Life-Cycle Consumption: Can Single Agent Models Get it Right?

Alexander Bick
Goethe University Frankfurt

Sekyu Choi
Universitat Autònoma de Barcelona*

February 15, 2011

[PRELIMINARY]

Abstract

In the quantitative macro literature, single agent models are heavily used to explain “adult equivalent” household data, a common example being household consumption deflated by some form of equivalence scale. In this paper, we study differences between the consumption of single agent models and the “adult equivalent” household consumption from a model where household size is taken into account. Using a theoretical model we prove that, under mild conditions, these predictions are different. Through a quantitative exercise, we then document that these differences can be substantial with a large variation over the life-cycle.

Keywords: Consumption, Life-Cycle Models, Households

JEL classification: D12, D91, E21, J10

*email: bick@wiwi.uni-frankfurt.de and sekyu.choi@uab.cat. We thank Dirk Krueger, Nicola Fuchs-Schündeln, Nezih Guner, Juan Carlos Conesa and many others for very helpful suggestions and comments. All errors are ours.

1 Introduction

To understand life-cycle consumption, economists seem to have converged to single agent models, given their tractability and ease of use. In the quantitative macroeconomic literature, a standard approach entails extracting “per-adult equivalent” (or equivalized) consumption facts from household survey data and use it as a target to be replicated by single agent models, which are calibrated using also equivalized household income. Some papers in this vein include Krueger and Perri (2006), Blundell et al. (2008), Kaplan and Violante (2009) and Guvenen and Smith (2010).¹

However, this approach faces the inherent challenge that consumption decisions might depend on household size and composition through non trivial channels. In this paper we ask about the differences between the predictions of a single agent model and the equivalized predictions of a model with household size and composition effects. Specifically, we are interested in predictions from the standard incomplete markets model for mean and cross sectional inequality of life-cycle consumption.

We perform our analysis by extending this framework to allow for deterministic changes in household size and composition during the life-cycle. Moreover, we let these changes affect optimal decisions on consumption and savings in a unitary model approach similar to Attanasio et al. (1999) and Gourinchas and Parker (2002). While these studies use a general “taste shifter” depending on household size and composition, we propose an explicit formulation combining (i) equivalence scales, which reflect economies of scale in household consumption and (ii) utility weights, which make explicit the value of having a household of different sizes and compositions. This setup accommodates both the case of single households (the *Single* model) and the case where household size varies during the life-cycle (the *Demographics* model). Obviously, we can have as many *Demographics* models as combinations between equivalence scales and utility weights.

Using this framework, we study differences between the predictions of the *Single* model, calibrated using equivalized household income, and the equivalized household consumption from the *Demographics* model (which uses total household income). Given that the *Single* model is just a particular case of what we label the *Demographics* model, we think of this framework as the mildest test to the single household approach.

¹Other papers present mixed empirical strategies. For example, Storesletten et al. (2004) use household income *per person* from the PSID, while trying to match the cross-sectional variance in *total* household consumption (without controlling for household size/composition). Another approach is in Fernández-Villaverde and Krueger (2010), who study durable consumption using an equilibrium life-cycle model. They use worker information to parameterize income profiles and contrast the results from their model with equivalized consumption from the CEX.

Throughout our analysis, we rely on equivalence scales to transform household data into “adult equivalent” data. These scales are functions of household size and composition and typically used to deflate total household information (like consumption and income) by a number less than the actual household size, in order to allow for economies of scale inside the household. This approach has gained importance in the macro literature: in the 2010 special issue of the Review of Economic Dynamics, equivalence scales are used to obtain consistent “individual” cross sectional facts for a wide range of countries.² Another example is Fernández-Villaverde and Krueger (2007), who discuss the properties of different types of equivalence scales and their effect on life-cycle consumption profiles.³

Our quantitative results show that the differences between equivalized consumption data from the *Demographics* model and individual data from the *Single* model can be substantial, both in terms of predicted levels and cross sectional inequality. The differences are increasing in the difference between utility weights (the extra utility households derive from consuming in a household with additional members) and the value of the equivalence scales. Intuitively, the larger this difference, the larger the motives for households in the *Demographics* model to allocate consumption in periods where household size is larger: the utility of doing so gains importance relative to the penalization in terms of the cost of sharing. Hence, the differences in predictions with respect to the *Single* model are going to be high since there equalization has only an effect through the calibrated individual income profile. Furthermore, differences in these predictions depend on the stage of the life cycle: when preferences for household size are high (low), equivalized consumption from *Demographics* model can be up to 15% (5%) higher (lower) than in the *Single* model when household size peaks (around age 35).

Since we are the first to consider (to the best of our knowledge) the notion of utility weights in the *Demographics* model, there is no direct evidence of their value in the literature so we perform our analysis considering a variety of cases. Nevertheless, using estimates from Attanasio et al. (1999) we can compute the value of the implied utility weights, under specific equivalence scales. We find that the empirical evidence in Attanasio et al. (1999) falls in our considered ranges for the utility weights and suggests that in most cases utility weights are larger than equivalence scales.

²See Krueger et al. (2010) for a general description and Heathcote et al. (2010) for the US economy.

³Other papers, like Gourinchas and Parker (2002), Blundell et al. (2008), Aguiar and Hurst (2009) and Guvenen and Smith (2010), obtain equalization of household data by introducing demographic controls (e.g., household size variables) in linear regressions where the dependent variable is the statistic of interest at the household level. This procedure resembles an equivalence scale, since the implied household effect from the regression approach can be translated trivially into an ad hoc equivalence scale.

The structure of the paper is as follows. In Section 2 we discuss our proposed preferences for the household and Section 3 we present theoretical predictions in a stylized two period framework. Sections 2 and 5 discuss the model we use to quantify these theoretical predictions. In Section 6 we discuss the quantitative features of the model and the calibration strategy while Section 7 shows our main quantitative results. In the last section we present our conclusions and extensions of our exercise.

2 Demographics and Preferences

We follow Attanasio et al. (1999) and Gourinchas and Parker (2002) and introduce exogenous and deterministic demographics via a taste shifter into the standard life-cycle model:

$$U(c, N_{ad}, N_{ch}) = \exp(\xi_1 N_{ad} + \xi_2 N_{ch}) u(c), \quad (1)$$

with N_{ad} being the number of adults and N_{ch} the number children in the household. In the subsequent analysis we will use an explicit formulation of the taste shifter:

$$u(c, N_{ad}, N_{ch}) = \delta(N_{ad}, N_{ch}) u\left(\frac{c}{\phi(N_{ad}, N_{ch})}\right). \quad (2)$$

Household utility $U(c, N_{ad}, N_{ch})$ is the product of the utility from adult equivalized consumption $u\left(\frac{c}{\phi(N_{ad}, N_{ch})}\right)$, where $\phi(N_{ad,t}, N_{ch,t})$ transforms total household consumption into an adult equivalent consumption, and the utility weight $\delta(N_{ad}, N_{ch})$ which differs across households of different size and composition. We are not the first to use such an explicit formulation of the taste shifter. Fuchs-Schündeln (2008) sets δ equal to ϕ which are then both a function of the average household size and an equivalence scale. In contrast, Hong and Ríos-Rull (2007) set $\delta = 1$ to one and estimate the equivalence scale ϕ for different household size. Both papers do not provide any further justification for their choice.

In the theoretical part of this paper, we will not rely on any specific functional form of ϕ . In general we might think of it as prototypical equivalence scale which are calculated with information on expenditure shares of households. As a concrete example consider the OECD equivalence scale:

$$\phi_{OECD} = 1 + 0.7(N_{ad} - 1) + 0.5N_{ch}. \quad (3)$$

According to Equation (3) it takes 1.7 \$ to generate the same level of welfare out of consumption for a two adult household that 1 \$ achieves for a single member household. The three mechanisms through which household

Table 1: Properties

	Equivalence Scale ϕ	Utility Weighting Factor δ
Level		
$N_{ad} = 1, N_{ch} = 0 :$	$\phi = 1$	$\delta = 1$
$N_{ad} > 1, N_{ch} \geq 0 :$	$\phi \in (1, N_{ad} + N_{ch})$	$\delta \in [1, N_{ad} + N_{ch}]$
Derivatives		
$\forall i = ad, ch$	$\frac{\partial \phi}{\partial N_i} \in (0, 1)$	$\frac{\partial \delta}{\partial N_i} \in [0, 1]$

size affects the intra-temporal rate of transformation between expenditures and services are family/public goods, economies of scale, and complementarities, see e.g. Lazear and Michael (1980). Note that while in Equation (3) each additional adult and child increases the equivalence scale by a constant number, these weights could also vary with household size and composition. Alternatively, ϕ could also be obtained from the coefficients on household size and composition dummies from a regression of consumption on various controls, as discussed in the introduction.

