How Much Insurance in Bewley Models?*

Greg Kaplan
New York University
gregkaplan@nyu.edu

Giovanni L. Violante
New York University, CEPR, and NBER
glv2@nyu.edu

First Draft: 12 February 2008 - This Draft: 21 May 2008

STILL PRELIMINARY

Abstract
The standard life-cycle incomplete-markets model, where households face idiosyncratic earnings shocks and trade a non-contingent asset, is a workhorse of quantitative macroeconomics. In this paper, we assess the degree of consumption insurance implicit in a plausibly calibrated version of the model, and we compare it to the data. On both actual and simulated data, we apply a technique recently developed by Blundell, Pistaferri and Preston (2006, BPP thereafter). We find that households in Bewley models have access to less consumption smoothing against permanent shocks than what is measured in the data . BPP estimate that 36% of permanent earnings shocks are insurable (i.e., do not translate into consumption growth), while in the model this insurance coefficient is only 17%. The life-cycle profile of insurance coefficients is sharply increasing and convex in the model, while BPP document that it is roughly flat in the data. Taken at face value, these results would suggest that macroeconomists should develop models with more avenues of insurance than a risk-free asset. Allowing for "advance information" about earnings changes in the model does not affect this conclusion. However, we also find that if earnings shocks are not modelled to be permanent, as assumed by BPP, but they display an autocorrelation coefficient between 0.90 and 0.95, then the degree of consumption smoothing in the model largely agrees with the empirical estimate. Finally, we uncover that the BPP estimator of the insurance coefficient has, in general, a small asymptotic downward bias. This downward bias becomes sizeable when borrowing limits bind frequently. However, in the presence of advance information about future earnings growth, the BPP insurance coefficient can be upward biased.

*We thank Richard Blundell and Luigi Pistaferri for helpful comments.

1 Introduction

In the last decade, incomplete-markets models have become a common tool for quantitative analysis in macroeconomics. The standard incomplete-markets model economy can be described as follows. A continuum of households with concave preferences over consumption streams populates the economy. Households face idiosyncratic exogenous labor income fluctuations. The only asset they can trade in order to smooth consumption is a non-state-contingent one-period security. A borrowing constraint limits the amount of debt they can accumulate. Borrowing and saving through the risk-free asset is often called "self-insurance". In equilibrium, the price of the asset is determined by supply and demand forces in a competitive market. Essentially, these models aggregate through competitive equilibrium the behavior of a continuum of independent consumers facing an "income fluctuation problem".[1]

This class of models was introduced by Bewley (1980) and Bewley (1983) to study properties of equilibria with fiat money. Once the numerical tools to compute recursive competitive equilibria became available, Imrohoroglu (1989), Huggett (1993), Aiyagari (1994) and Rios-Rull (1995) advocated and pioneered quantitative work in this framework. Krusell & Smith (1998) developed computational tools to approximate equilibria in the presence of aggregate shocks. The term "Bewley models" was first used in the Ljungqvist & Sargent (2004) textbook and, since then, it is becoming a common label for these models.

Bewley models have been used extensively in macroeconomics for a variety of questions: measuring the size of precautionary saving (Aiyagari (1994)), understanding the wealth distribution (Huggett (1996); Quadrini (2000); Castaneda, Diaz-Gimenez & Rios-Rull (2003); De Nardi (2004)), explaining the equity premium (Telmer (1993); Heaton & Lucas (1996); Storesletten, Telmer & Yaron (2007)), evaluating the distributional effects of fiscal policy (Krusell, Quadrini & Rios-Rull (1997); Ventura (1999); Domeij & Heathcote (2004)), measuring the welfare costs of business cycles (Krusell & Smith (1999)), and designing optimal fiscal policy (Erosa & Gervais (2002); Nishiyama & Smetters (2005); Conesa, Kitao & Krueger (2007)), only to cite a few.

The objective of this paper is to measure the degree of consumption insurance allowed to

[1] See Yaari (1976), Schechtman & Escudero (1977) and Chamberlain & Wilson (2000) for general theoretical results on the convergence of the optimal consumption sequence of an infinitely-lived household who can trade a safe asset at an exogenous interest rate, subject to a given lower bound on asset holding.

households in Bewley models and to compare it to existing empirical estimates for the US economy. The questions we address are: (i) How much consumption smoothing is there in Bewley models? And, (ii) Is this number high or low relative to that measured in the data?

Incomplete markets models with heterogeneous agents have rich predictions through the equilibrium distribution of allocations (e.g., consumption, wealth, and hours worked) in the cross-section, and over the life cycle. The focus of the existing literature has largely been on documenting these distributions, and on comparing them to the data. We argue that more emphasis should be put on measuring the degree of consumption smoothing implied by these models, for two reasons. First, imperfect insurance is at the heart of incomplete markets models. Thus it is useful to develop summary statistics for consumption insurance that allowe one to compare, parsimoniously, incomplete markets models along their most salient dimension. Second, often economists use these models for policy evaluation and design. The results of such experiment hinge crucially upon the range of insurance vehicles that households have access to, and on the amount of consumption smoothing they can afford. For example, a reform from a progressive to a flat tax system is judged on the basis of the gains from reduced distortions and the losses from lower redistribution. But the size of the latter margin depends on how much smoothing agents can do on their own, through private risk-sharing.

Our metric for consumption insurance is the fraction of permanent (random walk) and transitory (i.i.d.) idiosyncratic labor income (or, equivalently, earnings) shocks that does *not* translate into movements in consumption. We call this measure the *insurance coefficient* of income shocks. In a complete markets model, it is one for both shocks. In autarky, without storage, it is zero for both shocks. A third useful benchmark is the strict version of the permanent income hypothesis (PIH), where this insurance coefficient is close to one for transitory shocks and is exactly zero for permanent shocks (Hall (1978)).[2]

Our exercise is possible because we now have reliable empirical measures of partial consumption insurance. In a pair of recent papers, Blundell, Pistaferri and Preston (BPP, thereafter) have developed the necessary data set and methodology, and they have obtained empirical estimates of these transmission coefficients.

In BPP (2004), they have constructed a new panel data set for the US with household

[2]By strict version of the PIH, we mean the optimization problem faced by an infinitely-lived consumer with quadratic preferences and without borrowing limits, who faces a gross interest rate equal to the inverse of the discount factor.

information on income and nondurable consumption. The new data set is obtained by combining existing PSID data on income and food expenditures with repeated cross-sections from CEX on nondurable consumption. The key step is the imputation of a measure for nondurable consumption for each individual/year observation in PSID by exploiting the fact that food consumption is available in both data sets. From CEX, one can estimate a relationship between food and nondurable consumption expenditures–a food demand function–and then invert the demand function and implement the imputation procedure at the household level, based on the reported value for food consumption in PSID records.

In BPP (2008), the authors use this new data set to estimate insurance coefficients for permanent and transitory labor income shocks separately. We return to the details of the econometric procedure later. Here it suffices to say that their key contribution is in devising a way to identify separately the impact of permanent and transitory shock on consumption growth. They find that 36% of permanent shocks, and 95% of transitory shocks to disposable (i.e., post-tax) labor income are insurable, i.e. they are not passed onto consumption growth.

We calculate these transmission coefficients from an artificial panel simulated from a Bewley model. We choose a life-cycle version with capital, where the government imposes a progressive redistributive taxation system and a pension scheme that mimic those in the US economy. In the model, households can smooth shocks by borrowing, as long as their accumulated debt is below the "natural limit". They also save for life-cycle, and precautionary reasons, and thus their wealth helps in absorbing income shocks.[3] The calibration of the model is standard.

Our findings can be summarized as follows. First, in the model the insurance coefficient for transitory shocks to after-tax labor income is 92% , very close to the empirical estimate, whereas the insurance coefficient for permanent shocks is 17%, i.e. about half of the empirical estimate. The model-implied coefficients are found to be very robust in a series of sensitivity analyses: insurance coefficients in the model are above 30% only for values of risk-aversion above ten. Moreover, the life-cycle pattern of the of insurance coefficients for permanent shocks is sharply increasing and convex. This finding is strikingly at odds with the BPP estimates which display a roughly flat, if not slightly decreasing, age profile. The lifecycle discrepency between model and data is at young ages - the model generates roughly the

[3]Note that, because of the life-cycle, income shocks received towards the end of the working life have short-lived effects even when they are permanent.

correct level of insurance for older workers but too little insurance for workers in the early stages of their career.

Second, we are able to assess the ability of the insurance-coefficients estimator proposed by BPP to identify insurance for each type of shock. We find that it works very well for transitory shocks, but it tends to underestimate true coefficients for permanent shocks. For example, in the benchmark model the BPP estimate of the latter coefficient is 14% against a true value of 17%. The reason is that the estimation procedure, which is analogous to an instrumental variables approach, exploits an orthogonality condition between changes in consumption and a particular linear combination of past and future changes in income. The bias results from the fact that this orthogonality condition holds only approximately in the model. Quantitatively, on average the bias is not large (factor of 1.2), although systematically present. However, when households are near their borrowing constraint, this bias becomes very severe. In an economy with no borrowing, for example, the BPP estimator is less than half of the true coefficient. Overall, we conclude that, once we correct for this bias, empirical insurance coefficients will be even higher than estimated, so model and data are even more distant.

In light of these first two findings, one could conclude that macroeconomists should develop models with more avenues of insurance than just borrowing-saving through a risk-free asset. We explore two alternative ways in which standard Bewley models could generate less sensitivity of consumption to permanent shocks. First, we allow agents to have advance information about future shocks. Second, we change the specification of the statistical process for earnings.

We model advance information in two ways. We first let agents know a fraction of the permanent shock one period ahead of time. Next, we assume that earnings have an individual-specific deterministic trend which is known by the agent from age zero. In the first case, we show that the BPP estimator of insurance coefficients is, in essence, invariant to the amount of advanced information. In the second case, we argue that for plausibly calibrated heterogeneity in income profiles, the estimated insurance coefficients remain lower than in the data. Advanced information does not appear to bridge the gap between model and data.

When we choose a different specification of the income process, we can reconcile the model with the data. We posit that the persistent component of the income process is

$AR(1)$ instead of $I(1)$ as assumed by BPP. We first show that the BPP estimation method of insurance coefficients performs quite well even under this misspecification error, for degrees of persistence (ρ) between 0.90 and 0.95. Next, we document that for $\rho = 0.93$ the misspecified BPP estimate of the insurance coefficient for persistent shocks in the Bewley model is around 35%, i.e. close to its empirical value. However, the insurance coefficient for transitory shocks declines to 74% compared to 95% in the data. So, in a sense, the puzzle becomes that the model generates too little insurance with respect to transitory shocks.

The rest of the paper is organized as follows. Section 2 introduces a general framework for measuring insurance, and describes the BPP methodology as a special case. Section 3.1 outlines the version of the Bewley model we use for our experiments, and describes its parameterization. Section 4 contains the results from the benchmark economy and from a series of sensitivity analyses, including the one where we tighten borrowing constraints. Section 5 introduces advance information into the model. Section 6 analyzes the robustness of our conclusion to the degree of persistence of income shocks. In Section 7 we discuss our results in relation to the existing literature. Section 8 concludes the paper.

2 A framework for measuring insurance

2.1 Insurance coefficients

An income process Suppose that residual (i.e., deviations from a deterministic experience profile common across all households) log-earnings y_{it} for household i of age t can be represented as a linear combination of current and lagged shocks

$$y_{it} = \sum_{j=0}^{t} a_j' \mathbf{x}_{i,t-j} \tag{1}$$

where $\mathbf{x}_{i,t-j}$ is an $m \times 1$ vector of i.i.d. shocks and a_j is an $m \times 1$ vector of coefficients. Let $\sigma = (\sigma_1, ..., \sigma_m)'$ be the corresponding vector of variances for these shocks. This formulation is extremely general and incorporates, for example, linear combinations of ARIMA processes with fixed effects.

Insurance coefficients Let c_{it} be log consumption for household i at age t. We define the *insurance coefficient* for shock $x_{it} \in \mathbf{x}_{it}$ as

$$\phi^x = 1 - \frac{cov\left(\Delta c_{it}, x_{it}\right)}{var\left(x_{it}\right)}, \tag{2}$$

where the variance and the covariance are taken cross-sectionally over the entire population of households. One can similarly define the insurance coefficient at age t (denoted by ϕ_t^x) where the variance and covariance are taken conditionally on all households of age t. The insurance coefficient in (2) has an intuitive interpretation: it is the share of the variance of the x shock which does *not* translate into consumption growth.

