# 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

This Draft: February 2008

*PRELIMINARY AND INCOMPLETE*

**Abstract**
The standard life-cycle incomplete-markets model, where households face idiosyncratic income 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 than what is measured in the data. BPP find that 38% of permanent income shocks are insurable (i.e., do not translate into consumption growth), while in the model this insurance coefficient is 18%. Moreover, the life-cycle pattern of insurance coefficients is sharply increasing and convex in the model, while it is roughly flat in the data. Taken at face value, this result would suggest that macroeconomists should develop models with more avenues of insurance than a risk-free asset. However, we also find that if income shocks are not not modelled to be permanent, as assumed by BPP, but (even very) persistent, the degree of consumption smoothing implicit in Bewley models largely agrees with the data. Finally, we uncover that the BPP estimator for the insurance coefficient has, in general, a small downward bias. This bias though can be substantial for households near their borrowing limit, e.g., young or poor households.

# 1 Introduction

In the last fifteen years, 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 income fluctuations. The only asset they can trade in order to smooth consumption is a non-state contingent security. A borrowing constraint limits the amount of debt they can cumulate. 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 the stationary version of this framework. Krusell and Smith (1998) developed computational tools to compute equilibria in the presence of aggregate shocks. The term "Bewley models" was first used in the Ljungqvist and 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 et al. (2003); De Nardi (2004)), explaining the equity premium (Telmer (1993); Heaton and Lucas (1996); Storesletten et al. (2007)), evaluating the distributional effects of fiscal policy (Krusell et al. (1997); Ventura (1999); Domeij and Heathcote (2004)), measuring the welfare costs of business cycles (Krusell and Smith (1999)), and designing optimal fiscal policy (Erosa and Gervais (2002); Nishiyama and Smetters (2005); Conesa et al. (2007)), only to cite a few.

The objective of this paper is to quantify the amount of consumption insurance allowed to households in Bewley models and to compare it to empirical estimates for the US economy.

[1] See Yaari (1976), Schechtman and Escudero (1977) and Chamberlain and 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.

The questions we address are: (i) How much consumption insurance is there in Bewley models? and (ii) Is this high or low relative to that measured in the data?

The degree of consumption insurance is an important property of models with heterogeneous agents and is a key dimension of the data on which to assess the applicability of a particular model. Policy evaluation and optimal policy design are not independent of the insurance vehicles that households have access to, nor of 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 latter margin depends crucially on how much smoothing agents can do on their own.

Our metric for consumption insurance is the fraction of permanent and transitory idiosyncratic income 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. A third useful benchmark is the strict version of the permanent income hypothesis (PIH) where the insurance coefficient is close to one for transitory shocks and is exactly zero for permanent shocks (Hall (1978)).[2]

Our exercise is made possible by the fact that we now have reliable empirical measures of these insurance coefficients for permanent and transitory shocks. In a pair of recent papers, BPP have developed the data set and the empirical methodology to do just that.

In BPP (2004), they have constructed a new panel data set for the US with individual 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 based on the realized food consumption measure in PSID.

In BPP (2006), the authors use this new data set to estimate insurance coefficients for permanent and transitory shocks separately. We return to the details of the econometric procedure later. It suffices to say here that they assume an income process that is the

[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.

sum of a random walk and a transitory (i.i.d.) component. They find that 38% (53%) of permanent shocks and 95% (96%) of transitory shocks to disposable (pre-government) income are insured, i.e. they do not co-vary with consumption growth.

We repeat the same exercise done in BPP on an artificial panel data simulated from our 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 a pre-specified limit. They also save for life-cycle and precautionary reasons, thus their wealth helps in absorbing income shocks. One should not forget 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.

Our findings can be summarized as follows. First, the insurance coefficient of transitory shocks to after-tax labor income is 92% in the model and 95% in the data, whereas the insurance coefficient of permanent shocks is 18% in the model and 38% in the data. The model-implied coefficients are found to be very robust in a series of sensitivity analyses. Insurance coefficients in the model exceed 30% only for values of risk-aversion above ten.

Second, the life-cycle pattern of the BPP estimates of insurance coefficients for permanent shocks appear roughly flat, if not slightly decreasing. This empirical pattern is strikingly at odds with what is predicted by the model: a sharply increasing and convex pattern of insurance coefficients over the life-cycle of the agent.

Third, in general the BPP methodology works very well for transitory shocks, but it tends to underestimate true insurance coefficients for permanent shocks. For example, in the benchmark model the BPP estimate of the latter coefficient is 14% against a true value of 18%. 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 only holds approximately in their model. Quantitatively, on average the bias is not large (factor of 1.3), although systematically present. However, when households are near their borrowing constraint, this bias becomes very severe (factor of 2.5). Thus if borrowing constraints were prevalent in the US economy, one would conclude that, once we correct for this bias, true insurance coefficients are even higher than estimated, so model and data are even more distant.

Finally, we find that a different specification of the income process 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 ($\rho$) above 0.90. Next, we document that for $\rho = 0.95$ the BPP estimate of the insurance coefficient for persistent shocks in the Bewley model is around 34%, i.e. roughly its empirical value. However, the insurance coefficient for transitory shocks declines to 84% relative to 95% in the data.

The rest of the paper is organized as follows. Section 2 introduces a framework for measuring insurance and describes the BPP methodology. Section 3 outlines the version of the Bewley model we use for our experiments, and it describes the parameterization. Section 4 contains the results of the benchmark economy and of a series of sensitivity analyses, including the one where we tighten borrowing constraints. Section 5 analyzes the robustness of our conclusion to the degree of persistence of income shocks. In Section 6 we discuss our results in relation to the existing literature. Section 7 concludes the paper.

# 2 A framework for measuring insurance

## 2.1 Insurance coefficients

**An income process** Consider an economy where after-tax log-income $y_{it}$ 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 $\sigma^2 = (\sigma_1^2, ..., \sigma_m^2)'$ 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.

For much of the paper we will focus on a particular special case 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:

$$y_{it} = z_{it} + \varepsilon_{it}, \tag{2}$$

where $z_{it}$ follows a unit root process with innovation $\eta_{it}$, and where $\varepsilon_{it}$ is an uncorrelated income shock.[3] Let $\sigma_\eta$ and $\sigma_\varepsilon$ denote the variances of the two innovations. The individual shocks $\eta_{it}$ and $\varepsilon_{it}$ are i.i.d.. across the population. It follows that income growth can be written as

$$\Delta y_{it} = \eta_{it} + \Delta\varepsilon_{it}. \tag{3}$$

For most of the paper, we focus on this particular income process because this is the one chosen by BPP. It is also a very common income process, at least since the work by Abowd and Card (1989), who showed that this specification is parsimonious and yet fits income data particularly well. Nonetheless, we define insurance coefficients for the more general case and in Section 5, we verify the robustness of our results to more general specifications of the income process.

