

Income and Consumption: a Micro Semi-structural Analysis with Pervasive Heterogeneity¹

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

March 2017

¹We 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 joint dynamics of income and consumption is crucial for many research and policy issues such as 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. It is now well established that households have idiosyncratic income processes and ample evidence points to a high degree of heterogeneity in intertemporal preference parameters. In the literature on income processes, studies such as Baker (1997), Rubinstein and Weiss (2006), Guvenen (2009) and Browning *et al* (2010) document heterogeneity between individuals in, for example, the income trend and variance of income shocks. Likewise 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 find a large degree of heterogeneity in intertemporal preferences. There is also evidence that the within process heterogeneity in these parameters is co-dependent. For example, for income, the persistence of shocks is correlated with the variance of shocks and for consumption, the discount rate is correlated with risk aversion; see Browning *et al* (2010) and Alan and Browning (2010).

The paper offers three main contributions. First, we introduce co-dependence between household income process parameters and preference parameters. To the best of our knowledge, there is no study that models income and consumption jointly while allowing for heterogeneity in parameters as well as co-dependence *across* parameters of both processes.¹ The second

¹Blundell *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.

contribution is to allow the impact of income shocks on consumption to be heterogeneous across households and co-dependent with income and preference parameters. The third contribution is that we provide a methodology which allows us to theoretically postulate and then empirically quantify the extent of co-dependent heterogeneity in systems of processes. This framework is sufficiently flexible so that household level heterogeneity in both preference parameters and income process parameters can be accounted for, and conventional tests such as excess sensitivity can be conducted.

Co-dependence between the income and consumption processes is a priori 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 may 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* (2014) provide evidence that self-reported risk aversion measures correlate with the earnings risk of chosen occupations.

The presence of co-dependent heterogeneity in consumption and income has implications for both normative and positive analyses in economics. For 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, however, 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), where agents are assumed to be ex-ante identical.

In terms of positive analysis, consider 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 significantly 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.

Incorporating co-dependent ‘pervasive’ heterogeneity in consumption and income requires modelling the joint income and consumption processes for a given household. For (log) household income, we follow Browning *et al* (2010) and specify a trend stationary ARMA model with five household specific parameters. These parameters jointly capture the initial level, trend, long-run dynamics (AR(1) component), short-run dynamics (MA(1) component) and the variance of shocks. For consumption, using the first order condition (the exact Euler equation) obtained from a standard dynamic consumption model, we decompose shocks to the marginal utility of consumption into income and non-income components. Feeding contemporaneous income shocks into the marginal utility shocks in a parametric fashion, we establish a direct link between the income and consumption processes. Independent non-income shocks are modeled flexibly following Alan and Browning (2010). All the income model parameters and the consumption model parameters are assumed to be heterogeneous and co-dependent. In section 2, we lay out the details of this model of income and consumption. We then develop a parametric

factor structure that captures household level heterogeneity within and across two processes.

In section 3 we describe the longitudinal consumption and income information in the Panel Studies of Income Dynamics (PSID). In section 4 we present an indirect inference estimation procedure that requires simulation of the fully parametric model to generate simulated paths of income and consumption growth. Indirect 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 empirical co-dependence. To illustrate, consider how we assess the location and dispersion of income shock variances. For each household we run a regression of current income on lagged income and a trend and take the variance of the residuals. The distribution of these estimated variances cannot be used directly to generate an estimate of the distribution of the income shock variances. This is due to issues such as small-T dynamic regression bias, the presence of a moving average component and measurement error. Rather, the mean and variance of the individual estimates from the data and from the simulated data are matched; the logic of indirect inference is that the bias is the same for the two sources if we have the true data generating process. Similarly for the consumption process, consider estimating the effect of income shocks on consumption. We first run individual linear Euler equations for each household, including the income surprises on the right hand side. Then we use the mean and variance of the estimates of the income surprise coefficients to give ap's for the dependence of consumption changes on income shocks. Finally, we can use the correlation between the two sets of regression based estimates (income shock variances and the effect of an income shock on consumption) to capture the co-dependence between the income shock variance and the effect of income shocks on consumption.

We present our results 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 though 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. That is, some households are subject to large shocks which die out quickly and others have the opposite: small but persistent income shocks.

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 [5.3%, 8.0%, 9.3%] and [2.2, 7.6, 12.9] respectively. We find that 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. We also find a weak negative 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 6.9%, which indicates that even those who react most can still achieve some 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.

The strongest assumption we make in all of our analysis is that no household faces a liquidity constraint in any period. Although the sample we select is a sample of households who are less likely to be constrained and a test for ‘excess sensitivity’ fails to show any evidence of liquidity constraints, we acknowledge that some households may sometimes be constrained. We consequently round off the results section with an analysis of the bias that would be induced if

some households are sometimes constrained. The principal conclusion from this analysis is that constraints would lead to an under-estimate of the extent of heterogeneity in the population.

There are two broad implications of our results. First, as put forward above via a specific example of unemployment insurance, heterogeneity in impact of income shocks on consumption may fundamentally change the way normative analyses in economics are conducted. Second, 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). This is echoed in the recent dynamic public finance, which puts heterogeneity and uncertainty over future earnings at the heart of the analysis; see Kocherlakota (2010) and Farhi and Werning (2012). These implications 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. In the appendix A.1.1 we show how this ARMA representation is linked to the conventional permanent-transitory income model. For each household the log income process is:

$$y_{ht} = \{\mu_h(1 - \rho_h) + \alpha_h\rho_h\} + \rho_h y_{h,t-1} + (1 - \rho_h)\alpha_h(t - 1) + \nu_h(\xi_{ht} + \theta_h\xi_{h,t-1}) \quad (1)$$

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. We model the initial condition, y_{h1} , by a parametric model with two homogeneous parameters, as given in Appendix A.1.2, equation (A.1). Note that $\rho_h = 1$ implies a unit root income process with an idiosyncratic drift given by α_h and an $MA(1)$ stochastic component.

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 these parameters to be co-dependent. For example, as well as allowing for heterogeneity in the long run impact of an income shock, ρ_h , and the standard deviation of the shocks, ν_h , it may be that the two are correlated with, say, high variance households facing less persistent shocks.