To also give a hypothetical example to the utility weight $\delta(N_{ad,t}, N_{ch,t})$, which assigns a different value to the household utility of equivalized consumption across households of different size and composition, we could be specified in analogy to the OECD equivalence scale (Equation refequivalence scale oecd):

$$\delta(N_{ad}, N_{ch}) = 1 + \omega_{ad}(N_{ad} - 1) + \omega_{ch}N_{ch}. \quad (4)$$

The first adult has a utility weight of one whereas any further adult and child have the utility weights ω_{ad} and ω_{ch} . As with the equivalence scales the weights ω may vary with household size and composition.

Fernández-Villaverde and Krueger (2007) provide a list of six representative estimated household equivalence scales with properties being summarized in column one of Table 1. To the best of our knowledge, there is no empirical evidence on the utility weighting factor. As already mentioned Fuchs-Schündeln (2008) sets $\delta = \phi$ whereas Hong and Ríos-Rull (2007) set $\delta = 1$ for all household types. Using the same utility function as Attanasio et al. (1999) and their empirical estimates, δ can be backed out for a given equivalence scale ϕ . Given the analogy to equivalence scales, we believe that the utility weighting should have similar properties. However, as e.g. assumed in Hong and Ríos-Rull (2007) the utility weighting factor may remain constant

as the household increases, i.e. $\delta = 1 \forall N_{ad}, N_{ch}$, whereas equivalence scales always increase. On the other hand an additional household member could also increase the utility weighting factor by one, i.e. $\delta = N_{ad} + N_{ch}$, whereas equivalence scales always increase by less than one. These properties are outlined in column 2 of Table 1. Note that our subsequent analysis does not rely on the utility weight δ and the equivalence scale ϕ being only a function of the number of adults and children in a household. We could make them easily dependent on finer categories, as e.g. the number of children by different age groups. We just stick to the previous case as this approach has been chosen by Attanasio et al. (1999) and Gourinchas and Parker (2002) to introduce demographics into the life-cycle model. In addition, those who clean the empirical income and consumption series for household size and composition using equivalence scales rather than the regression approach usually use equivalence scales that only depend on the number of adults and number of children in the household.

3 A Two Period Model

We take the most simplistic approach to model household size following Attanasio et al. (1999), but allowing for a more structural formulation of the utility function as given by Equation (2). In particular, we assume that the true data is generated by such a model where demographics play a role. We will compare the equivalized consumption choices from this model with those implied by a model that abstracts from demographics but takes as input income data that is cleaned of family size effects through the application of an equivalence scale.

3.1 Basic Setup

We will analyze this question using a two-period model with deterministic demographics. We consider a household who lives for two periods and changes its size over time. In particular, we assume that the household size is one in the first period, e.g. a young person living alone, and larger one in the second period, e.g. a child is born. These assumptions and their implications for the utility weighting factor δ and an arbitrary equivalence scale ϕ , both satisfying the properties outlined in Table 1, can be summarized as follows:

$$\begin{aligned} N_1 = 1 &\Rightarrow \delta_1 = 1, \phi_1 = 1 \\ N_2 > 1 &\Rightarrow \delta_2 > 1, \phi_2 > 1, \end{aligned} \tag{5}$$

without restricting the ratio $\frac{\delta_2}{\phi_2}$. The household receives an income stream y_1 and y_2 , and can borrow and save at an interest r which without loss of

generality is set to zero. Similarly, the discount factor is set to one, i.e. from the perspective of period one the utility in period two is not discounted.

In the subsequent analysis we will label the data-generating process for consumption as the *Demographics* model and the one that abstracts from demographics in the utility function as the *Single* model.

In the *Demographics* model, household size (and composition) affect utility as outlined in Equation (2) and the available resources for consumption are given by household income such that the optimization problem is represented by

$$\max_{c_{1,D}, c_{2,D}} U = u(c_{1,D}) + \delta_2 u\left(\frac{c_{2,D}}{\phi_2}\right) \quad (6)$$

subject to

$$c_{1,D} + c_{2,D} = y_1 + y_2 \equiv Y_D. \quad (7)$$

In contrast, the *Single* model abstracts from demographics in the utility function and the available resources for consumption are given by equivalized income. The corresponding optimization problem is thus

$$\max_{c_{1,S}, c_{2,S}} U = u(c_{1,S}) + u(c_{2,S}) \quad (8)$$

subject to

$$c_{1,S} + c_{2,S} = y_1 + \frac{y_2}{\phi_2} \equiv Y_S. \quad (9)$$

In both models the natural borrowing constraint is imposed and the same utility function is used.⁴ Furthermore, in both specifications equivalized consumption is the unit over which households optimize which makes the setups comparable. For the *Single* model the consumption levels $c_{1,S}$ and $c_{2,S}$ in fact reflect equivalized consumption because income as an input to the optimization problem has been equivalized before. For the *Demographics* model it is obvious in the second period as the household receives the utility from equivalized consumption $\frac{c_{2,D}}{\phi_2}$ which is however also true in the first period because household size is one in period one such that $c_{1,D} = \frac{c_{1,D}}{\phi_1}$.

Result 1. *The equivalized consumption profile in the Demographics model and Single model only coincide if $\delta_2 = \phi_2$.*

This result can be immediately read of from the two Euler equations which

⁴The utility function exhibits the standard properties regarding curvature, i.e. $u'(c) > 0$, $u''(c) < 0$ and $u(c)''' > 0$, and satisfies the Inada conditions, i.e. $u'(c) = \infty$ and $u'(c) = 0$ as $c \rightarrow 0$ and $c \rightarrow \infty$ respectively.

for the *Demographics* model is given by

$$u'(c_{1,D}) = \frac{\delta_2}{\phi_2} u' \left(\frac{c_{2,D}}{\phi_2} \right) \Leftrightarrow u'(c_{1,D}) = \frac{\delta_2}{\phi_2} u' \left(\frac{a_{2,D}}{\phi_2} + \frac{y_2}{\phi_2} \right), \quad (10)$$

with $a_{2,D} = y_1 - c_{1,D}$ being the assets carried over from period one. The Euler equation for the *Single* model is

$$u'(c_{1,S}) = u'(c_{2,S}) \Leftrightarrow u'(c_{1,S}) = u' \left(a_{1,S} - \frac{y_2}{\phi_2} \right), \quad (11)$$

with $a_{2,S} = y_1 + c_{1,S}$ being the assets carried over from period one. Equation (11) predicts a flat equivalized consumption profile for the *Single* model. The equivalized consumption profile in the *Demographics* model is only flat if $\delta_2 = \phi_2$. If $\delta_2 > \phi_2$, the equivalized consumption profile is upward sloping, i.e. $c_{1,D} = \frac{c_{1,D}}{\phi_1} < \frac{c_{2,D}}{\phi_2}$, while the opposite is true for $\delta_2 < \phi_2$. In the *Demographics* model the benefit of consuming one additional unit consumption in the second period

1. is associated with the marginal utility of equivalized consumption in period one $\left[u' \left(\frac{c_{2,D}}{\phi_2} \right) \right]$
2. which accrues to all household members reflected through the utility weight factor $[\delta_2]$
3. has to be divided by the equivalence scale $[\phi_2]$ because each household member does not get the full unit to consume but only the fraction $\frac{1}{\phi_2}$.

Put differently, δ_2 provides an incentive to save because the household enjoys a larger utility from consuming in the second period, as each unit of equivalized consumption is enjoyed by more (weighted) individuals. However, in period two every unit of consumption has to be shared with more people which is reflected through the division with the equivalence scale ϕ_2 . This reduces the incentive to save, i.e. to consume the assets, in period two. If $\delta_2 = \phi_2$, the latter two effects cancel out and therefore equivalized consumption in period one equals equivalized consumption in period two. If $\delta_2 > \phi_2$, equivalized consumption in period two exceeds equivalized consumption in period one. Relative to period one, the absolute loss in consumption in period two through the equalization is in this case outweighed by the larger gain in household utility through the utility weight.

Interestingly, such a configuration provides an additional explanation for the hump observed in equalized consumption documented in Fernández-Villaverde and Krueger (2007).

