Quantitative Economics 7 (2016), 935–968 1759-7331/20160935

Euler equation estimation: Children and credit constraints

Thomas H. Jørgensen
Department of Economics, University of Copenhagen

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 the marginal utility of consumption if consumers face credit constraints. As a more constructive contribution, I propose a tractable approach to obtaining bounds on the effect of children on the marginal utility of consumption. I estimate these bounds using the Panel Study of Income Dynamics and find that conventional estimators yield point estimates that are above the upper bound. Children might, thus, not increase the marginal utility of consumption as much as previously assumed.

Keywords. Consumption, Euler equation estimation, credit constraints, children, demographics, life cycle, bounds.

JEL classification. D12, D14, D91.

1. Introduction

The effect of demographics on consumption patterns has 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] This study investigates what can be learned from Euler equation estimation of the effect of children on the marginal utility of consumption when households are potentially credit constrained. Although these estimators are now workhorses 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.

Thomas H. Jørgensen: thomas.h.jorgensen@econ.ku.dk
I am grateful to the handling co-editor and two anonymous referees for constructive comments on a previous version of the paper. I also benefited from discussions with 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, Jérôme Adda, Anders Munk-Nielsen, Mette Foged, Damoun Ashournia, Bart Capéau, and conference participants at the EEA, 2013, the CAM December workshop, 2013, the IAAE, 2014, and the ESEM, 2014. Financial support from the Danish Council for Independent Research in Social Sciences is gratefully acknowledged (FSE, Grant 4091-00040). All errors are my own.

[1] Irvine (1978) might be one of the first to suggest that the hump in consumption could be due to changes in household composition. Some important contributions to the literature on the effect of children 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).

DOI: 10.3982/QE492

The present study offers three contributions to this literature. First, I show how conventional Euler equation estimation methods in general produce biased estimates of the effect of children on the marginal utility of 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]

Second, I propose a tractable approach to obtaining bounds on the effect of children on the marginal utility of consumption that is robust to the presence of credit constraints. In particular, I propose to split the sample into young households, in which children might arrive, and older households, in which children might move. A *lower bound* can be estimated by comparing consumption growth of older households across different household composition while an *upper bound* can be estimated using the cohort-average number of children as an instrument, while restricting the sample to young households. I motivate the proposed bounds through an illustrative four-period model and a multiperiod (buffer-stock) life cycle model.

Finally, I find that conventional estimators yield point estimates of the effect of children on the marginal utility of consumption that are *above the upper bound*, estimated from the Panel Study of Income Dynamics (PSID). Existing point estimates from these conventional methods, reported in previous studies, are also above the estimated upper bound. In an influential study by Attanasio et al. (1999), the number of children is found to be important in describing the consumption behavior of U.S. consumers, using the Consumer Expenditure Survey (CEX). This is also the finding in Attanasio and Browning (1995), using the U.K. Family Expenditure Survey (FES), and Browning and Ejrnæs (2009) find that the number and age of children can explain completely the consumption age profile in the FES. The finding that conventional estimators yield point estimates above the upper bound suggests that household demographics might proxy for the inability (or unwillingness) of households to borrow against future income growth. While this concern has been recognized, it has not been analyzed empirically before.[3]

Equivalence scales are closely related to the effect of children on the marginal utility of consumption.[4] Pollak and Wales (1979) first noted that only equivalence scales *conditional* on demographic composition could be identified from observed consumption data. A large body of literature discusses the identification and estimation issues regarding conditional equivalence scales in great detail.[5] To the best of my knowledge,

[2]The fact that ignoring credit constraints produces 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.

[3]For example, Browning and Ejrnæs (2009) recognize that the importance of the number and age of children that they estimate might proxy for credit constraints.

[4]While equivalence scales are informative on how much consumption an additional household member requires to maintain the welfare level of a reference group (a single agent household, say), the parametrization employed herein—and in many existing studies—is how much an additional household member increases the marginal utility of consumption. Bick and Choi (2013) find that the implied consumption behavior from the two approaches is very similar.

[5]See, for example, Lewbel (1989), Blundell and Lewbel (1991), and Ferreira, Buse, and Chavas (1998). Lewbel (1997) provides a survey on the identification and estimation challenges related to equivalence scales.

none of the existing studies considers identification when households face potentially binding credit constraints. The present results imply that not even conditional equivalence scales can in general be uncovered from observed demand data by conventional estimators when households face potentially binding credit constraints.

The results have implications for a broad range of fields in economics. It is common practice by many researchers to rely on parameter estimates from external sources to calibrate economic models. Microeconometric estimates of the effect of children on the marginal utility of consumption (or equivalence scales) are no exception. For example, Cagetti (2003) uses the estimated effect of children on the marginal utility of consumption from Attanasio et al. (1999) when analyzing wealth accumulation and precautionary savings over the life cycle. Scholz, Seshadri, and Khitatrakun (2006) use equivalence scales from Citro and Michael (1995) when concluding that American households have saved "optimally" for retirement and Heathcote, Perri, and Violante (2010) use Organization for Economic Cooperation and Development (OECD) equivalence scales when studying inequality in the United States. It is important to understand the limitations of the econometric techniques employed to uncover these objects because the results of studies subsequently using these estimates might be greatly affected by externally calibrated parameters. The proposed bounds can inform researchers in which *range* the effect of demographic variables on the marginal utility of consumption may lie. The bounds could be used to perform sensitivity analysis to investigate the importance of demographic taste shifters in generating a particular result.

The present study relates to a recent strand of literature investigating the validity of Euler equation estimators. For example, Ludvigson and Paxson (2001) and Carroll (2001) argue that using the log-linearized Euler equation to estimate 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) argue that the contradictory results are due to differences in the time series dimension in the implemented Monte Carlo studies. Specifically, the bias in Euler equation estimators of the IES might be small when interest rates vary sufficiently over time and the time dimension is long, as in Attanasio and Low (2004). These studies all focus on the IES and ignore potentially binding credit constraints.

Empirical evidence suggests that credit constraints are important, especially for young consumers. Since Thurow (1969) suggested that borrowing constraints could explain the observed consumption age profile, a growing empirical literature finds evidence supporting an important role for credit constraints in explaining observed consumer behavior.[6] The results also generalize to cases in which consumers *do not* face an "explicit" credit constraint. If consumers instead face a positive probability of a zero income shock (as in, e.g., Carroll (1997) and Gourinchas and Parker (2002)), most results still hold. This is because risk averse consumers will instead face a "self-imposed"

[6]Some important contributions to this literature are due to Hall and Mishkin (1982), Zeldes (1989a), Jappelli (1990), Runkle (1991), Jappelli, Pischke, and Souleles (1998), Gross and Souleles (2002), Johnson, Parker, and Souleles (2006), Leth-Petersen (2010), Gross, Notowidigdo, and Wang (2014), and Crossley and Low (2014).

no-borrowing constraint (Schechtman (1976), Zeldes (1989b)) and consumption will respond substantially to transitory income shocks leading the log-linearized Euler equation to be a poor approximation.[7]

The rest of the paper proceeds as follows. The following section presents the constrained 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 the marginal utility of consumption when households face potentially binding credit constraints and suggests how bounds can be estimated using these methods. Section 4 reports estimated bounds using the PSID suggesting that estimates from conventional estimators are above the proposed upper bound. Section 5 discusses the robustness of the bounds and Section 6 concludes.

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 et al. (1999) and let children affect the *marginal utility* 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. Alternatively, the household composition could be included as a scaling of the level of consumption (equivalence scaling), as done, for example, in Fernández-Villaverde and Krueger (2007).[8]

The *constrained* consumption Euler equation is[9]

$$u'(C_t)v(\mathbf{z}_t;\theta)-\lambda_t=R\beta\mathbb{E}_t\big[u'(C_{t+1})v(\mathbf{z}_{t+1};\theta)\big], \tag{1}$$

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, and $u'(C_t)$ is the marginal utility from consuming C_t. Rearranging equation (1) yields

$$R\beta\mathbb{E}_t\bigg[\frac{u'(C_{t+1})}{u'(C_t)}\frac{v(\mathbf{z}_{t+1};\theta)}{v(\mathbf{z}_t;\theta)}\bigg]=1-\frac{\lambda_t}{u'(C_t)v(\mathbf{z}_t;\theta)}$$

in which the left hand side is a familiar term while the right hand side includes an additional part related to the shadow price of resources. Reformulating in "error form" yields

$$R\beta\frac{u'(C_{t+1})}{u'(C_t)}\frac{v(\mathbf{z}_{t+1};\theta)}{v(\mathbf{z}_t;\theta)}=\epsilon_{t+1}, \tag{2}$$

[7]This is the point of Carroll (2001) where he illustrates how this poor first (and second) order approximation of the nonlinear Euler equation results in poor estimates of the IES.

[8]See 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.

[9]See Zeldes (1989a, footnote 9) for a derivation.

where the structural Euler error, ϵ_{t+1}, satisfies

$$\mathbb{E}_t[\epsilon_{t+1}] = 1 - \frac{\lambda_t}{u'(C_t)v(\mathbf{z}_t;\theta)}.$$

From the Kuhn–Tucker conditions we know that $\lambda_t \geq 0$ in all time periods. Hence, the mean expectational error in (2) equals 1 only if consumers are not affected by the credit constraint. Generally, however, the expectational error in (2) is a function of information today and is potentially serially correlated. In the existing literature on intertemporal consumption allocation and the effect of children on the marginal utility of consumption, credit constraints are often ignored or assumed away. It is clear from (2), that estimators that ignore credit constraints suffer from something similar to an "omitted variable bias." Below, I discuss the two most common types of 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 $t = 1, \dots, T$. As is common in the literature, assume that utility is constant relative risk aversion (CRRA) with risk aversion parameter ρ and let $v(\mathbf{z}_t;\theta) = \exp(\theta'\mathbf{z}_t)$. Ignoring potentially binding credit constraints (i.e., imposing $\lambda_s = 0\ \forall s$) and inserting the functional form assumptions, a nonlinear (sometimes referred to as "exact") general method of moments (GMM) estimator of θ could be

$$\theta_{\text{GMM}} = \arg\min_{\theta}\left[\frac{1}{NT}\sum_{i}^{N}\sum_{t}^{T}\left(R\beta\left(\frac{C_{i,t+1}}{C_{i,t}}\right)^{-\rho}\exp(\theta\Delta\mathbf{z}_{i,t+1}) - 1\right)\cdot Z_{i,t+1}\right]^2, \tag{3}$$

such that θ_{GMM} is the parameter that satisfies the sample equivalent of $\mathbb{E}[(\epsilon - 1)'Z] = 0$, where Z contain instrument(s) assumed to be uncorrelated with the Euler residual. Ignoring measurement error, the estimator in (3) produces consistent estimates if a suitable instrument is available and, importantly, households do not face credit constraints.[10]

Using food consumption from the PSID, Alan, Attanasio, and Browning (2009) estimate that a child increases marginal utility by approximately 20 percent ($\hat{\theta} = 0.18$) from a similar estimator as (3). When allowing for measurement error in consumption they estimate an increase of as much as 145 percent ($\hat{\theta} = 0.9$).[11]

Most existing studies work with a log-linearized approximation of the Euler equation since it yields estimable equations linear in parameters that 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

$$\Delta \log C_{it} = \text{constant} + \rho^{-1}\theta'\Delta\mathbf{z}_{it} + \tilde{\epsilon}_{it}, \tag{4}$$

[10] Alan, Attanasio, and Browning (2009) supply modified GMM estimators to allow for measurement error.