Even if we have a trend stationary process ($\rho_h < 1$), an income shock can have a persistent impact on consumption through the consecutive revisions 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 as²:

$$\tau_h = \frac{1 + \theta_h}{1 - \rho_h} \quad (2)$$

In the subsequent analysis when we relate consumption changes to income shocks; the value of τ_h will be important since a higher value for τ_h implies that the revision to lifetime income that drives the consumption change is higher.

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 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 may change over time (as in Moffitt and Gottschalk (2012)) or over age (as in Blundell *et al* (2015)). These

²Strictly 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.

assumptions allow us to focus on 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.

2.2 Consumption

To model consumption we use the standard intertemporal consumption model and specify a consumption process based on the exact Euler equation. Our specification imposes a number of assumptions on the process. First, it treats the household as a unit which implies that husband and wife are assumed to have the same intertemporal allocation parameters. Second, we take an iso-elastic utility function with exponential discounting. Third, we do not allow for cross-sectional heterogeneity in the real interest rate r_t ; this rules out, for example, differences between borrowing and lending rates. Fourth, we assume that there are no liquidity constraints.

Of the many explicit and implicit assumptions we have made, the most problematic is that our households do not face any liquidity constraints. Without this, the Euler equation becomes an inequality and any estimation based on the Euler equation will no longer be valid. Ultimately a satisfactory approach to allowing for liquidity constraints or cross-section variation in interest rates requires better data than we currently have. At a minimum we require information on assets carried forward from period to period and the actual interest rate that households face. The closest the PSID has is a question of whether the household has two months income in liquid assets; this is a sufficient condition for not being constrained but it is not necessary. Restricting attention to periods when households report that they are carrying forward such liquid assets would lose a lot of observations. Given this, we make the assumptions stated. This is not quite an appeal to blind faith. First, the sample we draw is of continuously married households aged over 30, who are in the PSID for at least 15 years and, whose heads have education above high school; such households are less likely to be liquidity constrained. Second, in the empirical

analysis we specify a quasi-Lagrange multiplier (QLM) test to test for ‘excess sensitivity’ to anticipated income changes; this would be a symptom of a violation of the assumption of no liquidity constraints. We do not find any evidence of excess sensitivity. In Section 5.6, after we present our results, we discuss the likely biases that would arise if some households in our sample are sometimes constrained.

The consumption Euler equation for household h is given by:

$$E_t \left[\frac{1 + r_{t+1}}{1 + \delta_h} (C_{h,t+1})^{-\gamma_h} \right] = (C_{ht})^{-\gamma_h} \quad (3)$$

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 $t + 1$ and $E_t(\cdot)$ is the expectations operator conditional on information available at time t . Equation (3) can be written as:

$$\left(\frac{1 + r_{t+1}}{1 + \delta_h} \right) \left(\frac{C_{h,t+1}}{C_{ht}} \right)^{-\gamma_h} = \varepsilon_{h,t+1} \quad (4)$$

where $\varepsilon_{h,t+1}$ is a shock to the marginal utility of expenditure (mue) and $E_t(\varepsilon_{h,t+1}) = 1$.

2.3 Consumption shocks

To estimate the structural parameters, we need to simulate individual income and consumption paths. This can be done using conventional dynamic programming methods. However, allowing for pervasive individual level heterogeneity requires solving a life cycle model for each of hundreds of vectors of model parameters for the income and consumption processes. Embedding this in an optimization routine is currently infeasible. Instead, we follow Alan and Browning (2010) and employ *synthetic residuals* to simulate consumption paths using equation (4).³ This

³We have carried out a Monte Carlo experiment to validate the SRE method under pervasive heterogeneity. 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 allow for pervasive

is based on the finding in Alan and Browning (2010) that for a simulated population with heterogeneous iso-elastic preferences, the distribution of the pooled mue shocks is well approximated by a mixture of two log-normals.

To establish the link between income shocks and mue shocks, we extend Alan and Browning (2010) by decomposing the mue shock into a non-income shock and a function of the income shock. We define the total mue shock for a given household h and time t as:

$$\varepsilon_{ht} = \tilde{\varepsilon}_{ht} \cdot g(\nu_h \xi_{ht}; \lambda_h) \quad (5)$$

where $\tilde{\varepsilon}_{ht}$ is a ‘non-income mue shock’ and $g(\nu_h \xi_{ht}; \lambda_h)$ is an ‘income mue shock’ which depends on the contemporaneous income shock, $\nu_h \xi_{ht}$ and a sensitivity parameter λ_h ; the functional form of $g(\nu_h \xi_{ht}; \lambda_h)$ is given in the next paragraph. The two types of shocks are assumed to be independent and each to have a unit mean. The non-income mue shock captures ‘pure’ consumption shocks and shocks that arise from unanticipated changes in variables such as wealth and demographics. Appendix A.1.3 presents the details of the specification of the distribution of the non-income mue shocks, $\tilde{\varepsilon}_{ht}$.

For the function $g(\cdot)$ we assume the following form:

$$g(\nu_h \xi_{ht}) = \exp \left(-\frac{(\lambda_h \nu_h)^2}{2} - \lambda_h \nu_h \xi_{ht} \right) \quad (6)$$

This specification ensures a unit mean and implies that the dependence of the mue on the income shock is governed linearly by a single heterogeneous parameter, λ , such that:

$$\frac{\partial \ln(g)}{\partial(\nu_h \xi_{ht})} = -\lambda_h$$

correlated heterogeneity in preference and income parameters, a full-fledged structural modeling is currently infeasible. The full description of the Monte Carlo experiment is available on request or can be obtained from <https://sites.google.com/site/salancrossley/publications>.

This gives a direct link to the partial insurance parameter discussed in Blundell *et al* (2008). The semi-elasticity parameter λ_h is the idiosyncratic response of log mue to a contemporaneous income shock. Note that λ is not a structural parameter, but functionally depends on all the preference and income parameters, and therefore, heterogenous across households. For example, the value of λ_h will depend on the persistence of income shocks and the aversion to risk (see Kaplan and Violante (2010)). Theoretically, we should allow that the reaction to an income shock varies with age and/or time because of the time-varying cash on-hand; see Kaplan and Violante (2010). However, as we discuss later in our empirical analysis, we did not find any evidence of age dependence so we assume away 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.

