Saving Behavior around Intended and Unintended Childbirths*

Mette Ejrnæs[†] Thomas H. Jørgensen[‡]

May 12, 2016

Abstract

We investigate how fertility plans affect saving around childbirths. Households who plan on having children should accumulate wealth prior to childbirth to be able to transfer consumption from periods without children to periods with children. Households who do not plan on having (more) children should accumulate significantly less wealth. We estimate a model of fertility and optimal intertemporal consumption in which households have imperfect contraceptive control. In line with the theory, we find that households in British Household Panel Survey (BHPS) who plan to have children, save significantly more than those who do not plan to have children, and they reduce their saving rate when (intended) children arrive. Despite the forward looking behavior of households, the arrival of an unintended child tends to be associated with a smaller decrease in the saving rate, both in the data and in the estimated model. Thus, we can rationalize this saving behavior in the model and, finally, use it to estimate the non-pecuniary cost of an unintended childbirth (JEL: C54, D12, D14, D91, J13).

Keywords: Fertility, Children, Expectations, Unintended children, Savings, Consumption, Life cycle.

*We are grateful for helpful comments and suggestions from Eric French, Mariacristina de Nardi, Monica Costa-Dias, and Cormac O'Dea. We also benefited from comments on a previous version of the paper from seminar participants at the IFS. Jørgensen gratefully acknowledges financial support from the Danish Council for Independent Research in Social Sciences (FSE, grant no. 4091-00040). All errors are our own.

[†]Department of Economics, University of Copenhagen, Øster Farimagsgade 5, building 26, DK-1353 Copenhagen, and Centre for Applied Microeconometrics (CAM). Contact: mette.ejrnes@econ.ku.dk.

[‡]Department of Economics, University of Copenhagen, Øster Farimagsgade 5, building 26, DK-1353 Copenhagen, and Centre for Applied Microeconometrics (CAM). Webpage: www.tjeconomics.com. Contact: thomas.h.jorgensen@econ.ku.dk.

1 Introduction

This paper focuses on the saving behavior of couples around the time of birth of intended and unintended children. The consumption literature – not distinguishing between intended and unintended children – has documented that consumption increases while children are present in the household.[1] Similarly, the literature on female earning dynamics has documented that women's earnings decrease when a child arrives.[2] Thus, if households can perfectly plan their fertility, they should start saving while they are childless. The couple should then have children when they have accumulated enough wealth to smooth marginal utility through increased consumption when children arrive by running down accumulated wealth or borrowing against future income. In such a situation, childbirths will be strictly welfare improving.

Although this is a nice description of the ideal world, it may not be very realistic. In particular, households do not have full control over exactly when and how many children they will have. Some couples may be infertile and therefore not get their desired number of children. Some households may have unintended children either because they have children earlier than they intended or have more than they intended. Recent numbers from the US show that about 16 percent of all childbirths were mistimed and 7 percent of all births were completely unwanted (Mosher, Jones and Abma, 2012). Furthermore, the declining biological fecundity over age also restricts the timing of children, and couples may not be able to postpone childbearing until they have accumulated "enough" wealth. Finally, it may not be possible for households to borrow against future income when children arrive. In this situation, unintended childbirths will be associated with a welfare loss, similar in spirit to an uninsurable permanent income shock.

In this study, we analyze saving behavior around intended and unintended childbirths in a more realistic framework and investigate the non-pecuniary cost associated with an unintended childbirth. To do so, we estimate and calibrate a life cycle model of optimal consumption and saving with imperfect fertility control, using the British Household Panel Survey (BHPS). In particular, households in the model who wish to have children can exert an effort to conceive but will only give birth in the following period with a probability less than one. This implies that even if households intend to have children, they may not ever have them (infertile) or have them in a later period. Finally, households who do *not* wish to have (more) children face a small probability of conceiving a child and subsequently have an *unintended* birth. Both probabilities

[1] Although there is an ongoing discussion of how much of the hump-shape in consumption can be explained by children, most studies agree that children explain part of the hump (see e.g. (Attanasio, Banks, Meghir and Weber, 1999; Gourinchas and Parker, 2002; Fernández-Villaverde and Krueger, 2007; Browning and Ejrnæs, 2009))

[2] See Browning (1992), Waldfogel (1998) and Lundberg and Rose (2000).

(of intended and unintended childbirths) decline with age to emulate the declining biological fecundity. We impose a no-borrowing constraint, model household income as a stochastic process and allow children to have both transitory and permanent effects on household income.[3]

The BHPS data is well-suited for our purpose. A unique feature of the BHPS compared to other longitudinal surveys is that, in some waves, in addition to information on savings, income, and demographics, it contains questions about fertility plans. We utilize these questions to identify (some) unintended and intended births as empirical counterparts to the two types of childbirths in the theoretical model.[4] The longitudinal dimension of the BHPS enables us to follow the same household before and after childbirth and observe how income and saving behavior change.

Simulated data from the estimated model can reproduce the observed fertility and saving behavior around intended and unintended births in the BHPS reasonably well. In particular, saving rates are higher before birth than after and the decrease in the saving rate is largest for intended births. The first part of the result confirms previous findings from the Netherlands (Kalwij, 2005) and the latter confirms our predictions from the model. Also in line with our model, households who plan to remain childless save less than households who plan on having children.

We calculate the non-pecuniary cost of an unintended birth based on our estimated model. For a low skilled household that has an unintended birth, to be as well-off as they would be in the hypothetical scenario with perfect contraceptive control, they should be compensated with a lump-sum of one year's income. The similar compensation for a high skilled household should be around a quarter of their annual income. We find that the cost of an unintended birth for low skilled households is mainly driven by "unwanted" children in that they have too many children, while the cost for high skilled householdis driven by "mistimed" childbirths in the sense that they have children too early. This result indicates that there are considerable costs associated with an unintended birth, especially for low skilled households. We furthermore find that the risk of having unintended children leads to a lower realized fertility than the hypothetical situation with perfect contraceptive control.

While we are, to our knowledge, the first to investigate differences in saving behavior around intended and unintended births, the information on fertility plans in the BHPS has been used in previous studies. Grant (2010) investigates how income and consumption of childless households are affected by the intention to have children. Berrington (2004) uses the information in the BHPS in a descriptive analysis of the

[3]We do not explicitly model the labor supply, but allow household income to depend on the arrival of children. In the paper by Adda, Dustmann and Stevens (forthcoming) the focus is on labor supply, fertility and career cost of children in a dynamic life cycle model.

[4]We use answers to the question pre-childbirth to get the ex-ante expectations (for a detailed discussion of fertility expectations see Rosenzweig and Wolpin, 1993).

intention and fertility in England and Wales.[5] In line with our conclusions, a previous study on data from Peru shows that unintended children have a negative impact on sibling outcomes in contrast to intended children who have no effect (see Lordan and Frijters, 2013). This confirms the idea that parents can allocate resources to intended children while unintended children dilute the resources for the other children.

Our modeling framework is closely related to the model in Sommer (2014), where she investigates the link between fertility patterns and income risk in a life cycle model. She also explicitly models infertility but while we assume that households do not know about their infertility, they do in her model. Two other significant differences between the model in Sommer (2014) and our model are that she does not consider unintended children and we do not model the quality of children. Baudin, de la Croix and Gobbi (2015) also consider infertility, but they focus on the relation between childlessness, education and marriage in a static model. Our model is also related to the one used in Choi (forthcoming), who explicitly models unintended pregnancies and abortions, but does not consider infertility. Choi (forthcoming) uses his model to explain fertility choices and investigates how assets and income influence fertility decisions including abortions. None of the papers mentioned above focus on the savings behavior around births or allow for both infertility and unintended births. Furthermore, they are all silent about the non-pecuniary cost of unintended childbirths.

The paper is organized as follows. In the following section, we present the theoretical model and discuss how we model imperfect fertility control and declining biological fecundity. Section 3 contains a description of the BHPS data and in section 4 we calibrate the model to the data and estimate preferences for children. In section 5, we calculate the non-pecuniary costs of unintended children and discuss the welfare loss associated with imperfect fertility control. Section 6 concludes the paper.

2 Theoretical Framework

The framework used throughout this study is a version of the buffer-stock model pioneered by Deaton (1991) and Carroll (1992) augmented with endogenous fertility. Households work until an exogenously given retirement age, T_r, and die with certainty at age T after consuming all available resources. In all preceding periods, households chose optimal consumption, C_t, and contraceptive effort, e_t, to maximize their expected

[5] Fertility expectations have also been considered in a number of other studies: fertility expectations and employment (see Krishnan, 1988; Kingsbury and Greenwood, 1992 and Stolzenberg and Waite, 1977), fertility expectation and education (see Lynn, Schneider and Zhang, 2013); the stability of fertility expectations (see Hayford, 2009).

discounted future utility. Per-period utility is constant relative risk aversion (CRRA),

$$U(C_t, n_t) = g(n_t; \nu) C_t^{1-\rho} / (1-\rho) + k(n_t; \kappa)$$

with a multiplicative taste shifter, $g(\cdot)$, adjusting the marginal value of consuming C_t as a function of the number of children, n_t, and an additive constant adjusting the level of utility in the presence of children.

Households solve the dynamic optimization problem subject to the intertemporal budget constraint,

$$M_{t+1} = R(M_t - C_t) + Y_{t+1},$$

where R is the gross real interest rate, M_t is resources available for consumption at the beginning of period t and Y_t is the beginning-of-period income. Furthermore, households are not allowed to be net-borrowers, $M_t - C_t \geq 0$, $\forall t$. Below, we describe the model components in detail.

2.1 Fertility and Children

Households can have at most three dependent children, and fertility is imperfectly controlled by households. In each period, households choose the effort with which they try to conceive a child, $e_t \in \{0, 1\}$, where $e_t = 1$ indicates that the household chooses to *try* to conceive a child in the current period. The number of children evolves according to

$$n_{t+1}(e_t, n_t) = n_t + b_{t+1}(e_t, n_t) - l_{t+1}(n_t)$$

,where the event of a child leaving home ($l_{t+1}(n_t) = 1$) is binomial with probability p and n_t "trials" when the household is older than 35. Before that age, children cannot leave the household.

The birth of a new offspring is indicated by $b_{t+1}(e_t, n_t) = 1$ and is imperfectly controlled by exerting the effort e_t,

$$b_{t+1}(e_t, n_t) = \begin{cases} 1 & \text{with probability } \wp_t(e_t, n_t), \\ 0 & \text{with probability } 1 - \wp_t(e_t, n_t), \end{cases}$$

where $\wp_t(e_t, n_t)$ summarizes the age-dependent fecundity of the household along with potential imperfect contraceptive control. Specifically,

$$\wp_t(e_t, n_t) = \begin{cases} 0 & \text{if } n_t = 3, \\ \overline{\wp}_t & \text{if } n_t < 3 \text{ and } e_t = 1, \\ \underline{\wp} \cdot \overline{\wp}_t & \text{if } n_t < 3 \text{ and } e_t = 0, \end{cases}$$

where $\overline{\wp}_t$ denotes the biological fecundity of a household at age t, and by letting $\underline{\wp} > 0$ we allow for imperfect control and unintended childbirths.

2.2 Income Process

While working $(t < T_r)$, income is assumed to follow the stochastic process

$$\begin{aligned} Y_t &= \Pi(n_t, n_{t-1}) P_t \xi_t, \quad \log \eta_t \sim \mathcal{N}(-\sigma_\eta^2/2, \sigma_\eta^2) && (1) \\ P_t &= \Omega(n_t, n_{t-1}) G_t P_{t-1} \eta_t, \quad \log \xi_t \sim \mathcal{N}(-\sigma_\xi^2/2, \sigma_\xi^2), && (2) \end{aligned}$$

where P_t denotes permanent income, G_t is real gross income growth, η is a mean one permanent income shock, and ξ_t is a mean one transitory income shock. When retired, the income process is a deterministic fraction $\varkappa \leq 1$ of permanent income and permanent income is constant once retired, $Y_t = \varkappa P_{T_r}$, $\forall t \geq T_r$.

