Discussion Papers
Department of Economics
University of Copenhagen

No. 15-08

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

Thomas H. Jørgensen

Øster Farimagsgade 5, Building 26, DK-1353 Copenhagen K., Denmark
Tel.: +45 35 32 30 01 – Fax: +45 35 32 30 00
http://www.econ.ku.dk

ISSN: 1601-2461 (E)

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

Thomas H. Jørgensen†

May 17, 2015

Abstract

I study how children affect the marginal utility of non-durable consumption. I estimate by Maximum Likelihood a structural economic model of optimal intertemporal allocation of consumption in the presence of children using high quality Danish administrative longitudinal data. Contrary to existing studies, I allow income uncertainty, credit constraints, and post-retirement motives to affect household behavior while the number and age of all children can affect the marginal utility of consumption. I estimate that children have a negligible effect on the marginal utility of non-durable consumption. To reconcile these results with existing studies, typically estimating an important role for children while ignoring precautionary motives, I illustrate how ignoring precautionary motives increases the estimated importance of children. I interpret the results as indicating that precautionary motives might play a larger role than children in explaining the observed consumption age profile (JEL: D12, D14, D91).

Keywords: Consumption, Children, Precautionary saving, Life cycle, Structural Estimation.

†Department of Economics, University of Copenhagen, Øster Farimagsgade 5, building 26, DK-1353 Copenhagen, and Centre for Applied Microeconometrics (CAM). Webpage: www.tjeconomics.com. Contact: thomas.h.jorgensen@econ.ku.dk, Tlf.: 0045 35 32 39 94.

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 Introduction

This study is concerned with the effect of children on 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 consumption behavior has, therefore, received great attention over the last two decades with an important role for children as the most common finding.[1] 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 a significant bulk of studies analyze the effect of demographics on consumption, allowing for multiple consumption-savings motives simultaneously is rare.

The present study offers new insights to this literature. I estimate the effect of children on the marginal utility of consumption while allowing income uncertainty, credit constraints and post-retirement motives to also affect household behavior. A key difference from existing studies is that precautionary motives typically are excluded when the effect of children on consumption is analyzed.[3] Precautionary savings motives have, however, been illustrated to be important for wealth accumulation of particularly young households.[4] I estimate a life cycle model of intertemporal consumption and saving in the presence of children using high quality Danish administrative register data. Unlike must surveys typically used[5], the Danish data provide detailed longitudinal information on household characteristics, income, and, importantly, most assets and liabilities. Household net worth is a crucial determinant of optimal intertemporal allocation of consumption and savings, which I use to identify the parameters of interest through Maximum Likelihood estimation. Log consumption growth regressions suggest that the Danish data is comparable to data used in existing studies.

I find an economically negligible effect of children on the marginal utility of con-

[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, Banks, Meghir and Weber (1999); Fernández-Villaverde and Krueger (2007) and Browning and Ejrnæs (2009).

[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] Browning and Ejrnæs (2009) is a recent example. As argued in Jørgensen (2014), Euler equation estimation techniques are biased if risk averse households face sufficient precautionary motives such as credit constraints or a probability of zero income. Since almost all existing evidence on the effect of children on consumption is based on (log-linearized) Euler equation estimation, existing results rule out the alternative consumption/savings motive from income uncertainty and credit constraints.

[4] See, e.g., Gourinchas and Parker (2002) and Cagetti (2003).

[5] Commonly used surveys include the Consumer Expenditure Survey (CEX) and the Panel Study of Income dynamics (PSID).

sumption. In contrast, while ignoring precautionary motives, most existing studies find that children increase the marginal utility of consumption with 40-150 percent (see discussion below). To reconcile my results with existing studies I illustrate how ignoring precautionary motives increases the estimated importance of children in the Danish data. I interpret the results as indicating that precautionary motives might play a larger role than children in explaining the observed consumption age profile.[6]

In contrast to my findings it seems broadly accepted that children play an important role in generating the observed consumption profiles (Attanasio and Weber, 2010). Attanasio, Banks, Meghir and Weber (1999) estimate that children increase marginal utility with around 40 percent while the preferred specification in Alan, Attanasio and Browning (2009) implies that children increase marginal utility with around as much as 150 percent. Using the repeated cross section information on non-durable consumption in the Consumer Expenditure Survey (CEX), Attanasio, Banks, Meghir and Weber (1999) construct synthetic cohort panels (Browning, Deaton and Irish, 1985) and estimate a *log-linearized Euler equation* (i.e. log-consumption growth) with changes in the number of adults and number of children included as explanatory variables. The estimation of demographic effects on consumption by log-linearized Euler equations likely also removes variation in consumption stemming from precautionary behavior (Jørgensen, 2014). This is because changing household demographics coincide with high income growth and the inability or unwillingness of risk averse households to borrow against future income.

More recently, Browning and Ejrnæs (2009) find that the number and age of children can completely explain the hump in consumption over the life cycle. By estimating a log-linearized Euler equation on a synthetic cohort panel constructed from the Family Expenditure Survey (FES) they find that the marginal utility of consumption increases in the age of children. As they recognize, and as I have argued above, their results might proxy for precautionary motives. A novelty of the present study is that I, unlike Browning and Ejrnæs (2009), allow income uncertainty, credit constraints, and post-retirement motives to affect household behavior while the number and age of all children can affect the marginal utility of consumption. I find no evidence of an age effect for Danish households.

The results are also related to a growing strand of literature estimating models of intertemporal behavior 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] To be able to analyze whether the

[6]Gourinchas and Parker (2002) and Cagetti (2003) also find a significant role for precautionary motives for especially young households even *after* removing demographic variation.

[7]For example, Love (2010) assumes that children arrive with two years intervals and Sommer (2014)

results in Browning and Ejrnæs (2009) pertains when including income uncertainty, credit constraints and retirement it is, however, necessary to estimate a model in which the age of all children are potentially important.

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, including Browning and Ejrnæs (2009)[8], fertility is assumed to be exogenous, and, furthermore, labor supply and house purchases is not part of the economic model. I do, however, show that the effect of children on marginal utility only increases slightly when estimating the model on a sub-sample of permanently employed home owners.

The rest of the paper proceeds as follows. In the next section, I augment the standard buffer stock model of intertemporal consumption with the potential presence of children. Section 3 describes the Danish administrative registers and shows that commonly applied estimators yield similar results from the Danish data as existing studies, using UK or US data, report. Section 4 discusses how some of the parameters of the model are calibrated and section 5 discusses how the remaining parameters are identified. Section 6 presents the estimation results and model fit and investigates the robustness of the results. Finally, I conclude.

2 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, e.g., Attanasio, Banks, Meghir and Weber (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 in which they consume all available resources. In all preceding periods, households chose the level of consumption, C_t, that solves the optimization problem

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

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

assumes that there is 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 (2013) and Adda, Dustmann and Stevens (2015) only the age of the youngest child matters.

[8] Recent examples of exogenous (probabilistic) arrival of children include Love (2010); Hong and Ríos-Rull (2012); and Blundell, Dias, Meghir and Shaw (2013).

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, Banks, Meghir and Weber (1999) and let children affect the *marginal value* of consumption.[9]

Following Gourinchas and Parker (2002), bequest motives and survival and income uncertainty are omitted post retirement and γ in equation (1) is a parsimonious way of adjusting for all post retirement motives. I refer to γ as a "retirement motive" although it summarizes *all* potential ignored post-retirement motives.

Households solve (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, $A_t = M_t - C_t \geq 0, \forall t \geq T_r$, while working households can borrow up to a fraction of their permanent income $A_t \geq -\kappa P_t \, \forall t, \, \kappa \geq 0$.

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, $\log \eta_t \sim \mathcal{N}(-\sigma_\eta^2/2, \sigma_\eta^2)$ is a mean one permanent income shock, and ε_t is a mean one transitory income shock taking the value μ with probability $\wp$ and otherwise log normal,

$$\begin{aligned} \varepsilon_t &= \begin{cases} \mu & \text{with probability } \wp \\ (\tilde{\varepsilon}_t - \mu\wp)/(1-\wp) & \text{with probability } 1-\wp \end{cases} \\ \log \tilde{\varepsilon}_t &\sim \mathcal{N}(-\sigma_\varepsilon^2/2, \sigma_\varepsilon^2). \end{aligned}$$

When retired, the income process is a deterministic constant fraction $\varkappa \leq 1$ of permanent income at retirement, $Y_t = \varkappa P_{T_r}, \, \forall t \geq T_r$.

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.

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

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 behavior. To circumvent the computational cost of keeping track of the age of all children, strict assumptions on the timing of children are typically imposed.[10] 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.

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.

3 The Danish Register Data and Descriptive Analysis

I use high quality Danish administrative registers covering the entire population from the period 1987–1996. 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.

Household consumption is not observed in the registers and is, therefore, imputed 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 thus 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. The resulting consumption measure will, however, include some durables such as home appliances.

