels (row 3/columns A vs. row 1/columns C).

To illustrate the amplification mechanism, when the 1990s level of earnings risk is introduced into an economy with no infertility risk, the number of births by age 45 falls from 2.9 to 2.4, and the mean age at the first birth rises from 20.3 to 28.5 (row 3/columns A vs. row 3/columns C). Hence, in the absence of infertility risk, households respond to increased labor market risk by delaying the first and higher-order births, but also by reducing their desired family sizes. When infertility risk is additionally introduced, the loss of births becomes magnified, even as households now have children at younger ages relative to the model with no infertility risk. In particular, the number of births falls further from 2.4 to 1.9, even though the mean age at first birth declines from 28.5 to 24 (row 3/columns C vs. row 1/columns C). In all, even though infertility risk induces households to have children at younger, more fertile ages, some births are still lost to the realized infertility shocks.

6. Conclusions

Motivated by large shifts in uninsurable income risk over time, this paper studied the relationship between household fertility choices and idiosyncratic earnings uncertainty using a life cycle model of fertility choice. The documented linkage between earnings and fertility risks and household fertility 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 fertility trends.

Earnings risk is certainly not the only determinant of fertility choices. Education, career, changes in marriage and mating, and increasing contraceptive use have all been reported as important factors affecting household fertility choices. While it is currently hard to study all of these channels in a unified framework (largely due to computational constraints), it would be interesting to see whether changes in household risk could link some of these trends.[24] 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. Given possible inter-linkages and feedback effects amongst these channels, I view this as an important avenue for future work.

Acknowledgments

I would especially like to thank Jonathan Heathcote for his insightful comments and ongoing support, and to Victor Rios-Rull for always challenging me to make this project better. Alvaro Mezza, Martin Sommer and Paul Sullivan were instrumental as readers of the paper's various drafts. The current draft also greatly benefited from my discussions with Robert Cumby, Alessandra Fogli, Garance Genicot, Mark Huggett, Larry Jones, Greg Kaplan, Vasia Panousi, Ellen McGrattan, and Daniel Ringo; and with Mathias Doepke, Raquel Fernandez, John Knowles, Christos Koulovatianos, Maria Luengo-Prado, Alice Schoonbroodt, Kjetil Storesletten, and Michele Tertilt who provided valuable feedback in its early stages. Rebecca Wasyk speeded up my code. 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.

Appendix A. Supplementary data

Supplementary data associated with this article can be found in the online version at http://dx.doi.org/10.1016/j.jmoneco.2016.08.002.

[24] Santos and Weiss (2016) already study the interaction of two of these factors: marriage/mating and earnings risk.