The career cost of childbirth is captured through

$$\begin{aligned} \Pi(n_t, n_{t-1}) &= 1 + \pi_1 \mathbb{1}_{\{n_t=1, n_{t-1}=0\}} + \pi_2 \mathbb{1}_{\{n_t=2, n_{t-1}=1\}} + \pi_3 \mathbb{1}_{\{n_t=3, n_{t-1}=2\}}, \\ \Omega(n_t, n_{t-1}) &= 1 + \omega_1 \mathbb{1}_{\{n_t=1, n_{t-1}=0\}} + \omega_2 \mathbb{1}_{\{n_t=2, n_{t-1}=1\}} + \omega_3 \mathbb{1}_{\{n_t=3, n_{t-1}=2\}}, \end{aligned}$$

where $\pi = (\pi_1, \pi_2, \pi_3)$ and $\omega = (\omega_1, \omega_2, \omega_3)$ captures the transitory and permanent reduction in household income around the year of the first, second and third childbirth. This is a reduced-form implementation of the career cost of children. We imagine that the parameters in π and ω summarize several underlying mechanisms: Reduced labor market supply of the wife around and perhaps after childbirth and reduced wages.

2.3 Recursive Formulation

The model formulation allows us to normalize everything by permanent income.[6] We denote lower case variables as normalized and denote lower case functions as normalized by $P_t^{1-\rho}$. In the terminal period, $t = T$, households consume all available resources such that $v_T(m_T, n_T) = u(m_T, n_T)$. In preceding periods, after retirement $(T > t \geq T_r)$ all households are infertile and, thus solving the deterministic problem of intertemporal consumption,

$$v_t(m_t, n_t) = \max_{c_t \in (0, m_t]} \{U(c_t, n_t) + \beta v_{t+1}(m_{t+1}, n_{t+1})\}, \ t \geq T_r$$

where β is the discount factor.

[6] Specifically, we assume that $k(n_t; \kappa) P_t^{\rho-1}$ does not depend on permanent income. Since $\rho > 1$ this implies that children are inferior goods, which is consistent with empirical evidence (see e.g. Jones, Schoonbroodt and Tertilt, 2010, and references therein).

In periods in which households are fertile (determined by the biological fecundity), households choose optimal consumption and whether to exert effort to increase the chances of conceiving a child. The discrete choice over effort to conceive a child and the continuous consumption and savings choice can be partitioned into

$$v_t(m_t, n_t, \varepsilon_t) = \max\{v_t(m_t, n_t | e_t = 0) + \sigma_\epsilon \epsilon_t(0), v_t(m_t, n_t | e_t = 1) + \sigma_\epsilon \epsilon_t(1)\}$$

where *iid* effort-specific taste shocks, ϵ_t, also affect the level of utility and σ_ϵ^2 is proportional to the variance of these taste shocks. The effort-specific value functions are given as

$$\begin{aligned} v_t(m_t, n_t | e_t) &= \max_{c_t \in (0, m_t]} \{U(c_t, n_t) + \beta \mathbb{E}_t[\wp_t(e_t, n_t) \mathcal{G}_{t+1}^{1-\rho} V_{t+1}(m_{t+1}, n_{t+1}) \\ &\qquad + (1 - \wp_t(e_t, n_t)) \mathcal{G}_{t+1}^{1-\rho} v_{t+1}(m_{t+1}, n_{t+1})]\}, \end{aligned}$$

where

$$\mathcal{G}_{t+1} = \mathcal{G}_{t+1}(n_{t+1}, n_t) \equiv G_{t+1} \eta_{t+1} \Omega(n_{t+1}, n_t)$$

is a normalization factor stemming from the fact that we have normalized with permanent income.

Assuming that the choice-specific taste shock, ϵ, is extreme value type I distributed, the expected value function – with respect to the unobserved taste shocks – is given by the log-sum,

$$EV_t(m_t, n_t) = \sigma_\epsilon \log\left[\exp\left(v_t(m_t, n_t | e_t = 0)/\sigma_\epsilon\right) + \exp\left(v_t(m_t, n_t | e_t = 1)/\sigma_\epsilon\right)\right]$$

and the probability of a household trying to conceive a child is given as

$$\Pr(e_t = 1 | m_t, n_t) = \frac{\exp\left(v_t(m_t, n_t | e_t = 1)/\sigma_\epsilon\right)}{\exp\left(v_t(m_t, n_t | e_t = 0)/\sigma_\epsilon\right) + \exp\left(v_t(m_t, n_t | e_t = 1)/\sigma_\epsilon\right)}.$$

3 Data and Descriptive Analysis

We use the British Household Panel Survey (BHPS), which is a completed panel of 18 waves collected from 1991 through to 2009. We use this data set because it contains unique information on fertility plans. The BHPS consists of a household survey and an individual survey including a variety of questions. Some questions are asked in each wave, such as saving, income and household composition, while other questions, such as wealth and plans regarding having additional children in the future, are only asked in some waves. We combine these questions to analyze how plans about childbearing affect savings behavior in anticipation of having children. Given the fact that fertility plans are only reported in six waves we do not directly use this information in the estimation. Instead we use this information to test and validate our model. Therefore, we construct two samples based on the BHPS: a sample used for estimation and a sample used for validation.

3.1 Estimation sample

Sample Selection Criteria. We construct our sample by selecting cohabiting or married couples.[7] We restrict our sample to couples where one of the spouses is the reference person and delete same sex couples. Furthermore, we restrict the sample to couples where the women are aged between 20 to 62 years old. For the estimation sample, we impose two more restrictions to insure that we observe the start of the fertility process. We impose that the couple has to be childless the first time we observe them and that the women should be less than 35 years old the first time the couple is observed. The estimation sample contains 1704 households and 11,653 observations from 18 waves. In all the analyses, we split the sample according to the educational attainment of the husband. The sample contains 1301 low skilled households and 403 high skilled households. The main variables used for the estimations are demographic, income and saving.

Demographic and background variables. We define the number of children as the number of own biological children living in the household. High skilled individuals are defined as those who have completed the first or second stage of tertiary education (ISCED codes 5 or 6). Individuals responding that their race is "white" are classified as white, and married individuals are those who state that they are in a legal marriage. The age is defined as the age of the wife.

Wealth and Savings. In waves 5, 10 and 15, individuals are asked about their wealth stock. In all waves individuals are asked if and how much of their monthly income they typically save. In the estimation we have trimmed the data such that saving rates exceeding 100% have been trimmed. For the wealth data we have trimmed the data at the 1 and 99 percentiles.

Income. Individuals are asked about their most recent labor market gross pay and how long a period that pay covered. If their most recent pay was unusual, they are asked to report their usual pay. Based on the recent/usual pay and the reported weeks

[7] We refer to the woman and man in the household as the wife and husband throughout although not all couples are married.

that pay covered, annual gross labor market income is calculated. If only after tax net-income is reported, we scaled it with 40 percent to proxy for gross pay.

(a) Average number of children (b) Average income (c) Average amount saved

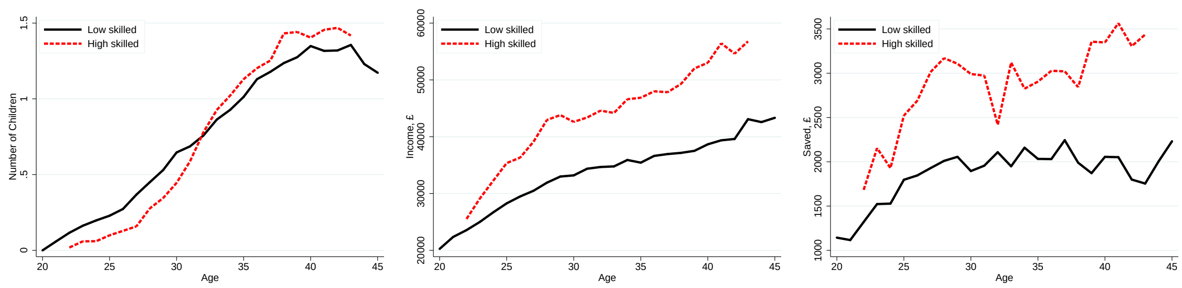

Figure 1 – Life Cycle Patterns for Children, Income and Saving.

Notes: Figure 1 presents the age profile of the average number of children, income and amount saved. Only cells with at least 50 observations are presented.

Figure 1a illustrates the age profile of the average number of children present in the household. The figure shows that the average number of children reaches its maximum of 1.5 children round the age of 40. The fertility in the BHPS sample might be lower than that for the UK population. Although the figure does not illustrate completed fertility, an average of 1.5 (own) children in the household at age 35 is lower than the completed fertility rates of 1.9 for the 1967 birth cohort (aged 24 in 1991), reported in Office for National Statistics (2013). The 1967 cohort had around 1.2 children at the age of 30, close to the average number of one child in BHPS households at age 30. For our sample the number of children is somewhat lower because we have younger cohorts and deselect couples who have children before the age of 20. Figure 1b shows that the average income for these households is increasing in age, especially at young ages. Furthermore, we see the income gap between low and high skilled is increasing over the life-cycle.

To further investigate how income and saving change around birth, we focus on the first birth. We have organized the observations according to the period of the first birth, as done in Kalwij (2005). To control for age and year effects we run a regression with age and wave dummies and report the average income and saving around the first birth in Table 1. As in Kalwij (2005), we find that the saving rate drops after the birth of the first child. We also find that income decreases after the birth of the first child.

Table 1 – Income and Savings around First Childbirth.

	Children	Income (£)	Savers (pct)	Annual saved (£)	Savings Rate
-7	0.0	35578.5	77.0	3619.4	9.1
-6	0.0	35516.3	77.4	3230.7	8.0
-5	0.0	36285.4	73.8	3353.9	8.2
-4	0.0	36270.8	73.7	3364.7	8.8
-3	0.0	36929.4	75.3	3516.3	9.0
-2	0.0	36463.5	74.9	3493.6	9.1
-1	0.0	35636.4	71.5	3507.2	9.0
0	1.0	33541.3	64.3	2397.1	6.3
1	1.1	29945.9	63.8	2032.9	6.2
2	1.3	29396.2	64.0	2068.2	6.1
3	1.6	28484.3	60.3	1855.9	5.8
4	1.8	28070.7	61.0	2011.6	6.1
5	1.9	27451.1	57.4	1781.5	5.8
6	2.0	26246.1	55.7	1418.7	5.6
7	2.0	25664.3	55.5	1602.9	6.4

Notes: Estimation is based on the households in the estimation sample where we observe the first birth. Age, wave, sex and race dummies are included in all regressions. The reference group is a non-white female aged 30 in wave 8.

3.2 Validation Sample: Fertility Plans

A unique feature of the BHPS is that households have been asked about their fertility plans in some waves.[8] This, combined with the fact that we can follow the same households over time, allows us to investigate whether households' behavior before and after childbirths depends on their fertility plans. In particular, the data allow us to examine how saving behavior differs around a childbirth depending on whether the birth was intended or unintended. We do not directly use this information in the estimation but use it instead for validation of the model, because the questions on fertility plans are not asked in every wave, and therefore, we would have a rather small estimation sample. To have a fairly large data set for validation we do not restrict the validation sample to households observed as childless before age 35. This sample contains 4459 households in which the wife is younger than 45 and, at some point we know their fertility plans. For a subset consisting of 731 households we observe a birth after the couple have stated their fertility plans.

Fertility Plans. Individuals aged 15-45 are asked in waves 2, 8, 11, 12, 13 and 17, whether they plan to have (more) children.[9] Specifically, respondents are asked whether

[8]The exact question reads *"Thinking about your plans for the future do you think you will have any (more) children?"*.

[9]Men are asked about their fertility plans until age 65.

(a) Across parity

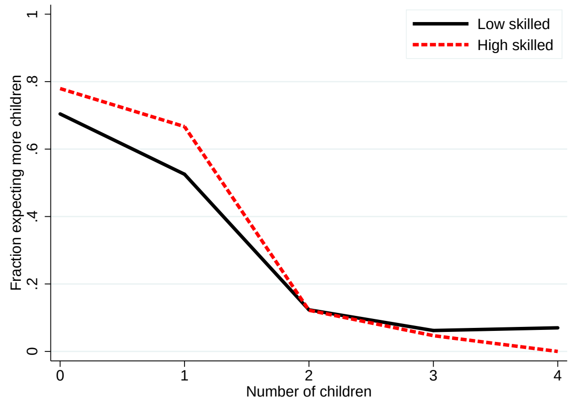

(b) Across age groups

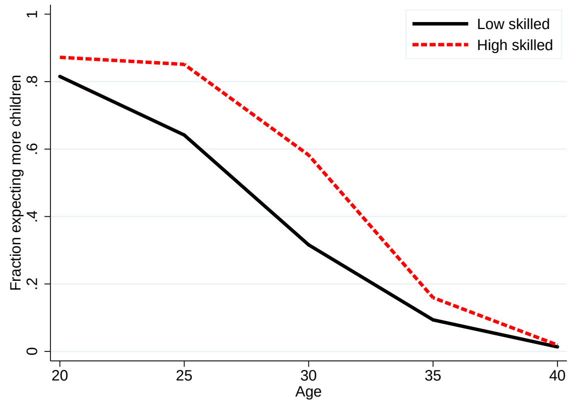

Figure 2 – Fertility Plans.

