

Jørgensen, Thomas H.

Doctoral Thesis

International migration: A destination country and migrant perspective

PhD Series, No. 172

Provided in Cooperation with:

University of Copenhagen, Department of Economics

Suggested Citation: Jørgensen, Thomas H. (2015) : International migration: A destination country and migrant perspective, PhD Series, No. 172, ISBN 978-87-91342-94-3, University of Copenhagen, Department of Economics, Copenhagen

This Version is available at:

<https://hdl.handle.net/10419/233436>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



PhD Dissertation

Thomas H. Jørgensen

Life-Cycle Consumption and Retirement within the Family

Supervisors: Mette Ejrnæs and Bertel Schjerning

Date of submission: October 30, 2014

Acknowledgments

My interest in applied microeconometrics was spurred by my student job in the Ministry of Economic and Business Affairs during my Master's studies. As a young student immigrating to Copenhagen from Jutland, the close connection with colleagues and the Danish register data excited me. The combination of access to the Danish register data and several academic courses on microeconometrics, taught by my supervisors to be, showed me how challenging and interesting applied econometrics can be.

I am deeply indebted to my supervisors, Professor Mette Ejrnæs and Associate Professor Bertel Schjerning. Without Bertel's profound knowledge and expertise within computation and estimation of dynamic stochastic programming models, I would never have been able to enter this interesting field of computational economics. Without the great supervision of Mette, the chapters in this dissertation would have been executed with much less confidence. Mette's ability to give precise and constructive suggestions have truly improved my PhD studies and the chapters in this dissertation. I am extremely happy that I have future projects together with both Mette and Bertel.

In general, the Department of Economics, University of Copenhagen, offered an excellent research environment. The academic discussions and the invaluable coffee breaks with close colleagues made everyday at the office much more fun and fruitful.

I also wish to thank Professor John Rust, who I visited in 2012 during my PhD. John was truly a great host and I enjoyed many academic and private conversations with him during my stay at Maryland. He inspired and continues to inspire my research in many ways. Obviously, his thorough and extremely applicable work is inspiring, but John also facilitated my participation in the (last) Initiative for Computational Economics (ICE) held in Chicago. I learned a lot from that summer workshop and my first chapter, which is now published, was initiated there. I am extremely excited and proud to be able to work together with John in future projects.

Finally, I wish to thank my family and friends, who all provided me with the ability to focus on life outside of the Economics Department. Important creativity stems from such inputs from social interactions. I wish to particularly thank my wife, Signe, and my daughter, Thea, for their ability to cope with me during ups and downs.

Although my research focus mainly on financial credit constraints, I have learned a lot about *family* credit during my PhD. It turns out this currency (and the exchange rate, particularly) may be at least as important as financial incentives. Signe has, on several occasions, granted additional credit – despite I was already close to (if not on) the limit. Without Signe's willingness to put up with me and her ability to provide a normal everyday life for Thea while I have been abroad for numerous conferences, my PhD would not have been the same. This dissertation is dedicated to her.

Thomas H. Jørgensen

Summary

This dissertation is comprised of four self-contained chapters concerned with estimation of life cycle models of consumption and retirement choices within families. The first two chapters are methodological while the two subsequent chapters are applications of the methods, techniques and challenges studied in the first two chapters.

Chapter 1, *“Structural Estimation of Continuous Choice Models: Evaluating the EGM and MPEC”*, investigates two recently proposed methods to solve and estimate stochastic dynamic programming problems: Mathematical Programming with Equilibrium Constraints (MPEC) and the Endogenous Grid Method (EGM). This chapter is purely methodological and provides evidence that the numerical solution method applied in the three following chapters (the EGM) is computationally efficient and accurate.

Chapter 2, *“Euler Equation Estimation: Children and Credit Constraints”*, shows that conventional estimators based on the consumption Euler equation produce biased estimates of the effect of children on consumption if potentially binding credit constraints are ignored. These estimators are intensively used in studies of intertemporal consumption behavior and the effect of demographics on consumption. I also show how these estimators can be used to estimate upper and lower bounds on the effect of children on consumption. Bounds estimated from the Panel Study of Income Dynamics (PSID) suggest that children affect consumption less than reported in existing studies.

Chapter 3, *“Life-Cycle Consumption and Children”*, estimates the effect of children on consumption for Danish and US households in the PSID. The findings in chapter two showed that “standard” estimators produce unreliable estimates of the effect of children on consumption if households face sufficiently strong precautionary motives. To overcome this challenge, I successively numerically solve the underlying dynamic model for all trial values of parameters. This approach has the benefit that several competing life cycle motives for consumption and saving, such as children, income uncertainty, credit constraints and retirement can be implemented simultaneously. The results suggest, as chapter two, that children do not affect non-durable consumption as much as previously assumed.

Chapter 4, *“Leisure Complementarities in Retirement”*, is concerned with later life-cycle choices of optimal saving and retirement of couples. Specifically, in this chapter, I investigate how husband and wives value joint leisure in retirement. To disentangle household level shocks that could drive the husband and wife to retire within close proximity from leisure complementarities, I formulate, solve and estimate using Danish register data a dynamic structural model of consumption and retirement in dual earner families. I find that leisure is valued twice as much if the spouse is also retired. Ignoring leisure complementarities in policy evaluations may lead to significant biased estimates of the effect on government surplus.

Resumé (in Danish)

Denne afhandling består af fire enkeltstående kapitler. Alle kapitler omhandler estimation of livscyklus-modeller for optimalt forbrug, opsparing og tilbagetrækning af par. De to første kapitler fokuserer på metoder til estimation af denne type modeller mens de to efterfølgende kapitler anvender disse metoder.

Kapitel 1, "*Structural Estimation of Continuous Choice Models: Evaluating the EGM and MPEC*", undersøger to nyligt foreslåede metoder til estimation af dynamiske modeller: Mathematical Programming with Equilibrium Constraints (MPEC) og Endogenous Grid Method (EGM). Dette kapitel fokuserer udelukkende på disse metoder og viser, at den metode (EGM) som benyttes i de følgende kapitler er beregningsmæssig hurtig og præcis.

Kapitel 2, "*Euler Equation Estimation: Children and Credit Constraints*", viser hvordan konventionelle estimatorer, baseret på Euler ligningen for forbrug, af effekten af børn på forbrug, ikke er middelrette, hvis husholdningerne er potentielt kredit-begrænsede. Disse estimatorer benyttes ofte i analyser af forbrugeradfærd over tid. I dette kapitel viser jeg desuden, hvordan disse estimatorer kan benyttes til at estimere en nedre og en øvre grænse for, hvor meget børn påvirker forbruget. Jeg estimerer disse øvre og nedre grænser for familier i det amerikanske Panel Study of Income Dynamics (PSID) og finder en effekt af børn på forbruget der er mindre end hidtil antaget.

Kapitel 3, "*Life-Cycle Consumption and Children*", estimerer effekten af børn på forbruget for danske og amerikanske familier. Resultaterne i kapitel 2 viste at standard estimatorer ikke kan anvendes til at besvare dette spørgsmål. For at overkomme disse udfordringer løser jeg den underlæggende økonomiske model for alle gæt af parametre og finder således de parametre, der passer bedst på de observerede data (strukturel estimation). Denne metode har den fordel, at mange alternative motiver for forbrug og opsparing kan inkluderes i modellen samtidigt. Specifikt så tillader jeg at husholdningers forbrug påvirkes af børn, indkomst usikkerhed, overlevelsesusikkerhed og kredit begrænsninger. Resultaterne antyder, som kapitel 2, at effekten af børn på forbruget er mindre end tidligere antaget.

Kapitel 4, "*Leisure Complementarities in Retirement*", beskæftiger sig med familiers adfærd senere i livet. Helt specifikt undersøger jeg i dette kapitel, hvordan danske par værdisætter fritid i fællesskab med deres ægtefæller. For at kunne adskille præferencer for fælles fritid fra fælles stød til ægtefæller, tillader jeg at stød til indkomst og helbred kan være korreleret mellem ægtefæller. Jeg finder at fritid sammen med en ægtefælle fordobler værdien af fritid. Ved at simulere en kunstig stigning af efterlønsalderen og folkepensionsalderen, finder jeg at hvis man udelader præferencerne for fritid, så underestimeres det offentlige overskud fra sådan en reform signifikant.

Contents

1	Structural Estimation of Continuous Choice Models: Evaluating the EGM and MPEC	1
2	Euler Equation Estimation: Children and Credit Constraints	11
3	Life-Cycle Consumption and Children	52
4	Leisure Complementarities in Retirement	93

Chapter 1

Structural Estimation of Continuous Choice Models: Evaluating the EGM and MPEC

This chapter, with only minor differences, is published as "Structural Estimation of Continuous Choice Models: Evaluating the EGM and MPEC," *Economics Letters*, 119(3), 287–290, 2013.

Structural Estimation of Continuous Choice Models: Evaluating the EGM and MPEC

Thomas H. Jørgensen[†]

October 30, 2014

Published as "Structural Estimation of Continuous Choice Models: Evaluating the EGM and MPEC," *Economics Letters*, 119(3), 287–290, 2013

Abstract

In this paper, I evaluate the performance of two recently proposed approaches to solving and estimating structural models: The Endogenous Grid Method (EGM) and Mathematical Programming with Equilibrium Constraints (MPEC). Monte Carlo simulations confirm that both the EGM and MPEC have advantages relative to standard methods. The EGM proved particularly robust, fast and straight forward to implement. Approaches trying to avoid solving the model numerically, therefore, seem to be dominated by these approaches (JEL: C61).

Keywords: Structural Estimation, Continuous Choice, Endogenous Grid Method (EGM), Mathematical Programming with Equilibrium Constraints (MPEC).

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

I wish to thank Kenneth Judd and participants at the Initiative for Computational Economics, especially Vera Molitor, Cormac O’Dea, and Kathrin Schlafmann. I have also recieved valuable comments from Anders Munk-Nielsen, Bo Honoré, and Bertel Schjerning.

1 Introduction

One of the novelties of structural models is the ability to perform counterfactual policy analysis. This requires – besides a realistic model – that researchers uncover the underlying structural parameters. Most existing approaches are notoriously slow and it is, therefore, tempting to calibrate parameters.

The Endogenous Grid Method (EGM) proposed by [Carroll \(2006\)](#) and Mathematical Programming with Equilibrium Constraints (MPEC) proposed by [Su and Judd \(2012\)](#) apply fundamentally different approaches aimed at overcoming the time consuming task of estimating structural models by, e.g., Time Iterations (TI). The EGM does this by a small but efficient modification of TI while MPEC abandon the “nested fixed-point” estimation structure, NFXP, which most other approaches follow.

The aim of this paper is to discuss a concrete implementation of these two recently proposed methods and supply new Monte Carlo evidence on performance in terms of speed, accuracy and practical implementation when estimating structural continuous choice models.¹ Hopefully, this will inspire estimation of more realistic models in terms of heterogeneity and uncertainty.

The paper proceeds as follows. Section 2 presents the model used in the analysis. Section 3 briefly discuss the estimation procedures, TI, the EGM and MPEC. Section 4 discuss data generation and present Monte Carlo results. Finally, Section 5 discuss and concludes the analysis.

2 The Model and DGP

I use the canonical model of [Deaton \(1991\)](#) where agents solve the infinite horizon problem

$$\begin{aligned} \max_{\{c\}_{t=0}^{\infty}} \quad & \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t u(c_t) \right], \\ \text{s.t.} \quad & a_{t+1} = R(a_t + y_t - c_t), \\ & a_t \geq 0 \forall t, \end{aligned}$$

where $0 < \beta < 1$ is the discount factor, R is the real gross interest rate, c_t is consumption in period t , a_t is assets at the beginning of period t , and $y_t \sim \mathcal{N}(\mu_y, \sigma_y^2)$ is stochastic income in beginning of period t . More complicated models could be formulated without changing the results. Preferences are assumed to be CRRA with relative

¹[Su and Judd \(2012\)](#) illustrate the applicability of MPEC to discrete choice models, using the bus-replacement model of [Rust \(1987\)](#) but do not consider explicitly continuous choice models.

risk aversion, ρ ,

$$u(c_t) = \frac{c_t^{1-\rho}}{1-\rho}.$$

It is convenient to formulate the state in this model as total cash-on-hand available in the beginning of period t as m_t , such that the state in the model evolves as

$$m_{t+1} = R(m_t - c_t) + y_{t+1}. \quad (1)$$

3 Estimation Approaches Considered

In this section, I provide a brief introduction to the implemented approaches. The first two, TI and the EGM, are based on the nested fixed point (NFXP) approach, in which the model is solved in an inner algorithm for a given set of trial values of parameters. An outer optimization algorithm estimates the structural parameters by varying these, leading to successively solving the structural model. The third approach, MPEC, abandons NFXP and formulates the solution of the model as equilibrium constraints when estimating the structural parameters.

The estimation framework adopted here is Maximum Likelihood. Without changing the results, a method of moments framework could be adopted where moments from the data are matched moments predicted from the model. It is assumed that panel data on consumption are observed with measurement error, such that

$$c_{it}^{data} = c(m_{it}^{data} | \rho) + \varepsilon_{it},$$

where $c(\cdot | \rho)$ is the consumption function predicted by the model and the measurement error is assumed *iid* Gaussian with mean zero and variance σ^2 .² The (mean) log likelihood function can be written as

$$\mathcal{L}(\rho; c, c^{data}, m^{data}) = -\log(\sigma) - \sum_i \frac{1}{NT_i} \sum_t \frac{1}{2\sigma^2} \left(c_{it}^{data} - c(m_{it}^{data} | \rho) \right)^2. \quad (2)$$

Since the consumption function in the present model has no closed form solution, $c(m|\rho)$ is found numerically. TI and the EGM find $c(m|\rho)$ for a given ρ and use that solution to evaluate the likelihood function. MPEC estimate $c(m|\rho)$ and ρ jointly. The solutions from each of the methods are indistinguishable, as shown in Figure 1.

I use $Q = 8$ Gauss-Hermite nodes (y^q) and weights (w^q) to approximate expectations with regard to labor market income, y . Consumption is approximated by 200 unequally spaced grid points over m_t , with more mass at the bottom of the distribu-

²Alternatively, the estimation could be framed as measurement error in the difference in log consumption or assets without changing the results.




Figure 1 – The Consumption Function, $c(m|\rho)$, from TI, the EGM and MPEC.

tion. In the EGM, the grid for m_t is determined endogenously, as discussed below. Linear interpolation is applied between grid points.

All approaches are implemented in MATLAB 2012b using the KNITRO solver for optimization (see [Byrd, Nocedal and Waltz, 2006](#)) on a laptop with Intel® Core™ i5-2520M CPU @ 2.50 GHz and 4GB RAM. Code are available from authors webpage.

3.1 Time Iterations (TI)

The Euler residual from the present model is as a nonlinear equation in consumption, c_t ,

$$\begin{aligned} \mathcal{E}(c_t|m_t) &\equiv R\beta\mathbb{E}[u_c(c_{t+1})|m_t] - u_c(c_t), \\ &\doteq R\beta \sum_{q=1}^Q w^q \check{c}_{t+1}(\underbrace{R(m_t - c_t) + y^q}_{m_{t+1}})^{-\rho} - c_t^{-\rho}, \end{aligned} \quad (3)$$

where $\check{c}_{t+1}(m_{t+1})$ represents a linear interpolation function. A numerical procedure, such as bisection or Newton iterations, is used to find optimal consumption that puts the residual in (3) to zero,

$$\begin{aligned} c_t^*(m_t) &: \mathcal{E}(c_t^*|m_t) = 0, \\ \text{s.t.} \quad &c_t \leq m_t. \end{aligned}$$

In order to find the stationary solution to the infinite horizon model, iterate over time until $\max_m \{|c_t^*(m) - c_{t+1}^*(m)|\} < 1.0E^{-7}$.

3.2 The Endogenous Grid Method (EGM)

The EGM proposed by [Carroll \(2006\)](#) modifies time iteration by defining the interpolation grid over end-of-period assets, a_t , instead of beginning-of-period cash-on-hand, m_t . This trick facilitates an analytical solution to optimal consumption today by inverting the Euler equation,

$$\begin{aligned} c_t^*(m_t) &= u_c^{-1}(R\beta\mathbb{E}[u_c(c_{t+1})|m_t]), \\ &\doteq \left(R\beta \sum_{q=1}^Q w^q \check{c}_{t+1} (Ra_t + y^q)^{-\rho} \right)^{-\frac{1}{\rho}}, \end{aligned} \quad (4)$$

where the rhs now is independent of c_t . Since no numerical methods are needed to find optimal consumption (contrary to time iteration), the method dramatically increases speed. Finding the stationary solution is done as for time iterations above.

Cash-on-hand today, m_t , consistent with end-of-period assets, a_t , and consumption, c_t^* , is determined endogenously as

$$m_t = c_t^*(m_t) + a_t.$$

The EGM perfectly tracks the credit constraint. This is because the lowest point in the grid over a_t , $\underline{a} = 1.0\text{E}^{-6}$, is (very close to) the point where agents are on the cusp of being credit constrained. This is illustrated in the right panel of Figure 1. Including the interpolation point $(m, c) = (0, 0)$ ensures the credit constrained level of cash-on-hand is handled correctly.

3.3 Mathematical Programming with Equilibrium Constraints (MPEC)

[Su and Judd \(2012\)](#) propose formulating the solution and estimation problem as a joint constrained maximization problem. The intuition is that NFXP spent most of the time solving models with high accuracy for “wrong” parameters. The behavior only needs to be optimal at the true parameters. Formalized as a nonlinear constrained optimization problem,

$$\begin{aligned} \max_{c, \rho} \quad & \mathcal{L}(\rho; c, c^{data}, m^{data}) \\ \text{s.t.} \quad & \end{aligned}$$

$$1 < \rho, \quad (5)$$

$$0 \leq c \leq m - c, \quad (6)$$

$$0 \geq \beta R \mathbb{E} [u'(c(R(m - c) + y))] - u'(c), \quad (7)$$

$$0 = (m - c) (\beta R \mathbb{E} [u'(c(\cdot))] - u'(c)), \quad (8)$$

where $\mathcal{L}(\cdot)$ is the likelihood function in (2), (5) is a lower bound on the risk aversion parameter, (6) are lower and upper bounds on the consumption parameters, (7) is the Euler residual formulated as a nonlinear inequality constraint, and (8) is a complementarity constraint, stating that if the credit constraint is not binding, the Euler equation must hold.

The consumption function is estimated along with the structural parameters. Hence, the number of parameters is the number of grid points used to approximate consumption in addition to the structural parameters. Here, that amounts to 201 parameters.

Convergence problems due to loose inner-loop stopping criteria are avoided completely. Inner loop iterations are simply not performed in MPEC. In practice, however, supplying good starting values for consumption parameters was necessary to obtain convergence to the right optimum.

4 Monte Carlo Comparison

To assess the performance of the approaches described in Section 3, synthetic data (5000 individuals in 10 time periods) are generated for value of $\beta \in \{.70, .95, .99\}$. To mitigate the influence of stochastic draws, I perform 50 Monte Carlo runs for each β .

Table 1 reports the Root Mean Squared Error (RMSE) and Monte Carlo Standard Error (MCSTD) along with average time used, the standard deviation of time use across MC runs, and the number of iterations used by each method. Iterations at level 1 refers to the outermost optimization, level 2 refers to iterations until convergence to the infinite horizon stationary solution, and level 3 refers to the innermost numerical procedure, finding the optimal consumption. The three methods differ in the levels of iteration. The EGM circumvents the inner most procedure while MPEC only operates on the outer level. All approaches are initialized using the same starting value for ρ .

As expected, TI is slowest overall and both TI and the EGM (which both rely on NFXP) is slowed by higher values of β . The EGM does, however, seem to be less sensitive to β relative to TI. MPEC should be roughly invariant to the level of the discount rate and the variation across β -values reflect the difficulties in supplying good starting values for consumption parameters rather than the effect of changing β . This instability is also reflected in the relatively large dispersion in time to convergence across MC runs (column 4) for MPEC. The large RMSE of 0.049 when $\beta = 0.7$ stems from MPEC not converging to right optimum in five of the MC runs.

The EGM and TI use the same number of level 1 and 2 iterations. The great speed gain from the EGM is clearly stemming from the elimination of the inner most searches for optimal consumption (level 3), that TI suffers from. MPEC use significantly more level 1 iterations due to the fact that 201 parameters are estimated in MPEC. Since

Table 1 – Monte Carlo Comparison.

β		RMSE	MCSTD	Time (secs)	Std. time	Iterations		
						level 1	level 2	level 3
.70	TI	0.002	0.002	26.0	0.48	5	142	147,900
	EGM	0.002	0.002	0.1	0.03	5	174	–
	MPEC	0.049	0.046	112.4	269.97	123	–	–
.95	TI	0.009	0.006	650.7	6.80	5	3,473	3,621,124
	EGM	0.006	0.006	1.9	0.05	5	3,636	–
	MPEC	0.009	0.006	93.7	37.00	94	–	–
.99	TI	0.000	0.000	1,682.6	15.74	6	9,215	8,475,336
	EGM	0.000	0.000	5.0	0.08	6	9,247	–
	MPEC	0.000	0.000	30.9	6.26	23	–	–

Notes: Based on 50 MC runs with $N \cdot T = 5000 \cdot 10$ simulated observations each run. Columns 3, 5, 6 and 7 are Monte Carlo averages. only ρ is estimated. $R = 1.05$, $\mu_y = 10$, $\sigma_y^2 = 100$ and 200 grid points are used to approximate consumption.

MPEC only operates on the outer level, the approach is considerably faster than TI.

The EGM outperforms MPEC on both speed and RMSE. The EGM is able to uncover the structural parameter in on average five seconds while MPEC uses around 30 seconds and TI uses 30 minutes to complete the same task. Due to the EGMs relatively straight forward reformulation of time iterations, this result is very encouraging.

5 Discussion

Through this analysis, two recent proposed approaches to structural estimation, the EGM and MPEC, have been evaluated. The theoretically appealing constraint optimization approach, MPEC, proved to be somewhat disappointing. Even if researchers apply state of the art solvers to problems supplied with (correct) gradients, hessian and sparsity pattern, the size limitation on the solvable problems is a significant constraint. Problems that are not sparse with large state space dimensions would require an intimidating amount of memory. This limitation is also recognized by [Su and Judd \(2012, p. 2215\)](#).

The size limitations of MPEC effectively rules out (realistic) finite horizon models since the number of parameters and constraints are the number of time periods multiplied the number of grid points in addition to the structural parameters, $T \cdot n + k$. Furthermore, using simulation based estimation methods, such as indirect inference or simulated method of moments are generally not feasible in the MPEC framework. A small perturbation in a consumption parameter requires (costly) re-simulation of

synthetic data.

The EGM proved very robust and fast. The small change to time iteration is very straight forward to implement. Furthermore, the EGM includes the exact point where agents are on the curb of being liquidity constrained, increasing accuracy. The EGM (as well as TI and MPEC) can also handle continuous-discrete choice models, see, e.g., [Fella \(2014\)](#) and [Iskhakov, Jørgensen, Rust and Schjerning \(2014\)](#) who generalize the EGM to handle discrete choices.

The fact that structural parameters can be estimated in a fraction of the time conventional methods require has widespread implications. Heterogeneous parameters and correlated uncertainty could be some of many “new” improvements in structural models. These features have often not been feasible to implement in structural estimation. This also means that several approaches trying to avoid solving the model numerically, such as non-linear GMM estimation ([Alan, Attanasio and Browning, 2009](#)) or Synthetic Residual Estimation ([Alan and Browning, 2010](#)), are dominated by the EGM and MPEC.

References

- 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.
- ALAN, S. AND M. BROWNING (2010): "Estimating Intertemporal Allocation Parameters using Synthetic Residual Estimation," *The Review of Economic Studies*, 77(4), 1231–1261.
- BYRD, R. H., J. NOCEDAL AND R. A. WALTZ (2006): "KNITRO: An Integrated Package for Nonlinear Optimization," in *Large-Scale Nonlinear Optimization*, ed. by G. di Pillo and M. Roma, pp. 35–59. Springer: New York.
- CARROLL, C. D. (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.
- FELLA, G. (2014): "A generalized endogenous grid method for non-smooth and non-concave problems," *Review of Economic Dynamics*, 17(2), 329–344.
- ISKHAKOV, F., T. H. JØRGENSEN, J. RUST AND B. SCHJERNING (2014): "Estimating Discrete-Continuous Choice Models: Endogenous Grid Method with Taste Shocks," Unpublished working paper.
- RUST, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55(5), 999–1033.
- SU, C.-L. AND K. L. JUDD (2012): "Constrained Optimization Approaches to Estimation of Structural Models," *Econometrica*, 80(5), 2213–2230.

Chapter 2

Euler Equation Estimation: Children and Credit Constraints

Euler Equation Estimation: Children and Credit Constraints

Thomas H. Jørgensen[†]

October 30, 2014

Abstract

I show that conventional estimators based on the consumption Euler equation, intensively used in studies of intertemporal consumption behavior, produce biased estimates of the effect of children on consumption if potentially binding credit constraints are ignored. As a more constructive contribution, I supply a tractable approach to obtaining bounds on the effect of children and estimate these bounds using the Panel Study of Income Dynamics (PSID). Results suggest that children might not affect household consumption in the same magnitude previously assumed (JEL: D12, D14, D91).

Keywords: Consumption, Euler Equation Estimation, Credit Constraints, Children, Demographics, Life Cycle, Bounds.

[†]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.

I am grateful to Mette Ejrnæs, Bertel Schjerning and Jeppe Druedahl for fruitful discussions. This paper also benefited from discussions with Martin Browning, Christopher Carroll, John Rust, Richard Blundell, Søren Leth-Petersen, Miles Kimball, Dennis Kristensen, Daniel Borowczyk-Martins, Kjetil Storesletten, Jérôme Adda, Anders Munk-Nielsen, Mette Foged, Damoun Ashournia 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. All errors are my own.

This paper extends the first part of the overly long working paper: Jørgensen, T. H. (2014): "Life-Cycle Consumption and Children," CAM Working Paper No. 2014-02, Department of Economics, University of Copenhagen.

1 Introduction

This study investigates what can be learned from Euler equation estimation of the effect of children on household consumption when households are potentially credit constrained. Although these estimators are now work horses in the analysis of intertemporal consumption behavior, little is known about their performance when households face potentially binding credit constraints, invalidating the standard Euler equation. Particularly, the effect of demographics on consumption and the extent to which children affect consumption behavior have received great attention the last two decades. Through numerous Euler equation estimations, a consensus has been reached in the literature that children are important drivers of consumption over the life cycle.¹

The present study offers three contributions to this literature. First, I show *how* conventional Euler equation estimation methods produce biased estimates of the effect of children on consumption if consumers face possibly binding credit constraints. This has not been subject to a thorough analysis and the volume of work in the field of intertemporal consumption behavior merits one.²

Secondly, I supply a tractable approach to obtaining bounds on the effect of children on consumption that allows households to be affected by constraints. Specifically, if the effect of children on consumption is *large*, the credit constraint likely restrains households from increasing consumption as much as desired had (additional) borrowing been possible, producing a downwards bias. To the contrary, if the effect of children is relatively *low*, conventional methods will overestimate the effect of children. Even if children does not affect consumption, the inability to borrow against future income growth produce a positive correlation between consumption growth and changes in household demographics because children often arrive while households are young and affected by credit constraints the most.

I propose to split the sample into young households, in which children might arrive, and older households, in which children might move. Comparing older households with and without children produce a lower bound for the reason discussed above. Using the cohort average number of children as instrument produces an upper bound due to the positive correlation between the growth in the average number of children and income growth of young households.

Finally, I find that the effects of children reported in the existing literature are

¹Irvine (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 are due to Browning, Deaton and Irish (1985); Blundell, Browning and Meghir (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).

²The fact that ignoring credit constraints produce biased Euler equation estimates is not new. Adda and Cooper (2003) show how Euler equation estimation of the intertemporal elasticity of substitution is overestimated if credit constraints are ignored.

above the proposed upper bound estimated from the Panel Study of Income Dynamics (PSID). In contrast to what I find, it seems broadly accepted that children play an important role in generating the observed consumption profiles. In an influential study by [Attanasio, Banks, Meghir and Weber \(1999\)](#), the number of children is found to be important in order to describe the consumption behavior of US consumers, using the Consumer Expenditure Survey (CEX). This is supported by the results in [Attanasio and Browning \(1995\)](#) using the UK Family Expenditure Survey (FES). [Browning and Ejrnaes \(2009\)](#) find that the number and age of children can explain completely the hump in consumption in the FES. However, all existing studies apply Euler equation estimation techniques ignoring potentially binding credit constraints. As I show, if the effect of children is relatively low the applied estimators overestimates the effect of children on consumption if households face potentially binding credit constraints.

The present study is related to a recent strand of literature investigating the validity of Euler equation estimation. For example, [Ludvigson and Paxson \(2001\)](#) and [Carroll \(2001\)](#) argue that using a log-linearized Euler equation for estimation of the intertemporal elasticity of substitution (IES) suffers from an omitted variable bias if consumers face sufficient income uncertainty. [Attanasio and Low \(2004\)](#) find, however, that the critique is unwarranted. Recently, [Alan, Atalay and Crossley \(2012\)](#) investigate how measurement error affects Euler equation estimation results and unify the seemingly contradictory results of [Ludvigson and Paxson \(2001\)](#) and [Carroll \(2001\)](#) with those in [Attanasio and Low \(2004\)](#). They argue that the contradictory results are due to differences in the time series dimension in the implemented Monte Carlo studies. The bias in Euler equation estimators of the IES might be small when interest rates vary sufficiently over time and the time dimension is (unrealistically) long, as in [Attanasio and Low \(2004\)](#). All these studies focus on the IES and, contrary to the present study, ignore potentially binding credit constraints.

A growing empirical literature finds evidence consistent with credit constraints being important for observed behavior. Interpreting the “excess sensitivity” in consumption growth to income as evidence of credit constraints, [Hall and Mishkin \(1982\)](#) estimate that around 20 percent of households in the PSID are credit constrained. If the excess sensitivity is due to credit constraints, households with high wealth levels should display significantly less effect of lagged income on consumption growth compared to low wealth households. This is what [Zeldes \(1989a\)](#) finds while [Runkle \(1991\)](#) does not find significant differences.

A second strand of literature exploits random variation to identify the importance of credit constraints. Using the random receipt timing of the 2001 federal income tax rebate in the US, [Johnson, Parker and Souleles \(2006\)](#) find that consumption in the CEX responds to the transitory income increase generated by the rebate. Including also the 2008 tax rebate, [Gross, Notowidigdo and Wang \(2014\)](#) show that low wealth

and low income households used their tax rebates to file for bankruptcy. Both results are consistent with an important role for credit constraints. [Gross and Souleles \(2002\)](#) analyze how changes in credit card debt limits increase debt holdings. They find that debt increases with 13 percent of the change in the debt limit. Using Danish register data, [Leth-Petersen \(2010\)](#) estimates the effect of an unanticipated reform in 1992 that allowed Danish house owners to use their house as collateral to take up consumption loans. He estimates that around 12 percent of Danish households, many of which were young households, were affected by credit constraints.

A third strand of literature identifies credit constrained households from direct survey measures on credit availability.³ Using information on whether a request for credit had been declined [Jappelli \(1990\)](#) estimates that around 19 percent of households in the Survey of Consumer Finances (SCF) are credit constrained. [Jappelli, Pischke and Souleles \(1998\)](#) extrapolate the likelihood of being credit constrained in the SCF to the PSID. Based on observable characteristics in both the SCF and the PSID, they find that the excess sensitivity of (food) consumption to lagged income of households who are more likely to be constrained is three times that of households who are less likely to be credit constrained. Recently, [Crossley and Low \(forthcoming\)](#) find that around 6-14 percent of job losers in the Canadian Out of Employment Panel (COEP) survey are credit constrained, depending on how households are classified as being constrained.

The present results generalize to cases in which consumers do not face explicit credit constraints. If there instead is a probability of receiving a zero-income shock (as in [Carroll, 1997](#) and [Gourinchas and Parker, 2002](#)), most results still hold. This is because risk averse consumers will instead face a “self-imposed” no-borrowing constraint stemming from the fear of receiving zero income in all future periods with consumption of zero as a consequence ([Schechtman, 1976](#); [Zeldes, 1989b](#)). In turn, consumption will respond substantially to transitory income shocks and the log-linearized Euler equation will be a poor approximation.⁴

The rest of the paper proceeds as follows. The following section presents the constrained consumption Euler equation and discusses the most commonly applied estimators derived from it when ignoring credit constraints. Section 3 illustrates how these estimators fail to uncover the effect of children on consumption when households face potentially binding credit constraints and suggests how bounds can be estimated using these methods. Section 4 shows that existing estimates of the effect of children on consumption are above the proposed upper bounds estimated using the PSID. Section 5 discusses the robustness of the bounds and section 6 concludes.

³[Thurow \(1969\)](#) suggested that borrowing constraints could explain observed consumption profiles.

⁴This is the point of [Carroll \(2001\)](#) where he illustrates how this poor first (and second) order approximation of the non-linear Euler equation results in poor estimates of the intertemporal elasticity of substitution. His result shows that the result in [Adda and Cooper \(2003\)](#) using an explicit no-borrowing constraint generalizes to cases with a self-induced constraint.

2 Euler Equation Estimation of Demographic Effects

Consider a life cycle model where consumers have time-separable utility over (a single) consumption good and are restricted in how much negative wealth they can accumulate. As most of the existing literature, I follow [Attanasio, Banks, Meghir and Weber \(1999\)](#) and let children affect the *marginal value* of consumption through a multiplicative taste shifter, $v(\mathbf{z}_t; \theta)$, in which \mathbf{z}_t contains variables describing household demographics and θ is their loadings. As is standard in the literature, I let $v(\mathbf{z}_t; \theta) = \exp(\theta' \mathbf{z}_t)$ throughout. Alternatively, the household composition could be included as a scaling of resources and consumption (equivalence scaling), as done in, e.g., [Fernández-Villaverde and Krueger \(2007\)](#).⁵

The *constrained* Euler equation is

$$\begin{aligned} u'(C_t)v(\mathbf{z}_t; \theta) - \lambda_t &= R\beta\mathbb{E}_t [u'(C_{t+1})v(\mathbf{z}_{t+1}; \theta) - \lambda_{t+1}] \\ &\Downarrow \\ R\beta \frac{u'(C_{t+1})v(\mathbf{z}_{t+1}; \theta)}{u'(C_t)v(\mathbf{z}_t; \theta)} &= \underbrace{\epsilon_{1,t+1} + \epsilon_{2,t+1}}_{\equiv \epsilon_{t+1}} \end{aligned} \quad (1)$$

where $\mathbb{E}_t[\cdot]$ denotes expectations conditional on information available in period t , λ_t is the shadow price of resources in period t , R is the real gross interest rate, β is the discount factor, C_t denotes consumption, $u(C_t) = C_t^{1-\rho}/(1-\rho)$ is the utility function, assumed to be constant relative risk aversion (CRRA) where ρ is the inverse of the IES. The structural Euler error, ϵ_{t+1} , satisfies

$$\begin{aligned} \mathbb{E}_t[\epsilon_{1,t+1}] &= 1, \\ \mathbb{E}_t[\epsilon_{2,t+1}] &= -\frac{\lambda_t - R\beta\mathbb{E}_t[\lambda_{t+1}]}{u'(C_t)v(\mathbf{z}_t; \theta)}. \end{aligned}$$

From the Kuhn-Tucker conditions we know that $\lambda_t \geq 0$ in all time periods. Hence, the mean expectational errors in (1) equals one only if consumers are not constrained in the current period and know with perfect certainty that the borrowing constraint will *not* bind in the future. In such a case, $\mathbb{E}_t[\epsilon_{2,t+1}] = 0 \forall t$. Generally, however, consumers are *not* certain that they will be unaffected by constraints and the expectational error in (1) is a function of information today,

$$\mathbb{E}_t[\epsilon_{t+1}] = f(C_t, \mathbf{z}_t) \neq 1,$$

and serially correlated through the presence of λ_t and λ_{t+1} in $\epsilon_{2,t+1}$.

⁵See [Bick and Choi \(2013\)](#) for an analysis of different approaches to and implied behavior from inclusion of household demographics in life cycle models. Alternative parametrizations would require reformulating the estimable equations accordingly.

In the existing literature on intertemporal consumption allocation and the effect of children on consumption, credit constraints are often ignored or assumed away. It is clear from (1), that estimators ignoring credit constraints suffer from something similar to an “omitted variable bias”. Below, I discuss the two most common estimators.

2.1 Conventional Euler Equation Estimators: Ignoring Constraints

Consider having longitudinal information on consumption and demographics for households $i = 1, \dots, N$ in time periods $t = 1, \dots, T$. Ignoring potentially binding credit constraints (i.e., imposing $\lambda_s = 0 \forall s$) and inserting the standard functional form assumptions mentioned above, a non-linear GMM estimator of θ could be

$$\theta_{GMM} = \underset{\theta}{\operatorname{argmin}} \left[\frac{1}{NT} \sum_i^N \sum_t^T \left(R\beta \left(\frac{C_{i,t+1}}{C_{i,t}} \right)^{-\rho} \exp(\theta \Delta \mathbf{z}_{i,t+1}) - 1 \right) \cdot Z_{i,t+1} \right]^2, \quad (2)$$

such that θ_{GMM} is the parameter that satisfy the sample equivalent of $\mathbb{E}[(\epsilon - 1)'Z] = 0$, where Z contain instrument(s) assumed uncorrelated with the Euler residual. Ignoring measurement error, the estimator in (2) produce consistent estimates if a suitable instrument is available and, importantly, households do not face credit constraints.⁶

Using food consumption from the PSID, [Alan, Attanasio and Browning \(2009\)](#) estimate the effect of children to be around .18 from a similar estimator as (2) and as large as .9 using estimators allowing for measurement error in consumption.

Most existing studies work with a log-linearized version of the Euler equation since it yields estimable equations linear in parameters which can easily be estimated with synthetic cohort panels ([Browning, Deaton and Irish, 1985](#)) and handle measurement error through instrumental variables estimation. The log-linearized Euler equation is

$$\Delta \log C_{it} = \text{constant} + \rho^{-1} \theta' \Delta \mathbf{z}_{it} + \tilde{\epsilon}_{it}, \quad (3)$$

where the first term is a constant as a function of structural parameters (β, ρ) and the interest rate (assumed constant throughout), the second term is the effect of children (times the IES) and the last term is a reduced form residual, $\tilde{\epsilon}_t = -\rho^{-1} \log \epsilon_t$.

In the influential study by [Attanasio, Banks, Meghir and Weber \(1999\)](#), θ and ρ is estimated from the CEX by a log-linearized Euler equation using lagged changes in \mathbf{z}_t as instruments along with lagged changes in income and consumption. The effect of the number of children is found to be around $\theta \approx .33$. Several studies use food consumption from the PSID to estimate versions of the log-linearized Euler equation, see, e.g., [Hall and Mishkin \(1982\)](#); [Runkle \(1991\)](#) and [Lawrance \(1991\)](#). The latter reports

⁶[Alan, Attanasio and Browning \(2009\)](#) supply modified GMM estimators to allow for measurement error while still ignoring possibly binding credit constraints.

estimates suggesting a value of θ of around 0.5. [Dynan \(2000\)](#), also using the PSID, estimates the effects of children to be around .7. [Browning and Ejrnaes \(2009\)](#) allow for a more flexible functional form of $v(\mathbf{z}_t; \theta)$ when estimating the effect of children consumption using the FES and find that the number and age of children can explain completely the hump in consumption.

Other estimators have been proposed to estimate Euler equations. For example, [Alan and Browning \(2010\)](#) propose a method in which they fully parameterize the Euler residuals and simulate these residuals and consumption paths simultaneously. Their Synthetic Residual Estimation (SRE) procedure does not allow for credit constraints in a coherent way. Since the GMM and log-linearized estimation methods are the conventional methods used in the literature, I focus exclusively on these.

Some empirical studies of intertemporal consumption behavior do recognize that credit constraints might affect household behavior. Potentially binding credit constraints are often handled by discarding households in which nothing is carried over from period t to $t + 1$ (see, e.g., [Alan, Attanasio and Browning, 2009](#)). This strategy is clearly not a satisfactory approach because expectations about the credit constraint potentially binding in future periods still affect present consumption behavior through $\mathbb{E}_t[\lambda_{t+1}]$. Determining at which level of wealth households are completely free of the credit constraint is not trivial.

3 Bias and Bounds from Euler Equation Estimation

In this section, I illustrate how conventional Euler equation estimators, (2) and (3), produce biased estimates of the effect of children on consumption and can be used to construct bounds of this parameter. I first formulate a four-period model from which I can derive analytical expressions for the log-linearized Euler equation estimator and show how bounds can be calculated from splitting the sample into young and older households. To confirm the results from the four-period model, I formulate and numerically solve a standard life cycle model of buffer-stock savings behavior. By simulating data from this model, I estimate the proposed bounds and show that they are very similar to the bounds from the four-period model.

The present exposition is based on the absolutely best of all circumstances in which *i*) a panel of consumers is available, *ii*) consumption is observed without measurement error, *iii*) researchers know the underlying model consumers solve, and *iv*) researchers know the preferences of consumers except the effect of children on consumption.

3.1 Evidence from A Four Period Model

Here, I setup a four-period model with an analytical solution to illustrate how Euler equation estimation performs when households face potentially binding credit constraints. In the initial period, $t = 0$, all households are childless. In period $t = 1$, the “young” stage, a child arrives, $z_1 = 1$, in p percent of the households and the remaining $1 - p$ percent remain childless, $z_1 = 0$. In period $t = 2$, the “old” stage, the child moves (if present in period one) such that $z_2 = 0$ for all households. Households die with certainty in the end of period $t = 3$ and consume all available resources in this terminal period.

Utility is CRRA and the taste shifter is assumed to be given by $v(\mathbf{z}_t; \theta) = \exp(\theta \mathbf{z}_t)$ with $\mathbf{z}_t \in \{0, 1\}$, and with baseline parameters of $\rho = 2$ and $\theta = 0.5$. To reduce unnecessary cluttering, the gross real interest rate and the discount factor both equal one, $R = \beta = 1$. Households receive a deterministic income of Y_t in beginning of every period. Income grows with G_1 between period zero and period one ($Y_1 = G_1 Y_0$) and is constant otherwise ($Y_t = Y_{t-1}$, $t = 2, 3$). The beginning-of-period resources available for consumption is the sum of household income and end-of-period wealth carried over from last period, $M_t = A_{t-1} + Y_t$.