Identification and estimation In general, how do we compute these insurance coefficients? In any given model, it is straightforward to calculate (2) by simulation, since the shocks are observable in the model. However, estimating (2) from the data poses a crucial difficulty: the individual shocks are not directly observed and cannot be identified from a finite panel of income data.[4]

Suppose one has panel data on households' income and consumption. Let $\mathbf{y}_i$ be the vector of income realizations for individual i at all ages $t = 2, ..., T$, and let $g_t^x(\mathbf{y}_i)$ index measurable functions of this income history, one for each t and for each shock x. Identification and estimation of insurance coefficients for shock x can be achieved by finding functions g_t^x such that

$$\begin{aligned} var\left(x_{it}\right) &= cov\left(\Delta y_{it}, g_t^x\left(\mathbf{y}_i\right)\right), \\ cov\left(\Delta c_{it}, x_{it}\right) &= cov\left(\Delta c_{it}, g_t^x\left(\mathbf{y}_i\right)\right), \end{aligned} \tag{3}$$

and then constructing ϕ^x as

$$\phi^x = 1 - \frac{cov\left(\Delta c_{it}, g_t^x\left(\mathbf{y}_i\right)\right)}{cov\left(\Delta y_{it}, g_t^x\left(\mathbf{y}_i\right)\right)}. \tag{4}$$

While verifying the first condition in (3) is straightforward, given an income process, verifying the second condition requires knowledge of how the consumption allocation in the data depends on the entire vector of realizations (past and future) of the shocks. Thus, it requires knowledge of the true data generating process (i.e., the model) for consumption.

This approach to identification is best thought of in terms of instrumental variables regressions. If $g_t(\mathbf{y}_i)$ satisfies the conditions in (3), then the resulting expression for $1 - \phi^x$ is equivalent to the coefficient from an instrumental variables cross-sectional regression of consumption changes on income changes, using $g_t(\mathbf{y}_i)$ as an instrument. In general, the

[4]Note that it is not sufficient to identify the variances of the different shocks, i.e the vector σ. Rather, the realizations of the shocks must be identified, household by household. With a very long sequence of observations, realizations may be identified using filtering techniques. However the short time-dimension of commonly available panel data sets makes filtering techniques unreliable in this context.

correct choice of instrument depends on the particular specification of the income process, and the underlying true model that determines consumption. To progress further, one has to make assumptions about both.

2.2 The BPP methodology

BPP income process BPP focus on a particular specification of the income process in (1), where $m = 2$, $\mathbf{x}_{it} = (\eta_{it}, \varepsilon_{it})'$, $a_0 = (1, 1)'$ and $a_j = (1, 0)'$ for $j \geq 1$. This corresponds to the sum of a random walk (permanent) and an i.i.d. component:

$$y_{it} = z_{it} + \varepsilon_{it}, \tag{5}$$

where z_{it} follows a unit root process with innovation η_{it}, and where ε_{it} is an uncorrelated income shock.[5] The individual shocks η_{it} and ε_{it} are i.i.d.. across the population. Let σ_η and σ_ε denote the variances of the two innovations.[6] It follows that income growth can be written as

$$\Delta y_{it} = \eta_{it} + \Delta\varepsilon_{it}. \tag{6}$$

This is a very common income process, at least since the work by Abowd & Card (1989), who showed that this specification is parsimonious and yet fits income data well. In Section 6, we verify the robustness of our results to more general specifications of the income process.

BPP consumption model BPP assume that the following pair of orthogonality conditions hold for the true consumption allocation:

$$cov\left(\Delta c_{it}, \eta_{i,t-1}\right) = cov\left(\Delta c_{it}, \varepsilon_{i,t-2}\right) = 0 \tag{7}$$

$$cov\left(\Delta c_{it}, \eta_{i,t+1}\right) = cov\left(\Delta c_{it}, \varepsilon_{i,t+1}\right) = 0. \tag{8}$$

The first assumption translates into short history dependence of the consumption allocation from shocks. The second assumption means that the agent has no advanced information

[5]BPP also allow for an MA(1) transitory shock, whereas for computational simplicity we restrict attention to an IID transitory shock. Since the estimated serial correlation on their MA(1) shock is only 0.11, none of their conclusions are affected by the exclusion of this component.

[6]One can easily generalize this process by allowing a different distribution for the initial draw of the permanent shock. For example, let $\eta_{i0} = z_0$ with variance σ_{z_0}. This generalization has no bearing on the methodology used to measure insurance coefficients, since the latter is based on growth rates, and hence the initial fixed effect is differenced out.

about future shocks.[7]

Under these assumptions, BPP propose a strategy to identify and estimate the insurance coefficients. For the transitory shock ε, they set $g_t^{\varepsilon}(\mathbf{y}_i) = \Delta y_{i,t+1}$ and note that

$$\begin{aligned} cov\left(\Delta y_{it}, \Delta y_{i,t+1}\right) &= -var\left(\varepsilon_{it}\right), \\ cov\left(\Delta c_{it}, \Delta y_{i,t+1}\right) &= -cov\left(\Delta c_{it}, \varepsilon_{it}\right), \end{aligned} \tag{9}$$

while for the permanent shocks η, they set $g_t^{\eta}(\mathbf{y}_i) = \Delta y_{i,t-1} + \Delta y_{it} + \Delta y_{i,t+1}$ and note that

$$\begin{aligned} cov\left(\Delta y_{it}, \Delta y_{i,t-1} + \Delta y_{it} + \Delta y_{i,t+1}\right) &= var\left(\eta_{it}\right), \\ cov\left(\Delta c_{it}, \Delta y_{i,t-1} + \Delta y_{it} + \Delta y_{i,t+1}\right) &= cov\left(\Delta c_{it}, \eta_{it}\right). \end{aligned} \tag{10}$$

Combining (2) with (9) and (10) confirms that these instruments do in fact correctly identify the insurance coefficients $(\phi^{\eta}, \phi^{\varepsilon})$. Note that only the orthogonality conditions in (8), i.e. "no advanced information", are required for the identification of the insurance coefficients for transitory shocks, while both (7) and (8) are needed for permanent shocks.

In what follows, we call ϕ^x_{BPP} the insurance coefficients estimated through the BPP methodology. When the orthogonality conditions hold, $\phi^x_{BPP} = \phi^x$, when they do not there will be a bias in ϕ^x_{BPP}.

Generality of the BPP approach The obvious question, at this point, is: how general are assumptions (7) and (8)? In the absence of advance information (8) holds. In addition, it is easy to see that complete markets and autarkic economies satisfy (7). Under complete markets, idiosyncratic shocks do not affect consumption, hence $cov\left(\Delta c_{it}, x_{it}\right) = 0$ and $\phi^x = 1$. In autarky, $\Delta c_{it} = \Delta y_{it}$, hence $cov\left(\Delta c_{it}, x_{it}\right) = var\left(x_{it}\right)$ and $\phi^x = 0$. Note that in these two extreme cases, the value of ϕ^x is independent of the durability of the shock. The life-cycle PIH, where agents have quadratic utility, live for T periods and can borrow and save at a constant risk-free rate generates the following rule for changes in consumption,

[7] To be precise, BPP assume that the log-linearized equilibrium consumption allocation of the true model implies the following approximate form for consumption growth

$$\Delta c_{it} \simeq \alpha_i^0 + \alpha_{it}^{\eta}\eta_{it} + \alpha_{it}^{\varepsilon}\varepsilon_{it} + \Psi_{it}$$

where where α_{it}^{η} and $\alpha_{it}^{\varepsilon}$ are marginal propensity to consume out of permanent and transitory shocks. The term Ψ_{it} is a residual component, possibly involving changes in the endogenous state variable (e.g., assets or promised continuation utility). BPP posit that $(\alpha_{it}^{\eta}, \alpha_{it}^{\varepsilon}, \Psi_{it})$ are all independent of income innovations at every lead and lag. This assumption is too strong. As explained, (7) and (8) which jointly are weaker assumptions, are enough.

when combined with the analogous income process to (5) specified in levels:

$$\Delta C_{it} = \eta_{it} + \chi_t \varepsilon_t,$$

where $\chi_t = \frac{r}{(1+r)} \frac{1}{1-(1+r)^{-(T-t+1)}}$.[8] Hence the PIH satisfies the BPP assumptions, and the insurance coefficients (defined in terms of *levels*) for a PIH economy are $\phi_t^\eta = 0$, and $\phi_t^\varepsilon = 1 - \chi_t$. These values imply no insurance against permanent shocks, and an insurance coefficient for transitory shocks that decreases monotonically towards zero as the end of life becomes nearer. Finally, one can verify that the BPP assumptions hold in the partial insurance economy developed by Heathcote, Storesletten and Violante (2007), and in the moral-hazard economy studied by Attanasio and Pavoni (2007). These examples demonstrate that, in a wide variety of economic environments, it is possible to justify consumption allocations that are consistent with (7) .

Estimation method Straightforward application of a minimum distance algorithm allows the estimation of the cross-sectional moments in (9) and (10). However, the model can only be estimated from panel data with at least four consecutive observations on both individual income and consumption. None of the currently available US surveys has this feature. BPP (2004) cleverly merged the CEX and PSID and constructed a long panel with nondurable consumption and income observations. BPP imputed a measure for nondurable consumption for each household/year observation in the PSID by exploiting the fact that food consumption is accurately recorded in both the PSID and CEX. From the CEX, they estimated a statistical relationship between food and nondurable consumption expenditures (a food demand function) and then inverted the demand function and implemented the imputation procedure based on the realized food consumption measure, and various demographics, in PSID.

Results BPP reach three main findings. First, when income is defined as after-tax household labor income, the insurance coefficient for permanent shocks ϕ^η is estimated to be 36%.[9] Second, the insurance coefficient for transitory shocks ϕ^ε is estimated to be 95%. Third, the age profile of ϕ^η is flat.

[8] We use upper-case letters to denote variables in levels, and lower-case letters to denote variables in logs.

[9] The definition of income is important. For example, we would expect insurance coefficients for pre-government income to be higher than for post-government income fluctuations, precisely because the government engages in some redistribution.

3 A Bewley model

In this section, we outline and calibrate a life-cycle Bewley economy. Next, we simulate an artificial panel of household income and consumption, and calculate the insurance coefficients in the model. By comparing them to the empirical values estimated by BPP we will learn whether the Bewley model has the "right" amount of insurance, and which features of the model govern the amount of insurance that it permits. Moreover, since in the model we can compute both the exact insurance coefficients and the BPP estimator, we will also learn about the reliability of the BPP methodology, i.e. we will find out if and when assumption (7) is violated.

Before outlining the model, a brief note on language. We have defined ϕ^x as *insurance* coefficients. In Bewley models, income shocks do not transmit one for one into consumption because consumers have access to non state-contingent borrowing and saving. And thus, in this context, one may prefer the term "self-insurance" or "smoothing" to insurance. For example, Cochrane (1991) advocates the strict use of the term insurance only to denote static inter-personal transfers. We have chosen to keep using the term insurance in the rest of the paper for four distinct reasons. First, the definition of ϕ^x is model-independent, and thus it can be used also in the context of models allowing for state-contingent income transfers. Second, unlike dissaving, borrowing involves an inter-personal transfer. Third, even within Bewley models, there may be important insurance mechanism over and above self-insurance, for example the government (that we will model), the family, and certain institutions such as consumer's bankruptcy (Livshits, McGee and Tertilt, 2003; Chatterjee et al., 2007). Fourth, BPP refer to their estimates as "partial insurance coefficients".

3.1 The Economy

The model economy is the life-cycle version of the standard incomplete markets model introduced by Huggett (1993). There is no aggregate uncertainty. The economy is populated with a continuum of households, indexed by i. Agents work until age T^{work}, after which they enter a period of retirement. The unconditional probability of surviving to age t is denoted by ξ_t. We assume that $\xi_t = 1$ for the first T^{work} periods, so that there is no chance of dying before retirement. After retirement, $\xi_t < 1$ and all agents die by age T with certainty.