[11] Because of the exponential effect of children on the marginal utility of consumption, the percentage increase is even larger for additional children.

where the first term is a constant as a function of structural parameters (β, ρ) and the interest rate (assumed constant here), 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 et al. (1999), θ and ρ are 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 in the marginal utility of consumption is found to be around $\theta \approx 0.33$, corresponding to a 39 percent increase in the marginal utility from the first child. Several studies use food consumption from the PSID to estimate versions of the log-linearized Euler equation; see, for example, Hall and Mishkin (1982), Runkle (1991), Lawrance (1991), and Dynan (2000). The latter estimates $\theta \approx 0.7$. Browning and Ejrnæs (2009) allow for a more flexible functional form of $v(\mathbf{z}_t; \theta)$ when estimating the effect of children on the marginal utility of consumption using the FES and find that the number and age of children can explain completely the hump in the consumption age profile.

Other estimators have been proposed to estimate Euler equations. For example, Alan and Browning (2010) propose a method in which they fully parametrize the Euler residuals and simulate these residuals and consumption paths simultaneously. Their synthetic residual estimation (SRE) procedure assumes away credit constraints. Since the GMM and log-linearized estimation methods are the conventional methods used in the literature, I focus exclusively on these.

3. Bias and bounds from Euler equation estimation

In this section, I illustrate how conventional Euler equation estimators, (3) and (4), produce biased estimates of the effect of children on the marginal utility of consumption if households face potentially binding credit constraints. I argue, however, that these methods 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. I assume here that a panel of consumers is available and only θ is to be identified.

3.1 *Evidence from a four-period model*

Here, I set up 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, $\mathbf{z}_1 = 1$, in p percent of the households and the remaining $1 - p$ percent remain childless, $\mathbf{z}_1 = 0$. In period $t = 2$, the "old" stage, the child moves (if present in period one) such that $\mathbf{z}_2 = 0$ for all households. Households die with certainty in the end of period $t = 3$ and consume all available resources.

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 notation, the gross real interest rate and the discount factor both equal 1, $R=\beta=1$. Households receive a deterministic income of Y_t in beginning of every period. Income grows with G_1 between period zero and period one ($Y_1=G_1Y_0$) and is constant otherwise ($Y_t=Y_{t-1}$, $t=2,3$). The beginning-of-period resources available for consumption are 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

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

subject to an explicit no-borrowing constraint, $A_t\geq 0\ \forall t$. Section S.1 in the supplementary material available in a file on the journal website, http://qeconomics.org/supp/492/supplement.pdf, solves the model analytically and reports the resulting optimal consumption functions. Replication files are available as supplementary files on the journal website, http://qeconomics.org/supp/492/code_and_data.zip.

Using the optimal consumption behavior from this model, Figure 1 presents consumption and wealth profiles for households initiated with no wealth, $A_{-1}=0$, income normalized to 1 ($Y_0=1$), and early life income growth of 8 percent, $G_1=1.08$. Panel 1(a) presents consumption profiles for models with a credit constraint (solid) and without a constraint (dashed) for households with children in period one (black) and without children (gray). Panel 1(b) illustrates the associated wealth profiles. Po-

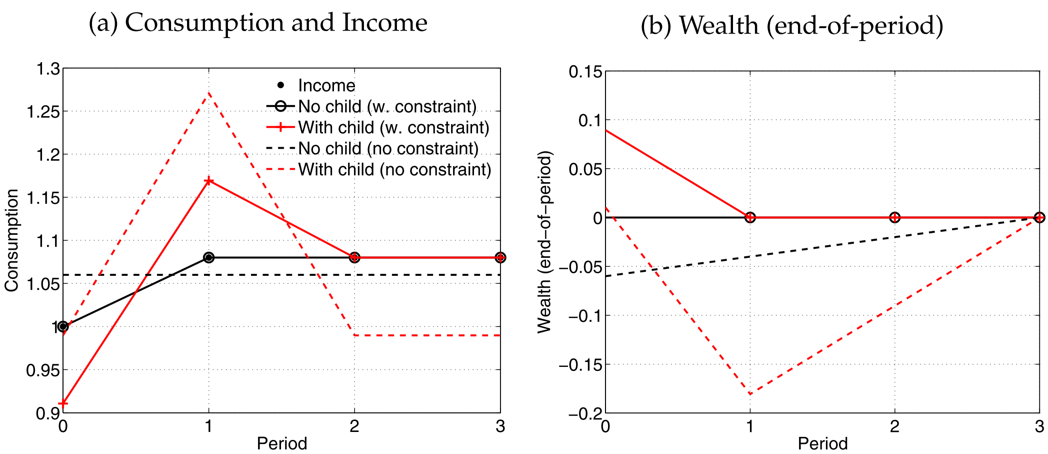

FIGURE 1. Consumption and wealth profiles from the four-period model with parameters $\rho=2$, $G_1=1.08$, and $\theta=0.5$. Households are initiated with zero wealth and initial income is normalized to 1, that is, $A_{-1}=0$ and $Y_0=1$, respectively. Panel (a) presents the income and consumption profiles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

tentially 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 that have a child 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. Both effects lead to a bias in the ordinary least squares (OLS) estimator of the effect of children on the marginal utility of consumption when comparing log consumption growth of households that have children with childless households. The OLS estimator using consumption growth from $t - 1$ to t from the log-linearized Euler equation (4) is given by

$$\hat{\theta}_{\text{OLS}}^{\text{young}} = (\Delta \log C_1|_{\mathbf{z}_1=1} - \Delta \log C_1|_{\mathbf{z}_1=0})\rho,$$
$$\hat{\theta}_{\text{OLS}}^{\text{old}} = -(\Delta \log C_2|_{\mathbf{z}_1=1} - \Delta \log C_2|_{\mathbf{z}_1=0})\rho.$$

Since the credit constraint invalidates the OLS estimator (see discussion above and below), a natural route is to search for usable instruments that might not be as affected by potentially binding credit constraints while being correlated with the growth in the number of children. Motivated by a large body of literature that estimates Euler equations using synthetic cohort panels, a candidate instrument could be the cohort-average number of children (p). Using the average realized fertility as an instrument effectively utilizes time variation in fertility over the life cycle to identify θ rather than using cross sectional variation in the arrival of children (as the OLS estimator does). In turn, this instrument is highly correlated with the number of children inside each household while it is much less affected by idiosyncratic credit constraints (and income uncertainty). The instrumental variable (IV) estimator using p as an instrument is[12]

$$\hat{\theta}_{\text{IV}}^{\text{young}} = \frac{1}{p}\big(p\Delta \log C_1|_{\mathbf{z}_1=1} + (1-p)\Delta \log C_1|_{\mathbf{z}_1=0}\big)\rho,$$
$$\hat{\theta}_{\text{IV}}^{\text{old}} = -\frac{1}{p}\big(p\Delta \log C_2|_{\mathbf{z}_1=1} + (1-p)\Delta \log C_2|_{\mathbf{z}_1=0}\big)\rho.$$

Section S.1 in the supplementary material derives explicit formulas for each estimator when using either young households or old households to estimate the effect of children on the marginal utility of consumption. The resulting estimators are

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

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

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

It is immediately clear from these estimators that both OLS and IV estimators will, in general, yield biased estimates of the true θ. Interestingly, as the effect of children goes toward 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 the marginal utility of consumption gets increasingly large, the positive bias part, $(1-p)/p \log(G_1)\rho$, of the young IV estimate, $\hat{\theta}_{\text{IV}}^{\text{young}}$, 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 to 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 results are intuitive. Young households will accumulate wealth in the initial period $t=0$, but not necessarily enough to ensure that the credit constraint is not binding in period $t=1$, in which a child arrives. Even if they do accumulate enough wealth, the fact that the childless households also increase consumption in period $t=1$ creates a downward bias in the OLS estimate. When children subsequently leave, households with children are likely to go from being constrained in period $t=1$ to unconstrained in period $t=2$ (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 downward biased. The IV estimator is upward biased because income is positively correlated with the (cohort-)average number of children in the early part of the life cycle and the credit constraint induces a positive correlation between income and consumption. Thus, the IV estimate from the young sample can be used as an upper bound.

Interestingly, for low levels of θ ($0 \le \theta \le \log(G_1)\rho$) the bounds are flat, illustrating how the inability to borrow against future income growth prevents identification of the effect of children on the marginal utility of 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 bias disappears when income is constant because the only source of consumption growth is, then, the presence of children. In turn, only households that have children will increase consumption and θ can be identified. The bias also disappears if households do not care about intertemporal smoothing of marginal utility ($\rho = 0$ and IES $= \infty$). Income growth and a finite intertemporal elasticity of substitution seem to be reasonable assumptions, however.

3.2 *Evidence from a multiperiod life cycle model*

To confirm the results from the four-period model, I set up a standard life cycle (buffer-stock) model, used extensively 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 et al. (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

$$\max_{C_t} \mathbb{E}_t\left[\sum_{\tau=t}^{T_r-1} \beta^{\tau-t} v(\mathbf{z}_t;\theta)u(C_\tau) + \gamma \sum_{s=T_r}^{T} \beta^{s-t} v(\mathbf{z}_s;\theta)u(C_s)\right]. \tag{5}$$

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 (5) 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 (5) 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 the 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 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,

$$Y_t = P_t \varepsilon_t \quad \forall t < T_r,$$
$$P_t = G_t P_{t-1} \eta_t \quad \forall t < T_r,$$

where G_t is the real gross income growth, P_t denotes permanent income, $\log \eta_t \sim \mathcal{N}(-\sigma_\eta^2/2, \sigma_\eta^2)$ is a mean 1 permanent income shock, and ε_t is a mean 1 transitory income shock taking the value μ with probability $\wp$ and otherwise equal to $(\tilde{\varepsilon}_t - \mu\wp)/(1-\wp)$, where $\log \tilde{\varepsilon}_t \sim \mathcal{N}(-\sigma_\varepsilon^2/2, \sigma_\varepsilon^2)$.[13] When retired, the income process is a deterministic fraction $\varkappa \leq 1$ of permanent income and permanent income is constant once retired, $Y_t = \varkappa P_t\ \forall t \geq T_r$ and $P_t = P_{t-1}\ \forall t \geq T_r$.

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 $\mathbf{z}_t$,

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

[13]This formulation allows for both an explicit and self-imposed credit constraint. Depending on the value of κ, $\wp$, and μ, either the explicit or the self-imposed constraint will be the relevant one. This is discussed further in Section S.2 in the supplementary 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-linarized Euler equation are 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.

TABLE 1. Parameter values used to simulate data.