Formally, households solve, for a given value of $\mathbf{z}_1 \in \{0, 1\}$, the problem,

$$\max_{C_0, C_1, C_2} \frac{C_0^{1-\rho}}{1-\rho} + \exp(\theta \mathbf{z}_1) \frac{C_1^{1-\rho}}{1-\rho} + \frac{C_2^{1-\rho}}{1-\rho} + \frac{(M_2 - C_2 + Y_3)^{1-\rho}}{1-\rho},$$

subject to a no-borrowing constraint, $A_t \geq 0$, $\forall t$. Appendix A in the online supplementary material solves the model analytically and reports the resulting optimal consumption functions.

Using the optimal consumption behavior from this model, Figure 1 presents consumption and wealth profiles for households initiated with no wealth in the initial period, $A_{-1} = 0$, period $t = 0$ income normalized to one ($Y_0 = 1$), and early life income growth of eight percent, $G_1 = 1.08$. Panel 1a presents consumption profiles for models with a credit constraint (solid) and without constraints (dashed) for households with children in period one (black) and without children (red). Panel 1b illustrates the associated wealth profiles. Potentially binding credit constraints affect the consumption and wealth profiles significantly.

Childless households increase consumption exactly as much as income grows and is in effect only potentially credit constrained in period $t = 0$ because they are unable to borrow against future income growth. Households in which a child arrives in period $t = 1$, on the other hand, might also be credit constrained in period $t = 1$ since they might want to increase consumption by more than their available resources.

The OLS estimator of the effect of children using consumption growth from $t - 1$ to

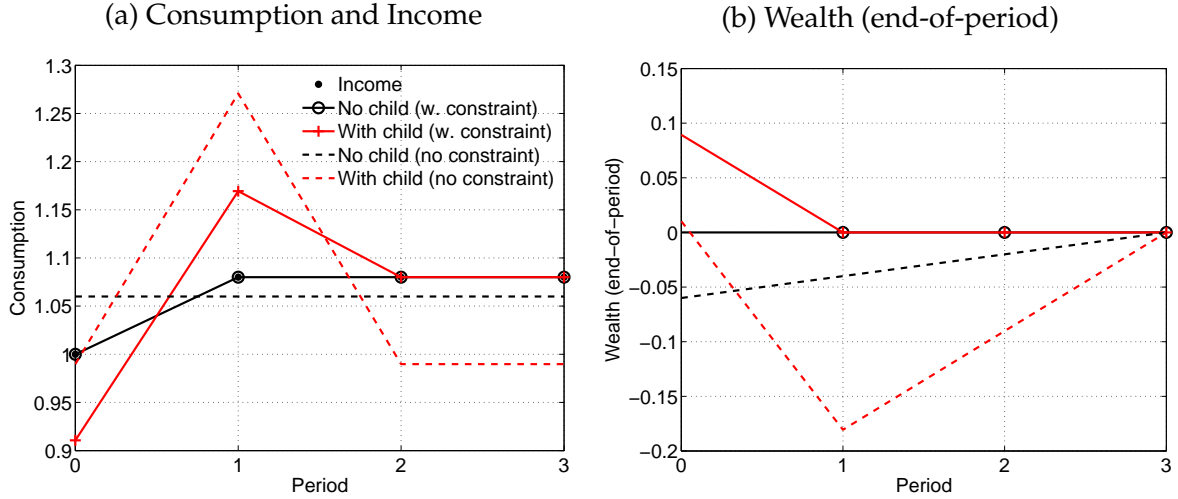


Figure 1 – Consumption and Wealth Profiles from the Four-period Model.

Notes: Figure 1 illustrates the consumption and wealth age profile from the four period model with parameters $\rho = 2$, $G_1 = 1.08$, and $\theta = 0.5$. Households are initiated with zero wealth and initial income is normalized to one, $A_{-1} = 0$ and $Y_0 = 1$, respectively. Panel a presents the income and consumption profiles for models with credit constraints (solid) and without constraints (dashed) for households with children in period one (black) and without children (red). Panel b illustrates the associated wealth profile.

t from the log-linearized Euler equation (3) is given by

$$\hat{\theta}_{OLS}^t = (\Delta \log C_t|_{z_1=1} - \Delta \log C_t|_{z_1=0})\rho.$$

Using the (cohort) average number of children as instrument ($Z = p$) should be less affected by idiosyncratic uncertainty and, thus, credit constraints. The IV estimator is⁷

$$\hat{\theta}_{IV}^t = \frac{1}{p}(p\Delta \log C_t|_{z_1=1} + (1-p)\Delta \log C_t|_{z_1=0})\rho.$$

Appendix A in the online supplementary material derives explicit formulas for each estimator when using either young households or old households to estimate

⁷Since there is only one cohort here (p is constant) no constant is included in the regression. Of course, in general, there will also be included a constant in such a regression.

the effect of children on consumption. The resulting estimators are

$$\begin{aligned}
 \hat{\theta}_{OLS}^{young} &= \begin{cases} \theta - \log(G_1)\rho & \text{if } \theta > \log(G_1)\rho, \\ 0 & \text{if } 0 \leq \theta \leq \log(G_1)\rho, \end{cases} \\
 &\leq \theta, \\
 \hat{\theta}_{IV}^{young} &= \begin{cases} \theta + (1-p)/p \log(G_1)\rho & \text{if } \theta > \log(G_1)\rho, \\ \log G_1\rho/p & \text{if } 0 \leq \theta \leq \log(G_1)\rho, \end{cases} \\
 &\geq \theta, \\
 \hat{\theta}_{OLS}^{old} &= \begin{cases} \rho \log\left(\frac{1+G_1}{G_1}\right) - \rho \log(1 + \exp(-\rho^{-1}\theta)) & \text{if } \theta > \log(G_1)\rho, \\ 0 & \text{if } 0 \leq \theta \leq \log(G_1)\rho, \end{cases} \\
 &\leq \theta, \\
 \hat{\theta}_{IV}^{old} &= \hat{\theta}_{OLS}^{old} \leq \theta.
 \end{aligned}$$

It is immediately clear from these estimators that neither OLS nor IV estimators will in general yield consistent estimates of the true θ .⁸ Interestingly, as the effect of children goes towards zero, the OLS (and IV) estimator using only older households comes close to the true value of θ . Similarly, as the effect of children on consumption gets increasingly large, the bias-part, $(1-p)/p \log(G_1)\rho$, of the young-IV estimate, $\hat{\theta}_{IV}^{young}$, becomes relatively less important. Therefore, I propose to split the sample into “young” households, in which children arrive, and “older” households, in which children leave, and use the IV estimate from the young sample as an upper bound and use the OLS estimate from the older sample as a lower bound.

The OLS estimate using young households could alternatively be used as a lower bound. However, as I show in the robustness exercise, if children arrive probabilistically, OLS from the young sample will underestimate the effect of children on consumption even if there is *no* credit constraint. I have chosen a lower bound that delivers the true effect of children on consumption when either there is no effect of children or no credit constraint, irrespectively if children arrive deterministically or probabilistically.⁹

The results are intuitive. Young households will accumulate wealth in the initial period zero, but not necessarily enough to ensure that the credit constraint is not binding in period one, in which a child arrives. Even if they do accumulate enough wealth, the fact that the childless households also increase consumption creates a downwards bias in the estimate. When children subsequently leave, households with children are likely

⁸The bias is constant, independent of the observations and does, therefore, not vanish asymptotically. Hence, this suggests that the estimators might not even be consistent. In the more realistic life cycle model studied below, no closed form expressions can be derived and arguments are made through Monte Carlo simulation of finite samples and only the bias can be illustrated.

⁹Note, as I will show in the robustness exercise, this is only correct for the non-linear GMM estimator. The log-linear estimator will *not* be able to uncover the true effect of children if there is no credit constraint but instead a probability of a low income shock.

to go from being constrained in period one to unconstrained in period two (since they prefer consumption when children are present). The resulting drop in consumption will be smaller compared to the situation without a constraint, resulting in the OLS estimator being downwards biased. The fact that income and the the average number of children are positively correlated in the early part of the life cycle produce an upwards bias in the IV estimator.

Interestingly, for low levels of θ ($0 \leq \theta \leq \log(G_1)\rho$) the bounds are flat illustrating how the inability to borrow against future income growth prevents identification of the effect of children on consumption. The bounds are tightened for lower levels of income growth (G_1) and lower levels of intertemporal smoothing (ρ).

The importance of the combination of income growth and a credit constraint is clear from the analysis of the four period model. If income is constant, the OLS and IV estimators using young households deliver the correct θ . The same is true if households do not care about intertemporal smoothing of marginal utility ($\rho = 0$ and $\text{IES} = \infty$). The assumptions of income growth and finite intertemporal elasticity of substitution seem reasonable, however.

3.2 Evidence from a Multi-Period Life Cycle Model

To confirm the results from the four-period model, I setup a standard life cycle (buffer-stock) model, used intensively for analysis of intertemporal consumption behavior. The model captures the main consumption and savings incentives of households over the life cycle prior to retirement. Specifically, the model is similar to those in [Attanasio, Banks, Meghir and Weber \(1999\)](#); [Gourinchas and Parker \(2002\)](#) and [Cagetti \(2003\)](#).

Households work until an exogenously given retirement age, T_r , and die with certainty at age T where they consume all available resources. In all preceding periods, households solve the optimization problem

$$\max_{C_t} \mathbb{E}_t \left[\sum_{\tau=t}^{T_r-1} \beta^{\tau-t} v(\mathbf{z}_t; \theta) u(C_\tau) + \gamma \sum_{s=T_r}^T \beta^{s-t} v(\mathbf{z}_t; \theta) u(C_\tau) \right]. \quad (4)$$

Following [Gourinchas and Parker \(2002\)](#), survival and income uncertainty are omitted post retirement and the parameter γ (referred to as the retirement motive) in equation (4) is a parsimonious way of adjusting for these elements. [Gourinchas and Parker \(2002\)](#) ignore the post-retirement consumption decisions and adjust the perfect foresight approximation by a parameter similar to γ through a retirement value function. Although I focus on consumption behavior *prior* to retirement, the potential presence of children at retirement forces the model to be specific about post retirement behavior.

Households solve (4) subject to the intertemporal budget constraint, $M_{t+1} = R(M_t - C_t) + Y_{t+1}$, where M_t is resources available for consumption in beginning of period t

and Y_t is beginning-of-period income. End-of-period wealth, $A_t = M_t - C_t$, must be greater than a fraction $-\kappa$ of permanent income in all time periods, $A_t \geq -\kappa P_t \forall t$, $\kappa \geq 0$. Following [Gourinchas and Parker \(2002\)](#), retired households are not allowed to be net borrowers, $A_t \geq 0, \forall t \geq T_r$.

Prior to retirement, income follows a transitory-permanent income shock process,

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

where G_t is the real gross income growth, P_t denotes permanent income and $\eta_t \sim \log \mathcal{N}(-\sigma_\eta^2/2, \sigma_\eta^2)$ is a mean one permanent income shock. ε_t is a mean one transitory income shock taking the value μ with probability \wp and otherwise distributed $(1 - \wp)\varepsilon_t \sim \log \mathcal{N}(-\sigma_\varepsilon^2/2 - \mu\wp, \sigma_\varepsilon^2)$.¹⁰ When retired, the income process is a deterministic fraction $\varkappa \leq 1$ of permanent income and permanent income grows with a constant rate of G_{ret} once retired, $Y_t = \varkappa P_t, \forall t \geq T_r$, and $P_t = G_{ret} P_{t-1}, \forall t \geq T_r$.

Households can have at most three children and no infants arrive after the wife turns 43 years old. For notational simplicity, the age of each child 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. When a child is aged 21 the child does not influence household consumption in subsequent periods regardless of the value of θ . Following [Browning and Ejrnæs \(2009\)](#), the arrival of an infant is deterministic in the sense that households know with perfect foresight how many children they will have and when these children arrive.¹¹

Unlike the simple four-period model, the life cycle model does not have an analytical solution. Therefore, to simulate synthetic data, I solve the model using the Endogenous Grid Method (EGM) proposed by [Carroll \(2006\)](#) with “standard” parameters presented in Table 1. The technical details of the solution method are provided in Appendix B in the online supplemental material. The solution is then used to generate data for 50,000 households from age 22 to 59 in each of the 1,000 Monte Carlo (MC) runs. All households are initiated at age 22 with zero wealth, $A_{21} = 0$, permanent income of one (normalization), $P_{22} = 1$, and no previous children, $\mathbf{z}_{21} = (\text{NC}, \text{NC}, \text{NC})$.

¹⁰This formulation allows for both an explicit and self-imposed credit constraint. Depending on the value of κ and \wp and μ , either the explicit or the self-imposed constraint will be the relevant one. This is discussed further in Appendix B in the supplemental material. In the baseline specification, $\kappa = 0$, $\wp = 0$ and $\mu = 0$ such that only the explicit credit constraint matters. I show in the robustness exercise, that the results regarding the log-linearized Euler equation is robust to letting $\wp = 0.003$ and $\mu = 0$ such that the self-imposed no-borrowing constraint is the relevant one rather than the explicit constraint.

¹¹In the robustness analysis in Section 5, I allow children to arrive probabilistically, as in [Blundell, Dias, Meghir and Shaw \(2013\)](#), and find that the bounds are robust to this alternative fertility process.

Children are distributed across households and age according to the observed arrival of children in the PSID, as illustrated in Figure 2b, and the income profile is calibrated to be concave (Figure 2a) and constant from age 40 to mimic empirical income profiles.

Table 1 – Parameter Values Used to Simulate Data.

G_t	R	σ_ε^2	σ_η^2	κ	\wp	μ	β	ρ	γ	\varkappa	G_{ret}	θ
Fig. 2a	1.03	.005	.005	0	0	0	.95	2	1.1	.8	1.0	$\in [0, 1]$

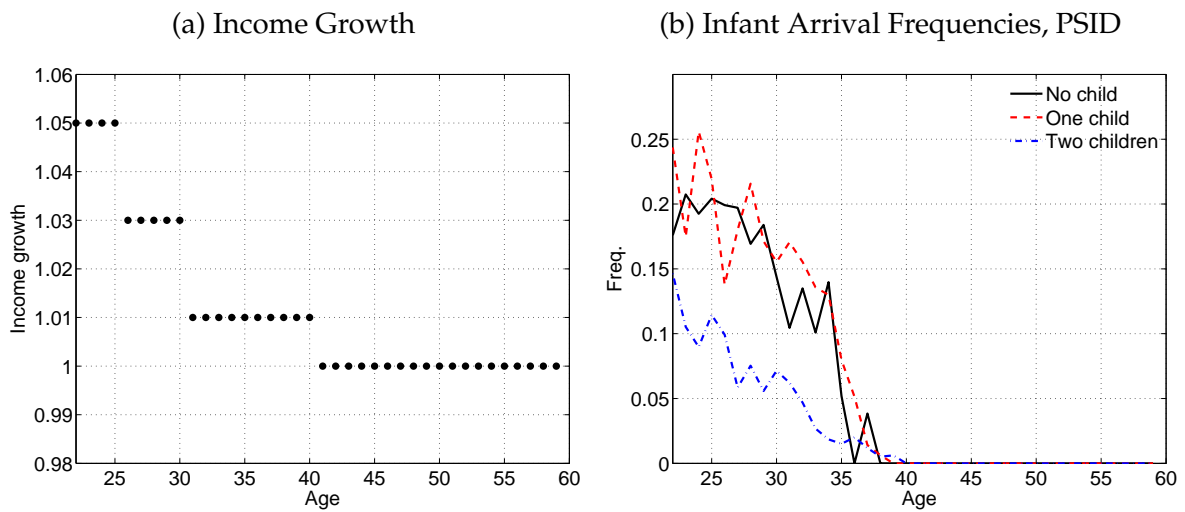


Figure 2 – Calibrated Income Growth and Arrival of Children.

Notes: Figure 2a reports how permanent income grows in the life cycle model while panel b shows how the arrival of children is calibrated using the PSID. The arrival of children is based on the PSID data described in Section 4.

Figure 3 presents simulated age profiles for income, consumption and wealth for different values of θ . All consumption profiles (even if children do not affect consumption) exhibit a hump when households are in the mid-40s, as typically observed in real data. If children affect consumption, the hump is more pronounced by a steeper consumption profile for young households and a subsequent larger decrease in consumption after the mid-40s. Income uncertainty, income growth and credit constraints produce an increasing consumption profile early in life, even if children do not affect consumption. The retirement motive produces an incentive (depending on the size of γ) to accumulate wealth for retirement later in life producing a downward sloping consumption profile after the mid-40s.

The consumption profiles are very similar for young households across θ -values. This is because credit constraints prevent households from borrowing against future income growth to increase consumption when children arrive – despite they would want to, had unlimited borrowing been possible. Hence, the effect of children would

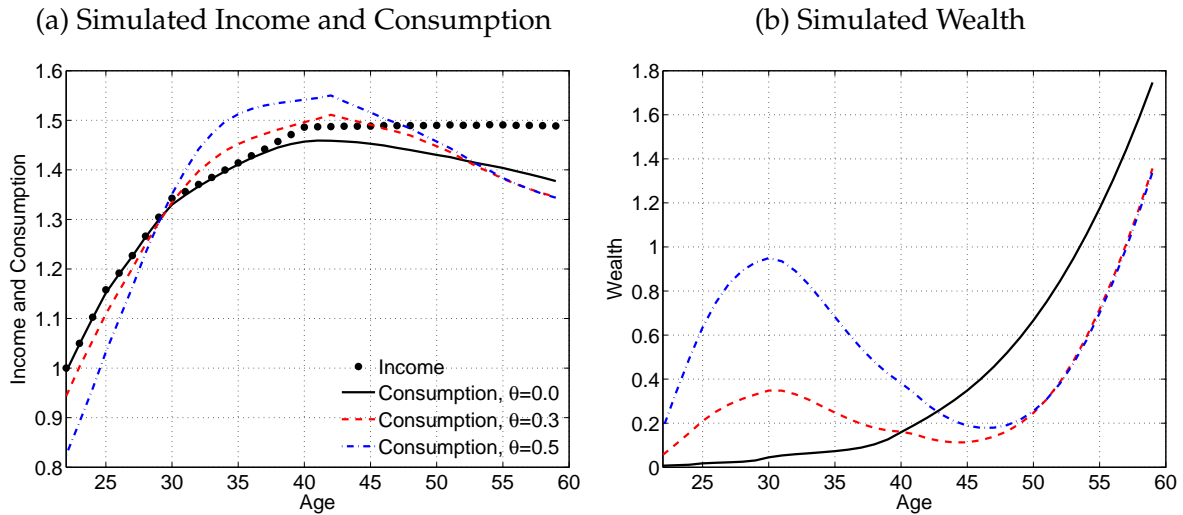


Figure 3 – Simulated Income, Consumption and Wealth Profiles.

Notes: Figure 3 illustrates the average age profile of income, consumption and wealth for 50,000 simulated households for different values of θ . Panel a shows how consumption profiles change relatively little across models with no effects of children, $\theta = 0$, through a model in which children are important, $\theta = 0.5$. Panel b shows how the wealth accumulation, on the other hand, is greatly affected by the importance of children. Particularly, a hump in the wealth profile emerges as children becomes more important.

in general be underestimated using young households, as shown earlier. Noticeably, young households accumulate large amounts of wealth in anticipation of children arriving in the future. When children subsequently arrive, wealth is almost depleted such that the credit constraint is binding for many households when children eventually leave. The relative drop in consumption from a *constrained* level to an (potentially) unconstrained level, when children leave, will in general be less than the relative change if households had never been constrained. Hence, the effect of children would be underestimated when only using older households as shown using the simple four-period model above.

Empirical age profiles of observed household wealth is typically not hump-shaped as illustrated in Figure 3 but rather monotonically increasing (Cagetti, 2003). This suggests that children might not be as important for consumption over the life cycle as previously found in the existing literature. I confirm this in section 4 below where I estimate the proposed bounds using the PSID.

Table 2 reports the average estimate of θ using all households, both young and old, from 1,000 MC runs and the standard deviation across these runs. For each run, data are simulated from the life cycle model for 50,000 households from age 22 through 59 and 20 random adjacent time-observations are drawn for each household from this simulation. It is clear that for low levels of θ , both the log-linearized and non-linear GMM estimators overestimate the effect of children on consumption while they un-

derestimate the effect if θ is large. This is true irrespectively if the actual change in number of children (Δz_t) are used in the estimation or the cohort average number of children ($\Delta \bar{z}_t$) is used as instrument.

Table 2 – Monte Carlo Results, Both Young and Old Households.

Instr.	$\theta = 0.0$		$\theta = 0.1$		$\theta = 0.5$		$\theta = 1.0$	
	LogLin	GMM	LogLin	GMM	LogLin	GMM	LogLin	GMM
Δz_t	0.015 (0.001)	0.006 (0.001)	0.086 (0.001)	0.078 (0.001)	0.227 (0.001)	0.221 (0.001)	0.397 (0.002)	0.374 (0.002)
$\Delta \bar{z}_t$	0.125 (0.002)	0.038 (0.001)	0.156 (0.002)	0.075 (0.001)	0.255 (0.002)	0.166 (0.002)	0.475 (0.003)	0.310 (0.003)

Notes: The average of all MC estimates and standard deviations (in parenthesis) across Monte Carlo runs are reported. All results are based on 1,000 independent estimations on simulated data from the life cycle model described in Section 3.2 with the parameters presented in Table 1. For each run, data are simulated for 50,000 households from age 22 through 59 and a random adjacent period of length time-observations are drawn from this simulation. All individuals are initiated at age 22 with zero wealth, $A_{21} = 0$, permanent income of one, $P_{22} = 1$, and no children. Children are assigned following the estimated arrival probabilities estimated from the PSID, reported in Figure 2b.

Figure 4 illustrates the proposed bounds based on the four period model in panel 4a and the multi-period life cycle model in panel 4b. The 45°-line represents the true value of θ while blue lines represent lower bounds and red lines represent upper bounds using the (cohort) average number of children as instrument. Young households are defined as those younger than 41.

The bounds derived from the simple four period model are very similar to the numerical bounds from the richer life cycle model. The bounds are fairly narrow for lower values of θ and the lower bound equals the true effect when $\theta = 0$ as expected. As the effect of children becomes larger, the bounds become wider and the upper bound is closest to the truth. The non-linear GMM estimator produces almost identical bounds as the log-linearized Euler equation, indicating that the nonlinear Euler equation ignoring credit constraints is an equally poor approximation to the true constrained Euler equation as the log-linearized Euler equation is.

The results show that the standard Euler equation estimators cannot in general estimate the effect of children on consumption when households face potentially binding credit constraints. Further, the multi-period life cycle model confirms that the proposed bounds are sensible in a more realistic framework. Below, I apply the bounds to the PSID and in section 5 I discuss the robustness of the proposed bounds.

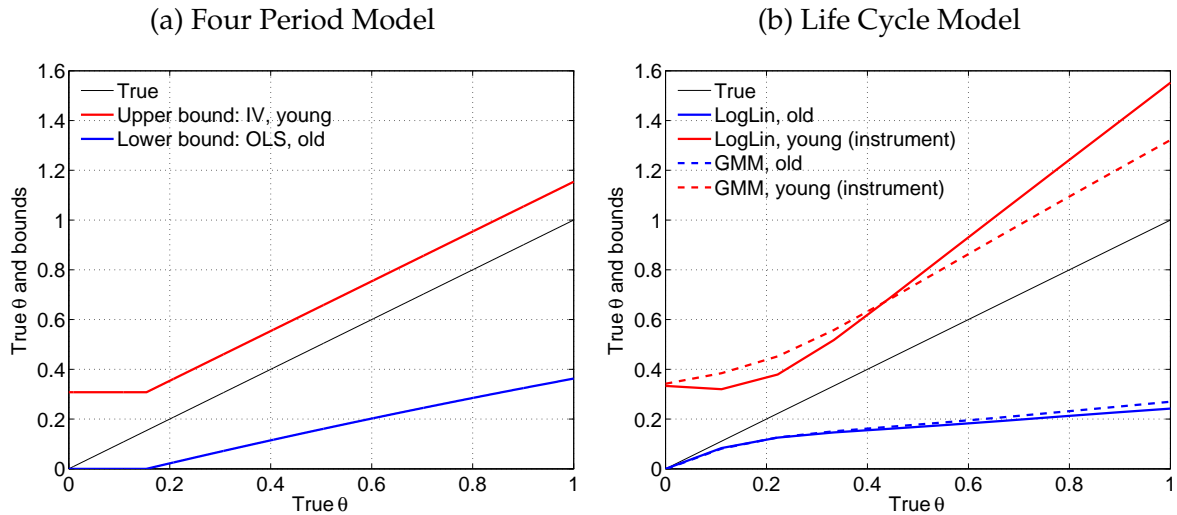


Figure 4 – Proposed Bounds, Four-Period Model and Life Cycle Model.

Notes: Figure 4 illustrates the proposed bounds based on the four period model in panel a and the life cycle model in panel b. The 45°-line represents the true value of θ while blue lines represent lower bounds and red lines represent upper bounds using the (cohort) average number of children as instrument. Solid lines are based on the log-linearized Euler equation (3) and dashed lines are based on the non-linear Euler equation estimated with the GMM estimator (2). Young households are defined as younger than 41.

4 Empirical Results from the PSID

The Panel Study of Income Dynamics (PSID) contains information on food consumption and has been used for a wide range of studies, including estimation of the effect of children on consumption. To study the evolution and link between income and consumption inequality over the 1980s, [Blundell, Pistaferri and Preston \(2008\)](#) impute total non-durable consumption for PSID households using food consumption measures in the CEX and the PSID. I use their final data set and refer the reader to their discussion of the PSID data.

The sample period is 1978 to 1992 and only male headed continuously married couples are used. The years 1987 and 1988 are not used because consumption measures were not collected those years. Since the present study focus on the effect of children on household consumption, I restrict the sample to cover households in which the wife is aged 20 to 59.¹² The supplementary low-income sub sample (SEO) is excluded from the analysis. All sample selection criteria leaves an unbalanced panel of 1,808 households observed for at most 13 periods in the final sample of in total 13,516 non-missing observations.¹³ Households are classified as high skilled if the male head has

¹²[Blundell, Pistaferri and Preston \(2008\)](#) use households in which the husband is aged 30 to 65.

¹³In an earlier working paper ([Jørgensen, 2014](#)), for tractability of an alternative estimation procedure, I restricted the sample further and removed year trends prior to estimation. The results presented here will, therefore, differ slightly from those reported in the earlier working paper.

ever enrolled in college, including college drop-outs.

Table 3 – Log-Linear Euler Equation Estimates, PSID.

	Low skilled		High skilled	
	OLS, age $\geq 45^\dagger$	IV, age $\leq 45^\ddagger$	OLS, age $\geq 45^\dagger$	IV, age $\leq 45^\ddagger$
<i>Food consumption</i>				
$\Delta\#kids$	0.031 (0.037)	0.122 (0.057)	0.035 (0.028)	0.127 (0.048)
Constant	-0.097 (0.032)	-0.049 (0.018)	-0.046 (0.027)	-0.016 (0.023)
Obs	1304	4425	1700	3351
R2	0.005	0.020	0.009	0.018
<i>Non-durable consumption (imputed)</i>				
$\Delta\#kids$	0.058 (0.063)	0.130 (0.079)	-0.033 (0.026)	0.047 (0.048)
Constant	-0.668 (0.116)	-0.101 (0.073)	-2.011 (0.028)	-1.993 (0.023)
Obs	1304	4425	1700	3351
R2	0.062	0.002	0.708	0.675

Notes: Reported are estimates of $\widehat{\rho^{-1}\theta}$ and a constant from a log-linear Euler equation estimation of food consumption in the top panel and total non-durable consumption, imputed by [Blundell, Pistaferri and Preston \(2008\)](#). Robust standard errors in parenthesis. All regressions include year-dummies. Households are classified as high skilled if the male head has ever enrolled in college, including college drop-outs. Age refers to the wife's age.

† This corresponds to the suggested lower bound of $\rho^{-1}\theta$.

‡ This corresponds to the suggested upper bound of $\rho^{-1}\theta$. The number of children is instrumented with the cohort-average number of children.

Table 3 reports the estimated bounds of $\rho^{-1}\theta$ for the PSID data using the log-linearized Euler equation (3) with year-dummies included in all regressions. Recall that the suggested lower bound on θ can be estimated using the change in number of children (Δz_t) while restricting the sample to include only older households and an upper bound can be found by using the cohort-average number of children ($\Delta \bar{z}_t$) as an instrument while restricting the sample to younger households.

To illustrate how these bounds provide valuable information, imagine having estimated the IES simultaneously (as done in, e.g., [Attanasio, Banks, Meghir and Weber, 1999](#) using interest rate variation) or having information of this parameter from elsewhere. For example, [Gourinchas and Parker \(2002\)](#) estimate $\rho \approx 0.87$ for high school graduates (low skilled by my definition) and $\rho \approx 2.29$ for college graduates (high skilled by my definition). Using these values, we could infer that the effect of children on total non-durable consumption is between .05 and .11 for low skilled and between -.08 and .11 for high skilled.

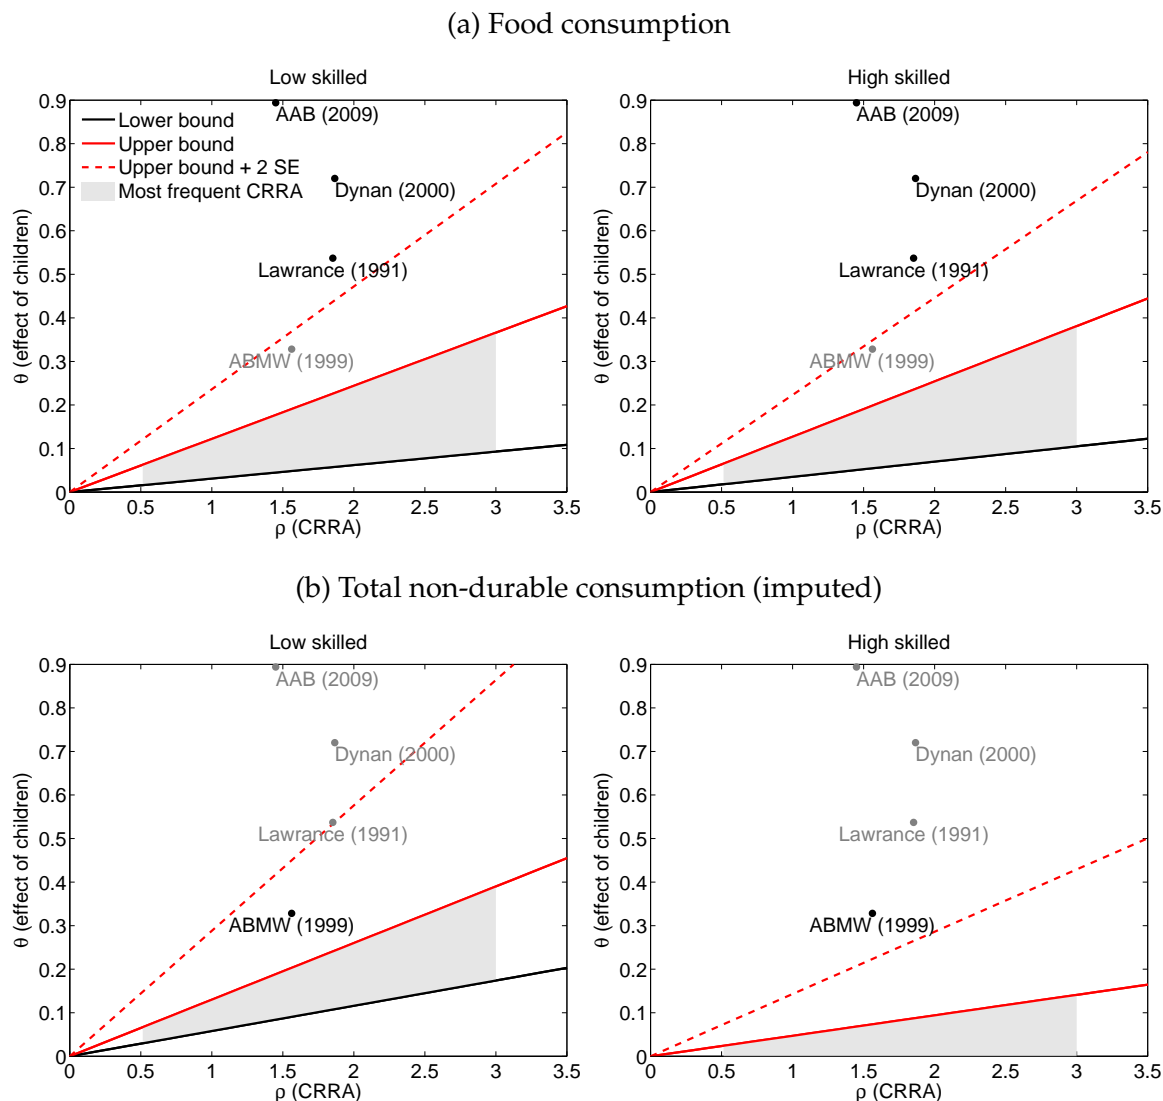


Figure 5 – Estimated Bounds of the Effect of Children on Consumption from the PSID.

Notes: Figure 5 reports the upper and lower bounds for low and high skilled households when varying ρ , the inverse of elasticity of intertemporal substitution. The top panel a reports results when using changes in log food consumption while the bottom panel b reports results when using changes in log non-durable consumption, imputed by [Blundell, Pistaferri and Preston \(2008\)](#). "Lawrance (1991)" refers to from [Lawrance \(1991\)](#), "ABMW (1999)" refers to [Attanasio, Banks, Meghir and Weber \(1999\)](#), "Dynan (2000)" refers to [Dynan \(2000\)](#) and "AAB (2009)" refers to [Alan, Attanasio and Browning \(2009\)](#). Gray dots illustrate that a different measure of consumption was used in the associated study.

Figure 5 reports how the upper and lower bounds for low and high skilled vary with the coefficient of relative risk aversion (the inverse of the IES). The top panel (panel 5a) reports results when using changes in log food consumption while the bottom panel (panel 5b) reports results when using changes in log non-durable consumption, imputed by [Blundell, Pistaferri and Preston \(2008\)](#).

The estimated effects of children on consumption reported in the existing literature are outside the proposed bounds. Specifically, the reported estimate of $\rho^{-1}\theta$ in the influential study by [Attanasio, Banks, Meghir and Weber \(1999\)](#) of 0.21 and their

estimated ρ^{-1} of .64 produce an effect of children on *non-durable consumption* in the CEX around .33. This is outside the upper bounds reported in Figure 5 for values of ρ in a neighborhood of their estimated value of 1.56. Likewise, Figure 5 maps the implied estimated effect of children on *food consumption* in the PSID reported in [Lawrance \(1991\)](#); [Dyran \(2000\)](#) and [Alan, Attanasio and Browning \(2009\)](#). The latter is based on a non-linear GMM estimator allowing for log-normal measurement error in consumption while the two former studies are based on the log-linearized Euler equation. All these estimates are outside the upper bound. Adding two times the standard error of the estimated $\widehat{\rho^{-1}\theta}$ reported in Table 3 widens the bounds significantly, but only the estimate from [Attanasio, Banks, Meghir and Weber \(1999\)](#) is now included in the bounds.

5 Robustness of Bounds

Choosing the age at which to split the sample into young and older households is not obvious. One choice could be to choose the age at which the average number of children starts to decline since the behavior of households should differ when children arrive from when they leave, c.f. the above discussion. Simply estimating different parameters related to when children arrive and move could be a route to pursue. Alternatively, the age at which average net wealth is significantly larger than average income could be chosen since around this point (on average) households are less affected by credit constraints.

A crucial assumption when calculating the bounds above is that of the researcher having knowledge on other structural parameters. Using the exact GMM estimation approach, both the discount factor, β , and the relative risk aversion, ρ , should be estimated simultaneously or qualified guesses on these parameters should be used. Log-linearized Euler equation estimation requires information only on the risk aversion parameter. This is a drawback but varying these preference parameters in “accepted” ranges produces a set of bounds with information on the size of the effect of children on consumption.

The bounds are robust to changing the calibrated parameters in Table 1. Figure 6 illustrates the bounds from models in which the discount factor, β , is .975 and .99. The bounds are robust and for larger values of the discount factor the bounds become increasingly tight. Especially for low values of θ , both the upper and lower bounds are close to the true value when households are more patient. This stems from the fact that more patient households put more emphasis on future marginal utility and, thus, accumulate more wealth prior to the arrival of children. In turn, the credit constraint has less bite.

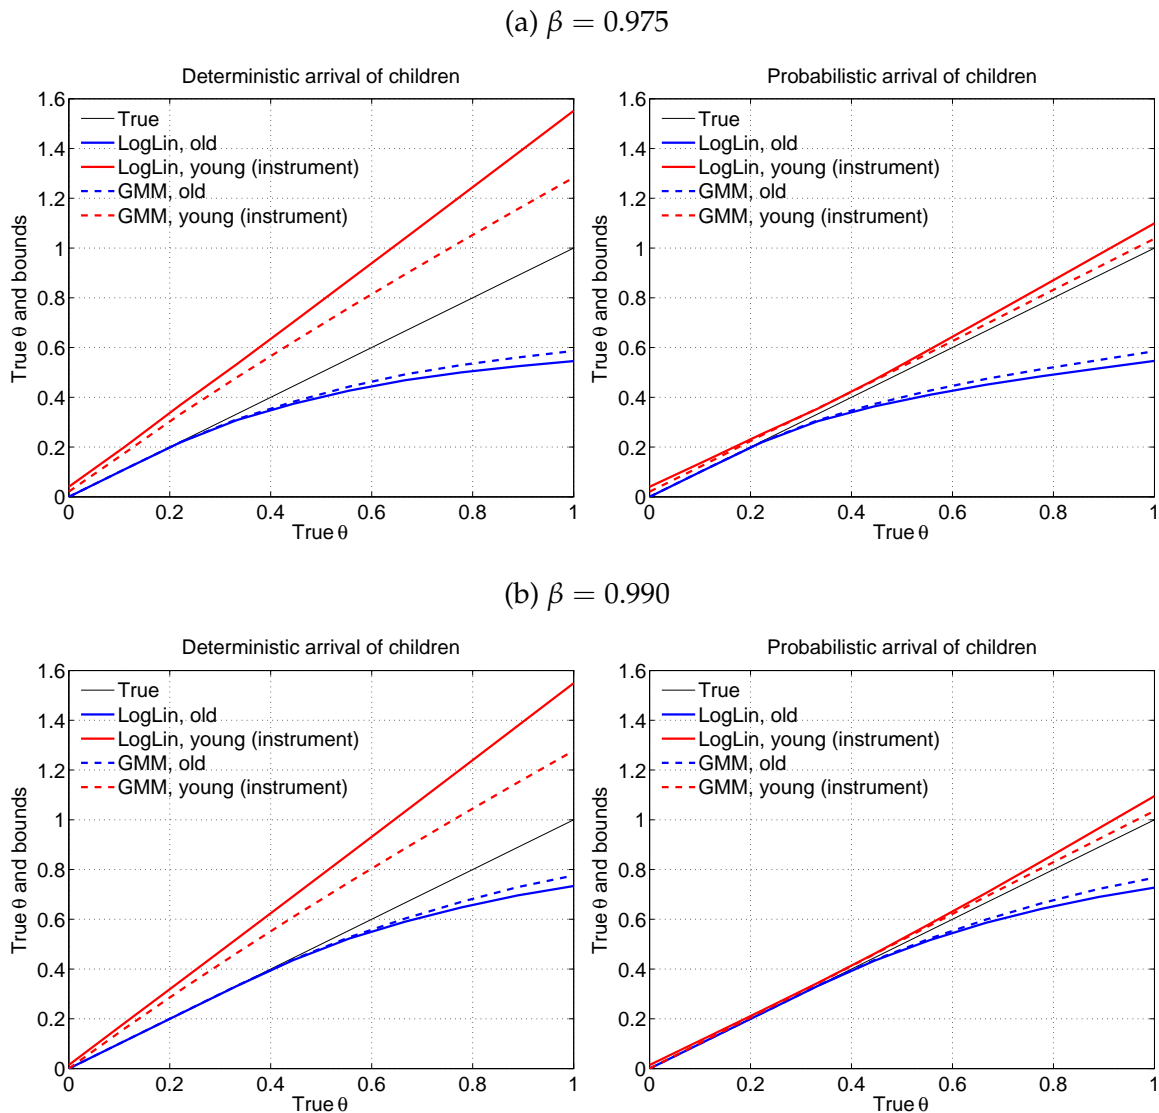


Figure 6 – Bounds, $\beta = \{0.975, 0.99\}$.

Notes: Figure 6 illustrates the proposed bounds based the life cycle model with a discount factor of .975 in panel a and .99 in panel b. The figures to the left illustrates the bounds from the baseline deterministic model, in which children are perfectly foreseen and figures to the right illustrates the bounds from a model in which children arrive probabilistically. The 45°-line represents the true value of θ while blue lines represent lower bounds and red lines represent upper bounds using the (cohort) average number of children as instrument. Solid lines are based on the log-linearized Euler equation (3) and dashed lines are based on the non-linear Euler equation estimated with the GMM estimator (2).

Below, I argue that the results are robust to alternative fertility processes, labor market costs of children and to replacing the explicit credit constraint with a zero-income shock. The robustness results are important since they stress that the true underlying effect of children on consumption is in between the lower and upper bound in realistic circumstances. An alternative route to estimating bounds could be to utilize the moment *inequality* rather than the equality in the GMM estimator (2). Assuming that an instrument is potentially positively correlated with the Euler residual, the inequality $\mathbb{E}[(\epsilon - 1)'Z] \geq 0$ could be used as a moment inequality to estimate bounds (Moon and

Schorfheide, 2009). This approach is very interesting for future research but I do not pursue that strategy here.¹⁴

5.1 Alternative Fertility Processes

All results have been derived assuming that children are perfectly foreseen. This assumption has primarily been deployed for tractability of the four period model since that model could then be solved analytically. Versions of the model in which children arrive probabilistically as in Blundell, Dias, Meghir and Shaw (2013) produce qualitatively unchanged results.