Households have time-separable expected utility and constant relative risk aversion intraperiod preferences given by:

$$E_0 \sum_{t=1}^{T} \beta^{t-1} \xi_t u\left(C_{it}\right)$$

During the working years, household receive stochastic after-tax income, Y_{it}, which is comprised of three components in logs:

$$\begin{aligned} \log Y_{it} &= \kappa_t + y_{it} \\ y_{it} &= z_{it} + \varepsilon_{it}, \end{aligned}$$

where κ_t is a deterministic experience profile that is common across all households and y_{it} is the stochastic portion of income; z_{it} is a permanent component and ε_{it} is a transitory component. The component z_{it} follows a random walk

$$z_{it} = z_{i,t-1} + \eta_{it},$$

where z_{i0} is drawn from an initial normal distribution with mean zero and variance σ_{z_0}. The shocks ε_{it} and η_{it} are zero-mean normally distributed with variances σ_ε and σ_η, and are independent over time and across households.

Retired households receive social security income $P\left(z_{T^{work}}\right)$ which is a function of the final permanent component of the households' income, i.e. $Y_{it} = P\left(z_{T^{work}}\right)$ for $t > T^{work}$.

Markets are incomplete: the only asset available to households is a single risk-free bond which pays a constant gross rate of return, $1+r$. We denote the amount of this asset carried over from time t to $t+1$ as $A_{i,t+1}$. After retirement, we assume that there exist perfect annuity markets where households use to insure against survival risk. Households begin their life with initial wealth A_{i0} and face a lower bound on assets $\underline{A} \leq 0$.

In the benchmark model, we treat Y_{it} as net household income after all transfers and taxes. Thus the budget constraint of households in this economy is simply

$$\begin{aligned} &C_{it} + A_{i,t+1} = \left(1+r\right) A_{it} + Y_{it}, && \text{if } t < T^{work} \\ &C_{it} + \tfrac{\xi_t}{\xi_{t+1}} A_{i,t+1} = \left(1+r\right) A_{it} + P\left(z_{T^{work}}\right), && \text{if } t \geq T^{work} \end{aligned}$$

We also consider a version of the economy where Y_{it} is treated as household labor earnings and we explicitly model the government sector. Government levies taxes on labor income only, using the non-linear tax rule, $\tau\left(Y_{it}\right)$. The tax function is calibrated to reflect the amount of redistribution in the US tax system. Following Hubbard, Skinner & Zeldes (1995), we introduce a consumption floor $\underline{C}$ to represent, in reduced form, means-tested programs such as Food Stamps and Temporary Assistance for Needy Families. In this version of the economy, the budget constraint reads:

$$C_{it} + A_{i,t+1} = (1+r) A_{it} + Y_{it} - \tau (Y_{it}) + B_{it}, \quad \text{if } t < T^{work}$$

$$C_{it} + \frac{\xi_t}{\xi_{t+1}} A_{i,t+1} = (1+r) A_{it} + P(z_{T^{work}}) + B_{it}, \quad \text{if } t \geq T^{work}$$

where

$$B_{it} = \max \left\{ \underline{C} - (1+r) A_{it} - Y_{it} + \tau (Y_{it}) + \underline{A}, 0 \right\}$$

is the government transfer guaranteeing a minimum level of consumption for every household.

Finally, it is useful to note that, in the benchmark case of the permanent-transitory process for disposable income, households behave close to the buffer-stock, no-debt consumers characterized by Carroll (1997)—the only difference being the retirement period and the looser borrowing limit.

3.2 Calibration

We calibrate the model parameters to reproduce certain key features of the US economy. Our parameterization is standard for this class of economies.

Demographics Households enter the labor market at age 25. We set $T^{work} = 35$ and $T^{ret} = 40$. Thus workers retire at age 60 and die with certainty at age 100. The survival rates ζ_t are obtained from the National Vital Statistics Reports (2004). We calibrate the deterministic age profile for income on PSID data. The estimated profile peaks after 21 years at roughly twice the initial value and then it slowly declines to about 80% of the peak value.

Preferences We choose a CRRA specification for $u(C_{it})$ with risk aversion parameter $\gamma = 2$. We explore the sensitivity of our results to values of γ in the range $[1, 15]$.

Discount factor and interest rate With a finite horizon, the size of the stock of accumulated assets directly affects the extent to which income shocks are insurable. Hence it is important to ensure that the wealth to income ratio in the model is similar to that in the US economy. Provided this is done, the interest rate has virtually no effect on the amount of insurance in the economy. We choose $\beta = 0.944$ to match an aggregate wealth-income ratio of 3.5. This ratio is a rough upper bound for the corresponding ratio in the US economy, for the bottom 99% of the wealth distribution over the period 1980-92, when wealth is defined as total net worth. Note that, with a Cobb-Douglas aggregate production function, a capital share of 0.3 and a depreciation rate of 5.5%, the equilibrium interest rate consistent with this aggregate wealth-income ratio would be roughly equal to 3%.

We also report results for wealth-income ratios of 1.5 and 2.5. An estimate of 1.5 corresponds roughly to the case where wealth is defined only as financial wealth (i.e., excluding housing wealth) of the bottom 99% among US households, and may be considered a lower bound on the aggregate wealth-income ratio over the 1980-1992 period.[10]

Income process For the benchmark income process, three parameters are required. These are the variance of the two shocks, σ_ε and σ_η, and the cross-sectional variance of the initial value of the permanent component σ_{z_0}. In our benchmark calibration we set the variance of permanent shocks to be 0.02, and the variance of transitory shocks to be 0.05. These correspond to the average values estimated by BPP (2006) for the period 1980-1992. The initial variance of the permanent shocks is set at 0.15 to match the dispersion of household earnings at age 25. We also report results from various sensitivity analyses on these values.

Initial wealth In the benchmark calibration we assume that all households start life with zero wealth, i.e. $A_{i0} = 0$. We also consider an environment in which initial wealth levels are drawn from a distribution calibrated to replicate the empirical distribution of wealth for young households in the data.[11]

Borrowing limit We consider two assumptions for the borrowing constraint. In our benchmark economy, we allow for borrowing subject only to the restriction that with probability one, households who live up to age T do not die in debt (i.e., the "natural debt limit"). This assumption represents an upper bound on the amount agents can borrow. We also study the insurance possibilities when the other extreme of no borrowing, $\underline{A}= 0$, is

[10] We focus on 1980-1992 surveys of the Survey of Consumer Finances (SCF) as this is the period to which the BPP estimates refer. Details of the calculations are available upon request. It is not clear *a priori* whether one should use net worth or financial wealth to calibrate the model. In the model assets can be liquidated and transformed into consumption within a year. Financial wealth has this feature. The liquidity of housing wealth depends on the local housing market. In practice, consumers can often borrow against a sizeable fraction (often 100%) of the value of their house, and thus focusing only on financial assets may seriously understate the resources available to the households for consumption smoothing.

[11] Precisely, we target the empirical distribution of financial wealth/earnings ratios in the population of households aged 20-30 in the SCF. We use a discrete distribution of 75 equally spaced points located between -2.375 and 19.375. We assume that the initial draw of earnings is independent of the initial draw of this ratio. This is consistent with an empirical correlation of 0.02.

imposed.[12]

Social security benefits Social security benefits are set as a concave function of the final permanent component of earnings – a function which is designed to mimic the progressivity of the actual US system. This is achieved by specifying that benefits are equal to 90% of the final permanent component up to a given bend point, 32% from this first bend point to a second bend point, and 15% beyond that. The two bend points are set at, respectively, 0.18 and 1.01 times average earnings, based on the US system in 1990. Benefits are then scaled up proportionately so that the ratio of total benefit payments to total labor income is the same as in the US economy. This ratio was 9.5% in 1990.

Tax function We assume that taxes are levied on labor income only, and comprise of two parts. The first part is a non-linear tax function of the form estimated by Gouveia & Strauss (1994) and used, for example, by Castaneda et al. (2003). The second part is a proportional tax on earnings that is intended to reflect employer and employee social security contributions. The explicit functional form is given by

$$\tau(Y) = \tau^{SS} Y + \tau^b \left[Y - \left(Y^{-\tau^\rho} + \tau^s \right)^{-\frac{1}{\tau^\rho}} \right].$$

The value for τ^{SS} is set at 0.124 which is the combined employer and employee rate for social security contribution. The values for τ^b and τ^ρ are taken from Gouveia & Strauss (1994). They are set at $\tau^b = 0.258$ and $\tau^\rho = 0.768$. The value for τ^s is then chosen so that the ratio of personal current tax receipts (not including social security contributions) to labor income is the same as for the US economy in 2006. This ratio is 17%.

Consumption floor Following Hubbard et al. (1995), we set $\underline{C}$ to \$4,000.

4 Results

All our results are based on simulating, from the invariant distribution of the economy, an artificial panel of 50,000 households for 71 periods, i.e. a life-cycle. We have verified

[12]In computing allocations, we discretize the two components of the income process. This implies that the realization of both components are always positive so that the natural borrowing limits are below zero. In a typical simulation of our economy with $\underline{a} = 0$, the no-borrowing limit is a binding constraint for about 11% of households. These are primarily young households: the fraction who are constrained decrease from 44% at age 25 to zero at age 52.

that increasing further the sample size does not lead to any change in the results.[13] To be consistent with the BPP approach, when computing insurance coefficients, log consumption and log earnings (c_{it}, y_{it}) are defined as residuals from a common age profile.

4.1 Insurance coefficients in the data and the model

We start by reporting insurance coefficients, for permanent and transitory shocks, in the data and the benchmark model based on the BPP methodology. These are shown in Table 1. In all tables, columns labelled "Model BPP" refer to the model's insurance coefficients calculated using the instrumental variables approach described in Section 2.1, i.e. ϕ^x_{BPP}. Columns labelled "Data BPP" report the BPP empirical estimate (with associated standard errors) from the merged PSID/CEX data set (1980-1992).

How much insurance does the Bewley model permit? We find that for disposable (i.e., after tax) household income, using the BPP methodology generates insurance coefficients of 0.14 for permanent shocks and 0.91 for transitory shocks. These figures compare to insurance coefficients of, respectively, 0.36 and 0.95 in the data on US households in the 1980s. Hence, whereas the model generates approximately the right amount of insurance with respect to transitory shocks, the amount of insurance against permanent shocks is substantially less than in the US economy.

The second row of Table 1 reports insurance coefficients with respect to shocks to pre-government earnings. Redistribution through taxes and public transfers mediates the impact of the shock in the second row but not in the first, hence one should expect a higher insurance coefficient here. We find that public insurance through the tax and transfer system has very little effect on the transmission of transitory earnings shocks to consumption, a result that lines up with the data. However, the US tax and transfer system appears to provide a great deal of insurance against permanent shocks: moving from disposable income to pre-tax labor earnings increases the empirical BPP estimates from 0.38 to 0.69. The same is qualitatively true in the model, where the insurance coefficient almost doubles.

In sum, whether one focuses on pre- or post-government measures of earnings, the Bewley

[13]The model is solved using the method of endogenous grid points (see Carroll (2006)) with 100 exponentially spaced grid points for assets. The decision rule is constrained to be linear between grid points. The permanent component is approximated using a discrete Markov chain with 39 equally-spaced points on a age-varying grid chosen to match the age-specific unconditional variances. The transitory component is approximated with 19 equally-spaced points.

model generates less insurance against permanent shocks, and roughly the same level of insurance against transitory shocks, relative to the US data.

Insurance coefficients by wealth level Following BPP, we define households in the bottom 20% of the wealth distribution as "Low Wealth", and the remainder as "High Wealth". The insurance coefficients for these two groups are reported in the last two rows of Table 1. With regard to transitory shocks, the table reveals that in the model, all households are equally well insured, while in the data there appears to be far less insurance for low-wealth households (0.72 in the data, 0.89 in the model). This could be, at least in part, due to the fact that we have assumed very loose borrowing constraints in the benchmark model. In the next section we analyze an economy with tighter borrowing constraints.