Result 2. *Life-time equivalized consumption in the Demographics model only coincides with life-time equivalized consumption in the Single model, if in the Demographics model the share of household consumption allocated to period two equals the share of household income received in period two.*

In order to show Result 2, it is helpful to define some variables. First, define on the household level prior to any equalization the share of household income y_2 received in period two relative to life-time household income as

$$\eta = \frac{y_2}{y_1 + y_2} = \frac{y_2}{Y_D}. \quad (12)$$

Second, for the *Demographics* model define similarly the share of household consumption allocated to period one relative to life-time household consumption

$$\zeta_D = \frac{c_{2,D}}{c_{1,D} + c_{2,D}}, \quad (13)$$

which is of course an endogenous object. Using these expressions life-time equivalized consumption from the *Demographics* model can be written as

$$C_D = c_{D,1} + \frac{c_{2,D}}{\phi_2} = \left[(1 - \zeta_D) + \frac{\zeta_D}{\phi_2} \right] (y_1 + y_2) = \left[(1 - \zeta_D) + \frac{\zeta_D}{\phi_2} \right] Y_D \quad (14)$$

and life-time equivalized income Y_S from the *Single* model as

$$Y_S = y_1 + \frac{y_2}{\phi_2} = \left[(1 - \eta) + \frac{\eta}{\phi_2} \right] (y_1 + y_2) = \left[(1 - \eta) + \frac{\eta}{\phi_2} \right] Y_D. \quad (15)$$

By the budget constraint (9) life-time equivalized consumption from the *Single* model ($C_S = c_{1,S} + c_{2,S}$) equals Y_S . Hence, the difference in life-time equivalized consumption between the *Demographics* and the *Single* model is given by

$$C_D - C_S = C_D - Y_S = (\eta - \zeta_D) \left(\frac{\phi_2 - 1}{\phi_2} \right) Y_D \quad (16)$$

which proofs Result 2. Whenever $\eta > \zeta_D$, the life-time equivalized consumption under the *Demographics* model is larger than under the *Single* model. The opposite is true for $\eta < \zeta_D$.

The intuition for this result can be explained best with a concrete example. Assume that the household income is zero in the first period ($y_1 = 0$), and positive in the second period ($y_2 > 0$) which implies that $\eta = 1$. In this case life-time equivalized income in the *Single* model is $\frac{y_2}{\phi_2}$ which by the budget constraint equals life-time equivalized consumption. In the *Demographics* model in turn, the household has the income y_2 available for consumption. By the Inada condition the household will consume in period one as well such that $c_{2,D} < y_2$ and thus $\zeta_D < \eta = 1$. Given that household size is

one in period one, in the calculation of life-time equivalized consumption in the *Demographics* model the equivalization matters only for period two consumption. Since $c_{2,D} < y_2$, “less” in absolute terms is lost through the equivalization in the calculation of life-time equivalized consumption in the *Demographics* model compared to the *Single* model.⁵ Essentially, Result 2 is the implication of a pure accounting exercise. The key driving force behind is that the households can shift consumption between periods but not income. If income and consumption allocations are not fully synchronized, then equivalization drives a wedge between equivalized consumption in the *Demographics* model and equivalized income.

Note that Result 2 and its derivation are completely independent of the relationship between δ_2 and ϕ_2 . These two functions of household size do however determine ζ_D and thus for a given η the relationship between the two equivalized consumption levels.

3.2 Income Inequality

In the following paragraph we introduce income heterogeneity, but not uncertainty, to demonstrate that the level differences between the two models also translate into equivalized consumption inequality differences. We will assume that the economy is populated with a mass of households of size one, where $\frac{1}{2}$ receives life-time household income Y_A and the other $\frac{1}{2}$ life-time household income Y_B .

Result 3. *Assume that life-time household income differs, i.e. $Y_D^A \neq Y_D^B$, but the share of household income in period two from life-time household income is the same, i.e. $\eta^A = \eta^B$. In this case any differences in life-time equivalized income between the *Demographics* and *Single* model (see Result 2) translate into differences in the level of life-time equivalized consumption inequality.*

Setting $Y_D^A = 1$ and $Y_D^B = \pi$ with $\pi > 1$, it is straightforward to show that the variances of life-time equivalized consumption in the two models

⁵More formally, rewrite life-time equivalized consumption in the *Demographics* model as

$$C_D = c_{1,D} + \frac{c_{2,D}}{\phi_2} = (y_2 - c_{2,D}) + \frac{c_{2,D}}{\phi_2} = y_2 - \frac{\phi_2 - 1}{\phi_2} c_{2,D},$$

which implies because $c_{2,D} < y_2$ that

$$C_D = y_2 - \frac{\phi_2 - 1}{\phi_2} c_{2,D} > y_2 - \frac{\phi_2 - 1}{\phi_2} y_2 = \frac{y_2}{\phi_2} = Y_S = C_S.$$

are given by

$$\text{Var}(C_S) = \left(\frac{\pi - 1}{2} C_S^A \right)^2 \text{ and } \text{Var}(C_D) = \left(\frac{\pi - 1}{2} C_D^A \right)^2. \quad (17)$$

Hence, the relationship between the two can be summarized as

$$\frac{\text{Var}(C_S)}{\text{Var}(C_D)} = \left(\frac{C_S^A}{C_D^A} \right)^2. \quad (18)$$

Whenever life-time equivalized consumption differs between the two models, i.e. $\eta \neq \zeta_D$, the level of life-time equivalized consumption inequality differs between the two models. If e.g. $\eta > \zeta_D$ and thus $C_S^A < C_D^A$, life-time equivalized consumption inequality is larger in the *Demographics* model. Through the equalization the absolute loss in equivalized income in the *Single* model is larger than the absolute loss in consumption in the *Demographics* model which reduces life-time equivalized consumption inequality more in the *Single* model.

3.3 Income Uncertainty

In the next paragraph, we analyze the role of income uncertainty. Assume there is a mass of households of size one. In the first period all household have the same income y_1 whereas income may take two values $y_2^L < y_2^H$ in period two with each state occurring with probability $\frac{1}{2}$. For the following results it is useful to define period two mean variables as $\bar{X} = \frac{1}{2} \sum_{i=L,H} X_{y_2=y_2^i}$.

Result 4. *In the presence of income uncertainty, the equivalized consumption profile is steeper in the Demographics model than in the Single model even if $\delta_2 = \phi_2$ and $\bar{\eta} = \bar{\zeta}_D$.*

In fact, we will show that Result 4 is always true whenever $\delta_2 \geq \phi_2$ and $\bar{\eta} \geq \bar{\zeta}_D$ and sometimes even when the inequality operators are reversed. In order to do so it is useful to start with the first-order condition for the *Single* model:

$$u'(c_{1,S}) = \frac{1}{2} \underbrace{\left[\underbrace{u'\left(y_1 - c_{1,S} + \frac{y_2^L}{\phi_2}\right)}_{c^{L2,S}} + \underbrace{u'\left(y_1 - c_{1,S} + \frac{y_2^H}{\phi_2}\right)}_{c^{H2,S}} \right]}_{\Gamma_S}. \quad (19)$$

Given the assumptions on the utility function ($u' > 0$, $u'' < 0$ and $u''' > 0$) the introduction of uncertainty introduces the standard precautionary sav-

ings motive which generates an upward sloping average consumption profile. The result is standard and induced by Jensen's inequality.⁶ The first-order condition for the *Demographics* model is given by

$$u'(c_{1,D}) = \frac{1}{2} \frac{\delta_2}{\phi_2} \underbrace{\left[u' \left(\underbrace{\frac{y_1 - c_{1,D}}{\phi_2} + \frac{y_2^L}{\phi_2}}_{c^{L2,D}} \right) + u' \left(\underbrace{\frac{y_1 - c_{1,D}}{\phi_2} + \frac{y_2^H}{\phi_2}}_{c^{H2,D}} \right) \right]}_{\Gamma_S} \quad (20)$$

For the moment consider the case of $\delta_2 = \phi_2$. Denote with $c_{1,S}^*$ the consumption allocation that solves Equation (19) for the *Single* model and set $c_{1,D} = c_{1,S}^*$. Hence, the savings brought to period two are the same in both models ($a_D = y_1 - c_{1,D} = y_1 - c_{1,S}^* = a_S$). However, consuming them in the *Demographics* model is associated with sharing reflected through the equivalization. As a consequence the uncertain part of equivalized cash on hand, i.e. equivalized assets plus equivalized income, is larger in the *Demographics* model. This implies that in the *Demographics* household more savings will be accumulated compared to the *Single* model and period one consumption will be lower, i.e. $c_{1,D}^* < c_{1,S}^*$.⁷ The difference in period one consumption is increasing in the ratio $\frac{\delta_2}{\phi_2}$ and even for values of $\delta_2 < \phi_2$, in the presence of uncertainty $c_{1,D}^*$ may still fall below $c_{1,S}^*$ because of the precautionary savings motive.