Although λ_h is not a structural parameter, it does have a useful interpretation. Taking the log of (4), substituting in (5) 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):

$$d \ln C_{h,t} = -\frac{1}{\gamma_h} \frac{\partial \ln(g)}{\partial (\nu_h \xi_{ht})} d(\nu_h \xi_{ht}) = \frac{\lambda_h}{\gamma_h} d(\nu_h \xi_{ht}) \quad (7)$$

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

$$\vartheta_h = \frac{\lambda_h}{\gamma_h} \quad (8)$$

Thus the change in consumption is larger the less the household is averse to fluctuations (low γ_h) and the higher is the sensitivity parameter λ_h . 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 type of shocks. In this study, we focus on household level heterogeneity and allow for the response to an income shock to vary systematically across households.

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 these eight household specific parameters as the *model parameters* to distinguish them from *distribution parameters* and *auxiliary 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 appendix A.1.4). We assume heterogeneity in the standard deviations of the measurement errors: m_h^y (income) and m_h^c (consumption). We allow for these variances to be correlated with the model parameters so that, for example, households with a low variance of income shocks report income more accurately. We also allow that the income and consumption measurement variances are correlated with each other. This gives two more heterogeneous parameters (m_h^y and m_h^c) to estimate in addition to the five income parameters and the three consumption parameters.

2.5 Accounting for heterogeneity

We model the joint distribution of the eight model parameters and the two parameters for measurement error variances using a factor structure with (standard normal) factors denoted by N_{kh} . The full model has ten factors (one for each model parameter), yielding a flexible correlational structure amongst the model parameters. The distribution parameters are denoted ϕ (location) and ψ (dispersion). The model parameters for the income process are specified as a five factor triangular model:

$$\begin{aligned}
\mu_h &= \phi_1 + \exp(\psi_{11}) N_{1h} \\
\alpha_h &= \phi_2 + \psi_{21} N_{1h} + \exp(\psi_{22}) N_{2h} \\
\rho_h &= \ell(\phi_3 + \psi_{31} N_{1h} + \psi_{32} N_{2h} + \exp(\psi_{33}) N_{3h}) \\
\theta_h &= 2 * \ell\left(\phi_4 + \sum_{j=1}^3 \psi_{4j} N_{jh} + \exp(\psi_{44}) N_{4h}\right) - 1 \\
\nu_h &= \exp\left(\phi_5 + \sum_{j=1}^4 \psi_{5j} N_{jh} + \exp(\psi_{55}) N_{5h}\right)
\end{aligned} \tag{9}$$

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} ($j \neq k$) 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:

$$\begin{aligned}
\delta_h &= 0.1 * \ell\left(\phi_6 + \sum_{j=1}^5 \psi_{6j} N_{jh} + \exp(\psi_{66}) N_{6h}\right) \\
\gamma_h &= 0.5 + 14.5 * \ell\left(\phi_7 + \sum_{j=1}^6 \psi_{7j} N_{jh} + \exp(\psi_{77}) N_{7h}\right) \\
\lambda_h &= \exp\left(\phi_8 + \sum_{j=1}^7 \psi_{8j} N_{jh} + \exp(\psi_{88}) N_{8h}\right)
\end{aligned} \tag{10}$$

The parameter restrictions are the result of a prior specification search and give $\delta \in (0, 0.1)$, $\gamma_h \in (0.5, 15)$ 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 standard deviations of measurement error by:

$$\begin{aligned} m_h^y &= \exp \left(\phi_9 + \sum_{j=1}^8 \psi_{9j} N_{jh} + \exp(\psi_{99}) * N_{9h} \right) \\ m_h^c &= \exp \left(\phi_{10} + \sum_{j=1}^9 \psi_{10,j} N_{jh} + \exp(\psi_{10,10}) * N_{10h} \right) \end{aligned} \quad (11)$$

for income and consumption respectively. With this structure we allow for the standard deviation 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 $\psi_{10,j}$.

Beside the set of homogenous parameters that capture initial conditions and non-income measurement shocks (see Appendix equations A.1 and A.2), the parameters ϕ and ψ , which describe the distribution of the model parameters are:

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

In our general factor model of the joint distribution 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 annual household income and information on food at home and food at restaurants.⁴ 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 their sample period. All our households are headed by males, and we select husbands whose 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 years⁵ and households with very low income ($< \$1$) or very large changes in consumption (more than 400%) or income (200%).⁶ 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.

⁴The 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.

⁵This implies that a household observed in 1968 should be observed at least until 1984 (17 calendar 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.

⁶By the last selection criteria we exclude nine households.

As a measure of income we use total family income deflated by the consumer price index. We take account of the fact that the measure of income refers to the previous calendar year. The consumption measure contains the total value of all food consumed by the household (including money spent 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 demographic variables, age and time dummies. This approach is problematic for three reasons. First, it does not have a sound theoretical base. 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 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. Details on how the simulation is implemented are given in appendix A.2.

4.1 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.⁷ Our construction of auxiliary parameters relies heavily on individual regressions for income and consumption growth. For example, for each household we regress income on a constant and a trend, and obtain a household specific estimate of the trend in income. The average of all household trend-estimates can then serve as an ap for the distribution parameter ϕ_2 in equation (9). Browning and Ejr n es (2013) provide a discussion of the advantages of using *individual regression based* (IRB) auxiliary models in the estimation of dynamic panel models. One problem that immediately arises for our data is that there are years

⁷The 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.

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. Below, we give a detailed account of the statistics we use to construct our ap's.

4.1.1 Income

Denote the first and last periods 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 time trend, t , for each household separately:

$$y_{ht} = b_{y1} + b_{y2} \cdot t + e_{ht} \text{ for } t = t_{hf}, ..t_{hl} \quad (12)$$

and record $(\hat{b}_{y1}, \hat{b}_{y2})$.⁸ The H estimates of $(\hat{b}_{y1}, \hat{b}_{y2})$ 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:

$$\hat{e}_{ht} = \text{constant} + b_{y3}\hat{e}_{ht-1} + u_{ht} \text{ for } t = t_{hf} + 1, ..t_{hl} \quad (13)$$

and record \hat{b}_{y3} , which will contain information on the AR parameter ρ_h . Of course, \hat{b}_{y3} is not an unbiased estimate of the AR parameter ρ_h due to the presence of short run dynamics and

⁸We suppress the index h to avoid triple indices.

the small- T sample. However, it does depend on the distribution of ρ_h , which is all we need for identification. Denote the expected value and residuals from this regression by \check{e}_{ht} and \hat{u}_{ht} , respectively. We then take the residuals from this regression and calculate the auto-correlation and the standard deviation:

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

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 $\{\hat{b}_{y1}, \hat{b}_{y2}, \hat{b}_{y3}, \hat{b}_{y4}, \hat{b}_{y5}\}$ 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.

