

Fertility and Consumption when Having a Child is a Risky Investment*

Pedro Gete[†] and Paolo Porchia[‡]

First Draft: August, 2010

This Draft: October, 2011

Abstract

This paper studies children as a risky asset associated to an investment option. Children provide utility but have a stochastic maintenance cost. We obtain several new results relative to models where children are deterministic goods, among which: i) Higher child risks diminish fertility and consumption. ii) Risk aversion speeds up fertility as households use the safe utility derived from a child as insurance against fluctuations in consumption. iii) Fertility is increasing in the correlation between income and child cost shocks. The household is reluctant to have children when positive cost shocks come together with bad income shocks. The opposite result happens when children hedge income shocks. iv) The sign of the correlation determines whether higher income volatility speeds up or delays fertility.

*We are very grateful to Gary Becker and Monika Piazzesi for their encouragement and useful advice. We also appreciate the helpful comments of Garance Genicot, Juan Pedro Gomez, Peter Hinrichs, Mark Huggett, Shareen Joshi, Anil Kashyap, Garen Markarian, Nancy Stokey and seminar participants at IE Business School, Georgetown-SFS and SED Annual Meetings 2011. Paolo Porchia gratefully acknowledges financial support from the Spanish Ministry of Science and Innovation under the Ramón y Cajal program (RYC-2010-07401).

[†]Georgetown University and IE Business School. Email: pg252@georgetown.edu

[‡]IE Business School. Email: paolo.porchia@ie.edu

1 Introduction

Having children is a risky investment because it is not possible to exactly know ex-ante the costs associated with raising a child (including the opportunity cost of the parents' time and the effects on their career paths) or the benefits that they will provide. A parent cannot predict, for example, how often her children will get sick, how much money and time that will cost, whether the child will need extra support at school or instead gets a full scholarship, or whether her child becomes a successful actor who brings in millions of dollars to the family. Moreover, insurance markets for many of the aforementioned contingencies, especially those related to the time costs for the parents, do not exist. Children are risky assets at most partially insurable. Consequently, childbearing adds another source of risk to households.

In this paper we study the consequences for fertility and consumption of taking into account that children are a source of risk that interacts with other risks borne by the parents. The existing literature has followed Gary Becker's seminal work in modelling children as deterministic durable goods, without considering that children are also a stochastic asset.¹ Considering children as risky assets associated to an investment option enlarges the set of effects that drive fertility.

We study the problem of an infinitely-lived household who has an initial level of assets and every period gets some stochastic income, that we assume exogenous for simplicity. She gets utility from consumption and can save at the risk free rate.² The decision to have a child is like the decision to exercise an option. Every period she can have a child or postpone the decision. If she has a child then she is acquiring an irreversible, durable and non-tradable asset that gives her utility, but implies stochastic exogenous costs.³ We focus on the uncertainty about children's costs and assume that the utility flow is deterministic.

Our main results are the following: i) Higher cost volatility diminishes fertility and consumption. Intuitively, risk averse households are less willing to invest in riskier assets. Higher risk results in higher precautionary savings. ii) Risk aversion speeds up fertility as households use the safe utility derived from a child as insurance against fluctuations in consumption. iii) Fertility is increasing in the correlation between income and child cost shocks. The household is reluctant to have children when positive cost shocks come together with bad income shocks (for example, households who must reduce hours of work and damage their careers when their child

¹Jones et al. (2008) and Tamura (2000) survey the literature on the economics of fertility.

²This is the only available asset. Thus, markets are incomplete.

³These are realistic assumptions given that once a child is born is very costly to get rid of her because of sentimental reasons and legal constraints. Moreover, children do not depreciate and it takes several years before they can live independently.

are sick). iv) The sign of the correlation determines whether higher income volatility speeds up or delays fertility.

These results come from the interplay of five effects that differ from the standard income and substitution effects: 1) As in any real options problem, having the option to time investment implies asymmetric convex payoffs ("in bad times do not exercise and wait for good times to come"). This makes the value of the option increasing in child cost volatility and encourages to delay fertility. 2) Higher cost volatility increases the risk of not exercising the option (the range of tomorrow prices is larger). This pushes risk averse households towards earlier fertility (exercise today at known price). 3) If the fertility option is exercised the child costs can be thought as an income shock. Incomplete markets and preferences displaying precautionary savings push the household to delay fertility to avoid new sources of risk. 4) A child is an alternative source of utility different from goods consumption. The utility derived from being parent can hedge fluctuations in utility from goods consumption. This channel pushes for early fertility to enjoy the children as an insurance mechanism. 5) When the income and cost shocks are correlated, the sign of the correlation determines if the shocks hedge or add up each other. Fertility speeds up when the shocks hedge each other.

Our analysis offer new insights that may help both in explaining facts and in the design of pronatalist policies. For example, Stetsenko (2010) documents that U.S. fertility changed from being countercyclical to being procyclical as female labor supply increased. This is consistent with our model if the increase in female labor supply implies a negative correlation between child costs and income shocks. On the policy side, our theory provides support for policies that reduce the uninsurable uncertainty from raising a child (such as child care programs), and the negative correlation between income and child costs shocks (as State-paid leaves of absence to insure the parent's career from children health shocks).

This paper contributes to the literature on the economics of fertility. Our innovation is to focus on the risks coming from children themselves. Moreover, we analyze the fertility decision from the perspective of the value of the fertility option.⁴ Our work is theoretical, we focus on characterizing numerically a set of results that could not be proved analytically. We checked the robustness of our results and provided comparative statics exercises to understand the effects of each parameter.

Our work complements the studies on the effects of income uncertainty on fertility. Several empirical papers have shown that fertility responds negatively to increases in income risk.⁵

⁴Ranjan (1999) is the only paper we are aware of that explicitly discusses having a child as an option, however in Ranjan's model, as in all existing models of fertility, children are deterministic.

⁵Adsera and Menendez (2011) focused on Latin America; Adsera (2005, 2006 and 2011) in Southern Europe;

Conesa (2000), Choi (2011), Da Rocha and Fuster (2006) and Sommer (2011) are different quantitative models in which income risk or unemployment delay fertility. In those models children are deterministic assets. Our results when children are risky and with shocks adding to income shocks suggest that this modification would reinforce the delays.