Disposable income includes all labor market and non-labor 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.

[10] 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 two years interval, and Sommer (2014) assumes that there is two types of children: Children living at home and children, who have left the household. Alternatively, Blundell, Dias, Meghir and Shaw (2013) assumes that only the youngest child matters.

[11] Love (2010); Hong and Ríos-Rull (2012) and Blundell, Dias, Meghir and Shaw (2013) also assume that children arrive probabilistically. The same is true in Adda, Dustmann and Stevens (2015) although in their model households essentially has some control over contraceptive choices.

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 percent 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 pre-retirement behavior, 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 to 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 four years and exclude households with more than three children.[12] Only households with children born when the wife was aged 15 through 43 are included in the analysis. Households in which one adult is self-employed or out of the labor 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. Table A1 in Appendix A in the supplemental material show how the selection criteria affect the number of observations available for estimation. 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.

Table 1 presents means and standard deviations for selected measures. The wives are on average 40.7 years old with on average one 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 35,000 USD only around a thousand dollars short of the average household income of 36,000 USD. Around 34 percent 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 in anticipation of the arrival of children in the future and almost deplete that wealth when they subsequently have children. When

[12]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	Std
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.

there is hardly any difference between wealth accumulation across the two groups this suggests that children might not be the primary explanation for the observed consumption age profile. I show supporting descriptive evidence of this assertion below.

3.1 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, using UK or US data. Particularly, the goal is not to estimate a statistically well-specified equation but rather to follow the approach in Attanasio, Banks, Meghir and Weber (1999) as closely as possible. Their estimates (and standard deviations) are reported in column five of Table 2 for easy comparison.

Under the rather strong assumptions of no income uncertainty and perfect markets (no constraints on borrowing), the coefficient in front of the growth in the number of children is $\rho^{-1}\theta$ (Carroll, 2001; Jørgensen, 2014). In turn, the estimates below can be interpreted as (a ratio of) structural parameters of a model in which households *do not* face income uncertainty, credit constraints and retirement 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)*

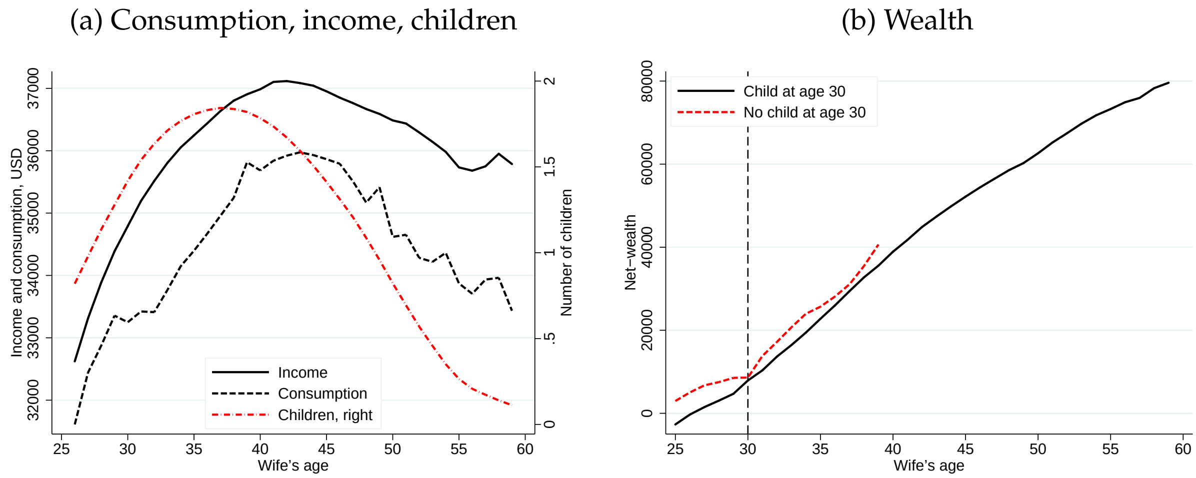

Figure 1 – Age Profiles of Income, Consumption, Children and Wealth.

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. The PSID numbers are in 1980 US dollars and the 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.

second to fourth lagged changes in log consumption and log income, and *iii)* a polynomial in age of the wife.[13] 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, Banks, Meghir and Weber, 1999) 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. To compare the descriptive regression results from the Danish data with these type of studies, I have collapsed the data into cohort panels (one year bands) and present the same estimation results using this data.

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, Banks, Meghir and Weber (1999) using the CEX. Table 2 shows that log consumption grows with around .49 when a child (younger than 18) arrives and around .56 when an additional adult (18 or older) is present compared to .21 and .45, respectively, reported in Attanasio, Banks, Meghir and Weber (1999). This suggests that the Danish imputed consumption measure is not

[13] 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.

Table 2 – Log Consumption Growth Regressions.

	Household panel		Cohort panel		
	OLS	IV	OLS	IV	ABMW†
$\Delta\#kids$	0.042	1.782	0.182	0.490	0.212
	(0.005)	(0.090)	(0.112)	(0.230)	(0.101)
$\Delta\#adults$	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 intstruments are on the cohort level.

† ABMW refers to the results reported in Table 1 of Attanasio, Banks, Meghir and Weber (1999) using the CEX to construct synthetic cohort panels (using five 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 $\theta = 0.212 \cdot 1.57 \approx .33$.

fundamentally different than the data typically used in existing studies.

The OLS estimate on the true household panel data is significantly smaller than the others. As argued in Jørgensen (2014), 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.[14] On the other hand, the synthetic cohort panel approach is potentially upwards biased since the average number of children is highly correlated with income growth and, thus, might proxy for precautionary motives. This motivates why I allow for all these motives simultaneously when estimating the economic model in Section 6.

The Danish data provide longitudinal information on the household level and I can, thus, investigate directly how consumption responds when a child arrives. Table 3 reports log consumption around the time of birth of the first child split by high and low skilled households. All regressions include age and year dummies and are relative

[14]Imagine the extreme case in which a household is constrained in all periods, and thus consumes all income. Even if such a household would have liked to increase consumption when a child arrive 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).

Table 3 – Log Consumption Around the Time of First Birth.

	Low skilled		High skilled	
	OLS	FE	OLS	FE
Year from birth: -3	0.001	-0.002	0.015	-0.046
	(0.019)	(0.023)	(0.023)	(0.042)
Year from birth: -2	0.013	0.012	0.004	-0.029
	(0.015)	(0.017)	(0.018)	(0.025)
Year of birth	0.029**	0.025	0.017	0.048*
	(0.011)	(0.013)	(0.013)	(0.022)
Year from birth: 1	0.037**	0.034*	0.023	0.089*
	(0.011)	(0.016)	(0.013)	(0.037)
Year from birth: 2	0.039**	0.044*	0.027	0.135*
	(0.012)	(0.021)	(0.014)	(0.054)
Year from birth: 3	0.018	0.024	0.008	0.152*
	(0.013)	(0.026)	(0.015)	(0.071)
Constant	Yes	Yes	Yes	Yes
Age dummies	Yes	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes
Household dummies	No	Yes	No	Yes
R^2	0.011	0.010	0.027	0.025
Obs	35,959	35,959	29,329	29,329

Notes: Reported are the estimated log consumption in periods before and after the birth of the first child relative to consumption in the year of birth. Robust standard errors in brackets. *: $p < .05$, **:$p < .01$, ***:$p < .001$.

to log consumption the year *before* arrival of the first child. Non-durable consumption of low skilled households hardly change around the arrival of the first child while high skilled households increase consumption with around 15 percent in subsequent years. The estimates are, however, not precisely estimated and only on a five percent confidence level can I reject that high skilled households are unaffected by the arrival of their first child. This result is supported below by the fact that childless households have an almost similar consumption age profile as households who have children.