**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

$$\phi^x = 1 - \frac{cov\left(\Delta c_{it}, x_{it}\right)}{var\left(x_{it}\right)}, \tag{4}$$

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 $\phi_t^x$) where the variance and covariance are taken conditionally on households of age $t$. The insurance coefficient in (4) 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 (4) by simulation, since the shocks are observable in the model. However, estimating (4) 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]

[3] 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 $\sigma_{z_0}$. This generalization has no bearing on the methodology used to measure insurance coefficients, since it is based on growth rates, and hence the initial fixed effect is differenced out.

[4] Note that it is not sufficient to identify the variances of the different shocks, i.e the vector $\sigma^2$. Rather, the realizations of the shocks must be identified, household by household. With an infinite history of data, these may be identified using filtering techniques, however the short duration of available panel data sets make this type of technique unreliable.

Suppose we have panel data on households' income and consumption. Let $\mathbf{y}$ be the vector of income realizations at all ages $t = 2, ..., T$ and let $g_t^x(\mathbf{y})$ 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

$$\begin{aligned} var(x_{it}) &= cov(\Delta y_{it}, g_t^x(\mathbf{y})), \\ cov(\Delta c_{it}, x_{it}) &= cov(\Delta c_{it}, g_t^x(\mathbf{y})), \end{aligned} \tag{5}$$

and then constructing an estimator for $\phi^x$ as

$$\hat{\phi}^x = 1 - \frac{cov(\Delta c_{it}, g_t^x(\mathbf{y}))}{cov(\Delta y_{it}, g_t^x(\mathbf{y}))}. \tag{6}$$

While verifying the first condition 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. To gain intuition, note that $1 - \hat{\phi}^x$ is simply the coefficient from a cross-sectional regression of consumption changes on individual income shocks. If $g_t(\mathbf{y})$ satisfies the conditions in (5) then the resulting expression for $1 - \hat{\phi}^x$ is equivalent to the coefficient from an instrumental variables regression of consumption changes on income changes, using $g_t(\mathbf{y})$ as an instrument.

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

**The BPP methodology** BPP assume that the following orthogonality conditions hold for the true consumption allocation:

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

Under this assumption, BPP propose a strategy to identify and estimate the insurance coefficients.[5] For the transitory shock, they set $g_t^{\varepsilon}(\mathbf{y}) = \Delta y_{i,t+1}$ and note that

$$cov(\Delta y_{it}, \Delta y_{i,t+1}) = -var(\varepsilon_{it}), \tag{8}$$

$$cov(\Delta c_{it}, \Delta y_{i,t+1}) = -cov(\Delta c_{it}, \varepsilon_{it}),$$

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

$$cov(\Delta y_{it}, \Delta y_{i,t-1} + \Delta y_{it} + \Delta y_{i,t+1}) = var(\eta_{it}), \tag{9}$$

$$cov(\Delta c_{it}, \Delta y_{i,t-1} + \Delta y_{it} + \Delta y_{i,t+1}) = cov(\Delta c_{it}, \eta_{it}).$$

Combining (4) with (8) and (9) confirms that these instruments do in fact correctly identify the insurance coefficients $(\phi^{\eta}, \phi^{\varepsilon})$. Note that the orthogonality conditions in (7) are required only for the identification of the insurance coefficients for permanent shocks, but they are not needed for transitory shocks.

Straightforward application of a minimum distance algorithm allows the estimation of the cross-sectional moments in (8) and (9). 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.

[5] 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 $\alpha_{it}^{\eta}$ and $\alpha_{it}^{\varepsilon}$ are marginal propensity to consume out of permanent and transitory shocks. The term $\Psi_{it}$ 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) is enough.

## 2.2 Generality of the BPP approach

The obvious question, at this point, is: how general is assumption (7)? We would like it to be as general as possible. The more general it is, the more likely it is that the true model (and hence the data) satisfy it, which means that the BPP estimates indeed have the structural interpretation of unbiased insurance coefficients. Below, we list a number of economies displaying equilibrium allocations that satisfy this property, sometimes exactly, sometimes only approximately.

**Complete markets and autarky** Complete markets economies feature full insurance and satisfy (7). Under complete markets, idiosyncratic shocks do not affect consumption, hence $cov(\Delta c_{it}, x_{it}) = 0$ and $\phi^x = 1$. In autarky, $\Delta c_{it} = \Delta y_{it}$. Therefore, the autarkic allocation satisfies condition (7). Under autarky, $cov(\Delta c_{it}, x_{it}) = var(x_{it})$ and $\phi^x = 0$. Note that in these two extreme cases the value of $\phi^x$ does not depend on the form of the income process (in particular the degree of persistence of shocks $x$).

**Permanent Income Hypothesis (PIH)** The textbook 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 consumption growth (in levels), when combined with the analogous income process to (2), specified in levels:

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

where $\chi_t = \frac{r}{(1+r)} \frac{1}{1-(1+r)^{-(T-t+1)}}$.[6] 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 has an upper bound of $(1+r)^{-1}$ and decreases monotonically towards zero as the end of life becomes nearer ($t \to T$).

**A Bewley economy** This is the economy we are mainly interested in, as in the rest of the paper we focus on this model. BPP prove that in a Bewley economy with loose borrowing constraints (i.e., which do not bind) and CRRA utility, consumption growth can be approximated as

$$\Delta c_{it} \simeq \pi_{it} \eta_{it} + \left( \frac{r}{1+r} \right) \pi_{it} \varepsilon_{it}, \tag{10}$$

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

where $r$ is the interest rate and $\pi_{it}$ is the share of human wealth (discounted expected future labor income) over total wealth, i.e. human and financial wealth, for individual $i$ at age $t$. In the infinite horizon case, $\pi_{it} \simeq 1$.

The argument developed by BPP has two potential shortcomings. First, the abstraction from binding borrowing limits. Second, the fact that $\pi_{it}$ is individual specific can be problematic. Since $\pi_{it}$ in general depends on the current and past idiosyncratic income shocks, (7) is violated. In Section 4.5, we show that in the presence of liquidity constraints the approximation is particularly poor and the BPP approach can significantly understate the true insurance coefficients.

