OXFORD BULLETIN OF ECONOMICS AND STATISTICS, 0305–9049
doi: 10.1111/obes.12170

Life-Cycle Consumption and Children: Evidence from a Structural Estimation*

THOMAS H. JØRGENSEN†,‡

†*Department of Economics, University of Copenhagen, Øster Farimagsgade 5, building 26 DK-1353, Copenhagen, (e-mail: thomas.h.jorgensen@econ.ku.dk)*
‡*Centre for Applied Microeconometrics (CAM)*

Abstract

I estimate by maximum likelihood a dynamic model of optimal intertemporal allocation of consumption in the presence of children using high-quality Danish longitudinal data. The number and age of all children can affect the marginal utility of consumption while income uncertainty, credit constraints and postretirement motives also influence household behaviour. While I estimate that children have a surprisingly small effect on the marginal utility of non-durable consumption, data simulated from the estimated model replicates similar correlations between log consumption growth and changing household composition as found in the Danish data and typically found in UK and US data. To reconcile the results with existing studies, I illustrate how ignoring precautionary motives increases the estimated importance of children. The results indicate that precautionary motives might play a larger role than children in explaining the observed consumption age profile.

I. Introduction

This study is concerned with the effect of children on the marginal utility of non-durable consumption over the life cycle. The average number of children and non-durable consumption share similar hump-shaped (inverted-U) age profiles. The extent to which children affect the marginal utility of consumption has, therefore, received great attention over the last two decades with an important role for children as the most common finding.[1]

JEL Classification numbers: D12, D14, D91
*I am grateful to Mette Ejrnæs, Bertel Schjerning, Martin Browning, Jeppe Druedahl, Søren Leth-Petersen, Christopher Carroll, John Rust, Richard Blundell, Orazio Attanasio, Hamish Low, Miles Kimball, Dennis Kristensen, Daniel Borowczyk-Martins, Kjetil Storesletten, Jérôme Adda, Anders Munk-Nielsen, Mette Foged, Damoun Ashournia, Signe Dyrberg and conference participants at the EEA in Gothenburg 2013, the CAM December workshop in Copenhagen, 2013, the IAAE in London, 2014, and the ESEM in Toulouse, 2014, for comments and suggestions. Financial support from the Danish Council for Independent Research in Social Sciences (FSE, grant no. 4091-00040) is gratefully acknowledged. All errors are my own.

[1] Irvine (1978) might be one of the first to suggest that the hump in consumption could be due to changes in household composition. Some important contributions to the literature on the effect of children on consumption are due to Blundell, Browning and Meghir (1994); Banks, Blundell and Preston (1994); Attanasio and Weber (1995); Attanasio and Browning (1995); Attanasio *et al.* (1999); Fernández-Villaverde and Krueger (2007) and Browning and Ejrnæs (2009).

The same consumption profile can, however, be rationalized by alternative life cycle motives such as precautionary motives or non-separability between consumption and leisure with very different policy implications.[2] Despite the amount of studies analysing the effect of demographics on consumption, allowing for multiple consumption-savings motives simultaneously is rare in this literature.

I estimate the effect of children on the marginal utility of consumption by a novel full-solution maximum likelihood (ML) estimator. The estimated dynamic model of intertemporal allocation of consumption allows the number and age of children to affect the marginal utility of consumption while income uncertainty, credit constraints and postretirement motives also influence consumption and saving behaviour. I estimate the model using Danish administrative register data giving detailed longitudinal information on household characteristics, income and, importantly, most assets and liabilities. Net worth is a crucial determinant of optimal intertemporal allocation of consumption and savings, which I use to identify the parameters of interest.

I find an economically negligible effect of children on the marginal utility of consumption. While I estimate that children have a surprisingly small effect on the marginal utility of non-durable consumption, the risk aversion coefficients and subjective discount factors are estimated in the range typically found. The estimated model fits the Danish data very well and simulated data from the estimated model reproduce correlations between log consumption growth and changing household composition as in the Danish data and typically found in UK and US data. Particularly, applying a similar log-linearized Euler equation estimator as in Attanasio *et al.* (1999) yields similar results for the Danish data and data simulated from the estimated model.[3] The estimation results in Attanasio *et al.* (1999) and the preferred specification in Alan, Attanasio and Browning (2009) imply that children increase marginal utility of consumption by around 40% and 150% respectively. Browning and Ejrnæs (2009) find that the number and age of children can completely explain the hump in consumption over the life cycle. Log-linearized Euler equation estimation methods do not, however, perform well if households face precautionary motives (Carroll, 2001; Ludvigson and Paxson, 2001; Alan, Atalay and Crossley, 2012; Jørgensen, 2016).

To reconcile my results with existing studies, I investigate how ignoring precautionary motives increases the estimated importance of children in the Danish data. As mentioned above, log-linearized Euler equation estimates (which assumes away precautionary motives) are in the same magnitude as found in UK and US data. This is also the case when simulating data from the model with small or no effects of children on the marginal utility of consumption.[4] Furthermore, I find that the full-solution estimates of the effect of

[2]Thurow (1969) shows how impatient consumers facing credit constraints can generate a hump in the consumption age profile and Heckman (1974) shows how non-separability between consumption and leisure could be yet another explanation for the observed consumption age profiles.

[3]Aggregating to cohort levels, the regression coefficient between log consumption growth and the growth in the number of children is around 0.2 in the Danish data and 0.1 in the simulated data based on the estimated model. The exact magnitudes differ slightly across educational groups. Attanasio *et al.* (1999) report a regression coefficient of around 0.2 from synthetic cohorts based on the CEX (using IV estimation in a slightly different setup).

[4]This result stems from the fact that the inability or unwillingness of risk averse consumers to borrow against future income growth produces a positive correlation between income growth and consumption growth. This is well known and the excess sensitivity of consumption to income has often been attributed to credit constraints. However, because children tend to arrive when income (and thus consumption) grows the most, this will produce a positive

children on the marginal utility of consumption increase towards those from Euler equation estimation when I estimate a version of the model with much less severe precautionary motives.

I interpret the results as indicating that precautionary motives play a larger role than children in explaining observed consumption behaviour. Gourinchas and Parker (2002) and Cagetti (2003) also find a significant role for precautionary motives for especially young households.[5] I find (like Gourinchas and Parker, 2002) that young households acts as buffer-stock savers and older households save for retirement and other life cycle motives. I find that a negligible part of saving is due to children. The current results should be viewed in light of the retained simplifying assumptions applied for tractability of the estimation procedure.[6] The results could be viewed as indicating that children is not *needed* to rationalize the Danish data when households face realistic precautionary motives.

The results have implications for a wide range of fields in economics if children have a smaller impact on the marginal utility of consumption than previously believed. It is common practice to calibrate economic models using externally estimated parameters and the effect of children on the marginal utility of consumption (or equivalence scales) are often calibrated from external microeconometric sources. For example, Cagetti (2003) uses the estimates from Attanasio *et al.* (1999) when analysing wealth accumulation and precautionary savings over the life cycle and Scholz, Seshadri and Khitatrakun (2006) use equivalence scales from Citro and Michael (1995) when analysing retirement savings of American households.

The results are also related to a growing strand of literature estimating models of intertemporal behaviour related to household demographics. While I estimate a model in which the age of all (three) children can affect the marginal utility of consumption, all existing studies in this strand of literature estimate models in which children can affect households in much more restricted ways.[7] Motivated by the results in Browning and Ejrnæs (2009), I allow the age of all children to be potentially important.

The rest of the paper proceeds as follows. In the next section, I augment a standard imperfect markets life cycle model with the potential presence of children. Section III presents the Danish administrative registers and section IV discusses how some model parameters are calibrated to the Danish environment. Section V presents the estimation results and model fit while section VI investigates the robustness of the results. Finally, I conclude.

bias of (log-linearied) Euler equation estimators when using synthetic cohort panels. I elaborate further on this in a companion paper (Jørgensen, 2016).

[5]While they include demographic variation over the life cycle through a time-varying taste shifter, identical for all households within an age-group, they remove demographic variation *across* households.

[6]To allow for an arbitrary children, age and scale effect on the marginal utility of consumption, I have made simplifying assumptions. As in many existing studies, fertility is assumed to be exogenous, and labour supply and house purchases is not part of the economic model. I do, however, show that the effect of children on the marginal utility only increases slightly when estimating the model on a sub-sample of permanently employed home owners.

[7]For example, Love (2010) assumes that children arrive with 2-year intervals and Sommer (2014) assumes that there are two types of children: children living at home and children who have left the household. The model in Hong and Ríos-Rull (2012) is independent of the age of dependents while in the recent working papers by Blundell, Dias, Meghir and Shaw (2016) and Adda, Dustmann and Stevens (forthcoming) only the age of the youngest child matters.

II. A Model of Consumption in the Presence of Children

The theoretical framework used throughout this study is purposely very similar to the underlying models in, for example, Attanasio *et al.* (1999) and Browning and Ejrnæs (2009). The model is based on the buffer-stock model pioneered by Deaton (1991) and Carroll (1992) and first structurally estimated in Gourinchas and Parker (2002). A novelty of this study is that I augment the standard buffer-stock model with the potential presence of children and allow the marginal utility of consumption to be affected by the number and age of all children.

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 the level of consumption, C_t, that solves the optimization problem

$$\max_{C_t} \mathbb{E}_t \left[\sum_{\tau=t}^{T_r-1} \beta^{\tau-t} v(\mathbf{z}_t;\theta) u(C_\tau) + \beta^{T_r-t} V_{T_r}(M_{T_r}, \mathbf{z}_{T_r}) \right], \tag{1}$$

where utility is CRRA, $u(C_t) = C_t^{1-\rho}/(1-\rho)$ and $v(\mathbf{z}_t;\theta)$ is a taste shifter in which θ is the loadings on the number and age of children, contained in $\mathbf{z}_t$. As most of the existing literature, I follow Attanasio *et al.* (1999) and let children affect the *marginal value* of consumption.[8]

Households solve equation (1) 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 in beginning of period t and Y_t is beginning-of-period income. Retirees are not allowed to be net-borrowers, $A_t = M_t - C_t \geqslant 0, \forall t \geqslant T_r$, while working households can borrow up to a fraction of their permanent income $A_t \geqslant -\kappa P_t \forall t, \kappa \geqslant 0$.

Income process

Income follows a stochastic process when working,

$$\begin{aligned} Y_t &= P_t \varepsilon_t, \forall t < T_r, \\ P_t &= G_t P_{t-1} \eta_t, \forall t < T_r, \end{aligned}$$

where P_t denotes permanent income, G_t is real gross permanent income growth, η_t is a mean one permanent income shock, $\log \eta_t \sim \mathcal{N}(-\sigma_\eta^2/2, \sigma_\eta^2)$, and ε_t is a mean one transitory income shock taking the value μ with probability $\wp$ and otherwise log normal,[9]

$$\varepsilon_t = \begin{cases} \mu & \text{with probability } \wp \\ (\tilde{\varepsilon}_t - \mu\wp)/(1-\wp) & \text{with probability } 1-\wp \end{cases}$$

[8] Alternatively, the household composition could be included as a scaling of resources or consumption (equivalence scaling), as done in, for example, Fernández-Villaverde and Krueger (2007). See Bick and Choi (2013) for an analysis of different approaches to and implied behaviour from inclusion of household demographics in life cycle models.