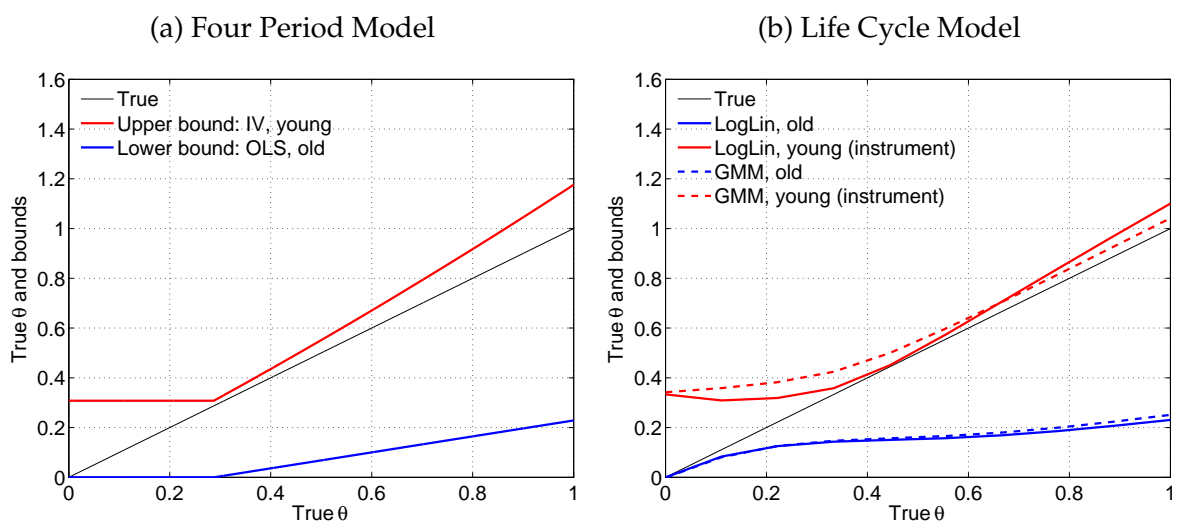


Figure 7 – Bounds, Probabilistic Arrival of Children.

Notes: Figure 7 illustrates the proposed bounds based on a model in which children arrive probabilistically rather than being perfectly foreseen as in the "deterministic" baseline model. Panel a illustrates the bounds from the four period model and the bounds from the life cycle model is illustrated in panel b. The 45°-line represents the true value of θ while blue lines represent lower bounds and red lines represent upper bounds using the (cohort) average number of children as instrument. Solid lines are based on the log-linearized Euler equation (3) and dashed lines are based on the non-linear Euler equation estimated with the GMM estimator (2). Young households are defined as younger than 41.

Figure 7 illustrates the proposed bounds based on the four period model in panel 7a and the life cycle model in panel 7b for the probabilistic version of the models. The bounds are very similar to those presented from the baseline model and the upper bound is close to the true effect of children when children arrive probabilistically.

In the probabilistic version, households are identical prior to arrival of children. In the four period model, all households save exactly the same in period zero, prior to a child potentially arrives in period one. In period one, households who do not receive a child increase consumption due to the fact that it has been revealed for them that they

¹⁴I am grateful to Dennis Kristensen for pointing this out to me.

will not have children. Their accumulated wealth from period zero is now distributed across remaining periods. This increase in consumption of childless households will bias the estimates downwards. This is true even if households do *not* face credit constraints and motivates the use of the OLS estimate from older households rather than the OLS estimate from young households to estimate a lower bound.

Children could, alternatively, be chosen endogenously. Endogenous fertility would significantly alter the economic environment and is typically *not* implemented in empirical work on the effect of children on consumption. It is important to stress that households in the deterministic life cycle model have strong incentives to accumulate wealth to finance increased consumption when children arrive. Thus, the baseline model is the one in which credit constraints has the *least* of an effect on the implied consumption behavior, compared to the probabilistic version. Further, the biological “constraint” on female fertility will interplay with the financial constraints and the latter is, thus, still likely to be important for household behavior in a model in which children are perfectly chosen by households (Almlund, 2013).

Importantly, the bounds will still be valid even if the financial credit constraint has less bite when households perfectly chose when to have children. Specifically, in the extreme case when households are perfectly able to break free of the constraint and is *never* affected by constraints, the lower bound will overlap with the true effect of children and the upper bound will be slightly above the true effect (see discussion in section 5.3).

5.2 Children and the Labor Market

As in the rest of the literature on the effect of children on consumption, income is assumed *independent* of household composition. If income depends on household composition, the results will change depending on in which ways children affect the labor market income of household members. The focus on the present study is on how Euler equation estimation fails to uncover the underlying effect of children if potentially binding credit constraints (either explicit or self-imposed) are ignored. Although allowing income to vary with household composition is an interesting avenue for future research, I have not pursued that here. The primary reason is that how children should affect labor market outcomes is not obvious and the results, in turn, would be hard to relate to existing estimates of the effect of children on consumption.

Children might, however, affect the number of hours worked. Calhoun and Espenshade (1988) estimate a substantial decrease in labor market hours of American females in response to childbearing. In a more recent working paper, Adda, Dustmann and Stevens (2012) analyze in a life cycle model of German households, the career cost of children and find that children can explain a substantial portion of the male-female

gender wage gap. In their model of fertility, occupational choice and labor supply, consumption is assumed linear in income. In turn, all households are constrained in their model illustrating that implementing *all* features into one model is not nearly as feasible as it is interesting.

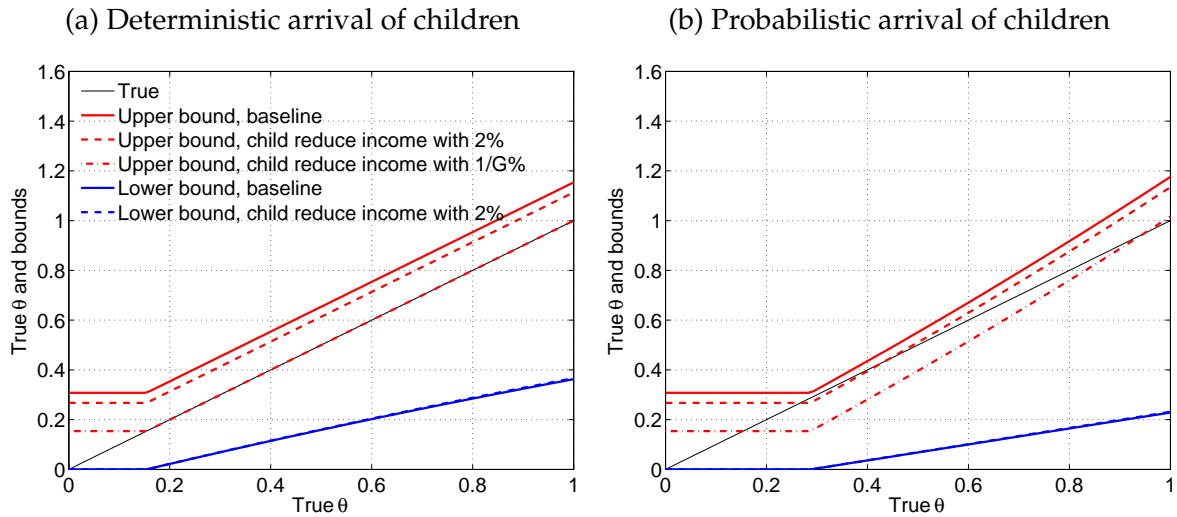


Figure 8 – Bounds if Children Reduce Labor Market Income.

Notes: Figure 8 illustrates the proposed bounds from the four period model when children reduce labor market income with 2 percent. Panel a illustrates the bounds from the model in which children arrive deterministically and the probabilistic version is illustrated in panel b. The 45°-line represents the true value of θ while blue lines represent lower bounds and red lines represent upper bounds using the (cohort) average number of children as instrument. Solid lines are based on the baseline case where children does not affect labor market income and dashed lines represent the extreme case where children reduce labor market income with 2 percent.

The bounds are still valid if children reduce permanent income, as suggested by the results above. This is true as long as children does not reduce permanent income by more than the permanent income growth as illustrated in Figure 8a. If children arrive deterministically and children reduce income to a degree that only childless households experience income growth, the upper bound equals the true effect. If children arrives probabilistically, however, the upper bound might be below the true effect, as illustrated by panel 8b.

5.3 Self-imposed No-borrowing vs. Explicit Credit Constraint

The results generalize to cases in which consumers do not face credit constraints. If risk averse consumers instead face a positive probability of receiving a zero-income shock (as in [Carroll, 1997](#) and [Gourinchas and Parker, 2002](#)), all results concerning the log-linearized Euler equation (3) still hold. This is basically because risk averse consumers will instead face a “self-imposed” no-borrowing constraint stemming from the

fear of receiving zero income in all future periods with consumption of zero as a consequence (Schechtman, 1976; Zeldes, 1989b; Carroll, 1992). In turn, consumption will respond substantially to negative income shocks if either explicit or self-imposed credit constraints affect consumers, increasing the variance in consumption growth. Because higher order moments (such as something like the variance of consumption growth, Carroll, 2001) enters the reduced form residual, $\tilde{\epsilon}$, log-linearized Euler equation estimation will not be able to uncover the effect of children on consumption. This result supports and extends the critique in Carroll (2001) on the inability of log-linearized Euler equation estimation to uncover the IES to the inability to uncover demographic effects on consumption.

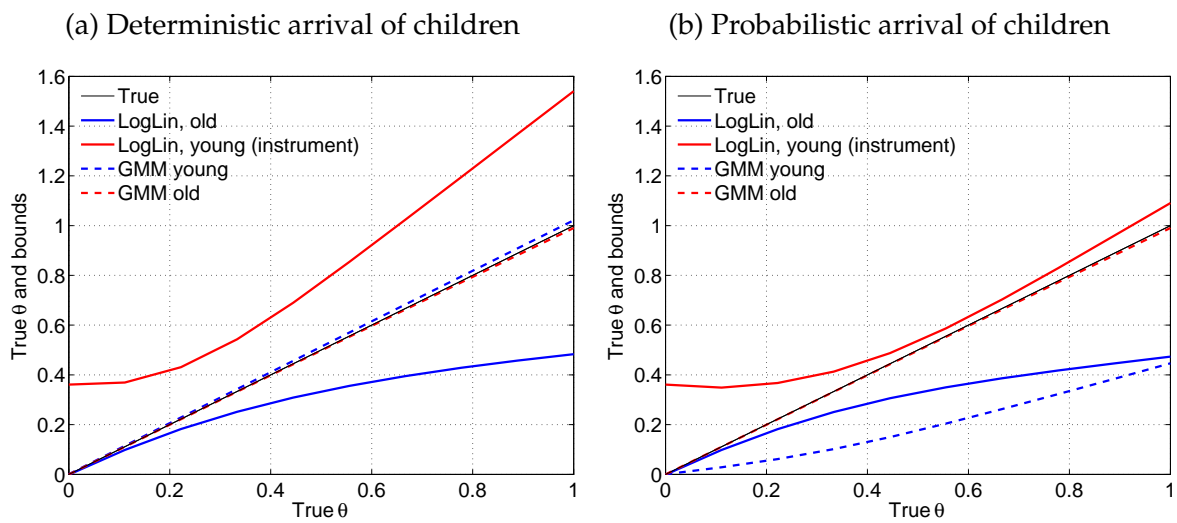


Figure 9 – Bounds, No Explicit Constraint but Self-Imposed No-borrowing.

Notes: Figure 9 illustrates the proposed bounds based on a model in which there is unlimited borrowing but a positive probability of receiving zero income. Panel a illustrates the bounds from the baseline deterministic model, in which children are perfectly foreseen and panel b illustrates the bounds from a model in which children arrive probabilistically. The 45°-line represents the true value of θ while blue lines represent lower bounds and red lines represent upper bounds using the (cohort) average number of children as instrument. Solid lines are based on the log-linearized Euler equation (3) and dashed lines are based on the non-linear Euler equation estimated with the GMM estimator (2). Young households are defined as younger than 41.

In absence of “explicit” credit constraints, the Euler equation in (1) has mean one because $\lambda_t = 0, \forall t$ and the exact GMM estimator is expected to produce unbiased estimates of the effect of children. Figure 9 illustrates the log-linearized Euler equation bounds along with GMM estimates from a model with a .3 percent risk of zero household income (as calibrated in Gourinchas and Parker, 2002). Panel 9a illustrates the bounds from the deterministic model while panel 9b illustrates the bounds from the probabilistic version of the model. It is clear that when children are perfectly foreseen (panel 9a), the non-linear GMM produces the correct estimate.

Interestingly, if children arrive probabilistically, using young households to uncover the effect of children on consumption is impossible even if there is no explicit constraint on borrowing, as shown by the fact that the blue dashed (GMM on young) in Figure 9b is significantly below the 45°-line. This stems from the feature of the probabilistic model that households who turn out childless have accumulated as much wealth in their youth as those households who eventually had children. Comparing consumption growth of households with and without children will, thus, underestimate the true effect of children, if children arrive probabilistically. Appendix A.4 in the online supplemental material proves this result (using the four period model): If children arrive probabilistically, using the cohort average number of children as instrument when estimating the log-linearized Euler equation (proposed upper bound) overestimates the effect of children even if there is *no* explicit credit constraint.

Table 4 – Monte Carlo Results, No Explicit Constraint.

Instr.	$\theta = 0.0$		$\theta = 0.1$		$\theta = 0.5$		$\theta = 1.0$	
	LogLin	GMM	LogLin	GMM	LogLin	GMM	LogLin	GMM
<i>Deterministic arrival of children</i>								
$\Delta \mathbf{z}_t$	0.015 (0.001)	-0.000 (0.003)	0.099 (0.001)	0.100 (0.003)	0.358 (0.001)	0.500 (0.011)	0.580 (0.002)	0.995 (0.024)
$\Delta \bar{\mathbf{z}}_t$	0.118 (0.002)	-0.000 (0.007)	0.167 (0.002)	0.104 (0.008)	0.311 (0.002)	0.518 (0.017)	0.530 (0.003)	1.024 (0.031)
<i>Probabilistic arrival of children</i>								
$\Delta \mathbf{z}_t$	0.015 (0.001)	-0.000 (0.003)	0.086 (0.001)	0.084 (0.003)	0.283 (0.001)	0.424 (0.011)	0.447 (0.001)	0.852 (0.023)
$\Delta \bar{\mathbf{z}}_t$	0.118 (0.002)	-0.000 (0.007)	0.166 (0.002)	0.100 (0.008)	0.272 (0.002)	0.501 (0.017)	0.497 (0.002)	1.000 (0.028)

Notes: The average of all MC estimates and standard deviations (in parenthesis) across Monte Carlo runs are reported. All results are based on 1,000 independent estimations on simulated data from the life cycle model described in Section 3.2 with the parameters presented in Table 1. For each run, data are simulated for 50,000 households from age 22 through 59 and a random adjacent period of length 20 time-observations are drawn from this simulation. All individuals are initiated at age 22 with zero wealth, $A_{21} = 0$, permanent income of one, $P_{22} = 1$, and no children. The results are based on a life cycle model in which there is no explicit constraint but instead a .3 percent risk of a zero income shock, producing a self-imposed no-borrowing constraint. In the top panel, children arrive with perfect foresight while in the bottom panel children arrive probabilistically, following the estimated arrival probabilities estimated from the PSID, reported in Figure 2b.

Table 4 reports the Monte Carlo results from pooling young and older households, using the life cycle model *without* an explicit credit constraint. The top panel reports results from the baseline model and the bottom panel reports the results if children arrive probabilistically. It is clear that the non-linear GMM estimator can uncover the

correct estimate while the log-linearized Euler equation cannot when children are perfectly foreseen. The results in the bottom panel, where children arrive probabilistically, confirm that in this case, the GMM estimator using both young and older households cannot uncover the true effect of children on consumption (unless it is zero). Using only older households, in which children leave, will, however, lead the GMM estimator to produce unbiased estimates when children arrive probabilistically (Figure 9b).

6 Concluding Discussion

Many studies use estimators derived from the consumption Euler equation. Especially the log-linearized Euler equation is popular since it yields estimable equations linear in parameters which can easily be estimated with synthetic cohort panels and handle measurement error through instrumental variables estimation. Although these estimators have now become work horses in the analysis of intertemporal consumption behavior, little is known about their performance when households face potentially binding credit constraints and the standard Euler equation, thus, no longer holds.

I have showed how both the non-linear and the log-linearized Euler equation estimators fail to uncover the true underlying effect of children on consumption when potentially binding credit constraints are ignored. Through splitting the sample into young households, in which children arrive, and older households, in which children leave, I propose a tractable approach to uncovering bounds of the effect of children on consumption using these conventional estimators.

Estimating the proposed bounds on PSID data indicates that *all*, to the best of my knowledge, existing estimates of the effect of children on consumption are above the upper bound. In turn, these results suggest that the importance of children in intertemporal consumption behavior, found in previous studies, might simply proxy for the inability of households to borrow against future income growth.

Arguably, the proposed bounds suffer from many of the same assumptions as most existing empirical literature analyzing the intertemporal consumption behavior. Particularly, it has been assumed throughout (and in the related literature) that fertility is exogenous and children do *not* affect labor market outcomes. Although the bounds are somewhat robust to these assumptions, they have been invoked for tractability and comparability with existing studies of the effect of children on consumption. Allowing for endogenous fertility and endogenous labor market supply with children affecting the dis-utility from work is extremely interesting for future research.

References

- ADDA, J. AND R. COOPER (2003): *Dynamic Economics: Quantitative Methods and Applications*. MIT Press.
- ADDA, J., C. DUSTMANN AND K. STEVENS (2012): "The Career Costs of Children," Working Paper 2012/01, European University Institute.
- ALAN, S., K. ATALAY AND T. F. CROSSLEY (2012): "Euler Equation Estimation on Micro Data," Working Paper 1221, Koc University-TUSIAD Economic Research Forum.
- 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.
- ALAN, S. AND M. BROWNING (2010): "Estimating Intertemporal Allocation Parameters using Synthetic Residual Estimation," *The Review of Economic Studies*, 77(4), 1231–1261.
- ALMLUND, M. (2013): "Essays on Credit Constraints, Education, and the Family," Phd thesis, Economics Department, University of Chicago.
- 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 H. LOW (2004): "Estimating Euler equations," *Review of Economic Dynamics*, 7(2), 406–435.
- 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.
- 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.
- 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.
- (1997): "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *The Quarterly Journal of Economics*, 112(1), pp. 1–55.
- (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.
- CROSSLEY, T. F. AND H. LOW (forthcoming): "Job Loss, Credit Constraints and Consumption Growth," *Review of Economics and Statistics*.
- DYNAN, K. E. (2000): "Habit Formation in Consumer Preferences: Evidence from Panel Data," *American Economic Review*, 90(3), 391–406.
- FERNÁNDEZ-VILLAYERDE, 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-VILLAYERDE, 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.

- GROSS, D. B. AND N. S. SOULELES (2002): "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data," *The Quarterly Journal of Economics*, 117(1), pp. 149–185.
- GROSS, T., M. J. NOTOWIDIGDO AND J. WANG (2014): "Liquidity Constraints and Consumer Bankruptcy: Evidence from Tax Rebates," *Review of Economics and Statistics*, 96(3), 431–443.
- HALL, R. E. AND F. S. MISHKIN (1982): "The Sensitivity of Consumption to Transitory Income," *Econometrica*, 50(2), 461–481.
- IRVINE, I. (1978): "Pitfalls in the Estimation of Optimal Lifetime Consumption Patterns," *Oxford Economic Papers*, 30(2), pp. 301–309.
- JAPPELLI, T. (1990): "Who is Credit Constrained in the U. S. Economy?," *The Quarterly Journal of Economics*, 105(1), pp. 219–234.
- JAPPELLI, T., J.-S. PISCHKE AND N. S. SOULELES (1998): "Testing For Liquidity Constraints In Euler Equations With Complementary Data Sources," *The Review of Economics and Statistics*, 80(2), 251–262.
- JØRGENSEN, T. H. (2014): "Life-Cycle Consumption and Children," Working Paper 2014-2, Centre for Applied Microeconometrics (CAM), University of Copenhagen.
- JOHNSON, D. S., J. A. PARKER AND N. S. SOULELES (2006): "Household Expenditure and the Income Tax Rebates of 2001," *American Economic Review*, 96(5), 1589–1610.
- LAWRANCE, E. C. (1991): "Poverty and the Rate of Time Preference: Evidence from Panel Data," *Journal of Political Economy*, 99(1), 54–77.
- LETH-PETERSEN, S. (2010): "Intertemporal Consumption and Credit Constraints: Does Total Expenditure Respond to an Exogenous Shock to Credit?," *American Economic Review*, 100(3), 1080–1103.
- LUDVIGSON, S. AND C. H. PAXSON (2001): "Approximation Bias in Linearized Euler Equations," *Review of Economics and Statistics*, 83(2), 242–256.
- MOON, H. R. AND F. SCHORFHEIDE (2009): "Estimation with overidentifying inequality moment conditions," *Journal of Econometrics*, 153(2), 136 – 154.
- RUNKLE, D. E. (1991): "Liquidity constraints and the permanent-income hypothesis: Evidence from panel data," *Journal of Monetary Economics*, 27(1), 73 – 98.
- SCHECHTMAN, J. (1976): "An Income Fluctuation Problem," *Journal of Economic Theory*, 12, 218–41.

THUROW, L. C. (1969): "The Optimum Lifetime Distribution of Consumption Expenditures," *The American Economic Review*, 59(3), pp. 324–330.

ZELDES, S. P. (1989a): "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, 97(2), pp. 305–346.

——— (1989b): "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *Quarterly Journal of Economics*, 104(2), 275–298.

Online Appendix (Not for Publication)

Euler Equation Estimation: Children and Credit Constraints

Thomas H. Jørgensen
 University of Copenhagen
 October 29, 2014

A Solving the Four-Period Model

All variables are normalized with income such that small letter variables are normalized ones. Hence, e.g., $m_1 = (A_0 + Y_1)/Y_1 = G_1^{-1}a_0 + 1$ since $Y_1 = G_1Y_0$. In all other periods, income is constant. This normalization facilitates solving the model analytically for *all* possible values of income. The resulting consumption function should be multiplied with current period income to give the un-normalized level of consumption, $C_t^* = Y_t c_t^*$. The consumption functions in periods $t = 1, 2, 3$ are independent of whether children arrive deterministically or probabilistically since it is assumed that children, if present in period $t = 1$, will move with certainty in period $t = 2$. Therefore, I first solve for optimal consumption in period $t = 1, 2, 3$ which are identical for the deterministic and probabilistic versions and then turn to the initial period consumption, prior to potential arrival of children. This analysis is split between the model in which children arrive deterministically and the model in which children arrive probabilistically.

In the terminal period, period three, all resources are consumed ($c_3^* = m_3$) and the *unconstrained* Euler equation linking period two and period three consumption is then

$$c_2^{-\rho} = m_3^{-\rho}$$

such that inserting normalized resources, $m_3 = m_2 - c_2 + 1$ and re-arranging shows that optimal consumption in period two is the minimum of available resources, m_2 , and $\frac{1}{2}(m_2 + 1)$. Since income does not fall between period one and two and because negative wealth is not allowed, $m_2 \geq 1$ and optimal consumption is then

$$c_2^*(m_2) = \frac{1}{2}(m_2 + 1). \tag{A.1}$$

In period one, a child may be present and the *unconstrained* Euler equation is given by

$$c_1^{-\rho} \exp(\theta \mathbf{z}_1) = c_2^{-\rho},$$

such that inserting normalized resources and re-arranging yields,

$$c_1^*(m_1|\mathbf{z}_1) = \min \left\{ m_1, \frac{m_1 + 2}{1 + 2 \exp(-\rho^{-1}\theta\mathbf{z}_1)} \right\}, \quad (\text{A.2})$$

where the constraint is binding if $m_1 \leq \underline{m}_1 \equiv \exp(\rho^{-1}\theta\mathbf{z}_1)$. Note, this is certainly the case if nothing is saved from period zero.

Optimal consumption in period $t = 0$ depends on whether children arrive deterministically or probabilistically in period one. I first derive optimal consumption in the case where children arrive deterministically and then turn to the probabilistic arrival of children.

A.1 Initial Period Consumption: Deterministic Arrival of Children

In the first period, the *unconstrained* Euler equation is

$$c_0^{-\rho} = G_1^{-\rho} \exp(\theta\mathbf{z}_1) c_1^{-\rho},$$

since income grows with a factor G_1 from period zero to period one. Since consumption in period one is potentially constrained, this has to be explicitly taken into account. First, assuming that period one consumption is less than available resources, $c_1 < m_1$, inserting the optimal consumption found in (A.2) and tedious re-arranging yields optimal consumption in this case,

$$c_0^*(m_0|\mathbf{z}_1)^{\text{det}}|_{c_1 < m_1} = \frac{m_0 + 3G_1}{3 + \exp(\rho^{-1}\theta\mathbf{z}_1)}. \quad (\text{A.3})$$

If, on the other hand, consumption in period one is constrained ($c_1 = m_1$), inserting this in the Euler equation and re-arranging yields,

$$c_0^*(m_0|\mathbf{z}_1)^{\text{det}}|_{c_1 = m_1} = \frac{m_0 + G_1}{1 + \exp(\rho^{-1}\theta\mathbf{z}_1)}. \quad (\text{A.4})$$

To determine which of the consumption functions is relevant, note that equation (A.3) would imply a too high level of consumption in period zero if ignoring, that at some point, consumption in period one would be constrained because “too little” is saved in period zero. Hence,

$$\tilde{c}_0^*(m_0|\mathbf{z}_1)^{\text{det}} = \min \left\{ m_0, \frac{m_0 + G_1}{1 + \exp(\rho^{-1}\theta\mathbf{z}_1)}, \frac{m_0 + 3G_1}{3 + \exp(\rho^{-1}\theta\mathbf{z}_1)} \right\},$$

where the level of period $t = 0$ resources implying that consumption in period one is constrained is the level of resources, $\underline{m}_0^1 = \exp(\rho^{-1}\theta\mathbf{z}_1)G_1$, that makes the expression

in (A.4) to be less than that in (A.3). Hence, when $m_0 \leq \underline{m}_0^1$ optimal consumption in period $t = 0$ is given by equation (A.4) and when $m_0 > \underline{m}_0^1$ optimal consumption is given by equation (A.3).

When households are initiated with zero wealth ($a_{-1} = 0$) available normalized resources in period zero is one, $m_0 = 1$, and only equation (A.4) is relevant since $m_0 = 1 \leq \underline{m}_0^1$ for all values of $\theta \geq 0$ and $G_1 \geq 1$. Therefore, assuming no initial wealth and deterministic arrival of children, optimal consumption in period zero is given by

$$c_0^*(m_0|\mathbf{z}_1)^{\text{det}} = \min \left\{ m_0, \frac{m_0 + G_1}{1 + \exp(\rho^{-1}\theta\mathbf{z}_1)} \right\}, \quad (\text{A.5})$$

where for $m_0 \leq \underline{m}_0^2 \equiv \exp(-\rho^{-1}\theta\mathbf{z}_1)G_1$, the constraint is binding and it is optimal to consume everything. This is very intuitive: If income growth is very high, resources next period is much higher and saving today is less attractive. On the other hand, if children affect marginal utility a lot (θ large), the level of resources should be very low before it is optimal not save anything for next period, in which a child arrives.

Note, focusing on the situation in which a child arrives in period one, if $m_0 \leq \underline{m}_0^2$ next-period resources is $m_1 = G_1^{-1}(m_0 - \frac{m_0+G_1}{1+\exp(\rho^{-1}\theta)}) + 1$ and we can check whether this is less than \underline{m}_1 which is the case as long as $\theta \geq 0$ and $G_1 \geq 1$. Hence, if $m_0 = 1 \leq \underline{m}_0$ we know that $m_1 \leq \underline{m}_1$ and $c_1^*(m_1|\mathbf{z}_1 = 1) = m_1$. If a child do not arrive, optimal consumption in all periods would be to consume available resources, since in period $t = 0$, borrowing against future income growth is not allowed. This is used when calculating the OLS and IV estimators below.

A.2 Initial Period Consumption: Probabilistic Arrival of Children

The analysis, if children arrive probabilistically, is slightly more complicated than the above where children arrive deterministically. The *unconstrained* Euler equation is here given by

$$c_0^{-\rho} = G_1^{-\rho} c_1^{-\rho} (p \exp(\theta) + 1 - p),$$

such that in case where period one consumption is unconstrained ($c_1 < m_1$), inserting optimal consumption from equation (A.2) and re-arranging yields

$$c_0^*(m_0)|_{c_1 < m_1} = \frac{m_0 + 3G_1}{1 + [p(\exp(\rho^{-1}\theta) + 2)^\rho + (1-p)3^\rho]^{\frac{1}{\rho}}}. \quad (\text{A.6})$$

However, if households are potentially credit constrained if a child arrives next

period, the model has in general no analytical solution because the Euler equation is

$$\begin{aligned} c_0^{-\rho} &= G_1^{-\rho} \left[c_1^{-\rho} (1-p) + p \exp(\theta) m_1^{-\rho} \right], \\ &= G_1^{-\rho} \left[\left[\frac{1}{3} (G_1^{-1}(m_0 - c_0) + 3) \right]^{-\rho} (1-p) + p \exp(\theta) \left[G_1^{-1}(m_0 - c_0) + 1 \right]^{-\rho} \right], \end{aligned}$$

with no general analytical solution for c_0 . To complete arguments, I solve for the optimal consumption numerically using the EGM proposed by [Carroll \(2006\)](#), and use that solution, denoted $c_0^*(m_0)|_{c_1=m_1}^{\text{num}}$. In turn, optimal period zero consumption is given by

$$c_0^*(m_0) = \min \left\{ m_0, c_0^*(m_0)|_{c_1=m_1}^{\text{num}}, \frac{m_0 + 3G_1}{1 + [p(\exp(\rho^{-1}\theta) + 2)^\rho + (1-p)3^\rho]^{\frac{1}{\rho}}} \right\}. \quad (\text{A.7})$$

Figure A.1a reports the consumption function in the deterministic case for the baseline parameters used herein ($p = 0.5$, $\rho = 2$, and $\theta = 0.5$) and Figure A.1b reports the consumption function in the probabilistic case. The numerical solutions to both models are reported to complete the solution and confirm the analytical results.

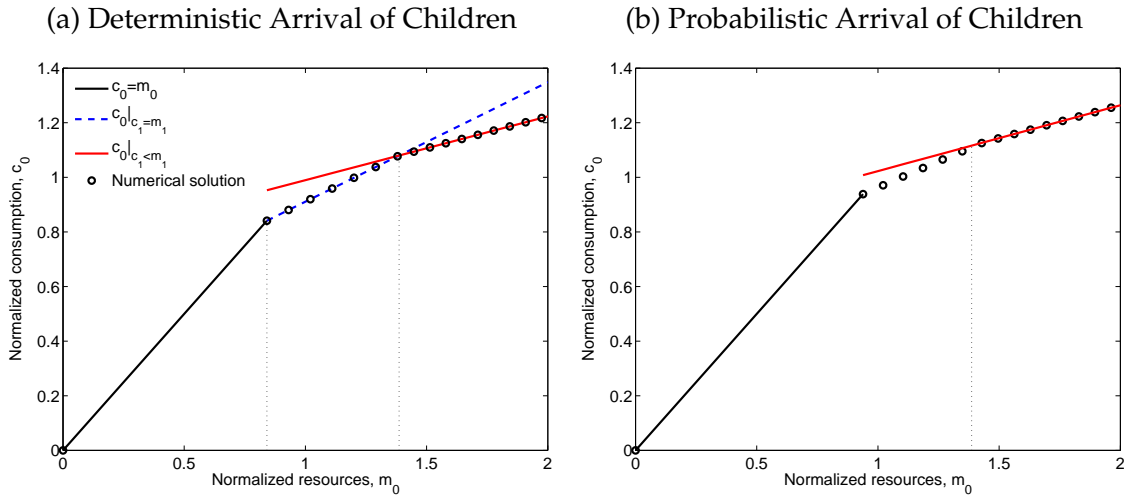


Figure A.1 – Period Zero Optimal Consumption Functions.

A.3 OLS and IV Estimates from the Four Period Model

We have that optimal consumption is given by

$$\begin{aligned} c_0^*(m_0|\mathbf{z}_1) &= \begin{cases} m_0 & \text{if } m_0 \leq \exp(-\rho^{-1}\theta\mathbf{z}_1)G_1 \\ \frac{m_0+G_1}{1+\exp(\rho^{-1}\theta\mathbf{z}_1)} & \text{else} \end{cases} \\ c_1^*(m_1|\mathbf{z}_1) &= \begin{cases} m_1 & \text{if } m_1 \leq \exp(\rho^{-1}\theta\mathbf{z}_1) \\ \frac{m_1+2}{1+2\exp(-\rho^{-1}\theta\mathbf{z}_1)} & \text{else} \end{cases} \\ c_2^*(m_2) &= \frac{1}{2}(m_2 + 1) \\ c_3^*(m_3) &= m_3. \end{aligned}$$

The OLS estimator is given as

$$\theta_{OLS}^t = (\Delta \log C_t|_{\mathbf{z}_1=1} - \Delta \log C_t|_{\mathbf{z}_1=0})\rho,$$

while the IV estimator, using the (cohort) average number of children as instrument, $Z = p$, is

$$\begin{aligned} \theta_{IV}^t &= \frac{\mathbb{E}[\Delta \log C_t' p]}{\mathbb{E}[p' p]}\rho, \\ &= \frac{1}{p}(p\Delta \log C_t|_{\mathbf{z}_1=1} + (1-p)\Delta \log C_t|_{\mathbf{z}_1=0})\rho. \end{aligned}$$

Inserting the optimal consumption for given set of parameters. Let $m_0 > \exp(-\rho^{-1}\theta\mathbf{z}_1)G_1$ (saves in period zero) and note that $m_0 = 1$, such that this implies that $\theta > \log(G_1)\rho$. The growth in log consumption is then (using the result that consumption is, then, constrained in period one)

$$\begin{aligned} \theta_{OLS}^{young} &= \begin{cases} \theta - \log(G_1)\rho & \text{if } \theta > \log(G_1)\rho \\ 0 & \text{if } 0 \leq \theta \leq \log(G_1)\rho \end{cases} \\ &\leq \theta \end{aligned}$$

Hence, OLS estimates will under estimate the true effect of children on consumption. The IV estimator is

$$\begin{aligned} \theta_{IV}^{young} &= \begin{cases} \theta + (1-p)/p \log(G_1)\rho & \text{if } \theta > \log(G_1)\rho \\ \log G_1\rho/p & \text{if } 0 \leq \theta \leq \log(G_1)\rho, \end{cases} \\ &\geq \theta \end{aligned}$$

such that IV estimation over-estimates the effect. However, as θ increases - for a fixed p and G_1 - the over estimation becomes potentially small.

Turning to older households, when children leave, the OLS estimate is

$$\theta_{OLS}^{old} = \begin{cases} \rho \log\left(\frac{1+G_1}{G_1}\right) - \rho \log(1 + \exp(-\rho^{-1}\theta \mathbf{z}_1)) & \text{if } \theta > \log(G_1)\rho \\ 0 & \text{if } 0 \leq \theta \leq \log(G_1)\rho, \\ \leq \theta \end{cases}$$

such that only if $\theta = 0$ will OLS produce a consistent estimate. Since consumption does not change between period one and two if there was no child in period one, the IV estimator is identical to OLS,

$$\theta_{IV}^{old} = \theta_{OLS}^{old}.$$

A.4 Upwards Bias of IV using Young Households Without Credit Constraints

Here, I show that in the model where children arrive probabilistically and there is *no* explicit credit constraints, there is still a (small) positive bias from IV estimation. Inserting optimal consumption in absence of credit constraints,

$$C_0 = Y_0 \frac{m_0 + 3G_1}{1 + [p(\exp(\rho^{-1}\theta) + 2)^\rho + (1-p)3^\rho]^{\frac{1}{\rho}}}, \quad C_1(\mathbf{z}_1) = Y_1 \frac{m_1 + 2}{1 + 2\exp(-\rho^{-1}\theta \mathbf{z}_1)},$$

into the IV estimator yields

$$\begin{aligned} \hat{\theta}_{IV}^{young} &= \log\left(Y_1 \frac{m_1 + 2}{1 + 2\exp(-\rho^{-1}\theta)}\right) - \log\left(Y_0 \frac{m_0 + 3G_1}{1 + [p(\exp(\rho^{-1}\theta) + 2)^\rho + (1-p)3^\rho]^{\frac{1}{\rho}}}\right) \\ &+ (1-p)/p \left(\log\left(Y_1 \frac{1}{3}(m_1 + 2)\right) - \log\left(Y_0 \frac{m_0 + 3G_1}{1 + [p(\exp(\rho^{-1}\theta) + 2)^\rho + (1-p)3^\rho]^{\frac{1}{\rho}}}\right) \right) \\ &= \rho^{-1}\theta - (\log(\exp(\rho^{-1}\theta) + 2) + (1-p)/p \log(3)) \\ &+ p^{-1} \left[\log\left(\frac{m_0 + 3G_1 - c_0}{m_0 + 3G_1}\right) + \log\left(1 + [p(\exp(\rho^{-1}\theta) + 2)^\rho + (1-p)3^\rho]^{\frac{1}{\rho}}\right) \right], \end{aligned}$$

where inserting again c_0 from equation (A.6) and re-arranging finally gives the IV estimator as

$$\begin{aligned} \hat{\theta}_{IV}^{young} &= \theta + \omega\rho, \\ &\geq \theta, \end{aligned}$$

where

$$\omega \equiv p^{-1} \left[\rho^{-1} \log \left(p \left(\exp(\rho^{-1}\theta) + 2 \right)^\rho + (1-p)3^\rho \right) - (p \log \left(\exp(\rho^{-1}\theta) + 2 \right) + (1-p) \log(3)) \right] \geq 0$$

such that defining $\omega_1 \equiv (\exp(\rho^{-1}\theta) + 2)^\rho$ and $\omega_2 \equiv 3^\rho$, the bias of the IV estimator can be seen to be the difference in the log-expected value and the expected log value;

$$\omega = p^{-1} \rho^{-1} (\log(p\omega_1 + (1-p)\omega_2) - (p \log \omega_1 + (1-p) \log \omega_2)),$$

which is always positive since the logarithm is a concave function.

B Solving the Life Cycle 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 $\hat{a}_t = (a_t, a_t^1, \dots, a_t^{Q-1})$ in which 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^* , is given by $m_t = \hat{a}_t + c_t^*(m_t, \mathbf{z}_t)$.

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

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

where $\check{c}_{t+1}^*(m_{t+1}, \mathbf{z}_{t+1})$ 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 lev-

els of wealth to approximate accurately the curvature of the consumption function. The number of points was chosen such that the change in 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, $c_{t+1}^* = m_{t+1}$ if m_{t+1} is lower than some threshold level, m_{t+1}^* . 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., $m_{t+1}^* = \underline{m}_{t+1}$. 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 2.

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 ε ,

and η , respectively, Ω_t can be found recursively as

$$\Omega_t = \begin{cases} R^{-1}G_{T_r}\varepsilon_{T_r}\eta_{T_r} & \text{if } t = T_r - 2, \\ R^{-1}(\min\{\Omega_{t+1}, \kappa\} + \varepsilon_{t+1})G_{t+1}\eta_{t+1} & \text{if } t < T_r - 2. \end{cases}$$

Proof. To see this, 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,

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

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,

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

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 $\underline{a}_t = -\min\{\Omega_t, \kappa\}$. \square

Chapter 3

Life-Cycle Consumption and Children

Life-Cycle Consumption and Children

Thomas H. Jørgensen[†]

October 30, 2014

Abstract

I study the effect children might have on marginal utility of consumption. To allow households to be simultaneously affected by children and other life cycle motives such as income uncertainty, credit constraints and retirement, I fully specify the economic environment households face when performing their consumption and savings decisions. Most existing studies ignore alternative life cycle motives when analyzing to which extent children, and household demographics in general, affect consumption behavior. I estimate the model using the Panel Study of Income Dynamics (PSID) for the US and high quality Danish administrative register data. Contrary to most existing studies, the results suggest that children do not affect non-durable consumption as much as previously assumed. In turn, precautionary motives seem more important than children in explaining non-durable consumption over the life cycle (JEL: D12, D14, D91).

Keywords: Consumption, Children, Precautionary Savings, 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.

I am grateful to Mette Ejrnæs, Bertel Schjerning, Jeppe Druedahl, Martin Browning, Christopher Carroll, John Rust, Richard Blundell, Orazio Attanasio, Hamish Low, Søren Leth-Petersen, Miles Kimball, Dennis Kristensen, Daniel Borowczyk-Martins, Kjetil Storesletten, Jérôme Adda, Anders Munk-Nielsen, Mette Foged, Damoun Ashournia 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. 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 number of children inside households and household consumption share similar hump-shaped age profiles. The extent to which children affect consumption behavior has, therefore, received great attention over the last two decades with great effects of children being 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 labor supply with very different policy implications.² Despite a significant body of literature analyze the effect of demographics on consumption, the relative importance of demographics and alternative life cycle motives is still unresolved.

The present study offers new insights into this literature. I estimate the effect of children on consumption while allowing households to also be affected by other alternative life cycle motives such as income uncertainty, credit constraints and retirement. A key difference from existing studies is that precautionary motives are typically excluded when the effect of children on consumption is analyzed.³ To allow for several life cycle motives simultaneously, I fully specify the economic environment households face when performing their consumption and savings decisions. I estimate the effect of children on consumption using both the Panel Study of Income Dynamics (PSID) and high quality Danish administrative register data to investigate how robust the results are across different data sources. Furthermore, the Danish registers allow for identification of childless households around completed fertility where surveys, such as the PSID, do not allow for identification of childless women.

I find that the effect of children on non-durable consumption is lower than reported in most existing studies. For Danish households, the estimated effect of children on consumption is economically negligible (or even negative) while, for the PSID sample, the estimated effect of children on consumption is close to (but less than) that reported in [Attanasio, Banks, Meghir and Weber \(1999\)](#). The results are robust in comparison with alternative model specifications and assumptions regarding measurement error

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 \(1994\)](#); [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 can explain the hump in consumption 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.

in consumption. The results suggest that precautionary motives are more important than children in explaining non-durable consumption over the life cycle. Interestingly, the age profile of non-durable consumption of Danish childless households is identical to households who have children. This is new evidence that children cannot be the primary driver of the observed hump-shaped consumption age profile.

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). In Attanasio, Banks, Meghir and Weber (1999), the number of children is found to be important in order to describe the consumption behavior of US consumers. Using the repeated cross section information on non-durable consumption in the Consumer Expenditure Survey (CEX), they 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 resulting estimates are interpreted in light of a life cycle model of intertemporal consumption behavior very similar to the one analyzed in the present study. Using the estimated parameters, they simulate data from alternative models with and without income uncertainty and demographics. Based on these age profiles, Attanasio, Banks, Meghir and Weber (1999) conclude that household demographics are important in order to fit the observed consumption age profile.

