Income and Consumption: a Micro Semi-structural Analysis with Pervasive Heterogeneity[1]

Sule Alan
Department of Economics
University of Essex, UK.
salan@essex.ac.uk

Martin Browning
Department of Economics
University of Oxford, UK.
martin.browning@economics.ox.ac.uk

Mette Ejrnæs
Department of Economics
University of Copenhagen, Denmark.
mette.ejrnes@econ.ku.dk

April 2016

[1]We would like to thank the editor, the referees, seminar participants at the University of Chicago, Yale University, NYU, University of Pennsylvania and University of Michigan for valuable comments and suggestions. We would also like to thank Jim Ziliak and Preston Brasher for their considerable help in setting up the PSID sample we use in the paper. All errors are our own.

Abstract

We develop a model of consumption and income that allows for pervasive heterogeneity in the parameters of both processes. Introducing co-dependence between household income parameters and preference parameters, we also allow for heterogeneity in the impact of income shocks on consumption. We estimate the parameters of the model using a sample from the PSID, covering the period 1968 to 2009. We find considerable co-dependent heterogeneity in the parameters governing income and consumption processes. Our results suggest a great deal of heterogeneity in the reaction of consumption to income shocks, highlighting the heterogeneity in the self-insurance available to households.

Keywords: preference heterogeneity; consumption; income.

JEL Classification: C33, D12, D31, J31

1 Introduction

An understanding of the relationship between income and consumption is crucial to many research and policy debates. Examples include: the efficacy of fiscal policy; the design of social insurance mechanisms; the determinants of saving over the short run and the long run and the tax treatment of different sources of income. In this paper we consider the income-consumption relationship with a particular focus on heterogeneity. It is now well established that households have highly idiosyncratic income processes and preference parameters (so called 'pervasive heterogeneity'). Moreover, the within process heterogeneity in these parameters is correlated ('co-dependence'). For example, for consumption, the discount rate and the coefficient of risk aversion may be correlated. We begin by providing two examples to motivate the potential importance of pervasive and co-dependent heterogeneity.

For our first example, consider a normative analysis of the benefits of unemployment insurance (UI). If there is limited heterogeneity then the benefits of UI will be much the same for everyone. If, for example, earnings variances and risk aversion are heterogeneous then the benefits of UI will also be heterogeneous and will be increasing in both parameters. Moreover, a positive correlation between these parameters will reinforce the heterogeneity in the benefits of UI, with some households benefiting a great deal and some very little. Such considerations could have a substantial impact on theoretical analyses of social insurance such as Huggett and Parra (2010).

Our second example is the positive analysis of the efficacy of fiscal stimulus policies which depends on the impact of income (shocks) on consumption. Introducing co-dependent heterogeneity in the reaction to an income shock can fundamentally change both theoretical and empirical analyses in this regard. Blundell *et al* (2008) model income and consumption simultaneously with limited allowance for heterogeneity in preferences and the income process. They find partial insurance against permanent income shocks but almost full insurance against transitory income shocks. Kaplan and Violante (2010) show that in a standard life-cycle model, the consumption reaction to an income shock depends on preference parameters such as risk aversion and on the parameters of the income process. This implies that heterogeneity in either preference parameters or income parameters will introduce heterogeneity in the consumption reaction to an income shock, and hence lead to heterogeneity in the degree of self-insurance available to households.

There are three main contributions in this paper. The first is that we introduce co-dependence between household income process parameters and preference parameters. The need for this is illustrated by the first example above. The second contribution is to allow for heterogeneity in the impact of income shocks on consumption, which is motivated by the second example above. The third contribution is a methodological contribution in that we provide a framework that allows us to discuss theoretically and empirically quantify the extent of co-dependent heterogeneity. This semi-structural parametric framework is sufficiently flexible to allow that household level heterogeneity in both preference parameters and income process parameters can be accounted for under a standard intertemporal consumption and saving model.

The presence of co-dependent heterogeneity in the earnings process has been documented in a number of studies (see Baker (1997), Rubinstein and Weiss (2006), Guvenen (2009) and Browning *et al* (2010)). Similarly, the presence of pervasive heterogeneity in intertemporal preference parameters is now well established. The experimental economics literature documents a great deal of preference heterogeneity. This literature offers reliable ways to elicit intertemporal allocation parameters via decision tasks given to individuals, often using real stakes; see Gneezy and Potter (1997), Holt and Laury (2002) for elicitation of risk aversion and Andersen *et al* (2006) and Andreoni and Sprenger (2012) for elicitation of time preferences. In an alternative approach, using household level consumption growth information, Alan and Browning (2010) estimate the joint distribution of discount rates and coefficients of relative risk aversion, and illustrate a large degree of heterogeneity in intertemporal preferences.

To the best of our knowledge, there is no study that models income and consumption jointly by allowing for pervasive heterogeneity in individual level parameters and co-dependence across parameters of both processes.[1] The need for such modelling is intuitively plausible. For example, patient individuals may select into jobs that have a high earnings growth rate, which partially motivated the framework developed in Mincer (1958). Equally plausible is that more risk averse individuals select into jobs with a low variance in earnings. As emphasized by Cunha *et al* (2005), the exact relationship between preferences and education and career choices will depend on the range of available earnings processes and on the environmental possibilities for shifting allocations across time and states. Cadena and Keys (2015) provide the most recent evidence on the link between time preferences and educational choices that impact on earnings processes. With regard to risk aversion, Bonin *et al* (2007) and Skriabikova *et al* (2012) provide evidence that self-reported risk aversion measures correlate with the earnings risk of chosen occupations.

In section 2, we develop a semi-structural parametric specification for the joint income and consumption processes for a given household. For household income, we follow Browning *et al* (2010) and specify a standard ARMA model with five parameters for each household. For consumption, we follow Alan and Browning (2010) and employ a semi-structural iso-elastic exact Euler equation approach. At the household level the only direct link between the two processes is that the consumption shock depends in part on the contemporaneous income shock. We then develop a parametric factor structure to capture heterogeneity across households. In doing this, we allow for co-dependence between all of the income and consumption parameters. In section 3 we describe the longitudinal consumption and income information in the Panel Studies of Income Dynamics (PSID). In section 4 we present a simulation based estimation procedure (indirect inference) that requires simulation of the fully parametric model. Indirect

[1] Blundell et al (2008) and Guvenen and Smith (2014) model consumption and income together. The former examine the link between consumption and income inequality. The latter use consumption information to pin down income process parameters. Neither study allows for pervasive heterogeneity.

inference is based on the construction of auxiliary parameters (ap's) that are matched between the actual data and the data from the simulated model. In generating ap's we rely heavily on regressions of income and consumption growth for individual households. This delivers a very rich empirical description of the two processes and their co-dependence.

The results are presented in section 5. We find considerable heterogeneity in the parameters of income and consumption processes, and, the main point of this paper, co-dependence between the parameters governing the two processes. With regard to quantifying the importance of income shocks for each household, our results suggest that even if all households face a trend stationary income process, for some households the long run effect of an instantaneous shock can be quite large and even a temporary shock in net income will result in a significant loss of life-time income. Offsetting this, we find a strong negative correlation between the income variance and the importance of income shocks on life time income. On the consumption side, our estimated first decile, median and ninth decile values of the discount rate and the coefficient of relative risk aversion are [4.9%, 7.7%, 9.2%] and [2.3, 7.48, 12.8] respectively. Moreover, they are positively correlated, implying that impatient households are more risk averse. With respect to co-dependence between preference and income parameters, we find that patient households have higher trends in income from age 30. However, we do not find any correlation between risk aversion and the variance of the idiosyncratic income process.

We also find that the reaction of consumption to income shocks is heterogeneous. We estimate that a 10% income shock raises consumption by a modest 1.9% for the median household. At the top end of this exposure distribution (the ninth decile) the value is 6.8%, which indicates that even those who react most can still achieve considerable consumption smoothing. We also find evidence of co-dependence; those with low risk aversion and/or a low variance of income shocks react more to income shock. Not surprisingly, we also find that households with more persistent income shocks have a bigger consumption reaction to an income shock.

A broad implication of our results is that pervasive co-dependent heterogeneity in income processes and preference parameters requires a more comprehensive approach to policy and welfare evaluation problems. Representative agent models and estimation of average treatment effects may not be appropriate when agents are ex-ante heterogeneous; see Heckman (2001). The new dynamic public finance, see Kocherlakota (2010) and Farhi and Werning (2012), puts heterogeneity and uncertainty over future earnings at the heart of the analysis. As another example, policies that involve delayed incentives (such as tax-deferred accounts to promote savings) may have a differential impact on individuals under discount rate heterogeneity. For example, under discount rate heterogeneity, the Atkinson and Stiglitz (1976) result that "savings should not be taxed" does not hold; see Saez (2002), Banks and Diamond (2008) and Diamond and Spinnewijn (2011) who show that a small saving tax is welfare improving under discount rate heterogeneity. These examples (and the two given above) highlight the importance of allowing for pervasive heterogeneity in theoretical and empirical analyses of the relationship between income and consumption.

2 Theoretical specification

2.1 The income process

For the dynamic specification of household income, we assume that log household income at age t for household h, y_{ht}, can be modelled as a general $ARMA(1,1)$ process with a linear trend. For each household the log income process is:

$$\begin{aligned} y_{ht} &= \{\mu_h(1-\rho_h)+\alpha_h\rho_h\}+\rho_h y_{h,t-1}+(1-\rho_h)\alpha_h(t-1) \\ &\quad +\nu_h\left(\xi_{ht}+\theta_h\xi_{h,t-1}\right) \end{aligned} \tag{1}$$

with $\xi_{ht}\sim N(0,1)$. The parameters μ_h and α_h capture the initial level and the trend respectively; ρ_h and θ_h determine the dynamics of the process where the AR parameter $\rho_h\in(0,1)$ captures the long run dynamics and the MA parameter $\theta_h\in(-1,1)$ captures the short run dynamics. Finally ν_h is the standard deviation of the income shock.

Even though we assume stationarity, a shock does have a positive impact on consumption through the consequent revision of future lifetime income. For future reference, we define the *long run cumulative impact of a shock*, denoted by the household specific parameter τ_h, in the standard fashion[2] as:

$$\tau_h = \frac{1+\theta_h}{1-\rho_h} \tag{2}$$

In the subsequent analysis we shall relate consumption changes to income shocks; the value of τ_h has an immediate impact on this since a higher value for τ_h implies that the revision of lifetime wealth that drives the consumption change is higher.

The formulation given in (1) allows each household to have its own set of parameters $\{\mu_h, \alpha_h, \rho_h, \theta_h, \nu_h\}$. Furthermore, we shall allow that these parameters are co-dependent. For example, as well as allowing heterogeneity in the long run impact of an income shock, τ, and the variance of the shocks, ν, it may be that the two are correlated with, say, high variance households facing more persistent shocks.

Although the model is fairly general it does impose strong assumptions. First, it assumes that there are no common macro shocks to the income process. Second, all parameters are are assumed to be time and age invariant. The latter precludes, for example, learning about the income process (as in Guvenen and Smith 2014) or that the variance of the shocks has changed over time (as in Moffitt and Gottschalk (2012)) or over age (as in Blundell *et al* (2015)). These assumptions are necessary to achieve the main objective of the paper, which is allowing for pervasive and co-dependent heterogeneity in *all* parameters. A challenge for future research in this area is to investigate whether features such as time varying variances are necessary if we allow for pervasive heterogeneity.

To model initial conditions we impose the stationarity conditions while allowing for nonsta-

[2] Strictly speaking, we should allow that the process is finite so that the value depends on age. However the approximation is good so long as the remaining lifetime is not too short. Note also that since we use a model of log income, this expression does not directly measure the impact on life time income.

tionarity of the distribution (Arellano (2003)). Specifically we set:

$$y_{h1} = b_0 + (\mu_h(1-\rho_h)) + \alpha_h\rho_h + \exp(b_1)\,\nu_h\left(\xi_{h1} + \frac{\theta_h+\rho_h}{\sqrt{1-\rho_h^2}}\xi_{h0}\right) \tag{3}$$

where $\xi_{h0} \sim N(0,1)$. Note that $(b_0, b_1) = (0,0)$ implies a stationary distribution.

2.2 Consumption

To model consumption we use the standard intertemporal consumption and saving model and specify the consumption process based on the exact Euler equation. Our specification imposes a number of assumptions on the process, which are fairly standard in the consumption literature. First, it treats the household as a unit which implies that husband and wife are assumed to have the same intertemporal allocation parameters. Second, it ignores liquidity constraints and imperfect markets. Third, it does not allow for cross-sectional heterogeneity in the real interest rate r_t. In the empirical analyses we specify quasi-Lagrange multiplier (QLM) tests which are designed to test for violations of the assumption of no liquidity constraints.