3.2 Childless Households

Figure 2 presents consumption and disposable income profiles for households with at least one child and childless households at completed fertility. Childless households are identified as households in which the wife is not registered as the mother to a child in 2010.[15] If the wife is not registered as a mother in 2010, I assume that the

[15]Virtually *all* childbirths after 1942 are matched to their mother. Only children born between January 1st 1942 and December 31th 1972 who either died or permanently immigrated to another country before

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.[16]

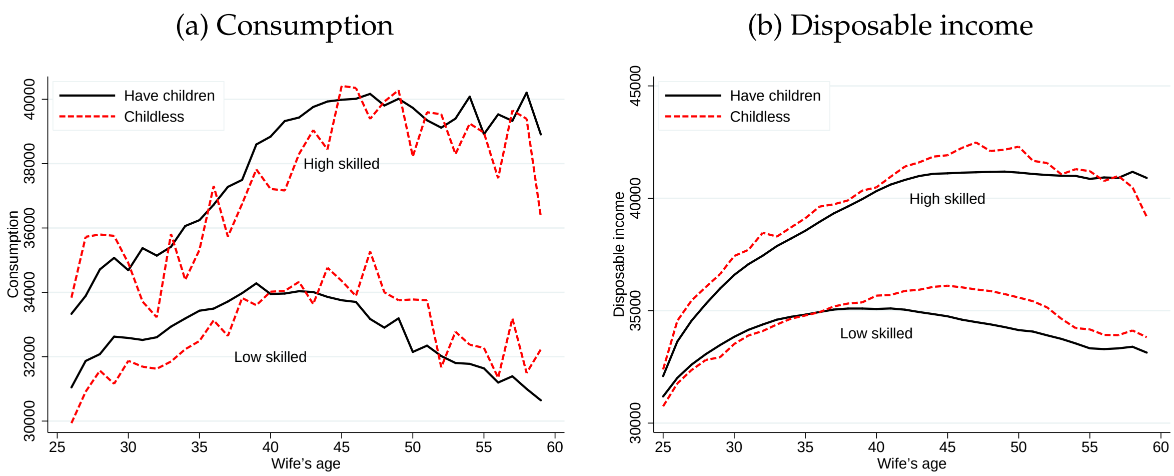

Figure 2 – Consumption and Disposable Income in Childless Households and Households with Children at Completed Fertility.

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.

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[17] 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

January 1st 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.

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

[17]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.

consumption/savings motives.[18]

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.

4 Calibrated Parameters

To keep the estimation procedure tractable, I reduce the number of parameters to be estimated by successively solving the structural model by calibrating some parameters in a first step.

Table 4 reports the values and sources for the calibrated parameters. The exogenous drop in permanent income when households retire, $\varkappa$, is calibrated to 90 percent 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.

Table 4 – Calibrated Parameters.

	Parameter	Value	Source
G_t	Permanent gross income growth	Fig. A1	Author's calculations: see text
R	Real gross interest rate	1.03	Gourinchas and Parker (2002)
κ	Credit limit	0.60	Author's calculations: see text
$\wp$	Probability of low income shock	0.10	Author's calculations: see text
μ	Size of low income shock	0.30	Author's calculations: see text
$\varkappa$	Replacement rate in retirement	0.90	The Danish Ministry of Finance (2003)

The social security system in Denmark seems compatible with a 10 percent risk of household income being reduced to 30 percent. Danish households are allowed to be net-borrowers by 60 percent of annual permanent income. These three values ($\kappa = .6$, $\mu = 0.3$ and $\wp = .01$) are somewhat arbitrary and have been chosen to provide a reasonable fit of the model for the bottom distribution of resources. Figure 6 below illustrates this by plotting within-percentile average consumption-income ratios against the household resources (also normalized by income). There is substantial variation

[18] If childless households have different preferences than households who have children, the age profiles might be similar due to these differences.

in consumption in the bottom distribution of resources for particularly young households and the calibrated parameters (along with the estimated preferences in Table 6) provide a good fit on average.

The permanent and transitory income shock variances 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

$$\begin{aligned} \hat{\sigma}_\eta^2 &= \text{cov}(\Delta\tilde{y}_t, \tilde{y}_{t+1} - \tilde{y}_{t-2}), \\ \hat{\sigma}_\varepsilon^2 &= -\text{cov}(\Delta\tilde{y}_t, \Delta\tilde{y}_{t+1}). \end{aligned}$$

Table 5 presents the estimated variance components. 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 an order of magnitude smaller than those typically estimated for the US.[19] 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, e.g., the US.[20] The Danish tax system is one of the most progressive tax schedules in the world with a marginal tax rate of more than 60 percent in 2010 for top earners. Around 40 percent where top earners in 2010. The progressive tax system reduces the dispersion in disposable income significantly. Finally, the Danish administrative registers also tend to be less noisy compared to surveys (Browning and Leth-Petersen, 2003), reducing the transitory income shock variance.

Table 5 – Permanent and Transitory Income Shock Variances.

	All		Low skilled		High skilled	
	Est	SE	Est	SE	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.

The income growth rate, G_t, can be estimated by taking logs of the income process

[19] Blundell, Pistaferri and Preston (2008) report $\sigma_\eta^2 \in [.0057, 0.0333]$ and $\sigma_\varepsilon^2 \in [.0190, 0.0753]$ depending on the combination of year, cohort and educational background, using the PSID. Gourinchas and Parker (2002), also using the PSID, calibrate $\sigma_\eta^2 = 0.0212$, $\sigma_\varepsilon^2 = 0.0440$.

[20] 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.

specified in section 2 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

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

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 C 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.

5 Identification of Remaining Parameters

An advantage and motivation for using the Danish register data in this study is the availability of reliable wealth and income data. This information allow me to evaluate the economic contingencies faced by households when performing their intertemporal allocation of consumption and savings. Furthermore, since the entire Danish population is available, millions of household-level observations provide variation in resources across household age and household composition.

As I will elaborate on below, the implemented Maximum Likelihood estimator minimizes the squared difference between household is imputed consumption measure and the predicted optimal consumption given household is level of resources and number and age of children. To establish intuition on what type of variation in the data identify the model parameters it is instructive to focus on the consumption functions; how optimal consumption depends on the level of resources. I first discuss how the parameters governing the taste shifter, θ, is identified and then turn to how the remaining parameters are identified below.

The Effect of Children on the Marginal Utility of Consumption, θ. Figure 3 illustrates how the parameters related to the effect of children on marginal utility of consumption, θ, are identified from the data. Panel (a) illustrates three hypothetical data points in the resources-consumption space. Both measures are normalized by permanent income. The blue star ($\star$) represents a 30 years old household with a child. The blue circle ($\bullet$) represents a 30 year old childless household, and the red circle ($\bullet$) represents a 35 year old childless household.[21]

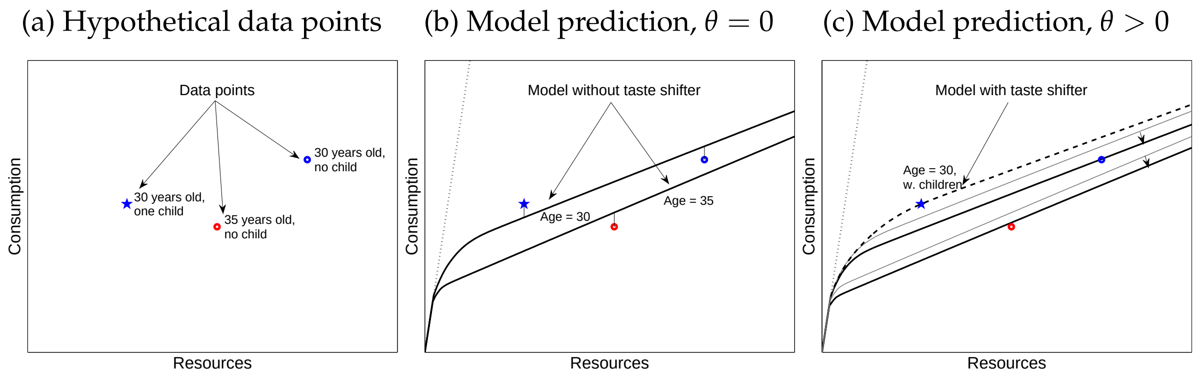

Figure 3 – Identification of the Effect of Children on Marginal Utility, θ.

Notes: Panel (a) illustrates three hypothetical data points in the permanent income adjusted resources/-consumption space. The blue star ($\star$) represents a 30 years old household with one child present. The blue circle ($\bullet$) represents a 30 year old childless household, and the red circle ($\bullet$) represents a 35 year old childless household. Panel (b) illustrates how structural parameters are estimated by *minimizing the squared vertical distance* between the observed data points and the model-predicted optimal consumption levels. Children are not allowed to affect household behavior in panel (b) while panel (c) introduce a positive effect of children on marginal utility.

Panel (b) illustrates how structural parameters are estimated by *minimizing the squared vertical distance* between the observed data points and the model-predicted optimal consumption levels. Children are not allowed to affect household behavior in panel (b) while panel (c) introduces a positive effect of children on the marginal utility of consumption. The positive effect results in higher optimal consumption for households who have children and lower level of consumption for fertile childless households (compared to panel (b)). The latter stems from the fact that households – also childless – form expectations over the arrival of children in the future. To be able to increase consumption in the event that a child arrives in the future, the 30 and 35 year old childless households should reduce their consumption today.

Variation in consumption (for a given level of resources and permanent income) across households with and without children identify θ. For a given number and age of children, differences in consumption across households in and after the fertile phase of the life cycle also help to identify the effect of children on the marginal utility of consumption. This is because once the fertile phase ends, the household *knows* that (more)

[21] These households could be the same households observed over several years, different households in the same year, or a combination of both.

children will never arrive while households in the fertile age assigns some positive probability of a child arriving. Finally, within-household time variation in consumption across periods with different number (and age) of children identify how much children affect the marginal utility of consumption.

