

Advances

Jeppe Druedahl* and Thomas H. Jørgensen*

Persistent vs. Permanent Income Shocks in the Buffer-Stock Model

DOI 10.1515/bejm-2016-0035

Previously published online June 29, 2016

Abstract: We investigate the effects of assuming a *fully permanent* income shock in a standard buffer-stock consumption model, when the true income process is only *highly persistent*. This assumption is computationally very advantageous, and thus often used, but might be problematic due to the implied misspecification. Across most parameterizations, and using the method of simulated moments, we find a relatively large estimation bias in preference parameters. For example, assuming a unit root process when the true AR(1) coefficient is 0.97, leads to an estimation bias of up to 30 percent in the constant relative risk aversion (CRRA) coefficient. If used for calibration, misspecified preferences could, for example, lead to a serious misjudgment in the value of social insurance mechanisms. Economic behavior, such as the marginal propensity to consume (MPC), of households simulated from the estimated (misspecified) model is, on the other hand, rather close to that from the correctly specified model.

Keywords: imperfect markets life cycle model; marginal propensity to consume; persistent and permanent income shocks; simulated method of moments.

JEL: D31; D91; E21.

1 Introduction

The degree of persistence of income shocks has potentially important implications for a broad range of economic decisions, and has therefore received significant

*Corresponding authors: **Jeppe Druedahl**, Department of Economics, University of Copenhagen, Øster Farimagsgade 5, Building 26, DK-1353 Copenhagen, Denmark, e-mail: jeppe.druedahl@econ.ku.dk. Website: <http://econ.ku.dk/druedahl>; and **Thomas H. Jørgensen**, Department of Economics, University of Copenhagen, Copenhagen, Denmark, e-mail: thomas.h.jorgensen@econ.ku.dk. Webpage: www.tjeconomics.com

empirical attention.¹ For a given income shock variance, more persistent shocks, e.g. imply higher effective uncertainty and thus affect the optimal level of wealth accumulation for both self-insurance and retirement purposes. We will remain agnostic about the exact degree of persistence in income, and instead investigate how severe a bias in preference estimates, and associated household behavior, researchers face when assuming a *fully permanent* income shock when the true data generating process only contains a *highly persistent* income shock.

We do this relying on the modern work horse buffer-stock consumption model pioneered by Deaton (1991, 1992) and Carroll (1992, 1997), which we describe in the next section. In that model, with constant relative risk aversion (CRRA) utility, the assumption of a permanent shock (unit root in log income) is computationally very beneficial because it implies that the model can be normalized with permanent income, essentially removing permanent income as a continuous state variable in the dynamic programming problem. If, on the other hand, income shocks are persistent but not permanent, the model cannot be normalized by income and has to be solved over a grid of income nodes at an increased computational burden. The trick of normalizing by permanent income also extends to more complex models with multiple states in which removing a continuous state variable by normalization is even more beneficial.

To quantify the bias, we assume that the researcher has access to the true average wealth age profile, insists of assuming a fully permanent income shock, and knows all the parameters of model with the exception of the preference parameters, which he or she estimates by simulated method of moments. He or she consequently estimates preferences to match the average wealth age profile, and draws conclusions about the economic behavior of households based on the estimated model with a misspecified unit root process and potentially biased preferences estimates.

It is not surprising that a misspecified model in general would lead to potentially flawed conclusions. However, the amount of empirical work based on the standard imperfect markets buffer-stock model (and its various extensions) with a unit root income process merits an analysis of the severity of these problems.² Since all work in economics are presumably simplifications of the real world, we should care about choosing our simplifying assumptions to minimize the error

¹ For early studies see in particular Lillard and Willis (1978); Lillard and Weiss (1979); MaCurdy (1982); Abowd and Card (1989) and Baker (1997), while more recent contributions, and additional references, are found in Meghir and Pistaferri (2004); Guvenen (2009); Browning et al. (2010) and Guvenen et al. (2015).

² A few examples include Gourinchas and Parker (2002); Cagetti (2003); Blundell et al. (2008); Kaplan and Violante (2010) and Blundell et al. (2013).

associated with them. Our analysis provide researchers with results enabling them to draw their own conclusions concerning this trade-off depending on the particular question in mind. We furthermore supply online MATLAB code used to generate all results in the paper with an user-friendly simulation suite in which researchers can analyze implied behavior of their own choice.

In our baseline parametrization, described in Section 3, we find a relatively large estimation bias in both the CRRA parameter and the discount rate. Particularly, in Section 4 we report an estimation bias for the CRRA parameter increasing from 6 percent, when the true level of persistence is 0.99, to about 28 percent when it is 0.97. For the discount rate the increase is from 4 percent to 15 percent. We further show that the results are robust in the sense that the bias is almost always large irrespectively of changing our assumptions about the preferences and the income process. Interestingly, however, the bias can be both positive and negative.

Our results affect many fields of economics. It is now standard procedure to calibrate models using externally estimated preference parameters, and the discount factor and/or CRRA parameter are no exceptions. If such estimates are based on a misspecified unit root income process, researches will likely use (significantly) biased estimates to calibrate their models. This would naturally affect the implied policy proscriptions. An excessively high CRRA parameter would, for example, all else equal imply an over-valuation of social insurance mechanisms. Likewise it would over-state the cost of business cycles.