We take an iso-elastic utility function, which leads to the consumption Euler equation for household h:

$$E_t\left[\frac{1+r_{t+1}}{1+\delta_h}\left(C_{h,t+1}\right)^{-\gamma_h}\right] = \left(C_{ht}\right)^{-\gamma_h} \tag{4}$$

where δ_h is the discount rate; γ_h is the coefficient of relative risk aversion; r_{t+1} is the real interest rate between periods t and $t+1$; $E_t(.)$ is the expectations operator conditional on information available at time t. Equation (4) can be written as:

$$\left(\frac{1+r_{t+1}}{1+\delta_h}\right)\left(\frac{C_{h,t+1}}{C_{ht}}\right)^{-\gamma_h} = \varepsilon_{h,t+1} \tag{5}$$

where $\varepsilon_{h,t+1}$ is a shock to the marginal utility of expenditure (mue) and $E_t(\varepsilon_{h,t+1}) = 1$. Given initial values, C_{h1}, and values $(\varepsilon_{h2}, \varepsilon_{h3}, ...)$ we recursively define consumption by:

$$C_{h,t+1} = C_{ht} \left\{ \left(\frac{1+\delta_h}{1+r_{t+1}} \right) \varepsilon_{h,t+1} \right\}^{-(\gamma_h)^{-1}} \tag{6}$$

2.3 Consumption shocks

To estimate the structural parameters, we need individual consumption paths, simulated using the solution of the full structural model. In our case of pervasive individual level heterogeneity, this implies solving numerous life cycle models with a large number of state variables using numerous combinations of preference and income process parameters, including their co-dependence. Instead, we follow Alan and Browning (2010) and employ *synthetic residuals* to simulate consumption paths using equation (6).[3] This is based on the finding in Alan and Browning (2010) that for a heterogeneous population, the distribution of shocks to the marginal utility of expenditure (mue) is well approximated by a mixture of two log-normals.

Here, we extend Alan and Browning (2010) by decomposing the mue shock into an income shock and a non-income shock. We define the total mue shock as:

$$\varepsilon_{ht} = \tilde{\varepsilon}_{ht} \breve{\varepsilon}_{ht} \tag{7}$$

where $\tilde{\varepsilon}_{ht}$ is the non-income shock and $\breve{\varepsilon}_{ht}$ is the income shock. The two types of shocks are assumed to be independent and each to have a unit mean. To model these two types of shocks, we introduce some extra parameters. Unlike the income and preference parameters, these parameters do not have a structural interpretation.

[3] We have carried out a Monte Carlo experiment to validate the SRE method. In the Monte Carlo experiment, we simulate income processes and we find the optimal consumption path using a standard life cycle model. The Monte Carlo experiment shows that the SRE method is able to recover the joint distribution of preference and income parameters. However, since conventional Euler equation methods cannot accommodate pervasive heterogeneity, we are not able to compare the performance of our method with that of those methods. Another competing approach is a full structural estimation where a fully specified model is solved by dynamic programming and its parameters are estimated via a type of indirect estimation methodology. Since we want to allow for pervasive correlated heterogeneity in preference and income parameters, we do not think a full-fledged structural modeling with such heterogeneity is currently feasible. The full description of the Monte Carlo experiment is available on request or can be obtained from https://sites.google.com/site/salancrossley/publications.

Not all shocks to consumption arise from income surprises. Some of them are 'pure' consumption shocks and some arise from unanticipated changes in variables such as wealth and demographics. To construct the non-income shock, we first define two unit mean log normals:

$$\epsilon_{ht}^{i} = \exp\left(-\frac{\ln(1+\sigma_i^2)}{2} + \sqrt{\ln(1+\sigma_i^2)}\eta_{ht}\right) \text{ for } i = a, b \tag{8}$$

where η_{ht}s are independent standard normals. Then we define the non-income shock to the marginal utility of expenditure by:

$$\begin{aligned}
\tilde{\varepsilon}_{ht} &= d\epsilon_{ht}^{a} \text{ with probability } \pi \text{ where } d \in \left(0, \pi^{-1}\right) \\
&= \left(\frac{1-\pi d}{1-\pi}\right)\epsilon_{ht}^{b} \text{ with probability } (1-\pi)
\end{aligned}$$

The parameter d allows that the two components of the mixture have different (positive) means and the second expression ensures that $\tilde{\varepsilon}_{ht}$ has a unit mean. Allowing for different means for the components gives us a flexible distribution with skewness and kurtosis different from a single log-normal. In our estimation step we could not reject that the mixing parameter was equal to one half so we impose $\pi = 0.5$ in all that follows. The parameters (σ_a, σ_b, d) are common to all households.

The response of consumption to a contemporaneous income shock depends on the degree of insurance that is available to the household. To capture this, we model the reaction to income shock as:

$$\breve{\varepsilon}_{ht} = \exp\left(-\frac{(\lambda_h \nu_h)^2}{2} - \lambda_h \nu_h \xi_{ht}\right) \tag{9}$$

where $\nu_h \xi_{ht}$ is the contemporaneous income shock in (1). By construction, this distribution has a unit mean. The parameter λ_h is the idiosyncratic response to an income shock; it is positive since the impact of a positive income shock on the mue is negative. Note that this parameter is not a structural parameter, but depends on the preference and income parameters (see Kaplan and Violante (2010)).[4]

Although it is not a structural parameter, we can provide a useful interpretation of λ_h. Taking logs of (6), substituting in from (7) and (9) and taking the derivative with respect to the shock to income, $\nu_h \xi_{ht}$, we have the following expression for the change in log consumption (holding everything else constant):

$$d \ln C_{h,t} = \frac{\lambda_h}{\gamma_h} d\left(\nu_h \xi_{ht}\right) \tag{10}$$

The increment in log consumption growth due to a one percent positive income shock is given by:

$$\vartheta_h = \frac{\lambda_h}{\gamma_h} \tag{11}$$

Thus the change in consumption is larger the less the household is averse to fluctuations (low γ_h) and the higher is the direct income effect parameter λ_h. This aspect of the model relates to models that consider partial self-insurance. For example, Blundell *et al* (2008) define the degree of self-insurance as the fraction of the income shock that is transmitted to consumption growth. Analogously, the term ϑ_h represents the *degree of exposure to income shocks*, where $\vartheta_h = 1$ is no insurance at all (a one percent income shock translates into a one percent increase in consumption) and $\vartheta_h = 0$ is full insurance. Blundell *et al* (2008) use a different model for income with very limited heterogeneity and decompose the income shock into a permanent component and a transitory component. They can therefore obtain the self-insurance parameter for each of the two types of shocks. In this study, we focus on the heterogeneity and allow that the responses to an income shock can vary systematically across households.

Our framework also allows us to quantify the importance of non-income consumption shocks relative to income shocks. Taking logs of (7) and using (9), we obtain the proportion of variance

[4] Strictly, the reaction to an income shock will vary with age and/or time; see Kaplan and Violante (2010). In our empirical analysis we did not find any evidence of significant age dependence so we ignore that possibility here. This finding is in line with Blundell et al (2008), who do not find significant age dependence in the partial insurance coefficient.

of log shock due to income as against the total variance of log consumption shocks:

$$\begin{aligned}\kappa_h &= \frac{var\left(\ln\left(\breve{\varepsilon}_{ht}\right)\right)}{var\left(\ln\left(\breve{\varepsilon}_{ht}\right)\right)+var\left(\ln\left(\tilde{\varepsilon}_{ht}\right)\right)} \\ &= \frac{\lambda_h^2\nu^2}{\lambda_h^2\nu^2+var\left(\ln\left(\tilde{\varepsilon}_{ht}\right)\right)}\end{aligned} \tag{12}$$

where, once again, we made use of the independence between $\breve{\varepsilon}_{ht}$ and $\tilde{\varepsilon}_{ht}$. This ratio is increasing in the effect of a shock on the mue (λ) and the income variance (ν^2).

In summary, the heterogeneous parameters for consumption are $\{\delta_h, \gamma_h, \lambda_h\}$. In the specification below we allow for these parameters to be correlated with each other and also with the income process parameters $\{\mu_h, \alpha_h, \rho_h, \theta_h, \nu_h\}$. We refer to the set of household specific parameters as the *model parameters* to distinguish from two other types of parameters described below.

2.4 Measurement error

There is believed to be substantial measurement error in reported consumption and income in surveys such as the PSID. We need to take this into account in simulating consumption and income processes to match them with their data counterparts. To this end, we assume non-classical measurement error structures for both consumption and income. Specifically, we assume that observed (levels of) income and consumption have log-normally distributed, unit mean multiplicative error components with idiosyncratic variances (details are given in the next section). Denote the standard deviations of measurement error for income and consumption by m_h^y and m_h^c respectively. Taking variables u_{ht}^y and u_{ht}^c which are independent standard normals, we assume that observed *levels* of income and consumption are given by:

$$\begin{aligned}Y_{ht}^{obs} &= Y_{ht}\exp\left(-\frac{\left(m_h^y\right)^2}{2}+m_h^y u_{ht}^y\right) \\ C_{ht}^{obs} &= C_{ht}\exp\left(-\frac{\left(m_h^c\right)^2}{2}+m_h^c u_{ht}^c\right)\end{aligned} \tag{13}$$

where Y_{ht} is defined as $\exp(y_{ht})$ from subsection 2.1 and C_{ht} is given by (6). This gives two more heterogeneous parameters (m_h^y and m_h^c) in addition to the five income parameters and the three consumption parameters.

2.5 Accounting for heterogeneity

We model the joint distribution of the ten model parameters using a factor structure with (standard normal) factors denoted by N_k. The full model has ten factors (one for each model parameter), yielding a flexible correlational structure amongst the model parameters. The model parameters for the income process are specified as a five factor model:

$$
\begin{aligned}
\mu_h &= \phi_1 + \exp(\psi_{11}) N_{1h} \\
\alpha_h &= \phi_2 + \psi_{21} N_{1h} + \exp(\psi_{22}) N_{2h} \\
\rho_h &= \ell(\phi_3 + \psi_{31} N_{1h} + \psi_{32} N_{2h} + \exp(\psi_{33}) N_{3h}) \\
\theta_h &= 2 * \ell\left(\phi_4 + \sum_{j=1}^{3} \psi_{4j} N_{jh} + \exp(\psi_{44}) N_{4h}\right) - 1 \\
\nu_h &= \exp\left(\phi_5 + \sum_{j=1}^{4} \psi_{5j} N_{jh} + \exp(\psi_{55}) N_{5h}\right)
\end{aligned} \tag{14}
$$

where $\ell(x)$ is the transformation $e^x/(1+e^x) \in (0,1)$ so that $\rho_h \in (0,1)$ and $\theta_h \in (-1,1)$. The ψ_{kk} terms pick up heterogeneity while the ψ_{kj} $(j < k)$ terms pick up any co-dependence among the income process parameters.

For consumption, we allow for co-dependence with the income parameters and additional heterogeneity in preference parameters as follows:

$$
\begin{aligned}
\delta_h &= 0.1 * \ell\left(\phi_6 + \sum_{j=1}^{5} \psi_{6j} N_{jh} + \exp\left(\psi_{66}\right) N_{6h}\right) \\
\gamma_h &= 0.5 + 14.5 * \ell\left(\phi_7 + \sum_{j=1}^{6} \psi_{7j} N_{jh} + \exp(\psi_{77}) N_{7h}\right) \\
\lambda_h &= \exp\left(\phi_8 + \sum_{j=1}^{7} \psi_{8j} N_{jh} + \exp(\psi_{88}) N_{8h}\right)
\end{aligned} \tag{15}
$$

The parameter restrictions are $\delta \in (0, 0.1)$, $\gamma_h \in (0.5, 15)$ and $\lambda_h \in (0, \infty)$. The presence of coefficients such as ψ_{6j} allow for the preference parameters to be correlated with the income parameters. For example, the discount rate, δ_h, is allowed to be correlated with the income trend, α_h through ψ_{62}.

To incorporate measurement error, we take two new factors N_9 and N_{10} and define measurement error standard deviations by:

$$
\begin{aligned}
m_h^y &= \exp\left(\phi_9 + \sum_{j=1}^{8} \psi_{9j} N_{jh} + \exp\left(\psi_{99}\right) * N_{9h}\right) \\
m_h^c &= \exp\left(\phi_{10} + \sum_{j=1}^{9} \psi_{10,j} N_{jh} + \exp\left(\psi_{10,10}\right) * N_{10h}\right)
\end{aligned} \tag{16}
$$

for income and consumption respectively. With this structure we allow for the variance of measurement errors in both income and consumption processes to be correlated with each other (through the term $\psi_{10,9}$); with the income parameters and preference parameters through ψ_{9j} and through $\psi_{10,j}$.

In summary, the full set of homogeneous parameters to be estimated are (see equations 3 and 8):

$$
b_0, b_1, \sigma_a, \sigma_b, d
$$

and the parameters describing the distribution of the model parameters are:

$$\phi_1, \phi_2, .., \phi_{10}, \psi_{11}, \psi_{21}, \psi_{22}, ..., \psi_{10,10},$$

In our general factor model of the joint distirbution of model parameters, there are 10 parameters for location (the ϕ_k's), 10 parameters for dispersion (the ψ_{kk}'s) and 45 parameters for co-dependence (the ψ_{kj}'s for $j < k$). We refer to these as *distribution parameters* since they characterize the joint distribution of the model parameters. We estimate these parameters by indirect inference, which requires simulating income and consumption paths for a given combination of model parameters. We lay out the details of the estimation procedure after we present our PSID sample in the next section.