Moreover, our work also complements the literature studying children as an insurance mechanism (Portner 2005 surveys theoretical and empirical work). This literature has focused on the insurance coming from children as another source of household income, or because children take care of ageing parents. We show two other ways in which children serve as insurance mechanism: 1) children hedge utility fluctuations because they provide a safe flow of utility. What is especially valued by low wealth households. 2) If children cost shocks are positively correlated with income shocks.

The paper proceeds as follows. Section 2 motivates why a child is a source of risks and how they correlate with income shocks. Section 3 describes the model. Section 4 discusses the theoretical results. Section 5 concludes. Details on the solution method are in the Appendix.

2 Why is a child a risky investment?

Children are sources of costs and benefits, both monetary and non-monetary (E.g. the opportunity cost of the time parents spend on them, or the affection that children give to the parents). Both costs and benefits have a stochastic component. These risks come from a variety of sources and change as the child grows. For example, newborns face risks of birth defects.⁶ Different defects imply different costs. For example, Parish et al. (2009) document wide heterogeneity in the costs for a low income family of having a child with special needs.⁷ Part of the variation is associated with variation across states in the income eligibility guidelines for Medicaid and SCHIP. I.e., with variation in the amount of insurance against the shocks.

For child risks to affect behavior they need to be noninsurable. In the U.S. that is the case with most of the time costs for the parents. For example, most child care centers ask the parents to take their child home when the child gets sick.⁸ The National Association for Sick

Bhaumik and Nugent (2006) on East Germany; Mira and Ahn (2002) in a sample of OECD countries; Sommer (2011) in the U.S. Kreyenfeld (2005) provides a survey.

⁶The Center for Disease Control and Prevention has quantified the probability distributions: <http://www.cdc.gov/ncbddd/birthdefects/data.html>

⁷Among these families with positive expending, 30% had expenses between \$250 and \$500, and 34% had expenses of more than \$500. Twenty-seven percent had expenditures from 5.6% to 25.8% of their total household income.

⁸Hotz and Miao (2011) report that in the spring of 2005, 7.2 million children under the age of 5 (36.9%)

Child Daycare says that there is a huge unmet need for sick child care. They estimate that each day more than 350,000 children younger than 14 years of age, with both parents working, are too sick to attend school or child care. And that working mothers are absent from their jobs from 5 to 29 days per year caring for ill children.⁹

Why a child gets ill may be endogenous to the parents' behavior. But, according to the 2008 National Health Interview Survey, exogenous shocks seem to be very important. This survey reports that roughly 30% of the children missed 1 or 2 school days, 30% missed 3-5 days and 10% more than 6 days. These numbers are quite robust to differences in income, race or education of the parents. Thus, parents with young children face exogenous risks that their kids get sick. A sick child implies time and money costs (most health insurance plans have deductibles and co-payments).

As the child gets older the risks evolve, health shocks are smaller (with accidents becoming tail events) and uncertainty in education expenses become the main risks. There is large heterogeneity on how expensive is to educate a child and part of it comes from uninsurable random factors as children's ability.

Are child costs risks correlated with income risks? Most of the times they are. Probably the most common example are children health shocks that force the parents to divert time from work. Hours not worked are usually not paid, and may harm the parent's career path. Thus, in this case a bad child cost adds to a negative income shock (in our model we will say they have negative correlation). It may also happen that the shocks hedge each other (positive correlation in our model). For example, the parents of child athletes receive negative income shocks as they sacrifice their careers to travel with their children, but this comes while their children receive negative cost shocks (they earn money for the family). Or, as another example, eligibility guidelines for public subsidies may imply that an increase in child costs allow the household to qualify for public support, thus generating positive correlation between child costs and family income.

In this paper we will not take a stand on which of the previous empirical facts is more likely. We will study the theoretical consequences of the different possibilities.

were in some form of non-relative care, of which 4.6 million (23.3%) were in some form of organized child care. Among employed mothers with a child under 5, 5.9 million of their children (52.1%) were in non-relative child care, of which 3.6 million (31.9%) were in an organized care facility.

⁹<http://www.nascd.com/>

3 Model

We analyze a discrete-time, partial equilibrium economy, populated by an infinitely-lived household whose utility function depends on her consumption of goods (c_t) and on having or not a child (I_t). The household starts without a child at time $t = 0$ and has the option to have it at any moment. We denote by I_t an indicator function that captures the fertility status at time t : it takes the value one if the household has had a child prior or at time t , while it takes the value zero if she has not. We assume that the utility function is separable into goods' consumption and the utility from the fertility status. We assume Constant Relative Risk Aversion for the consumption component:

$$u(c_t, I_t) = \frac{c_t^{1-\gamma}}{1-\gamma} - \frac{I_t}{1-\alpha}, \quad \alpha > 1, \gamma > 0, \quad (1)$$

where γ is the relative risk aversion coefficient and α is a parameter that captures the utility benefit from having a child.

Every period the household earns a stochastic labor income stream (y_t). Moreover, if she has had a child she must pay the stochastic costs (q_t) associated with child rearing. We assume that markets are not complete and the household cannot hedge against child cost or income risk because she can only transfer consumption over time via the riskless asset. We denote by r the constant one-period risk-free rate. Hence the wealth (W) of the household evolves as

$$W_{t+1} = (W_t + y_t - q_t I_t - c_t)(1 + r), \quad (2)$$

where both income and costs are denominated in units of the consumption good that we use as numeraire.

We model the dynamics of the stochastic processes as first order autoregressions:

$$y_{t+1} = a_y + b_y y_t + e_{t+1}^y, \quad (3)$$

$$q_{t+1} = a_q + b_q q_t + e_{t+1}^q, \quad (4)$$

$$(e_{t+1}^y, e_{t+1}^q) \stackrel{iid}{\sim} N \left(0, \begin{pmatrix} \sigma_y^2 & \sigma_y \sigma_q \rho \\ \sigma_y \sigma_q \rho & \sigma_q^2 \end{pmatrix} \right), \quad (5)$$

where ρ denotes the correlation coefficient between shocks to income and child costs, while σ_y and σ_q denote the volatilities.