[9] This specification of the transitory income shock follows that of Carroll (1997) and ensures that $\mathbb{E}_t[\varepsilon_{t+1}] = 1$, regardless of the values of μ, $\wp$ and $\sigma_{\tilde{\varepsilon}}^2$. Note, that by setting $\wp = 0$ the specification is close to that in Carroll (1997) where I multiply $\tilde{\varepsilon}$ with the factor $(1-\wp)^{-1}$ to ensure a mean of one while Carroll (1997) p.6) is not explicit about how he ensures that $\mathbb{E}_t[\varepsilon_{t+1}] = 1$.

$$\log \tilde{\varepsilon}_t \sim \mathcal{N}(-\sigma_\varepsilon^2/2, \sigma_\varepsilon^2).$$

When retired, the income process is a deterministic constant fraction $\varkappa \leqslant 1$ of permanent income at retirement, $Y_t = \varkappa P_{T_r}, \forall t \geqslant T_r$. This way of modelling the replacement rate in retirement is fairly standard and a similar approach is also applied in, for example, Kaplan and Violante (2010) and Blundell, Low and Preston (2013). While the latter assumes a replacement rate of zero, the former study implements a benefit system aimed at approximating the US pension scheme.[10]

Fertility

Households are fertile from age 15 to 43 and children arrive with a known probability distribution depending on the age of the wife, educational attainment and the number of children already present in the household.[11] Children leave home at age 21 and do not influence household consumption in subsequent periods. Households can have at most three children for which the age is contained in $\mathbf{z}_t$,

$$\mathbf{z}_t = (\text{age of child } 1_t, \text{ age of child } 2_t, \text{ age of child } 3_t) \in \{\text{NC}, [0, 20]\}^3,$$

where NC refers to 'no child' and the oldest child is denoted child one, the second oldest child as child two and the third oldest child as child three.

A novelty of this study is that I keep track of the age of *all* (three) children inside the household. To the best of my knowledge, this has not previously been done in dynamic models of intertemporal consumption and savings behaviour. To circumvent the computational cost of keeping track of the age of all children, strict assumptions on the timing of children are typically imposed.[12] Knowing the age of each child is, however, necessary to allow for an arbitrary child, age and scale effect of children on the marginal utility of consumption. To reduce the otherwise very large state space, I use the imposed ordering of children in $\mathbf{z}_t$: Because children leave the household at age 21, $\mathbf{z}_t$ evolves according to

$$\mathbf{z}_{1,t+1} = \begin{cases} \tilde{\mathbf{z}}_{2,t+1} & \text{if } \tilde{\mathbf{z}}_{1,t+1} \geqslant 21 \\ \tilde{\mathbf{z}}_{1,t+1} & \text{else} \end{cases}$$

$$\mathbf{z}_{2,t+1} = \begin{cases} \tilde{\mathbf{z}}_{3,t+1} & \text{if } \tilde{\mathbf{z}}_{1,t+1} \geqslant 21 \\ \tilde{\mathbf{z}}_{2,t+1} & \text{else} \end{cases}$$

$$\mathbf{z}_{3,t+1} = \begin{cases} NC & \text{if } \tilde{\mathbf{z}}_{1,t+1} \geqslant 21 \\ \tilde{\mathbf{z}}_{3,t+1} & \text{else} \end{cases}$$

[10] Blundell, Low and Preston (2013) also implement a version with $\varkappa = 0.5$ in their robustness checks. Kaplan and Violante (2010) p. 65) write that '*Benefits are then scaled proportionately so that a worker earning average labour income each year is entitled to a replacement rate of 45 percent...*'.

[11] Love (2010); Hong and Ríos-Rull (2012) and Blundell, Dias, Meghir and Shaw (2016) also assume that children arrive probabilistically. The same is true in Adda, Dustmann and Stevens (forthcoming) although, in their model, households essentially has limited control over contraceptive effectiveness.

[12] To reduce the computational complexity, Scholz and Seshadri (2009) assume that households choose the *number* of children to have at age 27, such that all children arrive simultaneously. Love (2010), on the other hand, assumes that children arrive with 2-year intervals, and Sommer (2014) assumes that there are two types of children: children living at home and children who have left the household. Alternatively, Blundell, Dias, Meghir and Shaw (2016) assumes that only the youngest child matters.

where the age of the children is re-arranged if the oldest child leaves the household. The age of each child (denoted with subscript j) evolves according to

$$\tilde{\mathbf{z}}_{j,t+1} = \begin{cases} 0 & \text{if } b_{t+1} = 1, \mathbf{z}_{j,t} = NC \text{ and } \mathbf{z}_{<j,t} \neq NC \text{ and } \mathbf{z}_{>j,t} = NC \\ \mathbf{z}_{j,t} + 1 & \text{if } \mathbf{z}_{j,t} \neq NC \\ NC & \text{else} \end{cases} \tag{3}$$

where the first case refers to a childbirth ($b_{t+1} = 1$) that occurs with an age and household size dependent probability, $\pi_{t+1}(\mathbf{z}_t)$ and the last two conditions in the first case ensures the relevant ordering of the children.[13] For example, if a household has one child and has another in the next period, we have that $b_{t+1} = 1$ and $\mathbf{z}_{1,t} \neq NC, \mathbf{z}_{2,t} = \mathbf{z}_{3,t} = NC$ and, thus $\tilde{\mathbf{z}}_{1,t+1} = \mathbf{z}_{1,t} + 1, \tilde{\mathbf{z}}_{2,t+1} = 0$ and $\tilde{\mathbf{z}}_{3,t} = NC$.

While I assume a stochastic process for childbirth, an alternative deterministic fertility process assumed in, for example, Gourinchas and Parker (2002); Cagetti (2003) and Browning and Ejrnæs (2009), is sometimes adopted. In their framework, households know with perfect foresight at the beginning of adulthood when and how many children they will have.[14]

Retirement

Following Gourinchas and Parker (2002), among others, V_{T_r} is a retirement value function, summarizing in a parsimonious way *all* postretirement savings motives. The retirement value function is assumed to be given as

$$V_{T_r}(M_{T_r}, \mathbf{z}_{T_r}) = \gamma \max_{C_j \in [0, M_j] \forall j = T_r, .., T} \left\{ v(\mathbf{z}_{T_r}; \theta) u(C_{T_r}) + \sum_{s=T_r+1}^{T} \beta^{s-T_r} v(\mathbf{z}_s; \theta) u(C_s) \right\}$$

where γ is a parsimonious way of adjusting for all postretirement motives such as survival and income uncertainty, retirement benefits and bequest motives. In the following, I refer to γ as a 'retirement motive' although it summarizes *all* potential ignored postretirement motives and should not be interpreted as a parameter measuring the importance of the retirement motive.

The specification in, for example, Kaplan and Violante (2010) and Blundell *et al.* (2013) is similar to assuming $\gamma = 1$. Naturally, and as I show in the robustness section, γ is closely related to the replacement rate κ and thus provides an additional degree of freedom in the model to correct for any potential miss-calibration of the replacement rate in retirement.

III. The Danish register data and descriptive analysis

I use high-quality Danish administrative registers covering the entire population from the period 1987–96. All information are based on third party reports with little additional self-reporting. All self-reporting are subject to possible auditing giving reliable longitudinal information on household characteristics, assets, liabilities and income.

[13] I use '$<j$' and '$>j$' subscripts to denote all indices below and above j respectively.

[14] In the robustness section, I estimate this alternative model and find similar results as in the baseline stochastic fertility model.

While the estimator outlined in section V match observed and predicted net-worth to uncover preference parameters, most existing studies use non-durable consumption (or food consumption) to analyse how children affect the marginal utility of consumption. To relate the Danish data, and the subsequent results, to the existing literature, I impute household consumption using a simple budget approach, $C_t = Y_t - \Delta A_t$, where Y_t is disposable income, A_t is end-of-period net wealth and ΔA_t proxies savings. This imputation method is evaluated on Danish data in Browning and Leth-Petersen (2003) and found to produce a reasonable approximation to survey responses in the Danish Family Expenditure Survey (DES). The resulting consumption measure will, however, include some durables such as home appliances.

Disposable income includes all labour market and non-labour market income net of all taxes. Transfers, such as child benefits and unemployment benefits, are also included to ensure that disposable income accurately measures the flow of resources available for consumption. Net wealth consists of stocks, bonds, bank deposits, cars, boats, house value for home owners and mortgage deeds net of total liabilities. The house value is assessed by the tax authorities for tax purposes and is included because it is impossible in the Danish registers to determine exactly which mortgages are related to the house and which are not.

Pension wealth is not included in the wealth measure. Information on pension accounts are not available for most of the cohorts studied here and the resulting net wealth is, therefore, slightly underestimated. However, pension funds are rather illiquid before retirement and only few withdraw pension funds prematurely. Heavy taxation leaves only 40% of prematurely withdrawn pension funds available for consumption purposes. Prematurely withdrawn pension funds are included in the measure of disposable income and since I focus on preretirement behaviour, exclusion of pension wealth is expected to have negligible effects on the results.

I restrict attention to stable married or cohabiting couples in which the wife is between 25 and 59 years old. This is to mitigate issues regarding educational and retirement choices. To increase homogeneity of households, I restrict the spousal age difference to be no more than 4 years and exclude households with more than three children.[15] Only households with children born when the wife was aged 15 through 43 are included in the analysis and households in which one adult is self-employed or out of the labour market are excluded from the analysis. Extreme or missing observations are also excluded from the analysis leaving an unbalanced panel of 201,618 households observed in at most nine periods with a total of 1,281,952 household-time observations. Financial measures are converted into 1987 US prices through regression and using an exchange rate of 5.5 DKK/USD. Households are classified as high skilled if either member holds at least a bachelor degree. Appendix B in the online supplemental material contains details on the sample selection criteria and their impact on the sample.

Table 1 presents means and SD for selected measures. The wives are on average 40.7 years old with on average 1 year older husbands. The Danish couples used here have on average 1.82 children at age 37. The imputed consumption measure is on average around

[15]This is exclusively for computational tractability of the economic model. Keeping track of the possible combinations of more than three children which can each be aged 0 through 21 would significantly increase computation time.

TABLE 1

Descriptive statistics

	Mean (SD)
Age, wife	40.749 (8.213)
Age, husband	41.978 (8.295)
Wealth	38,967 (49,635)
Non-durable consumption†	34,713 (18,674)
Disposable income	36,166 (6,712)
Number of children‡	1.817 (0.757)
High skilled	0.336 (0.472)
Number of observations	1,281,952

Notes: Year effects are removed by regression and financial measures are in 1987 US dollars using an exchange rate of 5.5 DKK/USD. †Non-durable consumption is imputed as disposable income net of changes in the wealth stock, as proposed by Browning and Leth-Petersen (2003). ‡The average number of children in households in which the wife is 37 years old. Based on 54,118 households with at most three children.

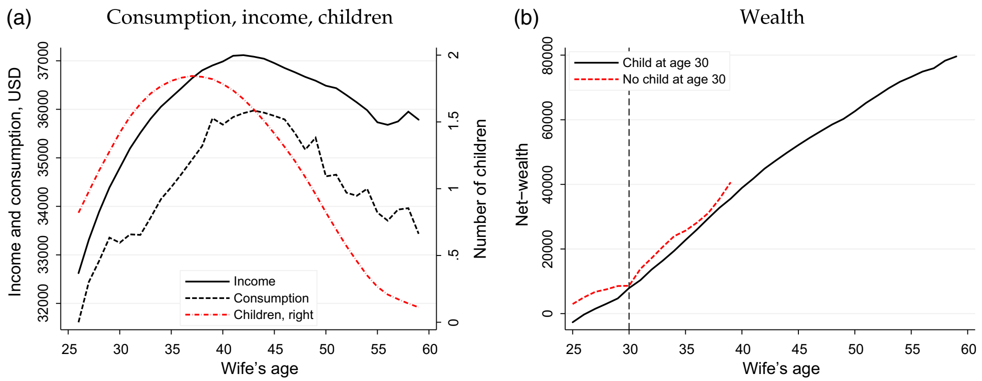

Figure 1. Age profiles of income, consumption, children and wealth.
Notes: Figure 1a illustrates average age profiles of income, non-durable consumption and number of children (on the right y-axis) for the Danish households. Year effects are removed by regression. Danish figures are in 1987 US dollars using an exchange rate of 5.5 DKK/USD. Non-durable consumption is imputed as the sum of disposable income net of changes in the wealth stock, as proposed by Browning and Leth-Petersen (2003). The right panel illustrates the average age profiles of household net-wealth of households who have no children at age 30 and households who have at least one child at age 30

35,000 USD only around a thousand dollars short of the average household income of 36,000 USD. Around 34% of households are classified as high skilled.