G_t	R	σ_ε^2	σ_η^2	κ	$\wp$	μ	β	ρ	γ	$\varkappa$	θ
Figure 2(a)	1.03	0.005	0.005	0	0	0	0.95	2	1.1	0.8	$\in [0, 1]$

where NC refers to no child, and the oldest child is denoted child 1, the second oldest child as child 2, and the third oldest child as child 3. 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.[14]

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 parameters presented in Table 1. The technical details of the solution method are provided in Section S.2 in the supplementary material. The solution is then used to generate data for 50,000 households from age 22 to 59 in each of the 1000 Monte Carlo (MC) runs. All households start at age 22 with zero wealth, $A_{21} = 0$, permanent income of 1 (normalization), $P_{22} = 1$, and no previous children, $\mathbf{z}_{21} = (\text{NC}, \text{NC}, \text{NC})$. Children are distributed across households and age according to the observed arrival of children in the PSID, as illustrated in Figure 2(b), and the income profile is calibrated to be concave (Figure 2(a)) and constant from age 40 to mimic empirical income profiles.

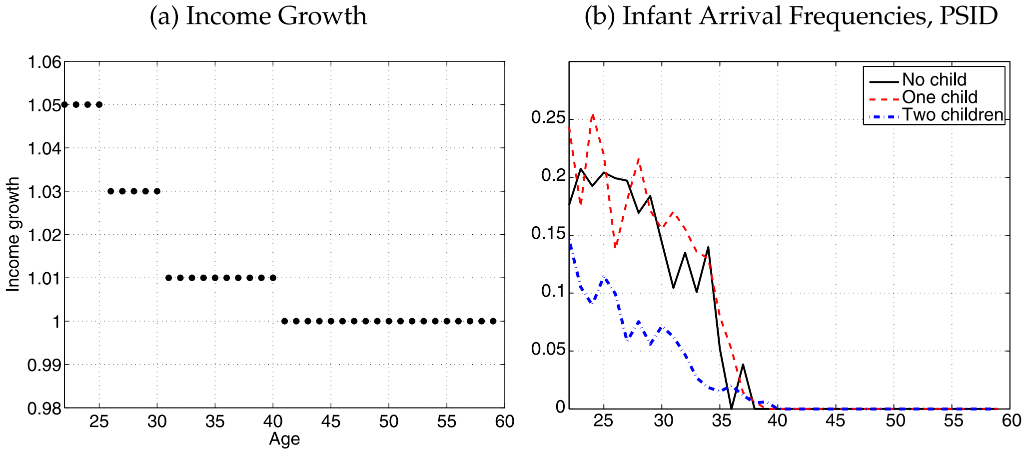

FIGURE 2. Calibrated income growth and arrival of children. Panel 2(a) reports how permanent income grows in the life cycle model while panel (b) shows how the arrival of children is calibrated using the PSID. The arrival of children is based on the PSID data described in Section 4.

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

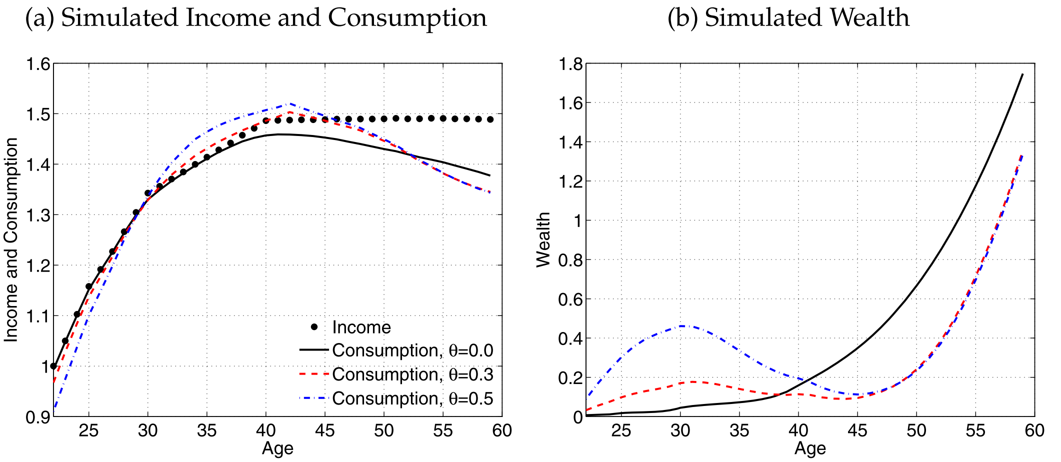

FIGURE 3. Simulated income, consumption, and wealth profiles. The figure illustrates the average age profile of income, consumption and wealth for 50,000 simulated households for different values of θ. Panel (a) shows how consumption profiles change relatively little across models with no effects of 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. In particular, a hump in the wealth profile emerges as children becomes more important.

The discount factor and interest rate is based on Gourinchas and Parker (2002) while the constant relative risk aversion, ρ, is set between those estimated in Gourinchas and Parker (2002) and Cagetti (2003). The transitory and permanent income shock variances are slightly lower than most studies. The permanent income shock variance is, however, in the range reported in Blundell, Pistaferri, and Preston (2008).[15] The replacement rate in retirement, $\varkappa$, is fixed at 80 percent, in between the median replacement rates of 70 percent in the English Longitudinal Study of Ageing (ELSA), reported in Banks, O'Dea, and Oldfield (2010), and the median replacement rate of 90 percent reported in The Danish Ministry of Finance (2003). The post-retirement motive is fixed at 1.1, slightly lower than the estimates in Jørgensen (2015) of a similar model using Danish data.

Figure 3 presents simulated age profiles for income, consumption and wealth for different values of θ. All consumption profiles (even if $\theta = 0$) 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

[15] The low income shock variances were originally chosen because the quality of the log-linearized Euler equation approximation decreases with the income shock variances.

growth to increase consumption when children arrive—despite wanting to, had unlimited borrowing been possible. Hence, the effect of children would, in general, be underestimated using young households, as shown earlier. Noticeably, young households accumulate large amounts of wealth in anticipation of having children in the future. When they subsequently have children they almost deplete their wealth such that the credit constraint is binding (or close to binding) for many households when children eventually move.

The accumulation of wealth in anticipation of children and decumulation of wealth when children reside produce hump-shaped wealth age profiles, illustrated in Figure 3. Empirical age profiles of household wealth (or assets) are typically not hump shaped but rather monotonically increasing (Cagetti (2003)) suggesting that children might not be as important for consumption over the life cycle as previously found in the existing literature. In Section 4 below, I confirm this result by showing that existing estimates of θ lie above the estimated upper bound.

Table 2 reports the average estimate of θ using all households, both young and old, from 1000 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 (LogLin) and nonlinear GMM estimators overestimate the effect of children on the marginal utility of consumption while they underestimate the effect if θ is large. This is true regardless of whether the actual change in number of children ($\Delta\mathbf{z}_t$) is used in the estimation or the cohort-average number of children ($\Delta\bar{\mathbf{z}}_t$) is used as the instrument.[16] All results are based on individual-level Euler equations.

Figure 4 illustrates the proposed bounds based on the four-period model in panel 4(a) and the multiperiod life cycle model in panel 4(b). The 45° line represents the true value of θ while black lines represent lower bounds and gray lines represent up-

Table 2. Monte Carlo results, both young and old households.

	$\theta = 0.0$		$\theta = 0.1$		$\theta = 0.5$		$\theta = 1.0$	
Instr.	LogLin	GMM	LogLin	GMM	LogLin	GMM	LogLin	GMM
$\Delta\mathbf{z}_t$	0.015	0.006	0.086	0.078	0.215	0.208	0.357	0.328
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
$\Delta\bar{\mathbf{z}}_t$	0.125	0.038	0.155	0.073	0.201	0.129	0.342	0.224
	(0.002)	(0.001)	(0.002)	(0.001)	(0.002)	(0.002)	(0.003)	(0.002)

Note: The average of all MC estimates and standard deviations (in parentheses) across Monte Carlo runs are reported. All results are based on 1000 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 1, $P_{22} = 1$, and no children. Children are assigned following the estimated arrival probabilities estimated from the PSID, reported in Figure 2(b).

[16] The results remained qualitatively unchanged when including lagged income as an additional regressor. The results are not reported but are available upon request.

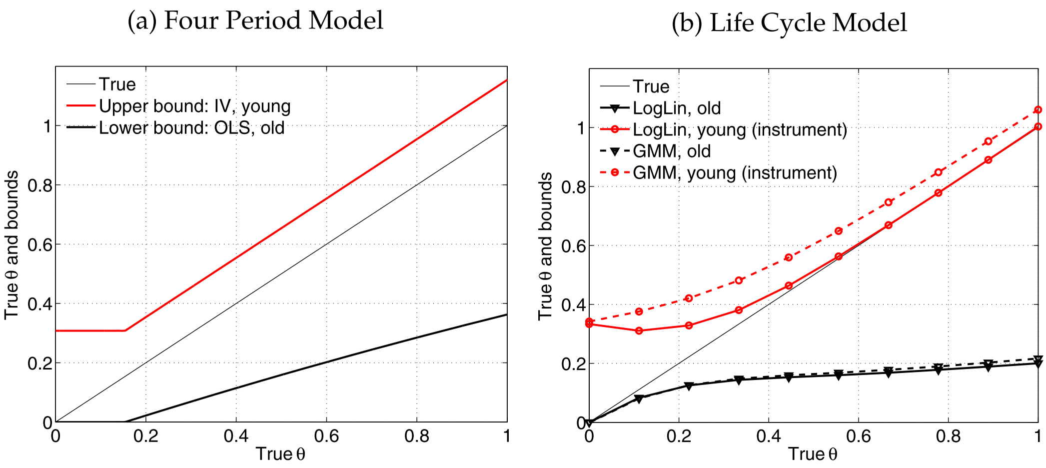

FIGURE 4. Proposed bounds for four-period model and life cycle model. The 45° line represents the true value of θ while black lines represent lower bounds and red/gray lines represent upper bounds using the (cohort-)average number of children as instrument. Solid lines are based on the log-linearized Euler equation (4) and dashed lines are based on the nonlinear Euler equation estimated with the GMM estimator (3). Young households are defined as younger than 41.

per 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 intuition is also the same. Although households are heterogeneous in when children arrive, households have not necessarily accumulated enough wealth when children arrive to be able to increase consumption as much as desired. Importantly, income growth (and the inability to borrow against it in previous periods) leads to increased current consumption irrespective of the arrival of children. Thus comparing households in which an additional child arrives ($\Delta \mathbf{z}_t > 0$) with those in which composition is unchanged ($\Delta \mathbf{z}_t = 0$) will result in an underestimation of θ. As argued in relation to the four-period model, using the cohort-average number of children as the instrument will be less affected by idiosyncratic credit constraints and income shocks while being highly correlated with the arrival of children. As income growth is positively correlated with the average number of children, the IV estimator using young households will be upward biased. In turn, using young households, in which children arrive, and the average number of children as instrument provides an *upper bound*.

Households decumulate their wealth while children reside at home. In general it might be that not *all* households have been able to accumulate enough wealth prior to the arrival of children to perfectly smooth marginal utility over the life cycle if they are not allowed to borrow. In turn, they are potentially on the credit constraint when children subsequently 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.[17] Hence, the effect of children would be underestimated when only using older households as shown using the simple four-period model above. Thus, using older households, in which children leave, can be used to estimate a *lower bound* of the effect of children on the marginal utility of consumption from changes in actual household composition.

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 nonlinear GMM estimator produces bounds almost identical to 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 is the log-linearized Euler equation.