The household maximizes time-additive expected utility over consumption and her parental

status, i.e. she decides two things: whether and when to have a child, and how much to consume and save every period. After she has had a child she only has to decide how much to save and consume. The problem of the household before having a child is:

$$H(W_0, q_0, y_0) = \max_{\{c_t, I_t\}} \mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t u(c_t, I_t) \right] \quad (6)$$

$$s.t. \quad (2) - (5). \quad (7)$$

3.1 Solving the model

The problem has a recursive nature. Conditional on having or not a child, the household's decisions depend on her wealth, income and on the cost of the child. These are the state variables of the problem. We denote by $J(W, q, y)$ the value function of a household who already had a child.¹⁰

$$J(W, q, y) = \max_c \left[\frac{c^{1-\gamma}}{1-\gamma} - \frac{1}{1-\alpha} + \beta \mathbb{E} [J(W', q', y')] \right] \quad (8)$$

$$s.t. \quad (9)$$

$$W' = (W + y - q - c)(1 + r), \quad (10)$$

$$y' = a_y + b_y y + e^y, \quad (11)$$

$$q' = a_q + b_q q + e^q. \quad (12)$$

The household without a child compares the value of having a child (the value function J), with the value of continuing without child. Her Bellman equation is:

$$H(W, q, y) = \max_I (1 - I) \left\{ \begin{array}{l} \max_c \left[\frac{c^{1-\gamma}}{1-\gamma} + \beta \mathbb{E} [H(W', q', y')] \right] \\ s.t. \\ W' = (W + y - c)(1 + r) \\ y' = a_y + b_y y + e^y \\ q' = a_q + b_q q + e^q \end{array} \right\} + I J(W, q, y). \quad (13)$$

As in the option pricing literature, we define the continuation region as the set of realizations of the state variables at which the household is better off without children and the fertility option is worth more alive. The fertility region is the complement of the continuation region. The

¹⁰An apostrophe identifies next-period quantities.

boundary between these regions is the critical level of costs below which the household is better off with a child, $\bar{q}(W, y)$. This decision (or fertility) boundary depends on current wealth and income:

$$\bar{q}(W, y) := \left\{ \sup q : \left\{ \begin{array}{l} \max_c \left[\frac{c^{1-\gamma}}{1-\gamma} + \beta \mathbb{E} [H(W', q', y')] \right] \\ \text{s.t.} \\ W' = (W + y - c)(1 + r) \\ y' = a_y + b_y y + e^y \\ q' = a_q + b_q q + e^q \end{array} \right\} < J(W, q, y) \right\}, \quad (14)$$

where $J(W, q, y)$ is decreasing in q . The household has a child once the stochastic cost process is equal to or smaller than $\bar{q}(W, y)$.

4 Results

In this Section we analyze the predictions of the model. Given the lack of closed form solutions we solve the model numerically to illustrate theoretical results that could not be proved analytically. We check the robustness of our results. Since fertility occurs when the cost is smaller than the boundary $\bar{q}(W, y)$, we refer to a smaller boundary as less or deferred fertility.

4.1 Parameterization

In selecting parameters our goal is not to solve a particular numerical exercise but to show our theoretical results. Thus, while we choose somewhat plausible values, we invite the reader not to focus on the precise numbers but more on the comparative statics exercises. They illustrate how changing the parameters affect the results.

Concerning preference parameters: we assume $\gamma = 3$ as benchmark coefficient of relative risk aversion. This is in the ballpark of those standard in macro models, and of those used in models with preferences additively separable in consumption and childbearing utility as, for example, Becker et al. (1990), or Jones et al. (2008). We checked thoroughly how risk aversion affects our results. We assume one period in our model to be one year and set the annual interest rate to 3%. As subjective discount rate β we assume $1/1.03$. Concerning α , the parameter which governs the deterministic utility from having a child, its value is crucial, especially for low wealth individuals with low consumption. As a benchmark we set it such that the utility

stream from a child equals 50% of the utility stream from consumption of a household with no fertility option. This value is arbitrary and we will discuss how its changes affect the results.

Concerning wealth and income we proceed in the following way: we set a_y such that the income long-term mean is normalized to 10, and consider different wealth levels in a grid ranging from 1/2 to 6 times that value. This range gives us room to explore the effects of wealth and precautionary savings in the fertility decision of the households. We assume income shocks to be fairly persistent, $b_y = 0.88$. Persistence is a consistent result in the empirical literature (see Sommer 2011 for a survey of recent papers estimating stochastic processes for households' income using PSID). Inspired by that literature, we set the income standard deviation σ_y to 0.19 and do comparative statics on it.

Concerning the parameters on the costs of a child: The U.S. Department of Agriculture provides estimates of expenditures on children from birth through age 17 (Lino 2011). This guides us to choose a_q , the long-term mean of the child cost process, to match a long term child cost mean of 18% of long term mean income. Then, concerning the most important parameters for our study: the persistence of the child cost shocks (b_q) and their volatility (σ_q), we start with an agnostic choice ($b_q = b_y$, $\sigma_q = \sigma_y$) and elaborate intensively of how altering them alters the results.¹¹ As benchmark correlation we use $\rho = 0$, we will discuss how altering the sign of this correlation matters.

Table 1 summarizes our benchmark parameterization. For each result we will discuss which parameters we are varying.

Insert Table 1 about here

4.2 Results

We start by analyzing the influence of cost volatility. Figure 1 reports the fertility boundary as a function of cost volatility for different levels of risk aversion and wealth. In Panel B the household is more than twice wealthier than in Panel A.

Insert Figure 1 about here

We find two results: i) higher cost volatility diminishes fertility, but the effect is less pronounced for wealthier households. ii) Higher risk aversion increases fertility. These results come from the interplay of several channels:

¹¹Section 2 showed that there are different ways to think on the risks coming from childrearing thus we believe it makes sense to study a broad range of parameters.

- 1) As in any real options problem, having the option to time investment implies asymmetric convex payoffs ("in bad times do not exercise and wait for good times to come"). This makes the value of the option increasing in the cost volatility for a risk neutral household. For a risk averse household this channel is more subtle. On one side the previous argument applies. The higher the risk aversion the higher the value of being able to time fertility (avoid investing in bad times). However, when cost volatility increases also the risk of not exercising the option increases. Intuitively, the household needs to choose between exercising today at price q or having the option to exercise tomorrow. When cost volatility increases waiting for tomorrow is riskier (the range of tomorrow prices is larger), what pushes a risk averse household towards earlier exercise of the option.

- 2) Once the option is exercised, the household starts bearing the uninsurable risk of the costs of raising a child. Her problem becomes a standard incomplete-markets consumption problem with stochastic income.¹² We can read this from equation (2) thinking on $(y_t - q_t)$ as income net of children costs. If the household has precautionary savings preferences ($\gamma > 0$), then she will reduce consumption for precautionary reasons once exposed to an additional source of net income risk. The precautionary savings effect is present both before and after the option exercise, but it is larger after the option exercise because the option to time fertility provides some insurance (it allows to avoid investing in bad times). Precautionary savings are increasing in risk aversion and child cost volatility, while decreasing in wealth. Higher precautionary savings push the household to reduce consumption and delay fertility.