From Result 2 we know that for $\eta \geq \zeta_D$ life-time equivalized consumption in the *Demographics* model is at least as large as than in the *Single* model implying that also in the case with uncertainty the following relationship has to hold for $\bar{\eta} \geq \bar{\zeta}_D$:

$$c_{1,D} + \frac{\bar{c}_{2,D}}{\phi_2} \geq c_{1,S} + \bar{c}_{2,S}. \quad (21)$$

As we have already shown for $\delta_2 \geq \phi_2 \Rightarrow c_{1,D} < c_{1,S}$ Equation (21) thus implies that $\frac{\bar{c}_{2,D}}{\phi_2} > \bar{c}_{2,S}$. Note that even for values of $\bar{\eta} < \bar{\zeta}$, the equivalized consumption profile in the *Demographics* model is steeper than in the *Single*

⁶The easiest way to illustrate this is by setting $y_2^L = (1 - \epsilon)y_1\phi_2$ and $y_2^H = (1 + \epsilon)y_1\phi_2$. In the absence of uncertainty, the natural borrowing constraint induces a flat consumption profile for the *Single* model. Now assume the household would choose current consumption such that the expected, or ex post the average, consumption profile is flat, i.e. $c_1 = E_1(c_2) = \frac{1}{2} \left[y_1 + E \left(\frac{y_2}{\phi_2} \right) \right] = y_1$. By Jensen's inequality, the left hand side of Equation (19) is smaller than the right hand side. In order to achieve equality $c_{1,S}$ has to decrease which generates the upward sloping average consumption profile.

⁷More technically, $c_{1,D} = c_{1,S}^* \Rightarrow \frac{c_{2,D}^i}{\phi_2} < c_{2,S}^i \forall i = 1, 2$ and hence by the properties of the utility function $\Gamma_S > \Gamma_D$. Equation (20) can thus only be satisfied if period one consumption is lowered such that $c_{1,D}^* < c_{1,S}^*$.

model as long as

$$\bar{\eta} > \underbrace{\bar{\zeta}_D \xi - \frac{\phi_2}{\phi_2 - 1}(\xi - 1)}_{< \bar{\zeta}_D}, \quad (22)$$

with $\xi = \frac{c_{1,S}}{c_{1,D}}$ and $\xi > 1$. The key message from Result 4 is that the equivalized consumption profiles do not coincide any longer for the case of $\delta_2 = \phi_2$ and $\bar{\eta} = \bar{\zeta}_D$.

3.4 Income Inequality and Income Uncertainty

In the previous subsection income heterogeneity/inequality in period two implied by the income uncertainty

- does not lead to any equivalized consumption inequality in period one as in both models all agents are the same in period one,
- maps one to one into equivalized consumption inequality in period two as in both models all agents enter have the same (model-dependent) level of assets.

Result 3 of course applies with respect to life-time equivalized consumption inequality. We now add to the income uncertainty and thus income heterogeneity in period two, income heterogeneity in period one. Simply assume that in each model there is a mass of households of size one, with $\frac{1}{2}$ having income y_1^L and $\frac{1}{2}$ y_1^H in period one and $y_1^L < y_1^H$. The two possible income states in period 2 are as before given by y_2^L and y_2^H , each occurring with probability $\frac{1}{2}$ independent of the income realization in period one.

Result 5. *The same level of equivalized consumption inequality in period one in the Demographics and Single model, implies a lower level of consumption inequality in period two in the Demographics model.*

This setup now permits only to make relatively general statements. It can be shown that the variance of equivalized consumption in period one is

$$\text{Var}(c_{1,i}) = \frac{1}{4}(c_{1,i}^H - c_{1,i}^L)^2 \quad \forall i = S, D, \quad (23)$$

and in period two for the *Demographics* model

$$\begin{aligned}\text{Var}\left(\frac{c_{2,D}}{\phi_2}\right) &= \frac{1}{4} \left[\left(\frac{1}{\phi_2}\right)^2 (y_1^H - y_1^L - [c_{1,D}^H - c_{1,D}^L])^2 + \left(\frac{y_2^H - y_2^L}{\phi_2}\right)^2 \right] \\ &= \frac{1}{4} \left[\left(\frac{a_D^H}{\phi_2} - \frac{a_D^L}{\phi_2}\right)^2 + \left(\frac{y_2^H - y_2^L}{\phi_2}\right)^2 \right]\end{aligned}\tag{24}$$

and for the *Single* model

$$\begin{aligned}\text{Var}(c_{2,S}) &= \frac{1}{4} \left[(y_1^H - y_1^L - [c_{1,S}^H - c_{1,S}^L])^2 + \left(\frac{y_2^H - y_2^L}{\phi_2}\right)^2 \right] \\ &= \frac{1}{4} \left[(a_S^H - a_S^L)^2 + \left(\frac{y_2^H - y_2^L}{\phi_2}\right)^2 \right].\end{aligned}\tag{25}$$

The key difference is that for a given level of equivalized consumption inequality in period one, i.e. $c_{1,D}^H - c_{1,D}^L = c_{1,S}^H - c_{1,S}^L$, the difference in equivalized assets in the *Demographics* model will be smaller than in the *Single* model which yields a smaller variance. However, the level of consumption inequality in period one reflects endogenous choices and will be in general different between the two models. Result 5 demonstrates the non-trivial interactions that the presence of demographics in preferences has on equivalized consumption inequality relative to a model without such a dependency.

3.5 Summary

In the absence of income inequality and uncertainty equivalized consumption data generated by a simple model with demographics can only be replicated with a model without demographics but using equivalized income under two strict conditions. First, the utility weight δ needs to equal the equivalence scale ϕ . Otherwise, the equivalized consumption profiles are different. Second, the share of household consumption allocated to period two (ζ_D) in the model with demographics needs to equal the share of household income received in period two (η). Otherwise, the level of life-time equivalized consumption and inequality differ. Once income uncertainty is introduced, even for the case of $\delta = \phi$ and $\eta = \zeta_D$ the *Single* model cannot replicate the equivalized consumption profile generated by the *Demographics* model.

4 CRRA Preferences

In quantitative life-cycle models CRRA preferences are the prevailing choice for the utility function. In the next paragraph we will discuss the role of the parameter of relative risk aversion in the *Demographics* model.

The following lines restate the objective function for the *Demographics* model given in Equations (6) and (7) using CRRA preferences:

$$\max_{c_{1,D}, c_{2,D}} U = \frac{c_{1,D}^{1-\alpha}}{1-\alpha} + \delta_2 \frac{\left(\frac{c_{2,D}}{\phi_2}\right)^{1-\alpha}}{1-\alpha} \quad (26)$$

subject to

$$c_{1,D} + c_{2,D} = y_1 + y_2 \equiv Y_D. \quad (27)$$

The implied Euler equation for equalized consumption is

$$c_{1,D}^{-\alpha} = \frac{\delta_2}{\phi_2} \left(\frac{c_{2,D}}{\phi_2}\right)^{-\alpha}, \quad (28)$$

In conjecture with the budget constraint, the optimal equalized consumption can be obtained

$$c_{1,D} = \frac{1}{1 + \left(\frac{\delta_2}{\phi_2}\right)^{\frac{1}{\alpha}} \phi_2} \text{ and } \frac{c_{2,D}}{\phi_2} = \frac{\left(\frac{\delta_2}{\phi_2}\right)^{\frac{1}{\alpha}}}{1 + \left(\frac{\delta_2}{\phi_2}\right)^{\frac{1}{\alpha}} \phi_2}. \quad (29)$$

Since CRRA preferences are just a special case of the general utility function discussed before, we can see here again that it is only the ratio $\frac{\delta_2}{\phi_2}$ which determines the profile of the equalized consumption. $\delta_2 = \phi_2$ implies a flat profile, $\delta_2 > \phi_2$ a positive slope, whereas the opposite is true for $\delta_2 < \phi_2$. In the latter two cases, the slope is tilted towards one as α increases. A large α means a low intertemporal elasticity of substitution and thus a low willingness to have differences in equalized consumption between the two periods.