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

$$\Delta c_{ht} = \text{constant} + b_{c2}r_t + b_{c3}\hat{u}_{ht} + w_{ht} \tag{14}$$

and record $(\hat{b}_{c2}, \hat{b}_{c3})$. Here, \hat{b}_{c2} captures the household specific elasticity of intertemporal substitution (the inverse of γ_h) and \hat{b}_{c3} captures the consumption response to contemporaneous income shocks. The estimates $(\hat{b}_{c1}, \hat{b}_{c2}, \hat{b}_{c3})$ characterize the distribution of the preference parameters δ_h , γ_h and the partial insurance parameter λ_h .

In the following we specify additional ap's used for estimation of the parameters of the non-income shock and the variances of the measurement error. Next, denote the Euler equation residuals by \hat{w}_{ht} and calculate the following the standard deviation and correlation coefficients:

$$\begin{aligned}
\hat{b}_{c4} &= \text{std}(\hat{w}_{ht}) \\
\hat{b}_{c5} &= \text{corr}(\hat{w}_{ht}, \check{e}_{ht}) \\
\hat{b}_{c6} &= \text{corr}(\hat{w}_{ht}, \hat{w}_{h,t-1}) \\
\hat{b}_{c7} &= \text{corr}(\hat{w}_{ht}, \hat{u}_{h,t-1})
\end{aligned} \tag{15}$$

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 (13)); this is to test 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 B.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.

To provide information on the shape of the distribution of the non-income mue shocks, we construct two extra ap's. These two ap's are based on the pooled residuals from the consumption changes, \hat{w}_{ht} . We calculate the skewness, $sk(\hat{w}_{ht})$ and kurtosis, $kurt(\hat{w}_{ht})$ of all residuals and add them as 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 age in period t , a_{ht} : that is, $corr(\hat{w}_{ht}, \hat{u}_{ht} \cdot a_{ht})$.⁹ 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 (A.1). 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, $m(y_{h30})$ and the standard deviation of the residuals, $std(y_{h30})$, 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 96 ap's to fit, which provide a rich characterization of the joint distribution of consumption and income parameters for our PSID sample. In the estimation step we use 82 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' (a test for liquidity

⁹Our Monte Carlo study confirms that this ap would indeed pick up a common age-dependence in λ .

constraints), and the α that picks up any age-dependence in λ_h .

Table 1 summarizes the above discussion. In particular, it lists the parameters we seek to estimate and the auxiliary parameters we use to identify them.

Model parameters:		Sources of auxiliary parameters	
Income model parameters			
μ	Mean at start of process	\hat{b}_{y1}	Income mean (corrected for trend)
α	Trend	\hat{b}_{y2}	Income trend
ρ	AR parameter	\hat{b}_{y3}	Autocorrelation in income residuals
θ	MA parameter	\hat{b}_{y4}	Autocorrelation corrected for AR(1)
ν	Shock standard deviation	\hat{b}_{y5}	Standard deviation of income residuals
Consumption model parameters			
δ	Discount rate	\hat{b}_{c1}	log consumption reg. on trend
γ	Coefficient of relative risk aversion	\hat{b}_{c2}	log growth consumption reg. on r_t
λ	Response of μ to income shocks	\hat{b}_{c3}	log growth consumption reg. on income shock
Measurement error parameters			
m^y	Variance in income	\hat{b}_{c7}	Correlation of cons. residuals and lagged income residuals
m^c	Variance in consumption	\hat{b}_{c6}	Correlation of cons. residuals and lagged cons. residuals
Initial income obs. (homogenous parameters)^a			
b_0	Mean of initial obs.	$m(y_{h30})$	Mean income at 30 (corrected for cohort effects)
b_1	Std. of initial obs.	$std(y_{h30})$	Std. of income at 30 (corrected for cohort effects)
Distribution of non-income μ shock (homogenous parameter)^b			
d	Skewness	$sk(\hat{w}_{ht})$	Skewness of consumption shock
σ_a, σ_b	Variance and kurtosis	\hat{b}_{c4}	Std. of consumption shock
		$kurt(\hat{w}_{ht})$	Kurtosis of consumption shock

Note: a) The specification of initial income observation, see Appendix A.1.2.

b) The specification of non-income μ shocks, see Appendix A.1.3.

Table 1: Model parameters and the statistics used for the auxiliary parameters

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 82 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} in equation (9)) 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 33 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 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 Cameron and Trivedi (2005). 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} (see equation (10)), has a low 't-value' of 1.35, but a high $\chi^2(1)$ QLR statistic of 7.5 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.

Given that our preferred model has 33 parameters to estimate and 82 ap's to match, we have 49 degrees of freedom. The over-identification (OI) test statistic is 79.59 so that the overall fit is marginal. The fits for most ap's are good; see Table B.1 in the online 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 (15). This has a data value of -2.73 and a simulated value of -3.18 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, although the simulated value is 0.56 , somewhat higher than the data value of 0.36 with a t-value of -1.77 (see the ap labelled as CS IQR in Table B.1 in the online Appendix). Our model is able to produce the increasing cross sectional variance in consumption, one of the most cited stylized facts in the literature.

The value for the QLM test for the 14 ap's not used in fitting is 15.2 , which has a $\chi^2(14)$ distribution. The first 13 ap's that we keep back for this test relate to excess sensitivity; the test statistics for excess sensitivity test is 12.02 , which has a $\chi^2(13)$ distribution. The low QLM statistic implies no evidence of excess sensitivity. This is the most direct evidence we have that liquidity constraints are not important for our sample (see Appendix Table A.2). The last ap, $corr(\hat{w}_{ht}, \hat{u}_{ht} \cdot a_{ht})$, is the QLM test which captures the potential age-dependence in the partial insurance parameter λ . The GF test (see Appendix Table A.2) indicates that our preferred specification, without an age-dependent λ , fits the ap reasonably well. This is also confirmed, when we estimated an extended version of the model with an age-dependent λ_h .¹⁰ The QLR statistics for the extended model against our preferred model is 1.9 , which has a $\chi^2(1)$. We therefore, conclude that we do not need age-dependence in λ_h . Hence, despite the fact that a standard life-cycle model implies an age-dependence in self-insurance, we find no statistically significant age-dependence in λ . This finding is consistent with Blundell *et al* (2008).¹¹

¹⁰We extend the model by introducing an extra parameter ψ_{sAge} such that $\lambda_{ht} = \lambda_h \exp(\psi_{sAge} \cdot (t - 1))$.