References

Becker, G.S., 1960. An economic analysis of fertility. In: Demographic and Economic Change in Developed Countries. NBER Chapters. National Bureau of Economic Research, Inc., pp. 209–240. http://www.nber.org/chapters/c2387.
Becker, G.S., Tomes, N., 1976. Child endowments and the quantity and quality of children. J. Polit. Econ. 84 (4), S143–162.
Card, D., 1991. Intertemporal Labor Supply: An Assessment. NBER Working Paper 3602. National Bureau of Economic Research.
Caucutt, E., Guner, N., Knowles, J., 2002. Why do women wait? Matching, wage inequality, and the incentives for fertility delay. Rev. Econ. Dyn. 5 (4), 815–855.
Chandra, A., Stephen, E.H., 1998. Impaired fecundity in the United States: 1982–1995. Fam. Plan. Perspect. 30 (1), 34–42.
Chetty, R., Szeidl, A., 2007. Consumption commitments and risk preferences. Q. J. Econ. 122 (2), 831–877.
Conesa, J.C., 2002. Educational Attainment and Timing of Fertility Decisions. Working Papers in Economics 78. Universitat de Barcelona, Espai de Recerca en Economia.
Da Rocha, J.M., Fuster, L., 2006. Why are fertility rates and female employment ratios positively correlated across O.E.C.D. countries? Int. Econ. Rev. 47 (4), 1187–1222.
De la Croix, D., Doepke, M., 2003. Inequality and growth: why differential fertility matters. Am. Econ. Rev. 93 (4), 1091–1113.
DeBacker, J., Heim, B., Ramnath, S., Panousi, V., Vidangos, I., 2013. Rising Inequality: Transitory or Persistent? New Evidence from a Panel of U.S. Tax Returns. Brookings Papers on Economic Activity Spring, pp. 67–122.
Fisher, J., Gervais, M., 2011. Why has home ownership fallen among the young? Int. Econ. Rev. 52 (3), 883–912.
Goldin, C., Margo, R., 1992. The great compression: the wage structure in the United States at mid-century. Q. J. Econ. 107 (1), 1–34.
Gottschalk, P., 1997. Inequality, income growth, and mobility: the basic facts. J. Econ. Perspect. 11 (2), 21–40.
Hamilton, B.E., Martin, J.A., Ventura, S.J., 2012. Births: preliminary data for 2011. Natl. Vital Stat. Rep. 61, 5.
Heathcote, J., Storesletten, K., Violante, G., 2014. Consumption and labor supply with partial insurance: an analytical framework. Am. Econ. Rev. 104 (7), 2075–2126.
Hotz, V.J., Miller, R.A., 1988. An empirical analysis of life cycle fertility and female labor supply. Econometrica 56 (1), 91–118.
Hubbard, R.G., Skinner, J., Zeldes, S.P., 1995. Precautionary saving and social insurance. J. Polit. Econ. 103 (2), 360–399.
Jones, L.E., Tertilt, M., 2008. An economic history of fertility in the U.S.: 1826–1960. In: Rupert, P. (Ed.), Frontiers of Family Economics. Elsevier. http://www.nber.org/papers/w12796.
Kopczuk, W., Saez, E., Song, J., 2010. Earnings inequality and mobility in the United States: evidence from social security data since 1937. Q. J. Econ. 125 (1), 91–128.
Levy, F., Murnane, R.J., 1992. U.S. earnings levels and earnings inequality: a review of recent trends and proposed explanations. J. Econ. Lit. 30 (3), 1333–1381.
Lino, M., 2008. Expenditures on Children by Families: 2007 Annual Report. Technical Report. U.S. Department of Agriculture, Center for Nutrition Policy and Promotion.
Mathews, T.J., Hamilton, B.E., 2009. Delayed Childbearing: More Women are Having their First Child Later in Life. National Center for Health Statistics Data Brief (21).
McLernon, D., 2010. Clinical effectiveness of elective single versus double embryo transfer: meta-analysis of individual patient data from randomized trials. BMJ Res. 341, c6945, http://dx.doi.org/10.1136/bmj.c6945.
Meghir, C., Pistaferri, L., 2004. Income variance dynamics and heterogeneity. Econometrica 72 (1), 1–32.
Munnell, A.H., Soto, M., 2006. What Replacement Rates do Households Actually Experience in Retirement? Working Paper 2005-10. Center for Retirement Research, Boston College.
Postlewaite, A., Samuelson, L., Silverman, D., 2008. Consumption commitments and employment contracts. Rev. Econ. Stud. 75 (2), 559–578.
Saks, R., Shore, S., 2005. Risk and career choice. Adv. Econ. Anal. Policy 5 (1), 1414.
Santos, C., Weiss, D., 2016. Why not to settle down already? A quantitative analysis of the delay in marriage. Int. Econ. Rev. 57 (2), 425–452.
Storesletten, K., Telmer, C.I., Yaron, A., 1998. The Risk Sharing Implications of Alternative Social Security Arrangements. GSIA Working Papers 252. Carnegie Mellon University, Tepper School of Business.
Storesletten, K., Telmer, C.I., Yaron, A., 2004. Cyclical dynamics in idiosyncratic labor market risk. J. Polit. Econ. 112 (3), 695–717.
Trussell, J., Wilson, C., 1985. Sterility in a population with natural fertility. Popul. Stud. 32 (2), 269–286.
Wallace, W.H.B., Kelsey, T.W., 2010. Human ovarian reserve from conception to the menopause. PLoS One 5, 1.