In many macroeconomic applications, however, the preference parameters are of little interest in their own right. Rather, economic outcomes such as the marginal propensity to consume (MPC) and the marginal propensity to consume out of permanent shocks (MPCP) are of great interest. We find that the MPC and the MPCP from the estimated misspecified model follow the same overall age profile as the true MPC and MPCP, and only have, what we judge to be, relatively minor deviations.

As we conclude in Section 5, the results suggest that if researchers are interested in the value of preference parameters, such as the relative risk aversion coefficient or discount rate, the income process specification will have large and significant implications for the conclusions drawn. On the other hand, if the interest is on overall patterns of economic behavior, such as the average MPC or MPCP, a (slightly) misspecified income process seems to affect the conclusions much less.

Two other studies have previously to some degree investigated the implications of misspecifying the income process in a buffer-stock consumption model. Kaplan and Violante (2010) shows that the methodology proposed by Blundell et al. (2008) to estimate *insurance coefficients* under the assumption of permanent

income shocks, is not additionally biased if the true data generating process only contains a highly persistent income shock. Karahan and Ozkan (2013) instead show empirical evidence of an age-dependent persistence of income shocks and discuss the effects of age-varying persistence in income shocks on the degree of self-insurance. Neither of these papers, however, provide an analysis of the implied bias in preference parameter estimates and simulated household behavior when solving and simulating household behavior from a model with a misspecified income shock process.

2 The theoretical framework

We rely on the canonical buffer-stock life-cycle model of Deaton (1991, 1992) and Carroll (1992, 1997). Below we present a version of the model which allows for transitory and persistent/permanent income shocks. Although this model is now one of the modern work horses in many fields of economics, we briefly present the model for completeness.

Households enter the labor market at period $t=1$, retire at the end of period T_R , and eventually die at the end of period T . The recursive form of the household problem is given by

$$V_t(P_t, L_t, M_t) = \max_{C_t \geq 0} \frac{C_t^{1-\rho}}{1-\rho} + \frac{1}{1+\phi} \mathbb{E}_t[V_{t+1}(P_{t+1}, L_{t+1}, M_{t+1})] \quad (1)$$

subject to an exogenous income process presented below, and the intertemporal budget constraint

$$A_t = M_t - C_t \quad (2)$$

$$M_{t+1} = RA_t + Y_{t+1} \quad (3)$$

where A_t is end-of-period assets, M_t is beginning-of-period market resources, Y_t is income, and R is the gross rate of return. Consumers are allowed to be net-borrowers up to a fraction of their permanent income P_t and end-of-period wealth, thus, has to satisfy

$$A_t \geq -\lambda_t P_t, \lambda_t = \begin{cases} 0 & t \geq T_R \\ \lambda & \text{else.} \end{cases} \quad (4)$$

Income process. In the beginning of each working period, households receive a stochastic income

$$Y_t = \Gamma_t L_t P_t \xi_t, \quad t \leq T_R \quad (5)$$

$$L_t = \begin{cases} 1 & \text{if } \alpha = 1 \\ L_{t-1}^\alpha \psi_t & \text{if } \alpha < 1 \end{cases} \quad (6)$$

$$P_t = \begin{cases} P_{t-1} \psi_t & \text{if } \alpha = 1 \\ P_{t-1} & \text{if } \alpha < 1 \end{cases} \quad (7)$$

$$\log \psi_t \sim \mathcal{N}(-0.5\sigma_\psi^2, \sigma_\psi^2) \quad (8)$$

where P_t is the permanent component of income, L_t is the persistent component of income with the AR(1) coefficient α , Γ_t is a normalization factor, ξ_t is a mean-one transitory shock to income, and ψ_t is a mean-one persistent ($\alpha < 1$) or fully permanent ($\alpha = 1$) shock to income. The transitory shock is given by³

$$\xi_t = \begin{cases} \mu & \text{with probability } p \\ (\epsilon_t - \mu p)/(1-p) & \text{with probability } 1-p \end{cases} \quad (9)$$

$$\log \epsilon_t \sim \mathcal{N}(-0.5\sigma_\xi^2, \sigma_\xi^2). \quad (10)$$

The normalization factor Γ_t is chosen to ensure that we always have $\mathbb{E}_0[Y_t] = (\Pi_{k=1}^t G_k) L_0 P_0$, where G_t is the deterministic age-dependent gross growth rate of income. To achieve this irrespectively of the degree of persistence of the income shocks we have⁴

$$\Gamma_t = \begin{cases} (\Pi_{k=1}^t G_k) e^{-\frac{1}{2} \left(\frac{1-\alpha^{2t}}{1-\alpha^2} - \frac{1-\alpha^t}{1-\alpha} \right)^2 \sigma_\psi^2} & \text{if } \alpha < 1 \\ \Pi_{k=1}^t G_k & \text{if } \alpha = 1 \end{cases} \quad (11)$$

In retirement, income is given by

$$Y_t = \kappa \Gamma_{T_R} L_{T_R} P_{T_R}, \quad t > T_R \quad (12)$$

where κ is the retirement replacement rate.

³ This specification of the transitory income shock follows that of Carroll (1997) and ensures that $\mathbb{E}_t[\epsilon_{t+1}] = 1$, regardless of the values of μ , p and σ_ξ^2 .