**A partial insurance economy** Heathcote et al. (2007) (HSV thereafter) outline an economy with partial insurance against permanent shocks and full insurance against transitory shocks where the BPP methodology correctly identifies the insurance coefficients. In the HSV economy,

$$\begin{aligned} \Delta y_{it} &= \tilde{\eta}_{it} + \omega_{it} + \Delta\varepsilon_{it}, \\ \Delta c_{it} &= \eta_{it} \end{aligned}$$

where $\tilde{\eta}_{it}$ and $\omega_{it}$ are, respectively, an uninsurable (i.e., it translates one for one into consumption growth) and an insurable (i.e., it does not affect consumption growth) permanent income shock. To make the income process comparable with (2), we define the combined permanent shock $\eta_{it} \equiv \tilde{\eta}_{it} + \omega_{it}$, whose age-varying variance is $\sigma^2_{\eta t} = \sigma^2_{\tilde{\eta} t} + \sigma^2_{\tilde{\omega} t}$. It is easy to verify that the BPP approach would correctly estimate $\phi^{\varepsilon}_t = 1$ and $\phi^{\eta}_t = \frac{\sigma^2_{\tilde{\eta} t}}{\sigma^2_{\tilde{\eta} t} + \sigma^2_{\omega t}}$.

**A private information economy** Attanasio and Pavoni (2007) develop a model where households have private information on their permanent and transitory income shocks and can hide savings from a social planner who chooses household consumption and production effort. Under specific functional forms, they can obtain closed-form solutions for the consumption allocation. Let $e_{it}$ denote individual effort. When the household production function is

$$y_{it} = \begin{cases} \eta_{it} + \varepsilon_{it} + a^{\eta} + a^{\varepsilon}, & \text{if } e_{it} < 0 \\ \eta_{it} + \varepsilon_{it} + b^{\eta} + b^{\varepsilon}, & \text{if } e_{it} \geq 0 \end{cases}$$

with $a^x \geq 1 > b^x$, and preferences are given by $u\left(c_{it}, e_{it}\right) = \exp\left(\left(1-\gamma\right) c_{it} - e_{it}\right)/1-\gamma$, then the constrained-efficient consumption growth is

$$\Delta c_{it} = \alpha^0_t + \frac{1}{a^{\eta}}\eta_{it} + \frac{r}{1+r}\frac{1}{a^{\varepsilon}}\varepsilon_{it}.$$

Once again, the BPP assumption is satisfied. The parameters $(a^{\eta}, a^{\varepsilon})$ measure the severity of the private information problem and, therefore, the degree of smoothing possible in the economy. The case $a^{\eta} = a^{\varepsilon} = 1$ corresponds to the infinite-horizon version of the Bewley economy $(\pi_{it} \simeq 1)$, while $a^{\eta} = a^{\varepsilon} = \infty$, corresponds to complete markets.

**A limited enforcement economy** *To be completed.*

These examples demonstrate that, in a wide variety of economic environments it is possible to justify consumption allocations that are, at least approximately, consistent with (7). This conclusion gives us some hope that the BPP estimates may capture correctly the empirical value of insurance coefficients. Strictly speaking, for our purposes, all we need is assumption (7) to hold in the Bewley model. Even though we argued above that, qualitatively, (7) does not hold, it is possible that quantitatively the bias is small.

In the rest of the paper, we outline and calibrate a life-cycle Bewley economy. We simulate a panel of artificial data on household income and consumption. Next, we 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. Since in the model we can compute both the exact insurance coefficients and the BPP estimates, we will also learn about the reliability of the BPP methodology.

To conclude, a note on language. We have defined $\phi^x$ as *insurance* coefficients. In Bewley models (or in the PIH) income shocks do not transmit one for one into consumption because consumers have access to non state-contingent borrowing and saving. And thus 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 use the term insurance in this paper for three reasons. First, the definition of $\phi^x$ 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 interpersonal transfer. Third, even within Bewley models, there may be important insurance mechanism over and above self-insurance, such as the government (that we will model) and the family.

# 3 A Bewley model

The model economy is the life-cycle version of the standard incomplete markets model introduced by Huggett (1993). There is no aggregate uncertainty and we abstract from modelling the production side of the economy. The economy is populated with a continuum of households, indexed by $i$, each of whom live up to age $T = T^{work} + T^{ret}$.

The unconditional probability of surviving to age $t$ is denoted by $\xi_t$. We assume that $\xi_t = 1$ for the first $T^{work}$ periods so that there is no chance of dying before retirement, which occurs at age $T^{ret}$. After retirement, $\xi_t < 1$ and all agents die by age $T$, so that $\xi_T = 0$.

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

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

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

$$y_{it} \equiv \log Y_{it} = \kappa_t + z_{it} + \varepsilon_{it},$$

where $\kappa_t$ is a deterministic experience profile that is common across all households; $z_{it}$ is a permanent component and $\varepsilon_{it}$ 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 $\sigma_{z_0}$. The shocks $\varepsilon_{it}$ and $\eta_{it}$ are zero-mean normally distributed with variances $\sigma_\varepsilon$ and $\sigma_\eta$, 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 that households use to insure themselves against mortality 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

$$C_{it} + A_{i,t+1} = (1+r)\,A_{it} + Y_{it}, \qquad \text{if } t < T^{ret}$$

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

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. We introduce a consumption floor $\underline{C}$ as a reduced form way of capturing the effects of means-tested programs such as Food Stamps and Temporary Assistance for Needy Families. In this case, the budget constraint reads:

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

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

where

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

is the government transfer guaranteeing a minimum level of consumption for every household, as modelled, for example, by Hubbard et al. (1995).

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.

## 3.1 Calibration

We calibrate the model parameters to reproduce certain key features of the US economy.

**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. 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** In our benchmark calibration, 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 $\gamma$ in the range $[1, 15]$.

**Discount factor and interest rate** For a given interest rate $r$, the discount factor $\beta$ uniquely determines the level of wealth that is accumulated by households. 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 thus set the interest rate exogenously at 3% and choose $\beta$ 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 interest rates of 2% and 4% and 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 of the bottom 99% among US households (excluding housing and the top 1%), and may be considered a lower bound on the aggregate wealth-income ratio over the 1980-1992 period.[7]

**Income process** For the benchmark income process, three parameters are required. These are the variance of the two shocks, $\sigma_\varepsilon$ and $\sigma_\eta$, and the cross-sectional variance of the initial value of the permanent component $\sigma_{z_0}$. 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

[7] 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.

these values.

**Initial wealth** In the benchmark calibration we assume that all households start life with zero wealth, $A_{i0} = 0$. We also consider an environment in which initial wealth levels are drawn from a distribution consistent with the financial wealth of young households (aged 20-30) in the Survey of Consumer Finances.[8]

**Borrowing limit** We consider two extreme 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. This case is closest to the assumptions imposed by BPP in deriving the approximation for optimal consumption growth. We also study the insurance possibilities when the other extreme of no borrowing, $\underline{A}= 0$, is imposed.[9]

**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 0.38 and 1.59 times average earnings, which is what they were in 2007 for the US system. 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 11.3% in 2006.