The amount of insurance against permanent shocks varies substantially with wealth levels, both in the model and the data. However, for both groups, the model generates less insurance than in the data: -0.22 vs 0.15 for low wealth households, and 0.15 vs 0.38 for high wealth households. The negative insurance coefficient for the wealth-poor may appear surprising: taken at face value, it means that consumption responds more than one for one to income shocks. However, as we will see below, there is a very severe downward bias in the BPP methodology for low wealth households.[14]

4.2 Accuracy of the BPP methodology

We now assess the accuracy of the BPP methodology for estimating insurance coefficients. This can be done by comparing the columns labelled "Model BPP" and "Model TRUE". This latter label refers to the model's insurance coefficients calculated directly from the realizations of the individual shocks. Table 1 reveals that whereas the BPP methodology works extremely well for transitory shocks, it tends to systematically underestimate the amount of insurance for permanent shocks. This suggests that the actual empirical insurance coefficients for permanent shocks may be higher than 0.36 (for disposable labor income) and

[14]BPP compute insurance coefficients also by education level, and find large differences between high-school graduates and college graduates in estimates of ϕ^{η}_{BPP}.If we interpret education differentials as differences in lifetime income, one simple way of capturing the effects of education in the model is to group people based on the level of their initial permanent component of income, z_0. When we group households this way, we find virtually no differences in insurance coefficients. Since z_t is permanent, shocks are multiplicative, and preferences are CRRA, consumption responds in the same way to shocks, independently of z_0. The natural borrowing constraint assumption is important. When the debt limit is exogenous, non-zero, and is not proportional to income, it will bind more often for low income households.

0.69 (for pre-tax earnings).

Failure of orthogonality conditions This downward bias in estimates of insurance coefficients for permanent shocks is exacerbated whenever the agent is close to the borrowing limit. Thus, it is particularly severe for households with low wealth holding, as shown above and for young households, as we will see below.

The reason for the large bias in ϕ_t^η is that the orthogonality conditions in (7) may fail when agents have low wealth holdings.[15] It turns out that both covariances in (7) contribute to the negative bias. However, the quantitatively more important term is $cov\left(\Delta c_{it}, \varepsilon_{i,t-2}\right) < 0$. To gain intuition for why this covariance may be negative near the borrowing limit, consider a household who receives a negative transitory shock at $t-2$ (i.e., $\varepsilon_{t-2} < 0$). Such a household would like to borrow (or dissave) to smooth the negative shock. However, for a household close to its borrowing limit, even a small reduction in wealth can have a large expected utility cost because of the possibility of becoming constrained in a future period. Thus smoothing entails an optimal drop in consumption at $t-2$ –the closer agents are to the borrowing constraint, the larger this drop. This leads to a *positive* expected change in consumption in the next period, i.e. $cov\left(\Delta c_{t-1}, \varepsilon_{t-2}\right) < 0$ as consumption returns to its baseline level. Since agents prefer smooth paths for consumption, this adjustment takes place gradually, and $cov\left(\Delta c_t, \varepsilon_{t-2}\right) < 0$ as well.

Figure 1 illustrates this argument graphically. It depicts the impulse response function of log consumption c_{it} to a negative transitory income shock of 1.5 standard deviations at age $t = 5$.[16] Consumption drops at impact, more so for agents with low wealth, and returns back to its long-run level only very slowly, thereby the negative correlation between Δc_t and ε_{t-2} which is the source of the downward bias.

Small-sample bias Even though we have mainly interpreted the data-model discrepancy in the BPP coefficients as failure of the orthogonality conditions assumed by BPP, there is an additional source of discrepancy. While in the model's simulations we use a very large

[15]Recall that this specific assumption is required for the identification of insurance coefficients for permanent shocks, but not for transitory shocks.

[16]Precisely, the impulse response function at time $t+j$ with respect to a transitory shock $\bar{\varepsilon} = -1.5\sqrt{\sigma_\varepsilon}$ occurring at t is constructed as:

$$\Psi_{t+j} = E_{t+j}\left[\ln\left(c_{i,t+j}\left(\bar{\boldsymbol{\eta}}_{t+j,}\boldsymbol{\varepsilon}_{t+j}\right)\right) - \ln\left(c_{i,t+j}\left(\bar{\boldsymbol{\eta}}_{t+j,}\bar{\boldsymbol{\varepsilon}}_{t+j}\right)\right)\right],$$

where $\boldsymbol{\varepsilon}_{t+j}$ is a vector of transitory shocks equal to 0 for all $j > 0$, and equal to $\bar{\varepsilon}$ for $j = 0$, while $\bar{\boldsymbol{\varepsilon}}_{t+j}$ and $\bar{\boldsymbol{\eta}}_{t+j}$ are vectors of zeros.

sample, the BPP estimates are based on a smaller sample of around 17,000 household/year observations, or roughly 1,300 households. To assess the magnitude of the small-sample bias, we have run 50 simulations of samples with 1,300 households each. The means of both the true and the BPP coefficients are virtually unchanged, so we conclude that the small-sample bias is negligible.

4.3 Insurance coefficients by age

Permanent shock In Figure 2, we plot age-specific insurance coefficients from the benchmark model for permanent and transitory shocks to disposable earnings. Although the corresponding coefficients for the US economy are not reported by BPP, they state in a footnote (BPP, page 25) that when they allow the insurance coefficient for permanent shocks to vary linearly with age, they find that the slope is negative (-0.0059) and not significantly different from zero.

Figure 2 reveals a very different scenario. The insurance coefficients for permanent shocks ϕ_t^η are mildly decreasing at young ages, but are increasing steadily after age 35 and are markedly convex. Permanent shocks rescale the entire earnings profile during work, and they also have an effect on retirement income. When these shocks hit early in life, agents tend to reduce consumption almost one for one, since they have a long horizon before them.

As agents accumulate financial wealth, for precautionary or life-cycle reasons, they consume more out of financial wealth and less out of human wealth (i.e., the expected discounted value of their earnings), so permanent shocks to earnings have a smaller impact on consumption. Moreover, as the time horizon shortens, permanent shocks become "less permanent" and begin to look more and more like transitory shocks, and hence accumulated wealth can be used to mitigate their impact on consumption.[17] This explains why insurance coefficients rise with age.

The reason for the small decline in ϕ_t^η over the first ten years is, again, related to the evolution of wealth. Agents face a deterministic age-profile for earnings which is increasing and concave. As a result, there is a strong incentive to borrow early in life to smooth consumption. With loose borrowing constraints, this means that the average wealth levels are decreasing and negative at young ages (they reach their bottom around age $30-35$). As

[17] Even at age $t = 60$ though, permanent shocks are not fully transitory since they affect income during retirement through the pension system. In absence of pensions, we would have that $\phi^\eta = \phi^\varepsilon$ at age 60.

explained, insurance coefficients for permanent shocks are increasing in financial wealth.[18]

Figure 2 also plots the BPP estimator (dashed line). Note the downward bias, especially evident in the age range 30-35, when many agents have negative wealth and are near their debt limit.

Transitory shock The insurance coefficients for transitory shocks ϕ_t^{ε} are above 0.9 at all ages and slightly decrease with age. The loose borrowing constraints imply that young households can smooth the effects of transitory shocks even though they have not accumulated much precautionary wealth. The decrease with age is due to the shortening time horizon. A negative transitory income shock is effectively "transitory" insofar as there are remaining future dates in which an offsetting positive shock may be received. Finally, note that the BPP estimator is extremely accurate at every age.

4.4 Sensitivity analysis on the benchmark economy

Table 2 reports a wide set of sensitivity analyses on the baseline economy. The second line of the table shows that allowing for an initial wealth distribution–calibrated on the asset holdings of the young in the SCF–has very little effect on the insurance coefficients.

Households with high levels of risk aversion are less tolerant of consumption fluctuation, thus as γ rises insurance coefficients for permanent shocks increase. However, only for values of γ beyond fifteen we reach insurance coefficients close to those estimated in the data.[19]

Smaller wealth income ratios (obtained by lowering β) map into smaller asset holdings that can be used to smooth income shocks. It is thus not surprising that when the ratio is reduced to 2.5 and 1.5, insurance coefficients are reduced considerably, generating an even

[18] The true insurance coefficient ϕ_t^{η} may go slightly negative over the first decade. A negative value for ϕ_t^{η} is obtained when $cov\left(\Delta c_{it}, x_{it}\right) > var\left(x_{it}\right)$, i.e. consumption responds more than one-for-one to a particular shock. The reason this may happen is due to the interaction of transitory shocks and permanent shocks in the model, as uncovered by Carroll (1997). With $\sigma_{\varepsilon}^2 > 0$, households will accumulate a target level of wealth which they use to buffer the effects of transitory shocks. When a positive permanent shock hits, transitory shocks become a smaller component of lifetime income, both in the current period, and in all future periods. Hence the utility cost of not being able to smooth transitory shocks falls. Households reduce the optimal level of wealth they desire to buffer transitory shocks. Consumption may thus respond to the full effect of the positive permanent shock, plus an additional amount that is the decrease in the optimal precautionary wealth level. A similar logic applies to negative permanent shocks. We have verified that when we simulate the model without transitory shocks $\left(\sigma_{\varepsilon}^2 = 0\right)$, then ϕ_t^{η} is always positive.

[19] With our preferences, the coefficient of relative risk aversion equals the intertemporal elasticity of substitution. We conjecture that it is the latter that matters for insurance coefficients. To know for sure, one would have to solve the model with Epstein-Zin preferences. This is beyond the scope of this project. It is conceivable that one could use this avenue to reconcile model with data, while maintaining plausible levels of risk aversion.

larger discrepancy with the data. Note how the bias in the BPP coefficient grows, as K/Y is reduced. Moreover, the amount of insurance in the model does not depend on the size of the shocks, when the latter is varied within a plausible range.

Finally, Table 2 shows clearly that our computed insurance coefficient against transitory shocks is extremely robust across different parameterizations.

Overall, none of the sensitivities we ran change our two main conclusions. First, under a permanent-transitory income process, the Bewley model allows less smoothing against permanent shocks than what is estimated in the data, and about the right amount against transitory shocks. Second, BPP estimates of insurance coefficient for permanent shocks are always downwards biased, while for transitory shocks they are extremely accurate.

4.5 A Bewley economy with tight borrowing constraint

We now contrast the benchmark economy with an alternative economy where agents are not allowed to borrow, i.e., $\underline{A} = 0$. We keep all parameter values unchanged, except for β which is set to generate a wealth-income ratio of 3.5. The results of this experiment are summarized in Table 3. In passing, we note that this economy coincides with Carroll's buffer-stock model.

The true insurance coefficient against permanent shocks is 0.19 just above the baseline value of 0.17. This economy has higher precautionary saving, since agents want to avoid hitting the tighter borrowing constraint, thus it has slightly more self-insurance. Interestingly, the insurance coefficients against transitory shocks is significantly lower: 0.80 compared to a baseline value of 0.92. In the baseline model, the loose borrowing limit is very helpful in smoothing transitory shocks, especially at young ages.

As one could have anticipated, the BPP methodology severely underestimates ϕ^{η}. The bias is of the order of a factor of two. The reason is that in this economy there is substantial bunching of agents at values of wealth exactly at or near the borrowing limit.

Figure 3 plots insurance coefficients by age. First, note that the declining portion of the age profile at young ages disappears. The discussion of Section 4 suggested that it was the combination of an increasing earnings profile and the ability to borrow that generates the decreasing profile of insurance coefficients at young ages. Second, the BPP methodology performs very well at old ages, when households have accumulated wealth, but the bias until age 35 is very large, and is responsible for the bias in the average coefficient. Turning to the transitory variance, when we impose a no-borrowing constraint, the age pattern of the transitory insurance coefficients changes dramatically: it starts at around 0.4 at age 25, and increases in a concave fashion to 0.9 by age 40. As explained, young workers have little wealth and cannot borrow. As such, they are unable to smooth transitory shocks until they have cumulated enough precautionary savings.

The bottom line is that if one takes the view that borrowing limits are somewhat tight in the US economy, then: 1) the BPP methodology grossly underestimates insurance coefficients for both permanent and transitory shocks, and 2) the Bewley model admits less insurance than the US economy both for permanent and for transitory shocks.

5 Advance information

In this section, we assess whether allowing for the possibility that agents know more about their future income growth than the econometrician can reconcile the higher insurance coefficients estimated in the data, compared with those computed in the benchmark model. The idea is that if part of a measured current income change was known to the agents in advance, then this change would have been already incorporated into consumption at the time it was learned, and would not affect current consumption growth.[20] The contemporaneous correlation of measured income growth with consumption growth would thus be lower than if all of the growth in income was news.