⁴ This is to ensure that the expected income profile is not affected by α such that when comparing consumer behaviors across different values of α , we only pick up differences due to the degree of persistence of shocks for a fixed expected income profile. See the supplemental material for the derivations.

Household behavior. The optimal consumption function is denoted by $C^*(P_t, L_t, M_t)$, and all unconstrained consumption choices satisfy the standard Euler equation,

$$C_t^{-\rho} = R \frac{1}{1+\phi} \mathbb{E}_t [C_{t+1}^{-\rho}] \quad (13)$$

For later use, the MPC is defined as

$$\text{MPC} \equiv \lim_{\Delta \downarrow 0} \frac{C^*(P_t, L_t, M_t + \Delta) - C^*(P_t, L_t, M_t)}{\Delta} \quad (14)$$

and the MPC out of permanent shocks (MPCP) is defined as

$$\text{MPCP} \equiv \lim_{\Delta \downarrow 0} \frac{C^*(P_t + \Delta, L_t, M_t + \Delta) - C^*(P_t, L_t, M_t)}{\Delta} \quad (15)$$

In the simulation exercises later, we calculate the MPC and MPCP using $\Delta=0.001$.

3 Calibration

Demographics. Households enter the labor market at age 25 ($t=1$), retire at age 60 ($T_R=35$) and die at age 80 ($T=55$).

Preferences. For our baseline results, we choose a relative risk aversion coefficient of $\rho=2$ and a discount rate of $\phi=0.03$. We explore the sensitivity of our results to alternative calibrations.

Borrowing/saving. For our baseline results we choose an interest of $R=\frac{1}{0.97}$, and restrict consumers to be holding non-negative net-wealth, $\lambda=0$.⁵

Income Process. The income growth rates are chosen to imply a standard concave income profile

$$G_t = \begin{cases} 1.08 & \text{if Age}_t \leq 30 \\ 1.05 & \text{if } 31 \leq \text{Age}_t \leq 35 \\ 1.03 & \text{if } 36 \leq \text{Age}_t \leq 45 \\ 1.01 & \text{if } 46 \leq \text{Age}_t \leq 50 \\ 1.00 & \text{if } 50 < \text{Age}_t \end{cases} \quad (16)$$

⁵ Households can receive a zero income shock with a positive probability, and the constraint will thus not be binding.

and the retirement replacement ratio is set to $\kappa=0.70$. Banks et al. (2010) report median replacement rates of 70 percent in the English Longitudinal Study of Ageing (ELSA).⁶ We set the transitory shock variance to $\sigma_\xi^2=0.01$, and allow for a low income shock of $\mu=0$ with probability $p=0.003$ (Gourinchas and Parker 2002).

In models with a permanent income shock, the variance of the permanent shock is often chosen targeting some measure of the cross-sectional dispersion in income. Changing the persistence of the shock, however, also changes the implied cross-sectional dispersion of log income for a given shock variance. Specifically, we have⁷

$$\text{Var}(\log(L_t P_t)) = \begin{cases} t\sigma_\psi^2 & \text{if } \alpha=1 \\ \frac{1-\alpha^{2t}}{1-\alpha^2}\sigma_\psi^2 & \text{if } \alpha<1 \end{cases}, \quad P_0=L_0=1 \quad (17)$$

We therefore adjust the variance of the persistent shock, when lowering α below one, to match the cross-sectional dispersion in log income at age 45 given by a standard assumption of a permanent shock variance of 0.01 when $\alpha=1$. Specifically, we choose a permanent shock variance of $\tilde{\sigma}_\psi^2=0.01$ and set

$$\sigma_\psi^2 = \begin{cases} \tilde{\sigma}_\psi^2 & \text{if } \alpha=1 \\ (k-25)\tilde{\sigma}_\psi^2 \frac{1-\alpha^2}{1-\alpha^{2(k-25)}} & \text{if } \alpha<1 \end{cases} \quad (18)$$

with $k=45$ controlling at what age the variance of log income is equalized across the parametrizations with persistent and permanent shocks. Particularly, note that $\sigma_\psi^2=\tilde{\sigma}_\psi^2$ for $k=26$, and that σ_ψ^2 is increasing in k and $\sigma_\psi^2>\tilde{\sigma}_\psi^2$ for all $k>26$.

Figure 1 shows the implied simulated age profiles for different α 's. Specifically, the left panel shows that the simulated income profile is *identical* across the different scenarios despite we impose different AR coefficients on the persistent shocks. This is insured through the specification of Γ_t in equation (11). The right panel shows the age profile of the variance of log income. Here, it is clear that the age profiles differ across the degrees of persistence, but that the variances are identical at age $k=45$ as implied by equation (18).

⁶ Blundell et al. (2008) experiment with replacement rates of zero and 50.

⁷ See the Supplemental material for the derivations.