Notes: Figure 2 presents the fraction planning to have additional children as a function of existing children and age. Only cells with at least 50 observations are presented. This graph is based on households in which the wife is aged 20-40.

they think they will have any (more) children and if so, how many. A significant percentage (85 percent) of respondents are classified as "inapplicable" in wave 11 compared to a stable 42 percent of inapplicable observations in waves 2, 8, 12, 13 and 17. Thus, throughout the analysis we discard the question in wave 11 and only use the fertility plans from the other five waves.

First we show how fertility plans vary with number of (existing) children and age of the women. In Figure 2 we see that between 70-80 percent of the childless households plan to have children.[10] We also see that the high skilled are more likely to plan to have children. For households *with* children, the fraction that plan to have additional children is much lower. This is particularly true for households who already have two or more children. The figure thus reflects a very strong two-children norm, since for households with two children only about 10 percent plan to have more children. We also see that the share that is planning to have more children declines with age. The decline is strongest for the low skilled, which is consistent with the fact that they have children earlier.

To show that the question on fertility plans actually has strong predictive power on future *realized* fertility, we have divided the sample according to the fertility plans and then investigate whether households who state they want more children also have more children.[11] In figure 3 we find that the average number of additional children is much higher for households stating that they planned to have children. For households

[10] These results are in line with Berrington (2004).

[11] Again, this confirms the pattern found in Berrington (2004).

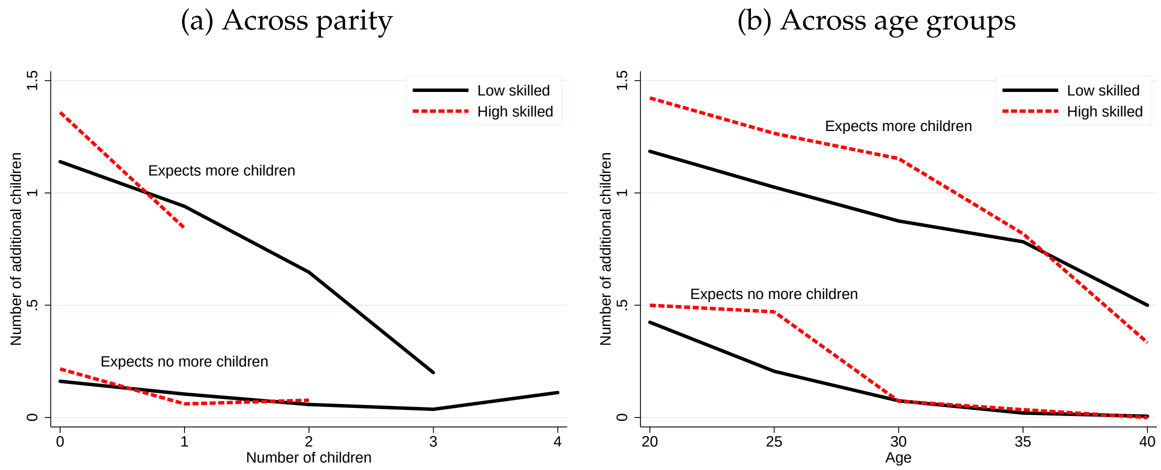

Figure 3 – Realized Fertility by Fertility Plans.

Notes: Figure 3 presents the fraction planning on having additional children as a function of existing children and age. Only cells with at least 50 observations are presented. This graph is based on households in which the wife is aged 20-40.

that are not planning to have more children, the additional number of children is low. However, for households aged between 20 and 25 a considerable share of those stating that they do not plan to have children actually do have children at a later stage, whereas of those who are above 30 and state that they do not intend to have more children, few will have additional children. Based on these data, we can also quantify the share of unintended births. We define unintended births as births that occur after the households have stated that they do not plan to have more children. The data indicate that about 15 percent of the births for low skilled households and 10 percent of the births for high skilled households are unintended. A study by Fleissig (1991) finds that the fraction of unintended births for married women in England and Wales was about 22 percent in 1989. Our numbers are slightly lower. One reason for the discrepancies could be that we do not have information about the expectations in each wave, and if households change their fertility plans between when the question was asked and the birth, our measure is imprecise. Furthermore, our sample selection also oversamples more stable couples, which could explain the lower number of unintended births. When looking at the pattern of unintended births for this sample, we see that unintended births are more likely for higher parity births. For first born children, only 3 and 6 percent of the births were unintended. This confirms the pattern also found for American couples (see Mosher, Jones and Abma, 2012).

As part of the validation exercise we will consider how saving is affected by intended and unintended births. In the following we take a closer look at how saving changes upon the arrival of a child, and we run a series of regression models. To increase the

Table 2 – Saving Rate around First Birth – Estimation Results.

	OLS			Censored regression (Tobit)		
	RE	FE	FE	RE	FE	FE
Number of children	-0.031***	-0.027***	-0.020**	-0.042***	-0.045***	-0.036***
	(0.004)	(0.008)	(0.006)	(0.005)	(0.012)	(0.010)
Unintended child	0.037	0.033	0.026	0.059*	0.055	0.047
	(0.020)	(0.021)	(0.021)	(0.027)	(0.036)	(0.026)
Married	0.014	0.018	0.012	0.026*	0.029	0.021
	(0.007)	(0.016)	(0.012)	(0.010)	(0.018)	(0.024)
Home owner	-0.010	-0.013	-0.012	-0.011	-0.018	-0.026
	(0.008)	(0.014)	(0.012)	(0.011)	(0.044)	(0.038)
Plans to be childless	-0.039			-0.067*		
	(0.020)			(0.029)		
High skilled, husband	-0.004			-0.002		
	(0.007)			(0.010)		
High skilled, wife	0.018*			0.027**		
	(0.007)			(0.010)		
White	-0.037*			-0.039		
	(0.017)			(0.023)		
Age, husband	0.000			0.001		
	(0.001)			(0.001)		
Constant	0.088*	0.076	0.088***	0.045		
	(0.043)	(0.046)	(0.020)	(0.059)		
Age dummies	Yes	No	Yes	Yes	No	Yes
Wave dummies	Yes	Yes	Yes	Yes	Yes	Yes
Obs	1274	1274	2169	1274	1274	2169

Notes: Estimation is based on households in which the wife is aged 20-39.

number of observations we include both educational groups in the regressions. We are, particularly, interested in how couples who did not plan to have a child are affected by a childbirth. To do this we select a sample of households where we observe at least one period before they have a birth and where we know their fertility plans. We focus on three periods around the childbirth. In Table 2 we estimate the impact of childbirth on saving rates. The first three columns refer to a linear regression model and the last three columns to a Tobit model. Columns 1 and 4 are random effect estimations while columns 2, 3 ,5 and 6 are fixed effect estimations.[12] The overall picture suggests that saving rates decrease with 2-4 percentage points when a child arrives, but the saving rate *increases* by up to 1.5 percentage points if the child was unintended. The regressions also show that those who do not plan to have children have 4-7 percentage points lower saving rates than those who intend to have children. If we instead use a dummy for

[12]The fixed effect Tobit model is estimated using the method in Honoré (1992).

saving as the dependent variable, we see the same pattern, namely that the fraction that saves drops after child birth, but if the child was unintended, the fraction actually goes up in some specifications (see Table A.1 in the online supplemental material).

4 Calibration and Estimation of the Model

We estimate or calibrate some of the model parameters and keep them fixed when estimating the preferences for children. Table 3 contains all these model parameters along with their fixed values and sources. Parameters governing the income process are not included in Table 3 but are deferred to Tables 4 and 5.

Table 3 – Calibrated Parameters. Excluding the Income Process.

	Parameter	Value	Source
R	Gross real interest rate	1.03	Gourinchas and Parker (2002)
β	Discount factor	0.97	Sefton, Van De Ven and Weale (2008)
T_r	Retirement age	62	Around the state retirement age of men and women in the UK
$\varkappa$	Replacement rate in retirement	0.7	Banks, O'Dea and Oldfield (2010)
$\overline{\wp}_t$	Biological fecundity	fig. 4	Trussell and Wilson (1985)
$\underline{\wp}$	Risk of unintended pregnancy	0.017	Own calculations, see text.
p	Probability of children moving	0.08	Own calculations, see text.
σ_ϵ	Variance of taste shocks	0.5	Own calculations, see text.

Preferences. Following Gourinchas and Parker (2002), we fix the real interest rate to three percent, $R = 1.03$. The discount factor, β, is set at $0.97 \approx 1/R$. This value of discounting is also preferred in Sefton, Van De Ven and Weale (2008) when calibrating a model of savings for retirement to UK survey data in the FES.

Fecundity and Contraceptive Control. Following Sommer (2014) we use the estimated biological infertility reported in Trussell and Wilson (1985). Fitting an exponential in age provides a good fit to their reported infertility rates over the life cycle, as reported in panel (a) of Figure 4. The resulting biological fecundity, $\overline{\wp}_t$, is illustrated in the right panel of the figure.

The likelihood that contraception fails is calibrated to a percentage of the biological fecundity, $\underline{\wp} = 0.017$. To calculate the probability $\underline{\wp}$, we use our validation sample which contains information about the fertility plans for a subset of respondents in the BHPS. We deflate the probability of an unintended birth by the measure of fecundity (see details in the online supplemental material).

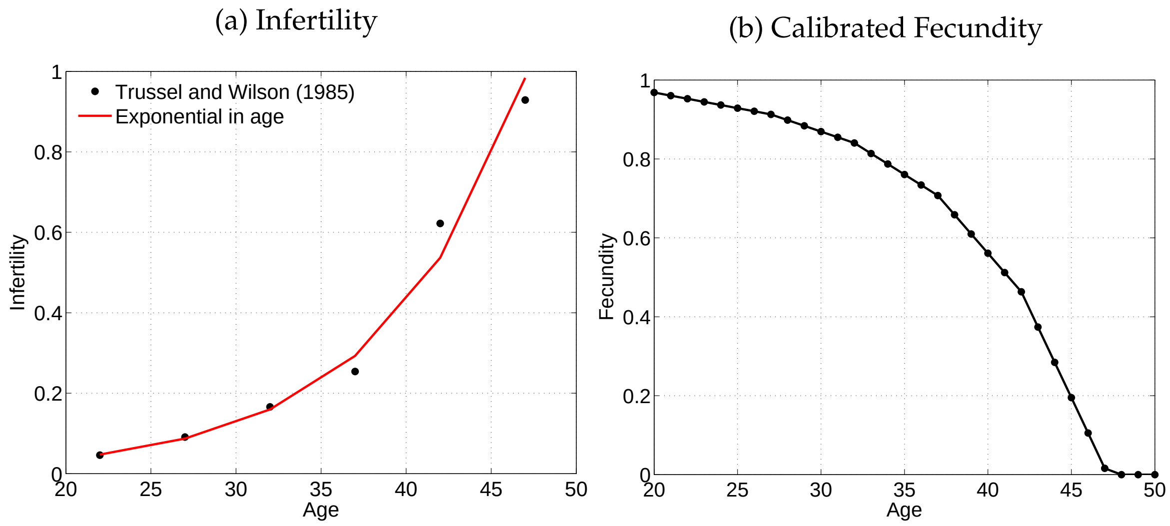

Figure 4 – Calibrated Biological Fecundity, $\overline{\wp}_t$.

Notes: The left panel reports the estimated infertility due solely to the process of aging, reported in Trussell and Wilson (1985, Table 8) (black dots) along with an exponential-in-age model of infertility (red line). The right panel illustrates the resulting biological fecundity, $\overline{\wp}_t$.

At age 35 children can leave the household. The likelihood that a child moves follows a binomial distribution with the likelihood of a child moving calibrated to $p = 0.08$ based on the observed likelihood of children moving in the BHPS.