Panel (a) in Figure 1 illustrates the life cycle profiles of income, non-durable consumption and the number of children (right axis). Income and consumption profiles peak around the mid-40s while the average number of children peak around age 37. Panel (b) in Figure 1 illustrates wealth profiles of households who have children at the age of 30 and households who are childless at the age of 30. The two wealth profiles are almost identical. If risk averse households are either unable or unwilling to borrow they should accumulate wealth

TABLE 2

Log consumption growth regressions

	Household panel		*Cohort panel*		*ABMW*†
	OLS	*IV*	*OLS*	*IV*	
Δ#*kids*	0.042	1.782	0.182	0.490	0.212
	(0.005)	(0.090)	(0.112)	(0.230)	(0.101)
Δ#*adults*	0.031	2.051	0.109	0.558	0.449
	(0.005)	(0.107)	(0.188)	(0.334)	(0.144)
Constant	0.006	−0.068	0.012	−0.045	0.045
	(0.001)	(0.003)	(0.011)	(0.015)	(0.010)
Obs	1,027,072	396,074	264	116	256

Notes: 'Household panel' refers to estimation results using the original household information and 'Cohort panel' refers to estimation results from synthetic cohort panels. Robust standard errors in brackets. The instrument set consists of (i) second and third lagged changes in the number of children and adults, (ii) second to fourth lagged changes in log consumption and log income and (iii) a polynomial in age of the wife. In the cohort panel regressions the instruments are on the cohort level.†ABMW refers to the results reported in Table 1 of Attanasio *et al.* (1999) using the CEX to construct synthetic cohort panels (using 5-year bands). They used quarterly data and included seasonal dummies and the interest rate in the regressions. The latter should identify ρ^{-1} if there is only negligible income uncertainty and perfect markets (no constraints on borrowing). They estimate $\rho \approx 1.57$ suggesting an estimate of $\theta = 0.212 \cdot 1.57 \approx 0.33$.

in anticipation of the arrival of children in the future and consume out of that wealth when they subsequently have children. This is not what the raw data suggest.

The Danish data facilitates the implemented estimator presented below in section V. Most existing surveys, as those used in, for example, Gourinchas and Parker (2002); Attanasio *et al.* (1999) and Cagetti (2003), do not contain quality panel data on income, wealth and household characteristics and would thus not be as well suited for the proposed estimator. Below, I investigate the similarities between the Danish register data and the results reported in existing studies.

Consumption around childbirth

Table 2 reports log consumption growth regressions with changes in the number of children and adults included as regressors. The purpose is to illustrate that applying commonly used estimators on the Danish data produce a correlation between log consumption growth and the growth in the number of children similar to existing studies. Particularly, the goal is not to estimate a statistically well-specified equation but rather to follow the approach in Attanasio *et al.* (1999) as closely as possible. Their estimates (and SD) are reported in column five of Table 2 for easy comparison. I include in Appendix C in the online supplemental material additional log-linearized Euler equation estimation results for alternative samples and instrument sets. The qualitative results remain unchanged across alternative specifications.

Under the rather strong assumptions of no income uncertainty and no constraints on borrowing, the coefficient in front of the growth in the number of children is $\rho^{-1}\theta$ (Jørgensen, 2016). In turn, the estimates below can be interpreted as (a scaled) effect of children on the marginal utility of consumption from a model in which households *do not* face precautionary savings motives.

The first two columns report OLS and IV estimates using the household panel information. Following Attanasio, Banks, Meghir and Weber (1999), the instruments used are (i) second and third lagged changes in the number of children and adults, (ii) second to fourth lagged changes in log consumption and log income and (iii) a polynomial in age of the wife.[16] The third and fourth columns report OLS and IV estimates from regressions on synthetic cohort panel data. Although individual panels are available for the Danish households, most existing studies (including Attanasio *et al.*, 1999 and Gourinchas and Parker, 2002) use repeated cross section survey data (such as the CEX) to construct synthetic cohort panels (Browning, Deaton and Irish, 1985) and estimate log consumption growth equations assuming homogeneity within cohorts. I have also collapsed the Danish data into cohort panels (1 year bands) and present the same estimation results using this cohort panel data for estimation.

Focusing on the cohort panel results, the correlation between log consumption growth and changes in the number of children is very close to (and even above) the estimates reported in the seminal paper by Attanasio *et al.* (1999) using the CEX. Table 2 shows that log consumption grows with around 0.49 when a child (younger than 18) arrives and around 0.56 when an additional adult (18 or older) is present compared to 0.21 and 0.45, respectively, reported in Attanasio *et al.* (1999). This suggests that the Danish imputed consumption measure is not fundamentally different than the data typically used in existing studies. The estimates are rather imprecise and the coefficient for Δ#*kids* in Attanasio *et al.* (1999) is only barely significant at a 5% confidence level.

The OLS estimates from the household panel data are significantly smaller, close to the log-linearized Euler equation estimates for food consumption in the panel study of income dynamics (PSID), reported in Alan and Browning (2010).[17] As argued in Jørgensen (2016), the OLS estimator on the household-level panel data is potentially downwards biased if risk averse households face potentially binding credit constraints or are otherwise unwilling to borrow.[18] On the other hand, the synthetic cohort panel approach is potentially upwards biased since the average number of children might proxy for precautionary motives. This motivates why I allow for all these motives simultaneously when estimating the economic model.

Childless households

Figure 2 presents consumption and disposable income age profiles for households with at least one child and childless households at completed fertility. Childless households are

[16] Attanasio, Banks, Meghir and Weber (1999) use similar instruments but also include lagged consumption and income growth along with lagged interest rates when estimating both ρ and $\rho^{-1}\theta$ from log consumption growth regressions.

[17] Kalwij (2005) finds a negative and insignificant effect using the Dutch Socio-Economic Panel (SEP).

[18] Imagine the extreme case in which a household is constrained in all periods, and thus consumes all their income. Even if such a household would have liked to increase consumption when a child arrives, they cannot due to the lack of wealth accumulation prior to the birth of their child. Therefore, log consumption growth would be unaffected by the arrival of a child. On the other hand, log consumption growth will track log income growth perfectly (excess sensitivity).

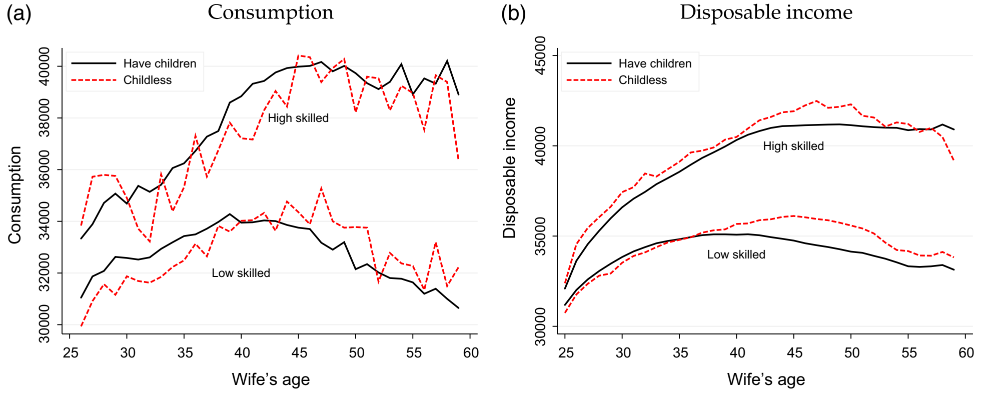

Figure 2. Consumption and disposable income in childless households and households with children at completed fertility.
Notes: Figure 2 illustrates consumption and income age profiles for households with children and childless households at completed fertility when the wife is aged at least 40. Childless households are identified as households in which the wife is not the mother to a child in 2010. If the wife is not registered as a mother in 2010, Figure 2 assumes that no children will arrive in that household

identified as households in which the wife is not registered as the mother to a child in 2010.[19] If the wife is not registered as a mother in 2010, I assume that the household will remain childless. This assumption is not overly restrictive since the youngest household in the data (aged 26 in 1996) will be 40 years old in 2010. Only few households have children at this age. Childless women could have adopted children or foster children from the current husband's previous marriage(s). Therefore, I restrict childless households to those without children registered as living at the same address as the couple at any point in the observed years.[20]

Childless households have almost identical income and consumption age profiles as those who have children at some point in their life. Income of childless households grows with a similar rate as households who have children until the wife is 40 and 45 years old for low- and high-skilled households respectively. Income continues to grow for around five additional years for childless households. Although there is few[21] childless households and the resulting age profiles are rather noisy, this pattern suggests that previous results that the number and age of children can completely explain the hump in consumption (Browning and Ejrnæs, 2009) might proxy for other consumption/savings motives.[22]

[19] Virtually *all* childbirths after 1942 are matched to their mother. Only children born between 1 January 1942 and 31 December 1972 who either died or permanently immigrated to another country before 1 January 1979 is *not* included in the Danish fertility registers. The youngest potential births used to identify childless households are in 1987 – (59 – 12) = 1940, assuming that fertility begins at age 12. In turn, almost *all* mothers used here will be matched to their children, if they have any.

[20] This does, however, not rule out the possibility that households defined as childless do foster children registered to be living at another address.

[21] There is more than 500 childless households in each age and educational group. Only high-skilled childless households are fewer than 500 after age 50 and at a minimum of around 200 high-skilled childless households at age 59.

[22] If childless households have different preferences than households who have children, the age profiles might be similar due to these differences. Furthermore, there might be significant differences in the consumption composition across the two groups. Unfortunately, the Danish data used herein do not allow me to decompose the consumption measure.

Unfortunately, I do not observe whether childlessness is caused by infertility or an active choice. Ideally the split should be between households who *intent* and do *not intent* to have children. Infertile households who intended to have children but realized their infertility late in their life would likely have accumulated wealth *as if* a child could arrive in the future. Infertility will, thus, tend to produce similar consumption age profiles for households who have children and childless households and the figures are therefore only suggestive. See for example Ejrnæs and Jørgensen (2006) on the saving behaviour around intended and unintended childbirths.

IV. Calibrated parameters

To keep the estimation procedure tractable, I reduce the number of parameters to be estimated by calibrating some parameters in a first step. Following Gourinchas and Parker (2002), the real gross interest rate is fixed at $R = 1.03$. The replacement rate in retirement, $\varkappa$, is calibrated to 90% based on the median couple reported in The Danish Ministry of Finance (2003). This implies a rather high level of income from transfers post retirement and stems from generous public transfers and private pension funds.[23]

The social security system in Denmark seems compatible with a 10% risk of household income being reduced to 30%. Danish households are allowed to be net-borrowers by 60% of annual permanent income. These three values ($\kappa = 0.6$, $\mu = 0.3$ and $\wp = 0.01$) are somewhat arbitrary and have been chosen to provide a reasonable fit of the model for the bottom distribution of resources. Figure 5 below illustrates this by plotting within-percentile average consumption-income ratios against the household resources (also normalized by income). There is substantial variation in consumption in the bottom distribution of resources for particularly young households and the calibrated parameters (along with the estimated preferences in Table 4) provide a good fit on average.