3 Data

We use the Panel Study of Income Dynamics (PSID) to estimate our model. The main advantage of the PSID is that it contains consumption and income information and it follows the same households over a long period. The survey contains detailed information on the annual household income and information on food at home and food at restaurants.[5] Our sample covers the periods between 1968 and 2009. An additional advantage of having data over such a long time period is that it gives us considerable intertemporal variation in real interest rates. The PSID is an annual panel survey from 1968−1997, switching to biannual from 1997 to 2009. Furthermore, no consumption information was recorded for the years 1968, 1973, 1988 and 1989.

We restrict our sample to households with married couples who stayed married throughout the sample period. All our households are headed by males, and we select husbands whose

[5]The use of food expenditure as a proxy to total expenditure is common in consumption studies as the PSID is the longest running panel available and it contains no information on household expenditure other than that of food. Alan and Browning (2010) and Browning and Crossley (2000) provide a formal justification for the use of food expenditure as a proxy for total expenditure. Another alternative would be to impute total expenditure using food expenditure as in Blundell *et al.* (2008). However, this would prevent us from using the years 1968-1979 in the PSID, since the CEX, which is used for imputation, is not available in those years.

education is above high school. We drop the periods in which the husband's age is below 30 or above 59. Finally, we exclude households that did not report food expenditure for at least 15 survey years[6] and households with very low income (<\$1) or very large changes in consumption (more than 400%) or income (200%).[7] Our final unbalanced panel has a minimum of 15, and a maximum of 26 survey years with a total of 583 households ($12,865$ observations). We assume that all households face a common real interest rate series calculated using the U.S. three-month treasury bill rates and the food price index.[8]

As a measure of income we use total family income deflated by the consumer price index. We also take into account of the fact that the measure of income refers to the previous calender year. The consumption measure contains the total value of all food consumed by the household (including money spend on food at home, food delivered, food out and the value of food stamps) deflated by the food consumption price index. We use the log of real income and real food expenditure for all our analyses below.

Note that demographics should be accounted for in this analysis and the common approach in this regard is to use a first round regression, where consumption and income are regressed on a set of variables including demographics, age and time dummies. This approach is problematic. First, it is an ad-hoc way of controlling for demographics and is not directly based on the theoretical model. Second, it removes important variation of age and time, which we would like to exploit in our estimation. Third, removing time effects can lead to bias in models with heterogeneity (for a detailed discussion, see MaCurdy (1982) and Browning and Ejrnæs (2013)). Given these reasons, we use a first round regression but only with a limited set of controls to deal with household composition. Specifically, we run a first round regression where we regress

[6]This implies that a household observed in 1968 should be observed at least until 1984 (17 calender years) to have 15 survey years with consumption information. This is because consumption is not reported in 1968 and 1973. Note that minimum 15 years refer to 15 years of observations, not 15 consecutive years as this is not possible given the issues mentioned in the PSID. However, as explained in Section 4 missing consumption information does not pose a problem for our estimation strategy.

[7]By the last selection criteria we exclude nine households.

[8]The use of a common interest rate series is fairly standard in the consumption literature. Our data do not contain information on specific interest rates individuals face. Depending on the form of heterogeneity in the interest rate individuals face, this could lead to biased estimates of the distribution of EISs and/or discount rates.

log real consumption and log real income on log household size, a dummy for whether children are present or not, and the age of the youngest child. Following MaCurdy (1982), we employ a fixed effects estimator. We use the residuals from these regressions in the subsequent analyses and will from now on refer to these residuals as income and consumption.

4 Empirical method

We estimate our model using indirect inference. Gouriéroux, Phillips and Yu (2010) provide a persuasive defence for using indirect inference in the context of estimating a fully parametric dynamic model for panel data. The advantages are: it is easy to use; it automatically corrects for the bias induced by the presence of the lagged dependent variable; it can automatically consider any statistics that previous researchers have used in estimation and it is simple to take account of features that arise from the sampling procedure, such as any imbalance in the panel or the change in sampling frequency in the PSID in 1997. The two principal steps in indirect inference are simulating from the parametric model and specifying a set of moments ('auxiliary parameters') that will be matched between sample data and simulated data.

4.1 Simulation

In the empirical implementation, we replicate each household R times to give $R * H$ simulated households. We first draw three sets of standard normal random numbers. The first set is for the income shocks, the ξ_{ht}'s in (3) and (1) for $t = 1, ..T$. The second set is for the consumption non-income shocks in (8), η_{ht} for $t = 2, ..T$. The final set is the factors, N_{kh}, for $k = 1, .., 10$; see (14)-(16). Once drawn, these random numbers are kept fixed in the estimation procedure.

For a given set of distribution parameters, we can construct model parameters from (14) and (15) and the factors N_{kh}. Based on the model parameters, we simulate income and consumption paths from 'age' 1 to age T. For the income paths we first calculate the initial income from (3); this gives $R * H$ values for y_{h1}. Then subsequent income paths are given recursively by (1) and the ξ_{ht}'s for $t = 2, ..T$.[9]

To simulate consumption growth paths we first simulate consumption shocks from (8); this uses the given values for $\{\sigma_a, \sigma_b, d\}$, the simulated values for λ_{ht} from (15) and the current income shocks, $\nu_h \xi_{ht}$. We set the initial value of consumption to unity[10] and construct levels sequentially, using (6) and the values for r_{t+1} (where t refers to age) and the simulated values for (γ_h, δ_h). Finally, measurement errors are added to the simulated incomes and consumptions, using (16).

In our sample we select on households that are aged 30 to 59 but many households are not observed at age 30 and/or at age 59. Moreover, many households appear after the first year of the PSID, 1968, or disappear before the last year, 2009. To take account of this unbalanced structure, we generate income paths for each replicated household for age $30-59$ and 'mask out' as missing the years between 1968 and 2009 as for the sample household that is being replicated. For example, suppose household h is born in year 1933 and is in the PSID from 1968 until 1994 so that the household is observed from age 35 to age 61. We select out the last two observations and thus have observations for age $35-59$ and years $1968-1992$. We simulate from age 30 until age 59 ($t = 1$ and $T = 30$ in the scheme of the previous subsection). Thus a path is modelled for this household from year 1963 until year 1992. We then drop the first 5 simulated values (1963 to 1967) and add missing values for the years 1993 to 2009. This procedure is valid since we do not have any year specific information that conditions the process. For consumption growth a similar procedure is followed, taking account of the fact that the real interest rate is year specific and needs to be made age specific for each household. In doing this one needs values for years outside the data period; for example, for the illustration in the previous paragraph we need values of the real rate for years 1963 to 1967.

[9] In practice, we start the income process from $t = -4$ to avoid awkward problems in modelling the first observations if we have a moving average process. We then discard the first five values to give a path from 1 to T.

[10] This choice of starting value distribution is irrelevant since the initial value plays no part in the simulated consumption growth path.

One further complication is that consumption is not recorded for the years 1968, 1973 and 1988 − 1989. When we have simulated years for consumption levels, we simply set the values for those years to missing. Finally, we have to take account of the fact that the PSID was an annual survey from 1968 until 1997, and then switched to a biannual survey, conducted in the odd years from 1999 until 2009. To deal with this, we set simulated values for those years to missing, just as in the original data. One of the great virtues of our indirect inference estimation procedure is that it allows us to take account of these survey features very cleanly. Basically, the simulated data is constructed to have exactly the same structure as the original data. This ensures that any bias in the moments induced by the peculiarities of sampling will be the same for the simulated sample as for the data sample.

4.2 Auxiliary parameters

Indirect inference requires the specification of a set of statistics which are known as *auxiliary parameters* (ap's). Estimation proceeds by comparing the ap's based on the sample with those based on the simulated data from the model. The estimated distribution parameters are determined by minimizing the weighted distance between the two sets of ap's. The ap's can be moments or functions of moments but could also be other statistics such as long or short run transitions. When choosing the ap's, one should ensure that the ap does have a probability limit as the number of cross-section units becomes large (but this probability limit does not have to be known, nor be anything of direct interest).

We choose the set of auxiliary parameters such that for each distribution parameter, there is at least one ap that is closely related.[11] Our construction of auxiliary parameters relies heavily on individual regressions for income and consumption growth. For example, for each household

[11] The ap does not have to be a consistent estimate of the distribution parameter, but it has to depend on it. However, this does not automatically ensure that the model is identified. We have investigated whether our choice of ap's is successful in identifying the structural parameters by performing a Monte Carlo experiment. The Monte Carlo experiment confirms that we can recover the true distributional parameters with reasonable precision.

we regress income on a constant and a trend, and obtain a household specific estimate of the trend in income. The average or median of all household trend-estimates can then serve as an ap for the distribution parameter ϕ_2 in equation (14). Browning and Ejrnæs (2013) provide a discussion of the advantages of using *individual regression based* (IRB) auxiliary models in the estimation of dynamic panel models. This requires us to run regressions on individual time paths of income and consumption growth. One problem that immediately arises for our data is that there are years in which some information is missing. To illustrate, consider a year in which consumption is not recorded (for example, 1973 for the PSID). To deal with the missing year, we linearly interpolate if the household is observed in 1972 and 1974 and set the value to missing if the household is not observed in either 1972 or in 1974. A similar interpolation is used for income and consumption after the survey switched to a biennial structure after 1997. If the auxiliary estimates were to be used directly, this would induce a bias of unknown form. In indirect inference, however, we circumvent this by using the same interpolation procedure for the simulated data.

4.2.1 Income

Denote the first and last ages at which household h is observed by t_{hf} and t_{hl} respectively. Based on our selection criteria discussed previously, we have at least 15 observations on any household. For estimating the ap's pertaining to income, we follow Browning and Ejrnæs (2013) and use a two step regression. In the first step, we regress log income, y_{ht}, on a constant and age for each household separately:

$$y_{ht} = b_{y1} + b_{y2} * t + e_{ht} \text{ for } t = t_{hf}, ..t_{hl} \tag{17}$$

and record $\left(\hat{b}_{y1}, \hat{b}_{y2}\right)$.[12] The H estimates of $\left(\hat{b}_{y1}, \hat{b}_{y2}\right)$ contain information on the distribution of income means μ_h and income trends α_h. In the second step, we regress the estimated residuals, $\hat{e}_{ht}$, on the lagged residuals:

$$\hat{e}_{ht} = \text{constant} + b_{y3}\hat{e}_{ht-1} + u_{ht} \text{ for } t = t_{hf} + 1, ..t_{hl} \tag{18}$$

and record $\hat{b}_{y3}$, which will contain information on the AR parameter ρ_h.[13] For the use in the next subsection, denote the expected value from this regression by $\check{e}_{ht}$. We then take the residuals from this regression and calculate the auto-correlation and the standard deviation:

$$\begin{aligned} \hat{b}_{y4} &= corr\left(\hat{u}_{ht}, \hat{u}_{h,t-1}\right) \\ \hat{b}_{y5} &= std\left(\hat{u}_{ht}\right) \end{aligned}$$

Here, $\hat{b}_{y4}$ captures the short run dynamics and contains information on the MA parameter, θ_h. Similarly, $\hat{b}_{y5}$ contains information on the distribution of the standard deviation ν_h. The joint distribution over H values of $\left\{\hat{b}_{y1}, \hat{b}_{y2}, \hat{b}_{y3}, \hat{b}_{y4}, \hat{b}_{y5}\right\}$ provide detailed information on (are 'bound to' in the indirect inference terminology) the joint distribution of the income process parameters $\{\mu_h, \alpha_h, \rho_h, \theta_h, \nu_h\}$ respectively.

4.2.2 Consumption.

For consumption, we follow a similar two step procedure. We first regress log consumption on a trend to give mean consumption growth for unit h; record this as $\hat{b}_{c1}$. The estimates of $\hat{b}_{c1}$ are intended to identify the distribution of discount rates as patience partially determines the trend in consumption. Next, take first differences of log consumption and regress this on the

[12] We surpress the index h to avoid triple indicies.

[13] Of course, $\hat{b}_{y3}$ is not an unbiased estimate of the AR parameter ρ_h due to short run dynamics and the small sample. However, $\hat{b}_{y3}$ still depends on the distribution of ρ_h, which is all we need for identification.

real interest rate and the estimated income shock from the previous sub-section:

$$\Delta c_{ht} = \text{constant} + b_{c2} r_t + b_{c3} \hat{u}_{ht} + w_{ht} \tag{19}$$

and record $\left(\hat{b}_{c2}, \hat{b}_{c3}\right)$. Here, $\hat{b}_{c2}$ captures the household specific elasticity of intertemporal substitution (the inverse of γ) and $\hat{b}_{c3}$ captures the consumption response to contemporaneous income shocks. These three sets of estimates characterize the distribution of preference parameters δ_h, γ_h and the partial insurance parameter λ_h.