**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 and 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 hence given by

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

[8]The initial wealth distribution is implemented as a discrete distribution on financial wealth / labor earnings ratios, with 75 equally spaced points between $-2.375$ and 19.375. We assume that the initial earnings draw is independent of the initial draw of this ratio. This is consistent with an empirical correlation of 0.02 in the SCF.

[9]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.

The value for $\tau^{SS}$ is set at 0.124 which is the combined employer and employee rate for social security contribution. The values for $\tau^b$ and $\tau^\rho$ are taken from Gouveia and Strauss (1994). They are set at $\tau^b = 0.258$ and $\tau^\rho = 0.768$. The value for $\tau^s$ 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 15.4%.

**Consumption floor** Following Hubbard et al. (1995), we set $\underline{C}$ to $\$4,000$. We also report results for a consumption floor of $\$8,000$.

# 4 Results

## 4.1 Insurance coefficients in the data and the model

We start by reporting estimated insurance coefficients in the data and the benchmark model for permanent and transitory shocks. These are shown in Table 1. In all tables, columns labelled "BPP" refer to insurance coefficients calculated using the instrumental variables approach described in Section 2.1, while columns labelled "TRUE" refer to insurance coefficients calculated directly from the realizations of the individual shocks.

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 0.92 for transitory shocks. These figures compare to insurance coefficients of, respectively, 0.38 and 0.94 in the data on US households during 1980-1992. 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, both in the model and the data.

However, against permanent shocks, the US tax and transfer system appears to provide a great deal of insurance: moving from disposable income to pre-tax labor earnings increases the empirical BPP estimates from 0.38 to 0.53. The same is true in the model, where the insurance coefficient increases from 0.14 to $0.25 - 0.27$, depending on the assumed level of the consumption floor provided through government safety nets.

Thus, whether one focuses on pre- or post-government measures of income, the Bewley model generates less insurance against permanent shocks, and roughly the same level of insurance against transitory shocks, relative to the data.

## 4.2 Accuracy of the BPP methodology

Our approach also allows us to 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".

Table 1 reveals that whereas the methodology works extremely well for transitory shocks, it tends to underestimate the amount of insurance for permanent shocks. This suggests that the true insurance coefficients for permanent shocks may be higher than 0.38 (for disposable labor income) and 0.53 (for pre-tax earnings). This downward bias in estimates of insurance coefficients for permanent shocks is a robust feature of the BPP methodology and, as we show below, it is exacerbated whenever wealth is close to the borrowing limit. Thus, it can be particularly severe for young households and for households with low wealth holding.

The reason for the large bias in $\phi_t^\eta$ is that the orthogonality conditions in (7) may fail when agents have low wealth holdings. Recall that these orthogonality conditions are only required for the identification of insurance coefficients for permanent shocks, and not for transitory shocks. 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 for a household close to the borrowing limit, consider a household who receives a negative transitory shock at $t-2$ ($\varepsilon_{t-2} < 0$). Such a household would like to borrow (or dissave) to smooth the effect of the negative shock. However for a household that is close to its borrowing limit, even a small reduction in wealth can have a large expected utility cost because of the possibility of being constrained in a future period. Thus smoothing expected marginal utility entails an optimal drop in consumption at $t-2$ - the closer they are to the borrowing constraint, the larger this drop.[10] This leads to a positive expected change in consumption at $t-1$, as consumption slowly returns to its baseline level.

[10] In a model with binding borrowing constraints, a household that is already at the constraint will incorporate the full size of the shock into consumption at $t-2$.

The larger the shock at $t-2$, the bigger this reversion and hence $cov\left(\Delta c_{t-1}, \varepsilon_{t-2}\right) < 0$. 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.

## 4.3 Insurance coefficients by age

In Figure 1 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 (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 1 reveals a vastly different scenario with regard to insurance in the Bewley model.

The insurance coefficients for permanent shocks are mildly decreasing at young ages, but are increasing steadily after age 35 and are markedly convex. The increase in the amount of insurance as households approach retirement is due to two features. First, as the time horizon decreases, permanent shocks become "less permanent" and begin to look more and more like transitory shocks. Hence accumulated wealth can be used to mitigate their impact on consumption. Second, older households have had time to accumulate a larger stock of wealth, and hence can absorb larger income shocks without changing their consumption too much.

The reason for the decline in $\phi_t^\eta$ over the first ten years and its taking negative values is subtle. In the model, agents face a deterministic age-profile for earnings which is increasing and concave. Hence, there is a strong incentive to borrow early in life, purely for life-cycle reasons. With loose borrowing constraints, this means that the average wealth levels are decreasing and negative at young ages. Insurance coefficients for permanent shocks are, in general, proportional to wealth.

A negative insurance coefficient is obtained when $cov\left(\Delta c_{it}, x_{it}\right) > var\left(x_{it}\right)$, which is the case when consumption responds more than one-for-one to a particular shock, e.g., $\Delta c_{it} > x_{it}$ in case of a positive shock. The reason that this may happen is due to the interaction of transitory shocks and permanent shocks in the model, as uncovered by Carroll (1997).[11] With $\sigma_\varepsilon^2 > 0$, households will accumulate a target level of wealth which they use to buffer

[11] We have verified that when we simulate the model without transitory shocks $\left(\sigma_\varepsilon^2 = 0\right)$, then $\phi_t^\eta$ is always positive.

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. So 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.

Finally, note that the insurance coefficients for transitory shocks 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.

## 4.4 Insurance coefficients by wealth and education

For the US data, BPP find that estimated insurance coefficients differ markedly across various subgroups in the population. In this section we investigate whether the low insurance coefficient for permanent shocks in the model, compared with the data, is being driven by a particular group of individuals, or whether the Bewley model generates too little consumption smoothing across the entire population.

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 first two rows of Table 2. Here we see that 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: $-0.01$ vs 0.04 for low wealth households, and 0.22 vs 0.44 for high wealth households. In terms of economic significance, perhaps one could conclude that the group where insurance against permanent shocks is most underestimated in the Bewley model is for rich households, since the poor de facto cannot smooth permanent shocks both in the data and the model.

Note also how severe is the bias in the BPP methodology for low wealth households—the insurance coefficient is enormously understated. This result is consistent with our intuition on the sources of bias discussed above.

With regard to transitory shocks, Table 2 reveals that in the model, households with all wealth levels are equally very well insured, while in the data there appears to be far less insurance for low-wealth households (0.72 in the data, 0.91 in the model). This could be 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.