¹¹Kaplan 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.

5.2 Marginal distributions of model parameters

Table 2 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 preference 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.

	10%	50%	90%
Panel 1: Income model parameters			
μ Mean at start of process	-0.13	0.05	0.22
α Trend ($\times 100$)	-1.07	-0.07	0.93
ρ AR parameter	0.51	0.85	0.97
θ MA parameter	-0.15	0.21	0.53
ν Shock standard deviation	0.06	0.13	0.26
Panel 2: Consumption model parameters			
δ Discount rate ($\times 100$)	5.30	7.96	9.31
γ Coefficient of relative risk aversion	2.23	7.55	12.94
λ Response of μ to income shocks	0.32	1.14	4.51
Panel 3: The effect of an income shock			
τ Long run effect of a shock on discounted income	2.35	7.94	40.32
ϑ Effect on consumption	0.05	0.19	0.69
κ Proportion of μ log shock due to income shocks ^a	0.04	0.23	0.67

Note: a) κ is defined in appendix equation (A.3).

Table 2: Marginal distributions of model parameters

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 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 about 8%, 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 7.55, 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 (2008) 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 , is also found to be very heterogeneous with some households hardly responding (the first decile value is 0.32) and others responding a lot (the ninth decile value is 4.51) to income shocks.¹²

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 2, Panel 3 presents the related estimates. The first of these estimates is the long run effect on income of an income shock, τ , 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 . From the estimates of τ it is clear that even though we have a trend stationary

¹²The value of this parameter does not have an immediate interpretation, see equation (8).

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 7.94 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 impact on consumption.

The second estimate is the consumption response to a positive income shock, ϑ as defined in (8). 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%.¹³ In our study, this parameter exhibits a great deal of heterogeneity. At the ninth decile, the impact on consumption is 0.69%, 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 comes naturally from our definition of mue shocks. Recall that we decompose these shocks into non-income and income shocks, which in turn allows us to estimate the relative importance of each type of shocks in the overall variation of mue shocks. The proportion of mue shocks that are due to income shocks, denoted as κ , is defined in Appendix, equation (A.3). 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.

¹³These numbers are for the "college sample" in Table 6 in Blundell *et al* (2008).

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. As can be seen in Panel 3, the estimated distribution of this parameter is very dispersed with a median value of 23%. 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 (67% at the ninth decile)¹⁴.

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

Table 3 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 significant 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 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 . Our most interesting

¹⁴Although 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 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	-0.16	-0.01
ρ			1.00	-0.03	-0.74	0.41
θ				1.00	0.00	0.11
ν					1.00	-0.37

Table 3: Correlations between income parameters

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 4. 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.

The parameter λ , the sensitivity of the μ to an income shock, is not a structural parameter, but depends on preference and income parameters. In our preferred model there is no independent factor for λ 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

	δ	γ	λ	ϑ	κ
δ	1.00	0.74	0.25	-0.22	0.45
γ		1.00	0.32	-0.26	0.44
λ			1.00	0.62	0.67
ϑ				1.00	0.34

Table 4: Correlations between preference parameters

latter is a key parameter for determining life time wealth accumulation. Households with a high discount rate accumulate lower net wealth which makes them 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 4).

The parameter $\vartheta = \lambda/\gamma$ gives the degree of exposure to income risk (see the discussion after equation (8)). 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).¹⁵ 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 . 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 .

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 5. 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

¹⁵Note that $1 - \vartheta$ almost corresponds to the partial insurance coefficient in Kaplan and Violante (2010).

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 choice. 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 evidence of excess sensitivity suggests that this is not a viable alternative explanation since income changes due to a trend are anticipated.

The estimated correlation coefficient between the coefficient of relative risk aversion γ and the dispersion of income ν is quite small, albeit it has the expected (negative) sign.¹⁶ The estimated sign is consistent 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 questions) are more likely to work in occupations with low earnings risk. Similarly, Skriabikova *et al* (2014) 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 (weakly) suggest that agents self-select into occupations

¹⁶Concerned about capturing important correlations in the model, we have 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 allows us 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.

	δ	γ	λ	ϑ	κ
μ	-0.02	0.00	-0.01	-0.01	-0.03
α	-0.61	-0.04	-0.01	0.00	-0.09
ρ	-0.06	-0.04	0.54	0.52	0.50
θ	-0.01	0.01	-0.01	-0.01	-0.03
ν	0.11	-0.09	-0.45	-0.39	-0.15
τ	-0.02	-0.04	0.76	0.83	0.35

Table 5: Correlations between income and consumption parameters

to mitigate the need for precautionary saving. One potential reason for our "weak 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 strongly co-dependent at the individual level, as suggested by the cited studies, household level data may not reveal this in its full strength. 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 Kaplan and Violante (2010) for theoretical results.

Finally, we find that the proportion of consumption shock variance due to income shocks 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 (11). In Table 6, 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 0.31 (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	2.15
	Consumption	1.17	4.62	36.2

Table 6: The ratio of noisy variance to the true variance

5.6 Liquidity Constraints

While we do not find evidence of excess sensitivity in our (selected) sample, we acknowledge that liquidity constraints may be prevalent in the population. In particular, households who have a high discount rate or a high trend in income are more likely to have periods in which

they are constrained in their ability to carry debt forward.¹⁷ We are not aware of any formal analysis of the bias in our estimates that would arise if we erroneously assume that households are unconstrained. Here we present an informal analysis of the likely biases for a sample that contains constrained as well as unconstrained households.

If a household is sometimes liquidity constrained the mean of the μ shocks in equation (4) will be less than unity. Such households will appear to have a lower discount rate than they actually have as consumption growth is higher than they would have otherwise chosen. We would thus underestimate the discount rate, δ , for these households. While not affecting the (patient) left tail of the marginal distribution of the discount rate, the estimate of this distribution would be less right-skewed than the true distribution. Consequently the mean and variance of the distribution will be under-estimated. There will also be a negative bias in the correlation between the discount rate and the trend in household income since a high income trend makes a constraint more likely and hence reduces the estimated discount rate for that household. Thus some of the negative correlation we find between α and δ could be the result of incorrectly assuming away liquidity constraints if our sample contains some constrained households.