The effect of children on the marginal utility of consumption cannot, in general, be identified from the observed consumption behavior when households face potentially binding credit constraints. Using the Euler equation to estimate θ will produce flawed results unless *all* households have accumulated enough wealth to be unaffected by the potential binding credit constraint. Table 3 reports Monte Carlo results from versions of the model in which the motives for saving are higher than in the baseline calibration. Not surprisingly, when credit constraints have less bite, the *unconstrained* Euler equa-

TABLE 3. Monte Carlo results for stronger saving motives.

	$\theta = 0.0$		$\theta = 0.1$		$\theta = 0.5$		$\theta = 1.0$	
Instr.	LogLin	GMM	LogLin	GMM	LogLin	GMM	LogLin	GMM
				More patient, $\beta = 0.99$				
$\Delta \mathbf{z}_t$	0.001	−0.000	0.101	0.100	0.482	0.485	0.740	0.768
	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)
$\Delta \bar{\mathbf{z}}_t$	0.009	−0.000	0.108	0.100	0.481	0.480	0.705	0.715
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
				More risk averse, $\rho = 3$				
$\Delta \mathbf{z}_t$	0.012	0.006	0.107	0.101	0.330	0.340	0.490	0.495
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
$\Delta \bar{\mathbf{z}}_t$	0.095	0.037	0.193	0.132	0.317	0.286	0.455	0.400
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)
				No income growth, $G_t = 1$				
$\Delta \mathbf{z}_t$	0.003	0.001	0.080	0.079	0.249	0.252	0.423	0.417
	(0.001)	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
$\Delta \bar{\mathbf{z}}_t$	0.023	0.003	0.067	0.053	0.206	0.183	0.393	0.326
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.002)

Note: The average of all MC estimates and standard deviations (in parentheses) across Monte Carlo runs are reported. In each panel, one parameter is changed to show how that affects the results. Remaining parameters are fixed at their baseline values in Table 1. See the notes to Table 2.

[17] If households do not deplete their wealth, they will be less likely to be affected by credit constraints. In this case the lower bound will still be valid and in the extreme case in which the constraints have *no* bite for *any* households in the data, the lower bound will equal the true effect of children on the marginal utility of consumption.

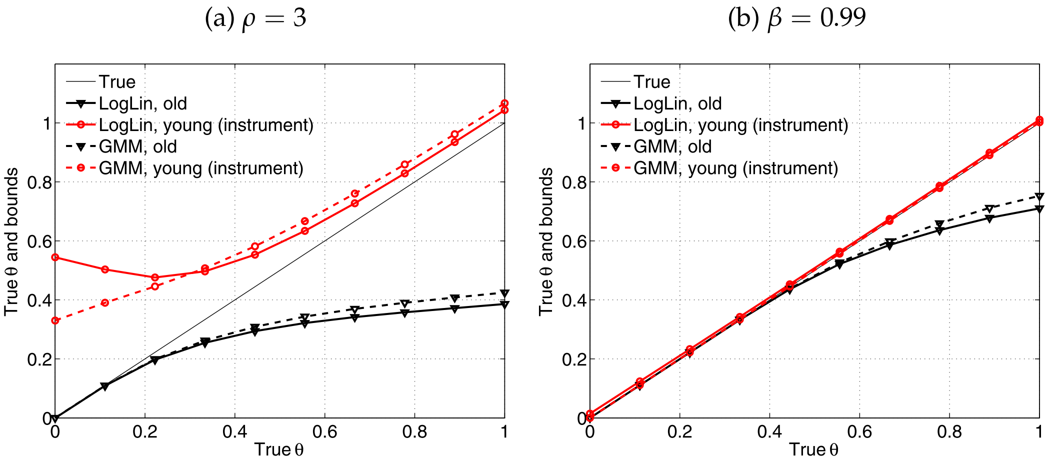

FIGURE 5. Bounds for alternative calibrations. This figure illustrates the proposed bounds based the life cycle model with a stronger motive for wealth accumulation than the baseline parametrization. In particular, $\rho = 3$ in panel (a) and $\beta = 0.99$ in panel (b). See the notes to Figure 4.

tion is a reasonable approximation and, thus, can in certain specifications uncover the true effect of children on the marginal utility of consumption. In particular, for lower values of θ, Euler equation estimation is able to uncover the true effect of children on the marginal utility of consumption. As noted in the Introduction, most empirical evidence point toward an important role for credit constraints in explaining the behavior of young households in particular.

While Euler equation estimation cannot, in general, point identify θ, if households face potentially binding credit constraints, the bounds are rather robust. In particular, even if all households have accumulated enough wealth to be unaffected by credit constraints, the bounds are still valid and rather tight, as illustrated in Figure 5. Although Euler equation estimation techniques *might* uncover the true point estimate under certain (special) cases, the fact that researchers have limited information on whether that special case is the relevant one should lead to the use of more robust methods. The proposed bounds are an attempt to provide researchers with a tractable, yet more robust, alternative. Below, I estimate the proposed bounds using the PSID and in Section 5, I discuss the robustness of the proposed bounds.

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 the marginal utility of consumption. To study the evolution and link between income and consumption inequality over the 1980s, Blundell, Pistaferri, and Preston (2008) impute total nondurable 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 1978–1992 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. I restrict the sample to cover households in which the husband is aged 20–59 and the supplementary low-income subsample (SEO) is excluded from the analysis.[18] All sample selection criteria leave an unbalanced panel of 1885 households observed for at most 13 periods with a total of 16,927 nonmissing observations. The results from nondurable consumption exclude the years 1978 and 1979, leaving a total of 14,144 households with nonmissing nondurable consumption.[19] Households are classified as high skilled if the male head has ever enrolled in college, including college dropouts.

Table 4 reports the estimated bounds of $\rho^{-1}\theta$ for the PSID data using the log-linearized Euler equation (4) 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 ($\Delta \mathbf{z}_t$) while restricting the estimation sample to include only older households and an upper bound can be found by using the cohort-average number of children ($\Delta \bar{\mathbf{z}}_t$) as an instrument while restricting the estimation sample to younger households.

Table 4. Log-linear Euler equation bound estimates, PSID.

	Low Skilled		High Skilled	
	OLS, age $\geq 45^{\dagger}$	IV, age $\leq 45^{\ddagger}$	OLS, age $\geq 45^{\dagger}$	IV, age $\leq 45^{\ddagger}$
	Food consumption			
$\Delta\#kids$	0.049	0.144	0.036	0.128
	(0.028)	(0.043)	(0.025)	(0.051)
Constant	0.108	0.092	0.043	0.055
	(0.031)	(0.015)	(0.022)	(0.017)
Obs.	1651	5776	2237	4515
	Nondurable consumption (imputed)			
$\Delta\#kids$	0.001	0.009	−0.023	0.044
	(0.028)	(0.043)	(0.027)	(0.062)
Constant	0.133	0.106	0.102	0.119
	(0.025)	(0.015)	(0.023)	(0.022)
Obs.	1415	4699	1807	3550

Note: 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 nondurable consumption, imputed by Blundell, Pistaferri, and Preston (2008), are reported. Robust standard errors are given in parentheses. All regressions include year dummies. Households are classified as high skilled if the male head has ever enrolled in college, including college dropouts. 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.

[18] Blundell, Pistaferri, and Preston (2008) use households in which the husband is aged 30–65.

[19] The imputed level of total nondurable consumption in 1978 and 1979 seems to be extreme compared to nondurable consumption in subsequent years. The level is around an order of magnitude higher in these two years without the same extreme pattern in food consumption. I did not investigate this further and chose to exclude these years from the analysis when using nondurable consumption.

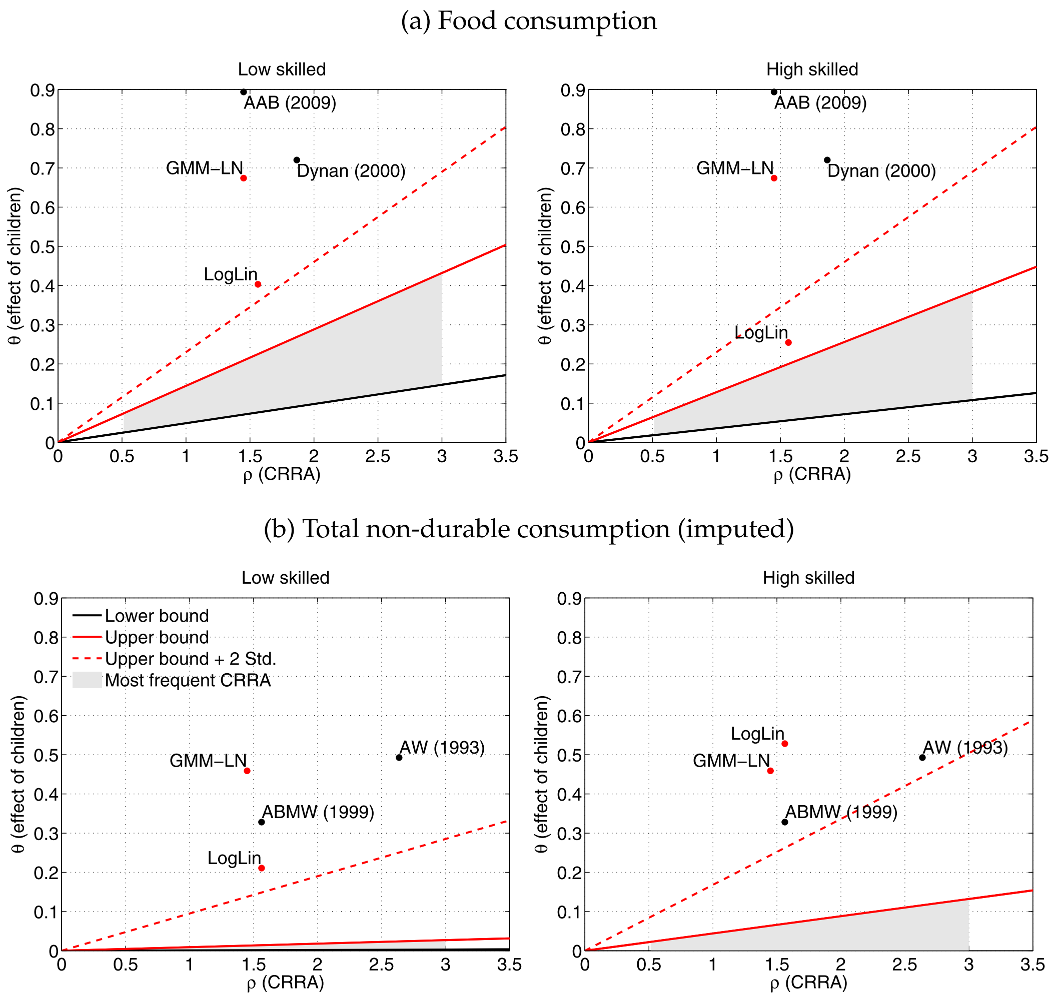

FIGURE 6. Estimated bounds of the effect of children on the marginal utility of consumption from the PSID. The figure reports the upper and lower bounds for low and high skilled households when varying ρ, the inverse of elasticity of intertemporal substitution. The top panel (a) reports results when using changes in log food consumption while the bottom panel (b) reports results when using changes in log nondurable consumption, imputed by Blundell, Pistaferri, and Preston (2008). LogLin and GMM-LN refer to estimation results from log-linearized Euler equation estimation and a nonlinear GMM estimator, respectively. AW (1993) refers to Attanasio and Weber (1993), ABMW (1999) refers to Attanasio et al. (1999), AAB (2009) refers to Alan, Attanasio, and Browning (2009), and Dynan (2000) is obvious.