The permanent and transitory income shock variances in Table 3 are estimated following the approach in Meghir and Pistaferri (2004). First, I run a regression of income on year dummies and the resulting log residual income, $\tilde{y}_t$, is used to calculate the permanent and transitory income shock variances as $\hat{\sigma}_\eta^2 = \text{cov}(\Delta\tilde{y}_t, \tilde{y}_{t+1} - \tilde{y}_{t-2})$ and $\hat{\sigma}_\varepsilon^2 = -\text{cov}(\Delta\tilde{y}_t, \Delta\tilde{y}_{t+1})$. The permanent income shocks are found to be more volatile for high-skilled households, a robust result in the literature. The variance of transitory income shocks is, however, often found to be lower for high-skilled households while I find the opposite here.

The income variances of Danish households are somewhat smaller than those typically estimated for the US.[24] This is most likely due to the generous social welfare system and progressive taxation in Denmark. Denmark has a relatively high 'minimum wage' of around \$20 per hour (in 2010) reducing the volatility in permanent and transitory income shocks compared to, for example, the US.[25] The Danish tax system is one of the most progressive tax schedules in the world with a marginal tax rate of more than 60% in 2010 for top tax

[23] In the robustness checks, I also estimate a version in which $\varkappa = 0.7$.

[24] Blundell, Pistaferri and Preston (2008) report $\sigma_\eta^2 \in [0.0057, 0.0333]$ and $\sigma_\varepsilon^2 \in [0.0190, 0.0753]$ depending on the combination of year, cohort and educational background, using the PSID. Gourinchas and Parker (2002), also using the PSID, calibrate $\sigma_\eta^2 = 0.0212$, $\sigma_\varepsilon^2 = 0.0440$.

[25] Strictly speaking, there is no minimum wage in Denmark but rather strong trade unions and the collective agreements determine the minimum wage for members within a certain union.

TABLE 3

Permanent and transitory income shock variances

	All Est (SE)	*Low skilled Est (SE)*	*High skilled Est (SE)*
$\hat{\sigma}_\eta^2$	0.0054 (0.000096)	0.0049 (0.000113)	0.0062 (0.000173)
$\hat{\sigma}_\varepsilon^2$	0.0072 (0.000156)	0.0059 (0.000167)	0.0095 (0.000315)

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

earners. Around 40% were top tax earners in 2010. The progressive tax system reduces the dispersion in disposable income significantly. Finally, administrative registers also tend to be less noisy compared to surveys (Browning and Leth-Petersen, 2003), reducing the transitory income shock variance.

The income growth rate, G_t, can be estimated by taking logs of the income process specified in section II and averaging over individuals, for a given age, t,

$$\frac{1}{N}\sum_i^N \Delta \log Y_{it} = \log G_t + \frac{1}{N}\sum_i^N \log \eta_{it} + \frac{1}{N}\sum_i^N \Delta \log \varepsilon_{it}.$$

Re-arranging and noting that the second term converges to $-\frac{1}{2}\sigma_\eta^2$ and the last term converges to zero as $N \to \infty$, gives a consistent estimate of the income growth rate as

$$\hat{G}_t = \exp\left(\frac{1}{N}\sum_i^N \Delta \log Y_{it} + \frac{1}{2}\hat{\sigma}_\eta^2\right). \tag{4}$$

Figure A1 in the supplemental material reports the estimated income growth rate profile for high- and low-skilled Danish consumers. As expected, high-skilled households have much higher income growth compared to low skilled. I use a moving average smoothed income growth rate throughout.

Permanent income, P_t, is uncovered by applying the Kalman Filter to each household's income process as described in Appendix F of the supplemental material. The arrival rate of infants are estimated as a simple logit model with age dummies for each educational group and number of children already present in the household. Figure A2 in the supplemental material presents the resulting age profiles.

V. Estimation results

A maximum likelihood estimator

To estimate how children affect the marginal utility of non-durable consumption, I formulate a continuous version of the Nested Fixed Point (NFXP) estimation approach, suggested by Rust (1987). For a given set of K structural parameters, Θ, the model is solved recursively using the Endogenous Grid Method (EGM) proposed by Carroll (2006) for all combinations of household composition. This yields optimal consumption as a function of resources, permanent income and household composition, $\{C_t^*(M_t, P_t, \mathbf{z}_t|\Theta)\}_1^T$. Consult Appendix D in the online supplemental material for details on how I solve the model numerically and Appendix E for a discussion of identification.

Let $\{M_{it}, P_{it}, C_{it}, \mathbf{z}_{it}\}$ denote for household i in period t *observed* resources, permanent income, consumption and the age of at most three children. The $i = 1,\ldots,N$ households are observed in $t = 1,\ldots T_i$ periods. I assume that imputed non-durable (normalized) consumption in the Danish registers is observed with additive iid Gaussian measurement error, ξ,

$$c_{it} = c_t^*(M_{it}, P_{it}, \mathbf{z}_{it}|\Theta) + \xi_{it}, \xi_{it} \sim \mathcal{N}(0, \sigma_\xi^2)$$

where small letter variables denote normalized measures (e.g. $c_t = C_t/P_t$). Because resources can only be consumed or saved, $m_{it} = c_{it} + a_{it}$, the difference between the predicted and observed level of wealth in the data $\xi_{it} = a_t^*(M_{it}, P_{it}, \mathbf{z}_{it}|\Theta) - a_{it}$ is (assumed) iid Normally distributed. The (mean) log-likelihood function is, thus,[26]

$$\mathcal{L}(\Theta, \sigma_\xi) = -\frac{1}{2}\frac{1}{N\bar{T}}\sum_{i=1}^{N}\sum_{t=1}^{T_i}\left\{\log(2\pi\sigma_\xi^2) + \frac{(a_t^*(M_{it}, P_{it}, \mathbf{z}_{it}|\Theta) - a_{it})^2}{\sigma_\xi^2}\right\} \tag{5}$$

where $\bar{T} \equiv \max_i T_i$ is the maximum number of time periods available in the data.

Optimal behaviour is found numerically and the likelihood function in equation (5) is an *approximation* to the exact likelihood function. Fernández-Villaverde, Rubio-Ramírez and Santos (2006) show that as long as the numerical approximation converges to the *unique* exact solution, the approximated likelihood function converges uniformly to the exact likelihood function. This provides the strong result that parameters estimated by maximizing the approximate likelihood are consistent and asymptotically normally distributed.[27] Table A1 in the supplemental material reports mean (and SD) of θ estimates from 50 independent simulations in which measurement error is added with a known variance of one. The estimation approach uncovers the true parameter, θ_0, in even small samples.

To calculate the likelihood function, I 'look up' in the numerically found optimal consumption function (interpolate) for a particular level of resources, M_{it}, permanent income P_{it} and household composition, $\mathbf{z}_{it}$ for all household-age observations in the data. The availability of accurate income and wealth data is key to the implementation of the current estimator.[28] The implemented estimator can also be viewed as a *nonlinear least squares estimator* (NLLS), where the nonlinear function of parameters is found numerically using dynamic programming techniques. In the robustness section I report qualitatively unchanged estimates from NLLS and an alternative estimator where I minimize the absolute difference rather than the squared difference.

An alternative commonly used estimation approach, also implemented in, for example, Gourinchas and Parker (2002) and Cagetti (2003), is based on simulation techniques. These approaches can often be implemented with less available information in the data.

[26] Note, that I weight each household's contribution by the number of time-periods they are observed, T_i, and refer to resulting objective function as the 'mean' log-likelihood for this reason. In an earlier working paper version of the current paper, I was not explicit about the fact that I match wealth levels rather than consumption levels. Because of the assumed identity $m_{it} = c_{it} + a_{it}$ the two estimators differ only by the amount of available data (because c_{it} is based on current and lagged values of a_{it}).

[27] Ackerberg, Geweke and Hahn (2009) correct a result (Proposition 2) of Fernández-Villaverde *et al.* (2006) stating that for the *approximated* likelihood to converge to the exact one the approximation error should decrease faster than the increase in observations. Ackerberg, Geweke and Hahn (2009) reassuringly show that this is *not* the case.

[28] A similar ML estimator is implemented in Palumbo (1999) and Jørgensen (2013).

While Gourinchas and Parker (2002) envision a similar scenario with multiplicative measurement error in consumption levels, they resort to a method of simulated moments to estimate parameters of their model using the CEX due to a lack of '...quality panel data on consumption, assets, and income for individual households' Gourinchas and Parker (2002) p. 58).[29] For a given value of Θ, they (i) solve the model numerically for a discrete grid of normalized resources (as I do), (ii) simulate synthetic data using an initial distribution of state variables and (iii) minimize the weighted quadratic difference between the average observed and simulated age profile of log consumption. Cagetti (2003) implemented a similar estimator but preferred matching median age profiles rather than average age profiles.

Estimated preference parameters

Table 4 presents the estimation results for low- and high-skilled Danish households. Columns (1) report estimates from a model without any household composition effects ($\theta = 0$). Columns (2) report estimates using a functional form of the taste shifter similar to existing literature, $v(\mathbf{z}_t, \theta) = \exp(\theta \text{Number of children})$, and, finally, columns (3) report estimates from a flexible functional form,

$$\begin{aligned} v(\mathbf{z}_t, \theta) = 1 &+ \theta_{11}\mathbf{1}_{\{\text{Age of child } 1 \in [0,10]\}} + \theta_{12}\mathbf{1}_{\{\text{Age of child } 1 \in [11,21]\}} \\ &+ \theta_{21}\mathbf{1}_{\{\text{Age of child } 2 \in [0,10]\}} + \theta_{22}\mathbf{1}_{\{\text{Age of child 1 and } 2 \in [11,21]\}} \\ &+ \theta_{31}\mathbf{1}_{\{\text{Age of child } 3 \in [0,10]\}} + \theta_{32}\mathbf{1}_{\{\text{Age of child 1,2 and } 3 \in [11,21]\}}, \end{aligned}$$

allowing for an arbitrary children, age and scale effect.

Danish households are not significantly affected by the presence of children. Although formal likelihood ratio (LR) tests reject that $\theta = 0$, the estimated effect of children on the marginal utility of consumption is economically negligible and even negative for low-skilled households. Contrary to the results in Browning and Ejrnæs (2009), I find no significant effects of the age of children on the marginal utility of consumption ($\hat{\theta}_{j2} \approx \hat{\theta}_{j1} \forall j = 1, 2, 3$).

Recall that Table 2 illustrates a positive correlation between consumption growth and arrival of children similar to that found in data from the UK and US from log-linearized Euler equation estimation. Log-linearized Euler equation estimation is, however, biased if risk averse households are either unwilling or unable to borrow and changing households composition is likely to proxy for precautionary motives (Jørgensen, 2016). Since I find, when including income uncertainty, credit constraints and a retirement motive, that children have no effect on the marginal utility of consumption, I interpret this as suggesting that precautionary motives are more important than children in explaining the observed consumption age profile. In the robustness checks below, I estimate a version of the model in which households are allowed to borrow significantly more and face lower income risk and find that children increase marginal utility more when estimating this (potentially misspecified) model.

[29] Note, that because I match wealth, not consumption, consumption is not needed here.