Next, denote the estimated consumption change by $\hat{w}_{ht}$ and calculate the following the standard deviation and correlation coefficients:

$$\begin{aligned} \hat{b}_{c4} &= std\left(\hat{w}_{ht}\right) \\ \hat{b}_{c5} &= corr\left(\hat{w}_{ht}, \check{e}_{ht}\right) \\ \hat{b}_{c6} &= corr\left(\hat{w}_{ht}, \hat{w}_{h,t-1}\right) \\ \hat{b}_{c7} &= corr\left(\hat{w}_{ht}, \hat{u}_{h,t-1}\right) \end{aligned} \tag{20}$$

Here, the standard deviation $\hat{b}_{c4}$ yields information on the variance of the non-income shock. The correlation coefficient $\hat{b}_{c5}$ picks up the correlation between the consumption change and expected income ($\check{e}_{ht}$ in (18)); this is to check for excess sensitivity of consumption to current income. The correlation coefficients $\hat{b}_{c6}$ and $\hat{b}_{c7}$ allow us to identify the measurement error in consumption and income respectively. To identify the variance of measurement error in consumption we follow a standard approach and use the correlation of consumption growth between period t and $t-1$; see Runkle (1991). In our set-up this correlation is captured by $\hat{b}_{c6}$. Identifying measurement error in the income process is less standard and the idea we employ here, to our knowledge, has not been used before. When considering only income processes, measurement errors are not separately identified from the short run dynamics of the process, the MA(1) parameter (Meghir and Pistaferri (2004)). In order to identify the measurement error in the income process, we exploit the fact that we also observe consumption and that consumption reacts to *true* income shocks, not to the measurement error in income. We discuss the identification of measurement error in detail in Appendix A.1. Given the 12 estimates for each household unit ($\hat{b}_{y1}$ to $\hat{b}_{y5}$ and $\hat{b}_{c1}$ to $\hat{b}_{c7}$), we then construct our auxiliary parameters as medians (12 ap's), interquartile ranges (12) and correlation coefficients (66) between the 12 variables, yielding a total of 90 regression based ap's.

We also construct an ap to capture the potential age-dependence of the partial insurance parameter, λ, as in Kaplan and Violante (2010). The age dependence in λ is captured by the correlation between consumption residuals and income residuals interacted with t: that is, $corr(\hat{w}_{ht}, \hat{u}_{ht} \cdot t)$.[14] The next ap we define aims to generate a well-known stylized fact in the consumption literature. This is that the cross sectional variance of consumption increases linearly over the life-cycle (Deaton and Paxson (1994)). To check that our model captures this feature we include the estimated trend in the cross sectional interquartile range over the life cycle as an additional ap.

Finally, our procedure also requires ap's for the distribution of the starting values given in (3). To construct these, we first regress log income at age 30 on the year of birth to take out cohort effects. We then record the estimated intercept and the standard deviation of the residuals we obtain from this regression. The only complication here is that we do not observe all households at age 30. We follow Browning *et al* (2010) and run the regression for the subsample of households observed at age 30; in our data this constitutes 56% of the sample. Note that in the simulation step we mask out the value at age 30 for replications of households that are not observed from that age, so that the same proportion of households is used for the simulated data.

With these extra 2 ap's we have a grand total of 94 ap's to fit, which we believe provide

[14] Our Monte Carlo study confirms that this ap would indeed pick up a common age-dependence in λ.

a rich characterization of the joint distribution of consumption and income parameters for our PSID sample. In the estimation step we use 80 of the ap's to fit the model. We keep back 14 ap's, 13 of which are associated with $\hat{b}_{c5}$ to provide a quasi-LM test for 'excess sensitivity' (to test for liquidity constraints), and the ap that picks up any age-dependence in λ_h.

5 Results

5.1 The fit of the model.

Our full model is a 10 factor model with 70 parameters to be estimated by matching to 80 auxiliary parameters. To estimate the model, we first performed an initial specific to general specification search. This starts with a very parsimonious model that only allows for very limited heterogeneity. This model fits very poorly. We then add parameters one at a time to deal with the worst fitting ap at each step. For example, the most parsimonious model fits very poorly the ap for the variability of the standard deviation of the income shocks; this is dealt with by including a distribution parameter (ψ_{55}) which controls the heterogeneity in the income standard deviations (the ν_h's). This fire-fighting approach is a reliable method for estimating large factor models with many parameters. Once we have a specification that cannot be significantly improved by adding further parameters, we conduct a final general to specific search to eliminate 'insignificant' parameters. This search resulted in a reasonably parsimonious model with 32 parameters and seven factors. In this preferred model, we have four factors for the income parameters; one additional factor for the preference parameters and two factors for the measurement error parameters.

The estimated parameters of the preferred model are given in Appendix Table A.1. In this table, we follow the convention and also present the standard errors, calculated using the delta method. As is well known, in non-linear models such standard errors are not invariant to the normalizations used and can be quite unreliable; see for example, Cameron and Trivedi

(2005), section 7.2.9. For this reason, we chose to rely on quasi likelihood ratio (QLR) statistics (comparisons of the fit of the restricted and unrestricted models) for our specification search, and exclude parameters accordingly. For example, the parameter governing the relationship between the income trend and the discount rate, ψ_{62}, has a low 't-value' of 1.07, but a high $\chi^2(1)$ QLR statistic of 6.1 and is therefore retained in the final preferred model.

The estimated values reported in Table A.1 have no immediate interpretations. Below we discuss the implications of these estimates in terms of characterizing the joint distribution of the model parameters and measurement error parameters.

Our preferred model has 32 parameters to estimate and 80 ap's to match, yielding 48 degrees of freedom. The over-identification (OI) test statistic is 72.24 so that the overall fit is marginal. The fits for most ap's are good; see Table A.2 in the Appendix. The worst fit, in statistical terms, is for the mean of $\hat{b}_{c6}$, the auto-correlation of the residuals from the Euler equation; see (20). This has a data value of -2.73 and a simulated value of -3.22 and a standard error for the difference of 0.14. The fit of the trend in cross sectional variation in consumption (see Deaton and Paxson (1994)) is reasonable (see the ap labelled as CS IQR in Table A.2), suggesting that our model is able to produce this important stylized fact.

The value for the QLM test for the 14 ap's not used in fitting is 16.1, which has a $\chi^2(14)$ distribution. The first 13 ap's that we keep back for this test relate to excess sensitivity; the low QLM statistic implies no evidence of excess sensitivity.[15] This is the most direct evidence we have that liquidity constraints are not important for our sample (see Appendix Table A.3). This result is not surprising as our sample contains households older than 30 who are less likely to be constrained. The last ap, $corr(\hat{w}_{ht}, \hat{u}_{ht} \cdot t)$, in the QLM test captures the potential age-dependence in the partial insurance parameter λ. The GF test (see Appendix Table A.3) indicates that our preferred specification, without an age-dependent λ, fits the ap reasonably well. Hence, despite the fact that a standard life-cycle model implies an age-dependence in

[15] The test statistics for excess sensitivity test is 13.11, which has a $\chi^2(13)$ distribution.

self-insurance, we find no statistically significant age-dependence in λ. This finding is consistent with Blundell *et al* (2008).[16]

5.2 Marginal distributions of model parameters.

Table 1 presents the marginal distributions of the heterogeneous model parameters. The table is divided into three panels. The first panel presents the income parameters, the second presents consumption parameters as well as the partial insurance parameter λ. The third panel presents the estimates for the three additional measures that are of interest, which we discuss in the next subsection.

For the income parameters the most striking result is the extent of heterogeneity in the standard deviation of the shock which ranges from 0.06 to 0.26; evidently some households have much more variable net income paths than others. A similar result is found for men's gross earnings using the PSID in Browning *et al* (2010) and using Danish data in Browning and Ejrnæs (2013). The upper value is particularly notable: a household with a standard deviation of 0.26 has a 2.5% probability of seeing its income drop by 40% from one year to the next, and a 2.5% probability of an increase of over 66%. In any discussion of social insurance this heterogeneity should play a critical role with high variance households valuing social insurance much more highly. There is also evidence of heterogeneity in trends and the $ARMA$ parameters. We find slightly less heterogeneity in the trend compared to individual earnings of men (see Browning *et al* (2010)). For the AR parameter, we find that most of the households are not close to having a unit root income process. In a recent study on Norwegian data, Blundell *et al* (2015) find an AR coefficient (assumed homogeneous) for disposable family income of 0.86, which is very close to our estimate of the median (0.85). The MA parameters are generally positive, which

[16] Kaplan and Violante (2010) discuss the lack of empirical support for the age-dependence in partial insurance and point to the fact that the simple life cycle model implies too much concentration of wealth at retirement compared to what is observed in the data. For example, a realistic model that allows for a bequest motive for the old and a specific saving motive for the young (such as a down-payment motive) would result in a flatter age-profile in the consumption reaction to income shocks.

contrasts with studies that do not explicitly control for measurement error. This is consistent with the result that measurement error induces a negative bias in the MA parameter.

Turning to the preference parameters (Panel 2), we first note that the estimated discount rate is heterogeneous with the median value of 7.7%, which is very close to the median discount rate estimated by Samwick (1998) and in line with previous studies using micro data on consumption, wealth and portfolio choice. The standard way of addressing discount rate heterogeneity has been to estimate discount rates for different education groups, assuming homogeneity within groups. The estimated range across education groups in Gourinchas and Parker (2002) is 3.94% to 5.93% and Cagetti (2003) estimates the range as 2% to 16%. All studies suggest a higher discount rate for the less educated. Alan and Browning (2010) is the only study that estimates individual specific discount rates using consumption data and, consistently with these studies, find higher median discount rate for the less educated (7.7%). As mentioned in the introduction, there is a growing literature that experimentally elicits individual discount rates using hypothetical or incentivised tasks that involve trade-offs between current and future consumption. Distributions of discount rates elicited experimentally are much higher than the estimates obtained from observational data (see Andreoni and Sprenger (2012) for the theoretical justification for this).

We also find considerable heterogeneity in the coefficient of relative risk aversion with the estimated median value of 7.48, which is consistent with Alan and Browning (2010). They find median coefficient of relative risk aversion to be 6.2 and 8.4 for the low and high educated respectively. These estimates are higher than those reported in most consumption studies which impose homogeneity. With regard to the heterogeneity in this parameter, as far as we are aware, all studies that allow for heterogeneity in risk tolerance find evidence of substantial differences across people. For example, the widely cited results in Barsky *et al* (1997) indicate considerable risk aversion (the modal group has a value between 4 and 16) but also considerable dispersion (23% have a coefficient of relative risk aversion of less than 2). Similarly, the experimental studies such as Andersen *et al* (2008) find considerable dispersion in risk tolerance parameters. Using a large representative sample who are asked directly about their attitudes to risk, Dohmen *et al* (2011) find considerable dispersion in responses. Similarly, Guiso and Piaella (2005) find a great deal of heterogeneity in an Italian survey that asks about the willingness to pay for a hypothetical lottery. They estimate a median coefficient of relative risk aversion of 4.8 with 90% of the sample being between 2.2 and 10.

Finally in Panel 2, the parameter that captures the direct impact of income shocks on the mue, λ_h, is also found to be very heterogeneous with some households hardly responding (the first decile value is 0.33) and others responding a lot (the ninth decile value is 4.47) to income shocks.[17]

5.3 Income shocks and expenditure reactions

One of our main contributions is to quantify the importance of income shocks at the household level. This contribution advances the literature that studies the way in which income and wealth shocks are transmitted to consumption as in Blundell *et al* (2008), Alan *et al* (2014) and DeNardi *et al* (2012). Table 1, Panel 3 presents the related estimates. The first of these estimates is the long run effect on income of an income shock, τ_h, as defined in (2). The AR and the MA parameters determine the dynamics of the income process and the cumulative impact of a shock, τ_h. From the estimates of τ it is clear that even though we have a stationary model for everyone, for some households the long run effect of an instantaneous shock can be quite large. The median value suggests that the cumulative impact of a shock is 7.9 times the value of the (transitory) instantaneous shock. This parameter is very dispersed with the most persistent having a value of about 40. This highlights the fact that for some households even a small net income shock might result in a significant loss in life-time income with a consequent consumption loss.

[17] The absolute value of this parameter does not have an immediate interpretation, see equation (11).

The second estimate is the consumption response to a positive income shock, ϑ_h as defined in (11). This parameter determines the amount of income shock transmitted to consumption. Recall that the parameter ϑ can also be interpreted in relation to partial insurance where the value one means no insurance at all and the value zero means full insurance. For the median household, a one percent income shock raises consumption by 0.19%. For comparison, Blundell *et al.* (2008) estimate a model without heterogeneity using a different income process, and they find that a one percent permanent income shock raises consumption by 0.41%, while a one percent transitory income shock by 0.02%.[18] In our study, this parameter exhibits a great deal of heterogeneity. At the ninth decile, the impact on consumption is 0.68%, which indicates that even those who react a lot can still achieve considerable consumption smoothing. This finding of heterogeneity suggests that similar income shocks can generate very different consumption responses across households, pointing to important positive and normative implications.

Our third estimate is the proportion of mue shocks that are due to income shocks, κ_h as defined in (12). This allows us to quantify the effect of different types of shocks have on different households. Theoretically, we do not expect contemporaneous income shocks to constitute a large part of consumption shocks for households with high net worth since for these households consumption is mainly financed by the accumulated financial wealth, not by the labor income. However, for these households some non-income shocks for example, wealth shocks stemming from asset price changes can be quite important (see Alan *et al* (2014)). In contrast, income shocks are likely to constitute a large part of consumption shocks for individuals with low wealth. The estimated distribution of this parameter is very dispersed with a median value of 24%. In our sample, we observe households for whom income shocks hardly matter (4% at the first decile) and households with considerable 'vulnerability' to income shocks (68% at the ninth decile)[19].