As for the coefficient of relative risk aversion, γ , we note that liquidity constrained households do not react to change in interest rates so that they would look more averse to fluctuations than they really are. Thus we would overestimate the location of gamma. As shown in Adda and Cooper (2003), the upward bias will be more pronounced for households with low γ as they will be more likely to hit their borrowing constraints, compressing the overall distribution. The likely bias in the correlation between the discount rate and the coefficient of relative risk aversion will be positive. This is because a high discount rate, which increases the probability of being constrained, will induce a positive bias in the estimation of γ . We report a positive

¹⁷There might also be other regions of the parameter space where households are constrained in some periods, but we focus on the most obvious parameters.

correlation when assuming no constraints but this could be partially due to erroneously ignoring constraints.

Finally, we consider the possible bias in the impact of an income shock on consumption, λ . We note once again that this is not a structural parameter and that it has a functional dependence on income and preference parameters. Kaplan and Violante (2010) consider an environment with homogeneous preference parameters and find that the partial insurance coefficient is lower for the “zero borrowing case” compared to the “natural borrowing constraints case” of the life cycle model. Since our ‘exposure’ parameter, λ , is negatively related to the partial insurance coefficient, we expect a positive bias in estimating λ . Once we allow for heterogeneity in δ , the bias we should consider is the bias in the correlation between δ and λ . This bias is positive. To see this, consider the very simple case of two households, one patient and the other impatient. The former is never constrained whereas the latter is sometimes constrained. The correlation between δ and λ is the slope of the line connecting the two households in δ - λ space. For the true values (if no one was constrained) this slope is positive. Once we allow for the possibility of a constraint, the estimated δ falls and the estimated λ rises. This unambiguously increases the slope with respect to the unconstrained case.

All said, the presence of constrained households in our sample would bias our results toward finding less heterogeneity in the intertemporal allocation parameters than there truly is.

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. More generally, we provide a methodology that

can quantify the extent of co-dependent heterogeneity in systems of processes with pervasive heterogeneity.

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 much 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 may 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.

The main limitation of the paper is that we do not allow for the possibility of liquidity constraints. We provide an informal analysis and conclude that if some of the households in our sample are sometimes liquidity constrained, this will bias us towards finding less heterogeneity in preference parameters than actually exists and hence cannot be the source of our finding of pervasive heterogeneity. Estimating income and consumption parameters under co-dependent heterogeneity in the presence of possibly co-dependent liquidity constraints requires much better data than we currently have. Most agents can borrow up to a limit, and this limit is likely to

be heterogeneous across individuals in a given period. We conjecture that our methods can be extended to incorporate liquidity constraints provided that, in addition to consumption and income, the household's per-period net worth is observed.

As better data become available, possibilities of future work our study generates abound. Future research that focuses on policy evaluations under pervasive heterogeneity and liquidity constraints would be especially promising. Finally, our model can be further enriched by explicitly accounting for other factors; examples include modelling fertility jointly with income and consumption and explicitly allowing for aggregate shocks.

References

- [1] Adda and Cooper (2003), "*Dynamic Economics: Quantitative Methods and Applications*", MIT Press.
- [2] Alan, S., and M. Browning (2010), "Estimating Intertemporal Allocation Parameters Using Synthetic Residual Estimation", *Review of Economic Studies* 77, 1231-1261.
- [3] Alan, S., T. Crossley and H. Low (2014), "Saving on a Rainy Day, Borrowing for a Rainy Day", manuscript.
- [4] 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).
- [5] Andersen, S. G. W. Harrison, M. I. Lau, and E. E. Rutstrom (2008), "Eliciting Risk and Time Preferences", *Econometrica*, 76(3), 583-618.
- [6] Andreoni, J. and C. Sprenger (2012), "Estimating Time Preferences from Convex Budgets.", *American Economic Review*, 102(7): 3333-56.
- [7] Arellano, M. (2003), *Panel Data Econometrics*. Oxford: Oxford University Press.
- [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] 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
- [10] 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.

- [11] Blundell R. , Luigi Pistaferri, Ian Preston (2008), "Consumption Inequality and Partial Insurance", *American Economic Review*, Vol. 98, No. 5 (Dec., 2008), pp. 1887-1921.
- [12] 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.
- [13] 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.
- [14] Browning, M. and T. Crossley (2000), "Luxuries Are Easier to Postpone: A Proof". *Journal of Political Economy*, vol. 108(5), pp.1022–1026.
- [15] Browning, M., M. Ejrnæs, J. Alvarez (2010), "Modelling Income Processes with Lots of Heterogeneity", *Review of Economic Studies* 77, 1353-1381.
- [16] Browning, M and M. Ejrnæs (2013), "Heterogeneity in the Dynamics of Labor Earnings", *The Annual Review of Economics*, vol.. 5. pp. 219-245
- [17] Browning, M., L. Hansen, and J. Heckman (1999), "Micro data and general equilibrium models", *Handbook of Macroeconomics* 1: 543-633.
- [18] 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.
- [19] Cagetti, M. (2003), "Wealth Accumulation Over the Life Cycle and Precautionary Savings", *Journal of Business and Economic Statistics*, 21(3), 339-353.
- [20] Cameron, C. and P. Trevedi (2005), *Microeconometrics*, Cambridge University Press, New York

- [21] 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.
- [22] Deaton, A. and C. Paxson (1994), "Intertemporal Choice and Inequality", *Journal of Political Economy*, vol. 102(3), pp. 437-67.
- [23] De Nardi, M., French, E. and D. Benson (2012) "Consumption and the Great Recession", *Economic Perspectives* 36.1
- [24] 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.
- [25] Eckel, C., C. Johnson and M. Claude (2005), "Saving Decisions of the Working Poor: Short-and Long-Term Horizons", *Field experiments in Economics*, 219.
- [26] Farhi, E., and I. Werning (2012) "Capital taxation: Quantitative explorations of the inverse Euler equation." *Journal of Political Economy* 120(3): 398-445.
- [27] Gneezy, U. and J. Potters. "An Experiment on Risk Taking and Evaluation Periods". *The Quarterly Journal of Economics*, vol. 112(2), pp. 631–645.
- [28] Gourieroux, C., P. C. Phillips, and J. Yu (2010), "Indirect Inference for Dynamic Panel Models", *Journal of Econometrics*, 157, 68-77.
- [29] Gourinchas, P. and J. Parker (2002), "Consumption Over the Life Cycle", *Econometrica*, 70(1), 47-89.
- [30] Guvenen, F. (2009), "An Empirical Investigation of Labor Income Processes", *Review of Economic Dynamics*, 12(1), pp. 58-79.
- [31] Guvenen F. and A. Smith (2014), "Inferring Labor Income Risk From Economic Choices: An Indirect Inference Approach", *Econometrica*, 82(6), pp. 2085-2129.