The Income Process. Table 4 reports estimates from the regression

$$\log Y_{it} = \phi_{0i} + \phi_{1w} + \phi_{2t} + \sum_{j=1}^{3} \omega_j \mathbb{1}_{\{n_{it}=j\}} + \sum_{j=1}^{3} \pi_j \mathbb{1}_{\{n_{it}=j, n_{it-1}=j-1\}} + \psi_{it}, \tag{3}$$

where the first three elements are household, wave and age dummies and the ω and π parameters contain the permanent and transitory effect of children, respectively.[13] The columns labelled (1) report estimates when allowing all three children to affect permanent income growth differently, while the preferred specification in the columns labelled (2) restrict $\omega_3 = \pi_2 = \pi_3 = 0$ such that only the first two children have a permanent effect on income and only the first child has a transitory effect. This preferred parsimonious income process seems to provide a reasonable fit to the observed income changes around childbirth (see figure A.1 in the supplemental material). We estimate that the first child permanently decreases the income by 5 percent for low skilled, and 11 percent for high skilled households. These effects are slightly lower than what is found in Sigle-Rushton and Waldfogel (2007).

The transitory and permanent income variances (σ_ξ^2, σ_η^2) reported in Table 5 are

[13] The estimated income process in (3) is an approximation of the income process in equations (1)–(2).

Table 4 – Career cost of childbirth.

	Low skilled		High skilled	
	(1)	(2)	(1)	(2)
ω_1	-0.048***	-0.050***	-0.113***	-0.111***
	(0.008)	(0.008)	(0.017)	(0.017)
ω_2	-0.046***	-0.040***	-0.019	-0.027
	(0.008)	(0.007)	(0.017)	(0.015)
ω_3	0.014		-0.046	
	(0.011)		(0.026)	
π_1	0.032***	0.033***	0.082***	0.080***
	(0.009)	(0.009)	(0.019)	(0.019)
π_2	0.011		-0.006	
	(0.009)		(0.020)	
π_3	-0.030*		0.014	
	(0.014)		(0.036)	
Constant	Yes	Yes	Yes	Yes
Wave dummies	Yes	Yes	Yes	Yes
Age dummies	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes
R^2	0.242	0.242	0.302	0.301
Obs.	29510	29510	5842	5842

Notes: Estimates are based on the specification in equation (3). Robust standard errors in brackets.

estimated using the methodology proposed in Meghir and Pistaferri (2004) on residuals from the log-income regression in equation (3). The estimates are similar to those typically estimated using US data.[14]

The income growth rate is based on the estimated ϕ_{2t} parameters in equation (3) and reported in Figure A.2 in the supplemental material. To allow for gross income growth over time, we add an annual growth of two percent. Finally, we assume that the first observed income level equals permanent income.

Retirement. Since we focus on individuals in their fertile years, the assumptions about the retirement parameters, $(T_r, \varkappa, \gamma)$, have only minor implications. We fix the exogenous retirement age to $T_r = 62$, in between the State Pension age for men (65) and women (60) to proxy the average retirement age *expectations* in the sample. We fix the replacement rate to 70 percent, $\varkappa = 0.7$, based on the median replacement rate in Banks, O'Dea and Oldfield (2010) from the English Longitudinal Study of Aging (ELSA).

[14] Blundell, Pistaferri and Preston (2008) report $\sigma^2_\eta \in [.0057, 0.0333]$ and $\sigma^2_\xi \in [.0190, 0.0753]$ depending on the combination of year, cohort and educational background, using the PSID. Gourinchas and Parker (2002), also using the PSID, estimate $\hat{\sigma}^2_\eta = 0.0212$, $\hat{\sigma}^2_\xi = 0.0440$.

Table 5 – Income Shock Variances.

	Low skilled	High skilled
σ_ξ^2 (Transitory)	0.016	0.022
	(0.004)	(0.005)
σ_η^2 (Permanent)	0.016	0.016
	(0.007)	(0.008)

Notes: Estimates are based on the approach in Meghir and Pistaferri (2004). Robust standard errors in parenthesis.

Taste Shocks. The variance of the Extreme Value Type I effort-specific taste shocks, σ_ϵ, is calibrated using a profiling strategy. Specifically, we estimate the remaining preferences, as described below, for a discrete grid of σ_ϵ between zero and one. The resulting estimates are included in Table A.2 in the online supplemental material. For each set of estimates, we evaluate the fit of the model, the criteria function and how well the model replicates the share of unintended childbirths. By this approach we found that fixing $\sigma_\epsilon = 0.5$ produces reasonable results.

4.1 Preferences for Children

We assume the following functional forms related to the value of children through the multiplicative taste shifter and additive constant,

$$\begin{aligned} g(n_t;\nu) &= 1+\nu n_t \\ k(n_t;\nu) &= \sum_{j=1}^{3}\kappa_j \mathbb{1}_{\{n_t=j\}} \end{aligned}$$

where $g(\cdot)$ is closely related to a common choice of $\exp(\nu n_t)$ but does not imply an exponentially increasing effect on the marginal utility of consumption from additional children. κ_j is the value of having j children compared to none and is thus the *cumulative* effect of children.

We estimate the remaining preference parameters, $\theta = (\rho, \kappa_1, \kappa_2, \kappa_3, \nu)$, by indirect inference estimation (Smith, 1993; Gouriéroux, Monfort and Renault, 1993) while fixing all the other calibrated parameters. Let λ be a $K \times 1$ vector of auxiliary parameters (AP) calculated based on observed data. Denote as $\overline{\lambda}(\theta)$ the average of the same K APs calculated from S simulated data sets from the calibrated life cycle model discussed above, for a given value of θ. For each value of θ, we solve the model using the method proposed in Iskhakov, Jørgensen, Rust and Schjerning (2015) and described in the online supplemental material. We estimate θ as the minimizer of the weighted squared

Table 6 – Estimated Preferences.

		Low skilled	High skilled
ρ	Risk aversion	1.117 (0.000)	1.299 (0.001)
κ_1	Utility value of 1 child	1.193 (0.001)	0.563 (0.001)
κ_2	Utility value of 2 children	1.256 (0.000)	0.554 (0.002)
κ_3	Utility value of 3 children	0.057 (0.005)	0.417 (0.001)
ν	Effect of children on marg. util.	0.136 (0.000)	0.161 (0.000)
Households		1301	403
Obj.		421.0	317.2

Notes: Asymptotic standard errors reported in brackets.

difference between the observed and simulated APs,

$$\hat{\theta} = \arg\min_{\theta \in \Theta} (\lambda - \overline{\lambda}(\theta))' W^{-1} (\lambda - \overline{\lambda}(\theta)),$$

where W is a diagonal $K \times K$ weighting matrix. The online supplemental material describes in detail how the estimator is implemented.

We use the share of savers for the age groups from 22 through to 45 (24 APs) to identify the constant relative risk aversion, ρ. Because ρ greatly affects the desire to smooth out consumption, the share that saves will be informative about the value of this parameter. To identify the value of children we use the average age profile of the number of children born to women aged 22 through to 40 (19 APs). The timing and number of children is largely controlled by the value of children in the model, κ, and the average age profile of the arrival of children is, thus, informative about these parameters. We use the change in saving rates at the arrival of the first, second, and third child (3 APs) to primarily identify ν. Because ν affects the marginal value of consumption, in the model it controls how much consumption and thus savings change at the arrival of children. In total, this yields $K = 46$ moments in the BHPS, which we estimate the five parameters in θ from. The estimated preferences are reported in Table 6.

The estimated constant relative risk aversion, ρ, is close to what is typically found in the literature. Particularly, Attanasio, Banks, Meghir and Weber (1999), using the Consumer Expenditure Survey (CEX), estimate a constant relative risk aversion of around 1.5 and Attanasio and Weber (1993) estimate ρ to be around 1.3 using the The British Family Expenditure Survey (FES). The effect of children on the marginal utility of consumption (ν) is slightly lower than the estimate in Attanasio, Banks, Meghir and Weber (1999).

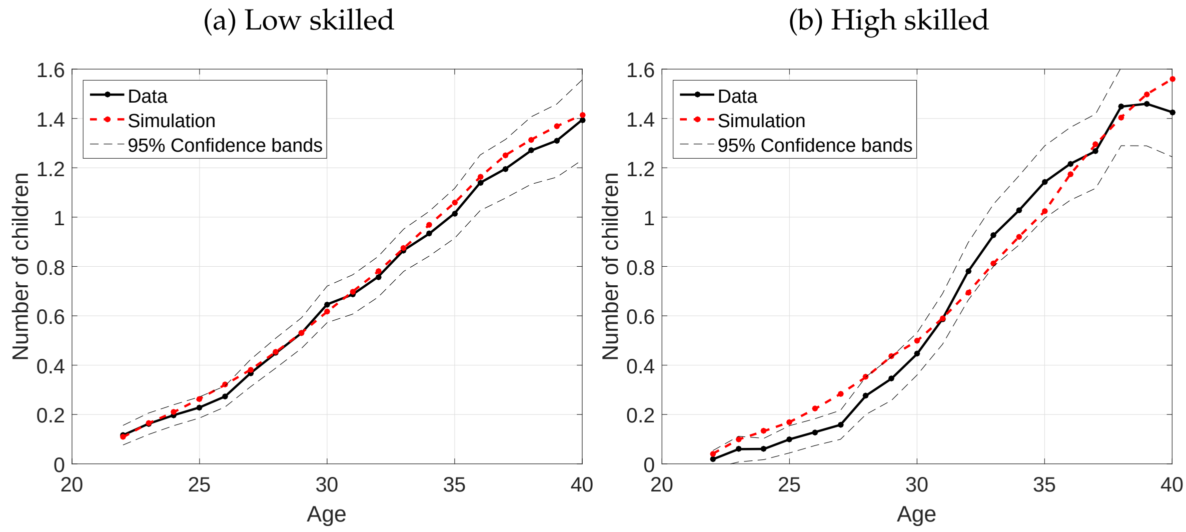

Figure 5 – Actual and Simulated Age Profile of the Number of Children in the Household.

We find that the value of additional children is decreasing in the number of existing children. The estimated value of children, κ, confirms the very strong two-children norm observed in our sample of the BHPS. Particularly for low skilled households.

Figure 5 illustrates the average number of children born in the estimation sample together with the simulated age profile based on the estimated model. The figures show that, especially for low skilled households, the model fits the number of childbirths well. For high skilled households, the model slightly over-predicts the number of births early in the life cycle.

4.2 External Validation: Unintended Childbirths

We investigate how well the estimated model of consumption, saving and fertility match saving behavior around intended and unintended childbirths observed in the BHPS. Especially, we use the validation sample, discussed above, for which we have information about fertility plans regarding having (more) children. All these moments have not been used to estimate the model and, therefore, we see this as an external validation of the model.[15]

Table 7 reports the share of unintended births across parities. The overall pattern in the simulated data is close to that found in the BHPS. For low skilled, however, the estimated model predicts a larger share of the third child being unintended than what we see in the data. The opposite is the case for high skilled.

[15]The simulation strategy is identical to the one employed when estimating the preference parameters. See also the online supplemental material.

Table 7 – Simulated Share of Unintended Childbirths.

	Low skilled		High skilled	
	Simulation	Data	Simulation	Data
1. child	0.074	0.062	0.089	0.028
		[0.036;0.087]		[0.003;0.053]
2. child	0.096	0.115	0.049	0.086
		[0.079;0.151]		[0.023;0.149]
3. child	0.730	0.435	0.085	0.516
		[0.353;0.516]		[0.320;0.702]

Notes: The share of unintended children simulated from the estimated model and observed in the BHPS in our validation sample. 95% confidence intervals are reported in square brackets.

Table 8 – Saving Rate Growth around Childbirths.

	Simulated		Data
	Low skilled	High skilled	BHPS
$\Delta NumChildren$	-0.062	-0.055	[-0.045;-0.020]
$\Delta NumChildren \cdot Unintended$	0.037	0.025	[0.026;0.059]
Constant	Yes	Yes	Yes
Age dummies	Yes	Yes	Yes

Notes: The estimates from BHPS (our validation sample) are obtained from several regression models, see table 2.

Table 8 reports the correlation between saving rate growth and the event of a childbirth. We distinguish between intended and unintended childbirth and find that the BHPS households tend to decrease saving when children arrive but if the pregnancy was unintended, the reduction in saving is much lower. The same pattern is replicated in the simulated data based on the estimated model.

To sum up, the validation exercise suggests that our model can reproduce the composition of unintended children at different parities and the saving behavior around first birth for unintended children. This is encouraging since these data moments have not been used in the estimation, thus giving our model some external validity.