The first-stage estimation of demographic effects on consumption in Attanasio, Banks, Meghir and Weber (1999) likely also removes variation in consumption stemming from precautionary behavior. This is because changing household demographics coincide with high income growth and the inability or unwillingness to borrow against future income (Jørgensen, 2014). Furthermore, the effects of demographics are implemented through an adjustment of the discount factor over the life cycle, which *all* households in their model face irrespectively of the actual demographic composition. In turn, this approach is not well suited to analyze the *relative importance of demographics* and precautionary motives.

The importance of children in explaining the life cycle profile of non-durable consumption is supported by earlier results in Attanasio and Browning (1995). Using the UK Family Expenditure Survey (FES), they also estimated log-linearized Euler equations on a synthetic cohort panel. Fernández-Villaverde and Krueger (2007) compare equivalence scaled consumption profiles and argue that around 50 percent of the hump in the consumption profile in the CEX is due to household demographics. Browning and Ejrnæs (2009) find that the number and age of children can completely explain the hump in consumption by estimating a more flexible functional form of the log-linearized Euler equation on a synthetic cohort panel, constructed from the FES.

Most existing evidence of the effect of children on consumption excludes alternative life cycle motives when estimating log-linearized Euler equations. As shown in

Jørgensen (2014), if households face sufficiently strong precautionary motives (credit constraints or a probability of zero-income) Euler equation estimators overestimate the effect of children on consumption if the effect of children is relatively low. Existing work might, therefore, suffer from an “omitted variable bias” if households face precautionary motives. The number of children (and adults) might proxy for these motives, explaining why I find that children have much lesser effect compared to existing studies when allowing also for precautionary motives. Hence, I interpret the results as indicating that precautionary motives are more important than children in explaining non-durable consumption over the life cycle.

The present results are related to a growing body of literature investigating the importance of precautionary motives. Gourinchas and Parker (2002) and Cagetti (2003) find that precautionary motives are important even after controlling for household demographics. An important caveat in those studies, however, is the adjustment of consumption to household demographic effects *prior* to estimation of models with precautionary savings motives. Gourinchas and Parker (2002) adjust the consumption measure for demographic effects via regression before calculating empirical moments to be matched by simulated data from their life cycle model. Cagetti (2003) uses the estimated demographic effects from Attanasio, Banks, Meghir and Weber (1999) to adjust the discount rate over the life cycle.

The richness of the present model framework, in which the age and number of children may affect household preferences, has previously precluded structural estimation of the effect of children on consumption.⁴ However, due to the Endogenous Grid Method (EGM), proposed by Carroll (2006), I am able to formulate and estimate a standard buffer-stock model in which households can be affected by the presence of children. The parameters are estimated using a continuous version of the Nested Fixed Point (NFXP) estimator proposed by Rust (1987) and found to perform well in Jørgensen (2013). The general estimation framework facilitates estimation of parameters under different assumptions about measurement error in consumption. Specifically, I allow the measurement error in the PSID to be multiplicative, heterogeneous across households and arbitrarily distributed.

The rest of the paper proceeds as follows. The next section describes the Danish administrative registers and the PSID used throughout the study. Section 3 formulates a standard life cycle model of consumption and savings in the presence of children and section 4 outlines the estimation strategy and discusses the calibration of some of the model parameters. Section 5 presents the estimation results and section 6 investigates the robustness of these results. Finally, Section 7 concludes.

⁴By structural estimation I refer to estimation routines in which the underlying economic environment is fully specified. Typically, this involves numerically solving for optimal consumption/savings behavior for a given set of trial values of parameters.

2 Danish Register Data and the PSID

2.1 The Danish Registers

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 care subsidies 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 amounts held in specific stocks are not known, only the total value of all stocks is. The house value is assessed by the tax authorities for tax purposes and is included because it is impossible in the Danish registers to determine exactly which mortgages are related to the house and which are not.

Pension wealth is not included in the wealth measure. Information on pension accounts are not available for most of the cohorts studied here and the resulting net wealth is, therefore, slightly underestimated. 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 while withdrawing pension funds after retirement only incur a tax of 40 percent. 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 continuously married and 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.⁵ Only households with children born when the wife was aged 15 through 43

⁵This 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 and, hence, make estimation of the model less tractable.

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. Danish 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 has at least a bachelor degree.

2.2 The PSID

The Panel Study of Income Dynamics (PSID) contains information on food consumption in and out of home. The PSID has been used intensively for several purposes, including estimation of the effect of children on consumption. To study the evolution and link between income and consumption inequality over the 1980s, [Blundell, Pistaferri and Preston \(2008\)](#) impute total non-durable consumption for PSID households using food consumption measures in the CEX. I use their data and measure of total consumption. I augment the PSID sample from [Blundell, Pistaferri and Preston \(2008\)](#) with information on the age of children inside the household through the age of birth of each child in the PSID.

The sample period is 1978-1992 and only male headed continuously married couples are used. The years 1987 and 1988 are not used since consumption measures were not collected those years. Since the present study focus on the effect of children on household consumption, and I want the sample to be comparable to the Danish data, I restrict the sample to cover 25 to 59 year old households and link this to the age of the wife.⁶ Further, as for the Danish data, I exclude households in which the age difference between husband and wife is larger than four years and more than three children is present. I also drop few households who have children before age 15 or later than 43. In turn, an unbalanced panel of 2,350 households observed for at most 13 periods are in the final sample of a total of 17,005 non-missing observations.

Household resources are composed of after tax household income of both spouses and household assets. The PSID does not contain information on liabilities and, therefore, I do not include house value in the measure of assets. The measure of resources in the PSID, therefore, differs from the net-wealth measure constructed for the Danish households but should still provide a good measure of consumption possibilities within a household. Households are classified as high skilled if the male head has ever enrolled in college, including college drop-outs. Assets, income and consumption are converted into 1980 US prices through regression.

⁶[Blundell, Pistaferri and Preston \(2008\)](#) use households in which the husband is aged 30 to 65.

2.3 Empirical Life Cycle Patterns

The first two columns in Table 1 present means and standard deviations for the Danish data while the third and fourth columns present statistics for the PSID sample. The data stem from two different countries in different years which should be kept in mind when comparing statistics across the two sources. Danish households are older with an average age of Danish wives of 40.7 compared to an average age of wives in the PSID of 38.5. The husband is around one year older than the wife in both samples. Despite the higher average age of Danish households, they have slightly fewer children (1.35 on average) than the US households (1.4 on average).

Table 1 – Descriptive Statistics.

	Danish registers		PSID	
	Mean	Std	Mean	Std
Age, wife	40.749	8.213	38.484	9.415
Age, husband	41.978	8.295	39.171	9.678
Wealth [†]	38,967	49,635	1,993	10,823
Non-durable consumption [‡]	34,713	18,674	11,903	20,080
Disposable income	36,166	6,712	26,874	24,502
Number of Children	1.346	0.922	1.432	1.039
High skilled	0.336	0.472	0.511	0.500
Number of observations	1,281,952		18,376	

Notes: 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.

[†] Wealth in the PSID consists only of assets while wealth in the Danish registers contain stocks, bonds, bank deposits, cars, boats, house value for home owners and mortgage deeds, net of total liabilities.

[‡] Non-durable consumption in the PSID is imputed by [Blundell, Pistaferri and Preston \(2008\)](#) based on food consumption in the CEX and the PSID. Non-durable consumption in the Danish data is imputed as the sum of disposable income net of changes in the wealth stock, as proposed by [Browning and Leth-Petersen \(2003\)](#).

The imputed consumption measure for Danish households is around 20,000 USD higher than that imputed for PSID households. Half of this difference can be explained by disposable household income being around 10,000 USD higher for Danish couples. Importantly, the Danish measures are in 1987 values with a fixed exchange rate of 5.5 while the PSID measures are from 1980. An annual inflation rate of 3 percent would imply an average disposable household income for PSID households of around 33,000 USD in 1987, very close to that of Danish households (36,166 in 1987 USD). The wealth measures cannot be compared across countries because net-worth and assets are included in the wealth measure for Danish and US households, respectively.

Figure 1 illustrates the life cycle profiles of income, non-durable consumption and the number of children for Danish and US households. Children seem to arrive and leave with very similar rates in the two countries. Income grows significantly more in the early part life of Danish households compared to that of households in the PSID. The same is true for the imputed non-durable consumption. Income and consumption profiles for Danish households peak around the mid-40s while the peak is slightly later for US households around age 50, in line with previous findings (see, e.g., [Attanasio and Browning, 1995](#) for the UK and [Gourinchas and Parker, 2002](#) for the US).

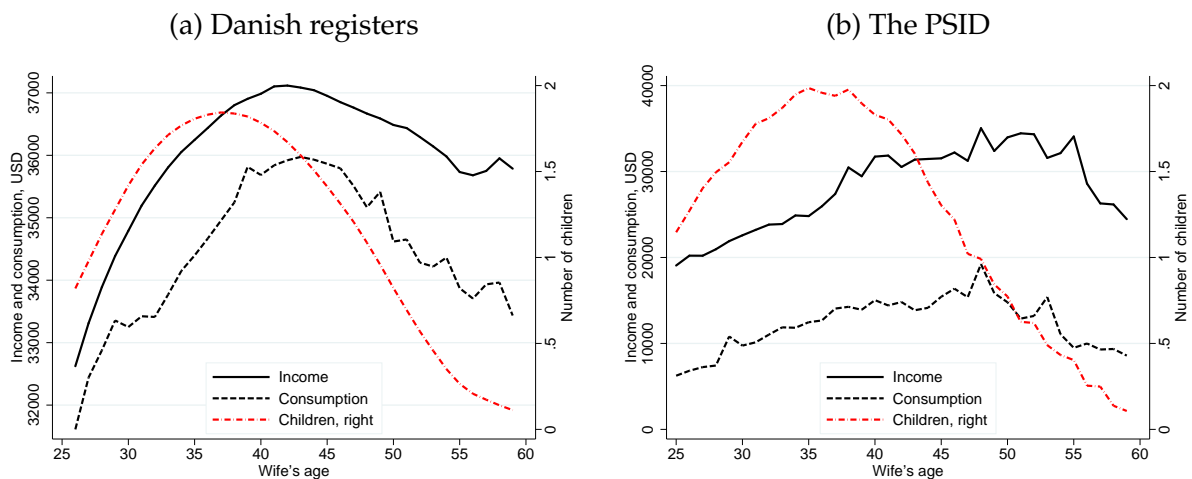


Figure 1 – Age Profiles of Income, Consumption and the Number of Children.

Notes: Figure 1 illustrates average age profiles of income, non-durable consumption and number of children (on the right y-axis) for the Danish and US 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 in the PSID is imputed by [Blundell, Pistaferri and Preston \(2008\)](#) based on food consumption in the CEX and the PSID. Non-durable consumption in the Danish data is imputed as the sum of disposable income net of changes in the wealth stock, as proposed by [Browning and Leth-Petersen \(2003\)](#).

Figure 2 illustrates wealth profiles of households who have children at the age of 30 and households who are childless at the age of 30. The overall wealth profile is monotonically increasing, at odds with children having a large effect on consumption, when risk averse households face sufficient precautionary motives ([Jørgensen, 2014](#)). If risk averse households are constrained either from an explicit credit constraint or from a self-induced borrowing constraint ([Schechtman, 1976](#); [Zeldes, 1989](#); [Carroll, 1997](#)), they will accumulate wealth in anticipation of the arrival of children in the future and almost deplete that wealth when they subsequently have children. In turn, this behavior produces a hump-shaped wealth age profile. The empirical wealth profiles in Figure 2 do not, however, show such a hump-shape suggesting that children might not be the primary explanation for the observed consumption age profile.

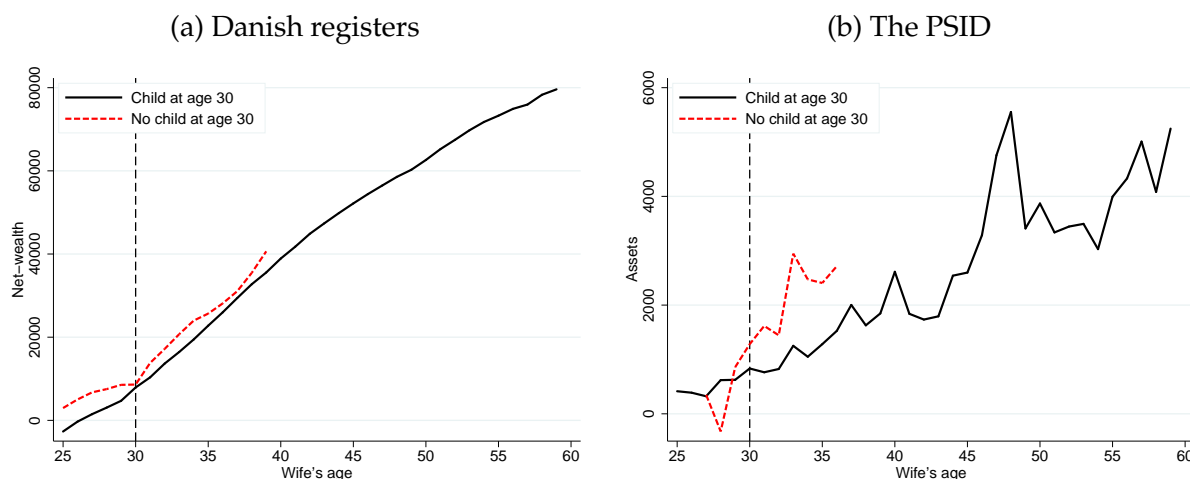


Figure 2 – Age Profiles of Wealth with and without Children at Age 30.

Notes: Figure 2 illustrates average age profiles of wealth for households with and without children at age 30, for the Danish and US 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. For the PSID, only assets (excluding house value) are used and cells with less than 60 observations are excluded. The Danish registers contain information on total net-wealth (including house value).

Households in the PSID who are childless at age 30 have slightly higher wealth levels when older than 30, compared to those households who have children at age 30. Assuming that these households remain childless for a few additional years, this suggests that children might increase marginal utility of consumption. If the presence of children increases the marginal utility of consumption, wealth would be expected to be accumulated prior to the arrival of children to smooth marginal utility of consumption across periods without children and periods with children.⁷ Figure 2 shows almost no evidence of such a pattern. However, households who expect to be childless throughout their entire life are included in the group of households who do not have children at age 30. Such households, who plan to remain childless, have no incentive to save for the arrival of children and might reduce the difference between the two groups in Figure 2.

The Danish registers allow for identification of childless households at completed fertility. Figure 3 presents consumption and disposable income profiles for Danish 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 2010.⁸ If the wife is not registered as a mother in 2010,

⁷This argument hinges on imperfect capital markets or households' desire to reduce borrowing due to precautionary motives. If households could borrow infinitely, and finds it optimal to do so, wealth accumulation would not be affected and consumption could respond perfectly to changing demographic composition within the household.

⁸Virtually *all* childbirths after 1942 are matched to their mother. Only children born between January

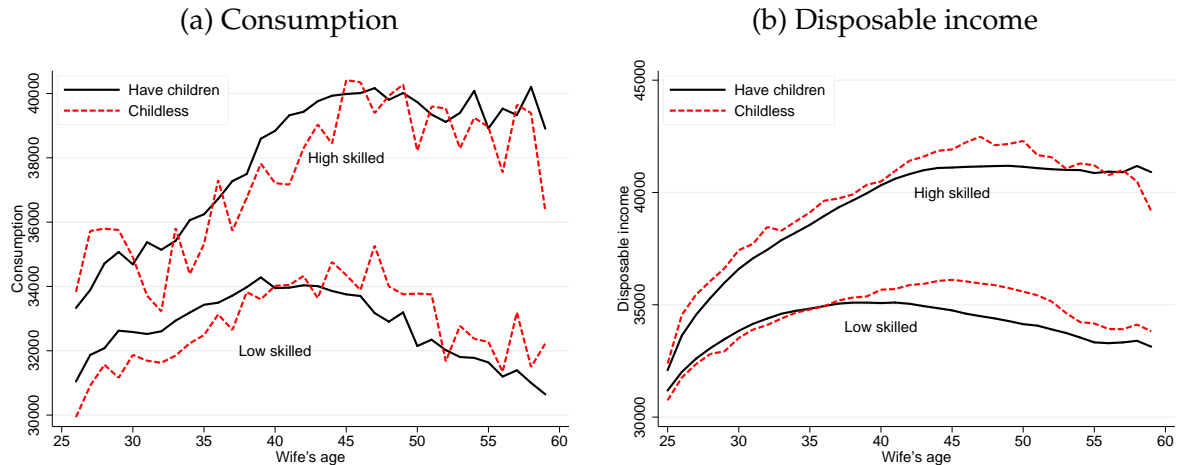


Figure 3 – Consumption and Disposable Income in Childless Households and Households with Children at Completed Fertility, Danish Registers.

Notes: Figure 3 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 3 assumes that no children will arrive in that household.

I assume that the household will remain childless. This assumption is not overly restrictive since the youngest household in the data (aged 26 in 1996) will be 40 years old in 2010. Only few households have children at this age. Childless women could have adopted children or foster children from the current husband's previous marriage(s). Therefore, I restrict childless households to those without children registered as living at the same address as the couple at any point in the observed years.⁹

Children do not explain the observed hump shape in consumption. 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 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 con-

1st 1942 and December 31st 1972 who either died or permanently immigrated to another country before 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.

⁹This does, however, not rule out the possibility that households defined as childless do foster children registered to be living at another address. Further, I do not know whether childlessness is caused by infertility or an active choice. Infertility of a household who otherwise would have had children, likely produce consumption and savings behavior similar to households who subsequently have children, prior to the infertile household learns about this infertility. This will tend to produce similar consumption age profiles for households who have children and childless households.

sumption (Attanasio and Browning, 1995; Browning and Ejrnæs, 2009) might proxy for precautionary motives.

3 A Model of Consumption in the Presence of Children

The framework used throughout this study is a version of the buffer-stock model pioneered by Deaton (1991) and Carroll (1992). Attanasio, Banks, Meghir and Weber (1999); Gourinchas and Parker (2002) and Cagetti (2003) employ models very similar to the present model and the framework is identical to that in Jørgensen (2014).

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 solve the optimization problem

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

where utility is CRRA, $u(C_t) = C_t^{1-\rho} / (1-\rho)$, and $v(\mathbf{z}_t; \theta)$ is a taste shifter where 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 through a multiplicative $v(\mathbf{z}_t; \theta)$. 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.

Following Gourinchas and Parker (2002), survival and income uncertainty are omitted post retirement and the parameter γ (referred to as the retirement motive) in equation (1) is a parsimonious way of adjusting for these elements.¹⁰ Gourinchas and Parker (2002) ignore the post-retirement consumption decisions and adjust the perfect foresight approximation by a parameter similar to γ through a retirement value function. Although I focus on consumption behavior *prior* to retirement, the potential presence of children at retirement forces the model to be specific about post retirement behavior as well. Specifically, ignoring the presence of children in retirement (as done in Gourinchas and Parker, 2002) would lead the model to underestimate optimal consumption in periods prior to retirement.¹¹ Since the focus of this study is on estimation of θ , it is essential that aspects related to children are properly handled.

¹⁰Survival is also certain prior to retirement.

¹¹This is because too much is saved in the last working period if the decrease in marginal value of consumption in the future, when a child moves, is ignored. Households who know that, while they are retired, a child will move, will incorporate the associated drop in household consumption already before retiring since less wealth is required to maintain a given level of consumption while retired.

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. While working, income is assumed to follow the stochastic process

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

where P_t denotes permanent income, G_t is real gross income growth, $\eta_t \sim \log \mathcal{N}(-\sigma_\eta^2/2, \sigma_\eta^2)$ 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 distributed $(1 - \wp)\varepsilon_t \sim \log \mathcal{N}(-\sigma_\varepsilon^2/2 - \mu\wp, \sigma_\varepsilon^2)$. When retired, the income process is a deterministic fraction $\varkappa \leq 1$ of permanent income and permanent income grows with a constant rate of G_{ret} once retired, $Y_t = \varkappa P_t, \forall t \geq T_r$. and $P_t = G_{ret} P_{t-1}, \forall t \geq T_r$.

In each period, households face the intratemporal budget constraint, $M_t = A_t + C_t$, such that end-of-period wealth, A_t , and consumption must equal the available resources in the beginning of the period. Further, net wealth must be greater than a fraction of permanent income in all time periods,

$$A_t \geq -\kappa P_t \forall t, \kappa \geq 0,$$

and, as in [Gourinchas and Parker \(2002\)](#), retired households are not allowed to be net borrowers, $A_t \geq 0, \forall t \geq T_r$.

3.1 Household Composition

The evolution of children, \mathbf{z}_t , is typically ignored. However, since θ is of primary interest in this study, I will be precise about the underlying process regarding the arrival and leaving of children. Most existing studies assume that the taste shifter is an exponential function in demographic variables, $v(\mathbf{z}_t; \theta) = \exp(\theta' \mathbf{z}_t)$. I will follow this functional form specification but also estimate a more flexible functional form for $v(\mathbf{z}_t; \theta)$ allowing for an arbitrary children, scale and age effect.

Individuals do not die, divorce or remarry such that households consist of the same husband and wife at all times. Households can have at most three children and once the wife is 43 years old the household cannot have any more children. The age of each

child is contained in \mathbf{z}_t ,

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

where “NC” refers to “No Child” and the oldest child is denoted child one, the second oldest child as child two and the third oldest child as child three.

A novelty of this study is that I keep track of the age of *all* children inside the household. To the best of my knowledge, this has not previously been done in dynamic models of intertemporal behavior. The recursive nature of the present model has precluded such a rich characterization of changes in household composition (Hotz and Miller, 1988). To circumvent the computational cost of keeping track of the age of all children, strict assumptions on the timing of children are typically imposed.¹² I circumvent the computational cost by solving the model by the EGM (Carroll, 2006). The EGM solves consumption models extremely fast and accurately (Jørgensen, 2013).

Knowing the age of each child is important to investigate how the number of children within different age groups might affect consumption. Browning and Ejrnaes (2009) show that the age of children have profound impact on consumption. Specifically, they find that younger children have a small negative effect while older children have a significant positive effect on consumption.

Following, e.g., Love (2010); Hong and Ríos-Rull (2012) and Blundell, Dias, Meghir and Shaw (2013), children arrive probabilistically with a known probability distribution depending on the age of the wife, educational attainment, and the number of children already present in the household. Children leave home at age 21 and do not influence household consumption in subsequent periods.

Alternatively, children could arrive deterministically with perfect foresight, as in Browning and Ejrnaes (2009). In the robustness analysis, I estimate a version of the model in which households know with perfect foresight how many children they will have and when these children arrive. The results are unchanged. Instead of exogenous arrival of children, the fertility choice could be endogenous in the model (Scholz and Seshadri, 2009; Sommer, 2014). It is not obvious that models of perfect contraception control are more realistic than probabilistic arrival of children (David and Mroz, 1989).¹³ For computational tractability, I have not pursued that here because solving such a model (with both continuous and discrete choices) typically involves computational methods (e.g., value function iterations) which are slower than the EGM applied

¹²For example, 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) 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. In her model of endogenous fertility, children leaves home probabilistically. Alternatively, Blundell, Dias, Meghir and Shaw (2013) assumes that only the youngest child matters.

¹³See Hotz, Klerman and Willis (1997) for a survey on the economics of fertility.

in the present study.¹⁴

In the baseline probabilistic case, not only households who have children are affected by the parameter θ . Households dynamically optimize their consumption behavior while incorporating expectations about the future. Thus, *all* younger households within the same age group who have no children will want to reduce their consumption today in anticipation of increased consumption when children might arrive in the future. In the deterministic case, the savings rates differ across households within age groups due to differences in when and how many children arrive over the life cycle. It is not obvious which is the most appropriate assumption (probabilistic or deterministic arrival on children) and the probabilistic version has been chosen as baseline since that model does not require knowledge on completed fertility, when estimating model parameters.

4 Estimation Strategy

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). The estimator is not new but the exposition here is a general framework in which different assumptions about measurement error in observed consumption measures easily can be used to derive a consistent estimator.

For a given set of K structural parameters, Θ , the model is solved recursively for all combinations of household composition. This yields optimal consumption as a function of resources, permanent income and household composition, $\{C_t^*(M_t, P_t, \mathbf{z}_t | \Theta)\}_1^T$.¹⁵ In principle, Mathematical Programming with Equilibrium Constraints (MPEC), proposed by Su and Judd (2012) could be used to estimate parameters. However, as discussed in Jørgensen (2013), 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.

Let $\mathcal{O} = (M, P, C, \mathbf{z})^{data} \in \mathcal{O}$ denote *observed* information where $\mathcal{O} \subset \mathbb{R}^{\dim(\mathcal{O})}$ and \mathcal{O}_{it} refers to household i in period t and \mathbb{R} are the real numbers. Define a function of the observed data and model solution, for a given set of parameters, as

$$\zeta_{it}(\Theta) \equiv \zeta(\mathcal{O}_{it}, C_t^* | \Theta),$$

in which observed data is used to infer the model predictions for each household-time observation. In a typical static linear model, $\zeta_{it}(\Theta) = y_{it} - \Theta' x_{it}$, would correspond to the residuals from the linear index ($\Theta' x_{it}$) of an observed outcome, y_{it} . Generally,

¹⁴Fella (2014) and Iskhakov, Jørgensen, Rust and Schjerning (2014) generalize the EGM to handle both discrete and continuous choices.

¹⁵Consult Appendix A for details on the solution method applied to solve the model described in Section 3.

for a given observation, \mathcal{O}_{it} , the associated prediction from the structural model can be found by interpolating the relevant policy function, referred to as $\check{C}^*(\mathcal{O}_{it}|\Theta)$ in examples below. Let

$$g_i(\Theta, \phi) \equiv g(\check{\xi}_i(\Theta), \phi),$$

denote a real-valued function taking as argument stacked time observations, $\check{\xi}_i(\Theta)$, in which ϕ contain K_ϕ additional parameters. All parameters are in a compact space, $(\Theta, \phi) \in \mathbb{C} \subset \mathbb{R}^{K+K_\phi}$, and $g : \mathbb{O} \times \mathbb{C} \rightarrow \mathbb{R}$ is, for all $\mathcal{O}_{it} \in \mathbb{O}$, continuous in (Θ, ϕ) .¹⁶

The proposed estimator solves the problem

$$\min_{(\Theta, \phi) \in \mathbb{C}} N^{-1} \sum_{i=1}^N g_i(\Theta, \phi), \quad (2)$$

and, assuming an *unique* solution exists, is consistent by the uniform weak law of large numbers. Under standard regularity conditions, this M-estimator is asymptotically normally distributed around the true parameter with asymptotic variance of $\mathbf{A}^{-1}\mathbf{B}\mathbf{A}^{-1}/N$, where $\mathbf{A} \equiv \mathbb{E}[\mathbf{H}(\Theta, \phi)]$ and $\mathbf{B} \equiv \mathbb{E}[\mathbf{s}(\Theta, \phi)'\mathbf{s}(\Theta, \phi)]$ (Wooldridge, 2002).

Optimal behavior is in general found numerically and the objective function in (2) is an *approximation* to the exact objective function. Fernández-Villaverde, Rubio-Ramírez and Santos (2006) show, in a likelihood framework, 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 (2) are consistent and asymptotically normally distributed even when the solution, C^* , is found numerically.¹⁷

To illustrate the flexibility of the estimator, I present concrete examples in which assumptions often invoked in the literature are implemented in the framework above. Example 1 illustrates how the estimator can estimate structural parameters if consumption is contaminated with additive normally distributed measurement error. Readers who feel uncomfortable with the normality assumption in Example 1 could think of the estimation problem as one of non-linear least squares. Alternative distributional assumptions could be implemented, or the absolute difference could be minimized ($g_i(\Theta) = \sum_t^T |\check{\xi}_{it}(\Theta)|$), yielding an estimator more robust to outliers.

Example 1 (Additive Normal Measurement Error). If consumption data is contami-

¹⁶Also, for each $(\Theta, \phi) \in \mathbb{C}$, g should be Borel measurable on \mathbb{O} .

¹⁷Akerberg, 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. Akerberg, Geweke and Hahn (2009) reassuringly show that this is *not* the case.

nated with *iid* additive $\mathcal{N}(0, \sigma_{\xi}^2)$ measurement error, then letting

$$\begin{aligned}\xi_{it}(\Theta) &= C_{it}^{data} - \check{C}^*(\mathcal{O}_{it}|\Theta), \\ g_i(\Theta, \sigma_{\xi}) &= \frac{T_i}{2} \log(2\pi\sigma_{\xi}^2) + \frac{1}{2\sigma_{\xi}^2} \sum_t^{T_i} \xi_{it}(\Theta)^2,\end{aligned}$$

produce structural parameters that maximize the likelihood of observed data being generated from the structural model.

Jørgensen (2013) shows that an estimation approach similar to that outlined in Example 1 can uncover parameters like the relative risk aversion, ρ , from similar models. For completeness, Table 2 reports mean (and standard deviation) of estimates of θ (assuming $v(\mathbf{z}_t; \theta) = \exp(\theta \text{NumChildren})$) from 50 independent simulations in which normal measurement error is added with a known variance of one. The estimation approach uncovers the true parameter, θ_0 , in even small samples.

Table 2 – 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 (0.034)	0.115 (0.062)	0.300 (0.099)	0.504 (0.080)
$N = 1000, T = 20$	-0.000 (0.017)	0.106 (0.039)	0.307 (0.078)	0.510 (0.057)
$N = 50000, T = 5$	0.000 (0.005)	0.099 (0.010)	0.299 (0.017)	0.501 (0.013)
$N = 50000, T = 20$	0.000 (0.002)	0.100 (0.004)	0.301 (0.010)	0.502 (0.009)

Notes: Table 2 reports means and standard deviations (in parenthesis) across 50 Monte Carlo runs. Data is simulated from the model described in Section 3 with parameters set to $\beta = .95$, $\rho = 2$, $R = 1.03$, $\kappa = 0$, $\varphi = 0$, $\mu = 0$, $\gamma = 1.1$, $\varkappa = 0.8$, $G_{ret} = 1$, $\sigma_{\varepsilon}^2 = 0.005$, 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.

Consumption is sometimes assumed to be observed with *multiplicative* measurement error (Gourinchas and Parker, 2002; Alan, Attanasio and Browning, 2009). The proposed framework can easily encompass this situation by letting $\xi_{it}(\Theta) = \log C_{it}^{data} - \log \check{C}^*(\mathcal{O}_{it}|\Theta)$ and letting $g(\cdot)$ correspond to a distributional assumption. If panel data is available, the measurement error can be allowed to vary systematically across households, as illustrated in Example 2 in which the multiplicative measurement error is heterogeneous and arbitrarily distributed.

Example 2 (Multiplicative Heterogeneous Measurement Error). If consumption data is contaminated with multiplicative measurement error systematically different across

households, then

$$\begin{aligned}\tilde{\zeta}_{it}(\Theta) &= \log C_{it}^{data} - \log \check{C}^*(\mathcal{O}_{it}|\Theta), \\ g_i(\Theta) &= \sum_t^{T_i} |\Delta \tilde{\zeta}_{it}(\Theta)|,\end{aligned}$$

produce consistent estimates of Θ , independent of the distribution of the measurement error.

4.1 Initial Estimations and Calibrations

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. Below, I discuss how these parameters are calibrated.

Table 3 reports the values and sources for the calibrated parameters. When retired, US households are assumed to experience a constant decrease in permanent income of 5 percent while income of Danish retirees is constant. This does not affect the results significantly since the value of retirement, γ , will adjust accordingly. The exogenous drop in permanent income when households retire, \varkappa , is calibrated to 90 percent in Denmark based on the median couple in the study by [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. Since the pension system is less generous in the US, I calibrate this drop to be larger, 20 percent, for the US.

Table 3 – Calibrated Parameters.

	Denmark		PSID	
	Value	Source	Value	Source
G_t	Fig. 5a	Own calculations: see text	Fig. 5b	Own calculations: see text
R	1.03	Gourinchas and Parker (2002)	1.03	Gourinchas and Parker (2002)
κ	0.6	Own calculations: see text	0.00	Self imposed constraint
\wp	0.10	Own calculations: see text	0.003	Gourinchas and Parker (2002)
μ	0.30	Own calculations: see text	0.00	Gourinchas and Parker (2002)
\varkappa	0.90	Ministry of Finance (2003)	0.80	Own calculations: see text
G_{ret}	1.00	Own calculations: see text	0.95	Own calculations: see text

Notes: Table 3 reports the values of some of the calibrated parameters of the life cycle model along with the relevant sources.

The low transitory income shock is calibrated such that with 0.3 percent probability US households receive zero income ($\mu_{US} = 0$ and $\wp_{US} = .003$), following [Gourinchas and Parker \(2002\)](#). This implies that households would never want to leave

zero resources to the next period in fear of having to consume zero with a dis-utility of negative infinity (Schechtman, 1976; Zeldes, 1989; Carroll, 1992). As a consequence, the explicit credit constraint does not affect the PSID consumers and κ_{US} is set to zero. The social security system in Denmark is more compatible with a 10 percent risk of income being reduced to 30 percent. Danish households are allowed to be net-borrowers by 60 percent of annual permanent income. These three values ($\kappa_{DK} = .6$, $\mu_{DK} = 0.3$ and $\wp_{DK} = .01$) are somewhat arbitrary and have been chosen to provide a reasonable fit in the bottom distribution of resources and do not influence the results significantly. Figure 4 illustrates this by plotting within-percentile average consumption-income ratios against the household resources (also normalized by income). There is substantial variation in consumption in the bottom distribution of resources for particularly young households and the calibrated parameters (along with the estimated preferences in Table 5) provide a good fit on average.

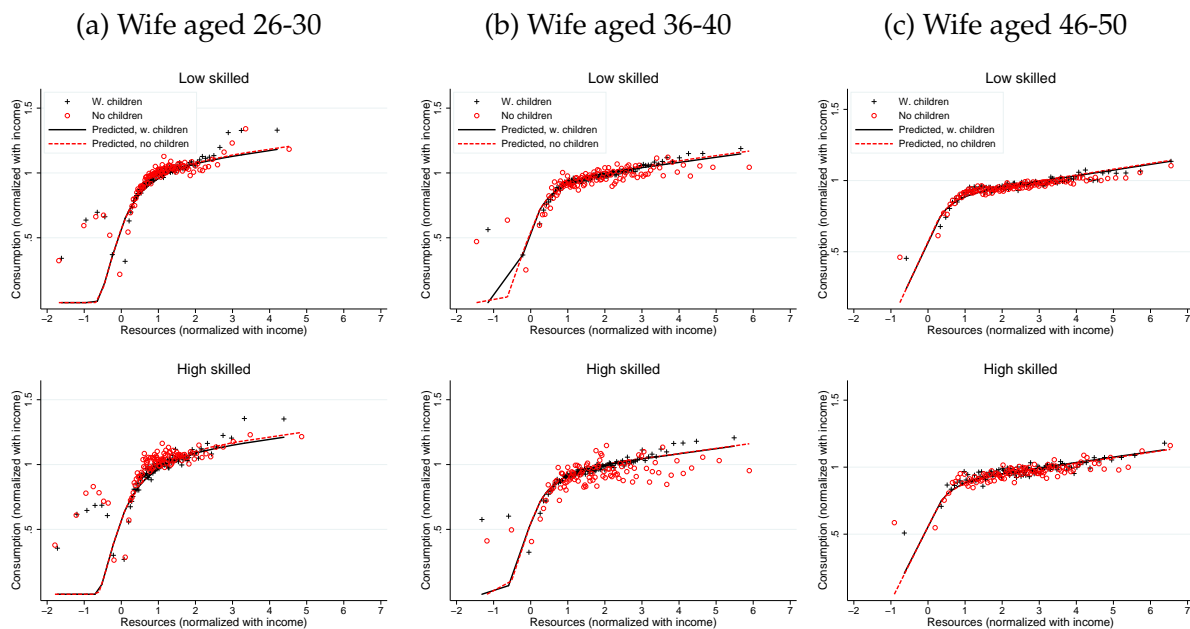


Figure 4 – Consumption Functions, Danish Data.

Notes: Figure 4 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 $\kappa_{DK} = .6$, $\mu_{DK} = 0.3$ and $\wp_{DK} = .01$ produce a good fit of the model to actual data.

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

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

Table 4 presents the estimated variance components for both data sources. 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. I find the opposite here. The permanent income shocks account for slightly more of the variation in income relative to the transitory shocks ($\hat{\sigma}_\eta > \hat{\sigma}_\varepsilon$) in the PSID while most existing studies report the opposite result.¹⁸ This is most likely due to the fact that I only remove year effects while most other studies include other “deterministic” components such as the number of children and household age (Carroll and Samwick, 1997). However, since I want to keep all of these aspects in the income and consumption measure, I believe it will be more comparable to also include variation from such factors in the variance measure. In the robustness checks, the estimated income shock variances from Gourinchas and Parker (2002) has been used without any significant changes to the results (Table 7).

Table 4 – Permanent and Transitory Income Shock Variances.

	All		Low skilled		High skilled	
	Est	SE	Est	SE	Est	SE
<i>Danish Registers</i>						
$\hat{\sigma}_\eta^2$	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)
<i>PSID</i>						
$\hat{\sigma}_\eta^2$	0.0785	(0.003898)	0.0756	(0.005973)	0.0815	(0.005004)
$\hat{\sigma}_\varepsilon^2$	0.0510	(0.004452)	0.0476	(0.005374)	0.0543	(0.007086)

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

The Danish income variances are an order of magnitude smaller than those for the US. 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. 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

¹⁸Blundell, Pistaferri and Preston (2008) report $\sigma_\eta^2 \in [0.0057, 0.0333]$ and $\sigma_\varepsilon^2 \in [0.0190, 0.0753]$ depending on the combination of year, cohort and educational background, using the PSID. Gourinchas and Parker (2002), also using the PSID, calibrate $\sigma_\eta^2 = 0.0212$, $\sigma_\varepsilon^2 = 0.0440$.

top earners. Around 40 percent were 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.

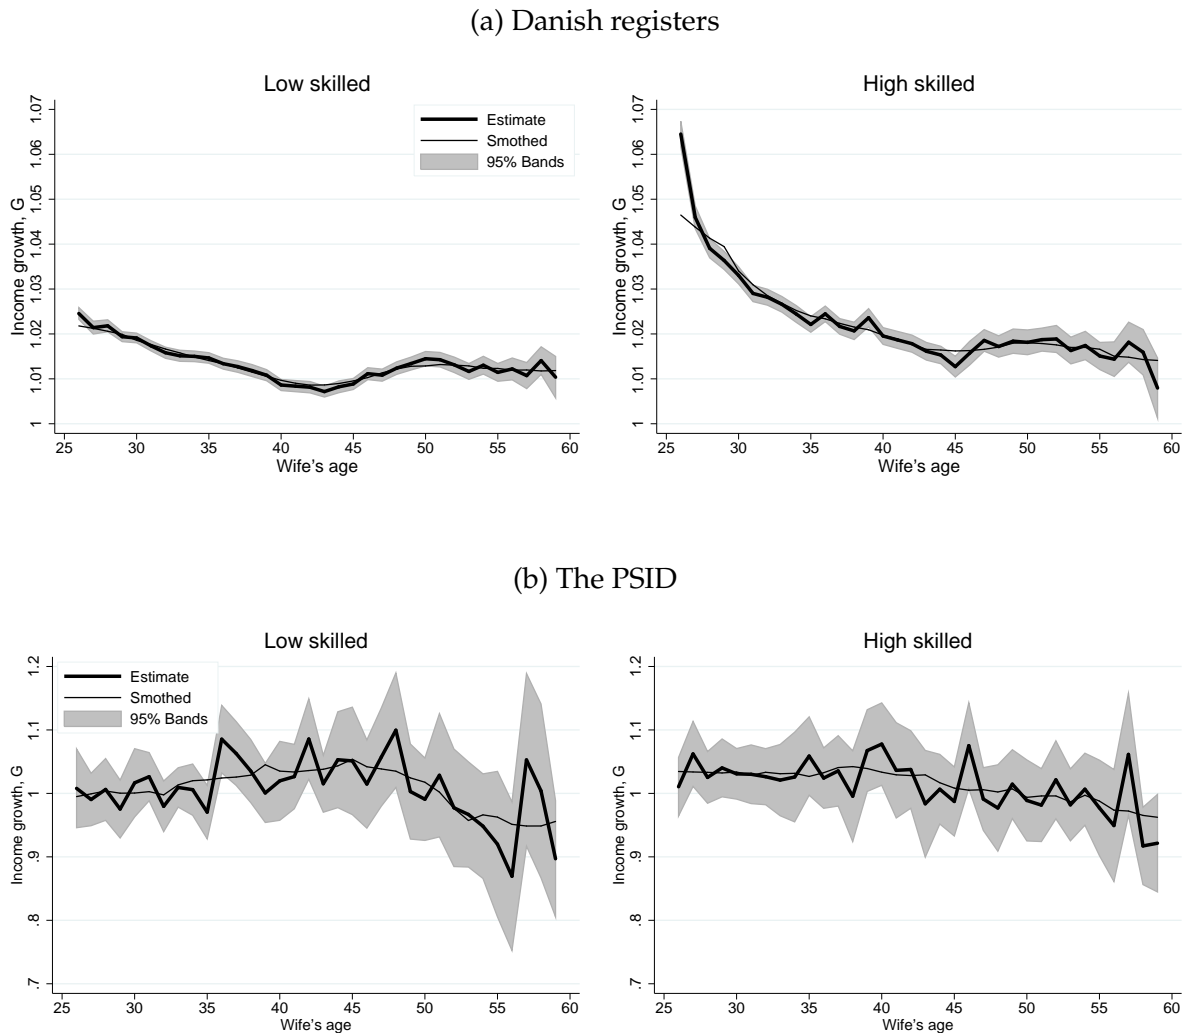


Figure 5 – Estimated Age Profiles of Income Growth, \hat{G}_t .

Notes: Figure 5 illustrates the estimated income growth age profile. Permanent income growth is estimated from equation (3).

The income growth rate, G_t , can be estimated by taking logs of the income process specified in section 3 and averaging over individuals, for a given age, t ,

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

where re-arranging and noting that the second term converges to $-\frac{1}{2}\sigma_\eta^2$ and the last

term converges to zero as $N \rightarrow \infty$, gives an estimate of the income growth rate as

$$\hat{G}_t = \exp \left(\frac{1}{N} \sum_i^N \Delta \log Y_{it} + \frac{1}{2} \sigma_\eta^2 \right). \quad (3)$$

Figure 5 reports the estimated income growth rate profile for Danish and US consumers, respectively. As expected, high skilled households have much steeper income growth than low skilled. There is significantly more noise in the PSID measures and I use, for both countries, the smoothed income growth rate.