BPP also report insurance coefficients for household with different education levels. We interpret education differences as differences in lifetime income. Hence 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, $z_0$. As evident from the third and fourth rows of Table 2, when we group households this way, we find little differences in insurance coefficients.

Why does $z_0$ not affect insurance coefficients? Three features of the model contribute to this outcome: (i) $z_t$ is permanent, so a higher $z_0$ means that income is increased now and in all future states by the same factor; (ii) shocks are multiplicative, so their variance is fixed as a fraction of average income; and (iii) preferences are CRRA. Together these imply that the optimal wealth level is not affected, only the *level* of consumption is effected, but consumption still responds in the same way to future shocks. Once again, the loose borrowing constraints are important. When the borrowing limit is exogenous, non-zero and is not proportional to income, it will bind more often for low income households.

## 4.5 Sensitivity analysis on the benchmark economy

Table 3 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 $\gamma$ rises insurance coefficients for permanent shocks increase. However, only for values of $\gamma$ beyond fifteen we reach insurance coefficients close to those estimated in the data.[12]

Clearly, smaller wealth income ratios (obtained by lowering $\beta$) map into smaller asset

[12] 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.

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 larger discrepancy with the data. Fixing the wealth income ratio at 3.5 and changing the interest rate has virtually no effect. Similarly, it appears that the amount of insurance in the model does not depend on the size of the shocks, when the latter is changed within a plausible range of values.

We also calculated insurance coefficients for the corresponding economy with quadratic preferences in order to assess the importance of the precautionary motive in accounting for the amount of insurance, holding constant the average level of risk aversion. With quadratic preferences $\phi^{\eta}$ decreases from 0.18 to 0.13, suggesting that precautionary savings account for just under one third of the amount of insurance in the benchmark model. Finally, Table 3 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 main conclusions. First, under a permanent-transitory income process, the Bewley model allows less insurance against permanent shocks than what is estimated in the data, and about the right amount against transitory shocks. Second, BPP estimates of insurance for permanent shocks are always downwards biased.

### 4.6 A Bewley economy with tight borrowing constraint

We now contrast the benchmark economy with one where agents are not allowed to borrow, i.e., $\underline{A} = 0$. We keep all parameter values unchanged, except for $\beta$, which is set to generate a wealth-income ratio of 3.5. The results of this experiment are summarized in Tables 4 and 5.

The true insurance coefficient against permanent shocks is 0.20 just above the baseline value of 0.18. The reason is that this economy has higher precautionary saving, since agents want to avoid hitting the borrowing constraint, thus slightly more self-insurance. Interestingly, the insurance coefficients against transitory shocks is significantly lower: 0.81 compared to a baseline value of 0.92. In the baseline model, the loose borrowing limit was very helpful in smoothing shocks, especially at young ages.

Finally, as anticipated, the BPP methodology severely underestimates $\phi^{\eta}$, by a factor of 2.5. The reason is that in this economy there is substantial bunching at values of wealth exactly at or near the borrowing limit.

Figure 2 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, indeed, 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 38 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 often little wealth and, as such, are unable to smooth transitory shocks.

In Table 5 we report insurance coefficients by wealth and education for the economy with tight borrowing constraints. The only substantial difference between these results and those in the benchmark economy is the insurability of transitory shocks for households with low wealth. The presence of borrowing constraints makes it difficult to smooth transitory shocks unless a substantial amount of precautionary wealth has already been accumulated.

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

## 4.7 The distribution of insurability

Our estimate of $\phi^x$ is obtained from cross-sectional variances and covariances. By construction, it is thus an *average* measure of the overall insurability of shocks in an economy. There may exist substantial heterogeneity in the extent to which different households are able to smooth the effects of income shocks, both in the data and in a model. A value of $\phi^\eta = \frac{1}{2}$ would be obtained for an economy inhabited by two types of households of equal measure, one group with access to complete markets, the other in autarky. But $\phi^\eta = \frac{1}{2}$ is also consistent with an economy inhabited by identical households in which there are two types of permanent shocks of equal variance, one insurable the other not, as in Heathcote et al. (2007). Although these two economies have the same aggregate insurance coefficients, the distribution of risk-sharing in the economy is very different.

Similarly to our definition for the average insurance coefficient (4), define the *individual insurance coefficient* for household $i$ at time $t$ as:

$$\phi_{it}^{x} = 1 - \frac{x_{it}\Delta c_{it}}{var\left(x_{it}\right)} \tag{11}$$

Cross-sectional heterogeneity in $\phi_{it}^{x}$ reflects the effects that heterogeneity in exogenous and endogenous state variables (such as accumulated wealth, and the realizations of income shocks) have on the transmission of income shocks to consumption fluctuations. The coefficient $\phi_{it}^{x}$ has the attractive property that its cross-sectional mean, conditional on age $t$, is simply the average insurance coefficient, $\phi_{t}^{x}$. One could thus interpret the results in the previous sections as an investigation of the properties of $E_t\left[\phi_{it}^{x}\right]$ in the data and the model.

Our first summary statistic for inequality in insurability is the conditional standard deviation of $\phi_{it}^{x}$. This is plotted in Figure 4. For both types of shocks inequality in insurability decreases with age. There is also substantially more heterogeneity in insurability against permanent shocks than against transitory shocks. Luck in the size of the income shocks received, especially permanent ones, matters much more when the household has a small stock of wealth to smooth consumption. At age 25 the standard deviation of $\phi_{it}^{\eta}$ is 1.4, compared with a mean of 0.18, while the standard deviation of $\phi_{it}^{\varepsilon}$ is 0.6, compared with a mean of 0.95. For both types of shocks, the standard deviation halves between age 25 and age 55.

An alternative measure of dispersion in insurability is one that captures the average extent to which shocks translate into consumption growth over the course of the life-cycle. We define the *lifetime individual insurance coefficient* as

$$\overline{\phi_{i}^{x}} = \frac{1}{T^{work}} \sum_{t=1}^{T^{work}} \phi_{it}^{x}. \tag{12}$$

One benefit of using the distribution of $\overline{\phi_{i}^{x}}$ as a measure of inequality in insurability against permanent shocks, over the distribution of $\phi_{it}^{x}$, is that the former distribution is a better indicator of how inequality in the ability to smooth consumption across households will translate into cross-sectional dispersion in household welfare.