4.3 Robustness

Arguably, our results might be subject to the concern that the risk of contraception failing, $\underline{\wp}$, is calibrated and fixed at 0.017. While we have made efforts to calibrate this unknown parameter using the biological fecundity reported in Trussell and Wilson (1985) together with the observed frequency of unintended births in the BHPS, here we investigate how the results depend on this key parameter.

In Table 9, we report estimates of the model parameters when reducing $\underline{\wp}$ to 0.005 and increasing it to 0.03. The results show that estimates of the risk aversion, ρ, are almost unaffected. For the low skilled households, we find that that the estimates of the utility of children, κ, vary with $\underline{\wp}$. When $\underline{\wp}$ decreases, the estimates of the utility of having three children, κ_3, increases and when $\underline{\wp}$ increases the estimates of the utility of having two (and three) children, κ_2, decrease. This implies that the share of unintended births also varies with $\underline{\wp}$.[16] Especially, for $\underline{\wp} = 0.005$, the share of unintended births is as low as 2.9 percent for low skilled households, while $\underline{\wp} = 0.03$ leads to 22.3 percent unintended births. For the high skilled, the parameter estimates of the utility of having children depend less on $\underline{\wp}$, but we still find that the share of unintended childbirths is as low as 2.2 percent for $\underline{\wp} = 0.005$ and 13.7 percent for $\underline{\wp} = 0.03$. This suggests that in order to calibrate $\underline{\wp}$ it is crucial to have information on the share of unintended births and that our calibrated value of $\underline{\wp}$ at 0.017 gives a reasonably good fit of the share of the unintended births for both educational groups.

Table 9 – Estimated Preferences, Various $\underline{\wp}$.

	Low skilled		High skilled	
	(1)	(2)	(1)	(2)
$\underline{\wp}$	0.005	0.030	0.005	0.030
ρ	1.166	1.105	1.356	1.307
	(0.001)	(0.001)	(0.002)	(0.001)
κ_1	0.864	1.738	0.338	0.530
	(0.001)	(0.001)	(0.002)	(0.002)
κ_2	0.886	0.014	0.307	0.586
	(0.002)	(0.002)	(0.001)	(0.002)
κ_3	0.207	0.176	0.231	0.399
	(0.000)	(0.000)	(0.000)	(0.000)
ν	0.139	0.136	0.108	0.175
	(0.000)	(0.000)	(0.001)	(0.001)
Obj.	450.2	403.0	336.1	303.7
Unintended (pct)	2.9	22.3	2.2	13.7
1st child.	2.2	9.3	2.7	18.7
2nd child	2.8	97.2	1.5	7.2
3nd child	18.1	80.3	2.2	15.0

Notes: Table 9 reports estimated preferences for various values of the risk of unintended pregnancies, $\underline{\wp}$. Asymptotic standard errors reported in brackets.

[16] In the estimation, we fit to the number of children, implying that the number of children remains almost constant across the calibrations.

5 The Non-Pecuniary Cost of Unintended Childbirths

Using the estimated parameters, we can evaluate the costs associated with unintended childbirths from imperfect contraceptive control. To evaluate how much better off households would be if they had *perfect contraceptive control*, we simulate data from the estimated model and compare the expected discounted utility to data simulated from a model in which households have perfect contraceptive control ($\underline{\wp} = 0$). All other parameters are fixed at their calibrated values, (reported in Section 4).

Figure 6 illustrates the simulated age profiles of the share of savers (left panel) and the realized fertility (right panel) from the baseline model (black solid line) and the hypothetical scenario in which households have perfect contraceptive control (red dashed lines). The parameters used are those estimated for low skilled households. First, we notice that the share of savers declines slightly in a model with perfect contraceptive control. The decline in savers stem from the reduced fertility risk and precautionary motives. More surprising is the fact that the realized fertility increases and households start having children earlier when unintended births do not occur. The reason for this result is that (low skilled) households have very strong preferences for two children, and low utility gains from the third child and would, therefore, prefer not to have the third child. In a world with imperfect contraceptive control, the households face the risk of ending up with too many children. Therefore, they might postpone efforts to conceive children and may stop having children after one child to reduce the risk of ending up with three children.

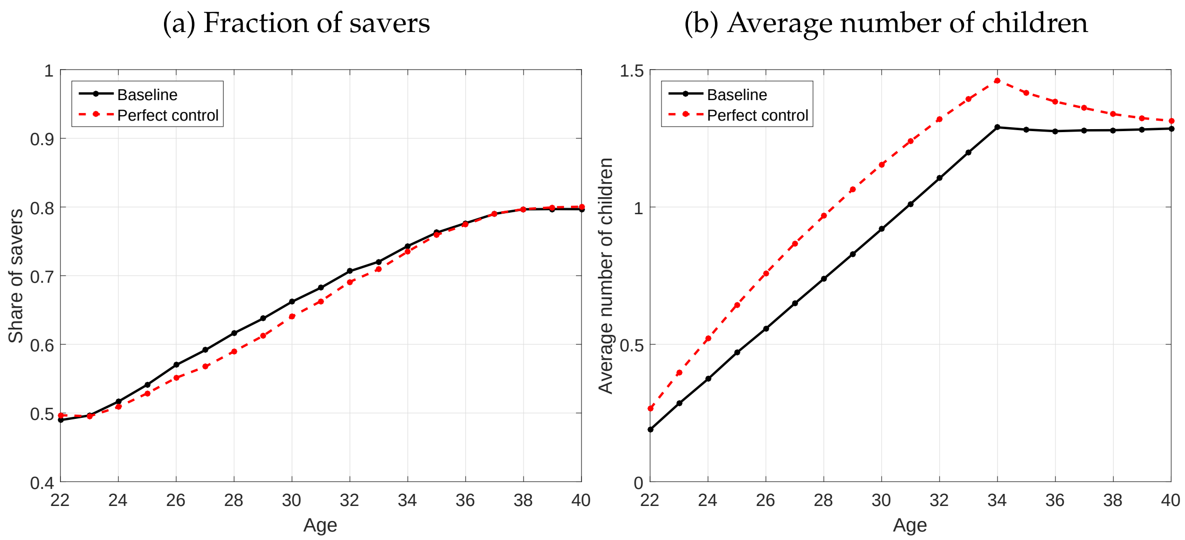

Figure 6 – Share of Savers over the Life Cycle.

Notes: Figure 6 illustrates the simulated age profiles of the share of savers (left panel) and the realized fertility (right panel) from the baseline model (solid lines) and the hypothetical scenario in which households have perfect contraceptive control (dashed lines).

To analyze the non-pecuniary cost associated with imperfect contraceptive control, we construct welfare measures as the expected discounted sum of utility. Specifically, we define welfare in the hypothetical case in which households have *perfect* contraceptive, denoted with superscript p, as

$$W^p(A_0, P_0) = \int\int\sum_{t=0}^{T}\beta^t v(n_t;\nu)u(C_t^\star(A_t, P_t|\underline{\wp} = 0)d\varepsilon d\eta$$

where A_0 and P_0 are initial wealth and initial permanent income, respectively.

Denote the period 1 realized income of households who – in the baseline model – would have a child of type $j \in \{\text{unintended}, \text{intended}, \text{all}\}$ as

$$\tilde{Y}_1 = Y_1\left(1 + \varrho_j\right),$$

where the income, Y_t, subsequently follows the process described in equations (1) and (2), and ϱ_j is the *compensation* received by a household in which a child of type j arrives.

The welfare of the baseline estimated model ($\underline{\wp} > 0$) is, as a function of the compensation parameter, ϱ_j, given as

$$W(A_0, P_0, \varrho_j) = \int\int\sum_{t=0}^{T}\beta^t v(n_t;\nu)u(C_t^\star(A_t, P_t, n_t|\underline{\wp} > 0, \varrho_j)d\varepsilon d\eta.$$

With these two welfare measures, we define the "welfare gap" between the baseline model and the hypothetical scenario with perfect contraceptive control as

$$\mathcal{G}(A_0, P_0, \varrho_j) \quad = \quad W(A_0, P_0, \varrho_j) - W^p(A_0, P_0),$$

such that the ϱ_j that sets $\mathcal{G}(A_0, P_0, \varrho_j) = 0$ will be a measure of how much income (as a percentage) households with various types of childbirths should be compensated by in beginning of adulthood to leave them as well off as in the hypothetical world with perfect contraceptive control. This measure is arguably artificial but provides us with a monetary measure indicating how costly these childbirths are to households. Table 10 presents the estimated compensation for all types of childbirths for both educational levels.[17]

The welfare loss experienced by those households who in fact have an unintended child in the baseline economy is informative about the non-pecuniary cost of an unintended childbirth. Comparing the welfare loss from households who have at least one

[17] To estimate the ϱ we approximate the integral by simulation and initialize all households with no children ($n_0 = 0$), wealth and permanent income of one ($A_0 = 0$ and $P_0 = 1$) and minimize the squared average (across simulations) welfare gap. We simulate 20 antithetic draws for all households (fixed at their respective values for low and high skilled households) per time-period and simulate them for 40 periods forward starting at age 20.

unintended child in the baseline scenario we can estimate $\varrho_{unintended}$ and interpret the estimate as the share of income that households who have an unintended child should receive as a lump sum in the beginning of their adult life (age 20) to make them just as well off as they would have been had they had perfect contraceptive control.

Table 10 – Estimated Non-Pecuniary Cost of Imperfect Contraceptive Control.

	Low skilled		High skilled	
Cost associated with (births)	Share	Est. cost	Share	Est. cost
Unintended childbirths (unintended)	0.12	1.06	0.09	0.25
By parity:				
1st child unintended		*0.27*		*0.23*
2nd child unintended		*0.56*		*0.20*
3nd child unintended		*5.04*		*0.42*
Risk of unintended childbirths (intended)	0.88	0.04	0.91	0.02
Imperfect contraceptive control (all)	1.00	0.26	1.00	0.07

We find that the non-pecuniary cost of having an unintended child is highest for low skilled households. Specially, we estimate the compensating variation of low skilled households to be around 1 times their annual income, while the compensating variation of high skilled households is somewhat lower at around 0.25 times their annual income. Interestingly, if we break up the compensation by parities, we see that the cost associated with having an unintended birth as the first child is about the same for the two groups at around 0.25 times of the annual income. This cost is mainly associated with mistimed childbirths because most households want to have at least one child and they can then adjust subsequent fertility choices. The main difference between educational groups is the cost of an unintended third birth. For the low skilled households, the compensation has to be five times the annual income. The high compensation reflects that low skilled households prefer not to have more than two children and an unintended third birth will lead to "too many children". So when an unintended third child arrives, it will have huge consequences on the welfare of that household. The last row (intended) shows that those households who do not have unintended births only need a minor compensation to compensate them for the extra risk they are facing. Thus, the major welfare loss from imperfect contraceptive control is incurred by those households who have unintended children.

As highlighted in the paragraph above, mistiming of births is associated with costs for the households because they have not accumulated the same level of wealth as households who have an intended birth. Our model framework allows us to calculate the wealth-to-income ratio the first time the households exert effort to conceive their first child in the simulated data. The optimal wealth-to-income ratio before having children is found to be 0.74 for low skilled and 0.91 for high skilled households, suggesting that only when low skilled households have accumulated around three fourths of their annual income, is it optimal to have their first child. The differences across educational groups might be due to the fact that high skilled households are more risk averse. Although the wealth-to-income ratio seems high, in our model this would include all types of savings, e.g., down payment on a house. The high wealth-income ratio may be one of the explanations for why so many households choose to postpone childbirth.

6 Concluding Discussion

The aim of this paper is to analyze the impact of intended and unintended births on household saving behavior and welfare. We achieve this by estimating a standard life cycle model where households choose consumption and savings and can imperfectly control fertility. The novel model framework allows us to identify intended and unintended births and analyze the welfare implications of the latter. We estimate the parameters of the model through indirect inference using the British Household Panel Survey. We select a sample of stable couples and use information on fertility, saving and income to estimate the model. A unique feature of the BHPS is that it contains information on fertility expectations. While we do not use fertility plans in the BHPS when estimating our model, we use this information as a validation check of the model predictions. We find that our estimated model can reasonably well reproduce the share of unintended births and the saving behavior around intended and unintended births.