Figure 6 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 6(a)) reports results when using changes in log food consumption while the bottom panel (panel 6(b)) reports results when using changes in log nondurable consumption. While the bounds do not differ significantly across educational groups, the bounds based on nondurable consumption are significantly lower than those based on food consumption.

Point estimates from four existing studies are illustrated in Figure 6 together with estimation results using the current PSID sample. In particular, log-linearized Euler equation estimation results (LogLin) and estimation results from a nonlinear Euler equation estimator (GMM-LN) allowing for log-normal multiplicative measurement error in consumption are included in Figure 6.

When implementing the log-linearized Euler equation estimator (LogLin) I follow closely the approach in Attanasio et al. (1999) and estimate by two-stage least squares (2SLS) using a synthetic cohort panel the effect of children on the marginal utility of consumption. The Appendix includes a detailed description of the implementation and estimation results. When implementing the GMM-LN estimator, I follow closely the approach suggested in Alan, Attanasio, and Browning (2009). As discussed in the Appendix, although I estimate similar effects of children on the marginal utility of consumption as reported in these studies, I estimate implausible low risk aversion parameters.[20] I thus report in Figure 6 estimated effects of children on the marginal utility of consumption while imposing risk aversion parameters from these studies.[21]

The estimated effects of children on the marginal utility of consumption estimated herein and reported in the existing literature are all above the estimated bounds. Specifically, the reported estimate of $\rho^{-1}\theta \approx 0.21$ in the influential study by Attanasio et al. (1999) and their estimated ρ^{-1} of 0.64 imply an effect of children on the marginal utility of *nondurable consumption* in the CEX of around $\theta \approx 0.33$. This is above the upper bound reported in Figure 6. Attanasio and Weber (1993) applied a similar estimation strategy using total expenditure in the FES and estimated an effect of children on the marginal utility of consumption that is also above the upper bound.

Panel (b) in Figure 6 maps the implied estimated effect of children on the marginal utility of *food consumption* in the PSID reported in Dynan (2000) and Alan, Attanasio, and Browning (2009). The latter is based on the GMM-LN estimator while Dynan (2000) estimate a log-linearized Euler equation allowing for habit formation. She finds no evidence of habit formation in food consumption and I, thus, interpret her estimates in light of the current model. All these estimates are outside the upper bound. Adding two times the standard error of the estimated $\widehat{\rho^{-1}\theta}$ reported in Table 4 widens the bounds significantly without including any existing estimates. Only the results from the log-linearized Euler equation estimation using food consumption of high skilled households in the current PSID sample yield an estimated effect of children on the marginal utility of consumption that is close to the upper bound.

[20] I follow Alan, Attanasio, and Browning (2009) and utilize the 3-month Treasury bill rate net of the (annual average) inflation based on the U.S. consumer price index to calculate the gross real interest rate.

[21] In particular, I use $\rho = 1.54$ from Attanasio et al. (1999) when plotting the log-linear Euler equation estimates from columns (2) in Table A.1 (LogLin) and $\rho = 1.45$ from Alan, Attanasio, and Browning (2009) when plotting the estimated θ from columns (2) in Table A.2 (GMM-LN).

5. Robustness of the bounds

The bounds require that researchers simultaneously identify or have knowledge on other structural parameters. In particular, using the log-linearized Euler equation to estimate bounds requires information on the risk aversion parameter while the exact GMM estimation approach also requires knowledge on the discount factor, β. This is a drawback, but varying these parameters in "accepted" ranges produces a set of bounds of the effect of children on the marginal utility of consumption.

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 from when children arrive to when they leave; compare to the above discussion. 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. Estimating different parameters related to when children arrive and move could be yet another route to pursue. An alternative route to estimating bounds could be to utilize the moment *inequality* rather than the equality in the GMM estimator (3). 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.[22]

Below, I argue that the proposed bounds are more robust to a host of alternative economic environments than the baseline model used throughout. In particular, I investigate the robustness of the bounds to (i) alternative fertility processes, (ii) self-imposed no-borrowing rather than explicit credit constraints, (iii) age effects of children, (iv) labor market costs of children, and (v) multiplicative measurement error in consumption.

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 et al. (forthcoming) produce qualitatively unchanged results. Figure 7 illustrates the proposed bounds based on the four-period model in panel 7(a) and the life cycle model in panel 7(b) for the probabilistic version of the models. The bounds are very similar to those presented from the baseline model.

In the probabilistic version, households are identical prior to the arrival of children (given household composition, age, wealth, and income). In the four-period model, all households save exactly the same in period $t = 0$, prior to a child potentially arriving in period $t = 1$. In this period, childless households increase consumption due to the fact that it has been revealed to them that they will remain childless and accumulated wealth

[22] I am grateful to Dennis Kristensen for pointing this out to me.

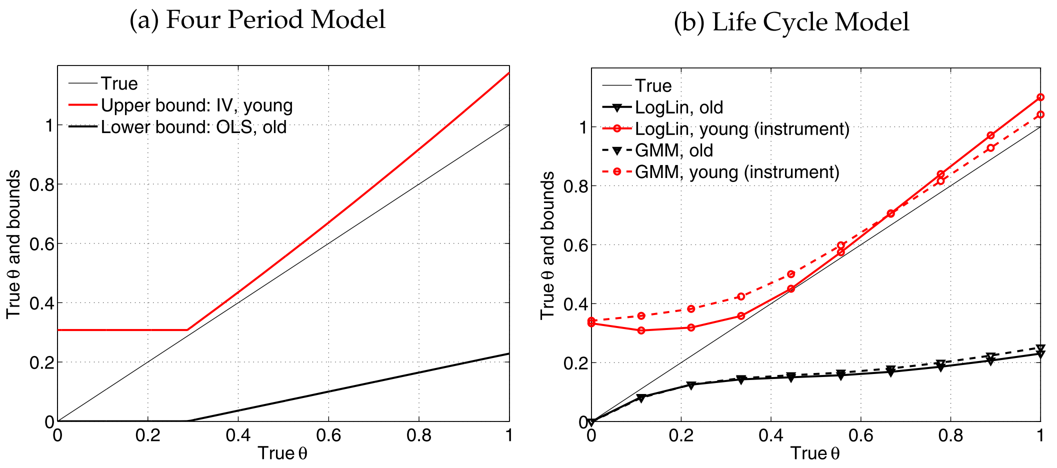

FIGURE 7. Bounds for the probabilistic arrival of children. The proposed bounds are based on a model in which children arrive probabilistically rather than being perfectly foreseen as in the "deterministic" baseline model. Panel (a) illustrates the bounds from the four-period model; the bounds from the life cycle model are illustrated in panel (b). The 45° line represents the true value of θ while black lines represent lower bounds and red/gray lines represent upper bounds using the (cohort)average number of children as instrument. Solid lines are based on the log-linearized Euler equation (4) and dashed lines are based on the nonlinear Euler equation estimated with the GMM estimator (3). Young households are defined as younger than 41.

from period $t = 0$ is distributed across remaining periods. This increased consumption of childless households will bias the estimate downward.[23]

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 the marginal utility of 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. Further, the biological "constraint" through reduced female fecundity will interplay with the financial constraints and the latter is, thus, still likely to be important for household behavior in a model in which fertility is perfectly controlled by households (Almlund (2013)).

5.2 *Self-imposed no-borrowing versus explicit credit constraint*

The results generalize to cases in which consumers do not face "explicit" credit constraints. If risk averse consumers instead face a positive probability of receiving a zero-income shock (e.g., as in Carroll (1997) and Gourinchas and Parker (2002)), all results concerning the log-linearized Euler equation (4) still hold. This is basically because risk averse consumers will instead face a "self-imposed" no-borrowing constraint stemming

[23]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.

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)) enter the reduced form residual, $\tilde{\epsilon}$, the log-linearized Euler equation estimation will not be able to uncover the effect of children on the marginal utility of consumption. This result extends the critique in Carroll (2001) on the inability of log-linearized Euler equation estimation to uncover the IES.

Table 5 reports Monte Carlo results from pooling young and older households, simulated from a version of the life cycle model with a 0.3 percent risk of a zero-income shock (as in Gourinchas and Parker (2002)) *without* an explicit credit constraint. It is clear that the nonlinear GMM estimator can uncover the correct estimate while the log-linearized Euler equation cannot when children are perfectly foreseen (top panel). The results in the bottom panel, in which children arrive probabilistically, illustrates that the GMM estimator using both young and older households could not uncover the true effect of children on the marginal utility of consumption (unless it is zero) when using the actual change in the number of children, $\Delta\mathbf{z}_t$. This stems from the feature of the probabilistic model that households are identical prior to arrival of children given household composition, age, wealth, and income. Households in which household composition remains constant from period $t-1$ to t will increase consumption due to income growth and an *unwillingness* to borrow against this income growth. In turn, comparing consumption growth of households with and without children will underestimate the true effect of

Table 5. Monte Carlo results for no explicit constraint.

	$\theta = 0.0$		$\theta = 0.1$		$\theta = 0.5$		$\theta = 1.0$	
Instr.	LogLin	GMM	LogLin	GMM	LogLin	GMM	LogLin	GMM
				Deterministic arrival of children				
$\Delta\mathbf{z}_t$	0.015	−0.000	0.100	0.100	0.352	0.498	0.554	0.992
	(0.001)	(0.003)	(0.001)	(0.003)	(0.001)	(0.011)	(0.002)	(0.025)
$\Delta\bar{\mathbf{z}}_t$	0.118	−0.000	0.165	0.099	0.265	0.496	0.405	0.989
	(0.002)	(0.007)	(0.002)	(0.008)	(0.002)	(0.017)	(0.002)	(0.031)
				Probabilistic arrival of children				
$\Delta\mathbf{z}_t$	0.015	−0.000	0.086	0.084	0.283	0.424	0.447	0.852
	(0.001)	(0.003)	(0.001)	(0.003)	(0.001)	(0.011)	(0.001)	(0.023)
$\Delta\bar{\mathbf{z}}_t$	0.118	−0.000	0.166	0.100	0.272	0.501	0.497	1.000
	(0.002)	(0.007)	(0.002)	(0.008)	(0.002)	(0.017)	(0.002)	(0.028)

Note: The average of all MC estimates and standard deviations (in parentheses) across Monte Carlo runs are reported. All results are based on 1000 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 20 time observations long is drawn from this simulation. All individuals are initiated at age 22 with zero wealth, $A_{21} = 0$, permanent income of 1, $P_{22} = 1$, and no children. The results are based on a life cycle model in which there is no explicit constraint but instead a 0.3 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 2(b).