While the profile of equalized consumption is entirely determined by $\left(\frac{\delta_2}{\phi_2}\right)^{\frac{1}{\alpha}}$, the profile of household consumption, i.e. $c_{2,D} - c_{1,D}$, is in turn determined by $\left(\frac{\delta_2}{\phi_2}\right)^{\frac{1}{\alpha}} \phi_2$. The household consumption profile is thus only increasing if

$$\alpha > 1 - \frac{\ln \delta_2}{\ln \phi_2}. \quad (30)$$

Since α is restricted to positive values, condition (30) is always satisfied for

$\delta_2 \geq \phi_2$. Any value of $\alpha > 1$, which includes the typical range of values used for calibration, also satisfies Equation (30) even if $\delta_2 < \phi_2$. Equation (30) essentially states the condition under which values of $\alpha < 1$ still generate an upward-sloping household consumption profile if $\delta_2 < \phi_2$.

5 Quantitative Model

Our goal here is not to build a quantitative model that is able to replicate all the empirical facts with respect to consumption. Our aim is to test and evaluate the implications of our theoretical analysis with a simple, stripped-down version of a standard incomplete markets life-cycle model.

Households start their economic life in period t_0 with zero assets. During their working life until period t_w they receive a stochastic income y_t in every period. There is no labor supply choice. From period $t_w + 1$ onwards households are retired and have to live from their accumulated savings during working life. We abstract from pensions. Life ends with certainty at age T and households do not leave bequests and cannot die with debt. Households have access to a risk-free bond a which pays the interest rate r . In this version of the paper, households can borrow up to the natural borrowing constraint (level of debt that they can always repay for sure) at the interest rate r . We leave the case of a zero borrowing constraint as an extension.

In the *Demographics* model household size changes over the life-cycle deterministically and is homogenous across all households. The maximization problem is given by

$$\max_{\{a_{t+1}\}_{t=t_0}^{T-1}} E_0 \sum_{t=t_0}^T \beta \delta_t u\left(\frac{c_t}{\phi_t}\right) \quad \text{subject to} \quad (31)$$

$$c_t + a_{t+1} \leq (1+r)a_t + y_t \quad (32)$$

$$a_{t+1} \geq a_{\min,t}. \quad (33)$$

where δ_t and ϕ_t are functions of household age, household size and composition (t , $N_{ad,t}$ and $N_{ch,t}$ respectively). The income process is given by:

$$\ln y_t = \varrho_t + \epsilon_t, \quad (34)$$

where ϱ_t is an age-dependent experience profile and

$$\epsilon_t = \rho \epsilon_{t-1} + \varepsilon_t \quad \text{with } \varepsilon_t \sim N(0, \sigma^2). \quad (35)$$

The Euler equation to this problem is given by

$$\frac{\delta_t}{\phi_t} u' \left(\frac{c_t}{\phi_t} \right) = \beta(1+r) \frac{\delta_{t+1}}{\phi_{t+1}} E_t \left[u' \left(\frac{c_{t+1}}{\phi_{t+1}} \right) \right]. \quad (36)$$

The structure of the *Single* problem is very similar. Demographics don't affect the utility function while income y_t is deflated by household size through equivalence scales ϕ_t :

$$\max_{\{a_{t+1}\}_{t=0}^{T-1}} E_0 \sum_{t=0}^T \beta u(c_t) \quad \text{subject to} \quad (37)$$

$$c_t + a_{t+1} \leq (1+r)a_t + \frac{y_t}{\phi_t} \quad (38)$$

$$a_{t+1} \geq a_{\min,t}, \quad (39)$$

with y_t following the same process as described in Equations (34) and (35).

The Euler equation to this problem is given by

$$u'(c_t) = \beta(1+r) E_t [u'(c_{t+1})]. \quad (40)$$

6 Quantitative Features of the Model

A model period is one year. Agents start life at age 25, retire when 65 and live with certainty until age 75. In this version of the model, and to maintain maximum comparability with our theoretical analysis, we set $\beta = 1$ and $\beta(1+r) = 1$. We set the CRRA coefficient α equal to 1.57, the value estimated and used by Attanasio et al. (1999). In the next section, we use as a benchmark the square root scale $\phi_t^{SQR} = \sqrt{N_{ad} + N_{ch}}$, and compare our results with the OECD and the Nelson scales. These choices follow closely the discussion of equivalence scales in Fernández-Villaverde and Krueger (2007). The OECD scale has the lowest economies of scale while the opposite is true for the Nelson scale. The square root scale is almost identical to the “Mean” scale in Fernández-Villaverde and Krueger (2007) which is their preferred choice.⁸

As for utility weights, we remain agnostic and compare three extremes: (i) $\delta_t = 1$ represents the case when households do not value household size; (ii) $\delta_t = N_t = N_{ad,t} + N_{ch,t}$ is the opposite, since households always enjoy having more members; and (iii) an intermediate case when $\delta_t = \phi_t$, i.e., we let the utility weight take the same value as the equivalence scale.

⁸For explicit formulations of the different equivalence scales used in empirical consumption literature, see Table 1 in Fernández-Villaverde and Krueger (2007).

6.1 Income

We use data from the Current Population Survey, from 1984 to 2003. We use the March supplements for years 1985 to 2004, given that questions about income are retrospective. We use total wage income (deflated by CPI-U, leaving amounts in 2000 US dollars), and apply the tax formula of Gouveia and Strauss (1994) to get after-tax income.⁹

We define a full time/full year (**FTFY**) worker, as someone who worked more than 40 hours per week and more than 40 weeks per year and earned more than \$2 per hour. We construct total household income $W_{i\tau}$ for household i observed in year τ , as the sum of individual incomes in the household for all households with at least one full time/full year worker. Then, we estimate the following regression:

$$\log\left(\frac{W_{i\tau}}{\phi_{i\tau}}\right) = D_{i\tau}^{age} \varrho + X_{i\tau}\gamma + \epsilon_{i\tau} \quad (41)$$

where $\phi_{i\tau}$ is an equivalence scale, $D_{i\tau}^{age}$ represents a set of age dummies of the head of household, ϱ and γ are estimated coefficients and ϵ are estimation errors. Note that for the *Demographics* model we use household income for the estimation, i.e. $\phi_{i\tau} = 1 \forall i, \tau$. We also control for cohort effects and time effects by introducing birth year and year dummies in $X_{i\tau}$.¹⁰

From this estimation, we are interested in the regression coefficients associated with age dummies of the household head (experience profiles in the model). In our exercise below, we use smoothed profiles, which we show in Figure 1 for different choices of equivalence scales.

From the estimation residuals, we calibrate the income process in (35). Our calibration procedure is standard and follows Storesletten et al. (2004): we pick values of ρ and σ in order to minimize the square difference between the profile of observed cross-sectional variances of income and the simulated one (given the chosen parameters). We also pick values of σ_0 , the standard deviation for the unconditional distribution of the first income shock ε_0

⁹Wage income in the CPS is pre tax income.

¹⁰Since year dummies are perfectly collinear with age and birth cohort dummies, we follow Aguiar and Hurst (2009) and include normalized year dummies instead, such that for each year τ

$$\sum_{\tau} \gamma_{\tau} = 0 \quad \text{and} \quad \sum_{\tau} \tau \gamma_{\tau} = 0$$

where $\{\gamma_{\tau}\}$ are the coefficients associated to these normalized year dummies. To compare life-cycle profiles across different cohorts/time periods, we normalize the estimated coefficients associated to age dummies by adding the effect of a particular cohort/time. More specifically, we picked the cohort corresponding to the median age observed at the last observed year.

Figure 1: Experience Profile for Households

Figure 2: Profiles for Household Size and Composition

Table 2: Calibrated Parameters: ρ , σ and σ_0

ρ	σ	σ_0
0.9906	0.0189	0.1575

in order to match the cross sectional variance of income for our first age group (25 years old). We present these values in table 2. We discretize this calibrated process using the Rouwenhorst method, using 20 points for the shock space. This methodology is specially suited for our case, given the high persistence of the process. This is discussed in (Kopecky and Suen, 2010).