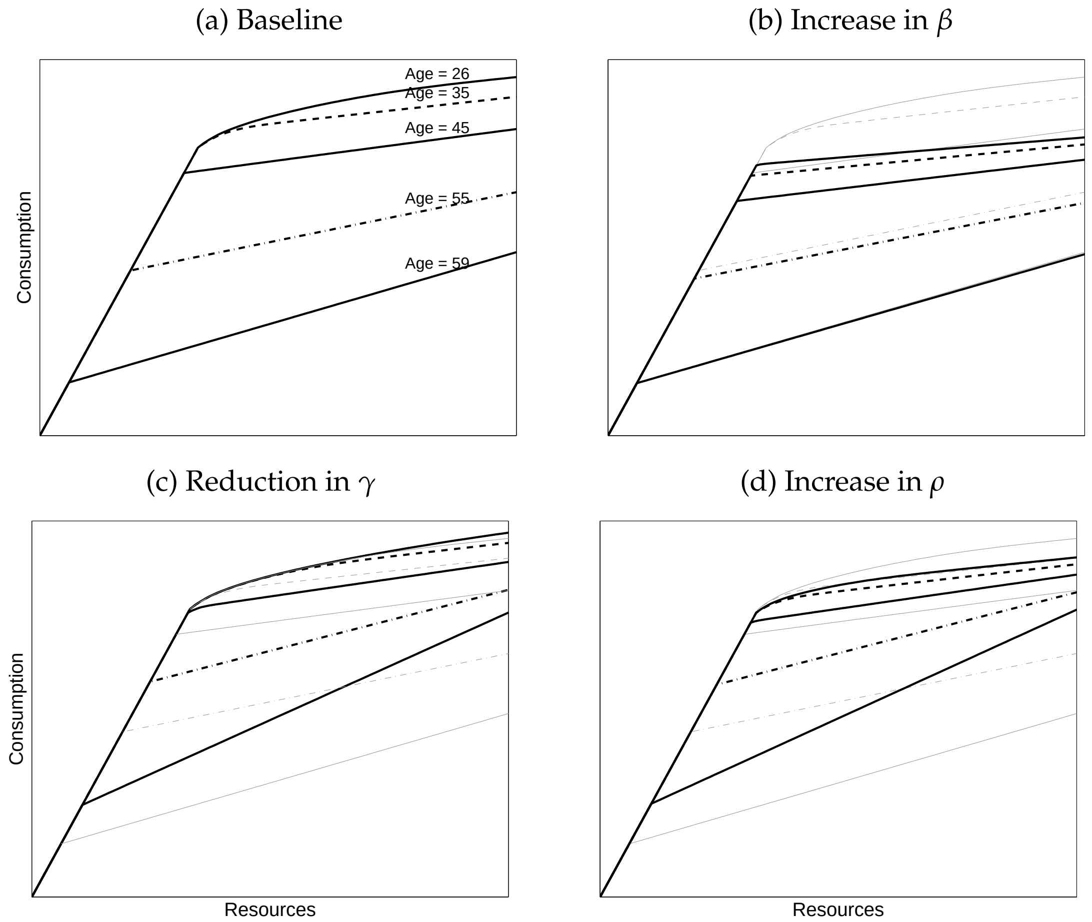

Figure 4 – Implied Consumption Functions for Different Parameter Values.

Notes: Panel (a) serves as benchmark and is reproduced in gray in all other panels. The effect of increasing the discount factor, β, is illustrated in panel (b) and panel (c) and (d) illustrate the effects of decreasing γ and increasing ρ, respectively.

The Discount Factor, β. To illustrate how the remaining structural parameters are identified, Figure 4 presents implied changes in the model predictions from changing the values for the risk aversion, ρ, discount factor, β, and value of consumption in retirement, γ. The implied behavior at ages 26, 35, 45, 55 and 59 are shown to illustrate the life cycle effects from changes in parameter values. The upper left panel (a) serves as a benchmark and is reproduced in gray in all other panels.

The effect of increasing the discount factor, β, is illustrated in panel (b). Since the change implies that households are more patient, young households decrease consumption significantly. A small change in the discount rate will affect young households while older households are not reacting noticeable. Importantly, increasing the discount factor will monotonically decrease the optimal level of consumption for all age groups. This is because the value of future utility from consumption is increased and it will be optimal to accumulate more wealth.

The Post-Retirement Motive, γ. Panel (c) in Figure 4 illustrates the effect of decreasing γ. When motives post retirement matters less, all age groups save less but particularly older households would want to consume more. This is intuitive because the value of consumption in retirement is most immediate for older households, provided a sufficient degree of impatience. In turn, a high (low) level of consumption of older households will identify a low (high) valuation of wealth post retirement, γ.

The Constant Relative Risk Aversion, ρ. Panel (d) in Figure 4 illustrates the effect of increasing the risk aversion parameter, ρ. A reader with good vision might be able to appreciate that younger households tend to respond more to a change in risk aversion, ρ, than to a change in the value of post retirement consumption, γ. The change in ρ and γ has purposely been chosen such to produce this almost similar picture. The relative risk aversion parameter will, thus, primary be identified through the consumption levels of young households with low levels of wealth. The consumption of households in the lower part of the distribution of resources are, thus, informative about the curvature of the consumption function, which the CRRA parameter influences in the model.

6 Estimation Results

To estimate how children affect 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, $\{C_t^{\star}(M_t, P_t, \mathbf{z}_t|\Theta)\}_1^T$.[22] In principle, Mathematical Programming with Equilibrium Constraints (MPEC), proposed by Su and Judd (2012) could be used to estimate the parameters. However, because the model in the present study is a life cycle model with a large state space, MPEC most likely would be much slower than the NFXP (Jørgensen, 2013).

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 $i = 1, \ldots, N$

[22] Consult Appendix B for details on how the EGM solves for optimal consumption.

households are observed in $t = 1, \dots T_i$ periods. I assume that imputed non-durable (normalized) consumption in the Danish registers is observed with additive *iid* Gaussian measurement error, ξ,

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

where small letter variables denote normalized measures (e.g., $c_t = C_t/P_t$). The (mean) log-likelihood function is, thus,

$$\mathcal{L}(\Theta, \sigma_\xi) = -\frac{1}{2}\sum_{i=1}^{N}\frac{1}{NT_i}\sum_{t=1}^{T_i}\left\{\log(2\pi\sigma_\xi^2) + \frac{(c_{it} - c_t^\star(M_{it}, P_{it}, \mathbf{z}_{it}|\Theta))^2}{\sigma_\xi^2}\right\}. \quad (3)$$

Because consumption levels are related to end-of-period wealth, a_t, through the intratemporal budget, $c_t = m_t - a_t$, the implemented estimator effectively match household net worth levels.

Optimal behavior is found numerically and the likelihood function in (3) 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.[23] Jørgensen (2013) shows that an estimation approach similar to that outlined above can uncover parameters like the relative risk aversion, ρ, from similar models. For completeness, Table A2 in the supplemental material reports mean (and standard deviation) 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.

Table 6 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) reports 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,

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

[23] Ackerberg, Geweke and Hahn (2009) correct a result (Proposition 2) of Fernández-Villaverde, Rubio-Ramírez and Santos (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.

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. As shown in the descriptive analysis, estimating a version of the model without precautionary motives and applying standard Euler equation estimators yield a strong correlation between consumption growth and arrival of children. 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 evidence 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 the income uncertainty are much less severe. I find again that children are found to increase marginal utility much more when excluding precautionary motives.

The marginal utility of consumption does not increase significantly in the age of children ($\hat{\theta}_{j2} \approx \hat{\theta}_{j1} \, \forall j = 1, 2, 3$). Contrary to these results, Browning and Ejrnæs (2009) finds that consumption increases significantly in the age of children by estimating log consumption growth regressions on synthetic cohort panels constructed from the FES. As argued in Jørgensen (2014) and recognized in Browning and Ejrnæs (2009), changing demographics likely proxy for credit constraints and precautionary savings motives. The estimates reported in Table 6 supports this explanation. In contrast to the results in Browning and Ejrnæs (2009), the estimated parameters do also not support economies of scale.

The Danish welfare system provides free health care, free schooling and significant childcare subsidies. For example, childcare is heavily subsidies and approximately 70 percent 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 subsidy* (in Danish "Statens Uddannelsesstøtte", abbreviated "SU") of around a thousand US dollars ($5,839$ Danish kroner) in 2014.[24] In turn, raising children in Denmark could be less costly than many other countries, 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 and I illustrate below how larger effects are estimated when ignoring precautionary motives.

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

[24] 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.

households. The discount factor, β, is estimated to be around .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 PSID, 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 .25.