Our results imply that there are substantial welfare costs associated with unintended births. We find both theoretically and empirically that saving rates around intended births decrease by about 3 to 6 percentage points. For unintended births, the decrease in the saving rate is much smaller than for intended births and the fall is less than 3 percentage points. To quantify the welfare costs, we calculate that a low skilled household that has an unintended birth should receive a lump sum of one-year's income to be as well off as they would be in the hypothetical scenario with perfect contraceptive control. For a high skilled household, the compensation is about ¼ of their annual income. We find that these high costs for low skilled households are mainly driven by the cost associated with "having too many children" rather than "having them too early". For the high skilled couples, the costs of unintended births are smaller and mainly caused by "having children too early".

While most of the literature on unintended births has focused on the consequences for young single women (teenage mothers), which arguably is an important issue, the consequences of unintended births for couples have received less attention. This is somewhat surprising because empirical studies have shown that around one out of five births among US and UK couples are unintended. This paper contributes to the literature by analyzing theoretically and empirically how unintended births affect the welfare of couples in a novel life cycle model with imperfect fertility control.

Our results highlight that there are considerable welfare implications associated with unintended births among couples and that even in our sample of stable couples, unintended births are as frequent as one in ten births. This suggests that improving contraceptive control among couples could lead to welfare improvements and perhaps also to a higher realized fertility. Also, policies aimed at compensating families with (unintended) children (e.g. child benefits) may be important to lower the cost of unintended births. Future research might, therefore, investigate the construction of welfare improving child benefits in a world with imperfect contraceptive control. Such policies might have long run effects on, for example, later life outcomes of unintended children and their siblings.

References

ADDA, J., C. DUSTMANN AND K. STEVENS (forthcoming): "The Career Costs of Children," *Journal of Political Economy*.

ATTANASIO, O. P., J. BANKS, C. MEGHIR AND G. WEBER (1999): "Humps and Bumps in Lifetime Consumption," *Journal of Business & Economic Statistics*, 17(1), 22–35.

ATTANASIO, O. P. AND G. WEBER (1993): "Consumption Growth, the Interest Rate and Aggregation," *The Review of Economic Studies*, 60(3), 631–649.

BANKS, J., C. O'DEA AND Z. OLDFIELD (2010): "Cognitive Function, Numeracy and Retirement Saving Trajectories," *The Economic Journal*, 120(548), 381–F410.

BAUDIN, T., D. DE LA CROIX AND P. GOBBI (2015): "Fertility and Childlessness in the United States," *American Economic Review*, 105(6), 1852–82.

BERRINGTON, A. (2004): "Perpetual Postponers? Women's, Men's and Couple's Fertility Intentions and Subsequent Fertility Behaviour," *Population Trends*, 117, 9–19.

BLUNDELL, R., L. PISTAFERRI AND I. PRESTON (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, pp. 1887–1921.

BROWNING, M. (1992): "Children and Household Economic Behavior," *Journal of Economic Literature*, 30(3), 1434–1475.

BROWNING, M. AND M. EJRNÆS (2009): "Consumption and Children," *The Review of Economics and Statistics*, 91(1), 93–111.

CARROLL, C. D. (1992): "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 2, 61–135.

——— (2006): "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economics Letters*, 91(3), 312–320.

CHOI, S. (forthcoming): "Fertility Risk in the Life-Cycle," *International Economic Review*.

DEATON, A. (1991): "Saving and Liquidity Constraints," *Econometrica*, 59(5), 1221–1248.

EISENHAUER, P., J. J. HECKMAN AND S. MOSSO (2015): "Estimation of Dynamic Discrete Choice Models by Maximum Likelihood and the Simulated Method of Moments," *International Economic Review*, 56(2), 331–357.

FERNÁNDEZ-VILLAVERDE, J. AND D. KRUEGER (2007): "Consumption over the Life Cycle: Facts from Consumer Expenditure Survey Data," *The Review of Economics and Statistics*, 89(3), 552–565.

FLEISSIG, A. (1991): "Unintended Pregnancies and the use of Contraception: Changes from 1984 to 1989," *British Medical Journal*, 302, 147.

GOURIÉROUX, C., A. MONFORT AND E. RENAULT (1993): "Indirect Inference," *Journal of Applied Econometrics*, 8, 85–118.

GOURINCHAS, P.-O. AND J. A. PARKER (2002): "Consumption Over the Life Cycle," *Econometrica*, 70(1), 47–89.

GRANT, C. (2010): "Demographics and Life-Cycle Consumption," Unpublished mimeo, Reading University.

HAYFORD, S. R. (2009): "The Evolution of Fertility Expectations over the Life Course," *Demography*, 46(4), 765–783.

HONORÉ, B. E. (1992): "Trimmed Lad and Least Squares Estimation of Truncated and Censored Regression Models with Fixed Effects," *Econometrica*, 60(3), 533–565.

ISKHAKOV, F., T. H. JØRGENSEN, J. RUST AND B. SCHJERNING (2015): "Estimating Discrete-Continuous Choice Models: The Endogenous Grid Method with Taste Shocks," Discussion Paper 15-19, Department of Economics, University of Copenhagen.

JONES, L. E., A. SCHOONBROODT AND M. TERTILT (2010): "Fertility Theories: Can They Explain the Negative Fertility-Income Relationship?," in *Demography and the Economy*, ed. by J. B. Shoven, chap. 2, pp. 43–100. University of Chicago Press.

KALWIJ, A. (2005): "Consumption and Income Around the Time of Births," *Review of Economics of the Household*, 3(1), 75–89.

KINGSBURY, N. M. AND L. GREENWOOD (1992): "Fertility Expectation and Employment across Three Female Cohorts," *Journal of Family and Economic Issues*, 13(1), 73–93.

KRISHNAN, V. (1988): "Occupational Status, Earnings, and Fertility Expectations: Development and Estimation of a Causal Model," *De Economist*, 136(3), 358–382.

LORDAN, G. AND P. FRIJTERS (2013): "Unplanned Pregnancy and the Impact on Sibling Health Outcomes," *Health Economics*, 22(8), 903–14.

LUNDBERG AND ROSE (2000): "Parenthood and the Earnings of Married Men and Women," *Labour Economics*, 7(6), 689–710.

LYNN, F. B., B. SCHNEIDER AND Z. ZHANG (2013): "The Changing Relationship between Fertility Expectations and Educational Expectations: Adolescents in the 1970s Versus the 1980s," *Journal of Family Issues*, 34(9), 1147–1174.

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

MICHAELIDES, A. AND S. NG (2000): "Estimating the Rational Expectations Model of Speculative Storage: A Monte Carlo Comparison of Three Simulation Estimators," *Journal of Econometrics*, 96(2), 231–266.

MOSHER, W. D., J. JONES AND J. C. ABMA (2012): "Intended and Unintended Births in the United States: 1982–2010," National Health Statistics Reports 55, U.S. Department of Health and Human Services.

OFFICE FOR NATIONAL STATISTICS (2013): "Statistical Bulletin: Cohort Fertility, 2012," Statistical bulletin, The UK Office for National Statistics (ONS).

ROSENZWEIG, M. R. AND K. I. WOLPIN (1993): "Maternal Expectations and Ex Post Rationalizations: The Usefulness of Survey Information on the Wantedness of Children," *The Journal of Human Resources*, 28(2), 205–229.

SEFTON, J., J. VAN DE VEN AND M. WEALE (2008): "Means Testing Retirement Benefits: Fostering Equity or Discouraging Savings?*," *The Economic Journal*, 118(528), 556–590.

SIGLE-RUSHTON, W. AND J. WALDFOGEL (2007): "The Incomes of Families with Children: A Cross-National Comparison," *Journal of European Social Policy*, 17(4), 299–318.

SMITH, A. A. (1993): "Estimating Nonlinear Time-series Models Using Simulated Vector Autoregressions," *Journal of Applied Econometrics*, 8, 63–84.

SOMMER, K. (2014): "Fertility Choice in a Life Cycle Model with Idiosyncratic Uninsurable Earnings Risk," Finance and Economics Discussion Series 2014-32, The Federal Reserve Board.

STOLZENBERG, R. M. AND L. J. WAITE (1977): "Age, Fertility Expectations and Plans for Employment," *American Sociological Review*, 42(5), 769–783.

TRUSSELL, J. AND C. WILSON (1985): "Sterility in a Population with Natural Fertility," *Population Studies*, 39(2), pp. 269–286.

WALDFOGEL, J. (1998): "Understanding the "Family Gap" in Pay for Women with Children," *The Journal of Economic Perspectives*, 12(1), 137–156.

Online Appendix (Not for Publication)

Saving Behaviour around Intended and Unintended Childbirths

Mette Ejrnæs Thomas H. Jørgensen
University of Copenhagen

May 12, 2016

A Additional Figures and Tables

(a) Low skilled (b) High skilled

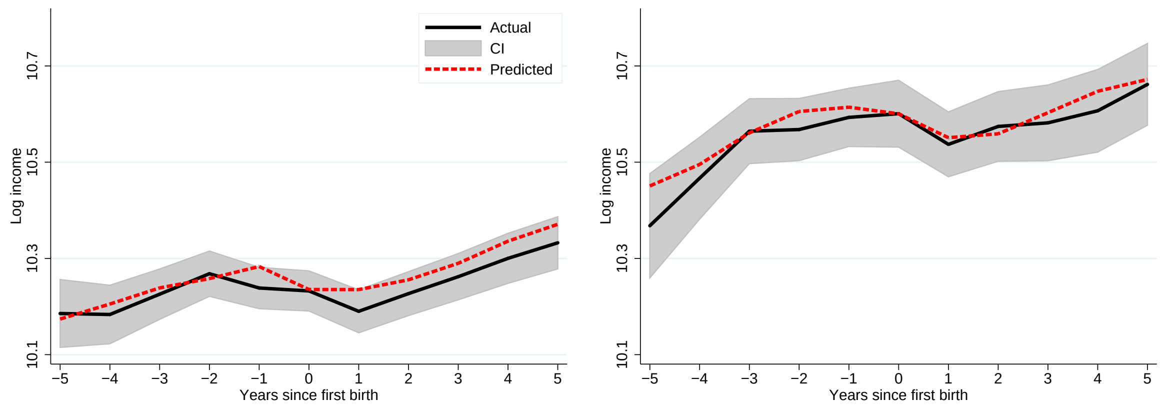

Figure A.1 – Income around First Birth.

Notes: Figure A.1 illustrates how well the income process approximates actual household income in the BHPS around the first childbirth.

(a) Low skilled (b) High skilled

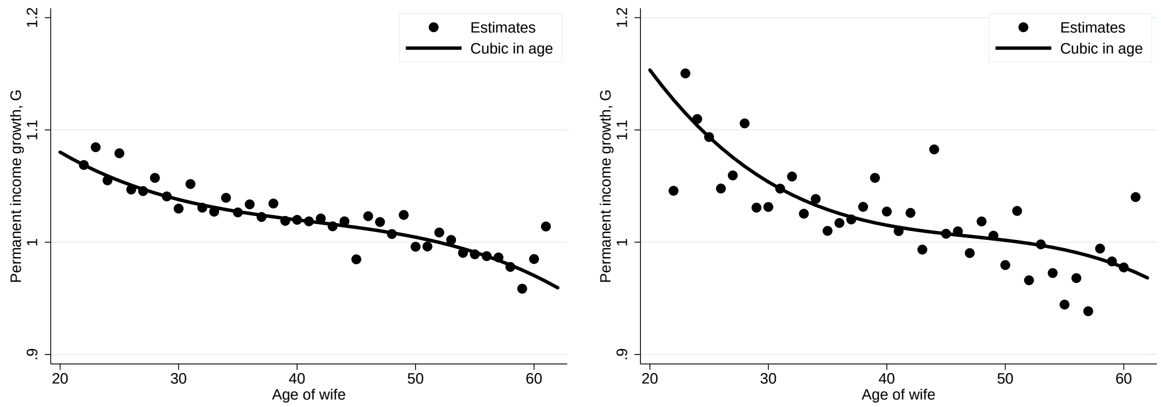

Figure A.2 – Age Profiles of Permanent Income Growth, $\hat{G}_t$.

Notes: Figure A.2a presents the age profile of the estimated permanent income growth, G_t. The dots are based on estimates of ϕ_{2t} from equation (3) with two percent annual growth added, while the solid line is a cubic polynomial in age from a median regression.

Table A.2 – Calibration of σ_ε.