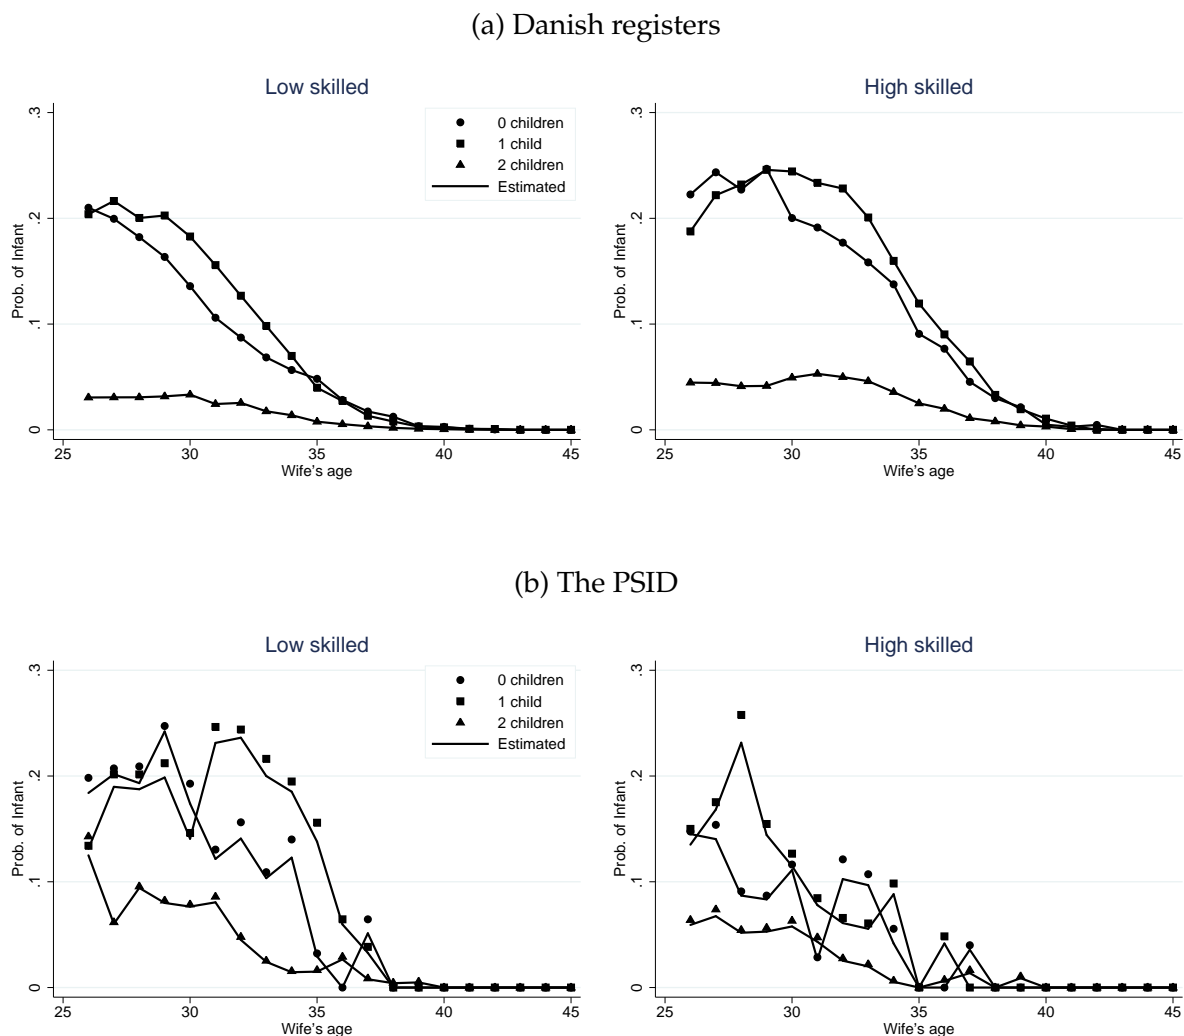


Figure 6 – Estimated Children Arrival Probabilities.

Notes: Figure 6 illustrates the estimated arrival rate of children in the Danish registers and the PSID. Arrival depend on the age of the wife and the number of children present in the household.

Permanent income, P_t , is uncovered by the Kalman Filter applied to each household's income process. Consult Appendix B for a description of the implementation. 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. The arrival probabilities using Danish registers and the PSID are presented in Figure 6.

5 Estimation Results

For both data sources, I use the raw non-durable consumption series only corrected for year-dummies through regressions to avoid removing valuable variation over the life cycle that might be correlated with children. The baseline assumption is that imputed non-durable consumption (normalized with permanent income) in the Danish registers is observed with additive normal-distributed measurement error while measurement error in the PSID is assumed logistically distributed. The results are robust to the distributional and additive assumptions (see Table 7).

The (mean) log-likelihood functions to be maximized are, thus,

$$\begin{aligned}\mathcal{L}_{DK}(\Theta, \sigma_{\xi}^2) &= -\sum_{i=1}^N \frac{1}{NT_i} \sum_{t=1}^{T_i} \left\{ \frac{1}{2} \log(2\pi\sigma_{\xi}^2) + \frac{\xi_{it}(\Theta)^2}{2\sigma_{\xi}^2} \right\}, \\ \mathcal{L}_{US}(\Theta, \sigma_{\xi}) &= -\sum_{i=1}^N \frac{1}{NT_i} \sum_{t=1}^{T_i} \left\{ \log(\sigma_{\xi}\sqrt{3}/\pi) + \frac{\xi_{it}(\Theta)}{\sigma_{\xi}\sqrt{3}/\pi} \right. \\ &\quad \left. + 2 \log \left[1 + \exp \left(-\frac{\xi_{it}(\Theta)}{\sigma_{\xi}\sqrt{3}/\pi} \right) \right] \right\},\end{aligned}$$

where $\xi_{it}(\Theta) \equiv (C_{it} - \check{C}^*(O_{it}|\Theta))/P_{it}$.

Several versions of the model are estimated for each data source. First, a model without any household compositional effects is estimated. Secondly, a functional form of the taste shifter similar to existing literature, $v(\mathbf{z}_t, \theta) = \exp(\theta \text{Number of children})$ is estimated, and, finally, a flexible functional form is implemented,

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

allowing for an arbitrary children, age and scale effect.

Table 5 presents the estimation results for low and high skilled Danish households. The estimated parameters yield a good fit of the model's predicted age profile to the Danish data, as reported in Figure 7. Table 6 presents the estimation results for low and high skilled US households in the PSID and Figure 8 reports the resulting fit. For the PSID households, the retirement motive is fixed at $\gamma = 1.3$ based on the estimates from the Danish data. This was done due to identification problems when using the rather

noisy observations in the PSID. Comparing the estimated measurement error variance for the PSID and the Danish registers, the error variance is around the same magnitude but stems from different distributions. In an earlier version of this paper, the model for PSID households was estimated assuming normally-distributed measurement error (as for the Danish data). The results indicated that the measurement error variance of the non-durable consumption measure imputed by [Blundell, Pistaferri and Preston \(2008\)](#) for PSID households is around three to seven times as noisy as the imputed consumption measure in the Danish registers.

The estimated relative risk aversion (CRRA) parameters, ρ , and discount factors, β , are in the range normally found in the literature. The discount factor is estimated to be around .97 for both low and high skilled Danish households and around .97 for low skilled and .94 for high skilled households in the PSID. [Gourinchas and Parker \(2002\)](#) also estimate a lower discount factor for high skilled (college graduates) compared to low skilled (high school graduates) in the PSID. Contrary, [Cagetti \(2003\)](#) estimates higher discount factors for college graduates than for high school graduates in the PSID and the Survey of Consumer Finances (SCF).

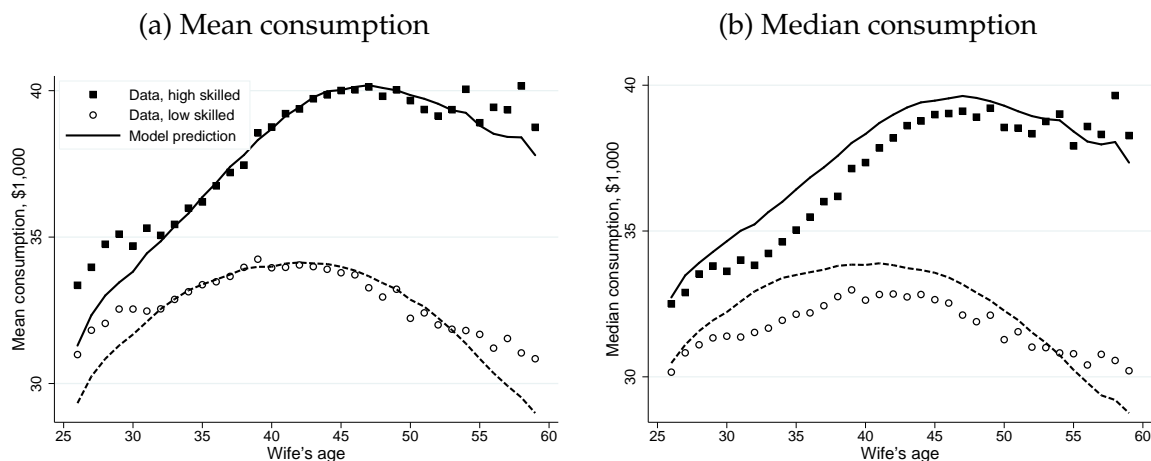


Figure 7 – Actual and Predicted Consumption profiles, Danish Data.

Notes: Figure 7 illustrates the mean (panel a) and median (panel b) age profile for actual (imputed) and predicted consumption in the Danish registers.

The estimated CRRA parameters for PSID households are around 1.3 and 1.8 for low and high skilled, respectively. [Gourinchas and Parker \(2002\)](#) and [Cagetti \(2003\)](#) also estimate larger relative risk aversion parameters for high skilled households. [Gourinchas and Parker \(2002\)](#) estimate $\rho \approx 0.87$ for high school graduates and $\rho \approx 2.29$ for college graduates while [Cagetti \(2003\)](#) estimates much higher relative risk aversion parameters. He estimates CRRA parameters of around 4 for college graduates and 3.5 for high school graduates. Constructing synthetic cohort panels using

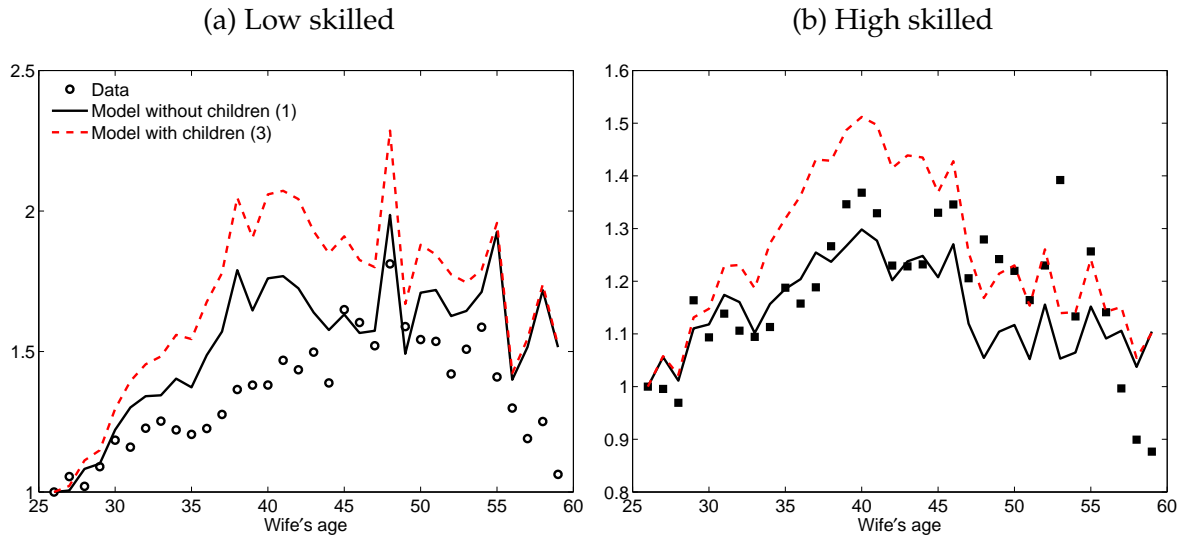


Figure 8 – Actual and Predicted Consumption Profiles, Relative to age 26, PSID.

Notes: Figure 8 illustrates the mean age profile for actual (imputed) and predicted consumption in the PSID. The measure is relative to consumption in age 26. Solid lines are based on a model *without* any effects of children, $\theta = 0$, and dashed red lines are predicted consumption from a the estimated baseline model, columns (3) in 6.

the CEX, [Attanasio, Banks, Meghir and Weber \(1999\)](#) estimate $\rho \approx 1.56$.

The estimates of ρ in Table 5 for Danish households are also in the range typically found using US and UK data. Specifically, ρ is estimated to be approximately 2.39 and 2.63 for low and high skilled Danish households, respectively. [Bingley and Lanot \(2007\)](#) estimates a dynamic discrete choice model of the retirement behavior of Danish couples. They find a very low ρ around zero implying risk-neutral consumers. Although they also use Danish register data, they focus on older households around retirement by restricting the analysis to households in which both spouses are older than 49. Further, they identify all parameters from the discrete retirement choice variation while I use the variation in (imputed) consumption. It seems likely that consumption contain more identifying variation about the relative risk aversion, ρ .

5.1 The Effect of Children on Consumption

Danish households are not significantly affected by the presence of children. Although formal Likelihood Ratio (LR) tests reject that $\theta = 0$ for Danish households, the effect is economically negligible and even negative for low skilled households. This result is not surprising in light of the descriptive evidence suggesting that childless Danish households exhibit a similar age profile as households who have children at some point in their lives.

For US households, the estimated effects of children on consumption are larger

and significant. The estimated effects of the *number of children* on consumption from the standard functional form assumption, $v(\mathbf{z}_t, \theta) = \exp(\theta \text{Number of children})$, are around 0.15 and 0.29 in columns (2) for low and high skilled, respectively. These estimates are close to but lower than the estimate of .33 ($= 0.21 / .64$) reported in [Attanasio, Banks, Meghir and Weber \(1999\)](#). Note, however, the fit of the estimated model *without* any effects of children (columns (1)) provides an equally good fit as the model in which children affect consumption (columns (3)), as shown in Figure 8. This suggests that children might not be the *primary* driver of the hump in the non-durable consumption age profile in the PSID.

The estimated effects of the number of children on consumption in columns (2) are inside the bounds reported in [Jørgensen \(2014\)](#). [Jørgensen \(2014\)](#) estimates lower and upper bounds of the effect of the number of children on consumption using roughly the same sample of PSID households as the present study. The estimated effect of the number of children on consumption for low skilled of .15 in Table 6 is inside the bounds in [Jørgensen \(2014\)](#) while the estimated effect of 0.29 for high skilled is above the upper bound. I find for high skilled households in the PSID that consumption increases in the age of children. The upper bound in [Jørgensen \(2014\)](#) is based on young households, in which children are on average younger, and the upper bound could, thus, be thought of as an upper bound on the effect of *young* children on consumption. The average effect on high skilled PSID households of having 1 through 3 young children is estimated to be approximately $(.138 + .070 + .152) / 3 = .120$. This is on the cusp of being within the upper bound of $0.047 \cdot 1.826 = 0.086$.¹⁹

Consumption is increasing in the age of children for high skilled households in the PSID. [Browning and Ejrnæs \(2009\)](#) finds that consumption increases significantly in the age of children while columns (3) in Table 6 show more moderate age-effects and primarily for high skilled ($\hat{\theta}_{j2} > \hat{\theta}_{j1} \forall j = 1, 2, 3$). The estimates are, however, imprecise and the increase in the likelihood function, compared to the standard functional form assumption in columns (2), is insignificant. In contrast to the results in [Browning and Ejrnæs \(2009\)](#), the estimated parameters do not support economies of scale in general. There seem to be economies of scale when having the second child but not when having the third child. For example, for high skilled US households, the effect of a third child older than 10 years increases the marginal value of consumption more than the first and second child in this age group did ($\hat{\theta}_{32} > \hat{\theta}_{12}, \hat{\theta}_{22}$). This is, however, not statistically significant and I cannot formally reject economies of scale.

The taste shifter can be thought of as an adjustment to the (constant) discount factor, β ([Attanasio, Banks, Meghir and Weber, 1999](#)). The fact that I estimate a lower

¹⁹The upper bound of θ for high skilled households in the PSID is calculated based on the estimated upper bound of $\widehat{\rho^{-1}\theta} \approx 0.047$ from [Jørgensen \(2014, Table 4\)](#) and the estimated $\hat{\rho} \approx 1.826$ from column (3) in Table 6 in the present study.

discount factor for high skilled households in the PSID could, thus, be due to the fact that children affect high skilled more than low skilled. This can explain why [Cagetti \(2003\)](#) finds that high skilled households have higher discount-factors while my results and those in [Gourinchas and Parker \(2002\)](#) suggest the opposite. [Cagetti \(2003\)](#) adjusts the discount factor for all educational groups using the same estimates from [Attanasio, Banks, Meghir and Weber \(1999\)](#).²⁰ If, as I find, high skilled households are more effected by the presence of children, the discount factor correction for high skilled in [Cagetti \(2003\)](#) will be *underestimated* and the estimated β for that group will, in turn, subsequently be *overestimated*.

The CRRA parameters decrease slightly with the effect children have on marginal utility of consumption. I interpret this as stemming from the fact that when marginal utility of consumption increases in the presence of children, $\theta > 0$, households are willing to substitute more consumption to future periods in which children arrive. This will decrease their preferences for smoothing consumption across time periods and yield lower an estimate of ρ .

Differences in institutions, data sources and imputation of non-durable consumption might explain why non-durable consumption of Danish households are unaffected by children while PSID households are affected. The Danish consumption measure effectively includes some durables such as dishwashers and other smaller household appliances while the US measure is imputed from food consumption measures in the PSID and CEX. The imputed PSID measure of non-durable consumption is, thus, closely related to food consumption, a consumption component likely more affected by the presence of children than total non-durable consumption.

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 2014.²¹ The living expenses of Danish households is, thus, not expected to increase as much as in the US when children arrive.

The effect of children on consumption of Danish households can, thus, be thought of as a lower bound. If children affect consumption in Denmark, children are very likely to affect consumption in other countries, such as the UK and the US. Since I find no effect of children, the lower bound is not very informative, however.

²⁰The discount factor adjustments in [Cagetti \(2003\)](#) do vary across educational groups but only due to differences in demographics, z_t , across educational groups. The loadings on these demographics, θ , are constant across educational groups.

²¹When 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.

6 Robustness Checks

The results are robust. Table 7 reports estimates from the PSID using four alternative specifications. Columns (1) report results when the absolute difference between observed and predicted consumption is minimized, relying not on the logistic assumption of measurement error. Columns (2) report results from minimizing the absolute difference in log-consumption growth, allowing for arbitrarily distributed multiplicative measurement error in consumption with heterogeneity in the measurement error across households. Columns (3) report *identical* estimation results from a model where the permanent and transitory income variances are calibrated to those found in [Gourinchas and Parker \(2002\)](#), $\sigma_\eta = 0.0212$ and $\sigma_\varepsilon = 0.0440$, respectively. Finally, columns (4) report results from a version of the model in which children arrive deterministically in the sense that they are perfectly foreseen by households.

Minimizing the absolute difference in consumption in columns (1) produce almost unchanged results. The relative risk aversion decreases slightly while the discount factor increases slightly. The effect of children is estimated to be smaller for low skilled and larger for high skilled, compared to the baseline. The differences are, however, insignificant. The results from minimizing the absolute difference in log-consumption growth (Example 2 in Section 4) in columns (2) are less comparable because the intertemporal discount factor was fixed at .95 due to identification difficulties in this specification. In this specification, children are estimated to have a negative and insignificant effect on household consumption.

The assumption of probabilistic arrival of children might be expected to reduce the estimates of the effect of children. In the probabilistic version, *all* households within a given age group have the same expected future children-related expenditures. The probabilistic model, therefore, suggests that *all* households should decrease consumption before having children in anticipation of children in the future. If this is not how households perceive the world, the estimate of θ would be forced downwards by households who *know* that they will have children, say, late in their life. Such households do, then, not save as much as the model would suggest when young and the only way for the probabilistic model to fit this behavior is to reduce how much children affect consumption.

Columns (4) in Table 7 report very similar estimation results, however, from a deterministic version of the model, in which households know with perfect foresight how many and when children will arrive, following [Browning and Ejrnaes \(2009\)](#). If anything, the effect of children is estimated to be less than in the probabilistic model. Since this model requires knowledge on *completed* fertility, I made the crude assumption that observed fertility is completed fertility. Alternatively, outside the scope of this paper, an estimated completed fertility could be used ([Browning and Ejrnaes, 2009](#)).

7 Concluding Discussion

I have estimated the effect of children on non-durable consumption using both the PSID for the US and high quality administrative register data of the entire Danish population. Results suggest that the effect of children on non-durable consumption might be smaller than previously assumed. Danish households are not significantly affected by the presence of children while PSID households are. The estimates for the PSID sample is, however, below those reported in the influential study by [Attanasio, Banks, Meghir and Weber \(1999\)](#).

The current study allows for several competing life cycle consumption and savings motives simultaneously. Specifically, children, income uncertainty, credit constraints, and retirement affect the dynamic decision on how much to consume and save in each period over the life cycle. Most existing studies on the effect of demographics on consumption exclude *all* motives except demographics and (not surprisingly, perhaps) find that children are important in explaining the observed hump in consumption age profiles. The number and age of children have been found to completely explain the life cycle profile of consumption ([Browning and Ejrnæs, 2009](#)).

By allowing for several competing consumption/savings motives simultaneously, the present study provides new insights on the *relative* importance of demographics and precautionary motives in explaining the life cycle consumption profile. The finding of small effect of children on non-durable consumption suggests that precautionary motives are more important than demographic effects. The Danish registers allow identification of childless households around completed fertility and the age profile of non-durable consumption of childless Danish households is almost identical to households who have children. This is new evidence that children *cannot* be the primary driver of the observed hump in the consumption age profile.

The estimated parametrization allows for an arbitrary children, age and scale effect. I find, as [Browning and Ejrnæs \(2009\)](#), that there is a tendency for an age effect. High skilled households in the PSID seem to increase non-durable consumption as children grow older. In contrast to the results in [Browning and Ejrnæs \(2009\)](#), the estimated parameters herein do not support economies of scale in general. I estimate economies of scale when having the second child but not when having the third child. The estimates are, however, rather imprecise.

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.

Table 5 – Estimated Preference Parameters, Danish Registers.

		Low skilled			High skilled		
		(1)	(2)	(3)	(1)	(2)	(3)
ρ	Risk aversion	2.316 (0.041)	2.363 (0.036)	2.385 (0.043)	2.639 (0.057)	2.626 (0.062)	2.634 (0.063)
β	Discount factor	0.965 (0.000)	0.964 (0.000)	0.964 (0.000)	0.973 (0.001)	0.973 (0.001)	0.972 (0.001)
γ	Retirement	1.454 (0.018)	1.492 (0.018)	1.491 (0.020)	1.251 (0.022)	1.245 (0.023)	1.265 (0.025)
σ_{ξ}	Meas. error	0.468 (0.000)	0.468 (0.000)	0.468 (0.000)	0.490 (0.001)	0.490 (0.001)	0.490 (0.001)
Taste shifter: $v(\mathbf{z}; \theta) = \exp(\theta' \mathbf{z})$							
θ	# of children		-0.017 (0.002)			0.004 (0.003)	
Taste shifter: $v(\mathbf{z}; \theta) = 1 + \theta' \mathbf{z}$							
θ_{11}	1. child ≤ 10			-0.004 (0.007)			-0.008 (0.010)
θ_{12}	1. child > 10			-0.031 (0.004)			0.002 (0.008)
θ_{21}	2. child ≤ 10			-0.034 (0.005)			-0.015 (0.008)
θ_{22}	2. child > 10			-0.006 (0.005)			0.000 (0.008)
θ_{31}	3. child ≤ 10			-0.005 (0.009)			0.022 (0.012)
θ_{32}	3. child > 10			0.019 (0.013)			0.021 (0.017)
	$-\mathcal{L}(\Theta)$	0.46536	0.46533	0.46529	0.49868	0.49863	0.49862
	$\max_j \partial \mathcal{L}(\Theta) / \partial \Theta_j $	$7.1e - 6$	$1.5e - 5$	$2.1e - 5$	$8.3e - 6$	$2.3e - 5$	$1.8e - 5$
	LR [p -val]		67[0.00]	153[0.00]		57[0.00]	68[0.00]
	# 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. 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.

Table 6 – Estimated Preference Parameters, PSID.

		Low skilled			High skilled		
		(1)	(2)	(3)	(1)	(2)	(3)
ρ	Risk aversion	1.540 (0.356)	1.432 (0.325)	1.272 (0.333)	2.087 (0.350)	2.283 (0.364)	1.826 (0.361)
β	Discount factor	0.967 (0.005)	0.975 (0.003)	0.976 (0.002)	0.934 (0.013)	0.939 (0.013)	0.949 (0.010)
γ^\dagger	Retirement	1.300	1.300	1.300	1.300	1.300	1.300
σ_{ξ}	Meas. err.	0.485 (0.001)	0.484 (0.001)	0.483 (0.001)	0.628 (0.001)	0.627 (0.001)	0.626 (0.001)
Taste shifter: $v(\mathbf{z}; \theta) = \exp(\theta' \mathbf{z})$							
θ	# of children	0.151 (0.046)			0.289 (0.093)		
Taste shifter: $v(\mathbf{z}; \theta) = 1 + \theta' \mathbf{z}$							
θ_{11}	1. child ≤ 10	0.201 (0.124)			0.138 (0.215)		
θ_{12}	1. child > 10	0.269 (0.118)			0.371 (0.204)		
θ_{21}	2. child ≤ 10	0.057 (0.086)			0.070 (0.184)		
θ_{22}	2. child > 10	0.075 (0.113)			0.208 (0.222)		
θ_{31}	3. child ≤ 10	0.258 (0.159)			0.142 (0.246)		
θ_{32}	3. child > 10	0.092 (0.195)			0.564 (0.454)		
$-\mathcal{L}(\Theta)$		0.39229	0.39074	0.39047	0.45887	0.45792	0.45750
$\max_j \partial \mathcal{L}(\Theta) / \partial \Theta_j $		$1.6e - 6$	$4.5e - 7$	$9.6e - 6$	$1.6e - 6$	$1.1e - 5$	$6.8e - 6$
LR [p -val]			45[0.00]	53[0.00]		33[0.00]	48[0.00]
# of observations		8,333	8,333	8,333	8,672	8,672	8,672

Notes: Standard errors are based on the outer product of scores. For some specifications, the Hessian was singular within numerical precision. This indicates that the objective function is almost flat and the estimates and standard errors for the PSID sample should be interpreted with this in mind. 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.

[†] For US households, the retirement motive is fixed across educational groups to be 1.3 based on results for Danish households in Table 5. This is done because simultaneous identification of ρ , β and γ failed using the PSID sample.

Table 7 – Robustness Checks, PSID.

	Low skilled				High skilled			
	Abs (1)	Het (2)	G&P (3)	Det (4)	Abs (1)	Het (2)	G&P (3)	Det (4)
ρ	1.148 (0.577)	2.411 (1.130)	1.272 (0.333)	1.265 (0.225)	1.956 (0.615)	1.930 (0.846)	1.826 (0.361)	1.812 (0.209)
β^\dagger	0.982 (0.004)	0.950	0.976 (0.002)	0.975 (0.002)	0.960 (0.013)	0.950	0.949 (0.010)	0.948 (0.005)
σ_ξ^\ddagger			0.483 (0.001)	0.487 (0.001)			0.626 (0.001)	0.615 (0.000)
θ_{11}	0.181 (0.207)	-0.355 (0.174)	0.201 (0.124)	0.107 (0.094)	0.309 (0.385)	0.115 (0.221)	0.138 (0.215)	0.013 (0.145)
θ_{12}	0.172 (0.173)	-0.362 (0.173)	0.269 (0.118)	0.170 (0.088)	0.385 (0.306)	0.078 (0.165)	0.371 (0.204)	0.233 (0.138)
θ_{21}	0.044 (0.137)	-0.084 (0.105)	0.057 (0.086)	0.097 (0.074)	-0.018 (0.266)	-0.054 (0.185)	0.070 (0.184)	0.119 (0.150)
θ_{22}	0.099 (0.187)	0.017 (0.125)	0.075 (0.113)	0.103 (0.084)	0.175 (0.325)	0.285 (0.264)	0.208 (0.222)	0.267 (0.161)
θ_{31}	0.234 (0.258)	-0.183 (0.133)	0.258 (0.159)	0.111 (0.080)	0.236 (0.408)	-0.241 (0.267)	0.142 (0.246)	0.135 (0.165)
θ_{32}	0.098 (0.294)	-0.119 (0.171)	0.092 (0.195)	0.041 (0.113)	0.661 (0.765)	0.087 (0.349)	0.564 (0.454)	0.288 (0.204)

Notes: Standard errors are based on the outer product of scores.

"Abs" in columns (1) refer to results from minimizing the absolute difference between observed and predicted consumption.

"Het" in columns (2) refer to results from minimizing the absolute change in log-differences between observed and predicted consumption, similar to example 2 in section 4.

"G&P" in columns (3) refer to a model in which the permanent and transitory income variances are calibrated to those found in [Gourinchas and Parker \(2002\)](#), $\sigma_\eta = 0.0212$ and $\sigma_\varepsilon = 0.0440$, respectively.

"Det" in columns (4) refer to a model in which children are deterministic in the sense that they are perfectly foreseen by households.

[†] The discount factor is fixed at .95 when minimizing absolute change in differences since I was unable to identify all parameters simultaneously in this specification.

[‡] Minimizing absolute differences (2) and absolute change in differences (3) does not yield an estimate of the measurement error variance.

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.
- 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.
- 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.
- BINGLEY, P. AND G. LANOT (2007): "Public pension programmes and the retirement of married couples in Denmark," *Journal of Public Economics*, 91(10), 1878–1901.
- 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.
- CARROLL, C. D. (1992): "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 2, 61–135.
- (1997): "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *The Quarterly Journal of Economics*, 112(1), pp. 1–55.
- (2006): "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economics Letters*, 91(3), 312–320.
- CARROLL, C. D. AND A. A. SAMWICK (1997): "The nature of precautionary wealth," *Journal of Monetary Economics*, 40(1), 41–71.
- DAVID, P. A. AND T. A. MROZ (1989): "Evidence of Fertility Regulation among Rural French Villagers, 1749-1789: A Sequential Econometric Model of Birth-Spacing Behavior (Part 1)," *European Journal of Population*, 5(1), 1–26.
- DEATON, A. (1991): "Saving and Liquidity Constraints," *Econometrica*, 59(5), 1221–1248.
- FELLA, G. (2014): "A generalized endogenous grid method for non-smooth and non-concave problems," *Review of Economic Dynamics*, 17(2), 329–344.
- FERNÁNDEZ-VILLAYERDE, 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-VILLAYERDE, 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.
- HOTZ, V. J., J. A. KLERMAN AND R. J. WILLIS (1997): "The economics of fertility in developed countries," in *Handbook of Population and Family Economics*, ed. by M. R. Rosenzweig and O. Stark, vol. 1, Part A of *Handbook of Population and Family Economics*, chap. 7, pp. 275–347. Elsevier.
- HOTZ, V. J. AND R. A. MILLER (1988): "An Empirical Analysis of Life Cycle Fertility and Female Labor Supply," *Econometrica*, 56(1), 91–118.
- IRVINE, I. (1978): "Pitfalls in the Estimation of Optimal Lifetime Consumption Patterns," *Oxford Economic Papers*, 30(2), pp. 301–309.
- ISKHAKOV, F., T. H. JØRGENSEN, J. RUST AND B. SCHJERNING (2014): "Estimating Discrete-Continuous Choice Models: Endogenous Grid Method with Taste Shocks," Unpublished working paper.
- 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.
- MINISTRY OF FINANCE (2003): *ældres sociale vilkår*. (In Danish).
- RUST, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55(5), 999–1033.
- SCHECHTMAN, J. (1976): "An Income Fluctuation Problem," *Journal of Economic Theory*, 12, 218–41.
- 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): "Constrained Optimization Approaches to Estimation of Structural Models," *Econometrica*, 80(5), 2213–2230.

THUROW, L. C. (1969): "The Optimum Lifetime Distribution of Consumption Expenditures," *The American Economic Review*, 59(3), pp. 324–330.

WOOLDRIDGE, J. M. (2002): *Econometric Analysis of Cross Section and Panel Data*. Massachusetts Institute of Technology, USA.

ZELDES, S. P. (1989): "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *Quarterly Journal of Economics*, 104(2), 275–298.

Online Appendix (Not for Publication)

Life-Cycle Consumption and Children

Thomas H. Jørgensen

University of Copenhagen

October 29, 2014

A 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 $\hat{a}_t = (\underline{a}_t, a_t^1, \dots, a_t^{Q-1})$ 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^* , is given by $m_t = \hat{a}_t + c_t^*(m_t, \mathbf{z}_t)$.

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

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

where $\check{c}_{t+1}^*(m_{t+1}, \mathbf{z}_{t+1})$ 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 in 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 $\text{age}_{2,t+1} = 0$.

A.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, $c_{t+1}^* = m_{t+1}$ if m_{t+1} is lower than some threshold level, m_{t+1}^* . 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., $m_{t+1}^* = \underline{m}_{t+1}$. 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

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

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,

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

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,

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

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 $a_t = -\min\{\Omega_t, \kappa\}$. \square

B 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

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

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,

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

where \mathfrak{S}_s denotes information known at time s . The updating step is given by

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

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, $\log Y_0 = \mathbb{E}[\log P_0 - \frac{1}{2}\sigma_\varepsilon^2 + \mathbf{v}_t | \mathfrak{S}_t]$, such that $\mu_{0|0} = \log Y_0 + \frac{1}{2}\sigma_\varepsilon^2$ and $\Sigma_{0|0} = \sigma_\varepsilon^2$.

Chapter 4

Leisure Complementarities in Retirement

Leisure Complementarities in Retirement

Thomas H. Jørgensen[†]

October 30, 2014

Abstract

I analyze which factors affect the valuation of joint leisure of elderly couples. It is well-documented that couples tend to retire within few years of each other but little is known about which factors affect leisure complementarities in retirement. I shed light on this by estimating a dynamic model of household consumption and retirement choices using high quality Danish administrative register data. To disentangle leisure complementarities from household level shocks, health and labor market income shocks are allowed to be correlated across household members. Individuals are heterogeneous in labor market income, eligibility for pension benefits, wealth, education, health, children and grandchildren. Results suggest an important role for leisure complementarities in retirement. Wives are found to value joint leisure more than their husbands and low skilled individuals value joint leisure more than high skilled. Children increase the value of joint leisure while grandchildren have a small negative effect. Own health status does not affect leisure complementarities while poor health of a spouse decreases the value of joint leisure (JEL: C63, D13, D91, J14, J26).

Keywords: Leisure Complementarities, Joint Retirement, Household Consumption, Labor Supply, 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.

I am deeply grateful to Bertel Schjerning, Mette Ejrnæs, John Rust, Bo Honoré, Áureo de Paula, Fedor Iskhakov, Eric French, Martin Browning, Christopher Carroll, Søren Leth-Pedersen, Peter Haan, Anders Munk-Nielsen, Damoun Ashournia, and Jeppe Druedahl for valuable comments and discussions. This project has benefited from participants at the CAM December workshop in Copenhagen, 2012, the Initiative for Computational Economics (ICE) at Chicago, 2012, the Workshop on Public Economics and Public Policy, Copenhagen, 2013, and the EEA in Toulouse, 2014. I also wish to thank the Danish Economic Council for initial data preparation. Any errors are my own.

1 Introduction

This study is concerned with complementarities in leisure of elderly households. Since [Hurd \(1990\)](#) documented that couples often retire within few years of each other, a growing literature has confirmed this finding.¹ Little is known, however, about how joint leisure in retirement is valued across groups and household characteristics. A better understanding of the determinants of leisure complementarities can guide optimal policy design towards more efficient reforms. Specifically, if leisure complementarities are important, policies that target incentives of *individuals*, such as increasing the early retirement age, have second order (spill-over) effects on the spouse who may or may not be directly affected by the reform.

The present study investigates the importance of and the heterogeneity in leisure complementarities in retirement. I do this by estimating a dynamic model of consumption and retirement of heterogeneous couples using high quality Danish administrative register data. Income and health shocks are allowed to be correlated across household members facilitating separation of complementarities in leisure from household-level shocks. Intra-household correlations in income and health shocks have, to the best of my knowledge, not been allowed in existing dynamic models of household retirement behavior. Ignoring common shocks will inflate the importance of leisure complementarities and is, thus, important to take into account when quantifying the importance of leisure complementarities.

Results suggest an important role for leisure complementarities in retirement. I estimate that leisure is valued twice as much if the spouse is also retired. Counterfactual policy simulations suggest that ignoring leisure complementarities when increasing the retirement age underestimates the government surplus by around 60-70 percent. The results stress the fact that performing event-studies of retirement reforms on individual's retirement responses (as in [Burtless, 1986](#)) will likely provide biased results. The response of an individual to a reform will be contaminated by spill-over effects from the spouse who may also be affected by the reform under study.

I estimate substantial heterogeneity in leisure complementarities. Specifically, wives are found to value joint leisure more than their husbands and low skilled individuals value joint leisure more than high skilled. Children are found to decrease the value of own leisure but increase the value of joint leisure. While grandchildren increase the value of own leisure, grandchildren decrease the value of joint leisure. Further, I estimate that the value of joint leisure is unaffected by own health status while poor spousal health decreases it.

Complementarities in leisure is typically found to be important ([Gustman and Stein-](#)

¹There is a large body of literature on joint (simultaneous) retirement of couples. Some important contributions are due to [Blau \(1998\)](#); [Gustman and Steinmeier \(2000, 2004\)](#); [Coile \(2004\)](#); [Jia \(2005\)](#); [Blau and Gilleskie \(2006\)](#) and [van der Klaauw and Wolpin \(2008\)](#).

meier, 2000, 2004; van der Klaauw and Wolpin, 2008; Casanova, 2010 and Honoré and de Paula, 2013). Using the Health and Retirement Study (HRS) and the National Longitudinal Survey of Mature Women (NLS), Gustman and Steinmeier (2000, 2004) estimate dynamic economic models of optimal retirement in dual earner households. They find that leisure complementarities are important. Especially the husband's utility from leisure is increased from his wife also being retired. Coile (2004) estimates probit models for married couples in the HRS and finds that financial incentives of spouses affect the retirement behavior, especially of married men. Using direct survey measures on the enjoyability of joint leisure, Coile (2004) shows how the asymmetry in the response to spousal incentives may be attributed to husbands valuing joint leisure more than their wives.

An, Christensen and Gupta (2004) estimates a bivariate mixed proportional hazards duration model using Danish couples. They formulate three durations; one for each household member and one for joint retirement, which are allowed to be correlated through observable characteristics and unobservable shocks. The joint retirement duration is interpreted as leisure complementarities and the correlations between individual durations are interpreted as intra-household correlated shocks. They find evidence that complementarities are more important than common shocks since they cannot reject that shocks to individual durations are uncorrelated across members. The model, however, does not yield easily interpretable economic parameters related to the household decision process.² Michaud and Vermeulen (2011) also find that complementarities are important in describing the retirement pattern of couples in the HRS by estimating a (static) collective model of couples' retirement.

Health has been a heavily studied driver of retirement of both singles and couples in the US. Blau (1998) estimates a random-effects multinomial probit model for married couples in the Retirement History Survey (RHS). He finds that wives are more likely to retire (independently and simultaneously with spouses) if the husband is in poor health while the reverse is not as pronounced. An, Christensen and Gupta (2004) find, as I do, that poor spousal health tend to *defer* retirement.

Children and grandchildren have also been identified as potential drivers of labor market supply and retirement. In a recent working paper, Warren (2013) estimates competing risks models of household retirement behavior using the Household, Income and Labour Dynamics in Australia (HILDA) data. Her results support the asymmetry in spousal health on retirement, found by Blau (1998). Further, females are found to retire earlier if a child is residing while that is not the case for males. Ho (forthcoming); Posadas and Vidal-Fernandez (2013) and Rupert and Zanella (2014) go one gener-

²The fact that the duration until *simultaneous* retirement is parametrized implies that interpretation of the parameters is fundamentally different from the present study. Their estimates could be interpreted as potentially saying something about complementarities in the *act* of retiring simultaneously.

ation further and investigate how grandchildren is related to the labor market supply of grandparents. [Ho \(forthcoming\)](#) finds that elderly individuals in the HRS are more likely to be employed when the first grandchild arrives. The effect is only statistically significant for married individuals. She finds evidence that these households are also more likely to provide financial help through increased labor market attachment. Using the Panel Study of Income Dynamics (PSID), [Rupert and Zanella \(2014\)](#) find the opposite result that grandparents *decrease* labor market hours especially around when they first become grandparents.

The existing literature has, however, not disentangled leisure complementarities from financial incentives and household level shocks. If income and/or health shocks are correlated across household members, adverse shocks are likely to induce couples to retire simultaneously. In that sense, correlated shocks are observationally equivalent to complementarities in leisure. I identify leisure complementarities by *i)* fully specifying the economic environment in which households make retirement decisions, while *ii)* allowing shocks to individual labor market income and health to be correlated across household members when estimating the model using Danish register data.

Danish households are well suited for studying retirement of couples. First, the Danish registers provide high quality longitudinal information on relevant variables such as wealth, income and household characteristics. The fertility register allows the linking of several generations. I use data on privately held pension wealth, information that is rarely available, allowing accurate measurement of the household contingencies as they were when consumption and labor supply decisions were made.

Secondly, Danish health care is universal and free. Only few health-related costs are not fully covered. This includes expenses for prescription medicine, the dentist and glasses although the former is heavily subsidized. The health care provision is independent of labor market participation (although financed through labor market income taxation). In turn, the Danish health care system allows significant modeling simplifications relative to the US where many households rely on employer-provided health care, distorting the retirement incentives ([Rust and Phelan, 1997](#); [French, 2005](#)). Since that is not the case for Denmark, health affects retirement primarily through the dis-utility or cost of labor market participation.

The rest of the paper is organized as follows. The next section describes the Danish pension and tax systems and Section 3 discusses the Danish register data. Section 4 presents a dynamic model of consumption and retirement and Section 5 discusses how the model is solved numerically. In Section 6, some model parameters are calibrated and Section 7 reports estimates of the remaining parameters. Using the estimated model, Section 8 illustrates how government surplus from counterfactual policy experiments are affected by leisure complementarities. Finally, Section 9 concludes.