Figure 5 shows a histogram of lifetime insurance coefficients for permanent and transitory shocks, calculated from simulated data for the benchmark model. Both histograms are plotted on the same scale. The distributions indicate substantial amounts of heterogeneity, especially for permanent shocks: the distribution of $\overline{\phi_{i}^{\eta}}$ is negatively skewed with a standard deviation of 0.21.[13]

From our previous discussion, it is clear that households with low values of $\overline{\phi_i^{\eta}}$ in the distribution of Figure 5 are those who spend a large portion of their life close to the borrowing limit, while households with high values of $\overline{\phi_i^{\eta}}$ are those who are lucky enough to receive large positive shocks early in life.

# 5 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 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 $a_j = (\rho^j, 0)' \ \forall \ j \geq 0$. 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 $\rho$ is available. This is a reasonable assumption, since $\rho$ can be identified using panel data on income alone, and thus can be estimated in a first stage.

Define the "twisted" growth of 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 $g_t^{\eta}(\mathbf{y}) = \rho\tilde{\Delta} y_t + \tilde{\Delta} y_{t+1} + \rho\tilde{\Delta} y$.[14] For the transitory shock, we have

[13] Of course, the main benefits to measuring insurability heterogeneity in the model accrue when there exist reliable estimates of the corresponding measures in the data. We plan to address the issue of nonparametric identification and estimation of the distribution of $\phi_{it}^x$ from the joint distribution of consumption and income growth in a future draft.

[14] Recall that the model with permanent shocks did not require us to take a stand on whether individuals have fixed unobserved differences in earnings (fixed effects). However, this is not true for the case with $\rho < 1$. In principal, insurance coefficients can still be estimated for the case of persistent shocks and fixed effects, using the following IV approach. (1) For transitory shocks, regress $\Delta c_t$ on $\tilde{\Delta}\Delta y_t$ using $\tilde{\Delta}\Delta y_{t+2}$ as an

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

and for the persistent shock,

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

Thus, in both cases, expression (6) yields a consistent estimator of $\phi^x$.

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 described above; 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.

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 $\rho$. 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 though 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 $\rho$ 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 $\rho = 0.95$,

instrument. (2) For persistent shocks regress $\Delta c_t$ on $\tilde{\Delta}\Delta y_t$ using $\rho\tilde{\Delta}\Delta y_t - \tilde{\Delta}\Delta y_{t+2} - \left[(1+\rho)^2+\rho^2\right]\tilde{\Delta}\Delta y_{t-2}$ as an instrument.

in the baseline economy the misspecified BPP estimate of $\phi^{\eta}$ 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 0.38.

Turning to insurance coefficients for the transitory shocks, here the model with persistent shocks is slightly less successful. For $\rho = 0.95$, the misspecified BPP estimate of $\phi^{\varepsilon}$ is 0.84 while the empirical counterpart is 0.94. 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 0.10 below the unbiased empirical value.

Figure 3 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 effective difference between a permanent and a highly persistent shocks is more pertinent for a young household with many periods ahead.

# 6 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, at the individual level, one can study how consumption growth responds to income growth. 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 and 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 and 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. A limit of the empirical work in Hall and Mishkin, relative to BPP, is that they only use food consumption from the PSID.

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 the 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 based on the quadratic utility case line up very well with the Hall and Mishkin number (see the last line in Table 3), 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.

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 and 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 and 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, the private information environment achieves insurance levels closer to the data.

Since the seminal work of Storesletten et al. (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 and 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 (Heathcote et al. (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 and Perri (2006) and Heathcote et al. (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 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.[15] Blundell and 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 et al. (2005), Guvenen (2007), Huggett et al. (2006), Primiceri and van Rens (2006)) argue that advance information about future income fluctuations can help in rationalizing the joint behavior of income and consumption inequality in the U.S. economy. What we called "insurance" could be relabelled as "advance information". In the next draft of the paper, we will assess how much advance information about future shocks one must introduce into a Bewley model in order to reproduce the values of $\phi^\eta$ in the data.

# 7 Conclusions

This paper is inspired by some intriguing empirical results reported by Blundell, Pistaferri and Preston (2006, BPP). BPP estimate that in the US economy, during 1980-1992, almost

[15] Heathcote et al. (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). Given the observed rise in male and female wage inequality, the model roughly generates the right rise in consumption inequality for the United States.

40% of permanent income shocks and over 95% of transitory shocks to post-government income can be insured away by households. These two numbers, we argue, should become central in macroeconomics. They represent a benchmark against which current macroeconomic models used for quantitative analysis should be compared in order to assess whether they admit the right amount of household insurance.

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 uncover 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.

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 reconcile model and data, the model must be modified in such a way to provide more insurability to younger households.

Third, we show 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 assess the accuracy of the estimation method proposed by BPP. We show 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.

Finally, we have outlined a framework for measuring the distribution of lifetime insurability of income shocks in a model economy. For the Bewley model, this distribution is non-trivial, suggesting that some households end up having far greater insurance possibilities than others. In future work will address the issue of identifying and estimating the distribution of insurability in the data, and comparing it with the implied distributions from the model.

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 disjuncture 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.

# References

Abowd, J. and 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. and 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. and Preston, I. (2004), 'Imputing Consumption in the PSID Using Food Demand Estimates From the CEX', *Mimeo, UCL* .

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

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

Campbell, J. and 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'.

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

Chamberlain, G. and 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. and Krueger, D. (2007), ‘Taxing Capital?: Not a Bad Idea After All!’, *Mimeo* .

Cunha, F., Heckman, J. and 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. and Paxson, C. (1994), ‘Intertemporal Choice and Inequality’, *The Journal of Political Economy* **102**(3), 437–467.

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

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

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

Gouveia, M. and 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. and 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. and Violante, G. (2004), 'The Macroeconomic Implications of Rising Wage Inequality in the United States', *CEPR wp* **4296**.

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

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

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

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

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

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

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

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

Storesletten, K., Telmer, C. and 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.

| | Permanent Shock | | | Transitory Shock | | |
|---|---|---|---|---|---|---|
| | Data BPP | Model BPP | Model TRUE | Data BPP | Model BPP | Model TRUE |
| Disposable Income | 0.38 | 0.14 | 0.18 | 0.94 | 0.92 | 0.92 |
| Pre-Govt Earnings: | | | | | | |
| No consumption floor | 0.53 | 0.25 | 0.28 | 0.94 | 0.94 | 0.94 |
| $4000 consumption floor | 0.53 | 0.25 | 0.28 | 0.94 | 0.94 | 0.94 |
| $8000 consumption floor | 0.53 | 0.28 | 0.30 | 0.94 | 0.95 | 0.94 |