To maintain full comparability with our simple theoretical model, we perform an *ex-post* equivalization procedure for the income process in the *Single* model: we use the same calibrated income profiles and shocks in both the *Demographics* and *Single* models (the calibration in Figure 1 and Table 2) and then feed the equivalized experience profiles to the *Single* model. Besides making the quantitative model more comparable to the theoretical model, this approach maintains the same shock structure across considered equivalence scales, making the comparison of biases more direct, since no extra “noise” is being introduced by different volatility parameters. This would be the case for an alternative approach, or an *ex-ante* equivalization: estimating Equation (41) with a particular equivalence scale $\phi_{i\tau}$, resulting in different age profiles and calibrated income shocks for the *Single* model. Since the income has already been equivalized *ex-ante*, there is no need to do so in budget constraint (38).¹¹

6.2 Family Structure

We use the March supplements of the CPS for years 1984 to 2003. For each household, we count the number of adults (individuals age 17+) and the number of children: individuals age 16 or less who are identified as being the “child” of an adult in the household. We restrict our sample to consider households with at most 2 adults and 4 children. We compute two separate profiles: one for number of adults and one for number of children. As above, we run dummy regressions to extract life-cycle profiles, where the considered age is that of the head (irrespective of gender) and control for cohort and year effects. After extracting these life-cycle profiles, we smooth them using a cubic polynomial in age, and restrict the number of children to zero after age 60. The results of this procedure are in Figure 2.

¹¹In our computations below, we do not find major differences between the *ex-post* or *ex-ante* equivalization strategies, so we show results only for the former. The results of these exercises are available on request.

7 Results

7.1 Deterministic Framework

We first present a deterministic version of the model, to highlight intuition and draw comparisons with the theoretical results in the previous section. Since there is no heterogeneity for any given age, in this subsection we present results on mean consumption profiles only.

7.1.1 Household Consumption

Figure 3 shows household consumption for different cases of household preferences, normalized at age 25. As a benchmark we also add the profile for the *Single* model. Recall that the output of the *Single* model is an equilibrated profile which is flat, since $\beta(1+r) = 1$. For the different cases of the *Demographics* model, consumption profiles exactly replicate the profile of household size and peak exactly when household size is the largest (see Figure 2). Recall from the discussion in Section 4 that for any ratio of $\frac{\delta}{\phi}$ the household consumption profile is increasing in household size if $\alpha > 1$. For the case when $\delta_t = N_t$ (utility weight equals total household size) the household values its size the most and accordingly the mean household consumption profile exhibits the largest “hump” early in life and the fastest decrease later in life. The opposite is true for the utility weight $\delta_t = 1$: households do not enjoy extra utility when its size grows, so we get the flattest consumption profile earlier in life and the slowest decline later, among all utility weight cases. Finally, $\delta_t = \phi_t$ is an intermediate case since in general, $\phi_t \in [1, N]$. As seen from Figure 3, the profile for total household consumption for this case lies exactly in between the two cases above.

All profiles intersect around age 50, which is the result of household size being equal at both that age and age 25 (first period of the model). Given the assumption of a natural borrowing constraint and no income uncertainty, periods with the same household size lead in all cases to the same consumption level, since all forces introduced by household size cancel each other. Just looking at the Euler equation (36) for the deterministic case, and substituting forward, it is clear that whenever two periods have the same household size, total consumption is the same, no matter what the utility weight is (again, using the fact that $\beta(1+r) = 1$).

Figure 3: Household Consumption Relative to Age 25

Figure 4: Equivalized Consumption Relative to *Single* Model

7.1.2 Equivalized Consumption

In Figure 4, we show equivalized consumption profiles from the *Demographics* model relative to the *Single* model.¹² We see that for every case of δ_t , equivalized consumption from the *Demographics* model *does not* coincide with the predictions of the *Single* model. Just as predicted by Result 1, the equivalized consumption profile generated by the *Demographics* model has the same shape as the one generated by the *Single* model if $\delta_t = \phi_t$. Whenever $\delta_t > \phi_t$, equivalized consumption is increasing in household size and the opposite is true for $\delta_t < \phi_t$. Even in the case where the utility weight is equal to the chosen equivalence scale ($\delta_t = \phi_t$) we get a biased profile of adult equivalent consumption: the equivalized consumption from the *Demographics* model is around 2.2% higher than in the single model on average (and in terms of lifetime equivalized consumption). The same happens in the case with $\delta_t = 1$ and $\delta_t = N_t$, with lifetime differences of equivalized consumption relative to the *Single* model around 3.5% and 1.0% respectively. This fact can be explained with Result 2 from the previous section: only when shares of household consumption and household income in each period and across models are the same, the *level* of equivalized consumption will be the same.¹³ That this condition is not met can be clearly seen in Figure 5, where we show that shares of household consumption per age over lifetime household consumption are equal across models only at age 25 and age 50; in terms of household income, the shares are all the same across models (by construction) but very different to the consumption shares at all ages.

7.2 Stochastic Framework

We now show simulation results from a model with the stochastic income process described in the previous section. We want to analyze simulated profiles of mean equivalized consumption, the associated implications for equivalized consumption inequality over the life-cycle and the differences between the *Demographics* approach and the *Single* model.

¹²Note that because the *Single* model predicts a flat consumption profile, the *Demographics* profiles relative to age 25 look exactly the same except of a level shift.

¹³In fact, in the multi-period case, also other combinations of household consumption shares could generate a life-time equivalized consumption in the *Demographics* model that coincides with the *Single* model as long as $\sum_{t=t_0}^T \frac{\zeta_{t,D}}{\phi_2} = \sum_{t=t_0}^T \frac{\eta_t}{\phi_t}$ with $\sum_{t=t_0}^T x_t = 1 \forall x = \zeta_D, \eta$. The *Demographics* model would however have to generate endogenously the right shares of household consumption in order to satisfy this condition.

Figure 5: Household Consumption and Income Shares

Note: Deterministic Income; Square Root Scale

7.2.1 Equivalized Consumption Means

Figure 6 shows the effect of uncertainty on the profile of mean equivalized consumption across different models relative to the respective consumption at age 25. Note that this figure is not the counterpart to Figure 3 which showed household consumption. As discussed in the context of Result 4, in the *Single* model consumption is increasing in age, reflecting the role of precautionary savings. This is true also for the rest of the models and in stark contrast to the deterministic case where the *Single* model profile was completely flat. The effect of utility weights is also present here: when $\delta_t = 1$, there is no utility from household size, hence the household consumes more when household size is low (later ages); moreover, consumption growth mirrors negatively the growth in household size. In the model with a deterministic income, equivalized consumption followed household size for $\delta_t = N_t$. In the presence of income uncertainty, the precautionary savings motive is stronger than this demographic force. The effect is reflected in the growth rate of equivalized consumption: it is positive until retirement, and has the steepest slope when household size increases. When household size starts to decrease, this slope starts decreasing as well. For the intermediate case ($\delta_t = \phi_t$), equivalized consumption tracks very closely the single household case, but the slope of the profile is marginally steeper. The latter effect is predicted by Result 4 and can be seen even more clearly in Figure 7 where we present equivalized household consumption for the three *Demographics* models, relative to the single household model prediction. All the relative effects we saw in the deterministic case remain the same. On the

Figure 6: Equivalized Consumption Relative to Age 25

Note: Stochastic Income; Square Root Scale

Figure 7: Equivalized Consumption Relative to *Single* Model

Note: Stochastic Income; Square Root Scale

other hand, average equivalized consumption relative to the single household model is 4.8%, 3.7% and 2.7% larger over the lifetime for the cases with $\delta_t = 1$, $\delta_t = \phi_t$ and $\delta_t = N_t$ respectively. These values are higher than in the deterministic case because of the different allocation of consumption over the life-cycle, compare Figures 5 and 8. Via this channel, the introduction of uncertainty increases the relative differences between the predictions of *Demographics* and *Single* models, all else constant.

7.2.2 Equivalized Consumption Inequality

With the introduction of stochastic income processes, we can compare equivalized consumption inequality for our different models over the life-cycle. In Figure 9 we see that all models imply increasing equivalized consumption inequality in the life cycle. However, how inequality evolves depends again on the choice of utility weights. For the case when $\delta_t = N_t$, equivalized consumption is higher when the household is larger (earlier in life), so equivalized consumption tracks income shocks more closely which can also be seen in Figure 8. Hence, the cross sectional inequality related to that model rises the fastest until around age 50. On the other hand, when $\delta_t = 1$ and households do not enjoy household size, households want to postpone consumption to periods when household size is small; hence, they accumulate more assets for every income shock, relative to the $\delta_t = N_t$ case. Hence, cross sectional consumption inequality is lower earlier in life for the $\delta_t = 1$ setup in comparison to when $\delta_t = N_t$. Later in life, there is a reversion of absolute levels of consumption inequality, since households in the $\delta_t = 1$ economy have accumulated assets in a more heterogeneous way, which leads to more heterogeneous consumption. The opposite is true for the $\delta_t = 1$ case, so the lines in the figure change their relative position.