TABLE 4

Estimated preference parameters

		Low skilled			High skilled		
		(1)	(2)	(3)	(1)	(2)	(3)
ρ	Risk aversion	2.316	2.363	2.385	2.639	2.626	2.634
		(0.041)	(0.036)	(0.043)	(0.057)	(0.062)	(0.063)
β	Discount factor	0.965	0.964	0.964	0.973	0.973	0.972
		(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)
γ	Retirement	1.454	1.492	1.491	1.251	1.245	1.265
		(0.018)	(0.018)	(0.020)	(0.022)	(0.023)	(0.025)
σ_ξ	Meas. error	0.468	0.468	0.468	0.490	0.490	0.490
		(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)
		Taste shifter: $v(\mathbf{z};\theta)=\exp(\theta'\mathbf{z})$					
θ	# of children		−0.017			0.004	
			(0.002)			(0.003)	
		Taste shifter: $v(\mathbf{z};\theta)=1+\theta'\mathbf{z}$					
θ_{11}	1. child $\leqslant$ 10			−0.004			−0.008
				(0.007)			(0.010)
θ_{12}	1. child > 10			−0.031			0.002
				(0.004)			(0.008)
θ_{21}	2. child $\leqslant$ 10			−0.034			−0.015
				(0.005)			(0.008)
θ_{22}	2. child > 10			−0.006			0.000
				(0.005)			(0.008)
θ_{31}	3. child $\leqslant$ 10			−0.005			0.022
				(0.009)			(0.012)
θ_{32}	3. child > 10			0.019			0.021
				(0.013)			(0.017)
$-\mathcal{L}(\Theta)$		0.46536	0.46533	0.46529	0.49868	0.49863	0.49862
$\max_j \lvert\partial\mathcal{L}(\Theta)/\partial\Theta_j\rvert$		$7.1e-6$	$1.5e-5$	$2.1e-5$	$8.3e-6$	$2.3e-5$	$1.8e-5$
LR [p-val]			67[<.001]	153[<.001]		57[<.001]	68[<.001]
Number of observations		851,249	851,249	851,249	430,703	430,703	430,703

Notes: Standard errors are based on the inverse of the hessian. Significant stars are not reported. Rather, the likelihood ratio (LR) test is a joint test of all taste-shifter parameters being zero, $\theta = 0$. In square brackets are reported the P-values from a χ^2 distribution with one or six degrees of freedom in columns (2) and (3), respectively.

The Danish welfare system provides free health care, free schooling and significant childcare subsidies. For example, childcare is heavily subsidized and approximately 70% of the cost of childcare is covered by the government. When children subsequently enter elementary school, the government covers completely the cost. Children older than 18 enrolled in at least high school receive a *monthly stipend* (in Danish ‘Statens Uddannelsesstøtte’, abbreviated ‘SU’) of around a thousand US dollars (5,839 Danish kroner) in 2014.[30] In turn, raising children in Denmark could be less costly than many other coun-

[30] When the child is younger than 20 the subsidy is subject to rather mild reductions depending on whether the child lives with its parents or alone and the parents income.

tries, contributing to the estimation of no children and age effects. Recall, however, that the correlation between log consumption growth and demographic changes are similar to other countries.

The remaining estimates are in the range of what is typically found. The relative risk aversion, ρ, is estimated to be around 2.4 for low-skilled and 2.6 for high-skilled households. The discount factor, β, is estimated to be around 0.97 for both educational groups. Gourinchas and Parker (2002) and Cagetti (2003) also estimate larger relative risk aversion parameters for high-skilled households. Gourinchas and Parker (2002), using the CEX, estimate $\rho \approx 0.87$ for high school graduates and $\rho \approx 2.29$ for college graduates while Cagetti (2003), using the Survey of Consumer Finances (SCF), estimates ρ to be around 3.5 for college graduates and 4.3 for high school graduates when matching median net-worth excluding housing. The measurement error variance is estimated to be around 0.22.

The γ, summarizing all postretirement motives, is estimated to be around 1.5 and 1.3 for low- and high-skilled respectively. The estimates should be viewed as summarizing a host of different potential saving and consumption motives related to life in retirement. Since the objective of this study has been to analyse how children affect behaviour prior to retirement, this parsimonious parametrization has been implemented. Taken at face value, the estimated γ's would suggest that the marginal utility of consumption should *increase* at retirement, contrary to most existing empirical evidence seem to suggest. Note, however, that if health, income and survival is uncertain in retirement while households gain utility from leaving bequests, consumers should (also prior to retirement) accumulate more wealth than the implemented model suggests. The observed level of wealth prior to and around retirement will result in a higher estimated value of γ in the parsimonious implementation here. I show in the robustness checks below that the postretirement motive does not seem to be driving the results and a lower replacement rate in retirement ($\varkappa$) would lead to a lower estimate of γ.

Inspired by the decomposition of savings motives in Gourinchas and Parker (2002), I decompose *simulated* savings into channels related to children, income risk and life-cycle/retirement motives.[31] Specifically, I define total saving for simulated individual i in period t as the discounted change in the simulated normalized wealth for the baseline estimated parameters,

$$S_{it} \equiv S_{it}(\hat{\theta}, \hat{\sigma}_{\varepsilon}^2, \hat{\sigma}_{\eta}^2, \hat{\wp}) = (a_{it}(\hat{\theta}, \hat{\sigma}_{\varepsilon}^2, \hat{\sigma}_{\eta}^2, \hat{\wp}) - a_{it-1}(\hat{\theta}, \hat{\sigma}_{\varepsilon}^2, \hat{\sigma}_{\eta}^2, \hat{\wp}))/R.$$

I denote simulated wealth as a function of a subset of the estimated parameters because I change the values of these parameters sequentially below, while leaving all other parameters, such as the risk aversion and discount factor, unchanged at their estimated values in columns (3) in Table 5.

To construct the counter factual level of saving in a scenario *without* effects of children, I solve the model while fixing $\theta = 0$ such that children have *no* effect on the marginal utility of consumption and simulate data from this version of the model. The corresponding simulated level of saving is, under this assumption, then denoted

$$S_{it}^c \equiv S_{it}(0, \hat{\sigma}_{\varepsilon}^2, \hat{\sigma}_{\eta}^2, \hat{\wp}).$$

[31] I simulate 100,000 households from age 25 through 59 and initialize all households without any wealth ($A_{i0} = 0$) and normalize permanent income to one $P_{io} = 1$.

TABLE 5

Log consumption growth and children: data and simulation

	Household panel		*Cohort panel*	
	Data	*Simulated*	*Data*	*Simulated*
Low skilled				
$\Delta kids$	0.035	0.054	0.223	0.106
	(0.006)	(0.003)	(0.087)	(0.007)
Constant	0.003	0.051	0.001	0.010
	(0.001)	(0.001)	(0.006)	(0.001)
Obs.	683,847	1,027,274	264	491
High skilled				
$\Delta kids$	0.052	0.087	0.262	0.128
	(0.007)	(0.005)	(0.097)	(0.009)
Constant	0.009	0.073	0.002	0.016
	(0.001)	(0.001)	(0.006)	(0.001)
Obs.	343,225	526,361	264	448

Notes: Robust standard errors in brackets. The dependent variable is log consumption growth ($\Delta \log C_{it}$) imputed from the Danish data and simulated from the estimated model using the parameters in columns (3) of Table 4.

Likewise, to construct the counter factual level of saving in a scenario without income uncertainty, I solve and simulate data from a version of the model in which income is perfectly deterministic, $\hat{\sigma}_{\varepsilon}^2 = \hat{\sigma}_{\eta}^2 = \hat{\wp} = 0$. The corresponding simulated level of saving under this scenario is

$$S_{it}^{i} \equiv S_{it}(\hat{\theta}, 0, 0, 0).$$

Figure 3 reports the average age profiles of these three measures. As expected from the very low estimated value of θ, the saving motive in relation to children is estimated to be extremely small. This is seen by the fact that the simulated saving from the model *without* any effects of children is effectively identical to (only slightly below) the baseline simulated saving over the life cycle.

If households did not face income uncertainty, on the other hand, they would tend to save substantially less, and consume their wealth in the beginning of life. This suggests that income uncertainty and precautionary motives are important drivers of wealth accumulation. While households are initiated with zero debt in the simulations, they are allowed to borrow up to 60% of permanent income. This explains the depletion of initial wealth, in this version of the model, until the beginning of the 30s. From the beginning of the 30s nothing is saved and wealth remains constant at the constraint until around the mid-40s where saving increases, due to the retirement motive.

Because saving motives related to children is almost absent, the average saving profile without income uncertainty, S_{it}^{i}, approximates other life cycle savings motives, such as retirement in the model. Likewise, the difference $S_{it} - S_{it}^{i}$ approximates what Gourinchas and Parker (2002) term buffer-stock saving. I find in Figure 4, like Gourinchas and Parker (2002), that buffer-stock saving motives dominate for young households while life cycle motives dominate for older households. The buffer-stock motives seem to dominate the

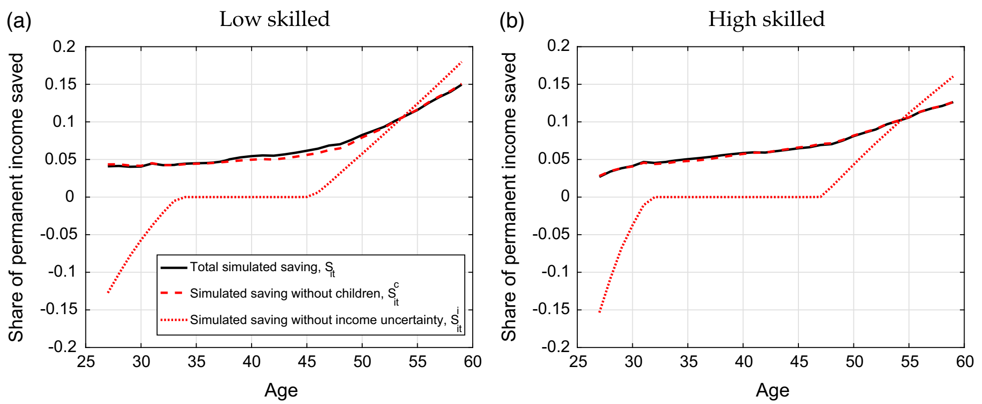

Figure 3. Decomposing saving motives.
Notes: The figure illustrates simulated age profiles of the amount saved (relative to permanent income) for the estimated model (solid), a version of the model in which children have no effect, $\theta = 0$, (dashed), and a version in which income is deterministic, $\sigma_{\varepsilon} = \sigma_{\eta} = \wp = 0$ (dotted). 100,000 households are simulated from 26 through 59, initialized with no wealth and permanent income of one

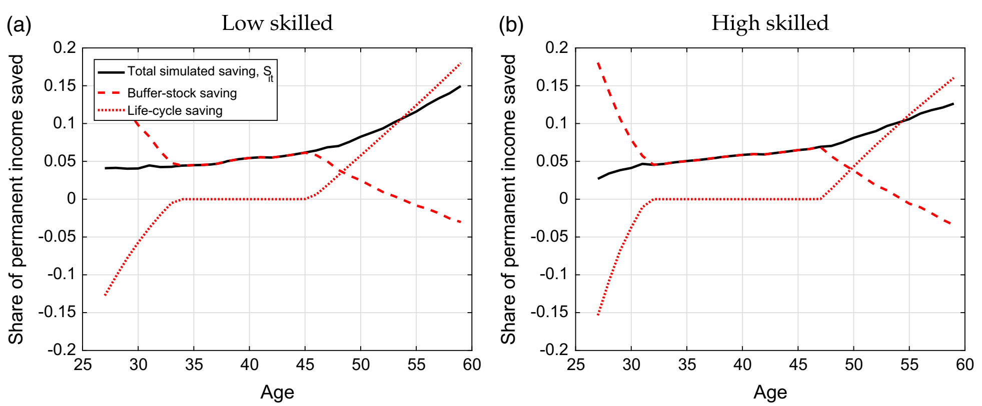

Figure 4. Decomposing saving motives: buffer-stock and life-cycle motives.
Notes: The figure illustrates simulated age profiles of the amount saved (relative to permanent income) for the estimated model (solid), saving due to buffer-stock motives, (dashed) and saving due to other life-cycle and retirement motives (dotted). 100,000 households are simulated from 26 through 59, initialized with no wealth and permanent income of one

retirement motive slightly longer than what Gourinchas and Parker (2002) find for the US, most likely stemming from the generous Danish retirement scheme. An important caveat here is that the model parameters used here was *not* chosen to match simulated age profiles.