Table 1: Benchmark Results: Model and Data

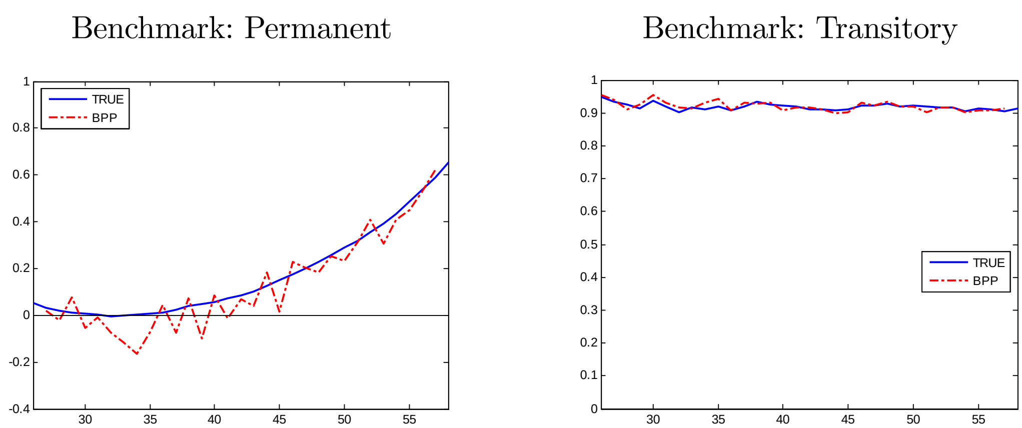


Figure 1: Insurance Coefficients By Age for Benchmark Model

| | Permanent Shock | | | Transitory Shock | | |
|---|---|---|---|---|---|---|
| | Data BPP | Model BPP | Model TRUE | Data BPP | Model BPP | Model TRUE |
| Low Wealth | 0.04 | −0.24 | −0.01 | 0.72 | 0.91 | 0.92 |
| High Wealth | 0.44 | 0.14 | 0.22 | 0.98 | 0.92 | 0.92 |
| Low Education | 0.08 | 0.14 | 0.18 | 0.91 | 0.92 | 0.92 |
| High Education | 0.67 | 0.15 | 0.19 | 0.96 | 0.91 | 0.91 |

Table 2: By Subsample

| Data | Permanent Shock | | Transitory Shock | |
|---|---|---|---|---|
| | 0.38 | | 0.94 | |
| | Model TRUE | Model BPP | Model TRUE | Model BPP |
| Benchmark | 0.18 | 0.14 | 0.92 | 0.92 |
| Initial Wealth Distribution | 0.18 | 0.14 | 0.92 | 0.92 |
| Risk Aversion: | | | | |
| $\gamma = 1$ | 0.16 | 0.12 | 0.92 | 0.92 |
| $\gamma = 5$ | 0.23 | 0.19 | 0.91 | 0.91 |
| $\gamma = 10$ | 0.32 | 0.25 | 0.89 | 0.90 |
| $\gamma = 15$ | 0.38 | 0.30 | 0.88 | 0.89 |
| Wealth Income Ratio: | | | | |
| $\frac{K}{Y} = 1.5$ | 0.09 | $-0.01$ | 0.82 | 0.82 |
| $\frac{K}{Y} = 2.5$ | 0.13 | 0.08 | 0.89 | 0.89 |
| Interest Rate: | | | | |
| $r = 4\%$ | 0.18 | 0.14 | 0.91 | 0.91 |
| $r = 2\%$ | 0.18 | 0.14 | 0.93 | 0.93 |
| Variance Permanent Shock: | | | | |
| $\sigma_\eta = 0.03$ | 0.18 | 0.15 | 0.91 | 0.91 |
| $\sigma_\eta = 0.01$ | 0.17 | 0.12 | 0.93 | 0.93 |
| Variance Initial Permanent: | | | | |
| $\sigma_{z_0} = 0.2$ | 0.18 | 0.13 | 0.92 | 0.92 |
| $\sigma_{z_0} = 0.1$ | 0.18 | 0.14 | 0.92 | 0.92 |
| Quadratic Preferences | 0.13 | 0.09 | 0.91 | 0.91 |

Table 3: Sensitivities

| | Permanent Shock | | | Transitory Shock | | |
|---|---|---|---|---|---|---|
| | Data BPP | Model BPP | Model TRUE | Data BPP | Model BPP | Model TRUE |
| Disposable Income | 0.38 | 0.08 | 0.20 | 0.94 | 0.80 | 0.81 |
| Pre-Govt Earnings: | | | | | | |
| No consumption floor | 0.53 | 0.22 | 0.29 | 0.94 | 0.89 | 0.89 |
| \$4000 consumption floor | 0.53 | 0.22 | 0.29 | 0.94 | 0.89 | 0.89 |
| \$8000 consumption floor | 0.53 | 0.26 | 0.33 | 0.94 | 0.89 | 0.89 |

Table 4: Results for Economy with Tight Borrowing Constraints: Model and Data

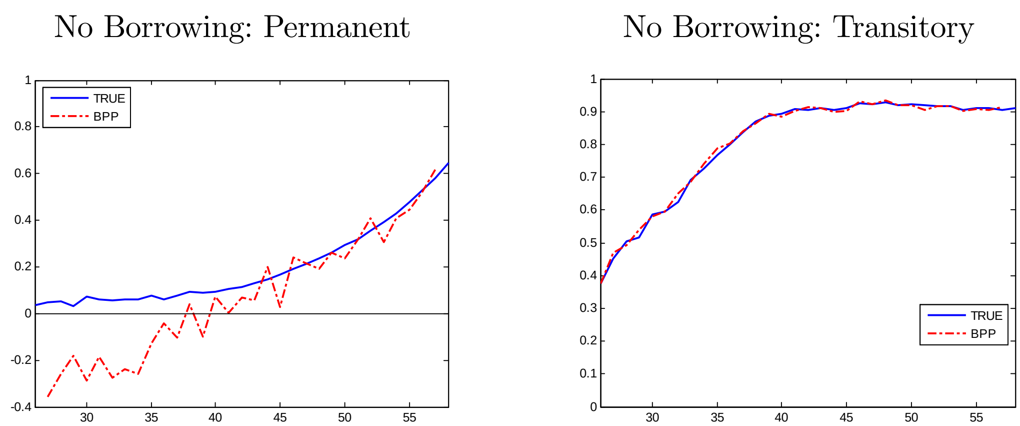


Figure 2: Insurance Coefficients By Age for Economy with Tight Borrowing Constraints

| | Permanent Shock | | | Transitory Shock | | |
|---|---|---|---|---|---|---|
| | Data BPP | Model BPP | Model TRUE | Data BPP | Model BPP | Model TRUE |
| Low Wealth | 0.04 | −0.31 | 0.04 | 0.72 | 0.46 | 0.46 |
| High Wealth | 0.44 | 0.06 | 0.24 | 0.98 | 0.88 | 0.88 |
| Low Education | 0.08 | 0.08 | 0.20 | 0.91 | 0.80 | 0.80 |
| High Education | 0.67 | 0.09 | 0.20 | 0.96 | 0.81 | 0.82 |