- [32] Guiso, L. and M. Paiella (2008), "Risk Aversion, Wealth and Background Risk", *The Journal of the European Economic Association*, 6(6), 1109-1150, 2008.
- [33] 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.
- [34] Heckman, J. (2001), "Micro Data, Heterogeneity, and the Evaluation of Public Policy: Nobel Lecture", *Journal of Political Economy*, vol. 109(4), pp. 673-748
- [35] 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.
- [36] Holt, C. A., and S. K. Laury (2002), "Risk Aversion and Incentive Effects", *American Economic Review*, 92, 1644-1655.
- [37] Kaplan, G. and G. Violante (2010), "How Much Consumption Insurance Beyond Self-Insurance?", *American Economic Journal: Macroeconomics*, vol. 2(4), pp. 53-87.
- [38] Kocherlakota, N. (2010), *The new dynamic public finance*, Princeton University Press.
- [39] 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
- [40] Meghir, C. and L. Pistaferri. (2004), "Income Variance Dynamics and Heterogeneity", *Econometrica*, 72(1), pp. 1-32
- [41] Mincer, J., (1958), "Investment in Human Capital and Personal Income Distribution", *Journal of Political Economy*, Vol. 66, No. 4, pp. 281-302.
- [42] 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

- [43] 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.
- [44] Runkle, D. E., (1991), "Liquidity Constraints and The Permanent Income Hypothesis", *Journal of Monetary Economics*, 27, 73-98.
- [45] Samwick, A. (1998), "Discount Rate Heterogeneity and Social Security Reform", *Journal of Development Economics*, 57 (October 1998), 117-146.
- [46] Shore, S. (2010), "For Better, for Worse: Intrahousehold Risk-Sharing over the Business Cycle", *Review of Economics and Statistics*, 92(3): 536-548.
- [47] Skriabikova, O., T. Dohmen, B. Kriechel. 2014. "New Evidence on the Relationship between Risk Attitudes and Self-Employment", *Labour Economics*, 30: 176-184.

A Appendix

A.1 Details on model specifications

A.1.1 ARMA representation

In a conventional income model, log income, y_{ht} for household h at age t , can be decomposed into three components: a deterministic component (a constant μ_h and a linear trend α_h), a persistent component, p_{ht} , and a transitory component, u_{ht} :

$$y_{ht} = \mu_h + \alpha_h \cdot (t - 1) + p_{ht} + u_{ht}$$

The persistent component is given by

$$p_{ht} = \rho_h p_{ht-1} + \zeta_{ht}$$

where ζ_{ht} is a persistent shock. This persistent-transitory model is a generalization of the widely used permanent-transitory model which imposes $\rho_h = 1$. Combining the two equations above, we have:

$$y_{ht} = \mu_h (1 - \rho_h) + \rho_h \alpha_h + \rho_h y_{ht-1} + [\alpha_h (1 - \rho_h)] (t - 1) + u_{ht} - \rho_h u_{ht-1} + \zeta_{ht}$$

If the persistent shock, ζ_{ht} , and transitory shock, u_{ht} , are assumed to be serially uncorrelated, then a general representation of this model is as an $ARMA(1, 1)$ model with a linear trend (as in equation (1) in the text):

$$y_{ht} = \{\mu_h (1 - \rho_h) + \rho_h \alpha_h\} + \rho_h y_{ht-1} + [\alpha_h (1 - \rho_h)] (t - 1) + \nu_h \xi_{ht} + \theta_h \nu_h \xi_{ht-1}$$

A.1.2 Initial conditions

To model initial conditions we impose the stationarity conditions while allowing for nonstationarity of the distribution (Arellano (2003)). Specifically we set:

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

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

A.1.3 Non-income mue shocks

To construct the non-income mue shock, we first define two unit mean log normals:

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

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

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

The parameter d allows that the two components of the mixture have different (positive) means and the second expression ensures that $\tilde{\epsilon}_{ht}$ 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.

Our framework also allows us to quantify the importance of non-income mue shocks relative to income mue shocks. Taking logs of (5) and using (6), we obtain the proportion of variance of log shock due to income as against the total variance of log consumption shocks:

$$\begin{aligned}\kappa_h &= \frac{\text{var}(\ln(g(\xi_{ht}; v_h, \lambda_h)))}{\text{var}(\ln(g(\xi_{ht}; v_h, \lambda_h))) + \text{var}(\ln(\tilde{\epsilon}_{ht}))} \\ &= \frac{\lambda_h^2 \nu_h^2}{\lambda_h^2 \nu_h^2 + \text{var}(\ln(\tilde{\epsilon}_{ht}))}\end{aligned}\tag{A.3}$$

where, we made use of the independence between $g(\xi_{ht}; v_h, \lambda_h)$ and $\tilde{\epsilon}_{ht}$. Note that this ratio is increasing in the sensitivity parameter (λ_h) and the income variance (ν_h^2).

A.1.4 Measurement error

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:

$$\begin{aligned} Y_{ht}^{obs} &= Y_{ht} \exp\left(-\frac{(m_h^y)^2}{2} + m_h^y u_{ht}^y\right) \\ C_{ht}^{obs} &= C_{ht} \exp\left(-\frac{(m_h^c)^2}{2} + m_h^c u_{ht}^c\right) \end{aligned} \quad (\text{A.4})$$

where Y_{ht} is defined as $\exp(y_{ht})$ from subsection 2.1 and C_{ht} is given by (4).