Model fit

The model fits the Danish register data very well. Particularly, the estimated model can replicate correlations in the raw data, presented in Table 2 and reported in the existing literature. Although I find minuscule effects of children on the marginal utility of consumption,

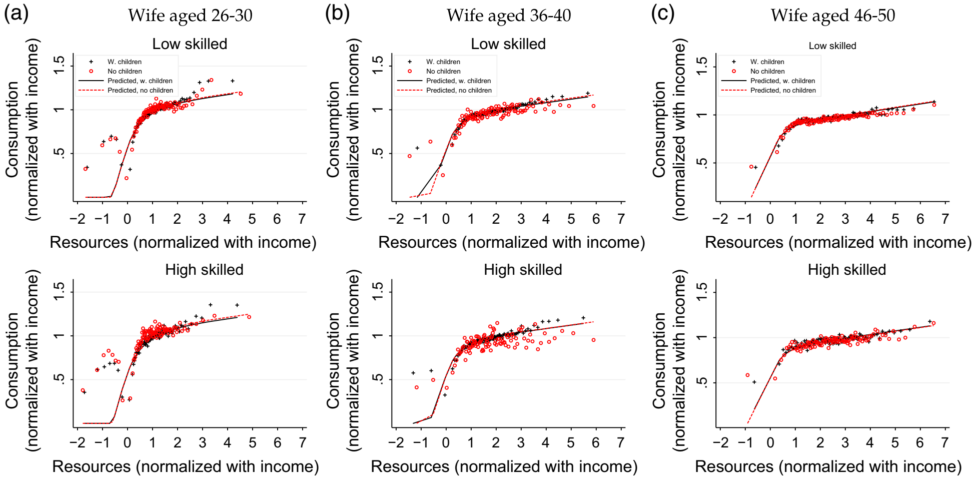

Figure 5. Consumption functions.
Notes: Figure 5 illustrates the average consumption in groups based on percentiles of available resources. Each dot represents a per cent of observations within each age and child/no child group. Average consumption and predicted consumption from the estimated model is plotted

Table 5 illustrates that simulated data from the estimated model generates similar correlations between log consumption growth and changing household composition.[32]

This is perhaps surprising at first glance. However, as I show in a companion paper (Jørgensen, 2016), Euler equation estimation fails to identify the effect of children on the marginal utility of consumption when households face potentially binding credit constraints. Particularly, using the cohort-average number of children as instrument (closely related to the synthetic cohort panel estimation) is found to be upwards biased due to a positive correlation between income and consumption growth when risk averse households are unable or unwilling to borrow. Therefore, I interpret the ability of the estimated model to replicate the estimates from log-linearized Euler equation estimations as a sign that precautionary motives (including credit constraints) are likely more important than children in explaining the observed consumption behaviour.

Figure 5 illustrates that the model fits the consumption functions remarkably well for childless households and households with children.[33] Childless households and households with children do not differ significantly. The figure also illustrates that the calibrated position of the credit constrain ($\kappa = 0.6$) and the low income shock ($\mu = 0.3$ with probability $\wp = 0.01$) provides a reasonably good fit for households in the lower end of the wealth distribution.

While the overall fit is very good, there is a tendency to a slight underestimation of consumption for households who are net borrowers. This is particularly evident for younger households, panel (a) in Figure 5, but does not seem to be due to underestimation of the

[32] Data are simulated by initializing households at the first non-missing observation in the data for wealth, permanent income and the age of all three children. Hence, I simulate the same number of households as is in the data in nine consecutive time periods. Because a significant fraction of households is observed in less than nine periods, the simulated data consist of more observations than the actual data cf. Table 5.

[33] Each point in Figure 5 is calculated as the average consumption within a given percentile of resources observed in the Danish data.

effect of children, θ. When estimating the model, consumption is set to zero for households with resources below the borrowing limit of 60% of permanent income. These households, however, consume significantly more than nothing, resulting in an underestimation of the level of consumption for these households. There does not seem to be a significant difference between childless households (red circle) and households with children (black plus) among this group with very low resources.

VI. Robustness

Weaker precautionary motives

The results indicate that precautionary motives are more important than children in explaining the observed life cycle profiles. I investigate the plausibility of this interpretation by estimating a model in which households face much less severe precautionary motives. Particularly, Table 6 reports estimation results from a situation in which consumers are allowed to borrow as much as desired while working and they face no low-income (unemployment) shock.[34] If introducing precautionary motives explain why I find minuscule effects of children, removing these motives should yield larger estimated effects of children, θ.

I find that children increase the marginal utility of consumption significantly more, with around 10–15%, when reducing the precautionary motives in the estimated model. The risk aversion parameter, ρ, is also estimated much higher because to rationalize the saving behaviour of households they must be much more risk averse now the borrowing limit is much higher and income uncertainty lower.

Measurement error distribution

The model parameters have been estimated through ML due to a Normality assumption on the measurement error. Rather than relying on this particular distributional assumption, I report in columns (1) in Table 7 the estimates from minimizing the absolute differences between observed and predicted wealth,

$$\mathcal{L}(\Theta) = \frac{1}{N\bar{T}} \sum_{i=1}^{N} \sum_{t=1}^{T_i} |a_t^*(M_{it}, P_{it}, \mathbf{z}_{it}|\Theta) - a_{it}|. \tag{6}$$

This essentially amounts to matching the median wealth level rather than the mean as suggested by Cagetti (2003). The estimated effect of children on the marginal utility of consumption is largely unchanged. As also found in Cagetti (2003), however, the risk aversion coefficients are estimated significantly lower when matching medians because medians are not affected by the wealth levels in the very top of the wealth distribution. When matching *average wealth* from a model that is not particularly well equipped to

[34] In the numerical implementation, I set $k = 10$ and $\wp = 0$. Since households are still not allowed to borrow while retired and income uncertainty is still present in the model through σ_η^2 and σ_ϵ^2, risk averse consumers will not accumulate infinite debt and the explicit borrowing limit, $k = 10$, is not binding. I fixed the retirement motive at their estimated values in Table 4 because I could not get reasonable estimates for both ρ and γ in this misspecified version of the model.

TABLE 6

Parameter estimates: unlimited borrowing and lower income uncertainty

	Low skilled		*High skilled*	
	(1)	*(2)*	*(1)*	*(2)*
ρ	5.636 (0.068)	5.699 (0.070)	7.177 (0.128)	7.112 (0.107)
β	0.958 (0.001)	0.960 (0.001)	0.948 (0.002)	0.952 (0.001)
σ_{ξ}	0.468 (0.000)	0.468 (0.001)	0.489 (0.001)	0.489 (0.001)
$v(\mathbf{z}_t;\theta)=\exp(\theta\mathbf{z}_t)$				
θ	0.073 (0.006)		0.087 (0.011)	
$v(\mathbf{z}_t;\theta)=1+\theta'\mathbf{z}_t$				
θ_{11}		0.154 (0.017)		0.135 (0.029)
θ_{12}		0.001 (0.020)		0.013 (0.019)
θ_{21}		0.094 (0.013)		0.158 (0.023)
θ_{22}		0.103 (0.012)		0.141 (0.022)
θ_{31}		0.080 (0.024)		0.203 (0.042)
θ_{32}		0.072 (0.035)		0.074 (0.054)
$-\mathcal{L}(\Theta)$	.46616	.46611	.49632	.49624
$\max_j \lvert\partial\mathcal{L}(\Theta)/\partial\Theta_j\rvert$	$8.3e-7$	$3.3e-5$	$2.7e-6$	$1.1e-5$
Number of observations	851,249		430,703	

Notes: Reported are estimates from a model with unlimited borrowing ($\kappa = 10$), $\wp = 0$ and $\mu = 0$. Other calibrated parameters are unchanged and the postretirement motive, γ, is fixed at their estimated values in columns (3) of Table 4. Standard errors in brackets are based on the inverse of the hessian.

describe the behaviour of the very wealthy, as the one herein, the only way to rationalize the observed wealth accumulation is to increase the risk aversion.

Throughout, I ignore the fact that permanent income, P_{it}, is estimated using the Kalman filter and kept fixed while estimating Θ.[35] If there is estimation bias in these parameters, the estimator would be affected. For example, the assumption of no serial correlation in the measurement errors might be invalid. To investigate whether this concern seems to be driving the results, I have implemented two alternative estimators.

First, rather than assuming that measurement error in wealth normalized by permanent income is normally distributed, I have implemented a nonlinear least squares type estimator minimizing the *levels* of observed and predicted wealth (not normalized by permanent income),

$$\mathcal{L}(\Theta)=\frac{1}{N\bar{T}}\sum_{i=1}^{N}\sum_{t=1}^{T_i}(A_t^*(M_{it},P_{it},\mathbf{z}_{it}|\Theta)-A_{it})^2 \tag{7}$$

which is consistent even if the measurement error is serially correlated. The results, reported in columns (2) of Table 7, are very similar to the baseline results.

[35] One approach to adjusting the covariance, and thus inference, of the estimated parameters for this first stage estimation, could be to formulate a joint likelihood function and estimate all parameters simultaneously using the first- and second-stage estimates as starting values. However, because the permanent income contain $\sum_i^N T_i$ parameters, this approach is infeasible in my setting. Other approaches could be to apply a bootstrap procedure or a Taylor expansion approach applied in, for example, Gourinchas and Parker (2002).

TABLE 7

Estimation results from alternative estimators

	Low skilled			*High skilled*		
	(1)	*(2)*	*(3)*	*(1)*	*(2)*	*(3)*
ρ	0.587	2.505	2.640	1.245	2.873	3.023
	(0.001)	(0.046)	(0.040)	(0.004)	(0.065)	(0.064)
β	0.974	0.963	0.961	0.982	0.974	0.970
	(0.000)	(0.001)	(0.001)	(0.000)	(0.001)	(0.001)
γ	1.010	1.562	1.610	0.921	1.275	1.410
	(0.000)	(0.024)	(0.024)	(0.005)	(0.025)	(0.043)
σ_ξ			0.468			0.491
			(0.000)			(0.001)
Taste shifter, $v(\mathbf{z};\theta)=1+\theta'\mathbf{z}$						
θ_{11}	0.001	0.006	−0.003	0.010	0.021	−0.008
	(0.001)	(0.006)	(0.007)	(0.014)	(0.010)	(0.012)
θ_{12}	−0.002	−0.021	−0.040	0.015	0.024	−0.008
	(0.001)	(0.004)	(0.005)	(0.008)	(0.007)	(0.012)
θ_{21}	−0.003	−0.025	−0.031	0.004	0.014	−0.007
	(0.001)	(0.005)	(0.006)	(0.007)	(0.008)	(0.009)
θ_{22}	0.003	−0.003	0.005	0.009	0.015	0.013
	(0.001)	(0.004)	(0.005)	(0.005)	(0.007)	(0.009)
θ_{31}	0.005	−0.000	0.005	0.015	0.033	0.041
	(0.002)	(0.008)	(0.010)	(0.008)	(0.012)	(0.015)
θ_{32}	0.006	0.020	0.031	0.004	0.024	0.037
	(0.001)	(0.011)	(0.014)	(0.014)	(0.016)	(0.020)

Notes: Standard errors in brackets. Standard errors in columns (3) are based on the inverse of the hessian. $^{(1)}$Based on minimizing the absolute distance between observed and predicted normalized wealth (see equation (6)). $^{(2)}$Based on minimizing the squared distance between observed and predicted wealth (see equation (7)). $^{(3)}$Based on the log-likelihood function in equation (5) but using an alternative measure of permanent income, $\bar{P}_t = \frac{1}{T_i}\sum_t^{T_i} P_{it}$, where P_{it} is the baseline estimated permanent income based on the Kalman filter.