The γ, summarizing all post-retirement 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 analyze how children affect behavior prior to retirement, this parsimonious parametrization has been implemented. Taken at face value, however, the estimated γs would suggest that consumption should *increase* at retirement, contrary to most existing empirical evidence seem to suggest. To reconcile my estimates with existing studies, recall that among other simplifications income, health and survival uncertainty has been assumed away post retirement. If, for example, health, income and survival is uncertain and households gain utility from leaving bequests, consumers should (also prior to retirement) accumulate more wealth than the implemented model suggests. In the current model this reduced consumption (increased saving) prior to and around retirement will result in a higher estimated value of γ.

Table 6 – 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 ≤ 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 ≤ 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 ≤ 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[$<.001$]
# 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.

6.1 Model Fit

The model fits the Danish register data very well. Figure 5 illustrates the age profiles of imputed consumption in the register data and optimal consumption predicted by the estimated model. Panel (a) shows that the mean age profiles are almost identical and panel (b) shows only minor differences between the median age profiles. The model slightly underestimates the average consumption level early in the life cycle for households younger than 30. This underestimation of consumption for young households is *not* due to underestimation of the effect of children on marginal utility, as seen in Figure 6.

(a) Mean consumption (b) Median consumption

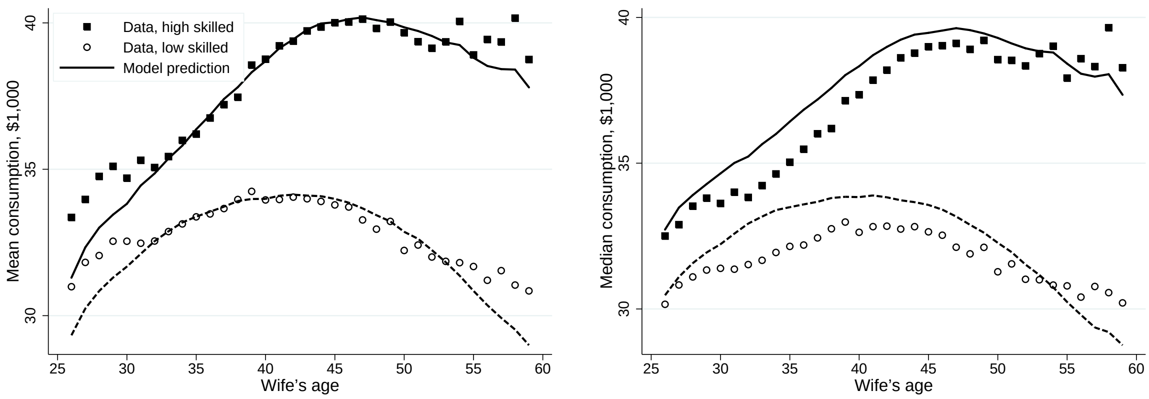

Figure 5 – Actual and Predicted Consumption profiles.

Notes: Figure 5 illustrates the mean (panel a) and median (panel b) age profile for actual (imputed) and predicted consumption in the Danish registers.

Figure 6 illustrates the fit of the model for childless households and households with children. Each dot is calculated as the average consumption within a given percentile of resources observed in the Danish data. Focusing on the left panel (a) in Figure 6 in which the households are aged 26-30 years old, it is evident that the underestimation of the average consumption for younger households, reported in figure 5, is not due to underestimation of the effect of children, θ. When estimating the model, consumption is set to zero for households with resources below the borrowing limit of 60 percent 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 circles) and households with children (black crosses) among this group with very low resources.

Figure 6 confirms the estimation results: childless households and households with children do not differ significantly. The figure also illustrates that the calibrated posi-

(a) Wife aged 26-30 (b) Wife aged 36-40 (c) Wife aged 46-50

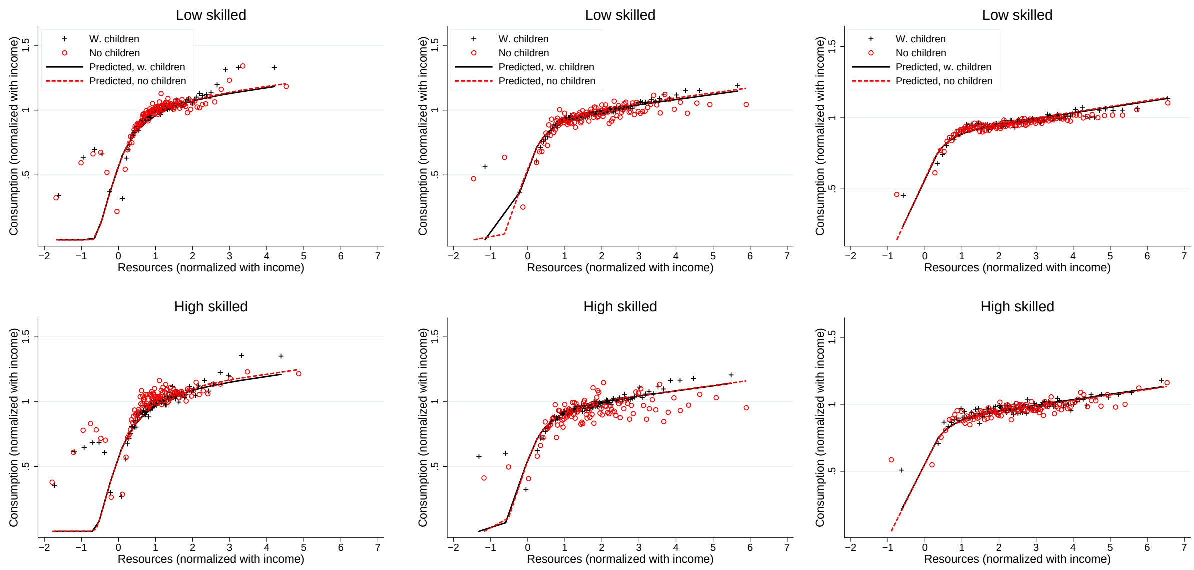

Figure 6 – Consumption Functions.

Notes: Figure 6 illustrates the average consumption in groups based on percentiles of available resources. Each dot represents a percent of observations within each age and child/no child group. Average consumption and predicted consumption from the estimated model is plotted to illustrate how the calibrated values of $\kappa = .6$, $\mu = 0.3$ and $\wp = .01$ produce a good fit of the model to the Danish data.

tion of the credit constrain ($\kappa = .6$) and the low income shock ($\mu = .3$ with probability $\wp = 0.01$) provides a good fit for households in the lower end of the resources distribution. Further, the model fits the consumption functions remarkably well over the life cycle.

6.2 Robustness

I interpret the results as indicating that precautionary motives are more important than children in explaining the observed life cycle profiles. I investigate this claim by estimating a model in which households face much less severe precautionary motives. Particularly, Table 7 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.[25] If introducing precautionary motives explain why I find minuscule effects of children, removing these motives should yield larger estimated effects of children.

I find that children increase the marginal utility of consumption significantly more,

[25]In the numerical implementation, I set $\kappa = 10$ 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, $\kappa = 10$, is not binding.

Table 7 – 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)	0489 (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
$\max_j \lvert\partial\mathcal{L}(\Theta)/\partial\Theta_j\rvert$	$8.3e-7$	$3.3e-5$	$2.7e-6$	$1.1e-5$
# 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 post-retirement motive, γ, is fixed at their estimated values in columns (3) of Table 6. Standard errors in brackets are based on the inverse of the hessian.

with around 8 to 9 percent, when reducing the precautionary motives in the estimated model. The estimates are, however, still much lower than reported in existing literature.[26] The risk aversion parameter, ρ, is also estimated much higher because to rationalize the saving behavior of households they must be much more risk averse now the borrow limit is much higher and income uncertainty lower. If the wealth accumulation was due to children, the θ should have increased much more than is has.

For computational tractability, income has been assumed independent of children throughout. A valid concern could be that labor market participation decisions are related to the presence of children (Heckman, 1974). If consumption and leisure are substitutes, and consumers tend to increase leisure time when having children (Calhoun and Espenshade, 1988), ignoring labor market supply will result in a downwards bias in the estimated effect of children on the marginal utility of consumption. To in-

[26] As argued above and in Jørgensen (2014), this is likely due to an upwards bias in the Euler equation estimators based on synthetic cohort panels, applied in most of the existing studies.

vestigate if reduced labor 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 8.

Table 8 – 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 ≤ 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 ≤ 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 ≤ 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
$\max_j \lvert\partial\mathcal{L}(\Theta)/\partial\Theta_j\rvert$		$6.1e-6$	$8.2e-6$	$6.8e-6$	$1.2e-5$	$1.9e-5$	$2.1e-5$
# 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.

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 percent for low skilled and 16 percent for high skilled. This might reflect preference heterogeneity or stem from a higher estimated value of retirement, γ.[27] Further, permanently working households may be in generally more stable environments leading to less wealth accumulation, captured by a higher estimate of ρ.

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 increase 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 consumption around and after childbirth will likely not increase in the same extend as the model (ignoring housing purchases) would suggest.[28] 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 8 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 homeowners.

Estimation results from only using the sub-sample of working home owners suggest that the first child increases the marginal utility with around 7 to 10 percent 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 (insignificant) economies of scale for this sub-sample.