- 3) A child is an alternative source of utility different from goods consumption. The parameter α governs how big this source is. The utility derived from being parent can hedge fluctuations in utility from goods consumption. This is an important characteristic of children when we think of them as an asset class. They are assets providing non-monetary utility. This channel pushes the household for early fertility to enjoy the children as an insurance mechanism.

Channel 2 explains why the fertility boundary is decreasing in the volatility of the cost of having a child. While for small volatility Channel 1 can partially offset Channel 2, as volatility increases Channel 2 pushes the household to delay fertility. Wealthier households do less precautionary savings (their Channel 2 is smaller) and Panel B reflects this in a fertility boundary with less steep as child risks increase.

¹²See for example Caballero (1991) and Wang (2006).

Channels 1 and 3 explain why risk aversion (γ) speeds up fertility. As child cost volatility increases, the higher the risk aversion, the less prone is the household to wait for more favourable costs conditions because there is higher cost uncertainty about tomorrow. Moreover, a child is a way to hedge future consumption uncertainty buffering negative income shocks. A child assures a constant flow of altruistic utility that offsets low consumption levels. Higher γ also exacerbates Channel 2, thus making fertility less attractive for more risk averse households, nonetheless Channels 1 and 3 dominate.

Figure 2 reports average consumption with the respect to the cost distribution. I.e., weighted averages of the consumption associated with each child cost level with weights being the unconditional probability of each cost level.¹³ Panels A and B focus on households who did not have yet a child. While Panels C and D report households with children. We compare two wealth levels.

Insert Figure 2 about here

In all four panels consumption decreases as child cost volatility increases. And, as Channel 2 predicts, the decrease is steeper for households with children and low wealth (Panel C). Among low wealth households we know from Figure 1A that those with higher risk aversion are more likely to have children, and those are who decrease consumption more in Figure 2A. They are substituting present consumption with the utility from being parents. When wealth is higher it is harder to see these mechanisms.

Decreasing α , the utility that parents enjoy from the children would make the child more attractive, what reinforces Channel 3. We see in Figure 3a that this pushes for higher fertility for any wealth level. Figure 3b shows that when α is smaller households substitute more utility from consumption with childbearing utility, thus as child cost risk increases (higher cost volatility) they increase their precautionary savings. As it is expected, the increase in precautionary savings from the extra desire to have children is larger for less wealthy households (Panel A vs Panel B of Figure 3b).

Insert Figure 3a about here

Insert Figure 3b about here

Figure 4a studies the effects of child cost persistence (b_q) on the fertility decision.¹⁴ The effects depend on wealth levels. For low persistence (small b_q) fertility is decreasing for lower

¹³Numerical details are discussed in the appendix.

¹⁴Rewriting the cost process (4) as:

$$\Delta q_{t+1} = (1 - b_q)(\bar{q} - q_t) + e_{t+1}^q$$

wealth households (increasing boundary) and decreasing for higher wealth households (decreasing boundary). The intuition is as follows: when b_q is low the cost process converges very fast to the long-run mean cost level. Poor households that would have a child at costs below this long run mean become more conservative, they know that any low cost level will be short-lasting. Wealthy household, which can afford children at mean cost levels, interpret the lack of cost persistence as an advantage, they can absorb short term high costs thanks to their wealth buffer, while they expect costs to decrease to the mean in a short amount of time. Thus, wealthier households take advantage of the high speed of mean reversion to have a child at higher cost levels and enjoy earlier childbearing utility.

Insert Figure 4a about here

Child cost persistence (b_q) alters the relation between fertility and child risk, as shown in Figure 4b. When b_q is closer to 1 the cost level at which the fertility option is exercised matters for long time. This reinforces the asymmetric convex payoffs discussed as Channel 1: the rewards from waiting to exercise in good times are higher because good times last longer. Thus, higher b_q delays fertility as cost volatility increases.

Insert Figure 4b about here

Higher b_q implies a substantial increase in precautionary savings. Figure 4c shows that average consumption with the respect to the cost distribution decreases very fast when persistence is high. Bad shocks last longer and precautionary households save more. When persistence is low shocks do not last long and consumption before fertility is almost insensitive to cost volatility.¹⁵

Insert Figure 4c about here

Figure 5 shows an important effect of child cost shocks. As discussed in Section 2, they can hedge or reinforce shocks to income. Thus, fertility is monotonically increasing in the correlation of income and cost shocks, regardless of the household's risk aversion.

Insert Figure 5 about here

with $\Delta q_{t+1} = q_{t+1} - q_t$ and $\bar{q} = a_q/(1 - b_q)$ we can identify $(1 - b_q)$ as the cost speed of mean reversion.

¹⁵The low cost consumption in Figure 4c looks flat but it is not exactly flat in reality, its variation is smaller than the step size of our grid. Thus, since the numerical algorithm cannot capture those small variations the consumption looks flat.

Intuitively, the household is reluctant to have children when positive child cost shocks come together with bad income shocks. A pronatalist government may encourage fertility by altering the correlation. For example, a negative correlation may be the outcome of child illnesses having negative effects on the parents' careers. State paid leaves in periods of high probability of child illnesses may break the correlation.

The correlation of income and cost shocks alters how income volatility affects fertility. This is studied in Figure 6.

Insert Figure 6 about here

When shocks are uncorrelated higher income volatility defers fertility (decreasing boundary in Figure 6) because of Channel 1 and 2. Higher income volatility pushes for higher precautionary savings, and the household is reluctant to add the additional source of risk of a child to an already volatile income. Moreover, when income volatility is high precautionary households value even more the option to time fertility and are more conservative in exercising it. If the correlation is negative the previous effects are reinforced because bad child cost shocks come together with bad income shocks. Thus for $\rho < 0$ high income volatility decreases fertility a lot. This is consistent with the empirical literature on the effects of income uncertainty on fertility cited in the introduction.

However, when costs hedge each other, higher income volatility may push for earlier fertility. We see this in Panel A of Figure 6. However, if the household is enough risk averse, the Channels 1 and 2 discussed before push for decreasing fertility even if the correlation is positive (Panel B of Figure 6). These channels are tampered if income is not very persistent. This is shown in Figure 7 that redoes Panel B of Figure 6 for $b_y = 0.4$ instead of $b_y = 0.88$.