2 The Danish Institutional Settings

The Danish pension system, naturally, plays a key role when Danish households form their labor market decisions. Furthermore, the Danish tax system allows couples to transfer unused tax deductions from one spouse to the other. All institutional settings used throughout are based on actual rules as they were in 2008. These rules were relevant for Danish birth cohorts born in 1940-1948, used to estimate the model.

The Danish retirement benefit system consists of two main elements: Early retirement pension (ERP, in Danish “*efterløn*”) and old age pension (OAP, in Danish “*folkepension*”). ERP is a voluntary program in which participants pay annual membership fees to obtain eligibility for annual ERP benefits of around \$30,000.³ The earliest age of eligibility is at age 60. OAP is available to *all* individuals at the age of 65 and is fully government financed. Below, I describe the Danish pension and tax systems as they are implemented in the present model. The supplemental material contains exact formulas for the implemented institutional settings.

2.1 The Early Retirement Pension (ERP) Scheme

The ERP scheme affects only financial incentives of *individuals* without any considerations about other household members. To be eligible for ERP benefits, an individual must have contributed to the program (approximately \$1,000 annually) for at least ten years.⁴ In effect, the government finances roughly 70 percent of the benefits, making the ERP scheme very popular.

To qualify for ERP, individuals must *be available to the labor market*. Particularly, retiring (permanently) before being eligible to ERP implies ineligibility to ERP in the future. For example, imagine a single female headed household being eligible for ERP at the age of 60. If she chooses to retire at age 59 she will waive five years of ERP benefits and have to finance any years of retirement before age 65 (OAP) herself. This labor market attachment requirement, in turn, results in almost no retirement prior to age 60 and a substantial spike in retirement at age 60.

The ERP benefits are means tested and depend on the *i*) level, *ii*) type and *iii*) whether the pension savings are administrated by the employer or the employee. Increasingly many employers of Danish wage workers contribute a percent (typically ranging from 12 to 17 percent) of the gross labor market income to individual pension funds. These pension funds are referred to as employer-administrated. Employee-administrated pension funds are funds which the employee decides to place in a pen-

³The ERP is subject to means testing and the actual payments are, thus, typically lower than \$30,000.

⁴The number of years varies across cohorts. Ten years of membership is the requirement for the cohorts used in this study while, for example, cohorts born in and after 1976, 30 years of membership is required for ERP eligibility.

sion savings account.

Three main *types* of retirement savings opportunities are available. First, *Lifelong Annuity (LA)*, in Danish »Livsvarige Pensionsordninger«, is an insurance guaranteeing a monthly payment when retired. The amount guaranteed (commitment value) is received until death and is therefore increased (decreased) if the owner postpone (advance) retirement. Secondly, *Annuitized Individual Retirement Arrangement (AIRA)*, in Danish »Ratepension«, is a pension balance committed by the owner to be distributed through annuities of 10 through 25 years, initiated no later than age 77. Thirdly, *Individual Retirement Arrangement with no restrictions (IRA)*, in Danish »Kapitalpension«, is an AIRA with no commitment to annuitize pension savings and no upper age limit to when the owner must withdraw the funds.

All private pension wealth is assumed to be IRAs in the present study. The six combinations of types of pension savings accounts in either employer or employee administration affects the ERP benefits differently. Due to a lack of disaggregation of the annual pension contributions to different types in the Danish registers, it is impossible to construct the different disaggregated balances. The ERP does not, however, distinguish between employer and employee administrated IRAs simplifying the implementation of the ERP in the economic model. All privately held pension wealth has been converted into IRA-equivalents by the Danish Economic Council.

If an individual has been eligible for ERP in at least two years before retiring, the level of ERP increases by ten percent and the privately held pension wealth is not means tested.⁵ This is referred to as the *two-year rule*. Particularly, individuals retiring *without* fulfilling the two-year rule receive 166,400DKK (\approx \$30,250) net of three percent of the IRA pension savings account balance at the point of retirement above a deduction of 12,600DKK (\approx \$2,300). Retirees fulfilling the two-year rule receive ERP benefits of 182,780DKK (\approx \$33,250) without any means testing.

2.2 The Old Age Pension (OAP) Scheme

The Danish OAP system is available to *all* Danish citizens aged 65 or above. The OAP, therefore, create a third (small) spike in retirement at age 65. It is primarily means tested based on individual labor market income but also labor market income of a potential spouse affects the benefits received. The OAP benefit differs across singles and couples and the labor market status of a potential spouse. Unlike the ERP, OAP is *not* means tested based on wealth holdings.

The baseline annual benefit (in Danish "grundbeløb") is approximately \$10,700 which is subject to a reduction of 30 percent of annual labor market earnings exceeding

⁵The requirement is wage work in at least 3,120 hours in at least two years after eligibility for ERP. Self-employed has similar but more vague requirements and are not included in the present study.

\$45,000. On top of the baseline benefit, an OAP supplement (in Danish “pensionstillæg”) of at most \$10,000 is available. The supplement depends on whether the recipient is *i*) single, *ii*) married to a non-OAP recipient, or *iii*) married to an OAP recipient.

If an individual is married or cohabiting with a partner who is *not* receiving OAP, half of the spouse’s annual labor market income exceeding \$20,000 is included in the measure of household income used to calculate the OAP benefits. The annual supplement of \$5,000 for couples (\$10,000 for singles) is reduced by 30 percent of the household income in excess of \$20,000. If the spouse is *also* receiving OAP, the annual supplement is reduced with only 15 percent of household income above \$20,000. The maximum amount of labor market income allowed before all OAP benefits voids also depend on the household composition. The online supplemental material contains detailed information and exact formulation of the implemented OAP and ERP schemes.

Figure 1a reports ERP and OAP benefits for single households with no private pension wealth and no labor market income in retirement. The figure illustrates the received benefits if retiring at age 60 (ERP), age 62 (two-year-rule), and age 65 (OAP).

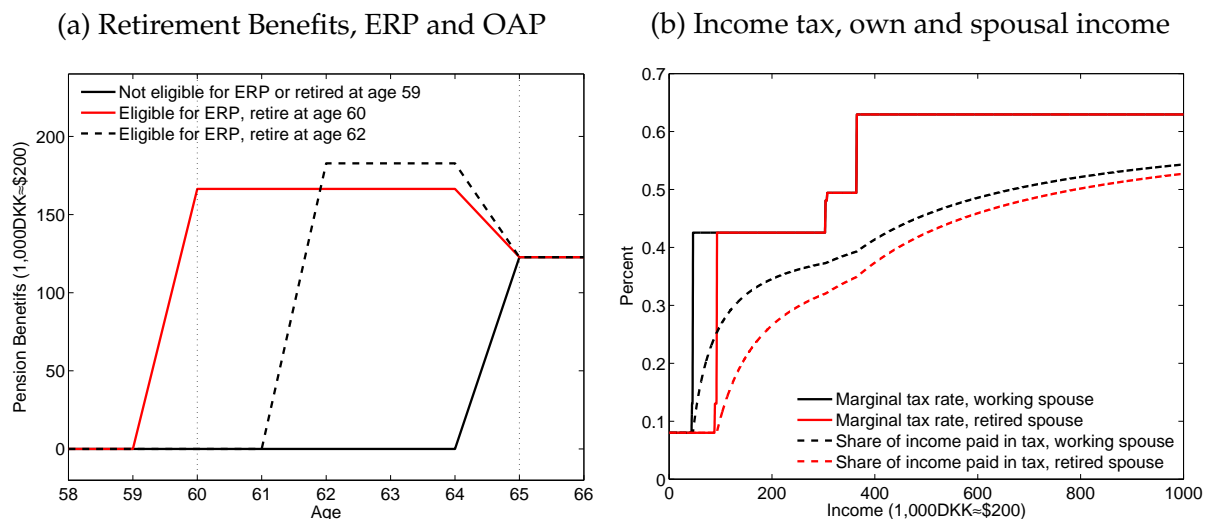


Figure 1 – Pension Benefits and Labor Market Income Taxation in Denmark.

Notes: Figure 1a illustrates how the Danish early retirement pension (ERP) and old age pension (OAP) scheme depend on the age of retirement and eligibility for ERP. The illustration is based on a single individual with no private pension wealth. Figure 1a illustrates the Danish tax schedule through marginal and share of income paid in tax as a function of labor market income. Unused labor market income deductions of 41,000 Danish kroner, or approximately 7,500 US dollars, from a potential spouse can be transferred to the other spouse. Black lines illustrate the tax schedule in the case of a spouse working (and using the deduction) while red lines illustrate the schedule if a spouse is retired (and has zero labor market income).

2.3 The Danish Tax System

The Danish tax scheme is one of the most progressive tax systems in the world. The system consists of three tax brackets in 2008; a “bottom tax”, a “middle tax” and a “top tax”. The marginal tax rate in the bottom bracket is around 43 percent, the marginal tax rate in the middle bracket is around 50 percent, and the marginal tax rate in the top tax bracket is around 63 percent.

The after tax income can be calculated by applying formulas given in supplementary material which also contains 2008 values for all relevant parameters governing the tax system. Figure 1b illustrates the Danish tax schedule through marginal tax rates along the share of income paid in tax as a function of labor market income.

If a spouse does not utilize the full labor market deduction of approximately \$7,500, the remainder is deductible to the spouse in the lowest tax bracket. This feature of the tax system increases the incentive to remain working if the spouse is retired. This is especially true for low income earners but high income earners are also affected by the transferable deduction from a retired spouse.

Contributions to IRAs are deductible in the income tax. If the owner initiates the distribution of funds after retirement, a 40 percent tax payment of the withdrawn amount will be collected by the government. If the funds are withdrawn earlier than the early retirement age, a tax of 60 percent is collected. Hence, the distribution of funds do not necessary start at the age of retirement although that is most common.

3 The Danish Administrative Register Data

This study uses Danish administrative register data for the entire Danish population in the years 1996–2008. The data provide high quality longitudinal information about all household members on labor market income, assets and liabilities, household characteristics such as marital status, age, number of annual visits to the general practitioner (GP), and educational attainment. The Danish fertility register, initiated in 1930, facilitates linking several generations providing information on potential grandchildren.⁶ These features make the Danish register data well suited for studying how household characteristics affect complementarities in leisure.

Household consumption is not observed in the registers and is, thus, imputed using a simple budget approach,

$$C_t = y_t^{disp} - \Delta a_t, \quad (1)$$

⁶Virtually *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 emigrated before January 1st 1979 is *not* included in the Danish fertility registers. The youngest potential births used to identify children and grandchildren are in $1996 - (68 - 12) = 1940$, assuming that fertility begins at age 12. In turn, almost *all* mothers used here will be matched to their children, if they have any.

where y_t^{disp} is disposable income, a_t is end-of-period net wealth, and Δa_t proxies savings. This imputation method is evaluated on Danish data in [Browning and Leth-Petersen \(2003\)](#) and found to produce a reasonable approximation.

I include information on private pension wealth held in pension funds. Individual pension wealth is based on information collected for the Danish tax authorities to calculate the individual ERP benefits (in Danish »Pensionsrettigheder«). Pension wealth information is collected for all individuals at the age of 59½ independent of their eligibility to receive ERP. The reduction in ERP from private pension wealth was introduced in 1999 such that pension wealth for individuals aged 61 or above in 2000 are not available. The oldest individuals in the data are, thus, 68 years old. The cohorts used throughout are those born between 1940-1948, both years included.⁷

To reduce household heterogeneity, that is not explicitly included in the model, I exclude several households from the analysis. Specifically, I restrict attention to individuals who are at least 57 years old with at most four years of age difference between spouses. Households with cohabiting children or with members who are not wage workers at the first data entry are excluded. In turn, I exclude all households in which one or more members was registered as self-employed. I also exclude households who are net-borrowers (excluding private pension wealth), have total wealth above 15 million Danish kroner (≈ 3 million US dollars), have negative labor market income, or annual labor market income above the top two percent. Furthermore, I exclude households in which at least one member of the household is eligible for the Danish equivalence of a defined benefit plan (DB) in the US (in Danish »Tjenestemandspension«) or leaves the workforce through disability pension. Although alternative exit routes out of the labor market may affect how retirement behavior responds to ERP and OAP reforms, I abstract from other exit routes here to maintain tractability of the model.⁸

Combined, these criteria yield a total of 848,992 household-time observations, summarized in Table 1. Throughout the analysis, income and wealth are measured in 2008 prices. The change in OAP basis pension (B_t) is used to adjust income and wealth to 2008 levels. This measure has been chosen to make the implemented retirement scheme for 2008 compatible with the years before 2008. The annual change (ΔB_t) has been around 3 percent in the years 1998-2008, close to the Danish inflation rate.

Individuals who do not work are classified as retired. Labor market status is observed at the end of November each year. Timing problems regarding labor market income can arise since an individual retiring, say, in the beginning of November has potentially earned nearly a full year of labor market income while being classified as retired by this definition. Alternatively, retirement could be defined based on the labor

⁷The pension wealth data (PERE) have kindly been made available to me by the Danish Economic Council who also performed initial data preparations.

⁸[Iskhakov \(2010\)](#) estimates a dynamic model of retirement allowing for alternative exit routes in Norway while abstracting from the simultaneous retirement decision of husbands and wives.

Table 1 – Descriptive Statistics.

	Couples				Singles			
	Males		Females		Males		Females	
	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
Net-wealth [†]	4790	(2920)	4790	(2920)	1899	(1652)	1856	(1658)
Pension wealth [†]	1236	(1355)	728	(945)	682	(951)	744	(977)
Labor income [†]	254	(214)	181	(146)	197	(179)	190	(157)
GP visits	12.149	(17.655)	13.732	(16.905)	11.718	(18.591)	15.177	(20.137)
Age	61.176	(2.562)	60.114	(2.483)	60.625	(2.805)	60.886	(2.840)
Retirement age	61.715	(1.682)	60.982	(1.358)	61.412	(1.805)	61.660	(1.817)
Eligible for ERP	0.959	(0.197)	0.972	(0.164)	0.951	(0.215)	0.970	(0.172)
High skilled	0.242	(0.428)	0.220	(0.415)	0.196	(0.397)	0.273	(0.445)
Children	0.938	(0.242)	0.938	(0.242)	0.558	(0.497)	0.797	(0.402)
Grandchildren	0.535	(0.499)	0.535	(0.499)	0.224	(0.417)	0.437	(0.496)
Households	73,844		73,844		27,412		40,642	
Observations	343,076		343,076		131,422		209,713	

Notes: Table 1 reports means and standard deviations (in brackets) for single and married males and females. Net-wealth refers to household level net-wealth and is, thus, identical for husband and wives. The same is true for the indicators of children and grandchildren for couples. The number of observations refers to the number of non-missing household-time observations in the selected population.

[†] Numbers are in thousands of Danish kroner. The exchange rate is roughly 5.5 USD/DKK.

market income (or a combination of status and income) such that individuals with labor market income less than some threshold are defined as retired. That classification strategy has not been pursued since this requires choosing the income threshold.

Determining the ERP eligibility status is, unfortunately, not straight forward. For the cohorts used here, membership of an unemployment insurance (UI) fund for *at least* 10 years before retirement is required for eligibility to ERP. Since the register data do not cover the complete working life of the cohorts used, eligibility for ERP cannot be precisely identified. [Ejrnæs and Hochguertel \(2013\)](#) show, however, that Danish entrepreneurs respond strongly to changes in the eligibility requirements. They show that almost all who plan on being eligible to ERP plan on being so at the age of 60. Therefore, I assume that if an individual is a member of an UI fund in at least one year, she is eligible for early retirement at age 60.

Table 1 shows that couples on average have 2.5 times as much net-worth compared to single households. Out of total wealth, married men have significantly more private pension wealth than all other groups. The same is true for labor market income. Married men earn annually on average 254,000 Danish kroner (\$46,000) while married women earn only 181,000 Danish kroner (\$33,000). Single men earn 197,000 Danish kroner while single females earn 190,000 Danish kroner on average.

Married men are on average 61.2 years old, one year older than married women and half a year older than singles. This might explain why married men earn most and have accumulated more pension wealth. Furthermore, 24 percent of married men are high skilled while only 22 percent of married women and 20 percent of single males are high skilled. Interestingly, 27 percent of single women are high skilled. This potentially reflect increased bargaining power of women with higher education (Konrad and Lommerud, 2003).

Fewer males are eligible for ERP. Around 95 percent of men are eligible for ERP while 97 percent of women are eligible. The average retirement age of singles is 60.7 and slightly higher for married women around age 61. The average retirement age of married men is high, around 61.75, potentially due to the higher labor market income and lower share being eligible for ERP. Furthermore, since husbands are roughly one year older than their wives, the fact that they retire almost one year older might reflect simultaneous retirement of couples.

Men visit the GP less often compared to women. The average number of annual visits to the GP ranges from 11.7 (single males) to 15.2 (single females) potentially reflecting better health of men. This is, however, at odds with the fact that men on average have higher death rates than women for a given age (Table 6). Alternatively, this pattern might reflect that men are more reluctant to visit the GP. This explanation suggests caution when using the number of GP visits as a proxy for health status.

Most couples have children. Only 4 percent of couples are not registered as parents and 54 percent of couples have grandchildren. 56 percent of single men have children and 22 percent have grandchildren while 80 percent of single women have children and 44 percent have grandchildren. The Danish fertility registers link almost all children to their mothers while information on the fathers are less complete. This might result in fewer child-father matches of single males.

3.1 Retirement Patterns of Danish Households

Table 2 reports estimates from linear probability models (LPMs) of individual retirement. The dependent variables equal to one in the year of retirement. The explanatory variables are own, spousal and household level characteristics. The financial incentives in the Danish pension scheme, described above, result in retirement “spikes” at age 60 (ERP age), 62 (two-year rule), and 65 (OAP age). High skilled individuals retire significantly later than low skilled do, potentially reflecting the higher opportunity costs of retiring through foregone labor market earnings of high skilled individuals. High skilled might also have different preferences. In fact, when estimating the model of household consumption and retirement outlined in the following section, I find that high skilled value leisure *less* than low skilled.

Table 2 – Linear Probability Model (LPM) of Individual Retirement.

	Couples				Singles			
	Husband		Wife		Male		Female	
<i>Female variables:</i>								
Age=60	-0.011***	(0.002)	0.365***	(0.002)			0.261***	(0.003)
Age=61	0.004*	(0.002)	0.140***	(0.002)			0.107***	(0.002)
Age=62	-0.006**	(0.002)	0.147***	(0.002)			0.166***	(0.003)
Age=63	-0.002	(0.003)	0.095***	(0.003)			0.117***	(0.002)
Age=64	0.000	(0.003)	0.038***	(0.002)			0.055***	(0.002)
Age=65	-0.008*	(0.003)	0.051***	(0.003)			0.090***	(0.003)
Age=66	-0.003	(0.004)	0.038***	(0.003)			0.054***	(0.003)
Age=67	-0.003	(0.005)	0.024***	(0.003)			0.027***	(0.002)
Age=68	0.001	(0.006)	0.023***	(0.004)			0.023***	(0.003)
High skilled	0.001	(0.001)	-0.008***	(0.001)			-0.010***	(0.001)
GP ∈ [1, 10]	0.000	(0.002)	-0.000	(0.002)			0.005*	(0.002)
GP > 10	-0.003	(0.002)	0.003	(0.002)			0.006**	(0.002)
<i>Male variables:</i>								
Age=60	0.189***	(0.002)	-0.011***	(0.001)	0.305***	(0.004)		
Age=61	0.087***	(0.001)	0.001	(0.001)	0.107***	(0.003)		
Age=62	0.187***	(0.002)	0.005**	(0.002)	0.138***	(0.003)		
Age=63	0.129***	(0.002)	0.012***	(0.002)	0.093***	(0.003)		
Age=64	0.055***	(0.002)	0.005	(0.003)	0.045***	(0.002)		
Age=65	0.067***	(0.003)	-0.009**	(0.003)	0.081***	(0.004)		
Age=66	0.048***	(0.003)	-0.007*	(0.003)	0.049***	(0.003)		
Age=67	0.030***	(0.003)	-0.012***	(0.003)	0.027***	(0.003)		
Age=68	0.028***	(0.004)	-0.011**	(0.004)	0.023***	(0.004)		
High skilled	-0.014***	(0.001)	-0.004**	(0.001)	-0.015***	(0.002)		
GP ∈ [1, 10]	0.004**	(0.001)	-0.001	(0.001)	-0.001	(0.002)		
GP > 10	0.010***	(0.001)	-0.003	(0.001)	0.005**	(0.002)		
<i>Household variables:</i>								
Spouse retired	0.089***	(0.002)	0.080***	(0.002)				
Same educ.	0.003*	(0.001)	0.003*	(0.001)				
Log wealth	-0.005***	(0.001)	-0.003***	(0.001)	-0.002***	(0.000)	-0.002***	(0.000)
Child	-0.008***	(0.002)	-0.001	(0.002)	-0.009***	(0.002)	-0.003	(0.002)
Grandchild	0.004***	(0.001)	0.006***	(0.001)	0.005*	(0.002)	0.007***	(0.001)
Constant	0.082***	(0.011)	0.039***	(0.010)	0.033***	(0.006)	0.029***	(0.006)
Observations	323,630		323,630		125,755		204,020	
R ²	0.077		0.193		0.131		0.099	

Notes: Table 2 reports estimates from a linear probability model (LMP) of individual retirement on own and potential spousal characteristics. Robust standard errors are reported in brackets and *, **, *** indicates significance on the 10, 5 and 1 percent level, respectively.

Poor health, measured as more than ten GP visits within a year, increases the likelihood of retirement. Interestingly, for couples, only husbands are affected by their own poor health and poor health of a spouse does not affect the individual retirement behavior of either member of the household. Blau (1998) finds that wives in the RHS are more likely to retire if the husband is in poor health. The health care system in Denmark is universal and free and the need for spousal care giving when a spouse falls ill is arguably less pronounced in Denmark compared to the US.

The effect of own and spousal health on leisure complementarities cannot, however, be inferred from Table 2. Intra-household correlation between spousal health will result in health outcomes of, say, the husband being informative of the *potential* health outcome of the wife. If couples match with partners who have similar health attitudes and conditions or if household members tend to experience the same health risks such as smoking or less exercise, health shocks will be positively correlated between husbands and wives. If poor health is a potential driver of retirement, as Table 2 and existing studies suggest, this intra-household correlation will lead the estimated effect of spousal health on individual retirement probabilities to be upwards biased. In turn, the effect health has on the value of own and joint leisure cannot be uncovered without explicit assumptions about the intra-household correlation of health shocks.

The estimated correlations in Table 2 indicate that higher wealth reduce the likelihood of retirement. This is greatly at odds with the intuition that decreasing marginal utility of consumption will lead to leisure being more preferable for higher levels of wealth. Net-wealth is most likely increased while working and decreased during retirement. This will lead retirees to have on average lower wealth levels than workers and does not imply that higher levels of wealth leads to postponed retirement.

Children seem to reduce the probability of retirement. Because the estimates are based on elderly households older than 56, fertility is completed and the number of children is constant within each household. The effect of children on retirement is, thus, identified from differences in the retirement rates through the observed time periods between households with children and childless households. The negative effect of children can be interpreted as indicating that a larger proportion of childless households (compared to households with children) are observed to be retired at the last period in which they are observed (68 is the maximum age in the sample). Again, the estimated effects are composed of several underlying effects. Households with children might differ in their preferences for leisure but also their labor market experience and the opportunity cost of retirement is different. Grandchildren increase the likelihood of retirement by roughly as much as children reduce it. Women tend to be slightly more likely to retire in the presence of grandchildren compared to men.

It is more likely that an individual retires if his or hers spouse is retired. Table 2 indicates that the probability of retirement increases with 8 to 9 percentage points for

wives and husbands, respectively, if the spouse is retired. [Gustman and Steinmeier \(2000, 2004\)](#) also find that husbands are more affected by their wives being retired than the wives are affected by their husbands being retired. When estimating the economic model of leisure complementarities, I find, however, that a retired spouse increases the value of leisure for wives more than the value of leisure for husbands.

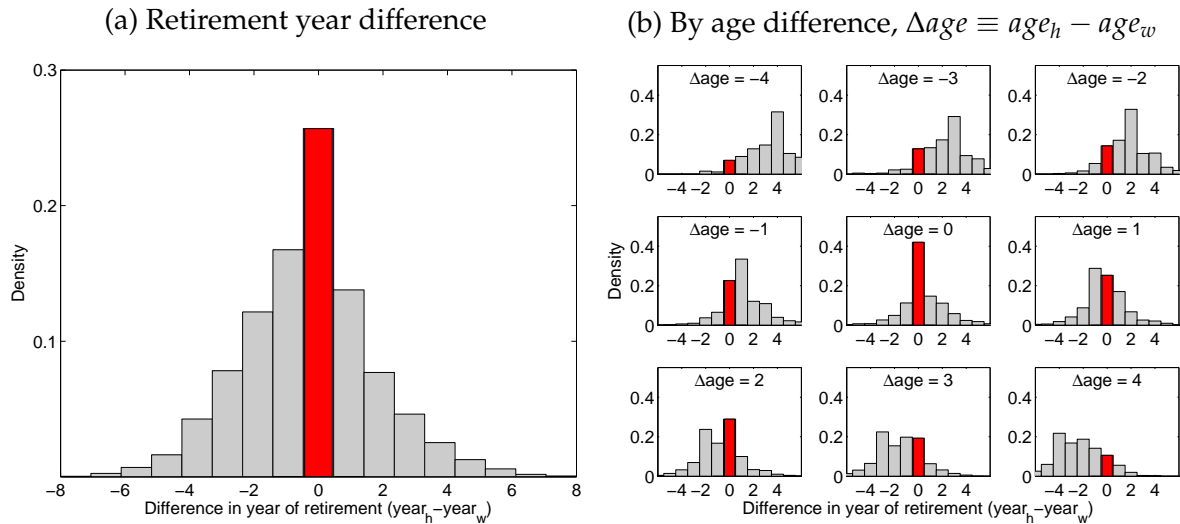


Figure 2 – Simultaneous Retirement.

Notes: Figure 2 illustrates the distribution of differences in retirement year in panel 2a. Figure 2b reports density plots of differences in retirement year within nine groups of households defined based on the age difference between the husband and wife, $\Delta age \equiv age_h - age_w$. Red bars mark simultaneous retirement of the husband and wife.

Figure 2 reports the *simultaneous retirement* pattern of Danish couples. The left panel presents the distribution of differences in year of retirement between husband and wives. There is a clear spike at zero (simultaneous retirement). Simultaneous retirement of couples could, however, be driven by households in which members are close in age. Figure 2b reports density plots of differences in retirement year within nine groups of households defined by the age difference between the husband and wife. Figure 2b shows that joint retirement is most pronounced for households close in age but is also observed for households who are further apart in age.

There is an asymmetry between households in which the husband is the youngest ($\Delta age < 0$) and households in which the wife is the youngest ($\Delta age > 0$). Particularly, when the wife is oldest she will tend to retire at the *latest* in the same year as her husband. This can be seen by the fact that there is almost no mass to the left of simultaneous retirement (marked by a red bar at zero) in figures where $\Delta age < 0$. There is considerably more mass to the right of simultaneous retirement in households where the husband is oldest ($\Delta age > 0$), suggesting that even if the husband is older than his wife, he might retire years later than her. This pattern could be explained by on

average higher income of the husband and, thus, a larger opportunity cost of retirement or husbands valuing joint leisure the most. Alternatively, and indistinguishably, husbands might simply dislike retiring *before* their wives due to, e.g., social norms.

Table 3 – Simultaneous Retirement (LPM).

	Estimate	Std.
Age difference (abs.)	-0.006***	(0.000)
Husband oldest	-0.003***	(0.000)
High skilled, husband	-0.001**	(0.000)
High skilled, wife	-0.001***	(0.000)
Same education	0.003***	(0.001)
Child	-0.001	(0.001)
Grandchild	0.001**	(0.000)
Poor health, husband [†]	0.001**	(0.000)
Poor health, wife [†]	-0.001*	(0.000)
Log wealth	0.000	(0.000)
Age, husband: 60	0.032***	(0.001)
Age, husband: 61	0.022***	(0.001)
Age, husband: 62	0.049***	(0.001)
Age, husband: 63	0.039***	(0.001)
Age, husband: 64	0.022***	(0.001)
Age, husband: 65	0.018***	(0.001)
Age, husband: 66	0.014***	(0.001)
Age, husband: 67	0.012***	(0.001)
Age, husband: 68	0.010***	(0.001)
Constant	0.006	(0.004)
Obs.	343,076	
R ²	0.017	

Notes: Results are based on a linear probability model (LPM) of the event of both spouses retiring in the same year. Standard errors in brackets are clustered on the household level and *, **, *** indicates significance on the 10, 5 and 1 percent level, respectively.

[†] Poor health is defined as more than ten visits to the GP within a given year.

To illustrate which factors might affect joint retirement, Table 3 reports estimates from a LPM of *simultaneous retirement* of couples. As suggested in Figure 2b, simultaneous retirement is more likely the closer the couple is in age. It is less likely that couples retire within the same year if the male is oldest, as discussed above. Low skilled households tend to retire simultaneously. If the husband and wife have the same educational level (low or high skilled), however, the likelihood of the couple retiring simultaneously increases. Interpreting the same educational level of both household members as a proxy for partner match quality, this indicates that couples who are a better match

enjoy leisure together more. Alternatively, this could be driven by intra-household correlation in labor market income, job security, or sector-specific shocks which both members are likely to be affected by since they have same educational background.

Children have no effect and grandchildren increase the likelihood of simultaneous retirement slightly. If the husband is in poor health, the likelihood of simultaneous retirement increases while it decreases if the wife is the one in poor health. Blau (1998) shows that wives are more likely to retire simultaneously if the husband is in poor health while the reverse is not as pronounced, in line with the estimates in Table 3.

Age dummies for the age of the husband suggest that especially when the husband is aged 62 is the likelihood of simultaneous retirement high. This is in line with the financial incentives because at age 62 the average husband will satisfy the two-year rule while the wife on average is eligible to receive ordinary ERP. If the husband is 63 years old the likelihood of simultaneous retirement is also higher. If both are working when the husband is 63 years old, the average couple will both satisfy the two-year rule and be eligible to receive higher levels of ERP with less severe means testing on their private pension wealth.

The raw correlations presented here do not necessarily translate into causal interpretable parameters or estimates of preferences such as leisure complementarities. To disentangle intra-household shocks from preferences over joint leisure, the following sections formulate, solve and estimate a rich life cycle model of consumption, saving and retirement choices of elderly Danish households.

4 A Dynamic Model of Household Consumption and Individual Retirement

Following van der Klaauw and Wolpin (2008), households maximize the expected discounted utility from *household* consumption, C_t , and leisure (retirement) for both husband and wife, $d_t^j \in \{0, 1\}$, $j = h, w$. $d_t^j = 1$ if household member j is *working* in the end of period t . Retirement is absorbing such that a retiree cannot re-enter the labor market once retired and only consumption choices are, thus, made in subsequent periods.⁹ Households die with certainty at age T and are forced to retire at age T_r . In each period, choices are based on the state of the household economy $(\mathbf{s}_t, \varepsilon_t)$, where \mathbf{s}_t contains states observable to both the household and the econometrician in beginning of period t while ε_t contains states only observed by the household.

The Bellman equation of working couples prior to forced retirement, $t < T_r$, is given

⁹This assumption is fairly standard in the literature and match the empirical regularity that very few is observed to re-enter the labor market once retired.

by

$$\begin{aligned}
 & V_t(\mathbf{s}_t, \varepsilon_t) = \\
 & \max_{C_t, d_t^h, d_t^w} \left\{ U(C_t, d_t^h, d_t^w, \mathbf{s}_t) + \sigma_\varepsilon \varepsilon_t (d_t^h, d_t^w) + \beta \mathbb{E}_t \left[\pi_{t+1}^h \pi_{t+1}^w V_{t+1}(\mathbf{s}_{t+1}, \varepsilon_{t+1}) \right. \right. \\
 & \quad \left. \left. + \pi_{t+1}^h (1 - \pi_{t+1}^w) V_{t+1}^h(\mathbf{s}_{t+1}^h, \varepsilon_{t+1}^h) + (1 - \pi_{t+1}^h) \pi_{t+1}^w V_{t+1}^w(\mathbf{s}_{t+1}^w, \varepsilon_{t+1}^w) \right. \right. \\
 & \quad \left. \left. + (1 - \pi_{t+1}^h)(1 - \pi_{t+1}^w) B(a_t) \right] \right\} \quad (2)
 \end{aligned}$$

where $U(\cdot)$ is a utility function, $B(\cdot)$ is a bequest function, β is the discount factor, π_s^j is the probability of household member j surviving to from period $s - 1$ to s , and $\mathbb{E}_s[\cdot]$ is the expectation operator conditional on information in period s . σ_ε is a scaling or “smoothing” parameter, determining how much of the observed retirement behavior is due to unobserved states. When $\sigma_\varepsilon \rightarrow \infty$, the choice of retirement is completely random and the variation in unobserved state variables determine optimal labor market participation. Contrary, as $\sigma_\varepsilon \rightarrow 0$, the observed state variables are the most important, and in the limit the model returns exact statements of when to retire. In between, $\sigma_\varepsilon \in (0, \infty)$, the model returns probabilistic statements on how likely it is to observe retirement given the state variables, \mathbf{s}_t , described below.

Households solve (2) subject to the household budget constraint,

$$C_t + a_t = Ra_{t-1} + \sum_{j \in h, w} \{T(y_t^j, y_t^{-j}) + P^j(\mathbf{s}_t)\}, \quad (3)$$

where a_t is end-of-period household wealth restricted to be non-negative in all periods, $a_s \geq 0, \forall s$. $T(y_t^j, y_t^{-j})$ is after-tax income of household member j depending also on the income of the spouse, y_t^{-j} , and $P^j(\mathbf{s}_t)$ is potential pension benefits (ERP and OAP) received by member j as a function of labor market status of each spouse, eligibility for early retirement and individual private pension wealth.

Couples do not divorce and singles do not (re)marry. This assumption is invoked for computational convenience since the model for single households can, then, be solved in a first step without any regard to optimal behavior of couples. The resulting bellman equation for a single individual j working in beginning of period t is

$$V_t^j(\mathbf{s}_t^j, \varepsilon_t^j) = \max_{C_t, d_t^j} \left\{ U^j(C_t, d_t^j, \mathbf{s}_t^j) + \sigma_\varepsilon \varepsilon_t (d_t^j) + \beta \mathbb{E}_t \left[\pi_{t+1}^j V_{t+1}^j(\mathbf{s}_{t+1}^j, \varepsilon_{t+1}^j) + (1 - \pi_{t+1}^j) B(a_t) \right] \right\},$$

where the continuation value of couples do not enter. The bellman equation of a single

individual j retired in beginning of period t is given by

$$V_t^j(\mathbf{s}_t^j) = \max_{C_t, d_t^j} \left\{ U^j(C_t, 0, \mathbf{s}_t^j) + \beta \mathbb{E}_t \left[\pi_{t+1}^j V_{t+1}^j(\mathbf{s}_{t+1}^j, \mathbf{e}_{st+1}^j) + (1 - \pi_{t+1}^j) B(a_t) \right] \right\},$$

due to the absorbing nature of retirement. Similarly for couples, if either of the spouses are retired in beginning of period t .

4.1 Preferences

Following [van der Klaauw and Wolpin \(2008\)](#), the bequest function is given by $B(a_t) = \gamma a_t$. Preferences are constant relative relative risk aversion (CRRA) and for singles given as,

$$U^j(C_t, d_t^j, \mathbf{s}_t^j) = \frac{C_t^{1-\rho} - 1}{1-\rho} + \alpha^j(\mathbf{s}_t^j) \mathbb{1}_{\{d_t^j=0\}},$$

where ρ is the risk aversion parameter and $\alpha^j(\mathbf{s}_t^j)$ summarizes the value of leisure when retired. Preferences of couples are a weighted sum of the member's individual utility,

$$\begin{aligned} U(C_t, d_t^h, d_t^w, \mathbf{s}_t) &= \lambda \left[\frac{(C_t/n_t)^{1-\rho} - 1}{1-\rho} + \alpha^h(\mathbf{s}_t^h) \mathbb{1}_{\{d_t^h=0\}} \left(1 + \phi^h(\mathbf{s}_t) \mathbb{1}_{\{d_t^h=0, d_t^w=0\}} \right) \right] \\ &+ (1-\lambda) \left[\frac{(C_t/n_t)^{1-\rho} - 1}{1-\rho} + \alpha^w(\mathbf{s}_t^w) \mathbb{1}_{\{d_t^w=0\}} \left(1 + \phi^w(\mathbf{s}_t) \mathbb{1}_{\{d_t^w=0, d_t^h=0\}} \right) \right] \end{aligned}$$

where λ is the relative weight on the preferences of the husband in the household decision process and $n_t = 1 + \nu(\text{ADULTS}_t - 1)$ is an equivalence scale.

Complementarities in leisure is measured by $\phi^j(\mathbf{s}_t)$. $\phi^j(\mathbf{s}_t)$ summarizes the *percentage* increase in utility from leisure when the spouse is also retired as a function of household characteristics. Importantly, utility from joint leisure is received in all periods in which both spouses is retired. Complementarities in leisure will, thus, tend to increase the tendency of couples to retire within few years of each other, as observed in the existing literature and in Figure 2 for Danish households.

4.2 State Variables and State Transitions

The relevant states of the economy, \mathbf{s}_t , is given by

$$\mathbf{s}_t = (age_t^h, age_t^w, d_t^h, d_t^w, y_t^h, y_t^w, \wp_t^h, \wp_t^w, e^h, e^w, h_t^h, h_t^w, elig^h, elig^w, ch, gt, a_{t-1}),$$

where

$age_t^j \in \{57, 110\}$:	Age of household member j ,
$d_{t-1}^j \in \{0, 1\}$:	Labor market status in beginning of period t . 0: retired, 1: working,
$y_t^j \in \mathbb{R}_+$:	Labor earnings of member j through period t
$\phi_t^j \in \mathbb{R}_+$:	Private pension funds in the end of period t of member j .
$e^j \in \{0, 1\}$:	Educational level of member j . 0: low skilled, 1: high skilled,
$h_t^j \in \{0, 1\}$:	Health status of member j in period t . 0: Poor, 1: Good,
$elig^j \in \{0, 1\}$:	Eligibility of member j to ERP at age 60,
$ch \in \{0, 1\}$:	Children related to household. 0: no, 1: yes,
$g_t \in \{0, 1\}$:	Grandchildren related to household. 0: no, 1: yes,
$a_{t-1} \in \mathbb{R}_+$:	Household wealth in beginning of period t .

The transition probability of observed state variables (including wealth) and unobserved state variables, ε , is assumed to be *conditionally independent* (CI). The joint transition density of the controlled process for $(\varepsilon_t, \mathbf{s}_t)$, thus, factors as

$$p(\varepsilon_{t+1}, \mathbf{s}_{t+1} | \varepsilon_t, \mathbf{s}_t, C_t, d_t) = p(\varepsilon_{t+1} | \mathbf{s}_{t+1}) p(\mathbf{s}_{t+1} | \mathbf{s}_t, C_t, d_t), \quad (4)$$

implying that all dynamics run through the observed state variables only. This assumption is common and applied for tractability. Below, I describe how the observed states, \mathbf{s}_t , transition over time.

Private pension wealth in beginning of period t , ϕ_t^j , is assumed to be a fraction, ζ_t , of total household wealth in the beginning of period t ,

$$\phi_t^j = a_{t-1} \zeta (age_t^j, age_t^{-j}, e^j, e^{-j}, elig^j, elig^{-j}, ch, a_{t-1}), \quad (5)$$

as a function of age, educational level, children, and net-wealth.

Underlying the binary health indicator, h , a latent (continuous) health indicator is assumed to follow a two-dimensional (Gaussian) Markov process,

$$h_{t+1}^{h*} = \delta_0^h + \delta_1^h h_t^{h*} + \eta_{t+1}^h, \quad (6)$$

$$h_{t+1}^{w*} = \delta_0^w + \delta_1^w h_t^{w*} + \eta_{t+1}^w, \quad (7)$$

where the husband's and wife's health shocks are potentially correlated,

$$\begin{pmatrix} \eta_t^h \\ \eta_t^w \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{h,h}^2 & \sigma_{h,hw} \\ \sigma_{h,hw} & \sigma_{h,w}^2 \end{pmatrix} \right),$$

accounting for adverse health shocks affecting both spouses.

This specification is motivated by partners matching on lifestyle attributes that might also affect health. This could be smokers matching with other smokers or if

people enjoying exercising also search for partners with a similar passion. Both examples suggest a positive correlation, $\sigma_{h,hw} \geq 0$, between health shocks. Incorporating this intra-household correlation in health shocks has, to the best of my knowledge, not been done in previous studies on the joint retirement of couples. Ignoring a positive (negative) intra-household correlation between health shocks would, however, result in over (under) estimation of leisure complementarities.

The labor market income process is given by a Mincer-type equation,

$$y_t^j = \exp(\gamma_0^j + \gamma_1^j e_j + \gamma_2^j age_t^j + \gamma_3^j (age_t^j)^2 + \gamma_4^j ch + v_t^j) \mathbb{1}_{\{d_t^j=1\}}, \quad (8)$$

where a polynomial in age proxy for potential labor market experience. As with health shocks, unforeseen shocks to labor market income is allowed to be correlated within the household,

$$\begin{pmatrix} v_t^h \\ v_t^w \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{y,h}^2 & \sigma_{y,hw} \\ \sigma_{y,hw} & \sigma_{y,w}^2 \end{pmatrix} \right).$$

Intra-household correlation between spousal income is typically ignored in existing studies of the joint retirement behavior of couples. For example, [Gustman and Steinmeier \(2000, 2004\)](#) do exclude income uncertainty while [Bingley and Lanot \(2007\)](#) and [Scholz and Seshadri \(2013\)](#) model *household* income and do, therefore, not account for the fact that couples are likely to be subject to similar labor market uncertainty. This would happen if couples match on, e.g., educational characteristics because couples would subsequently be likely to work in similar sectors of the economy and, thus, subject to some of the same labor market shocks. [Lewis and Oppenheimer \(2000\)](#) and [Nielsen and Svarer \(2009\)](#) find supporting evidence that US and Danish couples, respectively, match on similar educational background. Ignoring the (positive) intra-household correlation in income shocks will, thus, result in over estimation of the value of joint retirement and leisure complementarities.

