A Bayesian Approach to Estimating Household Income Risk and Consumption Insurance

Arpita Chatterjee* James Morley† Aarti Singh‡

September 20, 2016 §

Keywords: Bayesian estimation, income risk, consumption insurance, panel unobserved components model.

JEL codes: E32; E22; C32

*Corresponding author: School of Economics, UNSW Business School, University of New South Wales, Australia; Email: arpita.chatterjee@unsw.edu.au.

†School of Economics, UNSW Business School, University of New South Wales, Australia; Email: james.morley@unsw.edu.au.

‡School of Economics, University of Sydney, Australia; Email: aarti.singh@sydney.edu.au.

§We thank James Hansen, and participants at the Workshop of the Australasian Macroeconomics Society and Internal Brown Bag at University of Sydney for comments. The usual disclaimers apply.

Abstract

Motivated by the time series literature, we propose a panel unobserved components model of household income and consumption. Despite data limitations, this model can be estimated using full information Bayesian methods assuming a common distribution but independent shocks across households. Using U.S. data at an annual frequency, we find empirical support for a relatively simple specification with no persistence in or correlation across transitory components and no cointegration between household income and consumption. Notably, we can reject the implicit specification in Blundell, Pistaferri, and Preston's (2008) partial information moments-based approach and, despite similar estimates of income risk for the same dataset, our estimates of consumption insurance are higher than theirs. However, the Bayesian estimation reveals that consumption insurance is not particularly well identified in the data, as it can be heavily influenced by the prior on key model parameters. Subgroup estimation shows that our results regarding model specification are robust and that patterns of consumption insurance estimates are in line with what would be expected, with insurance being higher for older or more educated households.

1 Introduction

Blundell, Pistaferri and Preston (2008) (BPP hereafter) propose a panel model to estimate household income risk and consumption insurance without imposing a particular structural model of household's behavior and decisions. Using a partial information moments-based approach to estimation for U.S. panel data from the PSID and CEX Survey, they find a relatively low degree of consumption insurance in response to idiosyncratic shocks to permanent income. Their results have been challenged by Kaplan and Violante (2010), who suggest that the BPP estimation strategy will lead to a downward bias in consumption insurance, and the bias will be more pronounced for households that are borrowing constrained.

In this paper, we propose an alternative estimation strategy for a more general panel model of household income and consumption. In particular, we consider a full information Bayesian approach to estimation using a panel unobserved components (UC hereafter) model that is similar to a multivariate unobserved components model that has previously been estimated for quarterly aggregate income and consumption data by Morley (2007). Despite data limitations such as missing observations and a short time dimension to the panel given annual data, we are able to estimate the model assuming a common distribution but independent shocks across households and given reasonable priors on model parameters.

Using the same U.S. data considered in BPP, we find empirical support for a relatively simple specification with no persistence in or correlation across transitory components. Also, there appears to be no cointegration between household income and consumption. These results stand in contrast to those for the aggregate data, although this is perhaps not surprising given that common shocks have been removed from the data and idiosyncratic shocks are likely due to very different factors with different behaviors than the common shocks that drive the aggregate data.

Notably, we can reject the implicit specification in BPP for which transitory income shocks are assumed to have permanent effects on consumption and income growth follows an MA(2) process, contrary to what is implied by sample autocovariances and partial autocovariances. In addition, despite similar estimates of income risk, our estimates of consumption insurance are considerably higher than in BPP. Thus, our alternative approach potentially addresses the possible downward bias suggested by Kaplan and Violante (2010).

At the same time, a key finding from the Bayesian analysis is that consumption insurance is not particularly well identified. Specifically, estimates are quite sensitive to priors on key model parameters. However, differences in estimates of consumption insurance across different subgroups are much better identified, with consumption insurance being higher for older or more educated households.

The rest of this paper is organized as follows: Section 2 presents the general panel unobserved components model proposed in this paper, describes the data and priors and presents Bayesian estimates of alternative versions of the model; Section 3 compares our preferred model with the implied model estimated in BPP and examines the sensitivity of estimates to priors; Section 4 presents consumption insurance estimates for different subgroup of the panel and Section 5 concludes.

2 A Panel Unobserved Components model

In this section we specify a general panel UC model of household income and consumption and report our results for key specifications along with model comparison.