Secondly, I estimate the model parameters using slightly different estimates of permanent income. Particularly, I use the household-specific average, $\bar{P}_i = T_i^{-1}\sum_{t=1}^{T_i} P_{it}$, as an estimate of permanent income. The hope is that potential estimation bias will tend to net-out such that the average is less affected by bias. Columns (3) of Table 7 show, again, that the results are very similar to the baseline estimates albeit a slightly higher risk aversion coefficient and discount factor are estimated.

I interpret this set of robustness results as suggesting that the first stage estimation of permanent income is not driving the results. Decomposing savings motives, as done above for the baseline parameters, using the alternative parameters in Table 7, lead to the qualitatively unchanged result that children does not seem to be a major savings motive. See Figures A5–A7 in the online supplemental material for the figures associated with this exercise.

TABLE 8

Estimation results from alternative calibrations

	Low skilled			*High skilled*		
	(1)	*(2)*	*(3)*	*(1)*	*(2)*	*(3)*
ρ	1.024	2.298	2.170	2.012	2.544	2.415
	(0.011)	(0.041)	(0.064)	(0.057)	(0.057)	(0.063)
β	0.975	0.965	0.964	0.979	0.973	0.972
	(0.000)	(0.000)	(0.001)	(0.000)	(0.001)	(0.001)
γ	1.000	0.889	1.443	1.000	0.720	1.249
	—	(0.007)	(0.026)	—	(0.010)	(0.023)
σ_ξ	0.468	0.468	0.468	0.490	0.490	0.490
	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)
Taste shifter, $v(\mathbf{z};\theta)=1+\theta'\mathbf{z}$						
θ_{11}	0.029	−0.005	0.010	0.013	−0.008	0.004
	(0.003)	(0.006)	(0.006)	(0.008)	(0.010)	(0.009)
θ_{12}	−0.007	−0.032	−0.029	−0.001	0.000	−0.001
	(0.002)	(0.004)	(0.004)	(0.006)	(0.007)	(0.007)
θ_{21}	0.006	−0.032	−0.021	0.004	−0.013	−0.005
	(0.002)	(0.005)	(0.005)	(0.006)	(0.008)	(0.007)
θ_{22}	0.009	−0.006	−0.003	0.009	0.001	0.002
	(0.002)	(0.004)	(0.004)	(0.006)	(0.007)	(0.007)
θ_{31}	0.009	−0.003	0.010	0.026	0.023	0.037
	(0.004)	(0.009)	(0.009)	(0.010)	(0.012)	(0.012)
θ_{32}	0.010	0.024	0.026	0.018	0.023	0.028
	(0.005)	(0.012)	(0.012)	(0.013)	(0.017)	(0.016)
$-\mathcal{L}(\Theta)$	0.4660	0.4654	0.4654	0.4988	0.4987	0.4987

Notes: Standard errors in brackets are based on the inverse of the hessian. $^{(1)}$Based on a version of the model in which the postretirement motives have been fixed at $\gamma = 1$. $^{(2)}$Based on a version of the model in which the replacement rate in retirement are lower, at $\xi = 0.7$. $^{(3)}$Based on a version of the model in which each children present in the household reduces income by 1%.

Postretirement motives

As argued throughout, the parameter γ is a parsimonious way of accounting for *all* postretirement motives without explicitly having to model the postretirement decisions. The approach is similar to that of, for example, Gourinchas and Parker (2002) but their parametrization is slightly different. To ensure that the results are not driven by this simplification, I report in columns (1) of Table 8 estimation results from a version of the model in which I fix $\gamma = 1$ while estimating the remaining parameters. The effect of children on the marginal utility of consumption is almost unaffected.

The fact that I estimate $\gamma > 1$ could stem from a too high calibrated replacement rate in retirement, $\varkappa$, of 90%. If the replacement rate of the households used to estimate the model is actually lower, this would produce too little wealth accumulation in the model (because income drops with too little in retirement) and to fit the wealth level of households close to retirement, the value of post-retirement motives are estimated to be higher. Columns (2) of Table 8 report estimation results when assuming a lower replacement rate of 70% ($\varkappa = 0.7$). As expected, the estimated postretirement motives are now estimated significantly lower

– below one – while the effect of children on the marginal utility of consumption remains unaffected.

Labour market effects of children

The labour/leisure choice has been ignored throughout. This has been done for computational tractability and because all existing studies using Euler equation methods ignores the effect of fertility on labour market supply and household income. Here, I investigate in a *reduced form way* how important the effects of children on household income is in generating the results. Figure A4 in the supplemental material illustrates log household income around the arrival of the first child. While there is a significant drop in income around and after the arrival of the first child, the effect is economically small. Table A2 in the supplemental material suggests that income is around 1% lower in households who have one child compared to households without children and I estimate that the second and third child reduces this gap slightly.

Columns (3) in Table 8 report estimation results when assuming that each child reduces income with 1% while living at home. This amounts to an income process where $Y_t = P_t\eta_t[1 - 0.01 \cdot \sum_{j=1}^{3} \mathbf{1}_{\{\mathbf{z}_j \neq NC\}}]$. Although a 1% reduction in income per child aged 0 through 21 is a substantial cost of children, the estimated effects of children on the marginal utility of consumption are hardly affected.

A valid concern could be that labour market participation decisions are related to the presence of children. If consumption and leisure are substitutes, and consumers tend to increase leisure time when having children, ignoring labour market supply will result in a downwards bias in the estimated effect of children on the marginal utility of consumption. To investigate if reduced labour market supply around and after childbirth is driving the results, a sample of households in which both partners work at least 30 weeks a year in all observed years are used for estimation. Estimates based on this sample, which I refer to as 'Workers' are reported in columns (1) in Table 9.

As expected, the estimated effect of children on the marginal utility of consumption increases when estimating on the working households only. The estimates are, however, economically small and statistically insignificant. While the discount factor is unaffected, the relative risk aversion coefficient, ρ, has increased with around 9% for low skilled and 16% for high skilled. This might reflect preference heterogeneity or stem from a higher estimated value of retirement, γ.[36]

House purchases

Home purchase is another decision that is likely to be closely related to the presence of children. Figure A3 in the supplemental material illustrates that the share of home owners increases significantly in the beginning of the life cycle, around the same time as children tend to arrive. Further, there is a significantly lower fraction of home owners among childless households. If home purchases tend to be associated with children, observed saving around and after childbirth will likely not decrease in the same extend as the model

[36] The transitory and permanent income variances, and the permanent income growth rate along with the likelihood of children arriving are re-calibrated for all sub-samples.

TABLE 9

Estimation results from restricted samples

		Low skilled			*High skilled*		
		Working	*Owning*	*Both*	*Working*	*Owning*	*Both*
ρ	Risk aversion	2.603	3.101	3.277	3.075	3.021	3.445
		(0.065)	(0.053)	(0.091)	(0.098)	(0.073)	(0.119)
β	Discount factor	0.963	0.963	0.961	0.977	0.975	0.981
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
γ	Retirement	1.619	1.628	1.794	1.291	1.174	1.153
		(0.039)	(0.030)	(0.064)	(0.042)	(0.025)	(0.041)
σ_ξ	Meas. error	0.474	0.456	0.463	0.482	0.472	0.465
		(0.001)	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)
Taste shifter, $v(\mathbf{z};\theta)=1+\theta'\mathbf{z}$							
θ_{11}	1. child $\leqslant$ 10	0.045	0.039	0.095	0.008	0.052	0.069
		(0.012)	(0.010)	(0.018)	(0.018)	(0.014)	(0.024)
θ_{12}	1. child > 10	0.009	−0.012	0.036	0.040	0.042	0.082
		(0.008)	(0.006)	(0.011)	(0.013)	(0.010)	(0.016)
θ_{21}	2. child $\leqslant$ 10	−0.025	−0.015	−0.014	0.023	0.006	0.061
		(0.010)	(0.007)	(0.014)	(0.015)	(0.010)	(0.019)
θ_{22}	2. child > 10	−0.010	0.016	0.011	0.020	0.027	0.060
		(0.009)	(0.006)	(0.012)	(0.014)	(0.009)	(0.017)
θ_{31}	3. child $\leqslant$ 10	−0.021	0.005	−0.033	0.046	0.044	0.050
		(0.021)	(0.013)	(0.028)	(0.026)	(0.016)	(0.033)
θ_{32}	3. child > 10	0.007	0.025	0.005	0.028	0.018	0.037
		(0.027)	(0.018)	(0.036)	(0.033)	(0.021)	(0.041)
$-\mathcal{L}(\Theta)$		0.43699	0.47000	0.44571	0.45994	0.50558	0.47065
$\max_j \lvert\partial\mathcal{L}(\Theta)/\partial\Theta_j\rvert$		$6.1e-6$	$8.2e-6$	$6.8e-6$	$1.2e-5$	$1.9e-5$	$2.1e-5$
Number of observations		320,398	713,168	269,200	184,736	352,197	157,175

Notes: 'Working' refers to a sample in which both spouses in a household work at least 30 weeks a year in all years observed in the data. 'Owning' refers to a sample in which all households are home owners through the observed period. 'Both' refers to a sample in which households are home owners and working at least 30 weeks throughout the observed period. Standard errors in brackets are based on the inverse of the hessian.

(ignoring housing purchases) would suggest.[37] This will, in turn, result in a downwards bias in the estimated effect of children on the marginal utility of consumption. Columns (2) in Table 9 report estimation results using a sample of households who are homeowners in all observed years, which I refer to as 'Owners'. Columns (3), denoted 'Both', report estimates from a sub-sample of working home-owners.

Estimation results from only using the sub-sample of working home owners suggest that the first child increases the marginal utility by around 7–10% for low- and high-skilled households in this group. As expected, for this highly selected group, the effect of children is larger. There is still no sign of an age effect while I do find signs of economies of scale for this sub-sample.

[37] If a household goes from renting to owning, this would tend to increase net-worth because rent-costs are not included in the wealth measure. In turn, this will reduce the imputed measure of consumption around home purchases.

TABLE 10

Estimates from alternative deterministic fertility process

	Low skilled		*High skilled*	
	(1)	*(2)*	*(1)*	*(2)*
ρ	2.326 (0.052)	2.342 (0.045)	2.650 (0.061)	2.584 (0.060)
β	0.966 (0.001)	0.966 (0.000)	0.974 (0.001)	0.976 (0.001)
γ	1.428 (0.021)	1.421 (0.019)	1.220 (0.028)	1.187 (0.022)
σ_ξ	0.468 (0.000)	0.468 (0.000)	0.490 (0.001)	0.490 (0.001)
$v(\mathbf{z}_t;\theta)=\exp(\theta\mathbf{z}_t)$				
θ	0.017 (0.003)		0.021 (0.005)	
$v(\mathbf{z}_t;\theta)=1+\theta'\mathbf{z}_t$				
θ_{11}		0.070 (0.007)		0.081 (0.012)
θ_{12}		0.028 (0.004)		0.063 (0.007)
θ_{21}		−0.043(0.006)		0.003 (0.009)
θ_{22}		−0.048(0.005)		−0.014(0.008)
θ_{31}		0.046 (0.011)		0.046 (0.015)
θ_{32}		0.038 (0.013)		0.014 (0.017)
$-\mathcal{L}(\Theta)$	0.465340	0.465269	0.498612	0.498541
$\max_j \lvert\partial\mathcal{L}(\Theta)/\partial\Theta_j\rvert$	$2.0e-6$	$3.5e-5$	$7.3e-6$	$1.2e-6$

Notes: Reported are estimates from a model with perfect foresight over arrival of children. Other calibrated parameters are unchanged. Standard errors in brackets are based on the inverse of the hessian.