In order to assess the effects of advance information on insurance coefficients, it is necessary to take a stand on the particular form that the advance information takes. We examine two cases which make different assumptions about the timing of receipt of the advance information. In the first model, agents learn about a component of their permament shock to income one period in advance. One interpretation is that of receiving a signal about a future pay-rise, wage cut, promotion, or demotion, in the period before the change actually takes place. In the second model we allow agents to learn about a deterministic trend component of their future income profile upon entry in the labor market, i.e. at age $t = 0$. This model is a version of the heterogeneous income profiles model studied by Lillard and Weiss (1979), Baker (1997), Haider (2001), and Guvenen (2007), among others. One interpretation of this view is that, by choosing an occupation, an individual knows what income profile to expect,

[20] See, for example Cunha, Heckman and Navarro (2005), Guvenen (2007), and Primiceri and van Rens (2008).

on average.

We are also interested in knowing what the BPP estimator captures in these two cases. BPP state that, in the presence of advanced information, "the estimated coefficient has to be interpreted as reflecting a combination of insurance and information". We will show that, while this is exactly true in the second model, in the first only the insurance component is reflected in the estimates of ϕ^{η}_{BPP}.[21]

5.1 Advance information I: preempting permanent shocks

Consider a modification of the information set of the agent whereby the permanent change in income at time t, η_{it}, consists of two orthogonal components, η^s_{it}, and η^a_{it}. The component η^s_{it} becomes known to the agent at time t and affects income at time t. The component η^a_{it} was in the agent's information set already at time $t-1$, but is only incorporated into income at time t. The permanent component of earnings is hence given by

$$z_{it} = z_{i,t-1} + \eta^s_{it} + \eta^a_{it}$$

where $E\left[\eta^s_{it}\right] = E\left[\eta^a_{it}\right] = 0$, and the variances of the two components are changed so that σ_η is always kept constant at its baseline value of 0.02.

From the definition of insurance coefficient for permanent shocks,

$$\begin{aligned} \phi^{\eta} &= 1 - \frac{cov\left(\Delta c_{it}, \eta_{it}\right)}{var\left(\eta_{it}\right)} = 1 - \frac{cov\left(\Delta c_{it}, \eta^s_{it} + \eta^a_{it}\right)}{var\left(\eta^s_{it} + \eta^a_{it}\right)} \\ &= \left(1-\alpha\right)\phi^{\eta^s} + \alpha\phi^{\eta^a} \\ &\approx \left(1-\alpha\right)\phi^{\eta^s} + \alpha \end{aligned} \tag{11}$$

where ϕ^{η^s} and ϕ^{η^a} are "insurance coefficients" with respect to the two components of the permanent shock, and α is the share of the variance of permanent shocks which is known in advance (or, advance information ratio). The approximate equality in the last line holds because, when borrowing constraints are unimportant, as for the baseline economy, $cov\left(\Delta c_{it}, \eta^a_{it}\right) \approx 0$. It follows that our insurance coefficient ϕ^{η} is a combination of true smoothing $\left(\phi^{\eta^s}\right)$ and advance information, whose relative magnitude is regulated by α.

[21] It is often read in the literature that one testable implication of the advance information hypothesis is that changes in consumption at age t should be correlated with future changes in income at ages $t+j$, with $j>0$. BPP (section 3.2.3) argue that there is no evidence for this correlation in their data. However, it is difficult to know the power of such a test, particularly in the presence of measurement error and imputed consumption data. Moreover, this implication of advance information is accurate for the first model (when $j=1$), but not for the second, since there all the relevant knowledge is already acquired at age $t=0$.

When the BPP methodology is used to estimate insurance coefficients for permanent shocks, ignoring the downward bias discussed in Section 4.2, we have

$$\begin{aligned}
\phi^{\eta}_{BPP} &= 1-\frac{cov\left(\Delta c_{it},\Delta y_{i,t-1}+\Delta y_{it}+\Delta y_{i,t+1}\right)}{cov\left(\Delta y_{it},\Delta y_{i,t-1}+\Delta y_{it}+\Delta y_{i,t+1}\right)}=1-\frac{cov\left(\Delta c_{it},\eta^{s}_{it}+\eta^{a}_{i,t+1}\right)}{var\left(\eta^{s}_{it}+\eta^{a}_{it}\right)} \qquad (12)\\
&= 1-\frac{cov\left(\Delta c_{it},\eta^{s}_{it}\right)}{var\left(\eta^{s}_{it}\right)}\left[\frac{var\left(\eta^{s}_{it}\right)}{var\left(\eta^{s}_{it}+\eta^{a}_{it}\right)}\right]-\frac{cov\left(\Delta c_{it},\eta^{a}_{i,t+1}\right)}{var\left(\eta^{a}_{it}\right)}\left[\frac{var\left(\eta^{a}_{it}\right)}{var\left(\eta^{s}_{it}+\eta^{a}_{it}\right)}\right]\\
&= \left(1-\alpha\right)\phi^{\eta^{s}}+\alpha\left[1-\frac{cov\left(\Delta c_{it},\eta^{a}_{i,t+1}\right)}{var\left(\eta^{a}_{it}\right)}\right]\\
&\approx \phi^{\eta^{s}}.
\end{aligned}$$

As evident from the third line, the BPP estimator is a weighted average between the insurance coefficient against the current shock $\left(\phi^{\eta^{s}}\right)$ and a term which looks like an insurance coefficient against the component of the $t+1$ shock which is known at t. This last term enters the expression through the component $\Delta y_{i,t+1}$ of the BPP instrument, i.e. assumption (8) fails to hold.

Since, in the absence of severe borrowing constraints, consumption growth Δc_{it} should react equally to η^{s}_{it} and to $\eta^{a}_{i,t+1}$ (except for a minor difference due to discounting), we have $\phi^{\eta}_{BPP}\approx\phi^{\eta^{s}}$, as stated in the last line. We can thus conclude that the BPP coefficient is roughly independent of the amount of advance information. Whereas the true insurance coefficient, ϕ^{η}, reflects a combination of insurance and advance information as seen in (11), the BPP coefficient ϕ^{η}_{BPP}, and hence the estimates for the US economy, reflect only insurance.

Table 4 reports the true coefficients ϕ^{η}, the insurance coefficients with respect to current shock $\phi^{\eta^{s}}$, and the BBP estimator ϕ^{η}_{BPP} in the economies with loose constraints and with no borrowing, for different values of the advance information ratio α. When this ratio equals zero we are in the baseline model, when it's one we are in a model where agents fully learn their permanent shock one period ahead.

The true insurance coefficients for permanent shocks ϕ^{η} increase with α and are always larger than $\phi^{\eta^{s}}$ which confirms our derivation in (11) . The BPP estimator, instead, is roughly equal to $\phi^{\eta^{s}}$ and invariant to the amount of advanced information. In the zero-borrowing economy, the usual severe downward bias is also at work. To sum up, this form of advance information about permanent shocks cannot explain the gap between the values of the BPP insurance coefficient in the data and in the model.

For transitory shocks, Table 4 shows the opposite result. The true insurance coefficients

ϕ^{ε} are not affected by the presence of advanced information because the response of consumption growth to transitory shocks is not affected by the timing of news about permanent shocks. However, the BPP estimator ϕ^{ε}_{BPP} has an upward bias which increases with the size of α. To understand this bias, note that:

$$\begin{aligned}\phi^{\varepsilon}_{BPP} &= 1 - \frac{cov\left(\Delta c_{it}, \Delta y_{i,t+1}\right)}{cov\left(\Delta y_{it}, \Delta y_{i,t+1}\right)} = \phi^{\varepsilon} + \frac{cov\left(\Delta c_{it}, \eta^{a}_{i,t+1}\right)}{var\left(\varepsilon_{it}\right)} \\ &\approx \phi^{\varepsilon} + \left(1 - \phi^{\eta^{s}}\right)\alpha\left(\frac{\sigma^2_{\eta}}{\sigma^2_{\varepsilon}}\right)\end{aligned} \tag{13}$$

The upward bias results from a failure of the identification assumption (11), since a fraction of the future (permanent) income shock $\Delta y_{i,t+1}$ is known in advance and transmits to consumption. The second line uses the fact that $cov\left(\Delta c_{it}, \eta^{a}_{i,t+1}\right) \approx cov\left(\Delta c_{it}, \eta^{s}_{i,t}\right)$, and recall that $\sigma^2_{\eta}/\sigma^2_{\varepsilon}$ is constant throughout the simulations.

5.2 Advance information II: heterogenous income slopes

Consider a generalization of the income process in (5) that includes heterogeneous slopes in individual income profiles:

$$\begin{aligned} y_{it} &= \beta_i t + z_{it} + \varepsilon_{it} \\ z_{it} &= z_{i,t-1} + \eta_{it}, \end{aligned}$$

with $E\left[\beta_i\right] = 0$ in the cross-section, and $var\left[\beta_i\right] = \sigma_{\beta}$.[22] We assume that β_i is learned by the agents at age zero. In the experiments that follow we keep σ_{ε} as in the benchmark calibration, change the value for σ_{β} and set σ_{η} residually so that the overall cross-sectional variance of log income is unchanged from the benchmark. To have a sense of the size of advance information in each experiment, we report the fraction of the dispersion in log earnings at age 60 that is known by the agents upon entering the labor market.[23]

The results of this experiment are reported in Table 5. We find that the true insurance coefficients for permanent and transitory shocks $(\phi^{\eta}, \phi^{\varepsilon})$ are unchanged from the benchmark model. The reason is that the full effect of knowledge about β_i is incorporated into consumption from the outset. Thus dispersion in β_i translates into dispersion in the *level* of

[22] We retain the unit root specification for the permanent component of the income process, notwithstanding the empirical evidence for substantially lower persistence when heterogeneous slopes are present. We do this to provide a clean analysis of the effects of heterogeneous slopes without confounding the effects of lower persistence. We separately analyze the issue of persistence in Section 6.

[23] In general, the fraction of dispersion at age t known at birth is computed as $\left(\sigma_{\beta} t^2\right)/var\left(y_{it}\right)$.

individual expected consumption profiles. Our insurance coefficients, however, are a measure of how much consumption growth responds to contemporaneous shocks. This response is not affected by the presence of heterogeneous slopes known at $t = 0$.

Table 5 also shows that the downward bias in the BPP estimator decreases (and eventually becomes positive) as the amount of advance information is increased. Ignoring the usual sources of downward bias due to the failure of assumption (7) , the BPP insurance coefficient is given by

$$\begin{aligned} \phi^{\eta}_{BPP} &= 1 - \frac{cov\left(\Delta c_{it}, \Delta y_{i,t-1} + \Delta y_{it} + \Delta y_{i,t+1}\right)}{cov\left(\Delta y_{it}, \Delta y_{i,t-1} + \Delta y_{it} + \Delta y_{i,t+1}\right)} \\ &= 1 - \frac{cov\left(\Delta c_{it}, \eta_{it} + 3\beta_i\right)}{var\left(\eta_{it}\right) + 3var\left(\beta_i\right)} \\ &= \left(1-\alpha\right)\phi^{\eta} + \alpha\left[1 - \frac{cov\left(\Delta c_{it}, \beta_i\right)}{var\left(\beta_i\right)}\right] \\ &\approx \left(1-\alpha\right)\phi^{\eta} + \alpha \end{aligned} \tag{14}$$

where here we have defined $\alpha = 3var\left(\beta_i\right) / \left[var\left(\eta_{it}\right) + 3var\left(\beta_i\right)\right]$. The last row descends form the fact that, with loose borrowing constraints, Δc_{it} should be roughly invariant to β_i at any t. The BPP estimator displays an upward bias, and the bias is larger, the larger is α. However, even in the case where 80% of the variance of income at age 60 is known already at age 25, arguably an upper bound for advance information, the BPP coefficient is below the empirical estimate of 36%.

Finally, note that in the economy with tight borrowing constraints, the bias remains negative and worsens as one increases the amount of advance information in the economy. The reason is that, as σ_β grows, the economy is populated by a larger fraction of agents with steep income profile who would like to borrow against their future income, but they are liquidity constrained. As already explained, the source of downward bias in the BPP estimator gets severe when liquidity constraints are tight.