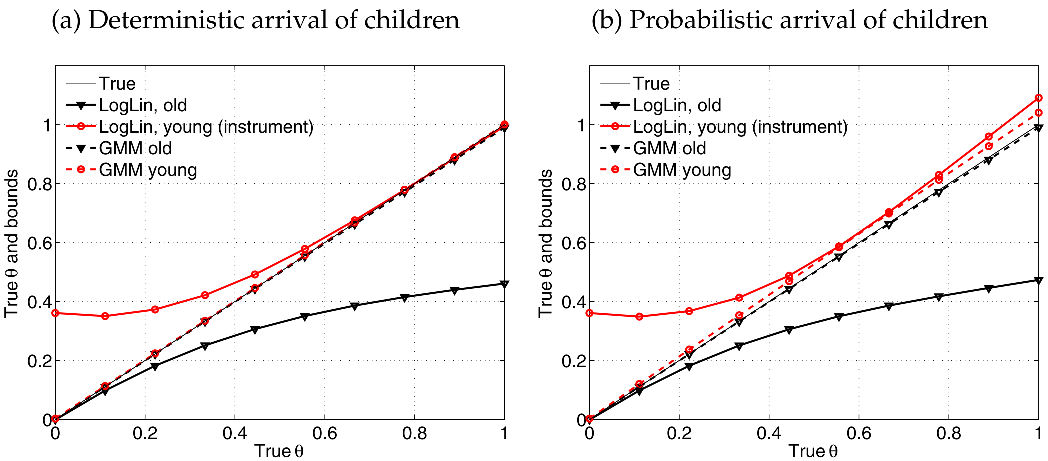

FIGURE 8. Bounds when there is no explicit constraint but self-imposed no-borrowing. The proposed bounds are based on a model in which there is unlimited borrowing but a positive probability of receiving zero income. Panel (a) illustrates the bounds from the baseline deterministic model in which children 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 black lines represent lower bounds and red/gray lines represent upper bounds using the (cohort-)average number of children as an instrument. Solid lines are based on the log-linearized Euler equation (4) and dashed lines are based on the nonlinear Euler equation estimated with the GMM estimator (3). Young households are defined as younger than 41.

children on the marginal utility of consumption if children arrive probabilistically.[24] Using only older households, in which children leave, or using the cohort-average number of children as an instrument will, however, lead the GMM estimator to produce unbiased estimates even if children arrive probabilistically.

Figure 8 illustrates the bounds based on the log-linearized and exact Euler equation. It is clear that while the bounds from the log-linearized Euler equation are still valid, they are, in general, different from the true value of θ. The bounds from the GMM estimator, using the exact Euler equation delivers bounds that are exactly (or very close to) the true value of θ. This is simply because when there is no explicit credit constraint, the exact Euler equation holds while the log-linearized Euler equation is a poor approximation due to income uncertainty. Recall, however, that the nonlinear GMM estimator in equation (3) is inconsistent if there is measurement error in consumption.

5.3 *Age effects of children*

For simplicity, I have followed the predominant specification in the existing literature throughout and assumed that the age of children does not affect the marginal utility of

[24]As pointed out by a referee, if there is sufficient heterogeneity in the arrival probability, the GMM estimator might be able to uncover the true effect of children on the marginal utility of consumption. Splitting the sample into two groups—one with the baseline arrival probabilities estimated from the PSID and another with a 50 percent lower arrival probability—did not change the results, however.

consumption. Browning and Ejrnæs (2009) find, however, that older children increase the marginal utility of consumption significantly more than younger children.

The proposed bounds generalize and can be used to estimate age effects of children on the marginal utility of consumption. Rather than letting $\mathbf{z}_t$ be the number of children, it could be the age of all children. For example, letting $\mathbf{z}_t$ contain age-group dummies and letting $\bar{\mathbf{z}}_t$ be the cohort average of these dummies could be used to uncover bounds on the effect of the number of children in different age groups.

A concern here could be that since the bounds are based on splitting the sample into younger and older households, the younger households will be less likely to have older children. For example, splitting the sample at age 45 will make the upper bound of children age 15 or older based on households that had children no later than age 30. Likewise, the lower bound of the effect of younger children aged, say, 0–2, will be identified of households having children after age 45. In turn, relatively large age groups would likely be required to empirically identify the bounds on the effect of children on the marginal utility of consumption.

Figure 9 illustrates that the proposed bounds using the log-linearized Euler equation can be used to uncover bounds on the effect of children in different age groups on the marginal utility of consumption. In particular, the results are based on a functional form of the taste shifter $v = \exp(\theta_y NumYoung + \theta_o NumOld)$, where *NumYoung* is the number of children below age 11 and *NumOld* is the number of children above age 10 (and below age 21). I impose $\theta_0 = 1.1 \cdot \theta_y$ and illustrate the estimated bounds for both parameters, θ_y (solid) and θ_o (dashed), in Figure 9.

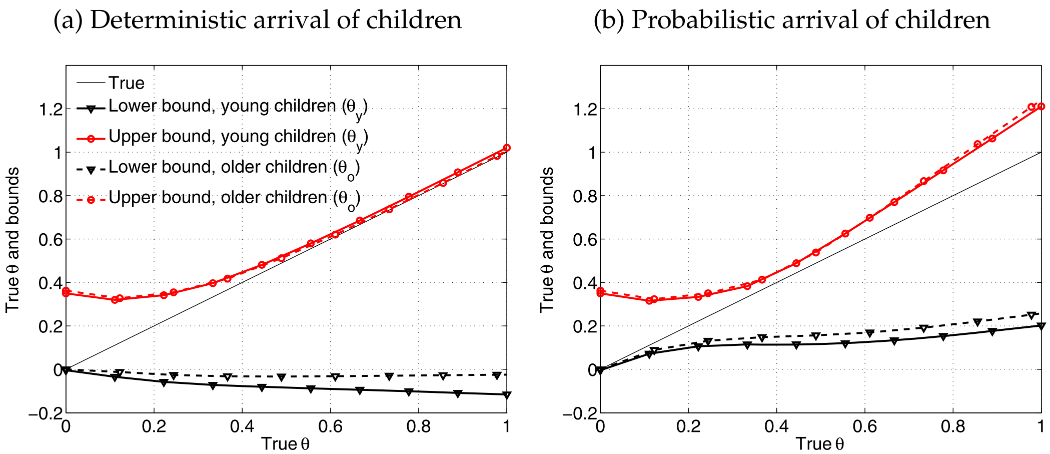

FIGURE 9. Bounds for child age effects. The proposed bounds when children more than 10 years old have a larger effect (10 percent) on the marginal utility of consumption than children younger than or 10 years old. Specifically, the Monte Carlo results are based on a functional form of the taste shifter $v = \exp(\theta_y NumYoung + \theta_o NumOld)$, where *NumYoung* is the number of children below age 11 and *NumOld* is the number of children above age 10 (and below age 21); $\theta_o = 1.1 \cdot \theta_y$.

5.4 *Labor market costs of children*

As in the rest of the literature on the effect of children on the marginal utility of consumption, income is assumed to be *independent* of household composition. If income depends on household composition, the results will change depending on the ways in which children affect the labor market income of household members. Children might, however, impose career costs and affect labor market outcomes. 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 (forthcoming) 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.

The bounds from the four-period model are still valid if children reduce permanent income, as suggested by the results above. This is true as long as children do not reduce permanent income by more than the permanent income growth as illustrated in Figure 10(a). 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 10(b).

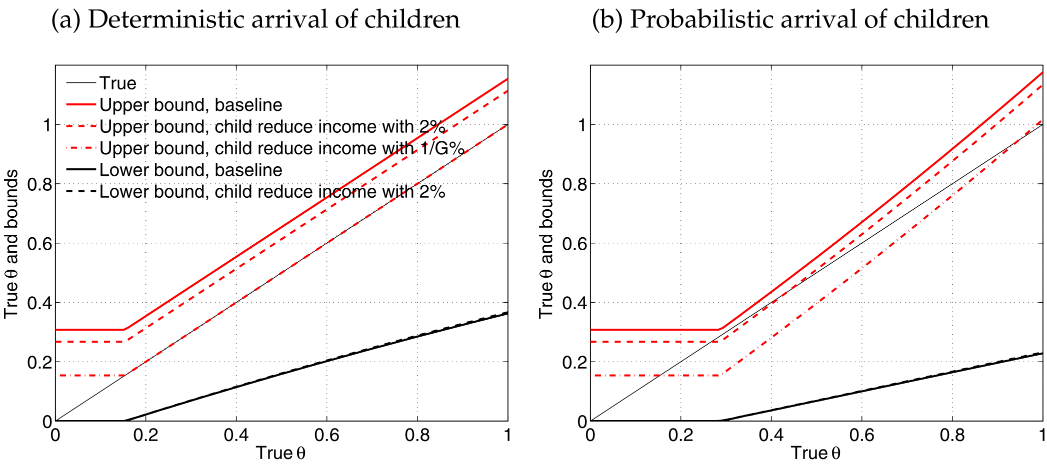

FIGURE 10. Bounds if children reduce labor market income. The proposed bounds are from the four-period model when children reduce labor market income by 2 percent. Panel (a) illustrates the bounds from the model in which children arrive deterministically; the probabilistic version is illustrated in panel (b). The 45° line represents the true value of θ, while black lines represent lower bounds and red/gray lines represent upper bounds using the (cohort-)average number of children as an instrument. Solid lines are based on the baseline case where children do not affect labor market income and dashed lines represent the extreme case where children reduce labor market income by 2 percent.

5.5 *Measurement error in consumption*

Runkle (1991) estimates that more than 70 percent of the variation in log consumption growth in the PSID is due to measurement error. An attractive feature of log-linearized Euler equation estimation is that measurement error can be handled using a standard IV approach. The GMM estimator in equation (3) is inconsistent in this case.

Imagine that consumption is observed with multiplicative measurement error, $C_t = C_t^\star \xi_t$, where $C_t^\star$ is the *true* consumption level and ξ_t is measurement error,

$$\xi_t = \nu_t(1 + \mu \mathbf{z}_t),$$

where measurement error is allowed to be correlated with the number of children through the parameter μ and ν the is log-normal independent and identically distributed (i.i.d.) measurement error component with mean and variance 1. The log-linearized Euler equation will have moving average-(1) (MA(1)) type errors in this case. To see this, write the constrained Euler equation in terms of the *observed* consumption,

$$R\beta\left(\frac{C_{t+1}}{C_t}\right)^{-\rho} \exp(\theta \Delta z_{t+1}) = \left(\frac{\xi_{t+1}}{\xi_t}\right)^{-\rho} \tilde{\eta}_{t+1},$$

with $\mathbb{E}_t[\tilde{\eta}_{t+1}] = 1 - \frac{\lambda_t}{(C_t/\xi_t)^{-\rho}\exp(\theta \mathbf{z}_t)}$. The log-linearized Euler equation then becomes

$$\Delta \log C_{t+1} = \text{constant} + \rho^{-1}\theta \Delta \mathbf{z}_{t+1} + \Delta \log(1 + \mu \mathbf{z}_{t+1}) + \Delta \log v_{t+1} + \eta_{t+1},$$

where $\eta_{t+1} \equiv -\rho^{-1} \log \tilde{\eta}_{t+1}$ and $\Delta \log v_{t+1}$ is the MA(1) error component.