A.2 Simulation

Indirect inference requires simulating from the parametric model. 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 (A.1) and (1) for $t = 1, \dots, T$. The second set is for the consumption non-income shocks in (A.2), η_{ht} for $t = 2, \dots, T$. The final set is the factors, N_{kh} , for $k = 1, \dots, 10$; see (9)-(11). 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 (9) and (10), 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 (A.1); this gives $R * H$ values for y_{h1} . Then subsequent income paths are given recursively by (1) and the ξ_{ht} 's for $t = 2, \dots, T$.¹⁸

To simulate consumption growth paths, we first simulate consumption shocks from (A.2); this uses the given values for $\{\sigma_a, \sigma_b, d\}$, the simulated values for λ_{ht} from (10) and the current income shocks, $\nu_h \xi_{ht}$. We set the initial value of consumption to unity¹⁹ and construct levels sequentially, using C_{h1} , and values for r_{t+1} (where t refers to age) and the simulated values for

¹⁸In 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 .

¹⁹This choice of starting value distribution is irrelevant since the initial value plays no part in the simulated consumption growth path.

(γ_h, δ_h) and $(\varepsilon_{h2}, \varepsilon_{h3}, \dots)$:

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

Finally, measurement errors are added to the simulated incomes and consumptions, using (11).

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.

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.

A.3 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 33 parameters and seven factors ($N_1, N_2, N_3, N_5, N_6, N_9, N_{10}$). These are given by:

$$\begin{aligned}
\mu_h &= \phi_1 + \exp(\psi_{11}) N_{1h} \\
\alpha_h &= \phi_2 + \exp(\psi_{22}) N_{2h} \\
\rho_h &= \ell(\phi_3 + \exp(\psi_{33}) N_{3h}) \\
\theta_h &= 2l(\ell(\phi_4 + \psi_{41} N_{1h})) - 1 \\
\nu_h &= \exp(\phi_5 + \psi_{52} N_{2h} + \psi_{53} N_{3h} + \exp(\psi_{55}) N_{5h}) \\
\delta_h &= 0.1 * \ell(\phi_6 + \psi_{62} N_{2h} + \exp(\psi_{66}) N_{6h}) \\
\gamma_h &= 0.5 + 14.5 * \ell(\phi_7 + \psi_{75} N_5 + \psi_{76} N_{6h}) \\
\lambda_h &= \exp(\phi_8 + \psi_{83} N_{3h} + \psi_{86} N_{6h}) \\
m_h^y &= \exp(\phi_9 + \psi_{95} N_{5h} + \exp(\psi_{99}) * N_{9h}) \\
m_h^c &= \exp(\phi_{10} + \psi_{10,6} N_{6h} + \psi_{10,9} N_{9h} + \exp(\psi_{10,10}) * N_{10h})
\end{aligned}$$

where $\ell(x)$ is the transformation $e^x / (1 + e^x) \in (0, 1)$. The model contains 10 mean parameters (ϕ_j), 18 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	0.0494	0.0512	0.9659
ϕ_2	-0.0007	0.0023	0.3049
ϕ_3	1.7690	0.2262	7.8214
ϕ_4	0.4414	0.0532	8.3005
ϕ_5	-2.0628	0.0454	45.4350
ϕ_6	1.3805	1.0735	1.2861
ϕ_7	0.0326	0.3649	0.0893
ϕ_8	0.1692	0.1362	1.2422
$\ln(\sigma_a)$	-1.5210	0.2721	5.5893
$\ln(\sigma_b)$	-2.4409	0.2665	9.1598
ψ_{11}	-1.9928	0.3164	6.2987
ψ_{22}	-4.8648	0.2796	17.3977
ψ_{33}	0.2547	0.1610	1.5816
ψ_{41}	0.5746	0.0805	7.1377
ψ_{52}	-0.1018	0.0404	2.5195
ψ_{53}	-0.4626	0.0394	11.7570
ψ_{55}	-0.9820	0.1395	7.0388
ψ_{62}	-0.5976	0.4435	1.3475
ψ_{66}	-0.2177	0.5739	0.3794
ψ_{75}	-0.2933	0.1356	2.1627
ψ_{76}	1.4678	0.1459	10.0603
ψ_{83}	0.9105	0.1113	8.1780
ψ_{86}	0.4953	0.1362	3.6353
ϕ_9	-2.2379	0.0840	26.6349
ϕ_{10}	-1.6657	0.0238	70.0247
ψ_{99}	-1.0263	0.1413	7.2648
$\psi_{10,10}$	-2.1127	0.2061	10.2518
$\psi_{10,9}$	0.3137	0.0404	7.7688
ψ_{95}	0.2878	0.0490	5.8720
$\psi_{10,6}$	0.4166	0.0281	14.8471
b_0	-0.4107	0.0921	4.4592
b_1	-0.6366	0.4052	1.5711
$d(mix)$	0.1837	0.0646	2.8450

Table A.1: Distribution parameters

AP	data	sim	se	t-val
$M(\hat{b}_{c5})$	-0.315	-0.164	0.082	-1.849
$S(\hat{b}_{c5})$	2.545	2.521	0.155	0.158
$C(\hat{b}_{c1}, \hat{b}_{c5})$	0.511	-0.119	0.441	1.429
$C(\hat{b}_{c2}, \hat{b}_{c5})$	-0.179	-0.678	0.434	1.150
$C(\hat{b}_{c3}, \hat{b}_{c5})$	-1.878	-2.497	0.431	1.438
$C(\hat{b}_{c4}, \hat{b}_{c5})$	0.327	0.556	0.427	-0.536
$C(\hat{b}_{c6}, \hat{b}_{c5})$	-0.609	-0.494	0.464	-0.248
$C(\hat{b}_{c7}, \hat{b}_{c5})$	-0.117	-0.413	0.430	0.688
$C(\hat{b}_{y1}, \hat{b}_{c5})$	0.056	-0.109	0.494	0.334
$C(\hat{b}_{y2}, \hat{b}_{c5})$	-0.042	-0.013	0.507	-0.059
$C(\hat{b}_{y3}, \hat{b}_{c5})$	0.180	-0.221	0.477	0.841
$C(\hat{b}_{y4}, \hat{b}_{c5})$	-0.112	-0.321	0.418	0.499
$C(\hat{b}_{y5}, \hat{b}_{c5})$	0.472	0.418	0.428	0.126
$corr(\hat{w}_{ht}, \hat{u}_{ht} \cdot a_{ht})$	0.193	-0.347	0.351	1.542
GF test $\chi^2_{(14)}$	15.21			
$M(\cdot)$: mean, $S(\cdot)$: std, $C(\cdot, \cdot)$: correlation				

Table A.2: AP used for goodness of fit test