7 Concluding Discussion

I have estimated the effect of children on the marginal utility of consumption while allowing income uncertainty, credit constraints and post-retirement motives to also affect household behavior. A novelty of this study is that precautionary motives are included while the number and age of all children can affect the marginal utility of consumption. To allow for several life cycle motives simultaneously, I estimate a life cycle model of intertemporal consumption and saving in the presence of children. The model is estimated by Maximum Likelihood (ML) using high quality Danish administrative register data, providing detailed longitudinal information on income, assets,

[27] 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.

[28] The imputed consumption is given as disposable income net of change in the household net-worth. 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.

liabilities and household characteristics. Household 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 that the effect of children on non-durable consumption is economically negligible. This is in stark contrast to most existing studies suggesting that children increase the marginal utility of consumption with 40-150 percent. In a recent study, Browning and Ejrnæs (2009) illustrates 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 suggested by e.g., Gourinchas and Parker, 2002 and Cagetti, 2003), the finding of an important role for children along an age effect might be proxy for precautionary motives. My results suggest that when allowing for precautionary motives *simultaneously* with an age and scale effect of children on the marginal utility of consumption, children play a minor role and I find no age effect of children.

The results do not seem to be driven by differences in the data sources used in the current and previous studies. Log (imputed) consumption growth in the Danish registers correlates with the arrival of children just as much as reported in the influential study by Attanasio, Banks, Meghir and Weber (1999) using the CEX. To reconcile my results with existing studies I illustrate how ignoring precautionary motives increases the estimated importance of children in the Danish data. I interpret the results as indicating that precautionary motives might play a larger role than children in explaining the observed consumption age profile.

Several interesting avenues for future research remains. 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. To allow for an arbitrary children, age and scale effect, several simplifying assumptions has been deployed. This has been done solely for computational tractability. For example, I ignore labor supply and home ownership choices although both likely are related to fertility. Noticeably, estimating the model on a sub-sample of permanently working home owners suggests only a small role for children as they at most increase marginal utility with around ten percent. Including fertility, labor supply, and home ownership choices in the current framework is extremely interesting, but also a computationally demanding task, which I leave for future research.

References

ACKERBERG, D., J. GEWEKE AND J. HAHN (2009): "Comments on "Convergence Properties of the Likelihood of Computed Dynamic Models"," *Econometrica*, 77(6), 2009–2017.

ADDA, J., C. DUSTMANN AND K. STEVENS (2015): "The Career Costs of Children," Working paper.

ALAN, S., O. ATTANASIO AND M. BROWNING (2009): "Estimating Euler equations with noisy data: two exact GMM estimators," *Journal of Applied Econometrics*, 24(2), 309–324.

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

ATTANASIO, O. P. AND M. BROWNING (1995): "Consumption over the Life Cycle and over the Business Cycle," *The American Economic Review*, 85(5), 1118–1137.

ATTANASIO, O. P. AND G. WEBER (1995): "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," *The Journal of Political Economy*, 103(6), 1121–1157.

——— (2010): "Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy," *Journal of Economic Literature*, 48(3), 693–751.

BANKS, J., R. BLUNDELL AND I. PRESTON (1994): "Life-cycle expenditure allocations and the consumption costs of children," *European Economic Review*, 38(7), 1391–1410.

BICK, A. AND S. CHOI (2013): "Revisiting the effect of household size on consumption over the life-cycle," *Journal of Economic Dynamics and Control*, 37(12), 2998–3011.

BLUNDELL, R., M. BROWNING AND C. MEGHIR (1994): "Consumer Demand and the Life-Cycle Allocation of Household Expenditures," *The Review of Economic Studies*, 61(1), 57–80.

BLUNDELL, R., M. C. DIAS, C. MEGHIR AND J. M. SHAW (2013): "Female Labour Supply, Human Capital and Welfare Reform," Working Paper 19007, NBER.

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

BROWNING, M., A. DEATON AND M. IRISH (1985): "A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle," *Econometrica*, 53(3), pp. 503–544.

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

BROWNING, M. AND S. LETH-PETERSEN (2003): "Imputing Consumption from Income and Wealth Information," *The Economic Journal*, 113(488), 282–301.

CAGETTI, M. (2003): "Wealth Accumulation Over the Life Cycle and Precautionary Savings," *Journal of Business and Economic Statistics*, 21(3), 339–353.

CALHOUN, C. A. AND T. J. ESPENSHADE (1988): "Childbearing and Wives' Foregone Earnings," *Population Studies*, 42(1), pp. 5–37.

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

——— (2001): "Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation)," *Advances in Macroeconomics*, 1(1), Article 6.

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

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

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

FERNÁNDEZ-VILLAVERDE, J., J. F. RUBIO-RAMÍREZ AND M. S. SANTOS (2006): "Convergence Properties of the Likelihood of Computed Dynamic Models," *Econometrica*, 74(1), 93–119.

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

HAMILTON, J. D. (1994): *Time Series Analysis.* Princeton University Press, Princeton, New Jersey.

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*, 64(1), pp. 188–194.

HONG, J. H. AND J.-V. RÍOS-RULL (2012): "Life Insurance and Household Consumption," *American Economic Review*, 102(7), 3701–30.

IRVINE, I. (1978): "Pitfalls in the Estimation of Optimal Lifetime Consumption Patterns," *Oxford Economic Papers*, 30(2), pp. 301–309.

JØRGENSEN, T. H. (2013): "Structural estimation of continuous choice models: Evaluating the EGM and MPEC," *Economics Letters*, 119(3), 287–290.

——— (2014): "Euler Equation Estimation: Children and Credit Constraints," Discussion Paper 14-25, Department of Economics, University of Copenhagen.

LOVE, D. A. (2010): "The Effects of Marital Status and Children on Savings and Portfolio Choice," *Review of Financial Studies*, 23(1), 385–432.

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

RUST, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55(5), 999–1033.

SCHOLZ, J. K. AND A. SESHADRI (2009): "Children and Household Wealth," Working paper, University of Wisconsin-Madison.

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.

SU, C.-L. AND K. L. JUDD (2012): "Constrainted Optimization Approaches to Estimation of Structural Models," *Econometrica*, 80(5), 2213–2230.

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*, 59(3), pp. 324–330.

Online Appendix (Not for Publication)

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

Thomas H. Jørgensen
University of Copenhagen
May 21, 2015

A Additional Figures and Tables

Table A1 – Selection Criteria.

	Household-time observations	Share of original sample
Original, all couples	16,268,110	1.000
Wife age	8,832,011	0.543
Husband age	5,874,962	0.361
No information on education	5,848,884	0.360
No more than 3 children	5,716,046	0.351
Both in labor force	2,238,076	0.138
No negative disposable income	2,231,204	0.137
Child outside allowed range	2,230,326	0.137
Extreme observations	1,522,994	0.094
Resources missing	1,281,952	0.079

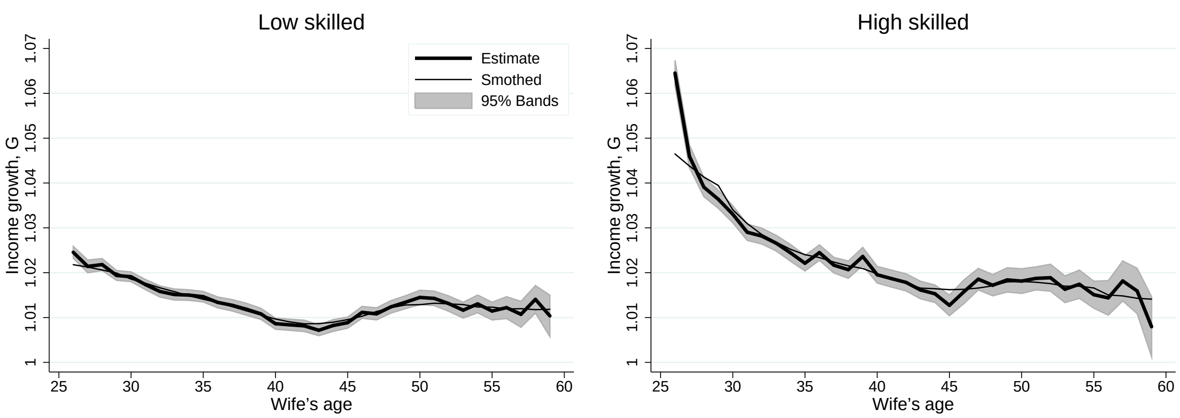

Figure A1 – Estimated Age Profiles of Income Growth, $\hat{G}_t$.

Notes: Figure A1 illustrates the estimated income growth age profile. Permanent income growth is estimated from equation (2).