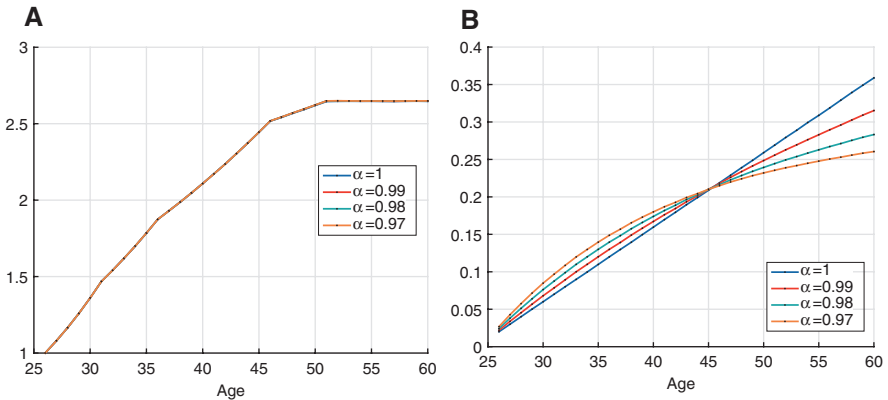


Figure 1: Income Age Profiles.

(A) Average Income Age Profiles, $\text{mean}(Y_t)$. (B) Variance of Log Income, $\text{var}(\log Y_t)$. Figure 1 shows age profiles of average income (left panel) and the cross-sectional variance of log income (right panel). All households are initialized with $P_0=1$ and $L_0=1$.

4 Results

We assume an *ideal* scenario in which the researcher observes the wealth of $N=50,000$ individuals in $T=34$ periods without any measurement error.⁸ Furthermore, the researcher knows the true values of all the remaining parameters except for either one of the preference parameters, ρ or ϕ , and the persistence of the income shock, α . We denote the parameter to be estimated as θ , which in turn can be either ρ or ϕ here.⁹

We assume that the researcher insists on assuming that the income process is permanent, $\alpha=1$, and calibrate the variance of the permanent income shock to achieve the same level of dispersion in log income as he or she observes in the data at age 45 (we discuss the robustness of our results to this assumption below). We will consider true values of $\alpha \in \{1, 0.99, 0.98, 0.97\}$, and use equation (18) to infer the implied shock variance σ_ψ^2 given our calibration of $\tilde{\sigma}_\psi^2=0.01$. Using equation (17) this implies, that the researcher, when choosing $\alpha=1$, and targeting the dispersion in log income at age 45, always sets $\sigma_\psi^2=0.01$.

⁸ This is to mitigate simulation and sampling error, which we will not focus on here.

⁹ The parameters ρ and ϕ are poorly identified jointly by the wealth moments used here (Cagetti 2003), and we thus abstract from joint estimation to minimize the influence of poor identification on our results. The poor identification using purely wealth moments stems from the fact that both higher risk aversion and lower discount rates generally both amplify wealth accumulation.

To emulate how applied researchers would likely go about estimating preference parameters using real data, we estimate θ by the method of simulated moments (Gouriéroux and Monfort 1997) and match the average wealth age profile from age 35 through 55.¹⁰ In turn, we estimate θ as

$$\hat{\theta} = \arg \min_{\theta} \left(\frac{1}{Q} \sum_{q=1}^Q \Lambda_q(\theta) - \Lambda \right)' W^{-1} \left(\frac{1}{Q} \sum_{q=1}^Q \Lambda_q(\theta) - \Lambda \right) \quad (19)$$

where Λ contains the moments from the data. We perform $Q=16$ simulations of the data for each trial value of θ , and as weighting matrix, W , we use a diagonal matrix with bootstrapped variances of the moments in the data on the diagonal.¹¹ We perform $S=100$ individual Monte Carlo estimations of θ , denoted $\hat{\theta}_s$, and Table 1 reports the bias in percentage terms, $\text{BIAS}(\hat{\theta}) \equiv (\text{MEAN}(\hat{\theta})/\theta - 1) \cdot 100$.

Table 1: Monte carlo estimation results. $\text{BIAS}(\hat{\theta}) \equiv (\text{MEAN}(\hat{\theta})/\theta - 1) \cdot 100$.

α :	Estimating $\theta=\rho$				Estimating $\theta=\phi$			
	1.00	0.99	0.98	0.97	1.00	0.99	0.98	0.97
$\tilde{\sigma}_{\psi}$:	0.10	0.10	00.10	00.10	0.10	0.10	0.10	0.10
σ_{ψ} :	0.10	0.11	0.12	0.13	0.10	0.11	0.12	0.13
$\rho=1.50, \phi=0.05$	-0.0	4.2	11.1	20.1	0.0	-3.3	-8.0	-13.0
$\rho=2.00, \phi=0.05$	-0.0	4.2	11.1	19.8	0.0	-4.0	-9.7	-15.6
$\rho=2.50, \phi=0.05$	-0.0	4.1	10.5	18.4	0.0	-4.7	-11.1	-17.8
$\rho=5.00, \phi=0.05$	-0.0	2.8	6.6	10.8	0.0	-6.8	-15.1	-23.6
$\rho=1.50, \phi=0.03$	-0.0	7.7	21.7	38.3	0.0	-2.8	-7.2	-11.9
$\rho=2.00, \phi=0.03$	-0.0	5.8	15.9	28.0	0.0	-3.6	-9.2	-15.1
$\rho=2.50, \phi=0.03$	-0.0	4.7	12.6	21.9	0.0	-4.4	-10.9	-18.0
$\rho=5.00, \phi=0.03$	-0.0	2.7	6.4	10.5	0.1	-8.3	-18.4	-29.1
$\rho=1.50, \phi=0.01$	0.0	-1.1	-3.3	-5.9	0.0	-1.5	-4.6	-8.9
$\rho=2.00, \phi=0.01$	0.1	-4.3	-9.6	-14.7	0.1	-3.4	-9.4	-17.0
$\rho=2.50, \phi=0.01$	0.0	-13.8	-21.9	-27.8	0.1	-5.5	-14.3	-25.1
$\rho=5.00, \phi=0.01$	-0.0	2.8	6.4	10.5	0.2	-17.8	-39.0	-62.0