Table 5: Results for Economy with Tight Borrowing Constraints: By Subsample

| | Permanent Shock | | | Transitory Shock | | |
|---|---|---|---|---|---|---|
| Data | 0.38 | | | 0.94 | | |
| | Model TRUE | Model BPP | Model BPP misspecified | Model TRUE | Model BPP | Model BPP misspecified |
| Benchmark: | | | | | | |
| $\rho = 0.99$ | 0.23 | 0.19 | 0.20 | 0.91 | 0.92 | 0.91 |
| $\rho = 0.97$ | 0.31 | 0.27 | 0.28 | 0.90 | 0.90 | 0.88 |
| $\rho = 0.95$ | 0.37 | 0.33 | 0.34 | 0.88 | 0.88 | 0.84 |
| $\rho = 0.90$ | 0.47 | 0.44 | 0.44 | 0.85 | 0.85 | 0.75 |
| No Borrowing: | | | | | | |
| $\rho = 0.99$ | 0.24 | 0.14 | 0.15 | 0.80 | 0.80 | 0.80 |
| $\rho = 0.97$ | 0.29 | 0.24 | 0.24 | 0.80 | 0.80 | 0.79 |
| $\rho = 0.95$ | 0.34 | 0.30 | 0.30 | 0.80 | 0.80 | 0.77 |
| $\rho = 0.90$ | 0.44 | 0.42 | 0.41 | 0.81 | 0.80 | 0.71 |

Table 6: Insurance Coefficients with Persistent Income Shocks

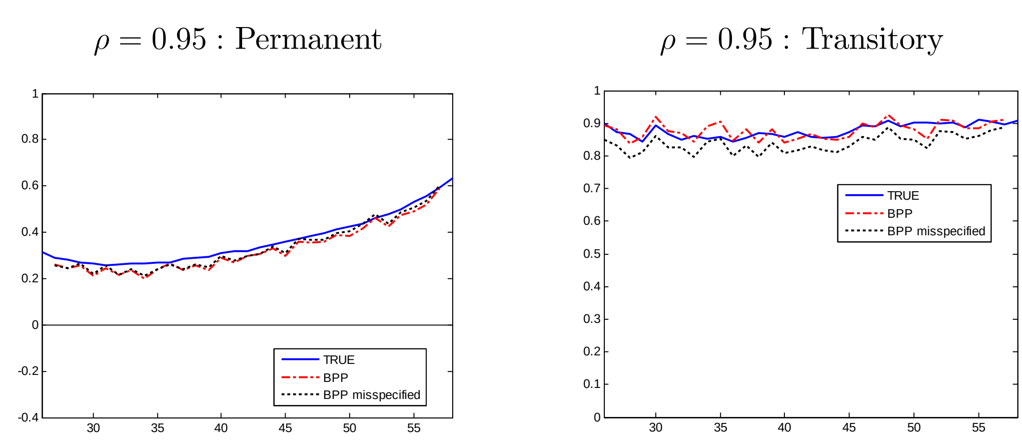


Figure 3: Insurance Coefficients By Age for Economy with $\rho = 0.95$

Mean and St. Dev. of $\phi_{it}^{\eta}$ by age
Permanent Shock

Mean and St. Dev. of $\phi_{it}^{\varepsilon}$ by age
Transitory Shock

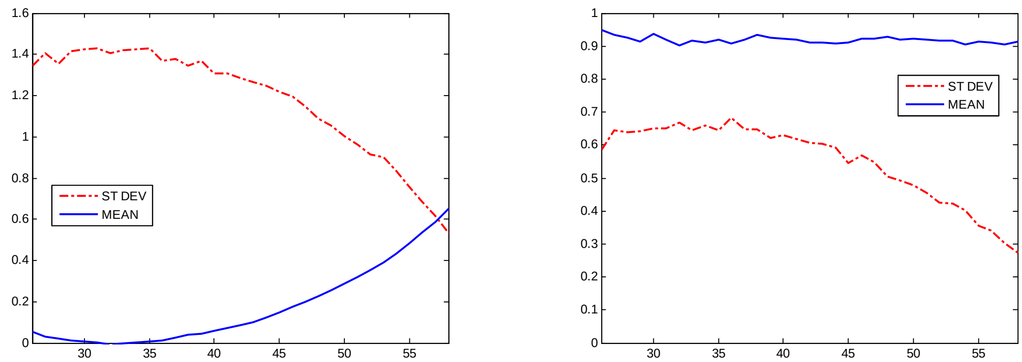


Figure 4: Mean and Standard Deviation of Individual Insurance Coefficients By Age

Distribution of $\overline{\phi_i^{\eta}}$ : Permanent Shocks

Distribution of $\overline{\phi_i^{\varepsilon}}$ : Transitory Shocks

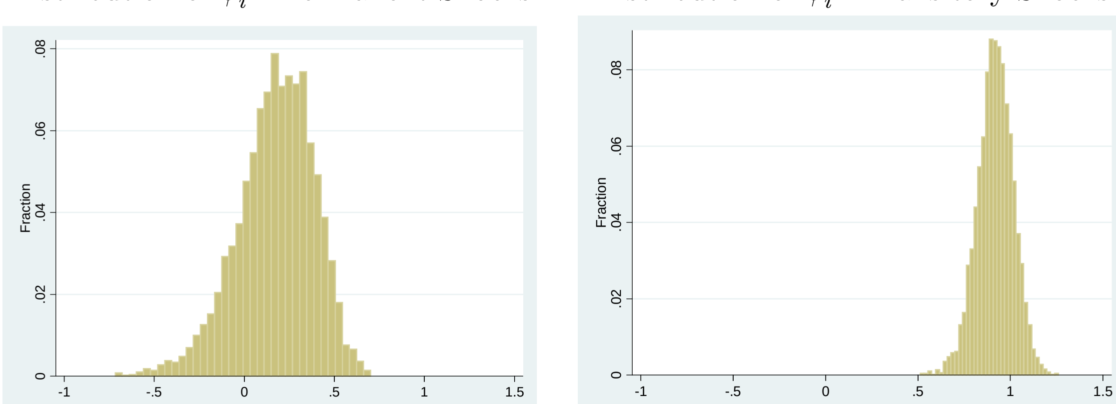


Figure 5: Distribution of Lifetime Individual Insurance Coefficients