Perfect Fertility Foresight

Throughout, the fertility process has been completely stochastic. As an alternative extreme, I here follow (among others) Gourinchas and Parker (2002); Cagetti (2003) and Browning and Ejrnæs (2009) and assume that children are deterministic in the sense that they are perfectly anticipated from the beginning of adulthood. Table 10 reports estimation results from an alternative fertility process where households in the model know with perfect foresight *when and how many* children they will have. I use in the data realized fertility as completed fertility. While children still arrive exogenously, this assumption leaves households with the advantage of knowing when children arrive and can, thus, adjust their optimal wealth path accordingly.

While the estimated effect of children on the marginal utility of consumption is still small, it is somewhat higher than in the baseline model. This is expected because this alternative fertility process gives households the ability to adjust their saving and consumption behaviour according to their *individual known fertility*. In contrast, the baseline stochastic arrival of children leads all households within an age and household size group to have the same fertility expectations and, thus, saving motives. Therefore, in the baseline model, even households who, for example, turn out to be childless would have accumulated wealth to be able to smooth out marginal utility in the case children arrive.

The 'true' fertility process is likely a combination of the two extreme fertility processes: There is some element of fertility planning but also stochastic elements from, for example, imperfect contraceptive control, infertility and declining fecundity (See e.g. Ejrnæs and Jørgensen, 2006).

VII. Concluding discussion

I have estimated the effect of children on the marginal utility of consumption while allowing income uncertainty, credit constraints and postretirement motives to also influence household behaviour. To the best of my knowledge, no existing study has allowed for as rich a specification of the effect of children on the marginal utility of consumption in a dynamic model of intertemporal consumption behaviour. The model is estimated by ML using high quality Danish administrative register data, providing detailed longitudinal information on income, assets, liabilities and household characteristics.

I find that the effect of children on the marginal utility of non-durable consumption is economically negligible. This is in stark contrast to most existing studies which, by applying Euler equation estimation techniques, find that children increase the marginal utility of consumption by around 40–150%. In a recent study, Browning and Ejrnæs (2009) illustrate that allowing the age of children to affect the marginal utility of consumption can completely explain the observed consumption age profile. As they note, however, if precautionary motives are important (as found by e.g., Gourinchas and Parker, 2002 and Cagetti, 2003), their results might proxy for precautionary motives. My results suggest that children play a minor role when allowing for precautionary motives *simultaneously*.

While I estimate that children have a surprisingly small effect on the marginal utility of non-durable consumption, data simulated from the estimated model reproduce correlations between log consumption growth and changing household composition as found in the Danish data and typically found in UK and US data. I interpret this result as indicating that children might proxy for precautionary motives when ignoring these. In support of this, I find that the estimated effect of children on the marginal utility of consumption increases significantly when estimating a version of the model with much less severe precautionary motives. The results are in line with the failure of Euler equation estimation techniques to uncover the effect of children on the marginal utility of consumption if risk averse households are unable or unwilling to borrow, studied in Jørgensen (2016).

Several interesting avenues for future research remains and, admittedly, the present analysis retains many of the same simplifying assumptions as the existing literature suffers from. For example, I ignore labour supply and home ownership choices although both likely are related to fertility. Including fertility, labour supply, and home ownership choices in the current framework is extremely interesting, but also a computationally demanding task, which I leave for future research. Furthermore, the small estimated effects of children on non-durable consumption likely camouflage significant shifts in the combination of consumption sub-components within a household. Specifically, the arrival of children may shift expenditures from luxury goods towards necessities while leaving total non-durable consumption almost unaffected.

Final Manuscript Received: November 2016

References

Ackerberg, D., Geweke, J. and Hahn, J. (2009). 'Comments on "Convergence Properties of the Likelihood of Computed Dynamic Models"', *Econometrica*, Vol. 77, pp. 2009–2017.

Adda, J., Dustmann, C. and Stevens, K. (forthcoming). 'The career costs of children', *Journal of Political Economy*.

Alan, S. and Browning, M. (2010). 'Estimating intertemporal allocation parameters using synthetic residual estimation', *The Review of Economic Studies*, Vol. 77, pp. 1231–1261.

Alan, S., Attanasio, O. and Browning, M. (2009). 'Estimating Euler equations with noisy data: two exact GMM estimators', *Journal of Applied Econometrics*, Vol. 24, pp. 309–324.

Alan, S., Atalay, K. and Crossley, T. F. (2012). Euler Equation Estimation on Micro Data, Working Paper 1221, Koc University-TUSIAD Economic Research Forum.

Attanasio, O. P. and Browning, M. (1995). 'Consumption over the life cycle and over the business cycle', *The American Economic Review*, Vol. 85, pp. 1118–1137.

Attanasio, O. P. and Weber, G. (1995). 'Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey', *The Journal of Political Economy*, Vol. 103, pp. 1121–1157.

Attanasio, O. P., Banks, J., Meghir, C. and Weber, G. (1999). 'Humps and bumps in lifetime consumption', *Journal of Business & Economic Statistics*, Vol. 17, pp. 22–35.

Banks, J., Blundell, R. and Preston, I. (1994). 'Life-cycle expenditure allocations and the consumption costs of children', *European Economic Review*, Vol. 38, pp. 1391–1410.

Bick, A. and Choi, S. (2013). 'Revisiting the effect of household size on consumption over the life-cycle', *Journal of Economic Dynamics and Control*, Vol. 37, pp. 2998–3011.

Blundell, R., Browning, M. and Meghir, C. (1994). 'Consumer demand and the life-cycle allocation of household expenditures', *The Review of Economic Studies*, Vol. 61, pp. 57–80.

Blundell, R., Pistaferri, L. and Preston, I. (2008). 'Consumption inequality and partial insurance', *American Economic Review*, Vol. 98, pp. 1887–1921.

Blundell, R., Low, H. and Preston, I. (2013). 'Decomposing changes in income risk using consumption data', *Quantitative Economics*, Vol. 4, pp. 1–37.

Blundell, R., Dias, M. C., Meghir, C. and Shaw, J. M. (2016). 'Female labour supply, human capital and welfare reform', *Econometrica*, Vol. 84, pp. 1705–1753.

Browning, M. and Ejrnæs, M. (2009). 'Consumption and children', *The Review of Economics and Statistics*, Vol. 91, pp. 93–111.

Browning, M. and Leth-Petersen, S. (2003). 'Imputing consumption from income and wealth information', *The Economic Journal*, Vol. 113, pp. 282–301.

Browning, M., Deaton, A. and Irish, M. (1985). 'A profitable approach to labor supply and commodity demands over the life-cycle', *Econometrica*, Vol. 53, pp. 503–544.

Cagetti, M. (2003). 'Wealth accumulation over the life cycle and precautionary savings', *Journal of Business and Economic Statistics*, Vol. 21, pp. 339–353.

Carroll, C. D. (1992). 'The Buffer-Stock Theory of Saving: some macroeconomic evidence', *Brookings Papers on Economic Activity*, Vol. 2, pp. 61–135.

Carroll, C. D. (1997). 'Buffer-stock saving and the life cycle/permanent income hypothesis', *The Quarterly Journal of Economics*, Vol. 112, pp. 1–55.

Carroll, C. D. (2001). 'Death to the log-linearized consumption Euler equation! (and very poor health to the second-order approximation)', *Advances in Macroeconomics*, Article 6, Vol. 1, pp. 1534–6013.

Carroll, C. D. (2006). 'The method of endogenous gridpoints for solving dynamic stochastic optimization problems', *Economics Letters*, Vol. 91, pp. 312–320.

Citro, C. F. and Michael, R. T. (1995). *Measuring Poverty: A New Approach*, National Academy Press, Washington, D.C.

Deaton, A. (1991). 'Saving and liquidity constraints', *Econometrica*, Vol. 59, pp. 1221–1248.

Ejrnæs, M. and Jørgensen, T. H. (2006). Saving Behavior around Intended and Unintended Childbirths, Working paper, University of Copenhagen.

Fernández-Villaverde, J. and Krueger, D. (2007). 'Consumption over the life cycle: facts from consumer expenditure survey data', *The Review of Economics and Statistics*, Vol. 89, pp. 552–565.

Fernández-Villaverde, J., Rubio-Ramírez, J. F. and Santos, M. S. (2006). 'Convergence properties of the likelihood of computed dynamic models', *Econometrica*, Vol. 74, pp. 93–119.

Gourinchas, P.-O. and Parker, J. A. (2002). 'Consumption over the life cycle', *Econometrica*, Vol. 70, pp. 47–89.

Heckman, J. (1974). 'Life cycle consumption and labor supply: an explanation of the relationship between income and consumption over the life cycle', *The American Economic Review*, Vol. 64, pp. 188–194.

Hong, J. H. and Ríos-Rull, J.-V. (2012). 'Life insurance and household consumption', *American Economic Review*, Vol. 102, pp. 3701–30.

Irvine, I. (1978). 'Pitfalls in the estimation of optimal lifetime consumption patterns', *Oxford Economic Papers*, Vol. 30, pp. 301–309.

Jørgensen, T. H. (2013). 'Structural estimation of continuous choice models: evaluating the EGM and MPEC', *Economics Letters*, Vol. 119, pp. 287–290.

Jørgensen, T. H. (2016). 'Euler equation estimation: children and credit constraints', *Quantitative Economics*, Vol. 7, pp. 935–968.

Kalwij, A. (2005). 'Consumption and income around the time of births', *Review of Economics of the Household*, Vol. 3, pp. 75–89.

Kaplan, G. and Violante, G. L. (2010). 'How much consumption insurance beyond self-insurance?', *American Economic Journal: Macroeconomics*, Vol. 2, pp. 53–87.

Love, D. A. (2010). 'The effects of marital status and children on savings and portfolio choice', *Review of Financial Studies*, Vol. 23, pp. 385–432.

Ludvigson, S. and Paxson, C. H. (2001). 'Approximation bias in linearized Euler equations', *Review of Economics and Statistics*, Vol. 83, pp. 242–256.

Meghir, C. and Pistaferri, L. (2004). 'Income variance dynamics and heterogeneity', *Econometrica*, Vol. 72, pp. 1–32.

Palumbo, M. G. (1999). 'Uncertain medical expenses and precautionary saving near the end of the life cycle', *The Review of Economic Studies*, Vol. 66, pp. 395–421.

Rust, J. (1987). 'Optimal replacement of GMC bus engines: an empirical model of Harold Zurcher', *Econometrica*, Vol. 55, pp. 999–1033.

Scholz, J. K. and Seshadri, A. (2009). Children and Household Wealth, Working paper, University of Wisconsin-Madison.

Scholz, J. K., Seshadri, A. and Khitatrakun, S. (2006). 'Are Americans saving 'optimally' for retirement?', *Journal of Political Economy*, Vol. 114, pp. 607–643.

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.

The Danish Ministry of Finance (2003). *ældres sociale vilkår*. (In Danish).

Thurow, L. C. (1969). 'The optimum lifetime distribution of consumption expenditures', *The American Economic Review*, Vol. 59, pp. 324–330.

Supporting Information

Additional supporting information may be found in the online version of this article:

Appendix A. Additional Figures and Tables
Appendix B. Construction of Estimation Sample
Appendix C. Additional Log-Linearized Euler Equation Estimates
Appendix D. Solving the Model
Appendix E. Identification
Appendix F. Permanent Income: The Kalman Filter