	Low skilled						High skilled					
	(1)	(2)	(3)	(4)	(5)	(6)	(1)	(2)	(3)	(4)	(5)	(6)
ρ	1.050	1.054	1.112	1.117	1.050	1.172	1.299	1.365	1.294	1.299	1.302	1.050
κ_1	2.910	2.404	1.318	1.193	2.990	1.011	0.451	0.396	0.504	0.563	0.537	2.998
κ_2	2.978	2.461	1.370	1.256	0.045	0.001	0.421	0.376	0.497	0.554	0.526	0.000
κ_3	0.001	0.871	0.138	0.057	1.117	0.388	0.138	0.176	0.314	0.417	0.435	1.331
ν	0.144	0.128	0.144	0.136	0.128	0.144	0.156	0.151	0.148	0.161	0.148	0.127
Obj.	419.2	417.7	422.5	421.0	466.2	454.8	304.3	311.7	313.8	317.2	325.0	351.9
σ_ε	1.000	0.700	0.600	0.500	0.400	0.100	1.000	0.700	0.600	0.500	0.400	0.100
Unint.	9.7	9.7	9.7	9.7	12.9	12.9	7.4	7.5	7.5	7.3	7.3	11.3
1. Unint.	7.1	7.3	7.4	7.4	5.0	5.0	8.4	8.9	8.9	8.9	9.3	3.8
2. Unint.	10.1	9.6	9.6	9.6	99.1	99.5	5.3	5.2	5.0	4.9	4.8	99.9
3. Unint.	82.9	70.8	65.1	73.0	18.3	15.6	9.5	9.3	9.2	8.5	7.6	19.9

Notes: Table A.2 reports estimates of the preferences for children for various values of σ_ϵ. The four last rows report the associated share of unintended childbirths also conditional on the parity.

Table A.1 – Saving (yes/no) around First Birth – Estimation Results.

	OLS			Logit		
	RE	FE	FE	RE	FE	FE
Number of children	-0.687***	-0.608	-0.422	-0.081***	-0.076	-0.059
	(0.192)	(0.343)	(0.251)	(0.022)	(0.042)	(0.033)
Unintended child	1.740*	2.058	1.920	0.221*	0.244*	0.216
	(0.835)	(1.061)	(1.051)	(0.108)	(0.113)	(0.113)
Married	0.702*	1.173	0.669	0.090*	0.152	0.105
	(0.334)	(0.762)	(0.512)	(0.041)	(0.088)	(0.065)
Home owner	0.088	-1.062	-0.697	0.008	-0.065	-0.066
	(0.367)	(0.680)	(0.511)	(0.045)	(0.077)	(0.064)
Plans to be childless	-2.340**			-0.301**		
	(0.869)			(0.110)		
High skilled, husband	-0.002			0.001		
	(0.343)			(0.040)		
High skilled, wife	0.723*			0.082*		
	(0.359)			(0.041)		
White	-0.316			-0.036		
	(0.792)			(0.091)		
Age, husband	0.014			0.002		
	(0.035)			(0.004)		
Constant	2.267			0.760**	0.747***	0.682***
	(2.044)			(0.234)	(0.138)	(0.110)
Age dummies	Yes	No	No	Yes	Yes	Yes
Wave dummies	Yes	Yes	Yes	Yes	Yes	Yes
Obs	1274	468	814	1274	1274	2169
R^2					0.090	0.051

Notes: Estimation is based on households in which the wife is aged 20-39.

B Estimation by Indirect Inference

We estimate the model parameters by indirect inference (Smith, 1993 and Gouriéroux, Monfort and Renault, 1993). Denote λ as a $K \times 1$ vector of auxiliary parameters (AP) calculated based on observed data. Denote as $\overline{\lambda}(\theta)$ the average of the same K APs calculated from S simulated data sets, from the calibrated life cycle model discussed above,

$$\overline{\lambda}(\theta) = \frac{1}{S} \sum_{r=1}^{S} \lambda_r(\theta).$$

We estimate θ as the minimizer of the weighted squared difference between the observed and simulated APs,

$$\hat{\theta} = \arg\min_{\theta \in \Theta} (\lambda - \overline{\lambda}(\theta))' W^{-1} (\lambda - \overline{\lambda}(\theta)), \quad \text{(B.1)}$$

where W is a diagonal $K \times K$ weight matrix with bootstrapped variances of the moments in the data on the diagonal, as suggested by Eisenhauer, Heckman and Mosso (2015). To minimize the effect of simulation noise, we let $S = 10$ and use antithetic draws.[18] For each trial value of the parameters θ, we numerically solve the model using the discrete-continuous endogenous grid method (DC-EGM), proposed in Iskhakov, Jørgensen, Rust and Schjerning (2015) and described below and simulate synthetic data from this model. To simulate data, we use the first observation in our data. For example, initial wealth, income, and number of children is used for all S simulations.

Because we have discrete choices and a relatively high dimensional optimization problem, the optimization problem is likely to be ill-behaved and we search for the *global* minimizer of (B.1). To do so, we apply a sequence of alternative optimizers, perform the sequence of estimators 10 times (where randomization ensures different starting values) and use the estimates yielding the lowest criteria function. For each of the 10 estimation runs, we start with MATLABs `particleswarm` which is a global optimization routine using randomization to search through the parameter space. We use 40 particles and switch to Nelder-Mead (`fminsearch` in MATLAB) using the best candidates from the particleswarm when the particaleswarm has either converged or used $30 \cdot \dim(\theta)$ iterations. We report the converged estimates with the lowest objective function out of the 10 sequences of estimators.

The simulated moments from the estimated model, $\overline{\lambda}(\theta)$, and the moments calculated from the BHPS, λ, are plotted against each other in Figure B.1. The model is *over-identified* and the simulated moments align with those in the BHPS. See also Figure 5 in the main text.

We report asymptotic standard errors, calculated as the squared root of the diagonal elements of the estimated asymptotic covariance matrix,

$$COV(\hat{\theta}) = N^{-1}\left(1 + \frac{1}{S}\right) AVar,$$

where

$$AVar = (g'W^{-1}g)^{-1}g'W^{-1}W_0W^{-1}g(g'W^{-1}g)^{-1},$$

in which g is the gradient of $\lambda - \overline{\lambda}(\theta)$ with respect to θ and W_0 is the optimal weighting matrix, which we calculate as the bootstrapped covariance matrix of the moments calculated from the data. If we used W_0 rather than W as the weighting matrix (optimal weighing), the asymptotic variance would collapse to $AVar = (g'W^{-1}g)^{-1}$.

[18] Michaelides and Ng (2000) find that this strategy of using a simulated sample 10 times as large as the empirical sample yields good small-sample properties.

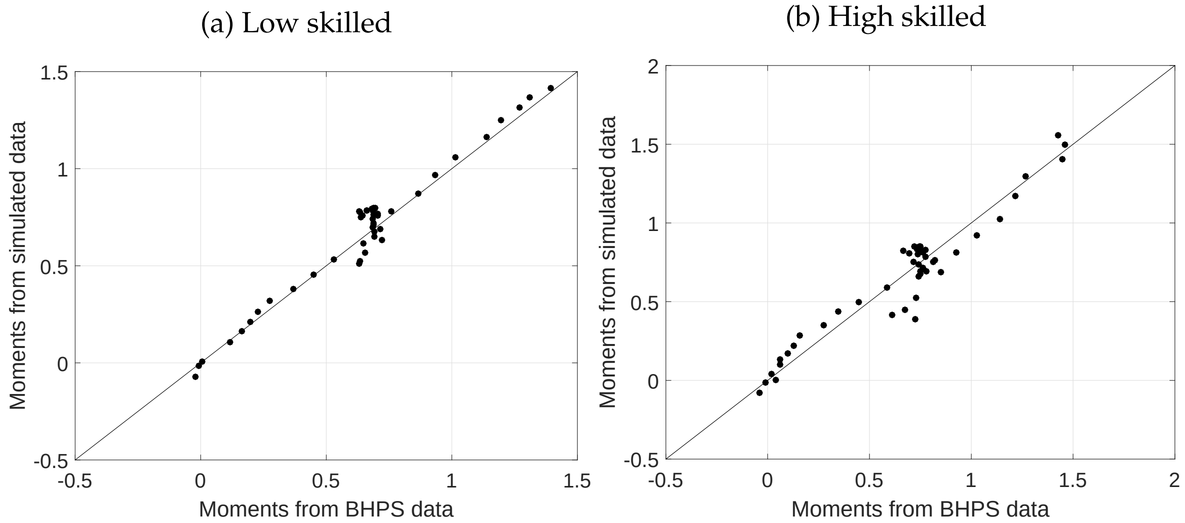

Figure B.1 – Moments in the BHPS and Simulated Data.

Notes: The figure illustrates on the y-axis the simulated moments against the moments in the BHPS data.

B.1 Imputing Wealth in the BHPS

We impute a measure of wealth from the BHPS to use to initialize households when simulating synthetic data in the indirect inference estimation approach. In waves 5, 10 and 15, individuals are asked about their wealth stock. The measures (bank holdings and savings accounts) are bracketed into five intervals and we use the average within a bracket as the level. This gives us an approximate level of wealth holdings in the three waves.

To impute the wealth holdings in the first wave in which a given household is observed, we add (subtract) state savings in each subsequent (previous) wave. While this is a very crude approximation strategy it serves only as a proxy for the initial distribution of wealth in the BHPS when simulating synthetic data.

B.2 Estimating the Risk-Parameter for Unintended Births

We obtain an estimate of the parameter $\underline{\wp}$ from the BHPS. In our model the probability of an unintended birth, when $e_t = 0$, is given by

$$\wp_t = \underline{\wp}\overline{\wp}_t.$$

In the validation sample, we can calculate the probability of unintended births (see Figure 3) for each age. To get enough observations we pool the two education groups and use 2-year bands. The estimates of $\wp_t$ are shown in Figure B.2. The estimate of $\underline{\wp}$ is found as the weighted average of the fractions of $\frac{\widehat{\wp}_t}{\overline{\wp}_t}$ from the age of 30 to 44:

$$\widehat{\wp} = \sum_t w_t \frac{\widehat{\wp}_t}{\overline{\wp}_t},$$

where $\overline{\wp}_t$ is the biological fecundity taken from Trussell and Wilson (1985).

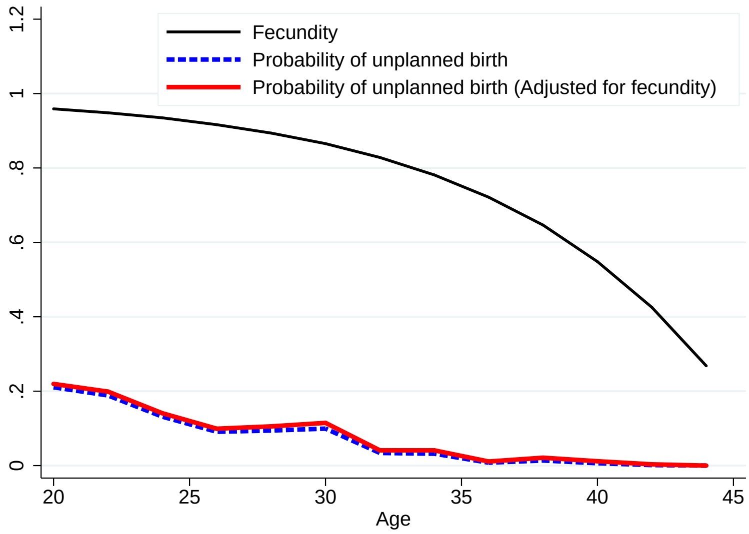

Figure B.2 – Probability of Unintended Birth and Fecundicity.

Notes: Figure B.2 shows the estimated probability for an unintended birth as a function of age (thin black and bold red lines) and the fecundity (blue bold dashed line). Biological fecundity is based on Trussell and Wilson (1985).

C Solving the Model

All relations are normalized by permanent income, P_t, and lower case variables denote normalized quantities (e.g., $c_t = C_t/P_t$). The model is solved recursively by backwards induction, starting with the terminal period, T. Within a given period, optimal consumption is found using the Discrete-Continuous Endogenous Grid Method (DC-EGM) generalization in Iskhakov, Jørgensen, Rust and Schjerning (2015) of the EGM proposed by Carroll (2006).