Table 1 reports Monte Carlo estimation results from $S=100$ independent simulated data sets consisting of $N=50,000$ individuals observed in $T=34$ periods. The bias is defined in percentage terms as $\text{BIAS}(\hat{\theta}) \equiv (\text{MEAN}(\hat{\theta})/\theta - 1) \cdot 100$.

10 Euler equation estimation can potentially provide unbiased estimates of ρ . If consumers are not affected by borrowing constraints and consumption is not contaminated with measurement error, ρ can be uncovered independently of the underlying income process using interest rate variation and the (exact) Euler equation. Alan and Browning (2010) provide exact Euler equation estimators that allow for multiplicative log-normal measurement error in consumption.

11 We use eight unique draws of income, and their eight antithetic draws.

4.1 Preference estimates

For a true $\alpha=1$, the model is correctly specified, and as expected we find $\text{BIAS}(\hat{\theta}) \approx 0$. For a true $\alpha < 1$, however, we find significant biases in the estimates of ρ and ϕ due to the incorrectly specified income process. Considering the baseline case of $\rho=2$ and $\phi=0.03$, we find that a true $\alpha=0.97$ implies a bias of 28 percent when estimating ρ , and a bias of 15 percent when alternatively estimating ϕ . The bias remains large and significant for alternative true values of ρ and ϕ , but vary substantially across the various parametrizations. The absolute bias is always strictly increasing in the size of the misspecification, $|1-\alpha|$, and when estimating ϕ the bias is always found to be negative, while for ρ the bias is positive unless the true ϕ is small where the bias in ρ becomes negative.

The explanation behind the found biases is that setting $\alpha=1$, when in truth $\alpha < 1$, has two opposite effects on the level of income risk in the estimated model relative to the true model. Firstly, the level of risk is *increased* because the shocks are more persistent, and secondly it is *decreased* because the shock variance is calibrated to ensure equality of income variances at age k (see equation 18). For our baseline choice of $k=45$, the second effect dominates; the under-stated income risk implies an under-accumulation of wealth, which when e.g. estimating the discount rate ϕ can only be countered by a lower discount rate, i.e. a higher discount factor, implying a negative bias in ϕ .

The sign of the bias when increasing ρ is more complicated because the effect from ρ on wealth accumulation runs through both the risk aversion channel and the inter-temporal substitution channel. The risk aversion channel always implies that a higher ρ amplifies the motive for wealth accumulation (precautionary saving), but the inter-temporal substitution channel can both decrease and increase wealth accumulation. Considering the deterministic case, the Euler equation (13) implies

$$\frac{C_{t+1}}{C_t} = \left(R \frac{1}{1+\phi} \right)^{\frac{1}{\rho}}$$

If ϕ is small such that $R \frac{1}{1+\phi} > 1$, increasing ρ flattens the steeply increasing consumption age profile. This logic extends to the stochastic case, and the inter-temporal substitution channel can thus imply that increasing ρ increases current consumption, and reduces wealth accumulation. For small enough ϕ , this can even dominate the risk aversion channel, and the bias in ρ thus becomes negative

to counter the under-accumulation of wealth implied by the misspecified income process. In Table 1 we see that when $\rho \in \{1.5, 2, 2.5\}$ the inter-temporal substitution channel becomes dominating at $\phi=0.01$, while the risk aversion channel continues to dominate at $\phi=0.01$ when $\rho=5$.

As we will show below in Table 4, the sign of the bias is also sensitive to our choice of k . For low k there is an overall *increase* in the level of risk due to the smaller adjustment of the income shock variance. This implies that the biases in ρ and ϕ must now instead imply less wealth accumulation changing their sign. The cut-off level of k , where there is an alternation in the sign of the bias, is furthermore affected by the choice of moments; this is illustrated in the supplemental material in the case of matching median levels of wealth rather than mean levels of wealth as done here.

Our results are related to the analysis in Kaplan and Violante (2010).¹² When lowering the persistence of the income shock, they adjust the income shock variance to keep the overall increase in the dispersion of log income from age 25 to age 60 constant (i.e. $k=60$ in our terminology). Simultaneously, they lower the discount factor, $\beta \equiv \frac{1}{1+\phi}$, to also match a constant wealth-to-income ratio of 2.5. This indicates that if they had estimated ϕ using a misspecified model with $\alpha=1$, they would have found a negative bias in ϕ (higher estimated β , lower estimated ϕ). This aligns with our results (for large enough k).

4.2 Implied household behavior

To study the effects of the misspecification on implied household behavior, we simulate panels of households from both the true model and the estimated model (with $\alpha=1$ and θ at its average estimated levels in Table 1). Figure 2 reports the *differences* between the estimated and the true models in the resulting MPC (left panels) and MPCPs (right panels) over the life-cycle.¹³ The differences in the average MPCs and MPCPs between the estimated and true models are shown in Tables 2 and 3.