If measurement error is independent of the number of children, $\mu = 0$, the MA(1) error will not affect the estimation of the bounds of θ, as illustrated in the left panel of Figure 11. The right panel of Figure 11 shows the bounds when households with children provide noisier measures of consumption ($\mu = 0.1$). The lower bound is now above the true value of θ for low values of θ because the omitted variable, $\Delta \log(1 + \mu \mathbf{z}_{t+1})$, is positively correlated with changes in the number of children.[25] The upper bound is robust to measurement error in consumption as long as measurement error is not *negatively* correlated with household composition.[26]

6. Concluding discussion

Many studies estimate the effect of children on the marginal utility of consumption by applying estimators derived from the consumption Euler equation. Especially the log-linearized Euler equation is popular since it yields estimable equations linear in parameters that can handle measurement error and be estimated with repeated cross section data by the construction of synthetic cohort panels. Although these estimators have

[25]The bias in the bounds can be approximated by noting that for small μ, $\log(1 + \mu \mathbf{z}_{t+1}) \approx \mu \mathbf{z}_{t+1}$ and the estimated θ is approximately equal to $(\rho^{-1}\theta + \mu)\rho$. For example, if $\theta = 0$ (while $\rho = 2$ and $\mu = 0.1$), the lower bound is approximately $0.1 \cdot 2 = 0.2$, close to the value of 0.19 reported in Figure 11(b).

[26]Similarly, the lower bound would be valid while the upper bounds would potentially not be valid if the correlation is negative ($\mu < 0$).

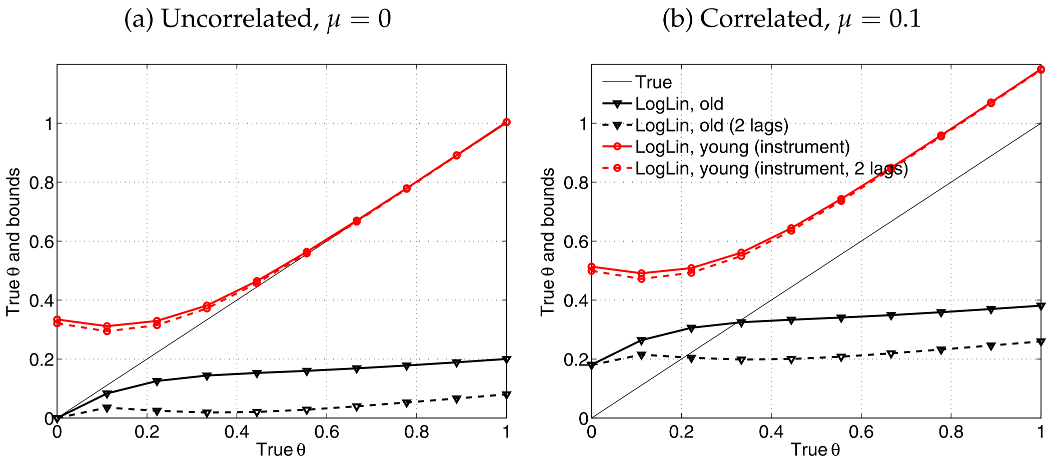

FIGURE 11. Bounds of the multiplicative log-normal measurement error in consumption. The proposed bounds are when observed consumption is measured with multiplicative log-normal measurement error. The 45° line represents the true value of θ while black lines represent lower bounds and red/gray lines represent upper bounds using the (cohort-)average number of children as an instrument. Solid lines are based on the proposed instruments while dashed lines illustrates bounds when using the cohort-average number of children lagged twice as an instrument. Young households are defined as younger than 41.

now become workhorses in the analysis of intertemporal consumption behavior, little is known about their performance when households face potentially binding credit constraints and, thus, the standard Euler equation no longer holds.

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

While Euler equation estimation cannot, in general, point identify the effect of children on the marginal utility of consumption, the proposed bounds are rather robust. Even in cases where Euler equation estimation techniques might work (if, for example, the estimation sample consists of sufficiently wealthy households), the bounds are also applicable. Researchers have limited information on whether that special case is the relevant one which should lead to the use of more robust methods. The proposed bounds are an attempt to provide researchers with a tractable, yet more robust, alternative.

Estimating the proposed bounds on PSID data shows that the point estimates from the conventional estimators are above the upper bound. Likewise, *all*, to the best of my knowledge, existing estimates of the effect of children on the marginal utility of 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. In particular, 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 the marginal utility of consumption.

Appendix: Euler equation estimation results from the PSID

I have implemented two Euler equation estimators for the effect of children on the marginal utility of consumption that both ignore credit constraints. In particular, I follow the approach in Attanasio et al. (1999) and estimate a log-linearized Euler equation and implement a nonlinear (exact) GMM estimator allowing for measurement error in consumption, as proposed in Alan, Attanasio, and Browning (2009). The aim is to implement these estimators using the current PSID sample while following their approach as closely as possible.

The implemented estimators both aim to identify the effect of children on the marginal utility of consumption, θ, along with the relative risk aversion parameter, ρ.[27] To identify the latter, I follow Alan, Attanasio, and Browning (2009) and use the (annual average) 3-month Treasury bill rate net of the (annual average) inflation based on the U.S. consumer price index to calculate the gross real interest rate, $R_t = (1 + r_t)$.[28]

A.1 *Log-linearized Euler equation estimates*

I estimate a log-linearized Euler equation following the approach in Attanasio et al. (1999) as closely as possible. That study used the repeated cross sections in the CEX to construct synthetic cohort panels based on 5-year birth cohort bands. Although the PSID contains longitudinal information on household level consumption, I collapse the panel into synthetic cohort panels because repeated cross sections are often used and this approach is common in the literature. The estimation equation is

$$\Delta \log C_{it} = \text{constant} + \rho^{-1} \log(R_t) + \rho^{-1} \theta_k \Delta\#Kids_{it} + \rho^{-1} \theta_a \Delta\#Adults_{it} + v_{it},$$

where i refers to a given cohort.[29] I follow Attanasio et al. (1999) and use as instruments for demographic variables (i) a polynomial in age, (ii) second to fourth lags of consumption, income growth and interest rates, and (iii) second to third lags of growth in the number of children and adults.

Table A.1 reports the estimation results using the PSID sample. The estimated correlations between log consumption growth and changes in the number of children and adults are close to those reported in Attanasio et al. (1999), reprinted in the rightmost column in Table A.1.

[27] The estimators in Alan, Attanasio, and Browning (2009) also estimates the discount factor, β.

[28] In turn, the ρ is identified for *time variation* only.

[29] I use 1-year bands to increase the number of observations in the PSID sample.

TABLE A.1. Log-linear Euler equation estimates.

	Low Skilled				High Skilled				
	Food		Nondurable		Food		Nondurable		
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	ABMW†
$\Delta\#Kids$	0.289	0.258	0.131	0.135	0.169	0.163	0.320	0.338	0.212
	(0.088)	(0.088)	(0.114)	(0.111)	(0.109)	(0.101)	(0.226)	(0.224)	(0.101)
$\Delta\#Adults$	0.437	0.479	0.288	0.279	0.461	0.453	0.487	0.466	0.449
	(0.203)	(0.193)	(0.228)	(0.236)	(0.189)	(0.193)	(0.263)	(0.259)	(0.144)
$\log(R_t)$		2.078		−28.604		0.900		−18.304	0.640
		(1.462)		(15.591)		(2.098)		(30.150)	(0.333)
Constant	0.033	−0.063	0.026	1.148	0.012	−0.029	0.025	0.744	0.045
	(0.009)	(0.067)	(0.011)	(0.613)	(0.015)	(0.094)	(0.024)	(1.179)	(0.009)
Obs.	152	152	76	76	160	160	80	80	256

Note: Reported are estimates based on synthetic cohort panels constructed from 1-year birth cohort bands. Estimates are based on 2SLS,with the instrument set containing (i) a polynomial in age, (ii) second to fourth lags of consumption, income growth and interest rates, and (iii) second to third lags of growth in the number of children and adults. Robust standard errors are given in parentheses. Households are classified as high skilled if the male head has ever enrolled in college, including college dropouts. Nondurable consumption in the PSID is imputed by Blundell, Pistaferri, and Preston (2008).

† Reprinted estimates from Attanasio et al. (1999). They used quartile CEX data and, thus, also included seasonal dummies. For readability these estimates are not reprinted here.

The risk aversion, ρ, is estimated implausibly low and even negative when using imputed nondurable consumption. The identification of ρ using Euler equation estimation techniques is, like demographic effects, threatened by credit constraints and income uncertainty (Carroll (2001), Adda and Cooper (2003)). The risk aversion is also very imprecisely estimated from the time variation in the interest rate. In turn, I cannot reject $\rho = 1.57$, as found in Attanasio et al. (1999). Excluding the interest rate does not significantly affect the estimated demographic effects on the marginal utility of consumption.

Using $\rho = 1.57$, the estimates in columns (2) in Table A.1 suggest that the first child increases the marginal utility of *food* consumption with around 50 percent for low skilled and around 30 percent for high skilled. Using nondurable consumption the effect reduces to 24 percent for low skilled and increases to 70 percent for high skilled. In general, these results are around and not significantly different from the 39 percent estimated in Attanasio et al. (1999).

A.2 *Exact Euler equation estimates*

Alan, Attanasio, and Browning (2009) show that if measures of consumption are contaminated with *multiplicative log-normal measurement error* with variance ν (and some arbitrary mean), the two nonlinear equations

$$u^1_{it+1} = \left(\frac{C_{it+1}}{C_{it}}\right)^{-\rho} R_{t+1}\beta \exp(\theta\Delta\#Kids_{it+1}) - \exp(\rho^2\nu),$$

$$u^2_{it+1} = \left(\frac{C_{it+2}}{C_{it}}\right)^{-\rho} R_{t+1}R_{t+2}\beta^2 \exp\bigl(\theta[\#Kids_{it+2} - \#Kids_{it}]\bigr) - \exp(\rho^2\nu)$$

Table A.2. Nonlinear Euler equation estimates.

	Food				Nondurables				
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)	AAB†
θ	0.830	0.674	0.336	0.368	0.829	0.459	0.019	0.407	0.894
	(0.051)	(0.162)	(0.184)	(0.291)	(0.047)	(0.201)	(0.117)	(0.644)	(0.51)
ν	0.085	0.071	0.069	−0.003	0.085	0.054	0.022	0.005	0.085
	–	(0.011)	(0.109)	(0.256)	–	(0.010)	(0.019)	(0.344)	(0.05)
ρ	1.45	1.45	0.272	0.362	1.45	1.45	0.320	−0.380	1.45
	–	–	(0.235)	(0.751)	–	–	(0.037)	(0.613)	(0.27)
β	0.99	0.99	0.99	0.980	0.99	0.99	0.99	0.928	0.99
	–	–	–	(0.016)	–	–	–	(0.014)	(0.10)
Obs.	10,957	10,957	8267	8267	8414	8414	7052	7052	19,317
Obj.	1.1e−04	8.8e−14	7.7e−05	4.9e−05	8.0e−04	1.2e−12	3.2e−04	1.5e−05	N.A.