Insert Figure 7 about here

When b_y is low, income shocks do not last long. Thus, there is lower need for precautionary savings. Figure 7 shows that in this case when the shocks hedge each other ($\rho > 0$) higher income volatility does not seem to affect fertility. The hedging benefit from having a child is offsetting the push for decreasing fertility from Channels 1 and 2 discussed before. Our analysis then predicts that households with high but volatile and nonpersistent income (hence very sensitive to the correlation between income and child risks) are those whose fertility rates are reduced more by a negative correlation.

So far we have focused on comparative statics on wealth rather on income levels. For a qualitative analysis this does not imply any loss of generality for very persistent income process (as in the benchmark parameterization of Table 1). Figure 8 explores the effects of different

levels of income persistence on fertility. Intuitively, households with bad income shocks (below average current income) have more children when income shocks are not persistent (fertility boundary decreases in b_y for low income household). The opposite trend happens for households with good income shocks (above average current income).

Insert Figure 8 about here

Finally, Figures 9 and 10 illustrate the standard income and substitution effects well known since Becker (1960). In Figure 9, when children are more expensive (higher long-run mean of costs) the household substitutes them for goods consumption. In Figure 10, since children are normal goods, fertility is increasing in wealth (Panel A), and income (Panel B).

Insert Figure 9 about here

Insert Figure 10 about here

5 Conclusions

In this paper we have studied the joint consumption and fertility decisions of a household who faces uninsurable shocks to both income and the cost of raising a child. Our model emphasizes two main determinants of fertility: 1) income and child risks and their correlation; 2) The trade-off between the safe utility flow brought by a child and the aversion to the risks that the child implies.

Our analysis can be extended in a number of directions. For simplicity, we have not considered life-cycle effects. This is a relevant factor in quantitative work as the ability to be fertile is age dependent. We have also abstracted from the fact that income is endogenous and fertility alters labor supply. Introducing these extensions into the model would allow to study career and fertility choices together. We leave these extensions, together with empirical tests of the theory, for future research.

References

Adsera, A.: 2005, "Vanishing Children. From High Unemployment to Low Fertility in Europe", *American Economic Review, Papers and Proceedings* **95**(2).

- Adsera, A.: 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.
- Adsera, A.: 2011, "The Interplay of Employment Uncertainty and Education in Explaining Second Births in Europe", *Demographic Research* **25**(16), 513–544.
- Becker, G.: 1960, "An Economic Analysis of Fertility", *Demographic and economic change in developed countries* **11**, 209–231.
- Becker, G., Murphy, K. and Tamura, R.: 1990, "Human Capital, Fertility, and Economic Growth", *Journal of Political Economy* **98**(5), 12–37.
- Bhaumik, S. and Nugent, J.: 2006, "Does Economic Uncertainty Affect the Decision to Bear Children? Evidence from East and West Germany".
- Caballero, R.: 1991, "Earnings Uncertainty and Aggregate Wealth Accumulation", *The American Economic Review* **81**(4), 859–871.
- Choi, S.: 2011, "Fertility Risk in the Life-Cycle".
- Conesa, J.: 2000, "Educational Attainment and Timing of Fertility Decisions".
- Da Rocha, J. and Fuster, L.: 2006, "Why Are Fertility Rates And Female Employment Ratios Positively Correlated Across O.E.C.D. Countries?", *International Economic Review* **47**(4), 1187–1222.
- Hotz, V. and Xiao, M.: 2010, "The Impact of Regulations on the Supply and Quality of Care in Child Care Markets", *American Economic Review* .
- Jones, L., Schoonbroodt, A. and Tertilt, M.: 2008, "Fertility Theories: Can They Explain the Negative Fertility-Income Relationship?", *NBER working paper 14266* .
- Kreyenfeld, M.: 2005, "Economic Uncertainty and Fertility Postponement. Evidence from German Panel Data".
- Lino, M.: 2011, "Expenditures on children by families, 2010", *US Department of Agriculture, Miscellaneous Publication No. 1528-2010* .
- Menendez, A. and Adsera, A.: 2011, "Fertility Changes in Latin America in the Context of Economic Uncertainty".

- Mira, P. and Ahn, N.: 2002, "A note on the changing relationship between fertility and female employment rates in developed countries", *Journal of Population Economics* **15**(4), 667–682.
- Parish, S., Shattuck, P. and Rose, R.: 2009, "Financial Burden of Raising CSHCN: Association with State Policy Choices", *Pediatrics* **124**(Supplement 4), S435.
- Portner, C.: 2005, "Risk and Household Structure: Another Look at the Determinants of Fertility".
- Ranjan, P.: 1999, "Fertility behaviour under income uncertainty", *European Journal of Population* **15**(1), 25–43.
- Sommer, K.: 2011, "Fertility Choice in a Life Cycle Model with Idiosyncratic Uninsurable Earnings Risk".
- Stetsenko, S.: 2010, "Essays on the Macroeconomics of Labor Markets", *University of Pennsylvania Dissertations* .
- Tamura, R.: 2000, "Growth, fertility and human capital: A survey", *Spanish Economic Review* **2**(3), 183–229.
- Tauchen, G.: 1986, "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions", *Economics letters* **20**(2), 177–181.
- Terry, S. and Knotek, E.: 2010, "Markov-Chain Approximations of Vector Autoregressions: Application of General Multivariate-Normal Integration Techniques", *Economics Letters* .
- Wang, N.: 2006, "Generalizing the permanent-income hypothesis: Revisiting Friedman's conjecture on consumption", *Journal of Monetary Economics* **53**(4), 737–752.

Tables and Figures

Table 1: Benchmark parameterization

$r = 0.03$	$\gamma = 3$	
$\alpha = 22.54$	$\beta = \frac{1}{1+r}$	$\rho = 0$
$a_y = 10(1 - b_y)$	$b_y = 0.88$	$\sigma_y = 0.19$
$a_q = 1.82(1 - b_q)$	$b_q = b_y$	$\sigma_q = \sigma_y$

Table 1: *Benchmark parameterization.* The values are discussed in Subsection 4.1.