[18] These numbers are for the "college sample" in Table 6 in Blundell et al (2008).

[19] Although we know of no previous estimates of κ with which to compare our results, it may be that the median value appears too low. However, we think this is an empirical fact and not related to the methodology we use. The reason why we think so is because when we perform our Monte Carlo experiments in a model where

		10%	50%	90%
Panel 1: Income model parameters				
μ	Mean at start of process	-0.13	0.04	0.22
α	Trend ($\times 100$)	-1.07	-0.05	0.97
ρ	AR parameter	0.50	0.85	0.97
θ	MA parameter	-0.15	0.22	0.54
ν	Shock standard deviation	0.06	0.13	0.26
Panel 2: Preference model parameters				
δ	Discount rate ($\times 100$)	4.86	7.69	9.20
γ	Coefficient of relative risk aversion	2.25	7.48	12.77
λ	Response of mue to income shocks	0.33	1.14	4.47
Panel 3: The effect of an income shock				
τ	Long run effect of a shock on discounted income	2.33	7.87	40.40
ϑ	Effect on consumption	0.05	0.19	0.68
κ	Proportion of mue log shock due to income shocks	0.04	0.24	0.68

Table 1: Marginal distribution of model parameters

5.4 The co-dependence between income and consumption parameters

We now turn to discussing the co-dependence between income and consumption parameters. First, we present the co-dependence within income and within consumption parameters, then we discuss our findings on the co-dependence across the two processes, which is the main point of the paper.

5.4.1 Co-dependence among income parameters

Our modeling with pervasive heterogeneity allows us to pin down the empirical association between all model parameters. Table 2 presents the estimated correlation coefficients amongst income parameters. The main difference from previous studies on individual earnings processes is that we do not detect a correlation between the trend α and the level parameter μ (see e.g. Baker (1997), Rubinstein and Weiss 2006 or Browning *et al.* (2010)). This lack of correlation in our study is likely due to the fact that we start to observe income at age 30. If much of the income growth takes place during the first few years in the labour market we would not

the only uncertainty is due to income and interest rate shocks, we find that income shocks explain about 95-99 percent of the variation in expectation shocks.

	μ	α	ρ	θ	ν	τ
μ	1.00	0.01	-0.03	1.00	0.00	0.11
α		1.00	0.05	0.00	-0.15	-0.01
ρ			1.00	-0.03	-0.74	0.41
θ				1.00	0.00	0.11
ν					1.00	-0.37

Table 2: Correlations between income parameters

expect to see much correlation at this age. Moreover, in contrast to Browning *et al* (2010), we find a negative correlation between the trend α and the variance parameter ν. Perhaps our most interesting novel finding regarding income is that there is a negative correlation between the variance of the income shock, ν, and the long run impact on life time income, τ. This implies that some households experience large but less persistent shocks while other households experience small but more persistent shocks.

5.4.2 Co-dependence between preference parameters

The correlations between the model parameters for consumption are displayed in Table 3. There is a strong positive correlation between the coefficient of relative risk aversion, γ, and the discount rate, δ, implying that impatient people are more risk averse. This is consistent with the results for the within education group correlations in Alan and Browning (2010). The empirical evidence on this correlation from the experimental literature is largely in agreement with our finding. Anderhub *et al* (2001) (using a sample of Israeli students) find a negative correlation between risk aversion and the discount *factor* which is consistent with our findings. Eckel *et al* (2005) conduct experiments with low income people in Montreal and find that 'risk averse individuals are also more present-oriented' which is again consistent with our findings. On the other hand, Harrison *et al* (2007) present results for a representative sample drawn from the Danish population and find no correlation.

As discussed previously, the parameter λ, the reaction of the mue to an income shock, is not a structural parameter, but depends on preference and income parameters. In our preferred model

	δ	γ	λ	ϑ	κ
δ	1.00	0.73	0.25	-0.20	0.44
γ		1.00	0.33	-0.24	0.55
λ			1.00	0.63	0.68
ϑ				1.00	0.27

Table 3: Correlations between preference parameters

there is no independent factor for λ_h so that all the heterogeneity in λ stems from the dependence on other structural parameters. Dependence of λ on the discount rate is particularly intuitive since the latter is a key parameter for determining life time wealth accumulation. Households with a high discount rate accumulate lower net wealth which makes them much more sensitive to income shocks. It is also plausible to expect that income shocks constitute a larger component of mue shocks for high discount rate households. This is consistent with our finding of a strong positive correlation between δ and the response to income shock λ, as well as the proportion of the mue shocks to income shocks, κ(see Table 3).

The parameter $\vartheta = \lambda/\gamma$ gives the degree of exposure to income risk (see the discussion after (11)). We find that ϑ is negatively correlated with the risk aversion parameter, γ, implying that the consumption of risk averse households is less affected by income shocks, possibly because they tend to accumulate more wealth due to the precautionary motive (governed by γ). This result also provides empirical support for the theoretical findings in Kaplan and Violante (2010).[20] However, we find that the correlation between the discount rate, δ, and ϑ is negative; the opposite of the intuitive positive correlation between δ and λ. The reason for this is because ϑ is negatively correlated with γ and the latter is strongly positively correlated with δ. The parameter ψ_{86} (the direct link between δ and λ) is positive, but for our sample, the magnitude of this correlation is not sufficient to outweigh the effect coming from γ.

[20] Note that $1 - \vartheta$ almost corresponds to the partial insurance coeffient in Kaplan and Violante (2010).

5.4.3 Cross process co-dependence

Turning to the co-dependence among income and preference parameters, we present the estimated correlation coefficients in Table 4. We find that the discount rate δ and the income trend α are negatively correlated. That is, impatient households have lower trends in income than patient households. As emphasized by Cunha, Heckman and Navarro (2005), the relationship between the 'choice' of an income process and intertemporal allocation preferences depends on the market environment. Our finding would be an immediate implication if there are imperfect capital markets. However, under our perfect capital markets assumption, whereby individuals can borrow and lend at the same rate, we do not expect impatient individuals to have flatter income profiles. An alternative (Mincerian) rationalization of our finding is that higher effort in the earlier years leads to a steeper income profile and patient people are more willing to exert such effort (and perhaps forgo immediate leisure possibilities) for the sake of future rewards. This explanation relies on impatience impacting on an unobserved variable (effort) which in turn calls for a future study of labour supply and human capital formation jointly with consumption profiles. Another issue regarding the negative correlation is that it largely reflects that households with a positive trend in income have a higher growth rate for consumption. This superficially looks as though what we are picking up is 'consumption tracking income'. However, the lack of any evidence of excess sensitivity suggests that this is not a viable alternative explanation.

To our surprise, we do not find any significant correlation between the coefficient of relative risk aversion γ and any of the income parameters; in particular, the dispersion of income ν.[21] This is at odds with the literature on occupational choice and earnings risk. For example, Bonin *et al* (2007) find that individuals with lower willingness to take risks (as measured by survey

[21] Concerned about capturing important correlations in the model, we run two separate MC experiments: one with the zero correlation and one with a negative correlation between the risk aversion parameter γ and the variance parameter ν of the income process. This is done to assess whether we can detect this correlation if it exists and recover the zero correlation when such a correlation does not exist. Our simulation results show that, in both MC cases, we can recover the parameters and the correlation coefficients.

questions) are more likely to work in occupations with low earnings risk. Similarly, Skriabikova *et al* (2012) show that after the transition to a market economy in Ukraine, workers who are more willing to take risk (again, measured via a survey question) switch to jobs with a higher earnings variance. Our results suggest that even risk averse agents do *not* use occupational choice to mitigate the need for precautionary saving. One possible explanation for the difference in conclusions is that we use observational information on consumption to measure risk attitudes rather than subjective survey questions. Assuming expected utility implies a strong link between risk aversion and reactions to changes in intertemporal prices; it would be of interest in future work to investigate the link between survey questions concerning risk attitudes and consumption behavior. Another potential reason for our "no correlation" result may be that our unit of observation is the household rather than an individual; see Shore (2010). Even if risk preferences and occupational choice are co-dependent at the individual level, as suggested by the cited studies, household level data may not reveal this. In this study we have ignored the possibility that husbands and wives may have different preference parameters and the related issue of how then to define household preferences. This raises a new set of issues which we leave to future work.

We find strong co-dependence between the degree of exposure to income shocks ϑ and the income parameters. The parameter ϑ is negatively correlated with the dispersion of income ν. This indicates that for those households with more volatile income the reaction to an income shock is smaller. This is consistent with households with high income shock variance building up buffer stocks to self-insure against income shocks. The reaction to an income shock ϑ is positively correlated with the persistence of income shock ρ and the long run persistence of shocks τ. This is consistent with persistent shocks having a larger impact on future discounted lifetime income; see Blundell *et al* (2008) for empirical evidence and Kaplan and Violante (2010) for theoretical results.

Finally, we find that the proportion of consumption shock variance due to income shocks

	δ	γ	λ	ϑ	κ
μ	-0.02	-0.01	-0.01	0.00	-0.03
α	-0.65	-0.03	-0.01	-0.01	-0.08
ρ	-0.06	-0.03	0.54	0.52	0.50
θ	-0.01	0.00	-0.01	0.00	-0.03
ν	0.11	0.01	-0.45	-0.44	-0.15
τ	-0.02	-0.05	0.75	0.87	0.34

Table 4: Correlations between income and consumption parameters

κ is positively correlated with the long run persistence of income shocks τ and negatively correlated with the variance of income shocks ν. This implies that income shocks are relatively less important for households with volatile income and less persistent shocks.

5.5 Measurement error

In our model we allow for non-classical idiosyncratic measurement error, see section 2.4 and equation (16). In Table 5, we present the ratio of the variances of noisy measure of income and consumption to the true variance. The table indicates considerable variance in the measurement error. At the median, the ratio for income is 1.23 indicating that the noisy measure is 23 percent higher than the true measure and at the ninth decile the variance of the noisy measure is more than twice as large as the variance of the true measure. The estimated median of the variance of the measurement error is close to the value Bound *et al* (1994) found in their PSID validation study. For consumption, the ratio of the variances is very large. At the median the ratio of the variances is four times as big as the variance of the true measure - as many others have concluded, the PSID consumption measure is very noisy.

In this specification we also allow for a correlation between the variances of the measurement errors in the two processes. The correlation is determined by the parameter $\psi_{10,9}$ which is estimated to be 0.33 (see Table A.1), indicating a positive correlation between the variances of the measurement errors. To our knowledge, this is the first piece of evidence that supports the plausible hypothesis that the accuracy of survey responses on consumption and income are positively correlated.

		10%	50%	90%
	Income	1.04	1.23	2.13
	Consumption	1.18	4.81	31.4

Table 5: The ratio of noisy variance to the true variance

6 Conclusion

We provide a framework for modelling income and consumption together whilst allowing for pervasive and co-dependent heterogeneity in both processes. At the household level we introduce a link between the two processes whereby the consumption shock depends in part on the contemporaneous income shock. We then develop a parametric factor structure to capture heterogeneity across households. In doing this, we allow for co-dependence between all of the income and consumption parameters.

Using a PSID sample from 1968 to 2009, we find considerable heterogeneity in income and consumption parameters, and co-dependence between the parameters governing the two processes. Our estimates of the intertemporal allocation parameters are very dispersed. Even though the estimated median values, considered in isolation, are similar to those documented in the literature, we posit that positive and normative analyses that focus on average values can be very misleading; see, for example, Browning, Hansen and Heckman (1999). We also find that the consumption reaction to an income shock is heterogeneous, implying a great deal of heterogeneity in the degree of self-insurance available to households. This particular finding has implications for welfare evaluations of social insurance and evaluations of the efficacy of stimulation policies.

Documenting the correlated heterogeneity in income and intertemporal allocation parameters is a novel endeavour in itself but the core contribution of our paper pertains to the usefulness of these estimates. They allow us to construct estimated quantities of crucial policy relevance, which were previously not available. Ignoring household level heterogeneity in these quantities may lead to misguided policy evaluations and welfare analyses. Although welfare evaluations and policy experiments are outside the scope of this paper, the framework we offer and the novel estimates we provide pave the way for such efforts.

Possibilities of future work our study generates abound. Future research that focuses on policy evaluations under pervasive heterogeneity would be especially promising. On the modeling side, our model can be further enriched by explicitly accounting for demographics; examples include modeling fertility jointly with income and consumption and explicitly allowing for aggregate shocks.

References

[1] Alan, S.,and M. Browning (2010), "Estimating Intertemporal Allocation Parameters Using Synthetic Residual Estimation", *Review of Economic Studies* 77, 1231-1261.

[2] Alan, S., T. Crossley and H. Low (2014), "Saving on a Rainy Day, Borrowing for a Rainy Day", manuscript.

[3] Anderhub, V., W. Guth, U. Gneezy and D. Sonsino (2001), "On the Interaction of Risk and Time Preferences: An Experimental Study" *German Economic Review*, 2(3).