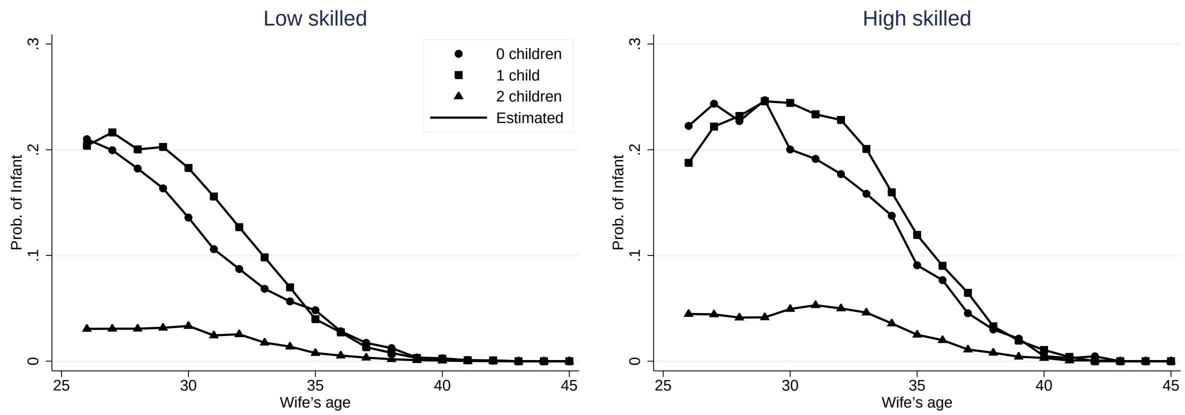

Figure A2 – Estimated Children Arrival Probabilities.

Notes: Figure A2 illustrates the estimated arrival rate of children in the Danish registers. Arrival depend on the age of the wife and the number of children present in the household.

Table A2 – Monte Carlo Results, Structural Estimation.

	$\theta_0 = 0.0$	$\theta_0 = 0.1$	$\theta_0 = 0.3$	$\theta_0 = 0.5$
$N = 1000$, $T = 5$	0.004	0.115	0.300	0.504
	(0.034)	(0.062)	(0.099)	(0.080)
$N = 1000$, $T = 20$	−0.000	0.106	0.307	0.510
	(0.017)	(0.039)	(0.078)	(0.057)
$N = 50000$, $T = 5$	0.000	0.099	0.299	0.501
	(0.005)	(0.010)	(0.017)	(0.013)
$N = 50000$, $T = 20$	0.000	0.100	0.301	0.502
	(0.002)	(0.004)	(0.010)	(0.009)

Notes: Table A2 reports means and standard deviations (in parenthesis) across 50 Monte Carlo runs. Data is simulated from the model described in Section 2 with parameters set to $\beta = .95$, $\rho = 2$, $R = 1.03$, $\kappa = 0$, $\wp = 0$, $\mu = 0$, $\gamma = 1.1$, $\varkappa = 0.8$, $\sigma_\varepsilon^2 = 0.005$, and $\sigma_\eta^2 = 0.005$. Income growth (G_t) is 1.05 when younger than 25, then 1.03 until age 30, and then 1.01 until age 40 where income is constant in all subsequent periods.

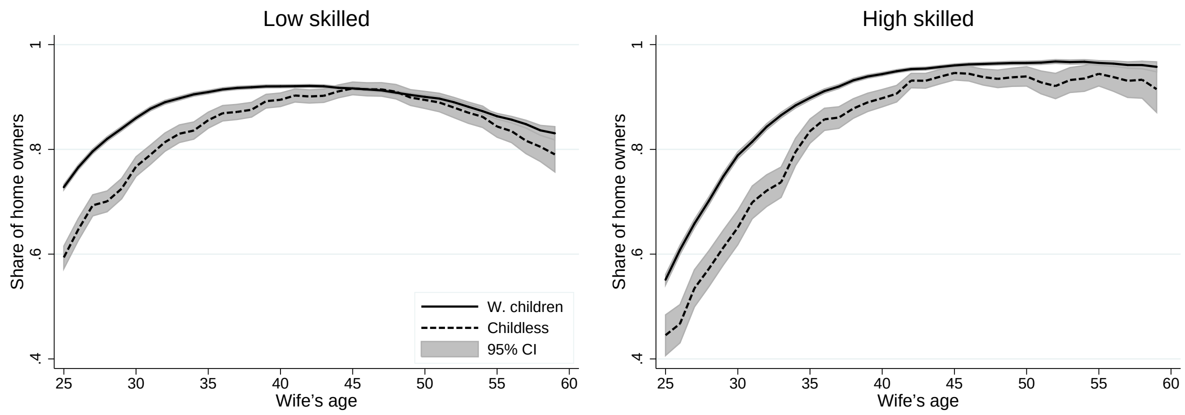

Figure A3 – Home Ownership Over the Life Cycle.

Notes: 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, I assume that the household will remain childless.

B Solving the Model

To reduce the number of state variables, all relations are normalized by permanent income, P_t, and small letter variables denote normalized quantities (e.g., $c_t = C_t/P_t$). The model is solved recursively by backwards induction, starting with the terminal period, T. Within a given period, optimal consumption is found using the Endogenous Grid Method (EGM) by Carroll (2006).

The EGM constructs a grid over end-of-period wealth, a_t, rather than beginning-of-period resources, m_t. Denote this grid of Q points as $\hat{a}_t = (\underline{a}_t, a_t^1, \ldots, a_t^{Q-1})$ in which $\underline{a}_t$ is a lower bound on end-of-period wealth that I will discuss in great detail below. The endogenous level of beginning-of-period resources consistent with end-of-period assets, $\hat{a}_t$, and optimal consumption, $c_t^\star$, is given by $m_t = \hat{a}_t + c_t^\star(m_t, \mathbf{z}_t)$.

In the terminal period, independent of the presence of children, households consume all their remaining wealth, $c_T = m_T$. In preceding periods, in which households are retired, consumption across periods satisfy the Euler equation

$$u'(c_t) = \max\left\{u'(m_t)\ ,\ R\beta\frac{v(\mathbf{z}_{t+1};\theta)}{v(\mathbf{z}_t;\theta)}u'(c_{t+1})\right\},\ \forall t \in [T_r, T],$$

where consumption cannot exceed available resources. When retired, households do not produce new offspring and the age of children ($\mathbf{z}_t$) evolves deterministically.

The normalized consumption Euler equation in periods prior to retirement is given by

$$u'(c_t) = \max\left\{u'(m_t + \kappa)\ ,\ R\beta\mathbb{E}_t\left[\frac{v(\mathbf{z}_{t+1};\theta)}{v(\mathbf{z}_t;\theta)}u'(c_{t+1}G_{t+1}\eta_{t+1})\right]\right\},\ \forall t < T_r,$$

such that when $\hat{a}_t > -\kappa$ optimal consumption can be found by inverting the Euler equation

$$c_t^\star(m_t, \mathbf{z}_t) = \left(\beta R \mathbb{E}_t\left[\frac{v(\mathbf{z}_{t+1};\theta)}{v(\mathbf{z}_t;\theta)}(G_{t+1}\eta_{t+1})^{-\rho}\check{c}_{t+1}^\star(\underbrace{(G_{t+1}\eta_{t+1})^{-1}R\hat{a}_t + \varepsilon_{t+1}}_{=m_{t+1}}, \mathbf{z}_{t+1})^{-\rho}\right]\right)^{-\frac{1}{\rho}},$$

where $\check{c}_{t+1}^\star(m_{t+1}, \mathbf{z}_{t+1})$ is a linear interpolation function of optimal consumption next period, found in the last iteration. Since $\hat{a}_t$ is the constructed grid, it is trivial to determine in which regions the credit constraint is binding and not. I will discuss this in detail below.

The expectations are over next period arrival of children ($\mathbf{z}_{t+1}$) and transitory (ε_{t+1}) and permanent income shocks (η_{t+1}). Eight Gauss-Hermite quadrature points are used for each income shock to approximate expectations. $Q = 80$ discrete grid points are used in $\hat{a}_t$ to approximate the consumption function with more mass at lower levels of wealth to approximate accurately the curvature of the consumption function. The number of points was chosen such that the change i the optimized log likelihood did not change significantly, as proposed in Fernández-Villaverde, Rubio-Ramírez and Santos (2006).

The arrival probability of a child next period is a function of the wife's age and number of children today, $\pi_{t+1}(\mathbf{z}_t)$. No more than three children can live inside a household at a given point in time and infants cannot arrive when the household is older than 43. The next period's state is therefore calculated by increasing the age of children by one and if the age is 21, the child moves. In principle, there is $22^3 = 10,648$ combinations three children can be either not present (NC) or aged zero through 20. To reduce computation time, children are organized such that child one is the oldest at all times, the second child is the second oldest and child three is the youngest child. To illustrate, imagine a household which in period t has, say, two children aged 20 and 17, $\mathbf{z}_t = (\text{age}_{1,t} = 20, \text{age}_{2,t} = 17, \text{age}_{3,t} = \text{NC})$, then, in period $t+1$, only one child will be present; $\mathbf{z}_{t+1} = (\text{age}_{1,t+1} = 18, \text{age}_{2,t+1} = \text{NC}, \text{age}_{3,t+1} = \text{NC})$, given no new offspring arrives. Had new offspring arrived, then $\text{age}_{2,t+1} = 0$.