The implied discrepancies in the MPCs and MPCPs are rather small compared to the relative large estimation biases in ρ and ϕ . The discrepancies are largest for the MPCs, where the fall over the life cycle happens a bit earlier

¹² See in particular their Section 5 and Table 5.

¹³ The levels are reported in Appendix Figure A.4.

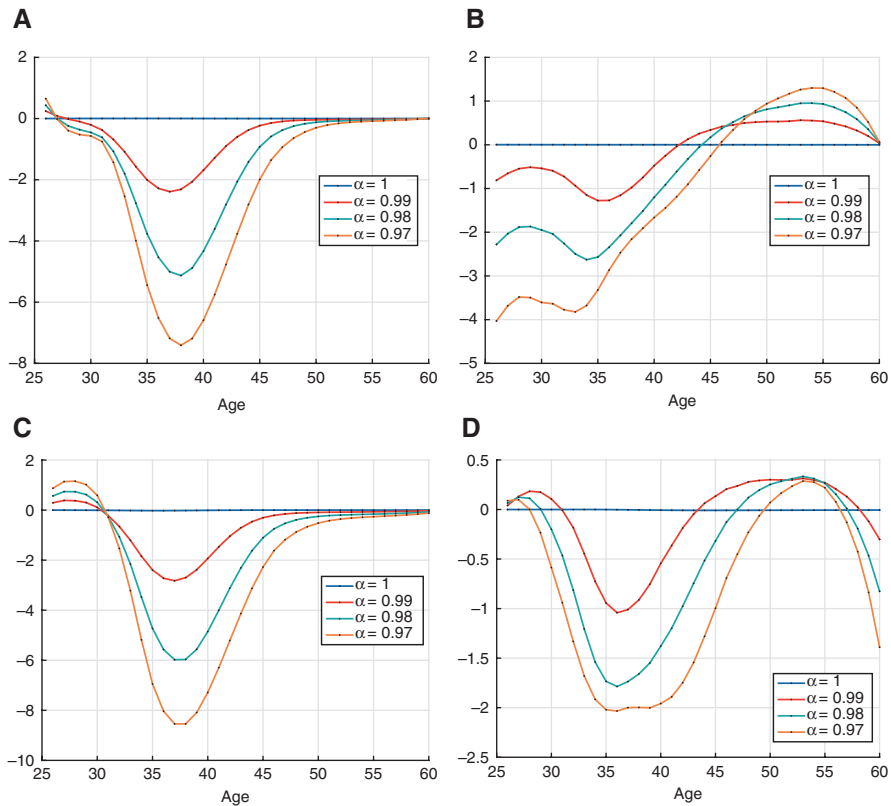


Figure 2: MPC and MPCP Age Profiles ($\rho=2$, $\phi=0.03$).

(A) MPC (estimating ρ). (B) MPCP (estimating ρ). (C) MPC (estimating ϕ). (D) MPCP (estimating ϕ). Figure 2 shows average age profiles of 500,000 simulated households. The differences between the estimated and the true model are shown. All households are initialized without any wealth.

when the persistence is misspecified. Setting $\rho=2$ and $\phi=0.03$, and taking the worst considered case of $\alpha=0.97$, the average MPC for the working age population is only 2.1 percentage points lower than its true value when estimating ρ , and only 2.3 percentage points lower when estimating ϕ . In relative terms this is less than ten percent as the average MPC for these preferences is about 30 percent. Still focusing on $\rho=2$ and $\phi=0.03$, the difference in the MPCP between the estimated model and the true model is even smaller; the MPCP is thus just 1.1 percentage points lower when estimating ρ , and 0.8 percentage points lower when estimating ϕ . With an average true MPCP of about 91, these deviations are minuscule.

Table 2: Differences in MPC.

α :	Estimating $\theta=\rho$				Estimating $\theta=\phi$			
	1.00	0.99	0.98	0.97	1.00	0.99	0.98	0.97
$\rho=1.50, \phi=0.05$	0.0	-0.9	-2.0	-3.1	-0.0	-1.0	-2.3	-3.5
$\rho=2.00, \phi=0.05$	0.0	-0.7	-1.6	-2.5	-0.0	-0.8	-1.9	-3.0
$\rho=2.50, \phi=0.05$	0.0	-0.6	-1.4	-2.2	-0.0	-0.7	-1.7	-2.6
$\rho=5.00, \phi=0.05$	0.0	-0.4	-0.9	-1.4	-0.0	-0.5	-1.1	-1.6
$\rho=1.50, \phi=0.03$	0.0	-0.7	-1.6	-2.5	-0.0	-0.7	-1.7	-2.6
$\rho=2.00, \phi=0.03$	0.0	-0.6	-1.4	-2.1	-0.0	-0.6	-1.5	-2.3
$\rho=2.50, \phi=0.03$	0.0	-0.5	-1.2	-1.8	-0.0	-0.6	-1.4	-2.1
$\rho=5.00, \phi=0.03$	0.0	-0.3	-0.8	-1.2	-0.0	-0.4	-0.9	-1.4
$\rho=1.50, \phi=0.01$	0.0	-0.4	-0.9	-1.5	-0.0	-0.4	-0.9	-1.4
$\rho=2.00, \phi=0.01$	0.0	-0.4	-1.0	-1.6	-0.0	-0.4	-0.9	-1.5
$\rho=2.50, \phi=0.01$	0.0	-0.5	-1.1	-1.7	-0.0	-0.4	-0.9	-1.5
$\rho=5.00, \phi=0.01$	0.0	-0.3	-0.6	-1.0	-0.0	-0.3	-0.8	-1.2

Table 2 shows the difference in the average MPC between the true and estimated models using 500,000 simulated households in each case. All households are initialized without any wealth.