If households have children they might also have grandchildren. Both children and grandchildren might affect their valuation of leisure time. Whether a household has grandchildren is denoted by $g \in \{0, 1\}$ where zero indicates “no grandchildren” and one indicates “grandchildren”. Grandchildren arrive probabilistically with a known probability. The likelihood of having a grandchild next period is given by

$$\Pr(g_{t+1} = 1 | \mathbf{s}_t) = \begin{cases} f(\widetilde{age}_t; \psi) & \text{if } ch = 1, g_t = 0, \\ 1 & \text{if } ch = 1, g_t = 1, \\ 0 & \text{else,} \end{cases} \quad (9)$$

where $\widetilde{age}_t = age_t^w$ unless it is a single male household, then $\widetilde{age}_t = age_t^h$. $f(\cdot, \psi)$ is a flexible function of age with ψ being a parameter vector.

5 Solving the Model

Here, I briefly discuss the solution strategy and refer to the supplemental material for a detailed description of the numerical solution. The life cycle model outlined above includes both a continuous consumption/savings choice and a discrete retirement choice. The joint decision problem can be separated into two decisions such that optimal consumption and saving is found conditional on a given retirement choice. For notational simplicity, I illustrate the idea using the model for single households. The bellman equation for a single individual j is given by

$$V_t^j(\mathbf{s}_t^j, \varepsilon_t^j) = \max_{C_t^j, d_t^j} \{U^j(C_t^j, d_t^j, \mathbf{s}_t^j) + \varepsilon(d_t^j) + \beta \mathbb{E}_t[\pi_{t+1}^j V_{t+1}^j(\mathbf{s}_{t+1}^j, \varepsilon_{t+1}^j) + (1 - \pi_{t+1}^j)B(a_t)]\},$$

where the optimization problem can be broken into two parts related to the consumption and the retirement decision,

$$\begin{aligned} v_t^j(\mathbf{s}_t^j, \varepsilon_t^j) &= \max_{d_t^j} \{v_t^j(\mathbf{s}_t^j, d_t^j) + \varepsilon(d_t^j)\}, \\ v_t^j(\mathbf{s}_t^j, d_t^j) &= \max_{C_t^j} \{U^j(C_t^j, d_t^j, \mathbf{s}_t^j) + \beta \int_{\mathbf{s}_{t+1}^j} \int_{\varepsilon_{t+1}^j} \left[\pi_{t+1}^j V_{t+1}^j(\mathbf{s}_{t+1}^j, \varepsilon_{t+1}^j) \right. \\ &\quad \left. + (1 - \pi_{t+1}^j)B(a_t) \right] p(d\varepsilon_{t+1}^j | \mathbf{s}_{t+1}^j) p(d\mathbf{s}_{t+1}^j | \mathbf{s}_t^j, C_t^j, d_t^j)\}. \end{aligned}$$

$v_t^j(\mathbf{s}_t^j, d_t^j)$ is a choice-specific value function related to a given labor market choice, d_t^j , conditional on optimal consumption under this labor market choice. $p(d\varepsilon_{t+1}^j | \mathbf{s}_{t+1}^j)$ and $p(d\mathbf{s}_{t+1}^j | \mathbf{s}_t^j, C_t^j, d_t^j)$ are the transition densities related to the unobserved and observed states, respectively.

When ε is Extreme Value type I distributed, as assumed here, the integral with respect to the unobserved state has a convenient closed form solution (Rust, 1987). Dropping j -superscripts for notational convenience, the choice-specific value functions becomes

$$v_t(\mathbf{s}_t, d_t) = \max_{C_t} \left\{ U(C_t, d_t, \mathbf{s}_t) + \beta \int_{\mathbf{s}_{t+1}} \pi_{t+1} EV_{t+1}(\mathbf{s}_{t+1}) p(d\mathbf{s}_{t+1} | \mathbf{s}_t, C_t, d_t) + \beta(1 - \pi_{t+1})B(a_t) \right\}, \quad (10)$$

$$EV_{t+1}(\mathbf{s}_{t+1}) = \sigma_\varepsilon \log \left\{ \sum_{k \in \{0,1\}} \exp(v_{t+1}(\mathbf{s}_{t+1}, k) / \sigma_\varepsilon) \right\}.$$

Since households optimize over a discrete retirement choice, the solution is typically found using value function iterations (VFI). VFI is, however, often time-consuming

because expectations over all future states (observed and unobserved) must be evaluated for a given *guess* of optimal consumption and labor market choice. This is done for all combinations of the state space. The state-space here is large because both spouses education, eligibility for ERP, health state, and age influence the household choices. Further, the household level of wealth and the presence of children and grandchildren also affect the decision process. To keep estimation of the rich household model tractable, I instead build on the endogenous grid method (EGM), proposed by [Carroll \(2006\)](#) to efficiently solve consumption models *without discrete choices*. The EGM relies on the first order condition (FOC) of the bellman equation (10) which, together with the standard envelope condition, yields consumption Euler equations.

When households also perform discrete choices, however, the FOC, and hence the Euler equation, is only *necessary but not sufficient* for an optimal consumption choice. Particularly, the choice-specific value function in (10) is not strictly concave and multiple levels of consumption might satisfy the Euler equation. [Fella \(2014\)](#) generalizes the EGM to handle discrete choices by utilizing the generalized envelope theorem in [Clausen and Strub \(2013\)](#). The approach in [Fella \(2014\)](#) is, however, framed in a setting in which there is *no* unobserved states, $\sigma_\varepsilon = 0$. [Iskhakov, Jørgensen, Rust and Schjerning \(2014\)](#) show that the EGM finds *all* solutions to the Euler equation when households perform both a continuous and a discrete decision with unobserved discrete-choice specific state variables, $\sigma_\varepsilon \geq 0$.

[Iskhakov, Jørgensen, Rust and Schjerning \(2014\)](#) show that their discrete-continuous EGM (DC-EGM) is well suited for estimation of models of consumption and labor market choices, similar to the present model. Particularly, they show how the variance of the unobserved state, σ_ε , smooths the problem. Furthermore, expectations about future income, grandchildren arrival, and health will further smooth the problem and potentially to such a degree that the multiplicity of the solution to the Euler equation vanishes. With “sufficient” smoothing relative to *i*) the difference in income from working and retiring, and *ii*) the value of leisure, the FOC will be *sufficient* even in cases with discrete choices and the standard EGM proposed by [Carroll \(2006\)](#) applies.

I have implemented the DC-EGM algorithm and find that it was never the case that EGM found several solutions to the Euler equation. This is basically because the Danish retirement pension scheme is very generous. The median drop in income when going from wage work to retirement is only around 10 percent and the associated drop in optimal consumption when retiring is relatively small. Section B in the online supplementary material discusses in some detail the DC-EGM approach.

6 Initial Estimations and Calibrations

As is common in the literature on intertemporal consumption allocation, I assume a fixed gross real interest rate of 3 percent, $R = 1.03$, (Gourinchas and Parker, 2002) slightly lower than an interest rate of five percent in van der Klaauw and Wolpin (2008). The discount factor is fixed at $\beta = .98$ and the bequest motive is calibrated to be $\gamma = 0.08$ based on initial investigations of the model for singles. Finally, I calibrate the intra-household power distribution between husband and wife such that utility of each household member enters equally, $\lambda = .5$. Below, I describe how several features of the model and state transition processes are estimated using the Danish register data.

Private Pension Wealth. Private pension wealth of household member j is assumed to be a fraction of *total* household wealth, $\varphi_t^j = a_{t-1}\zeta_t^j$. The fraction of household wealth, ζ_t^j , is a function of both household members characteristics as described in equation (5). I explicitly handle the fact that the share is bounded between zero and one by a censored regression approach (Tobin, 1958) ensuring that predicted private pension wealth cannot exceed total household wealth.

Table 4 – Private Pension Funds Share of Total Household Wealth, ζ .

	Couples				Singles			
	Husband		Wife		Male		Female	
Age	0.072	(0.023)	0.039	(0.020)	0.216	(0.069)	0.098	(0.056)
Age ² /100	-0.068	(0.019)	-0.037	(0.016)	-0.187	(0.056)	-0.091	(0.045)
High skilled	0.069	(0.001)	0.131	(0.001)	0.142	(0.004)	0.185	(0.003)
Child	0.026	(0.002)	-0.024	(0.002)	0.019	(0.003)	-0.032	(0.003)
Log wealth	8.864	(0.275)	4.290	(0.266)	12.057	(0.986)	10.062	(0.427)
(Log wealth) ²	-0.655	(0.019)	-0.327	(0.018)	-0.920	(0.075)	-0.732	(0.032)
(Log wealth) ³	0.016	(0.000)	0.008	(0.000)	0.023	(0.002)	0.018	(0.001)
Constant	-41.161	(1.502)	-19.000	(1.420)	-57.670	(4.662)	-47.565	(2.552)
sigma	0.194	(0.000)	0.145	(0.000)	0.324	(0.001)	0.350	(0.001)
Obs.	209035		160963		63500		108893	
R ²	0.507		0.221		0.185		0.144	

Notes: Table 4 reports estimation results from a bottom and top censored regression (Tobin, 1958) with robust standard errors in brackets. Results are based on 60-65 year old working individuals who are eligible for ERP actual share of private pension wealth relative to total household wealth, denoted ζ in the text.

Only the ERP system and not the OAP (available at age 65) are means tested based on private pension wealth. Therefore, only individuals eligible for ERP aged 60 through 65 years old are included in the estimation of ζ_t^j . Table 4 reports separate estimation results for males and females who are single and married. The predicted level of pension

wealth is quite close to the average age profile of actual pension wealth (not reported) supported by a relatively high pseudo R^2 of around 17 to 30 percent.

Labor Market Income. Households form expectations over each household member's future income, if working, as specified in equation (8). For single individuals, simple Mincer type log-wage equations with educational and children dummies along with a polynomial in age is estimated separately for males and females. For couples, intra-household correlation in income shocks are allowed through an iterated seemingly unrelated regression (SUR) approach. Table 5 reports the parameter estimates.

As expected, income is quadratic in age, which proxy for potential labor market experience. High skilled have on average higher labor market income and individuals with children tend to have a higher wage. The labor market income of married females is, however, lower if they have children. This is consistent with lower female human capital accumulation while on maternity leave. The intra-household correlation in income shocks is positive but low, around $\frac{\sigma_{hw}}{\sigma_h\sigma_w} \approx 0.04$.

Table 5 – Labor Market Income Process.

	Couples				Singles			
	Husband		Wife		Male		Female	
High skilled	0.262	(0.003)	0.318	(0.003)	0.230	(0.006)	0.248	(0.004)
Age _t ^j	0.629	(0.033)	0.544	(0.043)	0.934	(0.052)	1.036	(0.035)
(Age _t ^j) ² /100	-0.532	(0.001)	-0.453	(0.001)	-0.770	(0.001)	-0.856	(0.001)
Children ^j	0.060	(0.006)	-0.018	(0.006)	0.151	(0.005)	0.021	(0.004)
Constant	-5.999	(0.001)	-4.002	(0.001)	-15.956	(0.003)	-18.937	(0.002)
σ_j^2	0.288		0.347		0.544		0.399	
σ_{hw}	0.011		0.011					
R^2	0.047		0.050		0.030		0.037	
Obs.	175,293		175,293		81,697		131,811	

Notes: Table 5 reports the estimated labor market income process, specified in equation (8). Robust standard errors in brackets. The income equations for husband and wife are allowed to be correlated and estimated simultaneously by iterated seemingly unrelated (SUR) regression.

The income process is estimated using only working households because labor market income is assumed to be zero in retirement. In reality, labor market income in the year in which an individual have transitioned into retirement by the end of November, where the labor market status is recorded in the registers, is often greater than zero. Because information is annual and labor market status is observed at a given date in November each year, there will be a tendency to underestimate labor market income in the year of retirement. This underestimation of labor market income in the retirement

state will tend to inflate the estimated value of leisure. This approach is, however, a parsimonious description of the data while keeping the model tractable.

Arrival of Grandchildren. To estimate the arrival probability of the first grandchild, $\Pr(g_{t+1} = 1 | \mathbf{s}_t)$ in equation (9), I link all households in 2011 with their children and potential grandchildren in the Danish registers. The arrival probability is estimated using a linear probability model with age dummies on households with children and no grandchildren. If they do not have children at age 40 in 2011, I assume they will never have children and, hence, never have grandchildren. Figure 3 reports the estimated arrival probability of the first grandchild as a function of age. Results suggest that if households do not have grandchildren when they turn around 80 years old, they will never have grandchildren. This is a bi-product of the rapidly declining female fertility after age 40.




Figure 3 – Arrival Probability of First Grandchild, $\Pr(g_t = 1 | \mathbf{s}_{t-1})$, in Households with Children ($ch = 1$).

Survival Probability. The conditional survival probability from period $t - 1$ to t is given by one minus the death probability. The death probability is modeled as an exponential in age,

$$\pi_t^j = \min\{1, (1 - \min\{1, \exp(\pi_0^j + \pi_1^j \text{age}_t^j)\}) \cdot (\pi_2^j \mathbb{1}_{\{h_t^j=0\}} + \pi_3^j \mathbb{1}_{\{h_t^j=1\}})\},$$

and adjusted for health status, $h_t^j \in \{0, 1\}$, where $h_t^j = 1$ indicates “good health” and $h_t^j = 0$ indicates “poor health”. The baseline death probability is estimated on aggre-

gate numbers for the Danish population and reported in Table 6. The fit of the “model” is extremely good and, as expected, the death probability is always greater for males.¹⁰

Table 6 – Death Probability Estimates.

	Males	Females
	Estimate (SE)	Estimate (SE)
<i>constant</i> (π_0^j)	-10.338 (.036)***	-11.142 (.039)***
<i>age</i> (π_1^j)	.097 (.001)***	.103 (.001)***
\bar{R}^2	.996	.996
#Obs	245	245

Data is based on Statistics Denmark’s series BEF5 and FOD207 for the years 2006-2010. Robust standard errors reported. *: $p < .05$, **: $p < .001$, ***: $p < .0001$.

Poor health is calibrated to decrease the survival probability with two percent ($\pi_2^h = \pi_2^w = .98$) while good health increases the survival probability with two percent relative to the baseline survival probability ($\pi_3^h = \pi_3^w = 1.02$). The resulting survival probabilities are reported in Figure 4.




Figure 4 – Survival Probabilities as Function of Age, Health and Gender.

Health Transition Probabilities. The parameters of the latent health process, h_t^{j*} in equations (6) and (7), are estimated using the annual number of visits to the general

¹⁰The data used for estimation are the time tables BEF5 and FOD207 supplied by Statistics Denmark. The out-of sample predictions (individuals aged 99 or older) are in line with the actual probabilities of death since the oldest males in 2010 were 105 years old and the oldest females were 108 years old.

practitioner (GP) for each household member as a proxy for the latent health process. The results are reported in the top panel of Table 7 for singles and couples. Intra-household correlation in health shocks are allowed through an iterated seemingly unrelated regression (SUR) approach. There is a strong time dependence in the number of annual GP visits, reflected in an AR coefficient between .5 and .6. The intra-household correlation is positive but low, around 0.06.

Table 7 – Latent Health Process and Health Transition Probabilities.

<i>Latent Health Process, h_t^* (GP visits)[†]</i>				
	Couples		Singles	
	Husband	Wife	Male	Female
h_{t-1}^j	0.508 (0.002)	0.532 (0.002)	0.567 (0.003)	0.579 (0.002)
Constant	6.715 (0.031)	6.876 (0.029)	5.811 (0.053)	7.052 (0.042)
σ_j^2	226.852	209.574	256.221	273.689
σ_{hw}	13.397			
R^2	0.252	0.264	0.280	0.317
Obs.	243706	243706	91467	151993

<i>Binary Health Indicator Transition Probabilities[‡]</i>				
	Couples			
	\checkmark $h_t^h = 0, h_t^w = 0$	$h_t^h = 1, h_t^w = 0$	$h_t^h = 0, h_t^w = 1$	$h_t^h = 1, h_t^w = 1$
$h_{t+1}^h = 0, h_{t+1}^w = 0$	0.791	0.090	0.104	0.015
$h_{t+1}^h = 1, h_{t+1}^w = 0$	0.094	0.787	0.010	0.109
$h_{t+1}^h = 0, h_{t+1}^w = 1$	0.109	0.010	0.787	0.094
$h_{t+1}^h = 1, h_{t+1}^w = 1$	0.015	0.104	0.090	0.791

	Single Males		Single Females	
	\checkmark $h_t^h = 0$	$h_t^h = 1$	$h_t^w = 0$	$h_t^w = 1$
$h_{t+1}^j = 0$	0.916	0.084	0.922	0.078
$h_{t+1}^j = 1$	0.084	0.916	0.078	0.922

Notes: Table 7 reports estimation results of the latent health process approximated as annual visits to the general practitioner (GP) in the top panel. In the bottom panel, a transition matrix based on these estimates are presented for couples and singles. Results are based on 50–68 year old households.

[†] Parameter estimates of the latent health process, h^* , in eqs. (6) and (7) are based on Seemingly Unrelated Regression (SUR) from a two-equation process of annual visits the general practitioner (GP), one for both each spouse. For singles, OLS estimates are reported.

[‡] Based on the estimated latent health indicator processes, h^* , binary transition probabilities are constructed using the Markov-chain approximation method proposed in Terry and Knotek (2011). $h_t^j = 0$ indicates "poor health" while $h_t^j = 1$ indicates "good health". The transition direction is indicated with an arrow (\checkmark) and goes from the column to the row value.

To calibrate the transition probabilities across the binary health indicator, $h_t^j \in \{0, 1\}$, of each spouse, I apply the discrete Markov-chain approximation proposed in [Terry and Knotek \(2011\)](#).¹¹ The resulting transition probabilities are reported for couples and singles in the lower panel of Table 7. The estimated probabilities of remaining in a given health state are a bit low (around 90 percent) but quite close to the frequency-based estimates around 92 percent, reported in [Iskhakov \(2010\)](#). He also specifies a latent health indicator, when studying exit routes of Norwegian individuals.

[Iskhakov \(2010\)](#) argues that measurement error in health status classification biases his frequency-based probability of remaining in a given health state and finds larger estimates (around .97) when accounting for potentially miss-categorized health-statuses. This suggests that the health process used herein, reported in Table 7, might be slightly less persistent than the actual health status.

7 Estimation of Preference Parameters

To estimate the remaining preference parameters, $\theta = (\rho, \nu, \sigma_\varepsilon, \alpha(\mathbf{s}), \phi(\mathbf{s}))$, I combine information on the discrete choice of retirement with the continuous consumption choice of households. The estimation approach applied here is an extension with discrete choices to the method used in [Jørgensen \(2013, 2014\)](#) and studied in [Iskhakov, Jørgensen, Rust and Schjerning \(2014\)](#). Specifically, the assumption of Extreme Value Type I distributed unobserved shocks, ε , combined with a distributional assumption on measurement error in household consumption facilitate a joint maximum likelihood (ML) estimation approach.¹²

The Danish registers contain accurate measures of all relevant state variables. Specifically, wealth and income measures enable calculation of the available resources within households, a key determinant of consumption and retirement behavior. The imputed measure of consumption from income and wealth information in the Danish registers are, however, likely to be observed with measurement error. Assuming that measurement error is additive, C_{it}^{data} is observed household consumption at the end of period t , given as,

$$C_{it}^{data} = C_{it}^* + \zeta_{it},$$

where ζ_{it} is a noise term contaminating the “true” model-consistent consumption level, C_{it}^* . The model-consistent level of consumption is found numerically for a given set of model parameters, θ .

¹¹I use their online available Matlab code. [Tauchen \(1986\)](#) also proposed an extensively used approximation method which relies on the covariance matrix of the health shocks to be diagonal. The approach in [Terry and Knotek \(2011\)](#) allows for correlation between spouses, i.e., a non-zero off-diagonal element.

¹²[Fafchamps and Pender \(1997\)](#) apply a somewhat similar ML estimator in their analysis of (lack of) investments in irrigation wells in India.

Let $C_{t,\text{RetAge}}^*(\mathbf{s}_{it}^{\text{data}}; \theta | d_{it-1} = 0)$ denote the optimal consumption of household i in period t if retired at age RetAge , conditional on being retired at the beginning of period t .¹³ If working in beginning of period t , denote optimal consumption as $C_t^*(\mathbf{s}_{it}^{\text{data}}; \theta | d_{it-1} = 1)$. In what follows, I drop the “data” superscript and simply indicate observations from data with a household i subscript (time subscripts are on both data and model variables). The error-component of consumption can be expressed as a function of the parameters as,

$$\zeta_{it}(\theta) = \begin{cases} C_{it} - C_t^*(\mathbf{s}_{it}; \theta | d_{it-1} = 1) & \text{if } d_{it-1} = 1 \\ C_{it} - C_{t,\text{RetAge}}^*(\mathbf{s}_{it}; \theta | d_{it-1} = 0) & \text{if } d_{it-1} = 0 \end{cases}$$

Since unobserved states are (assumed) independent of the measurement error in consumption, the joint likelihood of observing a given retirement choice and household consumption level in household i in period t is given by

$$\ell_{it}(\theta) = \Pr(d_{it} | \mathbf{s}_{it}; \theta)^{\mathbb{1}_{\{d_{it-1}=1\}}} f(\zeta_{it}(\theta)),$$

where $\Pr(d_{it} | \mathbf{s}_{it}; \theta)$ is the probability of observing retirement choice d_{it} conditional on observed states, \mathbf{s}_{it} , and $f(\cdot)$ is the probability distribution of the measurement error related to the observed states and retirement choice. The Extreme Value type I assumption of unobserved states, ε , imply that choice probabilities (when working) are dynamic multinomial logits (Rust, 1987),

$$\Pr(d_{it} | \mathbf{s}_{it}; \theta) = \frac{\exp(v_t(\mathbf{s}_{it}, d_{it}; \theta) / \sigma_\varepsilon)}{\sum_{k \in \mathcal{D}(\mathbf{s}_{it})} \exp(v_t(\mathbf{s}_{it}, k; \theta) / \sigma_\varepsilon)},$$

where $v_t(\mathbf{s}_{it}, k; \theta)$ is the choice-specific indirect utility functions from a given choice k , conditional on subsequent optimal consumption/savings behavior given the choice k .

Let the measurement error in consumption be *iid* Normal with mean zero and variance σ_ζ^2 . The joint (mean) log-likelihood function is, then, given by

$$\tilde{\mathcal{L}}(\theta, \sigma_\zeta^2) = \sum_{i=1}^N \frac{1}{NT_i} \sum_{t=1}^{T_i} \left\{ \mathbb{1}_{\{d_{it-1}=1\}} \log(\Pr(d_{it} | \mathbf{s}_{it}; \theta)) + \log \left(\frac{1}{\sqrt{2\pi\sigma_\zeta}} \exp \left[-\frac{1}{2} \frac{\zeta_{it}(\theta)^2}{\sigma_\zeta^2} \right] \right) \right\},$$

where rearranging the first order condition with respect to σ_ζ^2 yields the standard ML estimator for the error variance, $\hat{\sigma}_\zeta^2(\theta) = \sum_{i=1}^N \frac{1}{NT_i} \sum_{t=1}^{T_i} \zeta_{it}(\theta)^2$. Inserting this estimator,

¹³Recall, the Danish ERP pension benefits depend on at what age and individual retired.

the concentrated (mean) log-likelihood function becomes

$$\mathcal{L}(\theta) = \Xi + \sum_{i=1}^N \frac{1}{NT_i} \sum_{t=1}^{T_i} \left\{ \mathbb{1}_{\{d_{it-1}=1\}} [v_t(\mathbf{s}_{it}, d_{it}; \theta) - EV_t(\mathbf{s}_{it}; \theta)] / \sigma_\varepsilon - \frac{1}{2} \log \left(\hat{\sigma}_\varepsilon^2(\theta) \right) \right\}, \quad (11)$$

where $\Xi = -\frac{1}{2} (\log(2\pi) + 1)$ is a constant and

$$EV_t(\mathbf{s}_{it}; \theta) = \sigma_\varepsilon \log \left[\sum_{k \in \mathcal{D}(\mathbf{s}_{it})} \exp(v_t(\mathbf{s}_{it}, k; \theta) / \sigma_\varepsilon) \right],$$

is the log-sum of choice-specific value functions. The parameter vector that maximizes (11), is the ML estimator of the model parameters. To compute this estimator in practice, I successively solve the model for all trial values of θ , as proposed in a discrete choice context by [Rust \(1987\)](#) in what he termed the Nested Fixed Point Maximum Likelihood Estimator (NFXP).

7.1 Estimation Results

Table 8 reports ML estimates of the model parameters. The standard error of the unobserved state variables, $\sigma_{\bar{\zeta}}$, is estimated to be approximately 0.4, suggesting that the model, rather than unaccounted elements, determines optimal retirement behavior.¹⁴

The estimated relative risk aversion, ρ , is just below one, rather low, compared to most existing literature. [van der Klaauw and Wolpin \(2008\)](#) estimates this parameter to be around 1.6 for low-income households older than 45 in the Health and Retirement Study (HRS). My estimate is, however, larger than those reported in [Bingley and Lanot \(2007\)](#), also using elderly Danish households. The estimated ρ in [Bingley and Lanot \(2007\)](#) is not statistically different from zero implying risk-neutral consumers.

Danish elderly couples enjoy almost the same level of welfare as singles from a given level of consumption. The estimated loading on the number of adults, ν , in the equivalence scale is significantly different from zero but low, around 0.05. The OECD equivalence scale is ten times higher, $\nu_{OECD} = 0.7$. The equivalence scales used in [Fernández-Villaverde and Krueger \(2007\)](#) are close to the estimate of 0.33 in [Hong and Ríos-Rull \(2012\)](#) based on the International Survey of Consumer Financial Decisions. My estimate of ν is surprisingly low and suggests that there is almost *no* cost to increasing the household with an additional adult.

Individual's value from own leisure time is allowed to be a function of their educational attainment, children and grandchildren, $\alpha^j(\mathbf{s}_t^j) = \alpha_0^j + \alpha_1^j e^j + \alpha_2^j ch + \alpha_3^j g_t$. I allow males and females to value leisure differently but restrict heterogeneity across other dimensions to be identical across males and females, $\alpha_s^m = \alpha_s^f$, $s = 1, 2, 3$. During

¹⁴All financial variables, such as income, wealth and pension benefits are in 100,000 Danish kroner.

Table 8 – Estimated Preferences.

Parameter		Estimate (SE)
ρ	Risk aversion	0.960 (0.006)***
ν	Equivalence scale	0.048 (4e-4)***
σ_ε	Unobserved states	0.435 (0.004)***
<i>Own leisure</i>		
α_0^m	Constant, males	0.160 (0.002)***
α_0^f	Constant, females	0.119 (0.002)***
α_1	High skilled	0.053 (0.002)***
α_2	Children	-0.036 (0.003)***
α_3	Grandchildren	0.061 (0.004)***
<i>Joint leisure</i>		
ϕ_0^m	Constant, males	1.187 (0.079)***
ϕ_0^f	Constant, females	1.671 (0.097)***
ϕ_1	High skilled	-0.621 (0.023)***
ϕ_2	Children	0.503 (0.075)***
ϕ_3	Grandchildren	-0.724 (0.073)***
ϕ_4	Poor health	0.214 (0.100)*
ϕ_5	Poor health, spouse	-0.342 (0.098)**
$\mathcal{L}(\theta)$		13.2686
$\max_i \{ \partial \mathcal{L}(\theta) / \partial \theta_i \}$		1.1e-5

Notes: Table 8 reports Maximum Likelihood (ML) estimated preference parameters of the model outlined in section 4. Standard errors (in brackets) are based on the inverse of the hessian. *, **, *** refers to $p < .05$, $p < .001$, $p < .0001$, respectively.

initial investigations, different parameters were estimated for males and females without significant differences. To reduce the number of parameters to be estimated in the NFXP routine, I restrict the parameters to be identical across household members.

I estimate that males value leisure more than females, $\hat{\alpha}_0^m > \hat{\alpha}_0^f$. This result stems from the fact that for the model to be able to explain the observed retirement pattern of males, who on average have higher incomes than females, the value of leisure must be greater for males. Contrary to the correlations reported in Table 2, the results suggest that high skilled individuals value leisure time more than low skilled. Households with children value leisure less and the presence of grandchildren almost offset the effect of children, in line with the correlations in Table 2.

Leisure complementarities are allowed to vary with educational attainment, chil-

dren, grandchildren and own and spousal health status,

$$\phi^j(\mathbf{s}_t) = \phi_0^j + \phi_1^j e^j + \phi_2^j ch + \phi_3^j g_t + \phi_4^j \mathbb{1}_{\{h_t^j=0\}} + \phi_5^j \mathbb{1}_{\{h_t^{-j}=0\}}.$$

Again, to reduce the number of parameters, only the constant is allowed to differ across gender and all other parameters are restricted to be identical across husband and wives, $\phi_s^m = \phi_s^f$, $s = 1, 2, 3, 4, 5$. Recall, the parametrization is such that leisure complementarities increases the value of own leisure time and the total value of leisure is, thus, $\alpha(\mathbf{s}_t)(1 + \phi(\mathbf{s}_t))$. The estimated parameters in $\phi(\mathbf{s}_t)$ should, thus, be interpreted as the *percentage increase* in the value of leisure when the spouse is also retired.

Couples enjoy leisure time together. The estimates of the baseline value of joint leisure for husbands and wives, $\hat{\phi}_0^m = 1.187$ and $\hat{\phi}_0^f = 1.671$, suggest that their value from leisure is increased with 100 to 150 percent, respectively, when a spouse is also retired. This suggests large complementarities in leisure of retired Danish couples. [Gustman and Steinmeier \(2000, 2004\)](#) estimates a parametrization very similar to mine.¹⁵ Using the NLS, [Gustman and Steinmeier \(2000\)](#) estimates the percentage increase in the value from leisure when a spouse is also retired in the range of 9-10 percent and 80-150 percent for wives and husbands, respectively. Their estimates are, however, not very precisely estimated. [Gustman and Steinmeier \(2004\)](#) estimates similar effects using the HRS. [Casanova \(2010\)](#) estimates, using an alternative parametrization, that joint leisure is valued as a nine percent increase in individual leisure *endowment*. Estimating interdependent durations for husbands and wives in the HRS, [Honoré and de Paula \(2013\)](#) find that the indirect utility associated with being retired increases by 10 percent if a spouse is already retired. They assume that retirement choices are outcomes of a Nash bargaining problem where each member commits at the age of 60 when to retire in the future. The model is static in the sense that individuals do not update the strategy dynamically over time and commits to the choice made at age 60.

Wives value joint leisure more than their husbands ($\hat{\phi}_0^m < \hat{\phi}_0^f$). Most existing studies suggest that husbands respond more to their wives incentives ([Coile, 2004](#)) and retirement status ([Gustman and Steinmeier, 2000, 2004](#)). As argued throughout, none of the existing studies allow for intra-household correlations in income and health and, thus, cannot disentangle common shocks from leisure complementarities.

The parametrization is such that for wives to enjoy the same *level* of utility from joint leisure as their husbands, the *percentage increase* (ϕ) must be greater. This is because I estimate $\hat{\alpha}_0^m > \hat{\alpha}_0^f$ and the (baseline) total value of joint leisure is $\hat{\alpha}_0^j(1 + \hat{\phi}_0^j)$. Comparing the total value of leisure when both spouses are retired, $.160 \cdot (1 + 1.187) = .35$ for hus-

¹⁵They parametrize utility as $C_t^\alpha / \alpha + e^{X_t^w \beta + \gamma_w L_t^h + \epsilon_w} L_t^w$ for wives and similarly for husbands. The term $e^{\gamma_w L_t^h}$ can be interpreted as the percentage increase in the wife's utility from leisure when her husband is also retired.

bands and $.119 \cdot (1 + 1.671) = .32$ for wives, it is clear that the difference is small. The *total value* from leisure when both are retired is in fact slightly larger for the husband. If husbands and wives are affected differently by children, grandchildren and health, the fact that I restrict both members to be affected identically from, e.g., poor health of a spouse might also distort the estimated baseline value of joint leisure, $\hat{\phi}_0^j$.

High skilled individuals value joint leisure less, $\hat{\phi}_1 < 0$, children increase the complementarity between the husband's and wife's leisure, $\hat{\phi}_2 > 0$, and grandchildren slightly decrease the complementarities, $\hat{\phi}_3 < 0$. All these parameters have opposite signs relative to how they affect own leisure time. The *total* marginal effect of, e.g., children on the value of leisure is a product of the two estimates.

Poor health does not seem to affect the complementarities in leisure significantly, $\hat{\phi}_4 \approx 0$. This is in line with the very small effects found on the probability of simultaneous retirement in Table 3. Health status, thus, seem not to affect complementarities once intra-household correlations in health status across husband and wife is controlled for. Poor health of a spouse, however, is found to have a slightly *negative* effect on the value of joint leisure, $\hat{\phi}_4 < 0$. Existing studies for the US suggest the opposite that poor spousal health increases the likelihood of retirement (Blau, 1998). Also using Danish data, An, Christensen and Gupta (2004) estimate that poor health of a spouse leads to deferred retirement. This in support of my results and is likely due to the free health care system in Denmark.

7.2 Model Fit

There is no mechanical age-effects (such as age-dummies) in the preferences for leisure. The predicted retirement patterns are exclusively a product of the implemented model and features of the Danish pension scheme together with the estimated preference parameters in Table 8. The retirement and wealth accumulation predicted from the model should, thus, be evaluated in that light. Allowing for changing preferences over different ages, as in, e.g., Bingley and Lanot (2007) might be tempting, and even reasonable, but would (in my view) lead to an *artificial* inflation of the model's capability to fit observed retirement frequencies for different age groups.

Figure 5 illustrates actual retirement frequencies and average retirement probabilities predicted by the model for *married couples*. The overall predicted retirement probabilities are close to the observed retirement frequencies. The fact that almost none retire prior to the ERP age of 60 is also predicted by the model and the retirement frequencies from age 62 through 68 are almost perfectly predicted. There is, however, a tendency to a slight overestimation of the retirement probability at age 61 for husbands and wives and an underestimation of the retirement probability at age 60 for wives.

Figure 6 illustrates actual retirement frequencies and average retirement probabili-

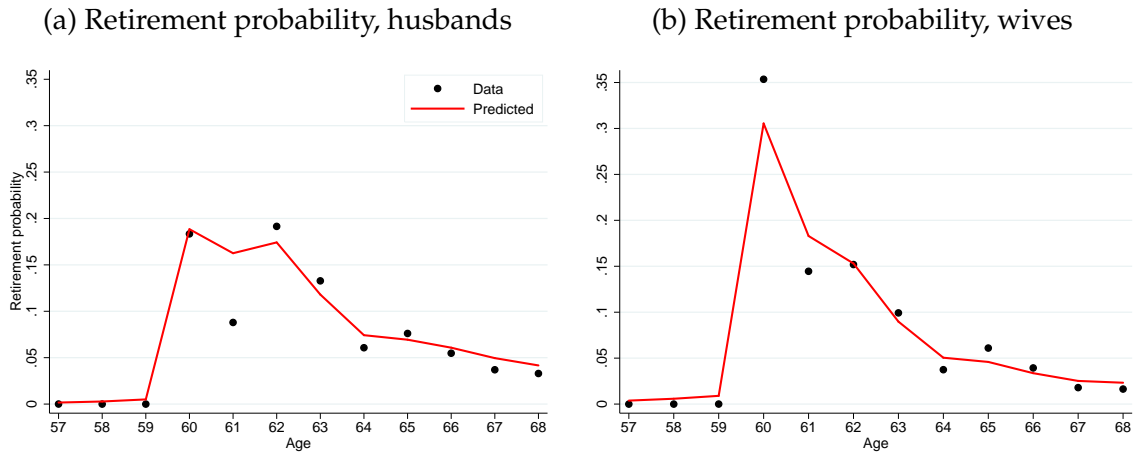


Figure 5 – Actual and Predicted Propensities to Retire, Couples.

Notes: Figure 5 illustrates actual retirement frequencies and average retirement *probabilities* predicted by the model for married males and females. Retirement probabilities are calculated for all individuals working in the beginning of the period and set to zero once retired.

ties predicted by the model for *singles*. The fit for singles is also good albeit the model has difficulties in explaining the large peak in retirement of single males at age 60. The peak at age 65 (OAP age) is larger for singles than for couples and the model slightly underestimates the retirement at that age for singles. As for couples, the model has difficulties in explaining the significant drop in retirement at age 61. At this age, individuals eligible for ERP at age 60 can qualify for the two-year rule by remaining employed an additional year until age 62. The effect it has on retirement behavior seem to be larger than what the financial incentives should suggest.

The inability of the model to explain the large peak at age 60 and the dramatic fall in retirement at age 61 could be suggestive of individuals being affected by social norms. For example, a social consensus could have been reached that one either retires at age 60 or wait until age 62. This is indeed likely to be the case, but giving such an interpretation from the estimated model requires that the model is sufficiently rich enough to describe all important financial incentives and individual heterogeneity. Although the estimated model of household consumption and individual retirement allows for a great amount of heterogeneity across income, wealth, educational attainment, household composition, and health, the inability of the model to fit the observed behavior perfectly at age 60 and 61 might also be due to additional excluded heterogeneity.

The model predicts the observed wealth accumulation very well. Figure 7 illustrates average age profiles of actual household net-worth and predicted net-worth based on the estimated model. The predicted age profiles are almost identical to the observed age profiles for couples and singles. There is, however, a tendency to a slight underestimation of the net-wealth of singles.

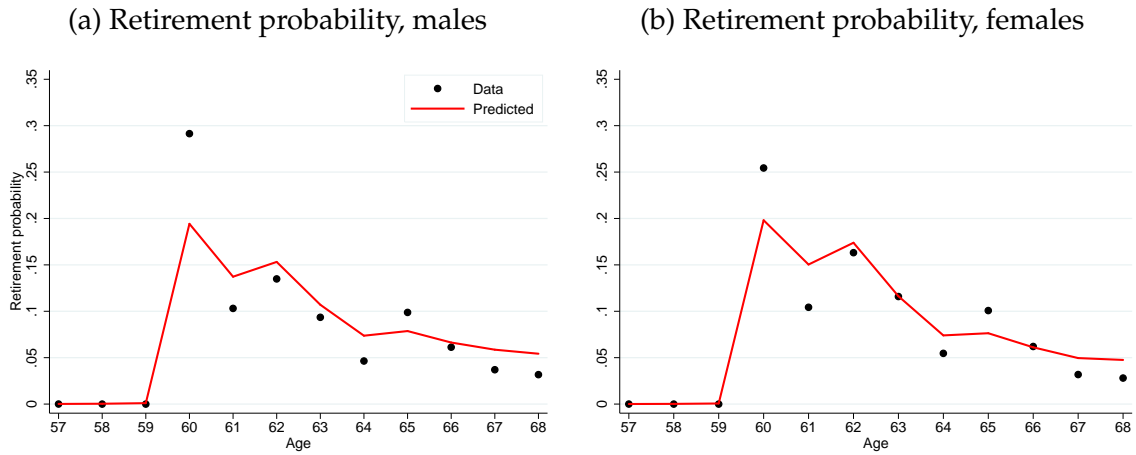


Figure 6 – Actual and Predicted Propensities to Retire, Singles.

Notes: Figure 6 illustrates actual retirement frequencies and average retirement *probabilities* predicted by the model for married males and females. Retirement probabilities are calculated for all individuals working in the beginning of the period and set to zero once retired.

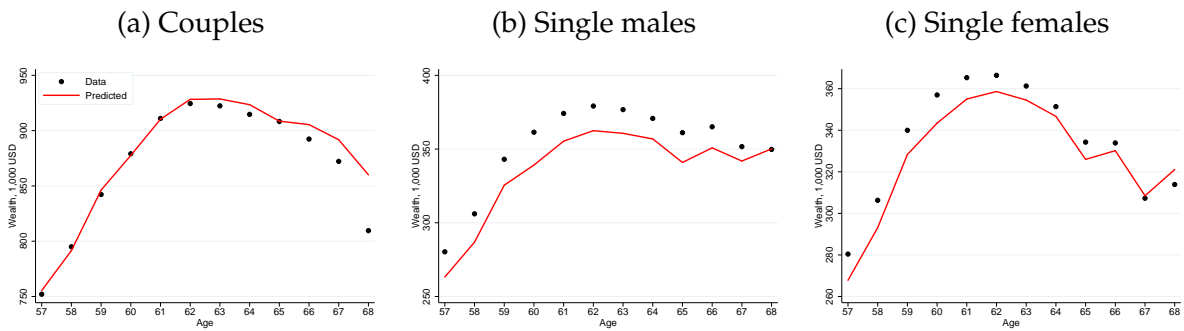


Figure 7 – Wealth Age Profiles, Actual and Predicted.

8 Counterfactual Policy Simulations

Hypothetical policy changes can be evaluated by simulating synthetic data from the estimated model. To illustrate the importance of leisure complementarities, I evaluate how increasing the retirement age affects retirement behavior of singles and couples and how government revenue is affected if leisure complementarities is ignored.

Data is simulated using the first available observation on wealth, health, income, children and educational attainment for each household member. In the following periods, health, grandchildren and income evolve according to the processes described in section 6. Optimal consumption, savings and retirement are determined by the estimated model. In turn, government income tax revenues and pension transfers can be inferred and used to evaluate changes in government surplus from policy changes.

Figure 8 illustrates the change in retirement behavior of couples and singles when the ERP and OAP age is increased with one year (from 60 to 61 and 65 to 66, respectively). The two-year rule is left unaffected. I report the percentage change compared to the fraction retiring at a given age prior to the policy change. The 100 percentage drop at age 60, thus, illustrates that *all* individuals retiring at age 60 under the old regime has delayed retirement after the reform. The small reductions in retirement frequencies before age 60 are unimportant since only a very small fraction of individuals retire prior to the early retirement age (Figures 6 and 5). This is also why the percentage change cannot be calculated for some retirement age groups; there was simply no individuals retiring at that age.