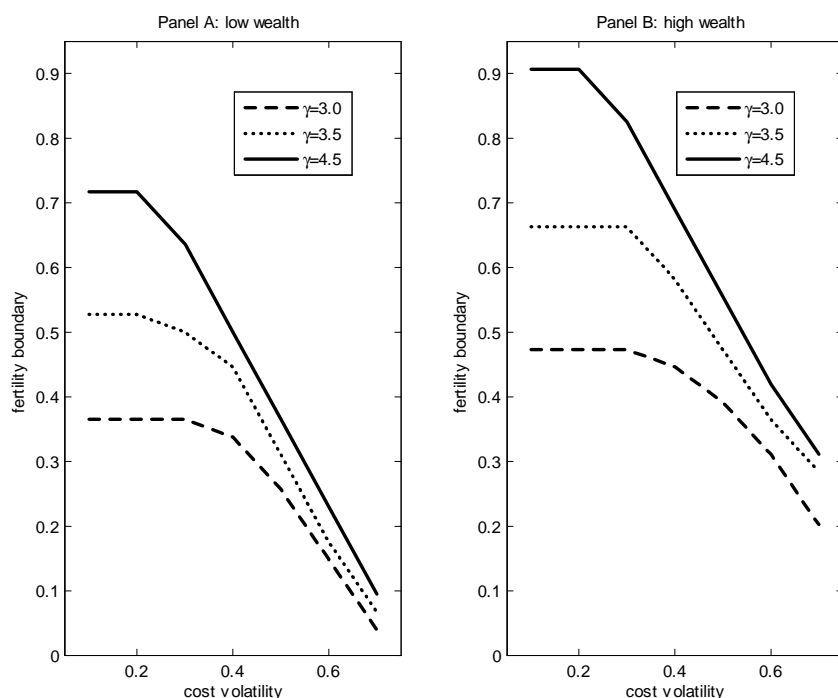


Figure 1: *Fertility as a function of cost volatility.* Panel A reports the fertility boundary (critical cost at which the household has a child) as a function of cost volatility, when household's wealth is the long-term expected income. Panel B reports the same variable when the household's wealth is 2.5 times long-term income. The solid, dashed, and dotted line correspond to a relative risk aversion coefficient γ of 3, 3.5, and 4.5, respectively. The fertility boundary is expressed as a fraction of long-term income. Both panels assume income at 65% of its long-run mean. The other parameters are as in Table 1.

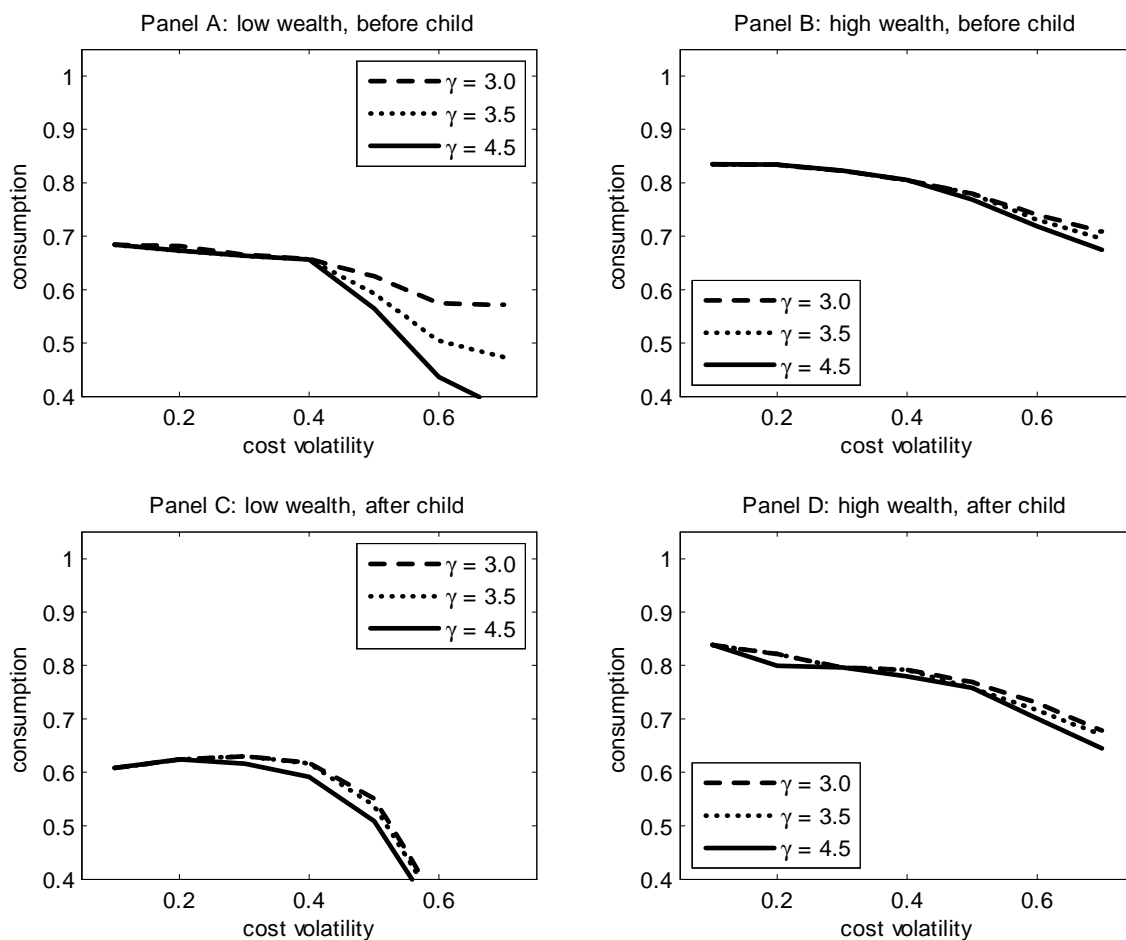


Figure 2: *Average Consumption as a function of cost volatility.* The figure reports average consumption for households who did not have a child yet (Panels A and B), and for those who already had a child (Panels C and D). The average is computed over child cost levels as discussed in the Appendix. In Panels A and C wealth is the income long-run mean, while in Panels B and D it is 2.5 times higher. The solid, dashed, and dotted line correspond to a relative risk aversion coefficient γ of 3, 3.5, and 4.5, respectively. Consumption is expressed as a fraction of long-term income and all panels assume income at 65% of its long-run mean. The other parameters are as in Table 1.




Figure 3a: *Fertility as a function of cost volatility for different levels of α .* In the solid line α is calibrated such that the utility stream from a child equals 25% of the utility stream from consumption of a household with no fertility option. In the dashed line to 30%. Wealth is the income long-run mean in Panel A, and 2.5 times higher in Panel B. The other parameters are as in Table 1.