[4] Andersen, S. G. W. Harrison, M. I. Lau, and E. E. Rutstrom (2008), "Eliciting Risk and Time Preferences", *Econometrica*, 76(3), 583-618.

[5] Andreoni, J. and C. Sprenger (2012), "Estimating Time Preferences from Convex Budgets.", *American Economic Review*, 102(7): 3333-56.

[6] Arellano, M. (2003), *Panel Data Econometrics*. Oxford: Oxford University Press.

[7] Atkinson, A. B. and J. E. Stiglitz (1976), "The Design of Tax Structure: Direct Versus Indirect Taxation", *Journal of Public Economics*, 6(1-2): 55-75.

[8] Baker, M. (1997), "Growth-Rate Heterogeneity and the Covariance Structure of Life-Cycle Earnings", *Journal of Labor Economics*, 15(2), pp. 338-75

[9] Banks, J. and P. Diamond (2008), "The Base for Direct Taxation", *Mirrlees Review*, Institute for Fiscal Studies, London.

[10] Barberis, N. M. Huang and R. Thaler (2006), "Individual Preferences, Monetary Gambles, and Stock Market Participation: A Case for Narrow Framing" *American Economic Review,* 96, 1069-1090.

[11] Barsky R. B., F. T. Juster, M. S. Kimball, M. D. Shapiro (1997), "Preference parameters and Behavioral Heterogeneity, An Experimental Approach in the health and Retirement Study", *Quarterly Journal of Economics*, May

[12] Blundell, R., M Graber and M. Mogstad (2015), "Labor Income Dynamics and the Insurance from Taxes, Transfers, and the Family", *Journal of Public Economics*, Volume 127, pp. 58-73.

[13] Blundell R. , Luigi Pistaferri, Ian Preston (2008), "Consumption Inequality and Partial Insurance", *American Economic Review*, Vol. 98, No. 5 (Dec., 2008), pp. 1887-1921.

[14] Bonin, H., T. Dohmen, A., Falk, D. Huffman and U. Sunde (2007), "Cross-sectional earnings risk and occupational sorting: The role of risk attitudes", *Labour Economics*, 14(6):926-937.

[15] Bound, J., C. C. Brown, G. J. Duncan, and W. Rodgers (1994). "Evidence on the Validation of Cross Sectional and Longitudinal Labor Market Data" *Journal of Labor Economics*, 12(3): 345-68.

[16] Browning, M. and T. Crossley. “Luxuries Are Easier to Postpone: A Proof”. *Journal of Political Economy,* vol 108(5), pp.1022–1026.

[17] Browning, M., M. Ejrnæs, J. Alvarez (2010), "Modelling Income Processes with Lots of Heterogeneity", *Review of Economic Studies* 77, 1353-1381.

[18] Browning, M and M. Ejrnæs (2013), "Heterogeneity in the Dynamics of Labor Earnings", *The Annual Review of Economics*, vol. 5. pp. 219-245

[19] Browning, M., L. Hansen, and J. Heckman (1999), "Micro data and general equilibrium models", *Handbook of Macroeconomics* 1: 543-633.

[20] Cadena, B. and B. J. Keys (2015), "Human Capital and the Lifetime Costs of Impatience", *American Economic Journal: Economic Policy*, vol. 7(3), pp. 126-53.

[21] Cagetti, M. (2003), "Wealth Accumulation Over the Life Cycle and Precautionary Savings", *Journal of Business and Economic Statistics*, 21(3), 339-353.

[22] Cameron, C. and P. Trevedi (2005), *Microeconometrics*, Cambridge University Press, New York

[23] Cunha, F., J. Heckman and S. Navarro (2005), "Separating uncertainty from heterogeneity in life cycle earnings." *Oxford Economic Papers* - New Series 57(2): 191-261.

[24] Deaton, A. and C. Paxson (1994), "Intertemporal Choice and Inequality", *Journal of Political Economy*, vol. 102(3), pp. 437-67.

[25] De Nardi, M., French, E. and D. Benson (2012) "Consumption and the Great Recession", *Economic Perspectives* 36.1

[26] Diamond P. and J. Spinnewijn (2011), "Capital Income Taxes with Heterogeneous Discount Rates", *American Economic Journal: Economic Policy*, 3 (4). pp. 52-76..

[27] Dohmen, T., A. Falk, D. Huffman, U. Sunde, J. Schupp, and G. G. Wagner (2011), "Individual Risk Attitudes: Measurement, Determinants, and Behavioral Consequences", *Journal of the European Economic Association*, 9(3):522-550.

[28] Eckel, C., C. Johnson and M. Claude (2005), "Saving Decisions of the Working Poor: Short-and Long-Term Horizons", *Field experiments in Economics,* 219.

[29] Farhi, E., and I. Werning (2012) "Capital taxation: Quantitative explorations of the inverse Euler equation." *Journal of Political Economy* 120(3): 398-445.

[30] Gneezy, U. and J. Potters. “An Experiment on Risk Taking and Evaluation Periods”. *The Quarterly Journal of Economics,* vol 112(2), pp. 631–645.

[31] Gourieroux, C., P.C. Phillips, and J. Yu (2010), "Indirect Inference for Dynamic Panel Models", *Journal of Econometrics*, 157, 68-77.

[32] Gourinchas, P. and J. Parker (2002), "Consumption Over the Life Cycle", *Econometrica*, 70(1), 47-89.

[33] Guvenen, F. (2009), "An Empirical Investigation of Labor Income Processes", *Review of Economic Dynamics*, 12(1), pp. 58-79.

[34] Guvenen F. and A. Smith (2014), "Inferring Labor Income Risk From Economic Choices: An Indirect Inference Approach", forthcoming, *Econometrica*.

[35] Guiso, L. and M. Paiella (2008), "Risk Aversion, Wealth and Background Risk", *The Journal of the European Economic Association*, 6(6), 1109-1150, 2008.

[36] Harrison, G. W, M. I. Lau, and E. E. Rutstrom (2007), "Estimating Risk Attitudes in Denmark: A Field Experiment", *Scandinavian Journal of Economics*, 109, 341-368.

[37] Heckman, J. (2001), "Micro Data, Heterogeneity, and the Evaluation of Public Policy: Nobel Lecture", *Journal of Political Economy*,vol. 109(4), pp. 673-748

[38] Huggett, M. and J. C. Parra (2010), "How Well Does the Us Social Insurance System Provide Social Insurance?", *Journal of Political Economy*, 118(1), pp. 76-112.

[39] Holt, C. A., and S. K. Laury (2002), "Risk Aversion and Incentive Effects", *American Economic Review*, 92, 1644-1655.

[40] Kaplan, G. and G. Violante (2010), "How Much Consumption Insurance Beyond Self-Insurance?", *American Economic Journal: Macroeconomics,* vol 2(4), pp. 53-87.

[41] Kocherlakota, N. (2010), *The new dynamic public finance*, Princeton University Press.

[42] MaCurdy, 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), pp. 83-114

[43] Meghir, C. and L. Pistaferri. (2004), "Income Variance Dynamics and Heterogeneity", *Econometrica*, 72(1), pp. 1-32

[44] Mincer, J., (1958), "Investment in Human Capital and Personal Income Distribution", *Journal of Political Economy*, Vol. 66, No. 4, pp. 281-302.

[45] Moffitt, R. A. and P. Gottschalk (2012), "Trends in the Transitory Variance of Male Earnings, Methods and Evidence", *Journal of Human Resources*, 47(1), pp. 204-36

[46] Rubinstein, Y. and Y. Weiss, "Post Schooling Wage Growth: Investment, Search and Learning", In: E. Hanushek and F. Welch, Editor(s), *Handbook of the Economics of Education*, Elsevier, 2006, vol 1, pp. 1-67.

[47] Runkle, D. E., (1991), "Liquidity Constraints and The Permanent Income Hypothesis", *Journal of Monetary Economics*, 27, 73-98.

[48] Saez, E., (2002), "The Desirability of Commodity Taxation under Non-Linear Income Taxation and Heterogeneous Tastes" *Journal of Public Economics*, 83(2): 217-30.

[49] Samwick, A. (1998), "Discount Rate Heterogeneity and Social Security Reform", *Journal of Development Economics*, 57 (October 1998), 117-146.

[50] Shore, S. (2010), "For Better, for Worse: Intrahousehold Risk-Sharing over the Business Cycle", *Review of Economics and Statistics*, 92(3): 536-548.

[51] Skriabikova, O., T. Dohmen, B. Kriechel (2012), "Risk attitudes and occupational choice", manuscript.

A Appendix

A.1 Identification of income measurement error

In general it is not possible to identify the measurement error in an income process without ruling out potentially important short run dynamics. Assume that the true income y^*_{ht} can be described by a general ARMA(1,1) and for simplicity ignore the deterministic part. The process for log income is given by

$$y^*_{ht} = \rho y^*_{ht-1} + \xi_{ht} + \theta\xi_{ht-1}.$$

We observe the income process with a classical measurement error m_{ht} :

$$\begin{aligned} y_{ht} &= y^*_{ht} + m_{ht} = \rho(y_{ht-1} - m_{it-1}) + \xi_{ht} + \theta\xi_{ht-1} + m_{ht} \\ &= \rho y_{ht-1} + \xi_{ht} + \theta\xi_{ht-1} + m_{ht} - \rho m_{ht-1}. \end{aligned}$$

The short run dynamics in the observed process will therefore be determined both by measurement error and by the short run dynamics in the true process θ (see Meghir and Pistaferri (2004) has an elaborated discussion of how to bound measurement error in a similar set-up). To show our identification strategy we start by observing that the parameter ρ is identified from the second order autocovariance

$$Cov(y_{ht}, y_{ht-2}) = \rho^2 V(y_{ht-2}).$$

For the short run dynamics we use the first order autocovariances and the variance

$$Cov(y_{ht} - \rho y_{ht-1}, y_{ht-1} - \rho y_{ht-2}) = \theta\nu^2 - \rho\sigma^2_m \tag{A.1}$$

$$Var(y_{ht} - \rho y_{ht-1}) = (1+\theta^2)\nu^2 + (1+\rho^2)\sigma^2_m \tag{A.2}$$

and cannot separately identify θ, ν^2 and σ^2_m unless we have additional information.

However, if we also have access to consumption information we can identify the short run dynamics and the measurement error separately. The idea is here that the consumption change $\Delta \ln C_{ht}$ reacts to an income shock, but only the true income shock ξ_{ht} and not the measurement error. This implies that

$$\Delta \ln C_{ht} = k + f(\nu\xi_{ht}),$$

We can now use two additional moment condition together (A.1) and (A.2) to identify ν^2, σ_m^2, θ and $Cov(\xi_{it}, f(\xi_{it}))$:

$$\begin{aligned} Cov(y_{ht} - \rho y_{ht-1}, \Delta \ln C_{ht}) &= Cov(\xi_{ht}, f(\xi_{ht})) \\ Cov(y_{ht} - \rho y_{ht-1}, \Delta \ln C_{ht-1}) &= \theta Cov(\xi_{ht-1}, f(\xi_{ht-1})) \end{aligned}$$

A.2 Estimation results

We started with a full model with ten factors (one for each model parameter) and have subsequently reduced the number of factors. The preferred model has 32 parameters and seven factors $(N_1, N_2, N_3, N_5, N_6, N_9, N_{10})$ and is given by:

$$\begin{aligned} \mu_h &= \phi_1 + \exp(\psi_{11}) N_{1h} \\ \alpha_h &= \phi_2 + \exp(\psi_{22}) N_{2h} \\ \rho_h &= \ell(\phi_3 + \exp(\psi_{33}) N_{3h}) \\ \theta_h &= 2l(\ell(\phi_4 + \psi_{41} N_{1h})) - 1 \\ \nu_h &= \exp(\phi_5 + \psi_{52} N_{2h} + \psi_{53} N_{3h} + \exp(\psi_{55}) N_{5h}) \\ \delta_h &= 0.1 * \ell(\phi_6 + \psi_{62} N_{2h} + \exp(\psi_{66}) N_{6h}) \\ \gamma_h &= 0.5 + 14.5 * \ell(\phi_7 + \psi_{76} N_{6h}) \\ \lambda_h &= \exp(\phi_8 + \psi_{83} N_{3h} + \psi_{86} N_{6h}) \\ m_h^y &= \exp(\phi_9 + \psi_{95} N_{5h} + \exp(\psi_{99}) * N_{9h}) \\ m_h^c &= \exp\left(\phi_{10} + \psi_{10,6} N_{6h} + \psi_{10,9} N_{9h} + \exp\left(\psi_{10,10}\right) * N_{10h}\right) \end{aligned}$$

where $\ell(x)$ is the transformation $e^x / (1 + e^x) \in (0, 1)$. The model contains 10 mean parameters (ϕ_j), 17 heterogeneity and co-dependence parameters (ψ_{ij}) and 5 homogeneous parameters $(b_0, b_1, \sigma_a, \sigma_b, d)$.