Table 3: Differences in MPCP.

α :	Estimating $\theta=\rho$				Estimating $\theta=\phi$			
	1.00	0.99	0.98	0.97	1.00	0.99	0.98	0.97
$\rho=1.50, \phi=0.05$	0.0	-0.6	-1.5	-2.4	-0.0	-0.6	-1.4	-2.2
$\rho=2.00, \phi=0.05$	0.0	-0.4	-1.2	-1.9	-0.0	-0.4	-1.0	-1.6
$\rho=2.50, \phi=0.05$	0.0	-0.2	-0.9	-1.4	-0.0	-0.2	-0.7	-1.0
$\rho=5.00, \phi=0.05$	0.0	0.4	0.5	0.5	-0.0	0.5	0.6	0.8
$\rho=1.50, \phi=0.03$	0.0	-0.3	-1.0	-1.6	-0.0	-0.3	-0.8	-1.3
$\rho=2.00, \phi=0.03$	0.0	-0.2	-0.6	-1.1	-0.0	-0.1	-0.5	-0.8
$\rho=2.50, \phi=0.03$	0.0	-0.0	-0.4	-0.7	-0.0	0.0	-0.2	-0.4
$\rho=5.00, \phi=0.03$	0.0	0.5	0.6	0.7	-0.0	0.5	0.7	0.9
$\rho=1.50, \phi=0.01$	0.0	0.1	0.0	-0.1	-0.0	0.1	0.0	-0.1
$\rho=2.00, \phi=0.01$	0.0	0.2	0.1	0.1	-0.0	0.2	0.1	-0.0
$\rho=2.50, \phi=0.01$	-0.0	0.3	0.4	0.4	-0.0	0.2	0.2	0.2
$\rho=5.00, \phi=0.01$	0.0	0.5	0.7	0.9	-0.0	0.5	0.8	1.0

Table 3 shows the difference in the average MPCP between the true and estimated models using 500,000 simulated households in each case. All households are initialized without any wealth.

The underlying reason behind these small discrepancies is that the average wealth and consumption age profiles are matched rather well in the estimation (see the supplemental material), and the changes in the shape of the consumption function implied by the biased preferences is then of second order. The

discrepancies also remain small for alternative values of ρ and ϕ , but the absolute differences are generally larger for smaller values of ϕ , where the true MPC and MPCP are also larger.

4.3 Robustness

Table 4 reports a number of robustness checks varying, respectively the age at which the variances of income are equalized (k) and the underlying variance of the permanent/persistent shock $\tilde{\sigma}_{\psi}^2$ (see equation 18).

Across all the considered cases, the discrepancies in the average MPC and the average MPCP remain small.

The results for the preference estimates are more complicated. The bias alternates in sign depending on k , which determines the size of the calibrated adjustment in the income shock variance, when the persistence of the income process is misspecified. For *low* values of k there is an *over*-statement of income risk (when assuming $\alpha=1$) inducing a counteraction *reducing* wealth accumulation (lower ρ , or higher ϕ). For high values of k there is an *under*-statement of income risk inducing a counteraction *increasing* wealth accumulation (higher ρ , or lower ϕ). The bias is only negligible for some unknown intermediate level of k . Changing the baseline persistent income shock variance, $\tilde{\sigma}_{\psi}^2$, to either a rather *high* level (0.03) or a *rather low* level (0.005), reduce the bias a bit, but it remains sizable.

In the supplemental material, we additionally show that our results are robust to matching the *median* wealth age profiles instead of average wealth age profiles.

Table 4: Robustness.

	BIAS($\hat{\theta}$)		DIFF(MPC)		DIFF(MPCP)	
	$\theta=\rho$	$\theta=\phi$	$\theta=\rho$	$\theta=\phi$	$\theta=\rho$	$\theta=\phi$
Baseline	28.0	-15.1	-2.1	-2.3	-1.1	-0.8
$k=26$	-19.0	12.6	-2.3	-2.0	-0.2	-0.3
$k=35$	0.6	-0.1	-2.2	-2.2	-0.6	-0.6
$k=55$	60.8	-30.6	-1.9	-2.4	-1.6	-0.9
$\tilde{\sigma}_{\psi}^2=0.005$	20.7	-11.3	-1.0	-1.5	-1.0	-0.7
$\tilde{\sigma}_{\psi}^2=0.02$	18.0	-13.9	-2.4	-2.6	-0.2	-0.1
$\tilde{\sigma}_{\psi}^2=0.03$	12.2	-11.5	-2.1	-2.3	0.6	0.6

Baseline parameters, ($\alpha=0.97, \rho=2, \phi=0.03$). See notes to Table 1. DIFF(MPC) is the difference between the average marginal propensity to consume (MPC) in the estimated and the true model. DIFF(MPCP) is the difference between the average marginal propensity to consume out of *permanent* shocks (MPCP) in the *estimated* and the *true* model.