Figure 3b: *Average consumption as a function of cost volatility for different levels of α .* The figure reports average consumption (as a fraction of income's long-term mean) for households who did not have a child yet. The average is computed over child cost levels as discussed in the Appendix.




Figure 4a: *Fertility as a function of persistence of childrearing cost (b_q)*. The figure reports the fertility decision boundary (critical cost at which the household takes the fertility decision, as a fraction of income long term mean) as a function of the autoregressive parameter of the cost process, b_q , while leaving the long term mean $a_q/(1 - b_q)$ unchanged. The dashed line plots a higher level of wealth than the solid line (70% of income long-run mean vs 60%). Income is at 65% of its long-run mean. The other parameters are as in Table 1




Figure 4b: *Fertility as a function of cost volatility for different levels of cost persistence.* Panel A reports the fertility decision boundary (critical cost, as a fraction of income long-term mean, at which the household has a child) as a function of cost volatility. The solid line has $b_q = 0.4$ while the dashed line $b_q = 0.88$, in both cases the cost long-run mean is unchanged. Wealth is the income long-run mean in Panel A, and 2.5 times higher in Panel B. The other parameters are as in Table 1.

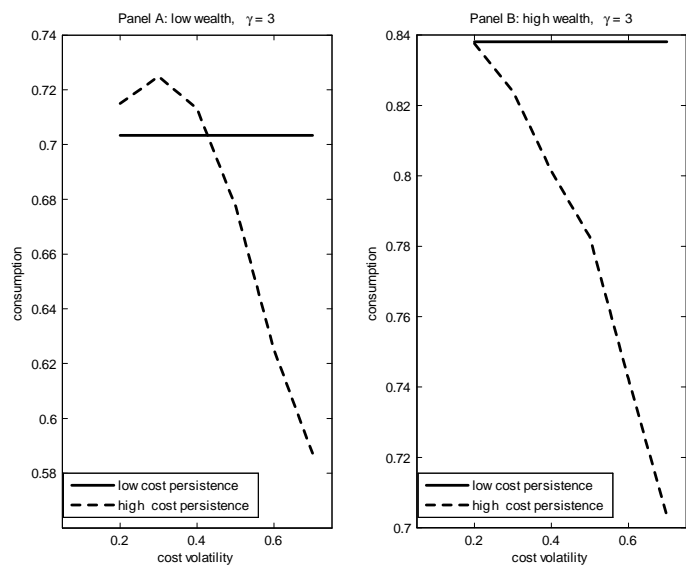


Figure 4c: *Average consumption as a function of cost volatility for different levels of cost persistence.* The figure reports average consumption (as a fraction of income's long-term mean) for households who did not have a child yet. The average is computed over child cost levels as discussed in the Appendix.




Figure 5: *Fertility as a function of the correlation between income and cost shocks.* This Figure reports the fertility decision boundary as a function of the correlation (ρ) between shocks to income and to child cost. Households’s relative risk aversion is $\gamma = 1.6$ in Panel A and 3 in Panel B. Solid lines plot current wealth at the income long-run mean while dashed lines at 2.5 times that value. The income is at its long-run mean. The other parameters are as in Table 1




Figure 6: *Fertility as a function of the volatility of income.* Panel A reports the fertility boundary as a function of volatility of income (σ_y). Households’s relative risk aversion is $\gamma = 1.6$ in Panel A and 3 in Panel B. Solid, dashed and dotted lines correspond to different levels of correlation between income and childrearing costs: $\rho = 0.8, 0,$ and $-0.8,$ respectively. In both panels income is at its long-run mean and wealth is 2.5 times this value. The other parameters are as in Table 1




Figure 7: *Fertility as a function of the volatility of income for lower income persistence.* This figure redoes Panel B of Figure 6 when $b_y = 0.4$ instead of $b_y = 0.88$.




Figure 8: *Fertility as a function of income persistence.* This figure reports the fertility boundary as a function of b_y , without altering the long-term mean of income. The solid line is for income being 65% of its long-term mean and the dashed line for income being 125% of its long-term mean. Wealth is at the income long-term mean. The other parameters are as in Table 1.




Figure 9: *Fertility as a function of the long-term mean of childrearing cost.* The figure reports the effects on the fertility boundary of altering the long-term mean of childrearing costs, while keeping unaltered the autoregressive parameter b_q . Wealth is at the benchmark income long-run mean, while income is 65% of its benchmark long-run mean. The other parameters are as in Table 1.




Figure 10: *Fertility as a function of income, and wealth.* This figure reports the fertility decision boundary as a function of wealth (Panel A) and income (Panel B), both expressed as fractions of income long-run mean. Parameters are as in Table 1.

Appendix. Numerical Method

We discretize the state-space by setting: i) an equispaced grid of $n_W = 50$ realizations for wealth, with lower bound $W_1 = 5$ and upper bound $W_{n_W} = 60$. As discussed in Section 4, these values correspond to 0.5 and 6 times, respectively, the income long-run mean. ii) An equispaced grid of $n_q = 25$ values for costs, with lower bound $q_1 = 0$ and upper bound $q_{n_q} = 9.2$. iii) An equispaced grid of $n_y = 25$ values for income, with upper and lower bounds $y_{n_y} = 12.5$ and $y_1 = 6.5$, respectively.

The realizations of (y, q) on the grids follow a two-state Markov chain whose transition probabilities approximate the transition density implied by the VAR process (3) – (4). They are computed according to the variant of Tauchen (1986) method proposed by Terry and Knotek (2010).

For each realization of y , q , and current wealth on the grids, we consider a grid of possible consumption values such that next period wealth, computed according to the budget constraint (2), takes values on the same grid as current period wealth. In other words, given present income y_z , cost q_v , wealth W_i , and next period wealth W'_j , we back out the implied consumption value from the budget constraint (2), thus obtaining the consumption grid:

$$c_j(W_i, q_v, y_z) = -\frac{W'_j}{1+r} + W_i + y_z - q_v \quad j, i = 1, \dots, n_W; \quad z = 1, \dots, n_y; \quad v = 1, \dots, n_q.$$

This methodology is without loss of generality provided that the grid for wealth is dense and wide enough.¹⁶

We compute the value function $J(W, q, y)$ over the 3-dimensional grid by value function

¹⁶We have compared the predictions of this methodology with those obtained with the methodology described below. Both yield the same results, thus we opted for the first one that is less computationally intensive:

- 1) We set a consumption grid c_j , $j = 1, \dots, n_c$.
- 2) We determine a grid for next period wealth according to the budget constraint:

$$W'_{ijzv} = (W_i - c_j + y_z - q_v)(1+r)$$

enforcing the transversality condition by imposing the wealth lower bound $W'_{ijzv} > W_1$.