Note: Reported are estimates based on a nonlinear estimator (GMM-LN) proposed in Alan, Attanasio, and Browning (2009). Robust standard errors are given in parentheses. Households are classified as high skilled if the male head has ever enrolled in college, including college dropouts. Nondurable consumption in the PSID is imputed by Blundell, Pistaferri, and Preston (2008).
Columns (1) present estimation results when fixing $\nu = 0.085$, $\rho = 1.45$, and $\beta = 0.99$ to the values found in Alan, Attanasio, and Browning (2009), reported in the rightmost column. Columns (2) present estimation results when fixing ρ and β, columns (3) present estimation results when fixing only β, and columns (4) estimate all parameters.
When estimating ρ (and β) in columns (3) and (4), the lagged interest rate is used as an instrument. Therefore, fewer observations are available for estimation in this case.
† Reprinted GMM-LN estimates from Alan, Attanasio, and Browning (2009).

can be utilized to estimate θ, β, ρ, and ν. The authors term this the GMM-LN estimator and it is the preferred estimator in their empirical application. When estimating ρ, I use the lagged interest rate along with a constant as instruments, as suggested by Alan, Attanasio, and Browning (2009). I only use the number of children here because Alan, Attanasio, and Browning (2009) do not distinguish between the ages of children.[30]

Table A.2 reports the estimation results using the PSID sample. I had substantial convergence problems and I, thus, pooled both educational groups. I also restricted the analysis to observations for which neither C_{it+1}/C_{it} nor C_{it+2}/C_{it} was less than 0.1. Dropping observations for which consumption fell to less than 10 percent of the previous consumption level improved the estimator's stability because these "extreme" observations would otherwise dominate the estimation process through $(\cdot)^{-\rho}$. The estimator turned out to be extremely sensitive to initial values and I report the estimates with the lowest objective function out of a sequence of different starting values.[31]

Columns (1) present estimation results when fixing $\nu = 0.085$, $\rho = 1.45$, and $\beta = 0.99$ to the values found in Alan, Attanasio, and Browning (2009), reported in the rightmost column of Table A.2. Columns (2) present estimation results when fixing ρ and β, columns (3) present estimation results when fixing β, and columns (4) estimate all parameters.

[30] In fact, Alan, Attanasio, and Browning (2009) use changes in family size but restrict attention to stable couples such that changing family size is extremely closely related to changing number of children.

[31] Starting values was generated as $(\beta, \rho, \theta, \nu)_0 = (1, 2, 1, 0.1) \cdot \delta_i$, where $\delta_i \in [0.01, 1]$, $i = 1, \dots, 30$.

I estimate a similar effect of children on the marginal value of consumption, θ, as reported in Alan, Attanasio, and Browning (2009). They estimate $\theta \approx 0.895$. When only estimating θ and keeping other parameters fixed at the values reported in Alan, Attanasio, and Browning (2009), I find $\theta \approx 0.83$, which drops to around 0.67 for food consumption and 0.46 for nondurable consumption when also estimating the measurement error variance, ν. In columns (4), when estimating all parameters simultaneously, I estimate low risk aversion parameters (as I did using the log-linear specification above) and low discount factors compared to Alan, Attanasio, and Browning (2009). When only fixing β and estimating ρ in columns (3), the effect of children is estimated significantly lower. This is because, as the results in columns (4) illustrate, the discount factor consistent with the current data is somewhat lower than the 0.99 imposed in columns (3). To compensate for a too large discount factor together with an implausible low risk aversion parameter, θ is estimated to be around 0.34 and 0.02 for food and nondurable consumption, respectively.

The preferred estimation results are those in columns (2). In this setting, both β and ρ are calibrated using the values estimated in Alan, Attanasio, and Browning (2009) and yield similar estimates of the effect of children and the measurement error variance reported in that study. When using interest rate variation to estimate β and ρ in columns (3) and (4), the resulting estimates seem implausible. The measurement error *variance* and risk aversion parameters are both estimated to be negative and the discount factor is estimated to be lower than reported in most studies. Note that if households face potentially binding credit constraints, the measurement error variance parameter, ν, will likely be affected by the shadow price of resources, ν. The interpretation of ν as a variance parameter is thus invalidated by potentially binding credit constraints, giving a rationale for a negative estimated ν.

References

Adda, J. and R. Cooper (2003), *Dynamic Economics: Quantitative Methods and Applications*. MIT Press, Cambridge. [936, 963]

Adda, J., C. Dustmann, and K. Stevens (forthcoming), "The career costs of children." *Journal of Political Economy*. [959]

Alan, S., K. Atalay, and T. F. Crossley (2012), "Euler equation estimation on micro data." Working Paper 1221, Koc University-TUSIAD Economic Research Forum. [937]

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. [939, 952, 953, 962, 963, 964, 965]

Alan, S. and M. Browning (2010), "Estimating intertemporal allocation parameters using synthetic residual estimation." *The Review of Economic Studies*, 77 (4), 1231–1261. [940]

Almlund, M. (2013), *Essays on Credit Constraints, Education, and the Family*. Ph.D. thesis, Economics Department, University of Chicago. [955]

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. [935, 936, 937, 938, 940, 944, 952, 953, 962, 963]

Attanasio, O. P. and M. Browning (1995), "Consumption over the life cycle and over the business cycle." *American Economic Review*, 85 (5), 1118–1137. [935, 936]

Attanasio, O. P. and H. Low (2004), "Estimating Euler equations." *Review of Economic Dynamics*, 7 (2), 406–435. [937]

Attanasio, O. P. and G. Weber (1993), "Consumption growth, the interest rate and aggregation." *The Review of Economic Studies*, 60 (3), 631–649. [952, 953]

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. [935]

Banks, J., C. O'Dea, and Z. Oldfield (2010), "Cognitive function, numeracy and retirement saving trajectories." *The Economic Journal*, 120 (548), 381–410. [946]

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. [936, 938]

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. [935]

Blundell, R., M. C. Dias, C. Meghir, and J. M. Shaw (forthcoming), "Female labour supply, human capital and welfare reform." *Econometrica.* [945, 954]

Blundell, R. and A. Lewbel (1991), "The information content of equivalence scales." *Journal of Econometrics*, 50, 49–68. [936]

Blundell, R., L. Pistaferri, and I. Preston (2008), "Consumption inequality and partial insurance." *American Economic Review*, 1887–1921. [946, 950, 951, 952, 963, 964]

Browning, M., A. Deaton, and M. Irish (1985), "A profitable approach to labor supply and commodity demands over the life-cycle." *Econometrica*, 53 (3), 503–544. [935, 939]

Browning, M. and M. Ejrnæs (2009), "Consumption and children." *The Review of Economics and Statistics*, 91 (1), 93–111. [935, 936, 940, 945, 958]

Cagetti, M. (2003), "Wealth accumulation over the life cycle and precautionary savings." *Journal of Business & Economic Statistics*, 21 (3), 339–353. [937, 944, 946, 947]

Calhoun, C. A. and T. J. Espenshade (1988), "Childbearing and wives' foregone earnings." *Population Studies*, 42 (1), 5–37. [959]

Carroll, C. D. (1992), "The buffer-stock theory of saving: Some macroeconomic evidence." *Brookings Papers on Economic Activity*, 2, 61–135. [956]

Carroll, C. D. (1997), "Buffer-stock saving and the life cycle/permanent income hypothesis." *The Quarterly Journal of Economics*, 112 (1), 1–55. [937, 955]

Carroll, C. D. (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. [937, 938, 956, 963]

Carroll, C. D. (2006), "The method of endogenous gridpoints for solving dynamic stochastic optimization problems." *Economics Letters*, 91 (3), 312–320. [945]

Citro, C. F. and R. T. Michael (1995), *Measuring Poverty: A New Approach*. National Academy Press, Washington. [937]

Crossley, T. F. and H. Low (2014), "Job loss, credit constraints and consumption growth." *Review of Economics and Statistics*, 96 (5), 876–884. [937]

Dynan, K. E. (2000), "Habit formation in consumer preferences: Evidence from panel data." *American Economic Review*, 90 (3), 391–406. [940, 952, 953]

Fernández-Villaverde, J. and D. Krueger (2007), "Consumption over the life cycle: Facts from consumer expenditure survey data." *The Review of Economics and Statistics*, 89 (3), 552–565. [935, 938]

Ferreira, M. L., R. C. Buse, and J.-P. Chavas (1998), "Is there bias in computing household equivalence scales?" *Review of Income and Wealth*, 44 (2), 183–198. [936]

Gourinchas, P.-O. and J. A. Parker (2002), "Consumption over the life cycle." *Econometrica*, 70 (1), 47–89. [937, 944, 946, 955, 956]

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), 149–185. [937]

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. [937]

Hall, R. E. and F. S. Mishkin (1982), "The sensitivity of consumption to transitory income." *Econometrica*, 50 (2), 461–481. [937, 940]

Heathcote, J., F. Perri, and G. L. Violante (2010), "Unequal we stand: An empirical analysis of economic inequality in the United States, 1967–2006." *Review of Economic Dynamics*, 13 (1), 15–51. Special issue: *Cross-Sectional Facts for Macroeconomists*. [937]

Irvine, I. (1978), "Pitfalls in the estimation of optimal lifetime consumption patterns." *Oxford Economic Papers*, 30 (2), 301–309. [935]

Jappelli, T. (1990), "Who is credit constrained in the U.S. economy?" *The Quarterly Journal of Economics*, 105 (1), 219–234. [937]

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. [937]

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. [937]

Jørgensen, T. H. (2015), "Life-cycle consumption and children: Evidence from a structural estimation." Discussion Paper 2015-08, Department of Economics, University of Copenhagen. [946]

Lawrance, E. C. (1991), "Poverty and the rate of time preference: Evidence from panel data." *Journal of Political Economy*, 99 (1), 54–77. [940]

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. [937]

Lewbel, A. (1989), "Identification and estimation of equivalence scales under weak separability." *Review of Economic Studies*, 56, 311–316. [936]

Lewbel, A. (1997), "Consumer demand systems and household equivalence scales." In *Handbook of Applied Econometrics* (M. H. Pesaran and P. Schmidt, eds.), Blackwell Publishers Ltd., London. [936]

Ludvigson, S. and C. H. Paxson (2001), "Approximation bias in linearized Euler equations." *Review of Economics and Statistics*, 83 (2), 242–256. [937]

Moon, H. R. and F. Schorfheide (2009), "Estimation with overidentifying inequality moment conditions." *Journal of Econometrics*, 153 (2), 136–154. [954]

Pollak, R. A. and T. J. Wales (1979), "Welfare comparisons and equivalence scales." *American Economic Review*, 69, 216–221. [936]

Runkle, D. E. (1991), "Liquidity constraints and the permanent-income hypothesis: Evidence from panel data." *Journal of Monetary Economics*, 27 (1), 73–98. [937, 940, 960]

Schechtman, J. (1976), "An income fluctuation problem." *Journal of Economic Theory*, 12, 218–241. [938, 956]

Scholz, J. K., A. Seshadri, and S. Khitatrakun (2006), "Are americans saving 'optimally' for retirement?" *Journal of Political Economy*, 114 (4), 607–643. [937]

The Danish Ministry of Finance (2003), *ældres sociale vilkår*. (In Danish). [946]

Thurow, L. C. (1969), "The optimum lifetime distribution of consumption expenditures." *American Economic Review*, 59 (3), 324–330. [937]

Zeldes, S. P. (1989a), "Consumption and liquidity constraints: An empirical investigation." *Journal of Political Economy*, 97 (2), 305–346. [937, 938]

Zeldes, S. P. (1989b), "Optimal consumption with stochastic income: Deviations from certainty equivalence." *Quarterly Journal of Economics*, 104 (2), 275–298. [938, 956]

Co-editor Petra E. Todd handled this manuscript.

Submitted September, 2014. Final version accepted December, 2015.