6 Persistent income shocks

Following BPP, until now we have focused on a particular income process that restricts shocks to be either fully permanent or fully transitory. There is no scope for income shocks that have lasting, but not permanent, effects on income. In this section we relax this assumption. One plausible explanation for why we find higher insurance coefficients in the data than in model is that, in reality, shocks are not purely permanent. Persistent shocks are easier to smooth, especially earlier in the life-cycle, by precautionary saving and borrowing.

Consider a variant of the income process whereby z_t follows an AR(1) process with parameter $\rho < 1$, rather than a random walk:

$$z_{it} = \rho z_{it-1} + \eta_{it}.$$

In the terminology of the more general model in equation (1), we now have $a_0 = (1, 1)$ and $a_j = (\rho^j, 0)' \; \forall \; j \geq 0$.

Identification With this income process the identification strategy of Section 2.1 is no longer valid. We propose two new $g_t^x(\mathbf{y})$ functions that identify the two insurance coefficients for $x = \{\eta, \varepsilon\}$. In order to do this, we assume that an external estimate of ρ is available. This is a reasonable assumption, since ρ can be identified using panel data on income alone, and thus can be estimated in a first stage.

Define the quasi-difference of log income as $\tilde{\Delta} y_t \equiv y_t - \rho y_{t-1}$. Identification of the two insurance coefficients can be achieved by setting $g_t^{\varepsilon}(\mathbf{y}) = \tilde{\Delta} y_{t+1}$ and $g_t^{\eta}(\mathbf{y}) = \rho^2 \tilde{\Delta} y_{t-1} + \rho \tilde{\Delta} y_t + \tilde{\Delta} y_{t+1}$. For the transitory shock, we have

$$\begin{aligned} cov\left(\tilde{\Delta} y_{it}, \tilde{\Delta} y_{i,t+1}\right) &= -\rho var\left(\varepsilon_{it}\right), \\ cov\left(\Delta c_{it}, \tilde{\Delta} y_{i,t+1}\right) &= -\rho cov\left(\Delta c_{it}, \varepsilon_{it}\right), \end{aligned}$$

and for the persistent shock,

$$\begin{aligned} cov\left(\tilde{\Delta} y_{it}, \rho^2 \tilde{\Delta} y_{i,t-1} + \rho \tilde{\Delta} y_{it} + \tilde{\Delta} y_{i,t+1}\right) &= \rho var\left(\eta_{it}\right), \\ cov\left(\Delta c_{it}, \rho^2 \tilde{\Delta} y_{i,t-1} + \rho \tilde{\Delta} y_{it} + \tilde{\Delta} y_{i,t+1}\right) &= \rho cov\left(\Delta c_{it}, \eta_{it}\right). \end{aligned}$$

Thus, in both cases, expression (4) yields a consistent estimator of ϕ^x, under exactly the same pair of assumptions (7) and (8).

Results In Table 6 we present insurance coefficients from the model with $\rho < 1$ estimated on simulated data in three different ways. The column headed "Model TRUE" reports insurance coefficients calculated using the realized values of the shocks; the column headed "Model BPP" reports estimates using the estimation procedure just described; and the column headed "Model BPP misspecified" reports the estimates that would obtain if one were to use the (invalid) instruments from the model with permanent shocks. This last column is the correct model counterpart of the data.

There are three important findings. First, the estimates obtained with the misspecified BPP instruments are very close to those obtained with the correct instruments, at all levels of ρ. This is true both for the model with and without tight borrowing constraint. Hence the bias that results from applying the instruments from the permanent shock case, on data generated by an AR(1) process, is not at all severe. It is thus justified to take the empirical BPP estimate of 0.38 seriously, even in the case it was estimated under a misspecified income process. Of course, it is still true that the BPP methodology underestimates true insurance coefficients, but the bias is not increased by the model misspecification.

Second, the insurance coefficient for persistent shocks quickly increases as ρ declines from 1 towards 0.90. Even with an autoregressive parameter as high as 0.95, the amount of insurance against persistent shocks in the model is roughly consistent with that in the data. This implies that a Bewley economy with a highly persistent (but not permanent) income process can generate the right level of insurance against persistent shocks. With $\rho = 0.95$, in the baseline economy the misspecified BPP estimate is 0.34 and in the economy with the no-borrowing constraint it is estimated to be 0.30, vis-a-vis an empirical estimate of 0.36.

It is also interesting to note that as ρ decreases, the downward bias in the BPP estimator vanishes. With shocks which are less durable than unit root, precautionary savings are more useful. Agents start accumulating wealth right away, and move far from the debt constraint early in life, which explains why the bias is now very small.

Third, turning to insurance coefficients for the transitory shocks, here the model with persistent shocks is less successful. For $\rho = 0.95$, the misspecified BPP estimate the insurance coefficient is 0.82, while the empirical counterpart is 0.95. Moreover, our simulations show that this estimator is downward biased, hence the Bewley model with persistent shocks may imply an insurance coefficient against transitory shocks which is at least 0.15 below the unbiased empirical value. The reason for why the model generates less smoothing with respect to transitory shocks is that now agents shift the use of savings from the smoothing of transitory shocks to the smoothing of persistent shocks, and are willing to tolerate larger fluctuations in transitory shocks.

Figure 4 plots the age profile of insurance coefficients for persistent and transitory shocks for the case $\rho = 0.95$. Relative to the case with permanent shocks, the age profile of insurance coefficients is now much flatter (and thus more consistent with the data). This is because the difference between a permanent and a highly persistent shocks is more pertinent for a young household with many periods ahead.

Finally, we note that there is a connection between the advanced information hypothesis and the degree of persistence of shocks. Several authors, e.g. Guvenen (2007), have pointed out that when the specification for the income process is augmented with heterogeneous slopes, the residual shocks are estimated to have lower persistence. Intuitively, the individual-specific linear profile captures, in a statistical sense, the high autocovariance in levels that is commonly present in panel data on earnings.

7 Relation to the literature

The investigation on how effectively households can smooth income fluctuations when making their consumption decisions is a classic research topic in economics. The key input of the analysis is the joint use of consumption and income data. Broadly speaking, one can observe co-movements between income and consumption at three different levels. First, one can study how consumption growth responds to income growth at the individual level. This is the approach taken by BPP and by this paper. Second, one can study the joint evolution of income and consumption inequality over the life-cycle. Third, one can investigate the joint dynamics of the cross-sectional dispersion of income and consumption in the time series.

We begin from the first approach. Dynarski & Gruber (1997) provide an overview of the theory and the data problems associated to the question. The typical exercise in the literature, including the one performed by Dinarski and Gruber, compares the sensitivity of individual consumption growth to income growth, i.e., the marginal propensity to consume (MPC) out of income, without attempting to unbundle permanent from transitory shocks.

One notable exception is Hall & Mishkin (1982) who exploit theoretical restrictions from the PIH to derive identification of the response of consumption growth to permanent and transitory shocks, separately. As in BPP, it is the covariance between consumption growth and lagged and future income growth that identifies the MPC out of the two shocks. The analysis of Hall and Mishkin focuses on the comparison between MPC out of transitory shocks in the PIH and the data. Based on simulations from a reasonably parameterized model, they report that the MPC in their model is around 0.08 while it is 0.30 in the data, a sign of excess sensitivity. Our calculations (and the BPP estimates) of the insurance coefficient can broadly be interpreted as one minus the MPC. While our simulations line up well with the Hall and Mishkin number, the BPP empirical estimate imply a much smaller MPC. The discrepancy may be due to the more comprehensive definition of consumption that BPP achieved through their imputation procedure. For example, a limit of the empirical work in Hall and Mishkin, relative to BPP, is that they only use food consumption from the PSID.[24]

Carroll (1997) extensively characterized the problem of a consumer facing a permanent-transitory income process and a zero natural borrowing limit. His key result is that households engage in, what he calls, "buffer-stock" saving behavior, i.e. they target a ratio of wealth to permanent income. As we showed, this characterization is very useful to understand some of our results in spite of some modelling differences in borrowing constraints, retirement period, government redistribution, etc. Interestingly, Carroll (2001) reports results from a set of simulations of his buffer-stock economy where he calculates the MPC out of permanent income shocks to be between 0.80 and 0.95, i.e. slightly higher than what our findings would imply, but in the same broad range.

Attanasio & Pavoni (2007) argue that the excess smoothness puzzle, i.e. the mild response of consumption to permanent shocks in excess of what is predicted by the PIH (Campbell & Deaton (1989)), can be rationalized by constrained-efficient allocations for an economy with private information. Their estimates for the United Kingdom show that, under the permanent-transitory income shock decomposition where, as we documented, the bond economy provides insufficient insurance to the agents, while the private information environment achieves insurance levels closer to the data.

Since the seminal work of Storesletten, Telmer & Yaron (2004), several authors (e.g., Kaplan (2007*a*)) have tried to assess whether the Bewley model can replicate the joint life-cycle path of income and consumption inequality. Storesletten et al. (2004) used the Deaton & Paxson (1994) cohort-based estimate of the growth in the variance of log consumption over the life cycle [0.15 from age 25 to age 55]. They concluded that the model broadly matches the data. This type of life-cycle analysis is plagued by the problem of separating time and cohort effects in identifying the age profile in the data. In addition, results can be sensitive to the time period under consideration. For example, more recent estimates

[24]We refer to BPP for a discussion of the differences between their approach and Hall and Mishkin.

(Heathcote, Storesletten & Violante (2005)) set the gradient of the consumption inequality over the life cycle to a third of the Deaton-Paxson estimate. As a result, what target one should use is not yet clear. By using year-by-year individual-level consumption and income growth, we don't face this problem.

A number of recent papers tries to explain the joint evolution of consumption and income inequality in the aggregate cross-section. Krueger & Perri (2006) and Heathcote, Storesletten & Violante (2007) argue that the small rise in household consumption inequality and the comparatively large rise in earnings inequality witnessed in the United States in the last thirty years cannot be replicated in a standard Huggett (1993) bond economy, but one needs a model economy with more insurance opportunities. In Krueger and Perri, the latter come from their limited enforcement environment, while in Heathcote, Storesletten and Violante they originate from explicit insurance against some permanent shocks, but not others, as explained in Section 2.1.[25] Blundell & Preston (1998) apply a "reverse-engineering" approach to the question by estimating the rise in transitory income dispersion that would make a PIH economy reproduce the observed rise in consumption inequality for the United Kingdom.

Finally, a series of recent papers (Cunha, Heckman & Navarro (2005), Guvenen (2007), Huggett, Ventura & Yaron (2006), Primiceri & van Rens (2006)) argue that advance information about future income fluctuations can help in rationalizing the joint behavior of cross-sectional income and consumption inequality in the U.S. economy.

8 Conclusions

This paper is inspired by the intriguing empirical results reported by Blundell, Pistaferri and Preston (2006, BPP). BPP estimate that in the US economy, during 1980-1992, 36% of permanent income shocks and 95% of transitory shocks to post-government income can be insured away by households, i.e. do not translate into contemporaneous consumption growth. These two numbers, we argue, should become central in macroeconomics. They represent a benchmark against which current incomplete-markets macroeconomic models used for quantitative analysis should be compared in order to assess whether they admit the right amount of household insurance.

[25] Heathcote, Storesletten & Violante (2004) augment the standard bond economy to allow for additional margins through which households can respond to individual wage shocks (education decision, labor supply, within-family insurance, tax and pension system). Taking as given the observed rise in male and female wage inequality, the model roughly generates the right rise in consumption inequality for the United States.

In this paper, we make a step forward in this direction by addressing two questions: 1) Is the Bewley model, arguably the workhorse of heterogenous agents macroeconomics, able to replicate such finding? And 2) Does the BPP methodology provide an unbiased estimator of insurance coefficients, if the US economy is accurately described by a Bewley model?

We have uncovered several interesting results. First, when the log-income process is assumed to be the sum of a permanent and a transitory component, then a plausibly calibrated Bewley model has substantially less insurance than the data against permanent shocks and about the same against transitory shocks. This finding is very robust. We have developed two alternative models of advance information, and we have argued that neither one can reconcile model and data.