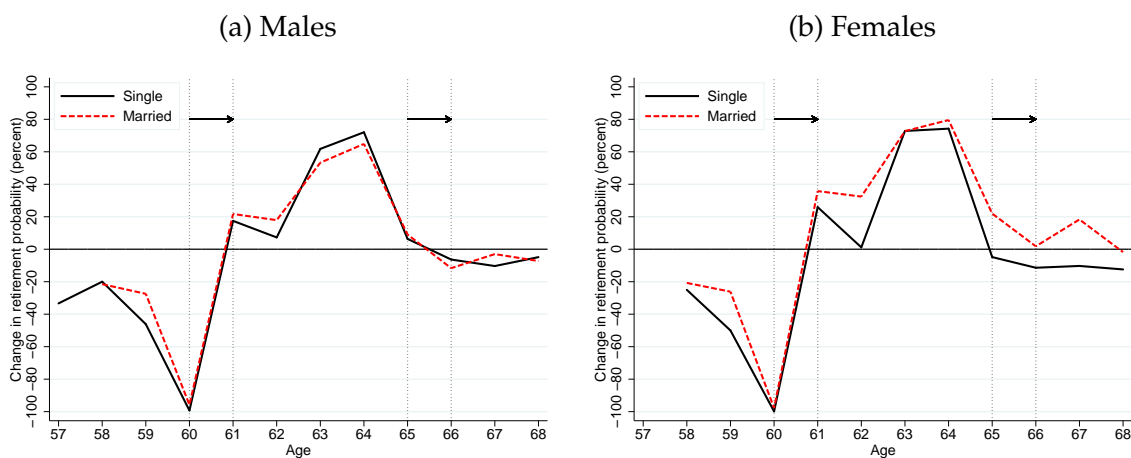


Figure 8 – Counter Factual Predictions: Increasing ERP and OAP age with one year.

Notes: Figure 8 illustrates the change in retirement behavior of couples and singles when the ERP age and old age retirement OAP age is increased with one year (from 60 to 61 and 65 to 66, respectively) while the two-year rule is left unaffected. The figures report the percentage change compared to the fraction retiring at a given age.

The changes in the retirement patterns are similar across gender and to some extent across marital status. Interestingly, however, married females particularly increase the retirement age to 61 (new ERP age) and 62. I interpret this difference as being caused by the fact that wives are on average one year younger than their husbands and, thus, more likely to retire as early as possible. There is hardly any significant response in the OAP age changing from 65 to 66 because most households have retired before the OAP age. The two-year rule is still available and early retirement pension is often higher than the old age pension. This implies that changing the OAP age alone will have little or no effect on the retirement pattern.

Figure 9 illustrates how the government surplus (labor income taxes net of pension transfers) is affected by increasing the retirement age. Both the ERP and OAP age is increased while the two-year rule is untouched, as in Figure 8. To illustrate how leisure

complementarities affect optimal policy design, the same exercise is performed using a version of the model where leisure complementarities is absent, i.e., $\phi_k^j = 0 \forall j, k$. Remaining parameters are fixed at their estimated values in table 8.

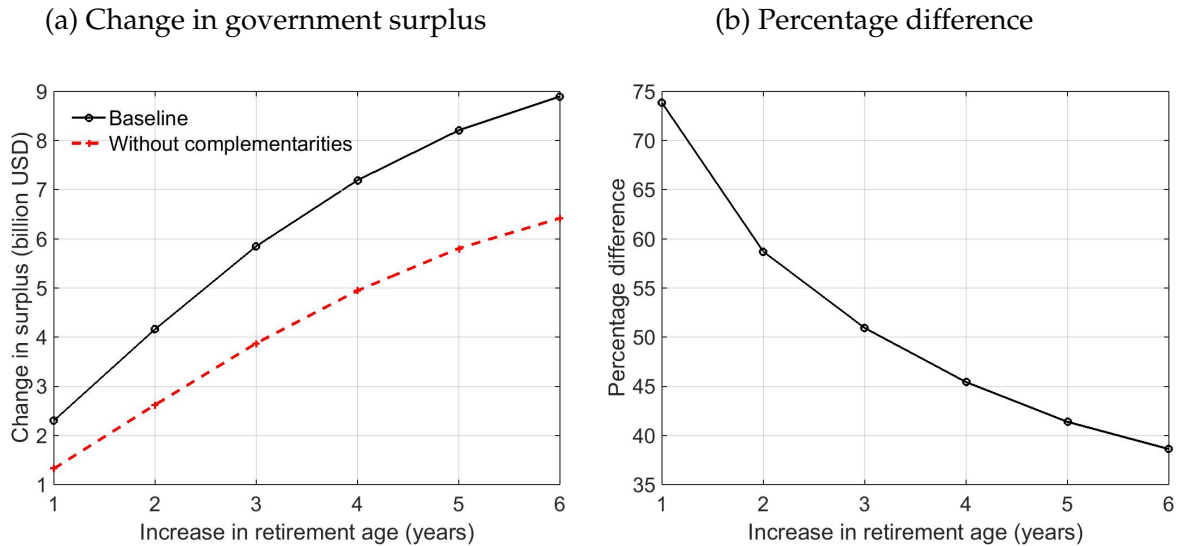


Figure 9 – Change in Government Surplus from Increased Retirement Age.

Notes: Figure 9a illustrates how the government surplus (labor income taxes net of transfers) is affected by increasing the age of retirement. Both the ERP and OAP age is increased and the two-year rule is still active. Figure 9b illustrates the percentage increase in government surplus from the baseline model *with complementarities* in leisure relative to the increase in government surplus from the model *without complementarities*. Hence, panel b illustrates the factor to be multiplied the government surplus if ignoring complementarities to get the surplus from the baseline model *with* leisure complementarities.

Leisure complementarities are important when evaluating alternative retirement policies. Figure 9a shows that increasing the retirement age with one year leads to an *underestimation* of the government surplus of around a billion US dollars (around 5 billion Danish kroner). The underestimation of government surplus when ignoring leisure complementarities increases in the number of years the retirement age is increased. Increasing the early retirement age to 64 (and old age retirement age to 69) leads to an underestimation of 2.25 billion US dollars.

Figure 9b shows the percentage increase in government surplus from the baseline model *with* complementarities in leisure relative to the increase in government surplus from the model *without* complementarities. Adding leisure complementarities increase the resulting government surplus with around 40-75 percent, depending on with how many years retirement age is increased. This is a substantial difference and, taken at face value, suggests that leisure complementarities are important to take into account when evaluating alternative policies.

The estimated government surplus is in a range similar to that predicted by the Danish Ministry of Finance when implementing a retirement reform in 2011. The re-

form passed in 2011 had many components but the main of which was a gradual four-year increase in the ERP age from 60 to 64 from 2014 through 2023 and a gradual two-year increase in the OAP age from 65 to 67 from year 2019 through 2022 ([The Danish Ministry of Finance, 2011](#)). The reform also included a small increase in the maximum ERP with additional means testing of especially private pension funds. Also, more transparent disability pension rules was included in the reform to ensure the benefits of poor-health individuals who was exiting the labor market through the ERP system rather than disability pension prior to the reform.

[The Danish Ministry of Finance \(2011\)](#) estimates that by 2020 the reform will have improved the government surplus by approximately 18 billion Danish kroner, or roughly 3.25 billion US dollars. Their estimate completely ignores leisure complementarities. The “without complementarities” calculation of the hypothetical increase in the ERP and OAP age in figure 9a shows that my estimates are of the same magnitude. For example, increasing the ERP and OAP age with three years instantaneously (rather than gradually) produces an estimated surplus of around 3.75 US dollars in my simulations.

The simulation results suggest that the Danish retirement reform passed in 2011 might improve the government surplus by 25 billion Danish kroner, 40 percent more than the 18 billion predicted by [The Danish Ministry of Finance \(2011\)](#). The simulation results are, of course, to be interpreted under the light of the assumptions imposed by the model and data selection criteria imposed throughout the analysis.

Policy reforms are typically very complex and work through numerous channels, of which many might not be included in the present model framework. Although these results are subject to a substantial amount of caveats, the results clearly show that ignoring the retirement pattern of couples and leisure complementarities in retirement likely will produce inaccurate policy advice.

9 Concluding Discussion

I have formulated, numerically solved and estimated a dynamic economic model of household consumption and individual retirement of Danish couples. The model is rich in household heterogeneity and includes the financial incentives for saving and retirement faced by Danish households through the Danish tax and retirement pension schemes. I use high quality Danish administrative registers to estimate preferences for own and joint leisure by Maximum Likelihood.

To disentangle leisure complementarities from intra-household correlated shocks, I allow health and labor market income to be correlated across husband and wife. This novel feature of the model together with the implemented Danish institutional settings allows me to separate leisure complementarities from financial incentives and household-level shocks. This has, to the best of my knowledge, not been implemented

in as rich dynamic models of dual earner optimal retirement and consumption, as the one estimated herein.

The results confirm the existing literature in the finding of an important role for leisure complementarities in retirement. I estimate that leisure is approximately twice as valuable if a spouse is also retired. I find that the value of leisure increases more for wives when their husbands are also retired than the reverse. This result contradicts most existing literature in which the husband is typically found to respond more to the labor market status of their wife, compared to the effect they have on their wife's value of leisure.

I estimate substantial heterogeneity in leisure complementarities. Specifically, wives are found to value joint leisure more than their husbands and low skilled individuals value joint leisure more than high skilled. Children are found to decrease the value of own leisure but increase the value of joint leisure. While grandchildren increase the value of own leisure, grandchildren decrease the value of joint leisure. Further, I estimate that the value of joint leisure is unaffected by own health status while poor spousal health decreases it.

Counterfactual policy simulations suggest that ignoring leisure complementarities underestimates the government surplus, when analyzing the effect of an increase in the retirement age. In turn, the results indicate that leisure complementarities in retirement is important for couples' behavior and, thus, for optimal policy design. Furthermore, the results stress the fact that performing event-studies of retirement reforms on retirement responses of *individuals* will likely provide biased results. The response of an individual to changes in the retirement scheme will be contaminated by potential spill-over effects from the spouse who may also be affected by the reform under study.

References

- AN, M. Y., B. J. CHRISTENSEN AND N. D. GUPTA (2004): "Multivariate mixed proportional hazard modelling of the joint retirement of married couples," *Journal of Applied Econometrics*, 19(6), 687–704.
- BINGLEY, P. AND G. LANOT (2007): "Public pension programmes and the retirement of married couples in Denmark," *Journal of Public Economics*, 91(10), 1878–1901.
- BLAU, D. M. (1998): "Labor Force Dynamics of Older Married Couples," *Journal of Labor Economics*, 16(3), 595–629.
- BLAU, D. M. AND D. B. GILLESKIE (2006): "Health Insurance and Retirement of Married Couples," *Journal of Applied Econometrics*, 21(7), 935–953.
- BROWNING, M. AND S. LETH-PETERSEN (2003): "Imputing Consumption from Income and Wealth Information," *The Economic Journal*, 113(488), 282–301.
- BURTLESS, G. (1986): "Social Security, Unanticipated Benefit Increases, and the Timing of Retirement," *The Review of Economic Studies*, 53(5), 781–805.
- CARROLL, C. D. (2006): "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economics Letters*, 91(3), 312–320.
- (2011): "Solution Methods for Microeconomic Dynamic Stochastic Optimization Problems," Lecture notes, Department of Economics, Johns Hopkins University, Baltimore.
- CARROLL, C. D. AND M. S. KIMBALL (1996): "On the Concavity of the Consumption Function," *Econometrica*, 64(4), 981–992.
- CASANOVA, M. (2010): "Happy Together: A Structural Model of Couples' Joint Retirement Choices," Working paper, Department of Economics, UCLA.
- CLAUSEN, A. AND C. STRUB (2013): "A General and Intuitive Envelope Theorem," Working Paper 248, Edinburgh School of Economics.
- COILE, C. (2004): "Retirement Incentives and Couples' Retirement Decisions," *Topics in Economic Analysis & Policy*, 4(1), article 17.
- EJRNÆS, M. AND S. HOCHGUERTEL (2013): "Is Business Failure Due to Lack of Effort? Empirical Evidence from a Large Administrative Sample," *The Economic Journal*, 123(571), 791–830.

- FAFCHAMPS, M. AND J. PENDER (1997): "Precautionary Saving, Credit Constraints, and Irreversible Investment: Theory and Evidence from Semiarid India," *Journal of Business & Economic Statistics*, 15(2), 180–194.
- FELLA, G. (2014): "A generalized endogenous grid method for non-smooth and non-concave problems," *Review of Economic Dynamics*, 17(2), 329–344.
- FERNÁNDEZ-VILLAYERDE, 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.
- FORSIKRING & PENSION (2008): *Sociale Ydelser 2008: Hvem, hvad & hvornår*. Forlaget Forsikring, (In Danish).
- FRENCH, E. (2005): "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour," *Review of Economic Studies*, 72(2), 395–427.
- GOURINCHAS, P.-O. AND J. A. PARKER (2002): "Consumption Over the Life Cycle," *Econometrica*, 70(1), 47–89.
- GUSTMAN, A. L. AND T. L. STEINMEIER (2000): "Retirement in Dual-Career Families: A Structural Model," *Journal of Labor Economics*, 18(3), 503–545.
- (2004): "Social Security, Pensions and Retirement Behaviour within the Family," *Journal of Applied Econometrics*, 19(6), 723–737.
- HO, C. (forthcoming): "Grandchild care, intergenerational transfers, and grandparents' labor supply," *Review of Economics of the Household*.
- HONG, J. H. AND J.-V. RÍOS-RULL (2012): "Life Insurance and Household Consumption," *American Economic Review*, 102(7), 3701–30.
- HONORÉ, B. E. AND A. DE PAULA (2013): "Interdependent Durations in Joint Retirement," Working Paper CWP05/13, The Institute for Fiscal Studies Department of Economics, UCL (CEMMAP).
- HURD, M. D. (1990): "The Joint Retirement Decision of Husbands and Wives," in *Issues in the Economics of Aging*, ed. by D. A. Wise, chap. 8, pp. 231–258. University of Chicago Press.
- ISKHAKOV, F. (2010): "Structural dynamic model of retirement with latent health indicator," *The Econometrics Journal*, 13(3), 126–161.
- ISKHAKOV, F., T. H. JØRGENSEN, J. RUST AND B. SCHJERNING (2014): "Estimating Discrete-Continuous Choice Models: Endogenous Grid Method with Taste Shocks," Unpublished working paper.

- JIA, Z. (2005): "Labor Supply of Retiring Couples and Heterogeneity in Household Decision-Making Structure," *Review of Economics of the Household*, 3, 215–233.
- JØRGENSEN, T. H. (2013): "Structural estimation of continuous choice models: Evaluating the EGM and MPEC," *Economics Letters*, 119(3), 287–290.
- (2014): "Life-Cycle Consumption and Children," Working Paper 2014-2, Centre for Applied Microeconometrics (CAM), University of Copenhagen.
- KONRAD, K. A. AND K. E. LOMMERUD (2003): "The Bargaining Family Revisited," *Canadian Journal of Economics*, 33(2), 471–487.
- LEWIS, S. AND V. OPPENHEIMER (2000): "Educational assortative mating across marriage markets: Nonhispanic whites in the United States," *Demography*, 37(1), 29–40.
- MICHAUD, P.-C. AND F. VERMEULEN (2011): "A collective labor supply model with complementarities in leisure: Identification and estimation by means of panel data," *Labour Economics*, 18(2), 159–167.
- NIELSEN, H. S. AND M. SVARER (2009): "Educational Homogamy: How Much is Opportunities?," *The Journal of Human Resources*, 44(4), 1066–1086.
- POSADAS, J. AND M. VIDAL-FERNANDEZ (2013): "Grandparents' Childcare and Female Labor Force Participation," *IZA Journal of Labor Policy*, 2(14), 1–20.
- RUPERT, P. AND G. ZANELLA (2014): "Grandchildren and Their Grandparents' Labor Supply," Working Paper 937, University of Bologna.
- RUST, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55(5), 999–1033.
- RUST, J. AND C. PHELAN (1997): "How Social Security and Medicare Affect Retirement Behavior In a World of Incomplete Markets," *Econometrica*, 65(4), 781–831.
- SCHOLZ, J. K. AND A. SESHADRI (2013): "Health Insurance and Retirement Decisions," Working Paper 2013-292, Michigan Retirement Research Center.
- TAUCHEN, G. (1986): "Finite state markov-chain approximations to univariate and vector autoregressions," *Economics Letters*, 20(2), 177 – 181.
- TERRY, S. J. AND E. S. KNOTEK (2011): "Markov-chain approximations of vector autoregressions: Application of general multivariate-normal integration techniques," *Economics Letters*, 110(1), 4 – 6.
- THE DANISH MINISTRY OF FINANCE (2011): *Aftaler om reformpakken 2020*. Danish Ministry of Finance, Foreign Policy Center (in Danish).

TOBIN, J. (1958): "Estimation of Relationships for Limited Dependent Variables," *Econometrica*, 26(1), pp. 24–36.

VAN DER KLAAUW, W. AND K. I. WOLPIN (2008): "Social security and the retirement and savings behavior of low-income households," *Journal of Econometrics*, 145(1-2), 21–42.

WARREN, D. (2013): "Retirement Decisions of Couples: The Impact of Spousal Characteristics and Preferences on the Timing of Retirement," Working Paper 41/13, Melbourne Institute of Applied Economic and Social Research.

Supplemental Material (not for publication)

Leisure Complementarities in Retirement

Thomas H. Jørgensen
University of Copenhagen
October 29, 2014

A Institutional Settings in Denmark

Combining the assumption that all pension wealth is held in private IRAs with the assumption of zero hours worked when retired, the retirement pension scheme can be formulated as

$$P(\mathbf{s}_t) = \begin{cases} \max\{0, \overline{\text{ERP}} - .6 \cdot .05 \max\{0, (\varphi_t^j - \underline{\text{ERP}})\}\} & \text{if } \text{elig}^j = 1, \text{ and } 60 \leq \text{RetAge}^j \leq 61, \\ \overline{\text{ERP}}_2 & \text{if } \text{elig}^j = 1, \text{ and } 62 \leq \text{RetAge}^j < 65, \\ OAP(d_t^{-j}, y_t^{-j}, \text{age}_t^{-j}) & \text{if } d_t^j = 0, \text{ and } \text{age}_t^j \geq 65, \\ 0 & \text{else,} \end{cases}$$

where RetAge^j denotes the retirement age of member j , \mathbf{s}_t contain all potential factors affecting pension benefits, φ_t^j is the private pension fund balance in IRAs, $\overline{\text{ERP}} = 166,400\text{DKK} \approx \$30,250$ is the maximum early retirement pension in 2008 if the two-year-rule is *not* fulfilled, $\overline{\text{ERP}}_2 = 182,780\text{DKK} \approx \$33,250$ is the maximum early retirement pension if the two year rule is fulfilled, and $\underline{\text{ERP}} = 12,600\text{DKK} \approx \$2,300$ is a deduction.

Old age pension, $OAP(d_t^{-j}, y_t^{-j}, \text{age}_t^{-j})$, depends heavily on the potential spousal labor market income, y^{-j} . It is, therefore, convenient to define an index

$$i = \begin{cases} 1 & \text{if single,} \\ 2 & \text{if couple and } d_t^{-j} = 1 \text{ or } \text{age}_t^{-j} < 65, \\ 3 & \text{if couple and } d_t^{-j} = 0 \text{ and } \text{age}_t^{-j} \geq 65, \end{cases}$$

such that $i = 1$ is a single household, $i = 2$ is a couple where the spouse is not receiving OAP, and $i = 3$ denotes a household in which both members receive OAP.¹⁶

The Danish OAP system is such that if the spouse of an retiree on OAP is *not* receiving OAP, the labor market income of the working spouse reduces the amount of OAP the retiree receives. This, in turn, promote simultaneous retirement of couples.

¹⁶In Danish, the terms used by [Forsikring & Pension \(2008\)](#) to denote these groups are “reelt enlige”, “gifte/samlevende”, and “samgifte”, respectively.

Specifically, the OAP can be calculated as the sum of two components,

$$OAP(d_t^{-j}, y_t^{-j}, age_t^{-j}) = OAP_B + OAP_A$$

where

$$\begin{aligned} OAP_B &= \mathbb{1}_{\{y_t^j < \bar{y}_B\}} \max\{0, (B - \tau_B \max\{0, y_t^j - \bar{D}_B\})\}, \\ OAP_A &= \mathbb{1}_{\{y_h < \bar{y}_i\}} \max\{0, (A_i - \max\{0, \tau_i(y_h - \bar{D}_i)\})\}, \\ y_h &= y_t^j + (y_t^{-j} - .5 \min\{\bar{D}_s, y^{-j}\}) \mathbb{1}_{\{i=2\}}. \end{aligned}$$

Table A.1 reports the relevant values for 2008 of the Danish OAP scheme.

Table A.1 – Old Age Pension Parameters.

Symbol	Value in 2008	Description
B	61,152 \approx \$10,700	Base value of old age pension
\bar{y}_B	463,500 \approx \$81,000	Maximum annual income before loss of OAP_B
τ_B	.3	Marginal reduction in deduction regarding income
\bar{D}_B	259,700 \approx \$45,500	Deduction regarding base value of OAP
\bar{D}_s	179,400 \approx \$31,500	Maximum deduction in spousal income
A_i	$\begin{cases} 61,560 \approx \$10,800 \\ 28,752 \approx \$5,000 \\ 28,752 \approx \$5,000 \end{cases}$	Maximum OAP_A , for $i = 1, 2, 3$.
\bar{y}_i	$\begin{cases} 153,100 \approx \$46,000 \\ 210,800 \approx \$37,000 \\ 306,600 \approx \$54,000 \end{cases}$	Maximum income before loss of OAP_A , for $i = 1, 2, 3$.
τ_i	$\begin{cases} .30 \\ .30 \\ .15 \end{cases}$	Marginal reduction in OAP_A , for $i = 1, 2, 3$.
\bar{D}_i	$\begin{cases} 57,300 \approx \$10,000 \\ 115,000 \approx \$20,000 \\ 115,000 \approx \$20,000 \end{cases}$	Maximum deduction regarding OAP_A , for $i = 1, 2, 3$.

A.1 The Danish Tax System

The after-tax income of an individual j with income y^j and potential spousal income of y^{-j} can be calculated based on the following equations:

$$\begin{aligned}
 \tau_{\max} &= \tau_l + \tau_m + \tau_u + \tau_c + \tau_h - \bar{\tau}, \\
 \text{personal income} &= (1 - \tau_{LMC}) \cdot y^j, \\
 \text{taxable income} &= \text{personal income} - \min\{WD \cdot y^j, \overline{WD}\}, \\
 \underline{y}_l &= \underline{y} + \max\{0, \underline{y} - y^{-j}\}, \\
 T_c &= \max\{0, \tau_c \cdot (\text{taxable income} - \underline{y}_l)\}, \\
 T_h &= \max\{0, \tau_h \cdot (\text{taxable income} - \underline{y}_l)\}, \\
 T_l &= \max\{0, \tau_l \cdot (\text{personal income} - \underline{y}_l)\}, \\
 T_m &= \max\{0, \tau_m \cdot (\text{personal income} - \underline{y}_m)\}, \\
 T_u &= \max\{0, \min\{\tau_u, \tau_{\max}\} \cdot (\text{personal income} - \underline{y}_u)\}, \\
 \text{after-tax income} &= (1 - \tau_{LMC}) \cdot y^j - T_c - T_h - T_l - T_m - T_u,
 \end{aligned}$$

where the values from 2008 along with descriptions are given in Table A.2.

Table A.2 – Tax System Parameters in 2008.

Symbol	Value in 2008	Description
$\bar{\tau}$.59	Maximum tax rate, »Skatteloft«
τ_{LMC}	.08	Labor Market Contribution, »Arbejdsmarkedsbidrag«
WD	.04	Working Deduction
\overline{WD}	12,300 \approx \$2,200	Maximum deduction possible
τ_c	.2554	Average county-specific tax rate (including .073 in church tax)
\underline{y}	41,000 \approx \$7,500	Amount deductible from all income
\underline{y}_m	279,800 \approx \$50,800	Amount deductible from middle tax bracket
\underline{y}_u	335,800 \approx \$61,000	Amount deductible from top tax bracket
τ_h	.08	Health contribution tax (in Danish »Sundhedsbidrag«)
τ_l	0.0548	Tax rate in lowest tax bracket
τ_m	0.06	Tax rate in middle tax bracket
τ_u	0.15	Tax rate in upper tax bracket

B Solving The Model

The stochastic dynamic programming model is solved by backwards induction using the EGM proposed by [Carroll \(2006\)](#) generalized to both continuous and discrete choices with unobserved states in [Iskhakov, Jørgensen, Rust and Schjerning \(2014\)](#). In a given period, optimal consumption and labor market status is found by a combination of Euler equation and value function iteration. Retirement is an absorbing state and re-entry into the labor market is, thus, not allowed once retired. This is a standard simplifying assumption in the literature and most Danish retirees stay out of the labor market, once they are retired.

It is convenient to recognize that a central determinant of household consumption and retirement choices is available resources within the household. I denote that as

$$M(\mathbf{s}_t) = Ra_{t-1} + \sum_{j \in h, w} \{T(y_t^j, y_t^{-j}) + P^j(\mathbf{s}_t)\},$$

such that the law of motion for available resources follows

$$M_{t+1} = R(M_t - C_t) + \sum_{j \in h, w} \{T(y_{t+1}^j, y_{t+1}^{-j}) + P^j(\mathbf{s}_{t+1})\},$$

where $M_t \equiv M(\mathbf{s}_t)$ is shorthand for the dependence of resources on state variables.

Before turning to how the model for couples is solved, it is instructive to discuss how the model for singles is solved since that solution is subsequently used when solving the model for couples. I will also introduce notation and numerical methods when discussing how the model for singles is solved and only briefly discuss how the model for couples is solved since solving the model for couples is basically the same as for singles just with more bookkeeping.

B.1 Singles

The bellman equation for a single individual j is given by

$$V_t^j(\mathbf{s}_t^j, \varepsilon_t^j) = \max_{C_t^j, d_t^j} \{U^j(C_t^j, d_t^j, \mathbf{s}_t^j) + \sigma_\varepsilon \varepsilon(d_t^j) + \beta \mathbb{E}_t[\pi_{t+1}^j V_{t+1}^j(\mathbf{s}_{t+1}^j, \varepsilon_{t+1}^j) + (1 - \pi_{t+1}^j)B(a_t)]\},$$

where the optimization problem can be broken into two parts related to the consumption and retirement decision,

$$\begin{aligned} V_t^j(\mathbf{s}_t^j, \varepsilon_t^j) &= \max_{d_t^j} \{v_t^j(\mathbf{s}_t^j, d_t^j) + \sigma_\varepsilon \varepsilon(d_t^j)\}, \\ v_t^j(\mathbf{s}_t^j, d_t^j) &= \max_{C_t^j} \{U^j(C_t^j, d_t^j, \mathbf{s}_t^j) + \beta \int_{\mathbf{s}_{t+1}^j} \int_{\varepsilon_{t+1}^j} [\pi_{t+1}^j V_{t+1}^j(\mathbf{s}_{t+1}^j, \varepsilon_{t+1}^j) \\ &\quad + (1 - \pi_{t+1}^j)B(a_t)] p(d\varepsilon_{t+1}^j | \mathbf{s}_{t+1}^j) p(d\mathbf{s}_{t+1}^j | \mathbf{s}_t^j, C_t^j, d_t^j)\}, \end{aligned}$$

where $v_t^j(\mathbf{s}_t^j, d_t^j)$ is choice-specific value functions conditional on optimal consumption/saving for a given labor market choice, d_t^j . In what follows, I will drop the j superscript to reduce clutter.

Rust (1987) proves that when ε is Extreme Value type I distributed, as assumed here, the integral with respect to the unobserved state has a convenient closed form solution. The choice-specific value function is, then,

$$\begin{aligned} v_t(\mathbf{s}_t, d_t) &= \max_{C_t} \{U(C_t, d_t, \mathbf{s}_t) + \beta \int_{\mathbf{s}_{t+1}} \pi_{t+1} EV_{t+1}(\mathbf{s}_{t+1}) p(d\mathbf{s}_{t+1} | \mathbf{s}_t, C_t, d_t) \\ &\quad + \beta(1 - \pi_{t+1})B(a_t, ch)\}, \tag{B.1} \\ EV_{t+1}(\mathbf{s}_{t+1}) &= \sigma_\varepsilon \log \left\{ \sum_{k \in \{0,1\}} \exp(v_{t+1}(\mathbf{s}_{t+1}, k) / \sigma_\varepsilon) \right\}. \end{aligned}$$

To reduce computation time, I work with the consumption Euler equation to solve for the choice-specific value functions. Note that the first order condition of (B.1) is given by

$$U^C(C_t, d_t, \mathbf{s}_t) - \beta \mathbb{E}_t[R\pi_{t+1} V_{t+1}^M(\mathbf{s}_{t+1}, \varepsilon_{t+1}) + (1 - \pi_{t+1})B^a(a_t)] = 0,$$

where U^C , V^M and B^a is partial derivatives w.r.t. consumption, available resources and wealth, respectively. Further, the envelope condition states that

$$V_t^M(\mathbf{s}_t, \varepsilon_t) = \beta \mathbb{E}_t[R\pi_{t+1} V_{t+1}^M(\mathbf{s}_{t+1}, \varepsilon_{t+1}) + (1 - \pi_{t+1})B^a(a_t)],$$

such that we must have $U^C(C_t, d_t, \mathbf{s}_t) = V_t^M(\mathbf{s}_t, \varepsilon_t)$. Rolling this one period forward, the consumption Euler equation becomes

$$U^C(C_t, d_t, \mathbf{s}_t) = \beta \mathbb{E}_t[R\pi_{t+1} U^C(C_{t+1}, d_{t+1}, \mathbf{s}_{t+1}) + (1 - \pi_{t+1})B^a(a_t)].$$

Importantly, the envelope condition is only valid and the Euler equation sufficient when the value function is concave. The discrete choice of retirement does, however,

introduce potential non-concavities in the value functions with related discontinuous consumption functions. Clausen and Strub (2013) provide envelope theorems for non-concave problems and Fella (2014) utilize these to generalize the EGM to handle discrete choices (as the retirement choice herein) when $\sigma_\varepsilon = 0$, i.e., without unobserved state variables. Iskhakov, Jørgensen, Rust and Schjerning (2014) show that the EGM finds *all* solutions to the Euler equation and provides evidence that their DC-EGM for estimating discrete and continuous choice models perform very well in Monte Carlo studies.

I have implemented the DC-EGM algorithm to calculating the so-called “upper envelope” of the value functions. Interestingly, however, the EGM did not find multiple solutions to the Euler equations. The Danish welfare system and generous pension systems ensures most households only experience a minor drop in income when retiring. In turn, the drop in consumption related to retirement is negligible and the Euler equation is, for all practical purposes in this application sufficient for optimal consumption. I now turn to how the numerical solution is implemented.

Retiring Singles

When retired, re-entry into the labor market is not allowed and labor market income is constantly zero. I utilize these assumptions to solve for optimal consumption post retirement by the Endogenous Grid Method (EGM) proposed by Carroll (2006). Specifically, I construct a grid over *end-of-period* wealth, $\hat{a}_t = (0, \dots, \bar{a})$ and invert the Euler equation to find the optimal consumption consistent with this level of end-of-period wealth. With end-of-period wealth and optimal consumption, using the budget constraint, beginning-of-period available resources, M_t , can be endogenously determined as $M_t = \hat{a}_t + C_t^*$. See also Jørgensen (2013) for a discussion of the EGM. The Euler equation is sufficient for optimal consumption for retired households because the households no longer has the discrete retirement choice.

Available resources, M_t , matters for consumption along with other household characteristics such as children and grandchildren. Therefore, with slight abuse of notation, I recognize that an important element is the available resources by including M_t as an argument in the consumption function,

$$C_t^*(M_t, \mathbf{s}_t | d_t = 0) = \left(\beta \mathbb{E}_t [R \pi_{t+1} (C_{t+1}^*(M_{t+1}, \mathbf{s}_{t+1} | d_{t+1} = 0))^{-\rho} + (1 - \pi_{t+1}) B^a(a_t)] \right)^{-\frac{1}{\rho}}$$

where the expectations are over potential arrival of grandchildren and health transition probabilities. There is no income uncertainty post retirement. Optimal consumption is

approximated as

$$\begin{aligned}
 C_t^*(M_t, \mathbf{s}_t | d_t = 0) &\doteq \left(\beta R \sum_{k=0}^1 \sum_{l=0}^1 \Pr(h_{t+1} = k | h_t) \Pr(g_{t+1} = l | \mathbf{s}_t) \pi_{t+1} \right. \\
 &\quad \times \left(\check{C}_{t+1}(\hat{M}_{t+1}, \mathbf{s}_{t+1} | d_{t+1} = 0, h_{t+1} = k, g_{t+1} = l) \right)^{-\rho} \\
 &\quad \left. + \beta(1 - \pi_{t+1}) B^a(\hat{a}_t) \right)^{-\frac{1}{\rho}}, \tag{B.2}
 \end{aligned}$$

where next-period available resources is calculated as $\hat{M}_{t+1} = R\hat{a}_t + P(\mathbf{s}_{t+1})$ (which is independent of the health and grandchildren states) and $\check{C}_{t+1}(\hat{M}_{t+1}, \cdot)$ is a linear interpolation function of optimal consumption found in last iteration. The interpolation is over available resources, M , since this is the “sufficient statistic” for consumption. The choice-specific value function is calculated by substituting optimal consumption into equation (B.1),

$$v_t(\mathbf{s}_t, 0) = U(C_t^*, 0, \mathbf{s}_t) + \beta \int_{\mathbf{s}_{t+1}} [\pi_{t+1} EV_{t+1}(\hat{a}_t) + (1 - \pi_{t+1}) B(\mathbf{s}_{t+1})] p(d\mathbf{s}_{t+1} | \mathbf{s}_t, C_t^*, 0), \tag{B.3}$$

where the $*$ superscript indicates that expectations are taken under optimal consumption this period.

In the last period, $T - 1$, all households are (forced) retired and consumption equals

$$C_T^*(M_T, \mathbf{s}_T) = \min\{M_T, (\beta\gamma)^{-\frac{1}{\rho}}\}$$

due to the bequest function and a leave-no-debt condition.

Working Singles

For working households it is slightly more complicated to solve for optimal consumption and choice-specific value functions because they can transition into retirement next period. To calculate the expected marginal utility of consumption next period, I calculate next-period choice-specific consumption and weight those with the probability of retiring, calculated from the last iteration. Specifically, consumption of a working

single is approximated as

$$\begin{aligned}
 C_t^*(M_t, \mathbf{s}_t | d_t = 1) &\doteq \left(\beta R \sum_{k=0}^1 \sum_{l=0}^1 \Pr(h_{t+1} = k | h_t) \Pr(g_{t+1} = l | \mathbf{s}_t) \pi_{t+1} \right. \\
 &\times \sum_{q=1}^Q \omega_q \sum_{d=0}^1 \Pr(d_{t+1} = d | \mathbf{s}_{t+1}^q) \\
 &\times (\check{C}_{t+1}(\hat{M}_{t+1}, \mathbf{s}_{t+1}^q | d_{t+1} = d, h_{t+1} = k, g_{t+1} = l))^{-\rho} \\
 &\left. + \beta(1 - \pi_{t+1})B^a(\hat{a}_t) \right)^{-\frac{1}{\rho}},
 \end{aligned}$$

where

$$\Pr(d_{t+1} = d | \mathbf{s}_{t+1}) = \frac{\exp(v_{t+1}(\mathbf{s}_{t+1}, d_{t+1} = d) / \sigma_\varepsilon)}{\sum_{k=0}^1 \exp(v_{t+1}(\mathbf{s}_{t+1}, d_{t+1} = k) / \sigma_\varepsilon)}, \quad (\text{B.4})$$

is the dynamic multinomial logit probability of retirement next period, found in the last iteration stemming from the Extreme value type I distributed ε (Rust, 1987). ω_q are Gauss-Hermite quadrature weights used to approximate the next-period labor market income if still working. If an individual chose to retire next period, labor market income is zero and optimal consumption found in (B.2) are used when interpolating next period optimal consumption. However, if the individual chose to work next period, not only available resources but also labor market income matters for consumption. Therefore, bi-variate linear interpolation in income and wealth space of next period consumption and value functions are used if the individual does not retire next period.

Value functions are typically highly non-linear for low levels of resources and, thus, potentially poorly approximated by linear interpolation. Since the value function “inherits” the curvature from the utility function (Carroll and Kimball, 1996) I interpolate $\tilde{v} = -v^{-1}$, which is “more” linear and bounded at zero from below. I then re-transform the resulting interpolated data, such that $\check{v} = -\tilde{v}^{-1}$, where $\tilde{\cdot}$ is a linear interpolation function. To further increase accuracy of the approximated consumption and value functions, the end-of-period wealth grid used when solving the model is unequally spaced, with more points at the lower end of the distribution. These ideas stems from Christopher Carroll’s lecture notes (Carroll, 2011).

With optimal consumption, the choice specific value function can be found as

$$v_t(\mathbf{s}_t, 1) = U(C_t^*, \mathbf{s}_t, 1) + \beta \int_{\mathbf{s}_{t+1}} [\pi_{t+1} EV_{t+1}(\mathbf{s}_{t+1}) + (1 - \pi_{t+1})B(\hat{a}_t)] p(d\mathbf{s}_{t+1} | \mathbf{s}_t, C_t^*, 1),$$

which, together with the value of retirement, found in equation (B.3), is used to calculate the probability of retirement this period by inserting these choice-specific value

functions in equation (B.4).

B.2 Couples

The Bellman equation of working couples prior to forced retirement, $t < T_r$, is

$$\begin{aligned}
 V_t(\mathbf{s}_t, \varepsilon_t) = & \\
 \max_{C_t, d_t^h, d_t^w} & \left\{ U(C_t, d_t^h, d_t^w, \mathbf{s}_t) + \sigma_\varepsilon \varepsilon_t (d_t^h, d_t^w) + \beta \mathbb{E}_t \left[\pi_{t+1}^h \pi_{t+1}^w V_{t+1}(\mathbf{s}_{t+1}, \varepsilon_{t+1}) \right. \right. \\
 & + \pi_{t+1}^h (1 - \pi_{t+1}^w) V_{t+1}^h(\mathbf{s}_{t+1}^h, \varepsilon_{t+1}^h) + (1 - \pi_{t+1}^h) \pi_{t+1}^w V_{t+1}^w(\mathbf{s}_{t+1}^w, \varepsilon_{t+1}^w) \\
 & \left. \left. + (1 - \pi_{t+1}^h)(1 - \pi_{t+1}^w) B(a_t) \right] \right\}, \tag{B.5}
 \end{aligned}$$

where the continuation value if one of the spouses is widowed enters along with the continuation value if both survives as a couple.

To find optimal consumption and choice-specific value functions, a similar approach as for singles is applied. Instead of a binary labor market choice, the household now chooses *both* members end-of-period labor market choice.

B.3 The DC-EGM

This section is based on the DC-EGM approach proposed in [Iskhakov, Jørgensen, Rust and Schjerning \(2014\)](#). Since the EGM finds *all* solutions to the Euler equation, we simply propose to use EGM to find all these solutions and then calculate the value function at these solutions to find which of the levels of consumption is associated with the highest value function. This is what [Clausen and Strub \(2013\)](#) refer to as the *upper envelope*. The approach consists of the following steps

1. Construct a grid of end-of-period wealth, $\vec{a} = (a_1, a_2, \dots, a_Q)$.
2. For all potential states next period, \mathbf{s}_{t+1} , and all potential discrete choices in the following period, d_{t+1} , calculate the next-period beginning-of-period resources, $m_{t+1}(\vec{a}, d_{t+1}, \mathbf{s}_{t+1})$ and interpolate the choice-specific value functions to get $v_{t+1}(d_{t+1}) = v_{t+1}(m_{t+1}(\vec{a}, d_{t+1}, \mathbf{s}_{t+1}), \mathbf{s}_{t+1}, d_{t+1})$ on the \vec{a} -grid.
3. Use $v_{t+1}(d_{t+1})$ to calculate the choice-probabilities, $\Pr(d_{t+1} | \mathbf{s}_{t+1})$, for each potential choice along with the expected continuation value, with respect to the

unobserved shock, ε ,

$$EV_t(\mathbf{s}_{t+1}) = \sigma_\varepsilon \log \sum_{d_{t+1} \in \mathcal{D}(\mathbf{s}_{t+1})} \exp(v_{t+1}(d_{t+1})/\sigma_\varepsilon).$$

4. For a given choice of end-of-period labor market status, d_t , solve for optimal consumption using the inverse of the Euler equation (EGM-step) where the choice-probabilities from step three is used to weight the likelihood of choice d_{t+1} the following period. Use the optimal consumption to generate the endogenous grid of resources, $\vec{m} = c^* + \vec{a}$.

5. Calculate the choice-specific value function related to the current-period choice of d_t ,

$$v_t(d_t) = U(c_t^*; d_t) + \int_{\mathbf{s}_{t+1}} \beta EV_t(\mathbf{s}_{t+1}) f(\mathbf{s}_{t+1} | \mathbf{s}_t)$$

6. Calculate the upper envelope of each of the choice-specific value functions. Since the EGM potentially has located several solutions to the Euler equation, the correct one has to be found. An algorithm to calculating the upper envelope is provided in the pseudocode below in Section B.4. Store the values of $v_t(d_t)$, $c_t^*(d_t)$, and $m_t(d_t)$ and move to next time-iteration, $t - 1$, and re-do from step 1.

B.4 Finding the (Secondary) Upper Envelope

The following pseudocode finds the upper envelope without adding additional grid points. Figure B.1 illustrates the upper envelope from a simple consumption and retirement model from [Iskhakov, Jørgensen, Rust and Schjerning \(2014\)](#). The figure illustrates how EGM finds all solutions of the Euler equation (the backward bending part of the consumption function) and how these solutions can be utilized to find the *optimal* level of consumption.

1. Find *all* points where resources start to fall. Denote these points for A.
2. Using the points in A, find the associated points where resources start to increase again. Denote these points for B.
3. Form a common grid of beginning-of-period resources, \vec{m} , and interpolate all B-A value functions and consumption functions on this grid.
4. With all solution-segments on the same grid, apply the max operator to find the upper envelope of all potential value function segments. Use the found upper envelope to back-out which part of the consumption function that solves the Euler equation is in fact the optimal one.




Figure B.1 – Upper Envelope, DC-EGM, Illustrative Model in *Iskhakov, Jørgensen, Rust and Schjerning (2014)*, $t = T - 4$.

Notes: Figure B.1 illustrates how the upper envelope is calculated using DC-EGM. The left panels illustrate the upper envelope of the (choice-specific) value function and the right panel illustrates the resulting consumption function. The *true* (VFI) solution is plotted together with the raw EGM solution and the DC-EGM proposed by *Iskhakov, Jørgensen, Rust and Schjerning (2014)*. Blue squares mark critical points where the resources from the raw solution start to fall, denoted "A" in the algorithm. Red circles mark where resources again start to increase, denoted "B" in the algorithm. Section B.4 describes how these points are used to find the upper envelope by constructing a common grid of resources and interpolating *all* line segments B-A, marked with solid black lines.