parameter	coef	se	t-val
ϕ_1	0.0452	0.0384	1.1795
ϕ_2	-0.0005	0.0035	0.1411
ϕ_3	1.7554	0.2222	7.9018
ϕ_4	0.4550	0.1594	2.8543
ϕ_5	-2.0603	0.0672	30.6773
ϕ_6	1.2097	1.9293	0.6270
ϕ_7	-0.0180	0.4045	0.0445
ϕ_8	0.1769	0.2898	0.6105
$\ln(\sigma_a)$	-1.5352	0.3926	3.9104
$\ln(\sigma_b)$	-2.4508	0.4647	5.2740
ψ_{11}	-2.0305	0.3084	6.5839
ψ_{22}	-4.8459	0.3540	13.6883
ψ_{33}	0.2607	0.1935	1.3479
ψ_{41}	0.5882	0.1195	4.9199
ψ_{52}	-0.0915	0.0323	2.8313
ψ_{53}	-0.4510	0.0595	7.5846
ψ_{55}	-0.9997	0.1617	6.1815
ψ_{62}	-0.6318	0.5893	1.0721
ψ_{66}	-0.2564	0.7817	0.3280
ψ_{76}	1.4590	0.1342	10.8748
ψ_{83}	0.8940	0.1151	7.7672
ψ_{86}	0.4984	0.2067	2.4117
ϕ_9	-2.2412	0.0664	33.7547
ϕ_{10}	-1.6583	0.0341	48.6655
ψ_{99}	-1.0168	0.1242	8.1875
$\psi_{10,10}$	-2.0052	0.1957	10.2460
$\psi_{10,9}$	0.3258	0.0538	6.0616
ψ_{95}	0.2859	0.0494	5.7878
$\psi_{10,6}$	0.3604	0.0476	7.5655
b_0	-0.4145	0.0563	7.3685
b_1	-0.6023	0.4940	1.2191
d(mix)	0.1829	0.0606	3.0158

Table A.1: Distribution parameters

AP	data	sim	se	t-val
$M(\hat{b}_{y1})$	-2.399	-2.402	0.388	0.007
$M(\hat{b}_{y2})$	0.043	0.041	0.007	0.280
$M(\hat{b}_{y3})$	4.263	4.405	0.172	-0.828
$M(\hat{b}_{y4})$	0.518	0.549	0.062	-0.506
$M(\hat{b}_{y5})$	1.695	1.712	0.037	-0.454
$M(\hat{b}_{c1})$	0.019	0.018	0.008	0.127
$M(\hat{b}_{c2})$	1.881	2.249	0.641	-0.574
$M(\hat{b}_{c3})$	0.677	0.634	0.122	0.349
$M(\hat{b}_{c4})$	2.300	2.290	0.053	0.194
$M(\hat{b}_{c6})$	-2.734	-3.217	0.143	3.371
$M(\hat{b}_{c7})$	0.325	0.321	0.100	0.043
$S(\hat{b}_{y1})$	10.666	10.776	0.567	-0.195
$S(\hat{b}_{y2})$	0.254	0.250	0.015	0.249
$S(\hat{b}_{y3})$	3.529	3.731	0.205	-0.983
$S(\hat{b}_{y4})$	1.919	1.829	0.092	0.982
$S(\hat{b}_{y5})$	1.323	1.199	0.092	1.351
$S(\hat{b}_{c1})$	0.204	0.204	0.009	-0.019
$S(\hat{b}_{c2})$	18.363	18.946	0.988	-0.590
$S(\hat{b}_{c3})$	4.544	4.351	0.288	0.670
$S(\hat{b}_{c4})$	1.131	1.109	0.071	0.313
$S(\hat{b}_{c6})$	3.035	2.843	0.187	1.025
$S(\hat{b}_{c7})$	2.909	3.140	0.164	-1.409

$M(.)$: mean, $S(.)$: standard deviation
Table continued

Table A.2: Auxiliary parameters

AP	data	sim	se	t-val
$C(\hat{b}_{y1}, \hat{b}_{y2})$	-9.302	-9.392	0.089	1.016
$C(\hat{b}_{y1}, \hat{b}_{y3})$	-0.544	0.049	0.439	-1.352
$C(\hat{b}_{y1}, \hat{b}_{y4})$	-0.422	0.295	0.440	-1.632
$C(\hat{b}_{y1}, \hat{b}_{y5})$	1.122	2.118	0.707	-1.409
$C(\hat{b}_{y2}, \hat{b}_{y3})$	0.972	0.403	0.437	1.304
$C(\hat{b}_{y2}, \hat{b}_{y4})$	0.737	-0.004	0.447	1.658
$C(\hat{b}_{y2}, \hat{b}_{y5})$	-0.658	-1.684	0.768	1.335
$C(\hat{b}_{y3}, \hat{b}_{y4})$	3.491	4.579	0.396	-2.747
$C(\hat{b}_{y3}, \hat{b}_{y5})$	-1.143	-0.795	0.450	-0.774
$C(\hat{b}_{y4}, \hat{b}_{y5})$	-0.544	-0.130	0.391	-1.057
$C(\hat{b}_{c1}, \hat{b}_{c2})$	-0.262	0.078	0.451	-0.755
$C(\hat{b}_{c1}, \hat{b}_{c3})$	0.250	-0.125	0.479	0.783
$C(\hat{b}_{c1}, \hat{b}_{c4})$	0.013	-0.083	0.509	0.189
$C(\hat{b}_{c1}, \hat{b}_{c6})$	0.507	0.703	0.457	-0.427
$C(\hat{b}_{c1}, \hat{b}_{c7})$	-0.178	0.384	0.412	-1.363
$C(\hat{b}_{c2}, \hat{b}_{c3})$	0.529	0.426	0.629	0.163
$C(\hat{b}_{c2}, \hat{b}_{c4})$	0.868	0.014	0.690	1.239
$C(\hat{b}_{c2}, \hat{b}_{c6})$	0.667	1.012	0.500	-0.692
$C(\hat{b}_{c2}, \hat{b}_{c7})$	0.438	0.406	0.397	0.082
$C(\hat{b}_{c3}, \hat{b}_{c4})$	-0.858	-0.040	0.544	-1.503
$C(\hat{b}_{c3}, \hat{b}_{c6})$	-0.348	-0.052	0.453	-0.653
$C(\hat{b}_{c3}, \hat{b}_{c7})$	-2.173	-2.218	0.420	0.107
$C(\hat{b}_{c4}, \hat{b}_{c6})$	-2.718	-3.016	0.459	0.649
$C(\hat{b}_{c4}, \hat{b}_{c7})$	0.254	0.240	0.406	0.035
$C(\hat{b}_{c6}, \hat{b}_{c7})$	0.556	0.016	0.393	1.376
$C(\hat{b}_{y1}, \hat{b}_{c1})$	-2.907	-1.948	0.475	-2.018
$C(\hat{b}_{y2}, \hat{b}_{c1})$	3.088	2.243	0.484	1.745
$C(\hat{b}_{y3}, \hat{b}_{c1})$	0.507	0.388	0.411	0.289
$C(\hat{b}_{y4}, \hat{b}_{c1})$	0.394	0.054	0.451	0.753
$C(\hat{b}_{y5}, \hat{b}_{c1})$	-0.021	0.317	0.426	-0.793
$C(\hat{b}_{y1}, \hat{b}_{c2})$	-0.045	-0.079	0.473	0.071
$C(\hat{b}_{y2}, \hat{b}_{c2})$	0.063	-0.068	0.471	0.277
$C(\hat{b}_{y3}, \hat{b}_{c2})$	0.000	-0.098	0.455	0.213
$C(\hat{b}_{y4}, \hat{b}_{c2})$	-0.261	0.061	0.404	-0.798
$C(\hat{b}_{y5}, \hat{b}_{c2})$	0.512	-0.124	0.422	1.506
$C(.,.)$: correlation				
Table continued				

Table A.2: Auxiliary parameters

AP	data	sim	se	t-val
$C(\hat{b}_{y1}, \hat{b}_{c3})$	-0.327	-0.137	0.384	-0.495
$C(\hat{b}_{y2}, \hat{b}_{c3})$	0.254	0.084	0.353	0.481
$C(\hat{b}_{y3}, \hat{b}_{c3})$	0.134	-0.002	0.465	0.292
$C(\hat{b}_{y4}, \hat{b}_{c3})$	0.006	0.047	0.397	-0.105
$C(\hat{b}_{y5}, \hat{b}_{c3})$	-0.039	-0.456	0.339	1.228
$C(\hat{b}_{y1}, \hat{b}_{c4})$	0.142	0.490	0.405	-0.857
$C(\hat{b}_{y2}, \hat{b}_{c4})$	-0.439	-0.490	0.377	0.136
$C(\hat{b}_{y3}, \hat{b}_{c4})$	-0.998	-1.193	0.412	0.472
$C(\hat{b}_{y4}, \hat{b}_{c4})$	-0.888	-1.131	0.393	0.619
$C(\hat{b}_{y5}, \hat{b}_{c4})$	1.680	1.192	0.442	1.105
$C(\hat{b}_{y1}, \hat{b}_{c6})$	-0.117	-0.015	0.384	-0.264
$C(\hat{b}_{y2}, \hat{b}_{c6})$	0.382	0.041	0.406	0.840
$C(\hat{b}_{y3}, \hat{b}_{c6})$	0.644	0.763	0.387	-0.307
$C(\hat{b}_{y4}, \hat{b}_{c6})$	1.225	0.532	0.418	1.658
$C(\hat{b}_{y5}, \hat{b}_{c6})$	-0.360	-0.189	0.427	-0.400
$C(\hat{b}_{y1}, \hat{b}_{c7})$	-0.025	-0.634	0.441	1.381
$C(\hat{b}_{y2}, \hat{b}_{c7})$	-0.064	0.557	0.439	-1.415
$C(\hat{b}_{y3}, \hat{b}_{c7})$	0.241	0.609	0.459	-0.800
$C(\hat{b}_{y4}, \hat{b}_{c7})$	-0.062	0.481	0.398	-1.367
$C(\hat{b}_{y5}, \hat{b}_{c7})$	0.531	0.070	0.391	1.181
Intercept$_{30}$	-0.168	-0.142	0.021	-1.194
Std ($\hat{e}_{30}$)	0.401	0.400	0.032	0.032
CS IQR	0.360	0.542	0.115	-1.584
OI test $\chi_{(48)}$	72.24			
$C(.,.)$: correlation				

Table A.2: Auxiliary parameters

AP	data	sim	se	t-val
$M(\hat{b}_{c5})$	-0.315	-0.124	0.082	-2.334
$S(\hat{b}_{c5})$	2.545	2.519	0.155	0.174
$C(\hat{b}_{c1}, \hat{b}_{c5})$	0.511	0.092	0.441	0.950
$C(\hat{b}_{c2}, \hat{b}_{c5})$	-0.179	-0.656	0.434	1.098
$C(\hat{b}_{c3}, \hat{b}_{c5})$	-1.878	-2.509	0.431	1.466
$C(\hat{b}_{c4}, \hat{b}_{c5})$	0.327	0.506	0.427	-0.419
$C(\hat{b}_{c6}, \hat{b}_{c5})$	-0.609	-0.441	0.464	-0.364
$C(\hat{b}_{c7}, \hat{b}_{c5})$	-0.117	-0.437	0.430	0.746
$C(\hat{b}_{y1}, \hat{b}_{c5})$	0.056	-0.079	0.494	0.273
$C(\hat{b}_{y2}, \hat{b}_{c5})$	-0.042	-0.054	0.507	0.022
$C(\hat{b}_{y3}, \hat{b}_{c5})$	0.180	-0.296	0.477	0.998
$C(\hat{b}_{y4}, \hat{b}_{c5})$	-0.112	-0.349	0.418	0.567
$C(\hat{b}_{y5}, \hat{b}_{c5})$	0.472	0.578	0.428	-0.248
$corr(\hat{w}_{ht}, \hat{u}_{ht} \cdot t)$	0.193	-0.304	0.351	1.418
GF test $\chi_{(14)}$	16.07			
$M(.)$:mean, $S(.)$: std, $C(.,.)$: correlation				

Table A.3: AP used for goodness of fit test

Online Appendix (Not for publication)

Monte Carlo experiments

To validate our methodology and show that we can recover the distribution of the model parameters, we perform a Monte Carlo experiment (MC). To make the exercise feasible, we use a simpler version of the model described in the paper, more specifically, a simpler version of the income process. We also ignore measurement error and taste shocks. We assume that for each period, household h finds the optimal consumption by maximizing the discounted expected utility subject to the 'natural' no borrowing constraint[22]

$$\begin{aligned} & \max_{C_{ht},..,C_{hT}} \frac{C_{ht}^{1-\gamma_h}}{1-\gamma_h} + \sum_{s=t+1}^{T} (1+\delta_h)^{-s} E_t(\frac{C_{hs}^{1-\gamma_h}}{1-\gamma_h}) \\ A_{ht+1} &= (1+r_{t+1})(A_{ht}+Y_{ht}-C_{ht}) \end{aligned}$$

where C_{ht}, Y_{ht} and A_{ht} are the consumption, income and assets of household h in period t. Note that the coefficient of relative risk aversion (γ_h) and the discount rate (δ_h) are assumed to be household specific. All households face the same interest rate r_t , which is assumed to follow an $AR(1)$ process such that:

$$r_{t+1} = \rho(1-\mu) + \rho r_t + \varepsilon_{t+1}$$

where μ and ρ are parameters of long-run mean and persistence, respectively and ε_{t+1} is iid normal.