Second, in the model there is a large amount of heterogeneity in insurance coefficients across age groups which is not present in the data: in the model the age profile is increasing and convex, whereas in the data it is flat. Hence, to bring together model and data, the model must be modified in such a way to provide more insurability to younger households. We have shown that allowing for a persistent $AR(1)$ shock ($\rho = 0.95$) instead of a pure permanent shock helps reconciling model and data. The overall degree of insurance is higher, and the age profile of insurance coefficients is flatter.

Fourth, we have assessed the accuracy of the estimation method proposed by BPP. We have shown that estimates of insurance coefficients are, in general, downward biased, with the bias exacerbated whenever households are close to their borrowing constraint. Put differently, the actual insurability of shocks may be higher than what was measured by BPP, especially for young and poor households.

This work suggests at least three important avenues for future research. The first one is to delve deeper into the actual durability of income shocks to households: we documented that modelling individual income fluctuations as a unit root or as a highly persistent first-order autoregressive process can make a significant difference in judging whether a bond-economy is a fair representation of the US economy in terms of household insurance opportunities. The appropriate statistical representation of individual income fluctuations is a long-standing question in labor and macro economics. Our work highlights an additional reason why a good answer to this question is paramount.

The second direction should explore whether endogenously incomplete markets models (e.g., environments with limited enforcement or private information) can replicate the two key BPP empirical estimates of household insurance against permanent and transitory shocks.

The misalignment between insurance coefficients in the model and the data is particularly acute for young individuals. Future research should try to identify additional sources of insurance against permanent shocks for the young, over and above borrowing and saving. Kaplan (2007*b*), who explores the role of co-residence decisions, is an example.

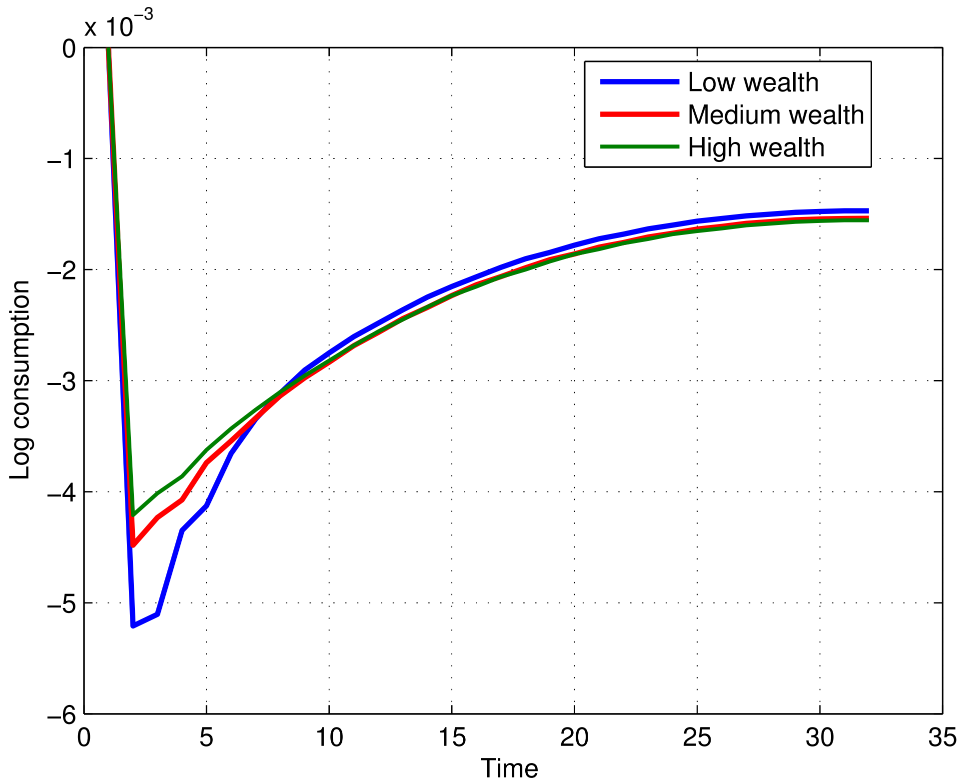

Figure 1: Response of log consumption to a negative transitory shock, for three different wealth levels.

	Permanent Shock			Transitory Shock		
	Data BPP	Model BPP	Model TRUE	Data BPP	Model BPP	Model TRUE
Disposable Income	0.36 (0.095)	0.14	0.17	0.95 (0.044)	0.91	0.91
Pre-Govt Earnings:	0.69 (0.057)	0.26	0.27	0.94 (0.031)	0.94	0.94
Low Wealth	0.15 (0.285)	−0.22	0.00	0.72 (0.114)	0.89	0.89
High Wealth	0.38 (0.100)	0.15	0.22	0.99 (0.041)	0.92	0.92

Table 1: Results from the benchmark model with natural borrowing limit

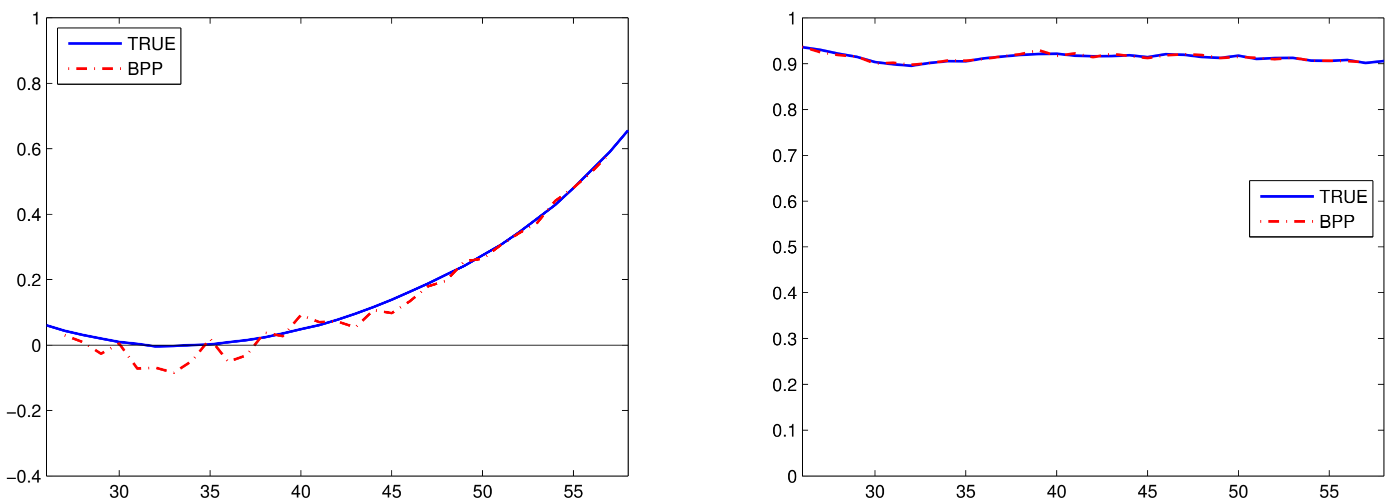

Figure 2: Age profiles of insurance coefficient for persistent shock (left panel) and transitory shock (right panel) in the benchmark model

	Permanent Shock		Transitory Shock	
Data	0.36 (0.095)		0.95 (0.044)	
	Model TRUE	Model BPP	Model TRUE	Model BPP
Benchmark	0.17	0.14	0.91	0.91
Initial Wealth Distribution	0.17	0.15	0.91	0.91
Risk Aversion:				
$\gamma = 1$	0.15	0.12	0.91	0.91
$\gamma = 5$	0.24	0.20	0.91	0.91
$\gamma = 10$	0.33	0.27	0.90	0.90
$\gamma = 15$	0.39	0.32	0.89	0.89
Wealth Income Ratio:				
$\frac{K}{Y} = 1.5$	0.10	-0.01	0.79	0.78
$\frac{K}{Y} = 2.5$	0.14	0.08	0.87	0.87
Variance Permanent Shock:				
$\sigma_\eta = 0.03$	0.20	0.17	0.90	0.90
$\sigma_\eta = 0.01$	0.16	0.13	0.93	0.93
$\sigma_\eta = 0.005$	0.16	0.13	0.93	0.93
Variance Initial Permanent:				
$\sigma_{z_0} = 0.2$	0.18	0.15	0.91	0.91
$\sigma_{z_0} = 0.1$	0.17	0.14	0.92	0.92
Variance Transitory Shock				
$\sigma_\varepsilon = 0.075$	0.18	0.14	0.90	0.90
$\sigma_\varepsilon = 0.025$	0.17	0.15	0.92	0.92

Table 2: Sensitivity analysis on the benchmark model

	Permanent Shock			Transitory Shock		
	Data BPP	Model BPP	Model TRUE	Data BPP	Model BPP	Model TRUE
Disposable Income	0.36 (0.095)	0.09	0.19	0.95 (0.044)	0.79	0.80
Pre-Govt Earnings:	0.69 (0.057)	0.23	0.28	0.94 (0.031)	0.88	0.89
Low Wealth	0.15 (0.285)	−0.36	0.03	0.72 (0.114)	0.44	0.45
High Wealth	0.38 (0.100)	0.07	0.23	0.99 (0.041)	0.87	0.87

Table 3: Results from the model with no borrowing

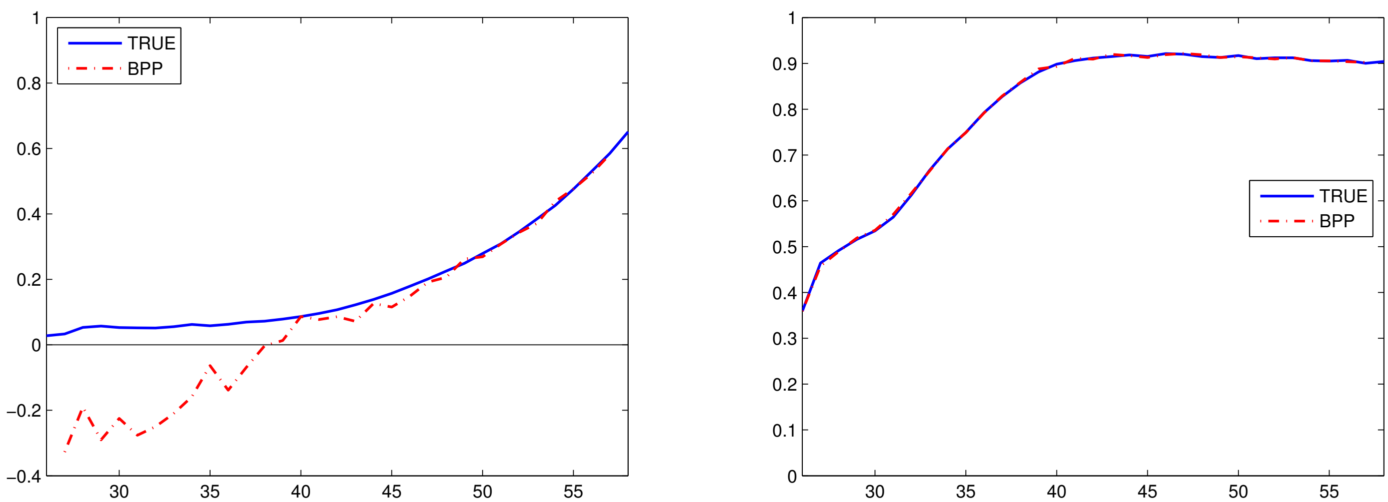

Figure 3: Age profiles of insurance coefficient for persistent shock (left panel) and transitory shock (right panel) in the model with no borrowing

	Permanent Shock			Transitory Shock	
Data	0.36 (0.095)			0.95 (0.044)	
	Model TRUE ϕ^{η}	Model TRUE (shock) ϕ^{η^s}	Model BPP ϕ^{η}_{BPP}	Model TRUE	Model BPP
Benchmark:					
$\alpha = 0.05$	0.21	0.17	0.17	0.91	0.93
$\alpha = 0.15$	0.29	0.17	0.15	0.91	0.96
$\alpha = 0.25$	0.37	0.17	0.16	0.91	0.99
No Borrowing:					
$\alpha = 0.05$	0.23	0.20	0.09	0.80	0.81
$\alpha = 0.15$	0.30	0.20	0.10	0.80	0.83
$\alpha = 0.25$					