5 Concluding discussion

We have investigated how severe a bias in preference estimates, and associated household behavior, researchers face when assuming a *fully permanent* income shock when the true data generating process only contains a *highly persistent* income shock. We find a relatively large bias in the estimation of preferences from even small and moderate misspecification of the income process. This suggests that if researchers care about the exact value of preference parameters, assuming a unit root in income might seriously affect their results. We also find the more positive result that the estimation bias in preferences seem to somewhat cancel out the effect of the misspecified income process when looking at household behavior such as the MPC or the marginal propensity to consume out of permanent income shocks (MPCP). In turn, if researchers care more about the economic behavior of consumers (compared to parameter estimates), the misspecification might not affect results too much, and the computational gain from assuming a unit root income process might outweigh the error.

While we have, for ease of exposition, focused on CRRA preferences, an interesting avenue for future research is to investigate how conclusions would be if consumers had habits over consumption or Epstein-Zin preferences. In both cases, risk aversion could be separated from the inter-temporal elasticity of substitution.

Acknowledgments: We thank the managing editor and two anonymous referees for their constructive comments. We also wish to thank Chris Carroll, Mette Ejrnæs and Greg Kaplan for helpful comments. Financial support from the Danish Council for Independent Research in Social Sciences (FSE, grant no. 5052-00086B and 4091-00040) is gratefully acknowledged. Part of this research was carried out while Jørgensen was visiting Princeton University in the fall 2015. Jørgensen thanks Bo E. Honoré for his exceptional hospitality and effort in organizing the stay.

References

- Abowd, J. M., and D. Card. 1989. "On the Covariance Structure of Earnings and Hours Changes." *Econometrica* 57 (2): 411–445.
- Alan, S., and M. Browning. 2010. "Estimating Intertemporal Allocation Parameters using Synthetic Residual Estimation." *The Review of Economic Studies* 77 (4): 1231–1261.
- Baker, M. 1997. "Growth-Rate Heterogeneity and the Covariance Structure of Life-Cycle Earnings." *Journal of Labor Economics* 15 (2): 338–375.

- Banks, J., C. O'Dea, and Z. Oldfield. 2010. "Cognitive Function, Numeracy and Retirement Saving Trajectories." *The Economic Journal* 120 (548): F381–F410.
- Blundell, R., L. Pistaferri, and I. Preston. 2008. "Consumption Inequality and Partial Insurance." *The American Economic Review* 98 (5): 1887–1921.
- Blundell, R., H. Low, and I. Preston. 2013. "Decomposing Changes in Income Risk Using Consumption Data." *Quantitative Economics* 4 (1): 1–37.
- Browning, M., M. Ejrnæs, and J. Alvarez. 2010. "Modelling Income Processes with Lots of Heterogeneity." *The Review of Economic Studies* 77 (4): 1353–1381.
- Cagetti, M. 2003. "Wealth Accumulation Over the Life Cycle and Precautionary Savings." *Journal of Business & Economic Statistics* 21 (3): 339–353.
- Carroll, C. D. 1992. "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence." *Brookings Papers on Economic Activity* 2: 61–156.
- Carroll, C. D. 1997. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *The Quarterly Journal of Economics* 112 (1): 1–55.
- Deaton, A. 1991. "Saving and Liquidity Constraints." *Econometrica* 59 (5): 1221–1248.
- Deaton, A. 1992. *Understanding Consumption*. Oxford: Oxford University Press.
- Gouriéroux, C., and A. Monfort. 1997. *Simulation-based Econometric Methods*. New York, NY: Oxford University Press.
- Gourinchas, P.-O., and J. A. Parker. 2002. "Consumption Over the Life Cycle." *Econometrica* 70 (1): 47–89.
- Güvenen, F. 2009. "An Empirical Investigation of Labor Income Processes." *Review of Economic Dynamics* 12 (1): 58–79.
- Güvenen, F., F. Karahan, S. Ozkan, and J. Song. 2015. "What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk?" NBER Working Paper 20913.
- Kaplan, G., and G. L. Violante. 2010. "How Much Consumption Insurance Beyond Self-Insurance?" *The American Economic Journal* 2: 53–87.
- Karahan, F., and S. Ozkan. 2013. "On the Persistence of Income Shocks Over the Life Cycle: Evidence, Theory, and Implications." *Review of Economic Dynamics* 16 (3): 452–476.
- Lillard, L. A., and R. J. Willis. 1978. "Dynamic Aspects of Earning Mobility." *Econometrica* 46 (5): 985–1012.
- Lillard, L. A., and Y. Weiss. 1979. "Components of Variation in Panel Earnings Data: American Scientists 1960–70." *Econometrica* 47 (2): 437–454.
- MacCurdy, T. E. 1982. "The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis." *Journal of econometrics* 18 (1): 83–114.
- Meghir, C., and L. Pistaferri. 2004. "Income Variance Dynamics and Heterogeneity." *Econometrica* 72 (1): 1–32.

Supplemental Material: The online version of this article (DOI 10.1515/bejm-2016-0035) offers supplementary material, available to authorized users.