B.1 Credit Constraint and Utility Induced Constraints

Since the EGM works with end-of-period wealth rather than beginning-of-period resources, credit constraints can easily be implemented by adjusting the lowest point in the grid, $\underline{a}_t$. The potentially binding credit constraint next period is implemented by the rule, $c_{t+1}^\star = m_{t+1}$ if m_{t+1} is lower than some threshold level, $\mathsf{m}_{t+1}^\star$. Including the credit constraint as the lowest point, $\underline{a}_{t+1} = -\kappa$, the lowest level of resources endogenously determined in the last iteration, $\underline{m}_{t+1}$, is the exact level of resources where households are on the cusp of being credit constrained, i.e., $\mathsf{m}^{\star}_{t+1} = \underline{m}_{t+1}$. This ensures a very accurate interpolation and requires no additional handling of shadow prices of resources in the constrained Euler equation, denoted λ_{t+1} in Section **??**.

Besides the exogenous credit constraint, κ, a "natural" or utility induced self-imposed constraint can be relevant such that the procedure described above should be modified slightly. This is because households want to accumulate enough wealth to buffer against a series of extremely bad income shocks to ensure strictly positive consumption in all periods even in the worst case possible.

Proposition 1. *The lowest possible value of normalized end-of-period wealth consistent with the model, periods prior to retirement, can be calculated as*

$$\underline{a}_t = -\min\{\Omega_t, \kappa\} \; \forall t \leq T_r - 2$$

where, denoting the lowest possible values of the transitory and permanent income shock as $\underline{\varepsilon}$, *and* $\underline{\eta}$, *respectively,* Ω_t *can be found recursively as*

$$\Omega_t = \begin{cases} R^{-1} G_{T_r} \underline{\varepsilon}_{T_r} \underline{\eta}_{T_r} & \textit{if } t = T_r - 2, \\ R^{-1}(\min\{\Omega_{t+1}, \kappa\} + \underline{\varepsilon}_{t+1}) G_{t+1} \underline{\eta}_{t+1} & \textit{if } t < T_r - 2. \end{cases}$$

Proof. Define $\underline{\mathbb{E}}_t[\cdot]$ as the *worst-case* expectation given information in period t and note that in the last period of working life, $T_r - 1$, households must satisfy $A_{T_r-1} \geq 0$. In the second-to-last period during working life, households must then leave a positive amount of resources in the worst case possible,

$$\begin{aligned} \underline{\mathbb{E}}_{T_r-2}[M_{T_r-1}] &> 0, \\ \underline{\mathbb{E}}_{T_r-2}[R A_{T_r-2} + Y_{T_r-1}] &> 0, \\ R A_{T_r-2} + G_{T_r-1} P_{T_r-2} \underline{\varepsilon}_{T_r-1} \underline{\eta}_{T_r-1} &> 0, \\ &\Updownarrow \\ A_{T_r-2} &> -\underbrace{R^{-1} G_{T_r-1} \underline{\varepsilon}_{T_r-1} \underline{\eta}_{T_r-1}}_{\equiv \Omega_{T_r-2}} P_{T_r-2}. \end{aligned}$$

Combining this with the exogenous credit constraint, κ, end-of-period wealth should satisfy

$$A_{T_r-2} > -\min\{\Omega_{T_r-2}, \kappa\} P_{T_r-2}.$$

In period $T_r - 3$, households must save enough to insure strictly positive consumption next period while satisfying the constraint above, in the worst case possible,

$$\begin{aligned}
\mathbb{E}_{T_r-3}[M_{T_r-2}] &> -\min\{\Omega_{T_r-2},\kappa\}\mathbb{E}_{T_r-3}[P_{T_r-2}], \\
RA_{T_r-3} + G_{T_r-2}P_{T_r-3}\varepsilon_{T_r-2}\eta_{T_r-2} &> -\min\{\Omega_{T_r-2},\kappa\}G_{T_r-2}P_{T_r-3}\eta_{T_r-2}, \\
&\Updownarrow \\
A_{T_r-3} &> -\underbrace{R^{-1}(\min\{\Omega_{T_r-2},\kappa\}+\varepsilon_{T_r-2})G_{T_r-2}\eta_{T_r-2}}_{\equiv \Omega_{T_r-3}}P_{T_r-3},
\end{aligned}$$

such that end of period wealth in period $T_r - 3$ should satisfy

$$A_{T_r-3} > -\min\{\Omega_{T_r-3},\kappa\}P_{T_r-3}.$$

Hence, we can find Ω_t recursively by the formula in Proposition 1 and calculate the lowest value of the grid of normalized end-of-period wealth as $\underline{a}_t = -\min\{\Omega_t,\kappa\}$. □

C Permanent Income: the Kalman Filter

Here, I give a brief description of the implementation of the Kalman Filter used to uncover household level permanent income. See, e.g., Hamilton (1994, ch. 13) for a detailed description of the Kalman Filter. Formulating the log income process on State Space form yields

$$\begin{aligned}
\mathbf{z}_{it} &= \mathfrak{A} + \mathfrak{B}\mathbf{x}_{it} + \mathbf{v}_{it}, \\
\mathbf{x}_{it} &= \mathfrak{C}_t + \mathfrak{D}\mathbf{x}_{it-1} + \mathbf{u}_{it},
\end{aligned}$$

where

$$\begin{aligned}
\mathbf{z}_{it} &= \log Y_{it}, \quad \mathfrak{A} = -\frac{1}{2}\sigma_\varepsilon^2, \qquad \mathfrak{B} = 1, \quad \mathbf{v}_{it} \sim \mathcal{N}(0,\sigma_\varepsilon^2), \\
\mathbf{x}_{it} &= \log P_{it}, \quad \mathfrak{C}_t = -\frac{1}{2}\sigma_\eta^2 + \log G_t, \quad \mathfrak{D} = 1, \quad \mathbf{u}_{it} \sim \mathcal{N}(0,\sigma_\eta^2),
\end{aligned}$$

and Y_{it} is observed household income, G_t, σ_ε^2, and σ_η^2 are known (estimated) parameters and $\log P_{it}$ is the unobserved log permanent income, I wish to uncover. For readability, I suppress i subscripts in what follows.

The Kalman Filter consists of a *prediction step* and an *updating step* where - given initial values, that I discuss below - the prediction step for the process at hand is,

$$
\begin{aligned}
\mu_{t|t-1} \equiv \mathbb{E}[\mathbf{x}_t|\Im_{t-1}] &= \mathfrak{C}_t + \mathfrak{D}\mu_{t-1|t-1}, \\
&= -\frac{1}{2}\sigma_\eta^2 + \log G_t + \mu_{t-1|t-1}, \\
\Sigma_{t|t-1} \equiv \mathbb{V}[\mathbf{x}_t|\Im_{t-1}] &= \mathfrak{D}\Sigma_{t-1|t-1}\mathfrak{D}' + \sigma_\eta^2, \\
&= \Sigma_{t-1|t-1} + \sigma_\eta^2,
\end{aligned}
$$

where $\Im_s$ denotes information known at time s. The updating step is given by

$$
\begin{aligned}
\mu_{t|t} \equiv \mathbb{E}[\mathbf{x}_t|\Im_t] &= \mu_{t|t-1} + K_t(\mathbf{z}_{it} - \mu_{t|t-1} - \mathfrak{A}), \\
&= \mu_{t|t-1} + K_t(\log Y_t - \mu_{t|t-1} + \frac{1}{2}\sigma_\varepsilon^2), \\
\Sigma_{t|t} \equiv \mathbb{V}[\mathbf{x}_t|\Im_t] &= (I - K_t\mathfrak{B})\Sigma_{t|t-1}, \\
&= (I - K_t)\Sigma_{t|t-1},
\end{aligned}
$$

where $\mu_{t|t} = \log \hat{P}_t$ is the "estimated" log permanent income and K_t is the *Kalman gain*,

$$
K_t = \Sigma_{t|t-1}(\Sigma_{t|t-1} + \sigma_\varepsilon^2)^{-1}.
$$

For each household, I identify the first year observed in the data (denoted $t = 0$) and use that observation as initial values for $\mu_{t|t}$ and $\Sigma_{t|t}$. Specifically, I assume that log income is at its population mean when first observed in the data, $\log Y_0 = \mathbb{E}[\log P_0 - \frac{1}{2}\sigma_\varepsilon^2 + \mathbf{v}_t|\Im_t]$, such that $\mu_{0|0} = \log Y_0 + \frac{1}{2}\sigma_\varepsilon^2$ and $\Sigma_{0|0} = \sigma_\varepsilon^2$.