Table 4: Results from model I with advance information: one period ahead preempting of permanent shocks

	Permanent Shock		Transitory Shock	
Data	0.36 (0.095)		0.95 (0.044)	
Fraction of $var\left(y_{i,60}\right)$ known at $t = 0$	Model TRUE	Model BPP	Model TRUE	Model BPP
Benchmark:				
20%	0.17	0.15	0.91	0.91
40%	0.17	0.16	0.91	0.91
60%	0.17	0.20	0.91	0.91
80%	0.17	0.29	0.92	0.92
No Borrowing:				
20%	0.19	0.07	0.79	0.79
40%	0.19	0.03	0.79	0.79
60%	0.19	−0.03	0.78	0.78
80%	0.19	−0.16	0.78	0.78

Table 5: Results from model II with advance information: heterogeneous earnings slopes known at age zero

	Persistent Shock			Transitory Shock		
Data	0.36 (0.095)			0.95 (0.044)		
	Model TRUE	Model BPP	Model BPP misspecified	Model TRUE	Model BPP	Model BPP misspecified
Benchmark:						
$\rho = 0.99$	0.23	0.21	0.20	0.91	0.91	0.90
$\rho = 0.97$	0.31	0.29	0.28	0.88	0.88	0.86
$\rho = 0.95$	0.37	0.35	0.34	0.86	0.86	0.82
$\rho = 0.93$	0.42	0.40	0.39	0.85	0.85	0.79
$\rho = 0.91$	0.45	0.44	0.43	0.84	0.84	0.75
No Borrowing:						
$\rho = 0.99$	0.23	0.16	0.15	0.80	0.79	0.79
$\rho = 0.97$	0.29	0.25	0.24	0.79	0.79	0.77
$\rho = 0.95$	0.34	0.32	0.30	0.79	0.79	0.76
$\rho = 0.93$	0.38	0.37	0.35	0.79	0.79	0.74
$\rho = 0.91$	0.42	0.41	0.39	0.79	0.79	0.71

Table 6: Results from the model with persistent earnings shocks

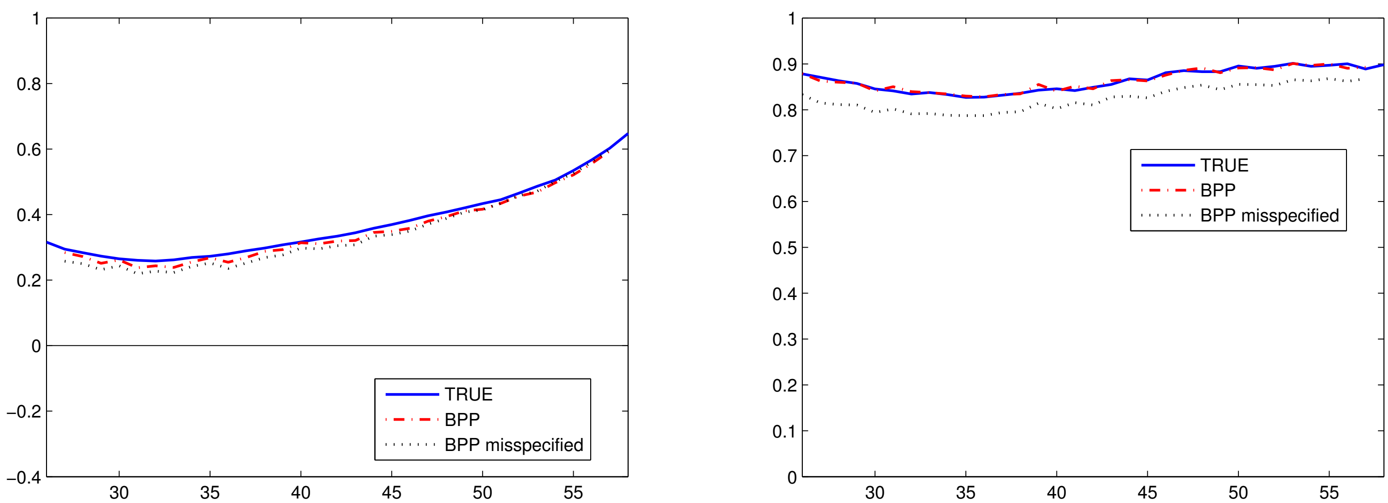

Figure 4: Age profiles of insurance coefficient for persistent shock (left panel) and transitory shock (right panel) in the model with persistent shocks ($\rho = 0.95$)

References

Abowd, J. & Card, D. (1989), 'On the Covariance Structure of Earnings and Hours Changes', *Econometrica* **57**(2), 411–445.

Aiyagari, S. (1994), 'Uninsured Idiosyncratic Risk and Aggregate Saving', *The Quarterly Journal of Economics* **109**(3), 659–684.

Attanasio, O. & Pavoni, N. (2007), 'Risk Sharing in Private Information Models with Asset Accumulation: Explaining the Excess Smoothness of Consumption', *Mimeo, UCL* .

Bewley, T. (1980), 'The Optimum Quantity of Money', *Models of Monetary Economics* pp. 169–210.

Bewley, T. (1983), 'A Difficulty with the Optimum Quantity of Money', *Econometrica* **51**(5), 1485–1504.

Blundell, R., Pistaferri, L. & Preston, I. (2004), 'Imputing Consumption in the PSID Using Food Demand Estimates From the CEX', *Mimeo, UCL* .

Blundell, R., Pistaferri, L. & Preston, I. (2006), 'Consumption Inequality and Partial Insurance', *Mimeo, UCL* .

Blundell, R. & Preston, I. (1998), 'Consumption Inequality and Income Uncertainty', *The Quarterly Journal of Economics* **113**(2), 603–640.

Campbell, J. & Deaton, A. (1989), 'Why is Consumption So Smooth?', *The Review of Economic Studies* **56**(3), 357–373.

Carroll, C. (1997), 'Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis', *The Quarterly Journal of Economics* **112**(1), 1–55.

Carroll, C. (2001), 'Precautionary Saving and the Marginal Propensity to Consume out of Permanent Income'.

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

Castaneda, A., Diaz-Gimenez, J. & Rios-Rull, J. (2003), 'Accounting for the US Earnings and Wealth Inequality', *Journal of Political Economy* **111**(4), 818–857.

Chamberlain, G. & Wilson, C. (2000), 'Optimal Intertemporal Consumption under Uncertainty', *Review of Economic Dynamics* **3**(3), 365–395.

Cochrane, J. (1991), 'A Simple Test of Consumption Insurance', *The Journal of Political Economy* **99**(5), 957–976.

Conesa, J., Kitao, S. & Krueger, D. (2007), 'Taxing Capital?: Not a Bad Idea After All!', *Mimeo* .

Cunha, F., Heckman, J. & Navarro, S. (2005), 'Separating uncertainty from heterogeneity in life cycle earnings', *Oxford Economic Papers* **57**(2), 191–261.

De Nardi, M. (2004), 'Wealth Inequality and Intergenerational Links', *Review of Economic Studies* **71**(3), 743–68.

Deaton, A. & Paxson, C. (1994), 'Intertemporal Choice and Inequality', *The Journal of Political Economy* **102**(3), 437–467.

Domeij, D. & Heathcote, J. (2004), 'On The Distributional Effects Of Reducing Capital Taxes', *International Economic Review* **45**(2), 523–554.

Dynarski, S. & Gruber, J. (1997), 'Can Families Smooth Variable Earnings?', *Brookings Papers on Economic Activity* .

Erosa, A. & Gervais, M. (2002), 'Optimal Taxation in Life-Cycle Economies', *Journal of Economic Theory* **105**(2), 338–369.

Gouveia, M. & Strauss, R. (1994), 'Effective Federal Individual Income Tax Functions: An Exploratory Empirical Analysis', *National Tax Journal* **47**(2), 317–39.

Guvenen, F. (2007), 'Learning your earning: Are labor income shocks really very persistent?', *American Economic Review* **97**(3), 687–712.

Hall, R. (1978), 'Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence', *The Journal of Political Economy* **86**(6), 971–987.

Hall, R. & Mishkin, F. (1982), 'The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households', *Econometrica* **50**(2), 461–481.

Heathcote, J., Storesletten, K. & Violante, G. (2004), 'The Macroeconomic Implications of Rising Wage Inequality in the United States', *CEPR wp* **4296**.

Heathcote, J., Storesletten, K. & Violante, G. (2005), 'Two Views of Inequality Over the Life-Cycle', *Journal of the European Economic Association* **3**(2-3), 765–775.

Heathcote, J., Storesletten, K. & Violante, G. (2007), 'Consumption and Labor Supply with Partial Insurance: An Analytical Framework', *Mimeo, NYU* .

Heaton, J. & Lucas, D. (1996), 'Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing', *Journal of Political Economy* **104**, 668–712.

Hubbard, R., Skinner, J. & Zeldes, S. (1995), 'Precautionary Saving and Social Insurance', *The Journal of Political Economy* **103**(2), 360–399.

Huggett, M. (1993), 'The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies', *Journal of Economic Dynamics and Control* **17**(5/6), 953–970.

Huggett, M. (1996), 'Wealth distribution in life-cycle economies', *Journal of Monetary Economics* **38**(3), 469–494.

Huggett, M., Ventura, G. & Yaron, A. (2006), 'Human capital and earnings distribution dynamics', *Journal of Monetary Economics* **53**(2), 265–290.

Imrohoroglu, A. (1989), 'Cost of Business Cycles with Indivisibilities and Liquidity Constraints', *The Journal of Political Economy* **97**(6), 1364–1383.

Kaplan, G. (2007*a*), 'Inequality and the Lifecycle', *Mimeo, NYU* .

Kaplan, G. (2007*b*), 'Moving Back Home: Insurance Against Labor Market Risk', *Mimeo, NYU* .

Krueger, D. & Perri, F. (2006), 'Does Income Inequality Lead to Consumption Inequality? Evidence and Theory', *Review of Economic Studies* **73**(1), 163–193.

Krusell, P., Quadrini, V. & Rios-Rull, J. (1997), 'Politico-economic equilibrium and economic growth', *Journal of Economic Dynamics and Control* **21**(1), 243–272.

Krusell, P. & Smith, A. (1998), 'Income and Wealth Heterogeneity in the Macroeconomy', *Journal of Political Economy* **106**(5), 867.

Krusell, P. & Smith, A. (1999), 'On the Welfare Effects of Eliminating Business Cycles', *Review of Economic Dynamics* **2**(1), 245–272.

Ljungqvist, L. & Sargent, T. (2004), *Recursive Macroeconomic Theory*, MIT Press.

Nishiyama, S. & Smetters, K. (2005), 'Consumption Taxes and Economic Efficiency with Idiosyncratic Wage Shocks', *Journal of Political Economy* **113**(5), 1088–1115.

Primiceri, G. & van Rens, T. (2006), 'Heterogeneous Life-Cycle Profiles, Income Risk and Consumption Inequality', *CEPR DP* **5881**.

Quadrini, V. (2000), 'Entrepreneurship, Saving, and Social Mobility', *Review of Economic Dynamics* **3**(1), 1–40.

Rios-Rull, J. (1995), 'Models with Heterogeneous Agents', *Frontiers of Business Cycle Research* pp. 98–125.

Schechtman, J. & Escudero, V. (1977), 'Some results on 'An income fluctuation problem'.', *Journal of Economic Theory* **16**(2), 151–66.

Storesletten, K., Telmer, C. & Yaron, A. (2004), 'Consumption and risk sharing over the life cycle', *Journal of Monetary Economics* **51**(3), 609–633.

Storesletten, K., Telmer, C. & Yaron, A. (2007), 'Asset pricing with idiosyncratic risk and overlapping generations', *Review of Economic Dynamics* **10**(4), 519–548.

Telmer, C. (1993), 'Asset-Pricing Puzzles and Incomplete Markets', *The Journal of Finance* **48**(5), 1803–1832.

Ventura, G. (1999), 'Flat tax reform: A quantitative exploration', *Journal of Economic Dynamics and Control* **23**(9-10), 1425–1458.

Yaari, M. (1976), 'A Law of Large Numbers in the Theory of Consumer's Choice Under Uncertainty', *Journal of Economic Theory* **12**(2), 202–217.