- 3) We compute the value function next period corresponding to the off-grid wealth levels by linear interpolation (weighted average of the value function at the closest wealth grid values)

$$J(W'_{ijzv}, q'_m, y'_l) = w_{jzv} J(\underline{W}'_{ijzv}, q'_m, y'_l) + (1 - w_{jzv}) J(\overline{W}'_{ijzv}, q'_m, y'_l)$$

where $\underline{W}'_{ijzv} \equiv \max_j W_j \leq W'_{ijzv}$, $\overline{W}'_{ijzv} \equiv \min_j W_j \geq W'_{ijzv}$ and

$$w^{jzv} = \frac{\overline{W}'_{ijzv} - W'_{ijzv}}{W'_{ijzv} - \underline{W}'_{ijzv}}$$

iteration, iterating the updating rule

$$J^a(W_i, q_v, y_z) = \max_j \left[\frac{\left(-\frac{W'_j}{1+r} + W_i + y_z - q_v\right)^{1-\gamma}}{1-\gamma} - \frac{1}{1-\alpha} + \beta \sum_{l=1}^{n_y} \sum_{m=1}^{n_q} p(q'_m, y'_l | q_v, y_z) J^{a-1}(W'_j, q'_m, y'_l) \right] \quad (16)$$

until the convergence criterion $\max_{i,v,z} \frac{|J^a(W_i, q_v, y_z) - J^{a-1}(W_i, q_v, y_z)|}{J^{a-1}(W_i, q_v, y_z)} < 0.001$ is satisfied. We assign an arbitrary small value to the value function whenever consumption is negative. $p(q'_m, y'_l | q_v, y_z)$ are the one-period transition probabilities of the two-state Markov chain for (q, y) .

The consumption rule of a household with children is

$$c_{j^*}(W_i, q_v, y_z) = -\frac{W'_{j^*}}{1+r} + W_i + y_z - q_v, \quad (17)$$

where j^* is the index achieving the maximum in (16) at the last iteration before convergence.

We compute the value function $H(W, q, y)$ over the 3-dimensional grid by value function iteration of

$$H^a(W_i, q_v, y_z) = \max \left[\max_j \left[\frac{\left(-\frac{W'_j}{1+r} + W_i + y_z\right)^{1-\gamma}}{1-\gamma} + \beta \sum_{l=1}^{n_y} \sum_{m=1}^{n_q} p(q'_m, y'_l | q_v, y_z) J(W'_j, q'_m, y'_l) \right], \right. \\ \left. \max_j \left[\frac{\left(-\frac{W'_j}{1+r} + W_i + y_z\right)^{1-\gamma}}{1-\gamma} + \beta \sum_{l=1}^{n_y} \sum_{m=1}^{n_q} p(q'_m, y'_l | q_v, y_z) H^a(W'_j, q'_m, y'_l) \right] \right], \quad 0 \leq t < T. \quad (18)$$

We use as convergence criterion $\max_{i,v,z} \frac{|H^a(W_i, q_v, y_z) - H^{a-1}(W_i, q_v, y_z)|}{H^{a-1}(W_i, q_v, y_z)} < 0.001$.

The consumption rule of the household without children is

$$c_{\hat{j}}(W_i, q_v, y_z) = \begin{cases} -\frac{W'_{\hat{j}}}{1+r} + W_i + y_z & \text{if } I(W_i, q_v, y_z) = 1 \\ -\frac{W'_{\hat{j}}}{1+r} + W_i + y_z & \text{if } I(W_i, q_v, y_z) = 0 \end{cases} \quad (19)$$

where \hat{j} is the index achieving the maximum in (18) at the last iteration before convergence.

The notation $I(W_i, q_v, y_z)$ emphasizes that the fertility decision depends on the current values of wealth, cost and income on the grids. The fertility boundary, expressed as critical

4) We iterate the updating rule

$$J^a(W_i, q_v, y_z) = \max_j \left[\frac{c_j^{1-\gamma}}{1-\gamma} - \frac{1}{1-\alpha} + \beta \sum_{l=1}^{n_y} \sum_{m=1}^{n_q} p(q'_m, y'_l | q_v, y_z) J^{a-1}(W'_{ijzv}, q'_m, y'_l) \right] \quad (15)$$

level of cost given income and wealth, is:

$$\bar{q}(W_i, y_z) := \max \left\{ q_v : \max_j \left[\frac{\left(-\frac{W'_j}{1+r} + W_i + y_z\right)^{1-\gamma}}{1-\gamma} + \beta \sum_{l=1}^{n_y} \sum_{m=1}^{n_q} p(q'_m, y'_l | q_v, y_z) H(W'_j, q'_m, y'_l) \right] < \right. \\ \left. \max_j \left[\frac{\left(-\frac{W'_j}{1+r} + W_i + y_z\right)^{1-\gamma}}{1-\gamma} + \beta \sum_{l=1}^{n_y} \sum_{m=1}^{n_q} p(q'_m, y'_l | q_v, y_z) J(W'_j, q'_m, y'_l) \right] \right\}, \quad 0 \leq t < T. \quad (20)$$

In the pictures reported in the manuscript, average consumption is computed as the unconditional expectation with respect the stationary marginal distribution of childrearing cost:

$$E [c(W, y, q)] = \int_{\mathbb{R}} c(W, y, q) f_{\infty}(q) dq, \quad (21)$$

where $f_{\infty}(q)$ is the Gaussian density with mean $\frac{a_q}{1-b_q}$ and standard deviation $\frac{\sigma_q}{\sqrt{1-b_q^2}}$.¹⁷ For households with child $c(W, y, q)$ is (17), while for households without child $c(W, y, q)$ is (19).

¹⁷We approximate this integral as $\sum_{j=1}^{n_q} c(W, y, q^j) w^j$, where $w^j = f_{\infty}\left(\frac{q^j + q^{j+1}}{2}\right) / \sum_{j=1}^{n_q-1} f_{\infty}\left(\frac{q^j + q^{j+1}}{2}\right)$, n_q is the number of discretization points on the cost grid.