In the terminal period, $t = T$, independent of the presence of children, households consume all their remaining wealth, $c_T = m_T$. In preceding periods, in which households are retired, and fertility is completed, consumption across periods satisfies the Euler equation

$$U'(C_t, n_t) = \max \left\{ U'(m_t, n_t) \ , \ R\beta U'(C_{t+1}, n_{t+1}) \right\}, \ \forall t \in [T_r, T],$$

where consumption cannot exceed available resources. The problem is standard, and we focus on how the model is solved in the fertile period of life.

The discrete choice over effort to conceive a child and the continuous consumption and savings choice can be partitioned into

$$v_t(m_t, n_t, \varepsilon_t) = \max\{v_t(m_t, n_t|e_t = 0) + \sigma_\epsilon \epsilon_t(0), v_t(m_t, n_t|e_t = 1) + \sigma_\epsilon \epsilon_t(1)\},$$

where an *iid* effort-specific taste shock, ε_t, also affects the level of utility. This shock is only relevant in periods in which the household is fertile. The effort-specific value functions are given as

$$\begin{aligned} v_t(m_t, n_t|e_t) &= \max_{c_t \in (0, m_t]} \{U(c_t, n_t) + \beta \mathbb{E}_t[\wp_t(e_t, n_t) \mathcal{G}_{t+1}^{1-\rho} V_{t+1}(m_{t+1}(e_t), n_{t+1}(e_t)) \\ &\qquad + (1 - \wp_t(e_t, n_t)) \mathcal{G}_{t+1}^{1-\rho} v_{t+1}(m_{t+1}(e_t), n_{t+1}(e_t))]\}, \end{aligned}$$

where $\mathbb{E}_t[\cdot]$ denotes expectations over future income shocks, ξ_{t+1}, η_{t+1}, and effort-specific taste shocks, ϵ_{t+1} and

$$\begin{aligned} m_{t+1}(e_t) &= \mathcal{G}_{t+1}^{-1} R(m_t - c_t) + \Pi(n_{t+1}(e_t), n_t) \xi_{t+1} \\ \mathcal{G}_{t+1} &\equiv \mathcal{G}_{t+1}(n_{t+1}(e_t), n_t) = G_{t+1} \eta_{t+1} \Omega(n_{t+1}(e_t), n_t) \end{aligned}$$

is the next period available resources and a normalization factor stemming from the fact that we have normalized by permanent income.

Assuming that the choice-specific taste shock, ϵ, is extreme value type I distributed, the expected value function – with respect to the unobserved taste shocks – is given by the log-sum,

$$EV_t(m_t, n_t) = \sigma_\epsilon \log\left[\exp\left(v_t(m_t, n_t|e_t = 0)/\sigma_\epsilon\right) + \exp\left(v_t(m_t, n_t|e_t = 1)/\sigma_\epsilon\right)\right],$$

and the probability of an household trying to conceive a child is given as

$$\Pr(e_t = 1|m_t, n_t) = \frac{\exp\left(v_t(m_t, n_t|e_t = 1)/\sigma_\epsilon\right)}{\exp\left(v_t(m_t, n_t|e_t = 0)/\sigma_\epsilon\right) + \exp\left(v_t(m_t, n_t|e_t = 1)/\sigma_\epsilon\right)}.$$

C.1 Effort-specific Euler Equations

The FOC wrt consumption, for a given discrete effort choice e_t, is

$$U'(c_t, n_t) - R\beta \mathbb{E}_t[\wp_t(e_t, n_t) \mathcal{G}_{t+1}^{-\rho} V'_{t+1}(m_{t+1}, n_{t+1}) + (1 - \wp_t(e_t, n_t)) \mathcal{G}_{t+1}^{-\rho} V'_{t+1}(m_{t+1}, n_{t+1})] = 0,$$

and the envelope theorem states that

$$V'_t(m_t, n_t) = R\beta \mathbb{E}_t[\wp_t(e_t, n_t) \mathcal{G}_{t+1}^{-\rho} V'_{t+1}(m_{t+1}, n_{t+1}) + (1 - \wp_t(e_t, n_t)) \mathcal{G}_{t+1}^{-\rho} V'_{t+1}(m_{t+1}, n_{t+1})]$$

such that we have

$$U'(c_t, n_t) = R\beta\mathbb{E}_t[\wp_t(e_t, n_t)\mathcal{G}_{t+1}^{-\rho}U'(C_{t+1}, n_{t+1}) + (1-\wp_t(e_t, n_t))\mathcal{G}_{t+1}^{-\rho}U'(C_{t+1}, n_{t+1})],$$

such that when $a_t > 0$, optimal consumption can be found by inverting the Euler equation

$$\begin{aligned} & c_t^\star(m_t, n_t|e_t) \\ & \quad = \left(R\beta\mathbb{E}_t[\wp_t(e_t, n_t)\mathcal{G}_{t+1}^{-\rho}\frac{v(n_{t+1})}{v(n_t)}C_{t+1}^{-\rho} + (1-\wp_t(e_t, n_t))\mathcal{G}_{t+1}^{-\rho}\frac{v(n_{t+1})}{v(n_t)}C_{t+1}^{-\rho}]\right)^{-\frac{1}{\rho}}, \end{aligned}$$

Note that the discrete choice might lead to a non-concave value function and the Euler equation is, thus, potentially only necessary but not sufficient for optimal consumption. Fortunately, the solution method applied here (the DC-EGM of Iskhakov, Jørgensen, Rust and Schjerning, 2015) finds *all* solutions to the Euler equation and we just need to discard non-optimal solutions afterwards. Below, we describe how we solve for optimal consumption using the DC-EGM.

C.2 Solving the model using the DC-EGM

The EGM constructs a grid over end-of-period wealth, $a_t = m_t - c_t$, rather than beginning-of-period resources, m_t. The discrete-continuous EGM (DC-EGM) proposed in Iskhakov, Jørgensen, Rust and Schjerning (2015) builds on the same idea. Denote this grid of Q points as $\vec{a}_t = (0, a_t^2, \ldots, a_t^Q)$ in which the lower bound of zero is where agents are on the cusp of being credit constrained and save nothing. The endogenous level of beginning-of-period resources consistent with end-of-period assets, $\vec{a}_t$, and optimal consumption, $c_t^\star$, is given by $m_t = \vec{a}_t + c_t^\star(m_t, n_t)$. The expectations are over – besides the potential arrival of a newborn child – next period transitory (ξ_{t+1}) and permanent income shocks (η_{t+1}). $\#_\eta = \#_\xi = 8$ Gauss-Hermite quadrature points are used for each income shock to approximate expectations. $\#_a = 200$ discrete grid points are used in $\vec{a}_t$ to approximate the consumption function with more mass at lower levels of wealth to approximate accurately the curvature of the consumption function.

Because the Euler equation is not satisfied at the constraint, we first describe how we find optimal consumption in the unconstrained region. Denote the probability that a child leaves the household at the beginning of the next period as $p(n_t)$ such that the law of motion for the number of children at the beginning of next period can be expressed as

$$n_{t+1}(e_t, n_t) = \begin{cases} n_t + 1 & \text{with probability } P_1(e_t, n_t) = \wp_t(e_t, n_t)(1 - p(n_t)) \\ n_t & \text{with probability } P_2(e_t, n_t) = \wp_t(e_t, n_t)p(n_t) \\ n_t & \text{with probability } P_3(e_t, n_t) = (1 - \wp_t(e_t, n_t))(1 - p(n_t)) \\ n_t - 1 & \text{with probability } P_4(e_t, n_t) = (1 - \wp_t(e_t, n_t))p(n_t) \end{cases}$$

where the first case is when no child leaves and a new arrives, the second case is when a new child arrives and an existing child leaves, the third is when no child leaves and no new child arrives, and the final and fourth case is when a child leaves while no new child is born. Denoting ω_j^η as the jth quadrature node, we approximate effort-specific consumption as

$$c_t^\star(m_t, n_t|e_t) \approx \left(R\beta\Xi_t(\vec{a}_t, n_t, e_t)\right)^{-\frac{1}{\rho}}$$

where

$$\begin{aligned} &\Xi_t(a_t, n_t, e_t) \\ = \quad & P_1(e_t, n_t) \sum_{i=1}^{\#_\varepsilon} \omega_i^\varepsilon \sum_{j=1}^{\#_\eta} \omega_j^\eta \sum_{e_{t+1}=0}^{1} \Pr(e_{t+1}|m_{t+1}, n_t + 1) \mathcal{G}_{t+1}^{-\rho} \frac{v(n_t + 1)}{v(n_t)} \check{C}_{t+1}(m_{t+1}, n_t + 1, e_{t+1})^{-\rho} \\ + \quad & (P_2(e_t, n_t) + P_3(e_t, n_t)) \sum_{i=1}^{\#_\varepsilon} \omega_i^\varepsilon \sum_{j=1}^{\#_\eta} \omega_j^\eta \sum_{e_{t+1}=0}^{1} \Pr(e_{t+1}|m_{t+1}, n_t) \mathcal{G}_{t+1}^{-\rho} \check{C}_{t+1}(m_{t+1}, n_t, e_{t+1})^{-\rho} \\ + \quad & P_4(e_t, n_t) \sum_{i=1}^{\#_\varepsilon} \omega_i^\varepsilon \sum_{j=1}^{\#_\eta} \omega_j^\eta \sum_{e_{t+1}=0}^{1} \Pr(e_{t+1}|m_{t+1}, n_t - 1) \mathcal{G}_{t+1}^{-\rho} \frac{v(n_t - 1)}{v(n_t)} \check{C}_{t+1}(m_{t+1}, n_t - 1, e_{t+1})^{-\rho} \end{aligned}$$

is the *expected marginal utility of next-period consumption* as a function of end-of-period wealth, a_t, number of children, n_t, and the effort, e_t. Functions denoted with $\check{\bullet}$ refer to interpolation functions. Because we fix a grid over end-of-period wealth and calculate $m_{t+1} = \mathcal{G}_{t+1}^{-1} R\vec{a}_t + Y_{t+1}$, we can find the optimal consumption in closed form and infer the effort-specific level of resources, m_t, consistent with end-of-period wealth and consumption.

The value function associated with the effort e_t is then calculated as

$$v_t(m_t, n_t|e_t) \approx U(c_t^\star(m_t, n_t|e_t), n_t) + \beta\Lambda_t(\vec{a}_t, n_t, e_t)$$

where

$$\begin{aligned}
\Lambda_t(a_t,n_t,e_t) &= P_1(e_t,n_t)\sum_{i=1}^{\#_\varepsilon}\omega_i^\varepsilon\sum_{j=1}^{\#_\eta}\omega_j^\eta\mathcal{G}_{t+1}^{1-\rho}\log\left[\sum_{e_{t+1}=0}^{1}\exp\left(\check{v}_{t+1}(m_{t+1},n_t+1|e_{t+1})\right)\right]\\
&+ (P_2(e_t,n_t)+P_3(e_t,n_t))\sum_{i=1}^{\#_\varepsilon}\omega_i^\varepsilon\sum_{j=1}^{\#_\eta}\omega_j^\eta\mathcal{G}_{t+1}^{1-\rho}\log\left[\sum_{e_{t+1}=0}^{1}\exp\left(\check{v}_{t+1}(m_{t+1},n_t|e_{t+1})\right)\right]\\
&+ P_4(e_t,n_t)\sum_{i=1}^{\#_\varepsilon}\omega_i^\varepsilon\sum_{j=1}^{\#_\eta}\omega_j^\eta\mathcal{G}_{t+1}^{1-\rho}\log\left[\sum_{e_{t+1}=0}^{1}\exp\left(\check{v}_{t+1}(m_{t+1},n_t-1|e_{t+1})\right)\right]
\end{aligned}$$

is the expected continuation value.

In the *constrained* region, $a=0$, we can easily infer the value function as (recall consumption is $c_t^\star(\bullet)=m_t$ here)

$$v_t(m_t,n_t|e_t)\approx U(m_t,n_t)+\beta\Lambda_t(0,n_t,e_t)$$

by constructing a grid over beginning of period resources in this region as $\vec{m}_t = (0,\ldots,\underline{c}_t^\star)$ where $\underline{c}_t^\star$ is the found consumption level associated with $a=0$.

Discarding Non-Optimal solutions. Because the Euler equation is potentially only necessary, we need to remove non-optimal solutions found in the EGM step outlined above. We do so by discarding point i if we can find some other point j for which we have a *lower* value with a *higher* level of resources. See also Iskhakov, Jørgensen, Rust and Schjerning (2015) for a description of the upper envelope algorithm, we apply.