In Figure 10 we show the evolution of equivalized assets and cross section inequality (standard deviations). In line with the discussion above, we see that both the level and inequality of assets rises the fastest for the case of $\delta_t = 1$ and the slowest when $\delta_t = N_t$.

In Figure 11 we compare the predictions for equivalized consumption inequality from the different *Demographics* models with respect to the *Single* model. The results are qualitatively the same as for mean equivalized consumption. The life-cycle profiles of inequality have the same shape as in Figure 7 but the magnitudes differ strongly: even in the case when the utility weights are equal to the equivalence scale $\delta_t = \phi_t$, the *Demographics* model predicts 5% more cross sectional inequality per period in equivalized consumption than the *Single* model. For the other utility weights the differences are even more pronounced.

Figure 8: Household Consumption and Income Shares

Note: Stochastic Income; Square Root Scale

Figure 9: Equivalized Consumption Inequality (S.D.)

Note: Stochastic Income; Square Root Scale

Figure 10: Equivalized Assets

(a) Means

(b) Inequality (S.D.)

Note: Stochastic Income; Square Root Scale

Figure 11: Standard Deviation of Equivalized Consumption Relative to *Single* Model

Note: Stochastic Income; Square Root Scale

Figure 12: Coefficient of Variation of Equivalized Consumption

Note: Stochastic Income; Square Root Scale

7.2.3 Equivalized Consumption Inequality: a Scale-Free Measure

Our discussion so far has focused on age specific means and standard deviations of equivalized consumption across models. We now propose a third statistic, the coefficient of variation (CV), to assess the differences in consumption inequality. The CV is the ratio of age specific standard deviations and age specific means, and thus is a normalized measure of dispersion in equivalized consumption. We think that this is an important measure to study, since age profiles of equivalized mean consumption vary markedly across models.¹⁴ Figure 12 demonstrates that all models have a similar CV profile, and that this relative inequality measure is increasing in the utility weight. Contrast that information with the absolute inequality measured by the SD in Figure 9.

Figure 13 shows the CV for the *Demographics* model, under different utility weights, along consumption means and standard deviations, all relative to the *Single* model.

We already explained at length the pattern for equivalized consumption means and standard deviations over the life-cycle for the different utility weights. For $\delta_t \neq \phi_t$, the CV follows a reversed profile. Whenever consumption (and inequality) are low relative to the *Single* model (over the life-cycle), the CV is large and vice versa.

For $\delta_t = 1$, in the *Demographics* model households want to shift consumption away from periods with large household size and save a lot during these periods even if being hit with large income shocks. Compared to the *Single* model, this implies a relatively larger equivalized consumption inequality (measured by the CV) even conditional on the lower equivalized consumption mean and absolute inequality (measured by the SD). In contrast, later in life (around age 47-48) the larger amount of accumulated assets allows to insure against income shocks better relative to the *Single* model.

The opposite is true for the profile of the CV in the case that $\delta_t = N_t$. In the *Demographics* model, households want to allocate consumption to periods where household size is large. During these time periods, the CV is the lowest. Bad income shocks are smoothed by eating up savings and by borrowing against future income. Later in life, this behavior results in a higher CV. Over the entire life-cycle, relative equivalized consumption inequality as measured by the CV is higher in the *Demographics* model than in the *Single* model. Hence, the higher life-time resources do not only lead to a larger absolute equivalized consumption inequality (as measured by the SD) but also to a larger relative equivalized consumption inequality (as

¹⁴In particular, recall Result 3: any difference in life-time equivalized consumption between the *Demographics* and *Single* model, maps one to one into differences in life-time equivalized inequality (measured by the standard deviation or variance).

Figure 13: Equivalized Consumption: Means and Inequality

Note: Stochastic Income; Square Root Scale

Figure 14: Number of “Adult Equivalents” for Different Equivalence Scales

Note: For explicit formulations of the different equivalence scales, see Table 1 in Fernández-Villaverde and Krueger (2007).

measured by the CV).

This latter effect is also present for the case of $\delta = \phi$, i.e. the CV is again higher in the *Demographics* model in all periods. The equivalized consumption profile is steeper earlier in life in the *Demographics* model because households accumulate more savings. This is because households in the *Demographics* model share savings among its member if they eat them up to smooth income shocks.¹⁵ However, given household size and equivalence scales, they are less willing to do so early in the life-cycle. This explains the increase of the CV early in life and the decrease later in life as higher savings can be used for insurance. This whole line of argumentation has of course to be seen relative to the *Single* model.

The main insight from this discussion, is that even in this scale-free measure, we detect differences between the predictions of the *Single* model and the equivalized *Demographics* model. After controlling for differences in mean consumption, implied inequality in the *Demographics* models can be up to 2% higher than from the *Single* model. Furthermore, incentives for self-insurance change over the life-cycle, not only because of the degree of uncertainty but also because of the interplay between uncertainty and household size and composition which (depending on the utility weight δ) provides additional motives for consumption-savings decisions.

¹⁵Compare Result 3.

7.3 Different Equivalence Scales

In Figure 14 we compare the implied value of equivalence scales, for different household sizes. The vertical axis measures the number of “adult-equivalent” members in the household. Below, we expand our analysis for the OECD and Nelson scales. These choices follow closely the discussion of equivalence scales in Fernández-Villaverde and Krueger (2007) as the OECD scale has the lowest economies of scale while the opposite is true for the Nelson scale. Furthermore, the square root scale “SR” is almost identical to the “Mean” scale in Fernández-Villaverde and Krueger (2007) which is their preferred choice.

In Figure 15 we show equivalized consumption for the different utility weights and equivalence scales considered. All the profiles are in terms of consumption from the single agent model. Recall from Figure 14, the OECD scale has the lowest economies of scale; put differently, is the closest in value to N . On the other hand, the Nelson scale has the highest economies of scale and is the closest to 1. The square root scale falls between the two. From Figure 15, we see that the ratio (δ_t/ϕ_t) , i.e. the difference between utility weights of additional family members and equivalence scales, determines the size of the difference between the *Demographics* and the *Single* model over the life-cycle: the bigger the difference between this ratio and one, the bigger the increase in the bias over the life-cycle introduced by the equalization of household data (as long as δ is not set exactly to ϕ).

For example, when we consider the OECD equivalence scale and the $\delta_t = N_t$ as utility weight (a case when the ratio is closest to one, Figure 15c), differences in equivalized consumption across the *Demographics* and *Single* models are milder than when we consider the $\delta_t = 1$ case (when the ratio deviates the most from one, Figure 15a). The opposite is true for the Nelson scale: equivalized consumption from the model with utility weights of $\delta_t = N_t$ are the ones that result in the most extreme differences with the *Single* model. The intuition behind the shape of these figures is the same that we discussed previously for the square root scale.

When the utility weight is exactly equal to the equivalence scale ($\delta_t = \phi_t$), the profiles of equivalized household and single household consumption are very similar throughout. However, the relative value of equivalized lifetime income is lower the higher the equivalence scale. Hence, for the OECD scale, the relative difference between equivalized household and single household consumption is greater than for the square root and Nelson scales.

As seen from Figure 15b, the best case scenario for the *Single* model is when the utility weight equals the equivalence scale ($\delta_t = \phi_t$), see Figure 15b) and the equivalence scale is the one provided by Nelson: in that case, the differences between the predictions of equivalized consumption across

Figure 15: Equivalized Consumption Relative to *Single* Model

Note: Stochastic income

models is the smallest (around 1% each period). This is due to the fact that the Nelson scale admits the highest economies of scale among all considered equivalence scales, meaning that the effect of household size on household income and consumption is very mild. In simple terms, given that with the Nelson scale $\phi_t \approx 1$ for all t , and $\delta_t = \phi_t$, in this case both the *Demographics* and the *Single* models behave very similarly.