For the income process we assume that log household income $\log Y_{ht}$, at age t for household

[22] That is, households are not allowed to end their lives with debt but they are allowed to borrow or save in all periods before T.

h, follows a unit root process such that:

$$\log Y_{ht} = \log Y_{ht-1} + \nu_h \xi_{ht} \qquad t \geq 2, \xi_{ht} \sim iiN(0,1) \tag{A.3}$$

$$\nu_h \xi_{ht} \sim iiN(-0.5\nu_h, \nu_h) \tag{A.4}$$

The income process contains only one household specific model parameter; namely the standard deviation of the income shocks, ν_h.

Thus we have a total of three household specific parameters: ν_h, γ_h and $\delta_{h.}$. We assume that the following joint distribution for these three model parameters:

$$\nu = \exp(\phi_1 + \exp(\psi_{11}) N_1) \tag{A.5}$$

$$\delta = 0.1 \frac{\exp(\phi_2 + \psi_{21} N_1 + \exp(\psi_{22}) N_2)}{1 + \exp(\phi_2 + \psi_{21} N_1 + \exp(\psi_{22}) N_2)} \tag{A.6}$$

$$\gamma = 9 \frac{\exp(\phi_3 + \psi_{31} N_1 + \psi_{32} N_2 + \exp(\psi_{33}) N_3)}{1 + \exp(\phi_3 + \psi_{31} N_1 + \psi_{32} N_2 + \exp(\psi_{33}) N_3)} \tag{A.7}$$

where N_1, N_2 and N_3 are independent standard normals.

To complete our specification for the SRE estimation we also need to consider parameters for which we do not know the true value. The first of these is the dependence of consumption on income shocks, λ_h, which we define by:

$$\lambda_h = \exp(\phi_4 + \psi_{41} N_1 + \psi_{42} N_2 + \psi_{43} N_3).$$

In the estimation, we allow λ_h to be age-dependent:

$$\lambda_{ht} = \exp(\psi_{4Age} \cdot (t-1)) * \lambda_h$$

so that λ_h is the dependence in the first period. We also estimate a homogenous parameter for the variance of the SRE non-income shock ϕ_5. This gives a total of 15 distribution parameters to be estimated.

The MC exercise is based on repeatedly estimating the nine known distribution parameters and the six unknown parameters:

$$\phi_1, \phi_2, \phi_3, \psi_{11}, .., \psi_{33}, \phi_4, \psi_{41}, .., \psi_{43}, \psi_{4Age}, \phi_5.$$

The MC steps

Below we describe the steps for the MC exercise. Each estimation is performed with a sample size of 600 households and 40 time periods; the number of MC replications is 100.

1. We first take values of the nine distribution parameters $\{\phi_1, ...\psi_{33}\}$; details of how these are chosen are given in the next subsection. We then use these values and equations $(A.5)$ to $(A.7)$ to simulate 600 values for $\{\nu_h, \delta_h, \gamma_h\}$. These are the (known) distribution parameters that we seek to recover.

2. For each set of model parameters $\{\nu_h, \delta_h, \gamma_h\}$, we solve the dynamic program for 60 periods via standard policy function iteration and obtain policy functions for all 600 households. To do this, we discretize income using a 10-point Gaussian quadrature, and interest rate process, following Tauchen (1986) using 10 nodes. For the latter we assume that the AR parameter is equal to 0.6, the the long-run mean is equal to 0.05 and the standard deviation of interest rate shocks is 0.025. Note that discretizing the income process gives a positive effective lower bound for income growth in each period.

3. With all 600 policy functions in hand, assuming zero assets in the initial period, we simulate consumption paths using household specific income paths and a common interest rate path. This gives income and consumption paths for 600 types based on a conventional dynamic program.

4. Treating a given sample of 600 simulated households as actual data, we calculate the 'data'

ap's to be matched in the SMD estimation. Since we use a simpler income process, for income we only use:

$$\hat{b}_{h,y1} = std(\Delta \log Y_{ht}).$$

Ap's for consumption are based on the coefficient of household-specific regressions $\hat{b}_{h,c1}$ to $\hat{b}_{h,c4}$ (as described in the paper). The ap's are then constructed as the median of $\hat{b}_{h,y1}, \hat{b}_{h,c1} - \hat{b}_{h,c4}$, the interquartile range of $\hat{b}_{h,y1}, \hat{b}_{h,c1} - \hat{b}_{h,c3}$ and the correlations between $\hat{b}_{h,y1}, \hat{b}_{h,c1} - \hat{b}_{h,c3}$. This gives $5 + 4 + 6 = 15$ ap's. A final ap is used to capture the age dependence in λ_{ht}; constructed as the correlation between consumption residuals, $\hat{w}_{ht}$, and $t \cdot \hat{u}_{ht}$, where $\hat{u}_{ht}$ is the income residual, see Equations 19 and 18. This gives a total of 16 ap's to match.

5. Next we estimate the distribution parameters using the SRE procedure described in the paper. The model is over-identified with one degree of freedom.

6. From the estimated distribution parameters $\hat{\phi}_1, \hat{\phi}_2, \hat{\phi}_3, \hat{\psi}_{11}, .., \hat{\psi}_{33}$, we construct our model parameters $\hat{\nu}_h, \hat{\gamma}_h, \hat{\delta}_h$. We also construct parameters relating to the impact of income shock and variance of non-income shock: $\hat{\lambda}_h, 1 - \hat{\vartheta}_h$ and $\hat{\kappa}_h$. For the model parameters and these extra parameters, we calculate the mean, standard deviation and correlations for the sample of 600 households.

7. We repeat step $3 - 6$ 100 times.

Results of the MC

We perform two MC experiments. In the first experiment (A), we choose the distribution parameters close to those estimated in the data (PSID). As we found that the correlation between ν_h and γ_h was close to zero in the PSID, we set ψ_{31} in equation $(A.7)$ to zero. In the second experiment (B), we assume a negative correlation between ν_h and γ_h and a higher level of risk aversion but keep the remaining parameters close to what found in the data. The latter experiment is performed to ensure that our procedure can detect the correlation between ν_h and γ_h, if it exists.

In table A.4 and A.5 we present the results. Table A.4 presents the estimated distribution parameters and table A.5 the distribution of the implied model parameters. The results show that we fit the distribution parameters, in particular the distribution of ν and γ, quite well. We slightly underestimate ϕ_2 in both experiments, which also pushes the estimated mean of δ to be 0.05 while the true mean is 0.06 (see Table A.5). The correlations between ν, δ and γ are well captured. We see that in experiment A (with no correlation between ν and γ) the correlation between ν and γ is estimated to be -0.08 (ψ_{31} the mean estimate is -0.02). In experiment B, where the true correlation coefficient is -0.33 ($\psi_{31} = -0.3$), our estimate is -0.28 (ψ_{31} the mean estimate is -0.25 and significantly different from zero). This suggests that our method is able to recover the correlation between the model parameters and that we can detect a correlation between e.g. ν and γ if it exists.

Our simulation exercise also provides us with the information of the 'non-structural' parameters λ and κ. Recall that λ is closely related to the partial insurance coefficient defined in Blundell et al (2008). The partial insurance coefficient can in our terminology be defined as $1 - \vartheta = 1 - \lambda/\gamma$. Kaplan and Violante (2010), show that the partial insurance coefficient depends on the structural parameters ν and γ and on age. In our framework, we allow for this flexible dependence through the parameters $\psi_{4Age}, \phi_4, \psi_{41}, \psi_{42}$ and ψ_{43}. Our estimated partial insurance coefficient is 0.17 at age 1 (See Table A.6). Given the negative age dependence of λ, we show that the partial insurance coefficient increases in age and at age 40 the mean insurance coefficient is 0.42. These estimates are consistent with findings of Kaplan and Violante (2010). They estimate the average insurance coefficient to be about 0.23 and show also that the coefficient is increasing in age. Moreover, we find a strong positive correlation between the insurance coefficient and the coefficient of relative risk aversion, a finding that is also consistent with Kaplan and Violante (2010). Kaplan and Violante explain their results as older households and households with higher risk aversion accumulating more wealth and therefore being better insured. Supporting this explanation, we find a small but positive correlation between income variance and insurance coefficient. In addition to the results in Kaplan and Violante (2010), we also consider the correlation between the discount rate and the insurance coefficient. This correlation is significantly negative in experiment A and positive but not significant in experiment B.[23] The negative correlation is consistent with households with a high discount rate accumulating less wealth and therefore have a lower degree of insurance.

Finally, we examine the estimated variance of the non-income shocks in consumption shocks. In this set up, we estimate that the income shock explains $95 - 99$ percent of the variation of the consumption shock at age 1 (see Table A.7). This is not surprising, since we do not allow for taste shocks and the only uncertainty is generated by income shocks and interest rate shocks. The non-income shocks in our model consist of unanticipated interest rate shocks and approximation errors. It is therefore encouraging to see that these approximation errors are small and only account for 1-5 percent of the variation in consumption shocks. This result is particularly encouraging since it indicates that although we do not know the true relationship between income and consumption shocks or the true dependence between preference parameters and insurance parameter our approximation works reasonably well.

In sum, our MC results indicate that our estimation methodology is valid and that we in fact can recover the structural distribution parameters. Moreover, our investigation of the insurance coefficient and the ratio of income shocks provide additional support for the validity of our estimation methodology, since our estimates are consistent with those found in the previous literature.

[23]Notice that $1 - \vartheta = 1 - \lambda/\gamma$ and that in table A.4, we see that the estimate of ψ_{42} is positive and significant in experiment A and negative and insignificant in experiment B.

	A			B		
	True	Mean	Std.	True	Mean	std.
ϕ_1	-2.3	-2.29	0.007	-2.3	-2.29	0.007
ϕ_2	0.5	0.20	0.068	0.5	0.14	0.133
ϕ_3	-1.0	-1.02	0.037	0.0	-0.10	0.113
ψ_{11}	-2.3	-2.16	0.081	-2.3	-2.18	0.096
ψ_{21}	0.2	0.23	0.058	0.2	0.28	0.136
ψ_{22}	-0.5	-0.74	0.072	-0.5	-1.18	0.377
ψ_{31}	0.0	-0.02	0.045	-0.3	-0.25	0.128
ψ_{32}	-0.1	-0.10	0.032	0.7	0.60	0.131
ψ_{33}	-0.5	-0.51	0.068	-0.3	-0.73	0.275
ϕ_4	.	1.83	0.067	.	1.68	0.103
ϕ_5	.	-3.44	0.052	.	-2.38	0.053
ψ_{41}	.	-0.20	0.056	.	-0.10	0.057
ψ_{42}	.	0.88	0.069	.	-0.42	0.315
ψ_{43}	.	-0.35	0.105	.	-0.21	0.332
ψ_{4A} $(*10)$	.	-0.09	0.038	.	-0.12	0.006

Table A.4: MC results for the distribution parameters

	A			B		
	True	Mean	Std.	True	Mean	std.
mean(ν)	0.10	0.10	0.001	0.10	0.10	0.001
mean(δ)	0.06	0.05	0.002	0.06	0.05	0.003
mean(γ)	3.54	3.50	0.067	5.50	5.28	0.226
std(ν)	0.01	0.01	0.001	0.01	0.01	0.001
std(δ)	0.01	0.01	0.001	0.01	0.01	0.003
std(γ)	1.02	1.02	0.066	1.95	1.64	0.211
corr(ν,δ)	0.34	0.45	0.086	0.34	0.62	0.191
corr(ν,γ)	-0.08	-0.05	0.076	-0.33	-0.28	0.151
corr(δ,γ)	-0.18	-0.15	0.074	0.50	0.37	0.157

Table A.5: MC results for the model parameters

	A		B	
	Mean	Std	Mean	Std
mean($1-\vartheta$) at age 1	0.17	0.006	0.18	0.011
std($1-\vartheta$)	0.13	0.007	0.09	0.011
$Corr(1-\vartheta,\nu)$	0.14	0.054	0.16	0.080
$Corr(1-\vartheta,\delta)$	-0.70	0.057	0.54	0.376
$Corr(1-\vartheta,\gamma)$	0.45	0.080	0.64	0.217

Table A.6: Partiel insurance

	A		B	
	Mean	Std	Mean	Std
mean(κ)	0.985	0.001	0.954	0.002
std(κ)	0.01	0.001	0.028	0.005

Table A.7: The ratio of variance of income shocks to total variance of consumption shocks

References

[1] Blundell R., Luigi Pistaferri, Ian Preston (2008), "Consumption Inequality and Partial Insurance", American Economic Review, Vol. 98, No. 5 (Dec., 2008), pp. 1887-1921.

[2] Kaplan, Greg, and Giovanni L. Violante (2010), "How Much Consumption Insurance beyond Self-Insurance?" American Economic Journal: Macroeconomics, 2(4): 53-87.

[3] Tauchen, George. 1986. “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions” Economic Letters. 20: 177–181.