In terms of equivalized consumption inequality, the resulting figures are very similar to Figure 15 for the standard deviations and the logic behind is similar. Given incomplete markets and idiosyncratic income shocks, periods in which households want to consume more, consumption is more subject to reflect income shocks, and thus, consumption inequality is higher. We therefore only present the CVs in Figure 16. For the cases where $\delta_t \neq \phi_t$, we still find that the further the ratio δ_t/ϕ_t is from one, the more extreme the differences in the predictions of the *Demographics* versus the *Single* model over the life-cycle. E.g. as argued before, for $\delta_t = 1$, the large value of the OECD scale compared to the Nelson scale makes the households in the *Demographics* model much less willing to consume when household size is large and thus, to consume out of savings. As a consequence, the CV relative to the *Single* model increases much more early in life for the OECD scale compared to the Nelsons scale which is however reversed later in life. From here we conclude that the single approach to modeling consumption, not only introduces absolute but also relative differences in inequality. Moreover, these differences vary over the life-cycle.

8 Discussion: value of δ and ϕ

We have discussed at length different cases for utility weights and equivalence scales and the sources of differences between what *Demographics* and *Single* models predict. However, the question still remains with regard to which values we should consider for empirical work. For example, Hong and Ríos-Rull (2007) set the utility weights to one,¹⁶ while Fuchs-Schündeln (2008) sets it equal to the equivalence scale. In both papers no further justification for this choice is provided.

Given the utility function chosen by Attanasio et al. (1999) and their parameter estimates for ζ_1 , ζ_2 and α in equation (1) allows us to back out δ for a given equivalence scale ϕ , by comparing the preferences in (1) and in (2):

$$\exp(\xi_1 N_{ad} + \xi_2 N_{ch}) = \left(\frac{\delta(N_{ad}, N_{ch})}{\phi(N_{ad}, N_{ch})^{1-\alpha}} \right) \quad \forall N_{ad}, N_{ch}. \quad (42)$$

¹⁶However, they consider two agent households and stochastic transitions between marital states, making the comparison to a unitary model difficult.

Figure 16: Coefficient of Variation Relative to *Single* Model

For $N_{ad} = 1$ and $N_{ch} = 0$, our setup implies $\delta = \phi = 1$ whereas the preference parameter in Attanasio et al. (1999) is $\exp\zeta_1$. We therefore normalize the utility function (1) from their setup by dividing it with this factor such that

$$\delta(N_{ad}, N_{ch}) = \exp(\zeta_1[N_{ad} - 1] + \zeta_2 N_{ch}) \phi(N_{ad}, N_{ch})^{1-\alpha} \quad (43)$$

Note that it is not possible to uniquely pin down individual household member weights but only the sum of the weights given by δ . In Figure 17, we show the calculated δ_t/ϕ_t ratios, for different household sizes and various equivalence scales. In addition, we plot the two extreme cases we considered in the quantitative analysis, $\delta_t = 1$ and $\delta_t = N_t$, and a utility weight of $\delta_t = \phi_t$, represented by a flat line equal to 1, as reference points.

For the OECD scale, probably the most common choice as equivalence scale (for example, see Krueger et al. (2010)) the estimates by Attanasio et al. (1999) imply a ratio δ_t/ϕ_t close to one. In terms of equivalized consumption means and inequality (relative to the single model) this case is associated with the top profile in Figure 15b and thus, the largest deviation from the *Single* model among the three considered scales. For the Nelson scale, the implied δ_t/ϕ_t ratio comes closest to the $\delta_t = N_t$ case which implies the largest deviation from the single model in terms of means (and volatilities) depicted in Figure 15c by the profile that exhibits the most extreme differences with respect to the *Single* model.

Figure 17 shows that for the Square Root scale, the δ_t/ϕ_t ratio is right between the bands and is higher than one. Since the ratio is clearly different from one (the case with the least amount of bias introduced by the equivalization of *Demographics* models) our results raise concerns about the reliability of the single household framework.

Our main conclusion from this section is that, given the limited empirical evidence available, predictions from the single household approach are likely biased, no matter the choice of equivalence scale used.

9 Discussion and Future Work

We want to stress that our analysis, although being mostly quantitative, is still of theoretical nature. E.g. deviating from the assumption that the discount factor β was set to one and the interest rate to zero, or more generally from $\beta(1+r) = 1$, will interact with the life-cycle consumption and savings decisions and might downplay or accelerate the differences between the two models. The question remains as to how important would these biases be in a state of the art model of consumption and income which we will consider for future work. Furthermore, we still need to explore additional factors which might affect the consumption-savings problem of the household, and

Figure 17: Ratio $\frac{\delta}{\phi}$ for Various Utility Weights and Equivalence scales

Note: Calculations based on households with two adults and the respective number of children.

thus, create differences to predictions of single agent models: zero borrowing constraints, stochastic demographic profiles, heterogeneous income and demographic profiles, among others.

10 Conclusions

In this paper we analyze the differences between the predictions of a single agent model and the equivalized predictions of a model with household size (and composition) effects. Specifically, we are interested in predictions from the standard incomplete markets model for mean and cross sectional inequality of life-cycle equivalized consumption.

We perform our analysis by extending the standard consumption-savings framework to allow for changes in household size and composition during the life-cycle. We propose an explicit formulation combining equivalence scales, which reflect economies of scale in household consumption and utility weights, which make explicit the value of having a household of different size and composition.

Using this framework, we study the potential sources of differences between the predictions of the *Single* model, calibrated using equivalized income, and the equivalized consumption from the *Demographics* model (which uses total household income). Given that the single agent model is just a particular case of what we label the *Demographics* model, we think of this framework as the mildest test to the single household approach.

Our quantitative results show that the differences between equivalized consumption data from the *Demographics* model and individual data from the *Single* model can be substantial, both in terms of predicted levels and cross sectional inequality. This result calls into question the standard approach of concentrating on single agent models.

References

- Aguiar, M. and Hurst, E. (2009). Deconstructing Lifecycle Expenditure, *Working paper*, University of Rochester and University of Chicago.
- Attanasio, O. P., Banks, J., Meghir, C. and Weber, G. (1999). Humps and Bumps in Lifetime Consumption, *Journal of Business & Economic Statistics* **17**: 22–35.
- Blundell, R., Pistaferri, L. and Preston, I. (2008). Consumption Inequality and Partial Insurance, *American Economic Review* **98**: 1887–1921.

- Fernández-Villaverde, J. and Krueger, D. (2007). Consumption over the Life Cycle: Facts from Consumer Expenditure Survey Data, *Review of Economics and Statistics* **89**(3): 552–565.
- Fernández-Villaverde, J. and Krueger, D. (2010). Consumption and Saving over the Life Cycle: How Important are Consumer Durables?, *Macroeconomic Dynamics* **Forthcoming**.
- Fuchs-Schündeln, N. (2008). The response of household saving to the large shock of german reunification, *American Economic Review* **98**(5): 1798–1828.
- Gourinchas, P.-O. and Parker, J. A. (2002). Consumption over the life cycle, *Econometrica* **70**: 47–89.
- Gouveia, M. and Strauss, R. (1994). Effective federal individual income tax functions: An exploratory empirical analysis, *National Tax Journal* **47**(2): 317–339.
- Güvenen, F. and Smith, A. A. (2010). Inferring labor income risk from economic choices: An indirect inference approach, *Working paper*, University of Minnesota.
- Heathcote, J., Perri, F. and Violante, G. L. (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.
- Hong, J. H. and Ríos-Rull, J.-V. (2007). Social security, life insurance and annuities for families, *Journal of Monetary Economics* **54**(1): 118 – 140. Carnegie-Rochester Conference Series on Public Policy: Economic Consequences of Demographic Change in a Global Economy April 21-22, 2006.
- Kaplan, G. and Violante, G. L. (2009). How Much Consumption Insurance Beyond Self-Insurance?, *Working paper*, University of Pennsylvania and New York University.
- Kopecky, K. and Suen, R. (2010). Finite state markov-chain approximations to highly persistent processes, *Review of Economic Dynamics* **13**(3): 701–714.
- Krueger, D. and Perri, F. (2006). Does Income Inequality Lead to Consumption Inequality? Evidence and Theory., *Review of Economic Studies* **73**(1): 163 – 193.
- Krueger, D., Perri, F., Pistaferri, L. and Violante, G. L. (2010). Cross-sectional facts for macroeconomists, *Review of Economic Dynamics* **13**(1): 1 – 14. Special issue: Cross-Sectional Facts for Macroeconomists.

- Lazear, E. P. and Michael, R. T. (1980). Family size and the distribution of real per capita income, *The American Economic Review* **70**(1): pp. 91–107.
- Storesletten, K., Telmer, C. I. and Yaron, A. (2004). Consumption and risk sharing over the life cycle, *Journal of Monetary Economics* **51**(3): pp. 609–633.