2.1 Model setup

Following Friedman and Kuznet (1954), the estimated household income process typically has a permanent random walk component, a transitory component that dies away, and zero correlation between the two components. See, for example, Moffitt and Gottschalk (2002), Storesletten, Telmer and Yaron (2004), Guvenen (2007), Blundell, Pistaferri and Preston (2008), Primiceri and van Rens (2009), Low, Meghir and Pistaferri (2010), and Heathcote, Perri and Violante (2010), among many others. However, it is straightforward to show that, if the zero correlation is incorrectly assumed, the model mis-specification will bias the estimate of permanent risk, a key ingredient in heterogeneous agent quantitative macro models.[1] In the time series literature using aggregate GDP data, Morley, Nelson and Zivot (2003) clearly establish that the assumption of zero correlation can be rejected for U.S. quarterly real GDP data, while Morley (2007) finds evidence in favor of correlated permanent-transitory shock using aggregate U.S. income and consumption data in a multivariate unobserved components model.[2] Motivated by these results, the model presented below allows for a correlated unobserved components with ARMA dynamics for the transitory components.

Our panel unobserved components model decomposes residual income and residual consumption into a permanent component and a transitory deviation from the permanent component for individual i:[3]

$$y_{i,t} = \tau_{i,t} + (y_{i,t} - \tau_{i,t}), \tag{1}$$

$$c_{i,t} = \gamma_\eta \tau_{i,t} + \kappa_{i,t} + (c_{i,t} - \gamma_c \tau_{i,t} - \kappa_{i,t}). \tag{2}$$

The permanent components are specified as follows:

$$\tau_{i,t} = \mu_{\tau,i} + \tau_{i,t-1} + \eta_{i,t}, \qquad \eta_{i,t} \sim i.i.d.N(0, \sigma_\eta), \tag{3}$$

$$\kappa_{i,t} = \mu_{\kappa,i} + \kappa_{i,t-1} + u_{i,t}, \qquad u_{i,t} \sim i.i.d.N(0, \sigma_u), \tag{4}$$

[1] See Ejrnaes and Browning (2014) for more details on the specifics of the bias and how to estimate models without assuming that the shocks are uncorrelated.

[2] Note, however, this uses total income, not just labor income as is considered in this paper.

[3] BPP compute residual income and residual consumption after removing the impact of observables such as education, race, family size, number of children, region, employment status, year and cohort effects, residence in large city, income recipient other than husband and wife from total household disposable labor income

The transitory components follow stationary invertible ARMA(p,q) processes:

$$\Phi_y(L)(y_{i,t} - \tau_{i,t}) = \lambda_{y\eta}\eta_{i,t} + \Psi_y(L)\epsilon_{i,t}, \tag{5}$$

$$\Phi_c(L)(c_{i,t} - \gamma_\eta \tau_{i,t} - \kappa_{i,t}) = \lambda_{c\eta}\eta_{i,t} + \lambda_{c\epsilon}\epsilon_{i,t} + \Psi_c(L)v_{i,t}, \tag{6}$$

where the $\Phi_y(L)) = (1 - \phi_{y,1}L - \phi_{y,2}L^2 - ... - \phi_{y,p}L^p)^{-1}$ and $\Psi_y(L)) = (1 - \psi_{y,1}L - \psi_{y,2}L^2 - ... - \psi_{y,q}L^q)^{-1}$ are the ARMA(p,q) specification for transitory residual income. Residual consumption also has the same lag structure. In our model, instead of directly allowing the shocks to be correlated, we assume that the permanent and transitory components are correlated. As a result, $\lambda_{y\eta}$ and $\lambda_{c\eta}$ are the impact coefficients for transitory income and consumption in response to the permanent shock to income and $\lambda_{c\epsilon}$ is the response of transitory consumption to transitory income shocks.

The permanent income shock in our empirical model is $\eta_{i,t}$ and it can be interpreted as shocks to health, promotion or other shocks that result in a change in permanent income. Examples of permanent shock to consumption, in addition the permanent shocks to income, could be taste and preference shocks or permanent shocks to non-labor income, such as wealth shocks. The transitory consumption shock could capture measurement error, such as that which could be due to the imputation of non-durable consumption. The model assumes time-invariant volatilities of shocks, although we could, in principle, test for and allow structural breaks in these parameters.[4]

Based on our panel UC model, we can solve for consumption growth for household *i*:

$$\Delta c_{i,t} = \gamma_\eta \eta_{i,t} + u_t + (1 - L)\Phi_c(L)^{-1}(\lambda_{c\eta}\eta_{i,t} + \lambda_{c\epsilon}\epsilon_{i,t} + \Psi_c(L)v_{i,t}), \tag{7}$$

which suggests that changes in consumption depend on the full history of

[4]For example, BPP use the same panel data to look at changes in income and consumption inequality over time. We leave this for future research and focus on estimating income risk and consumption insurance.

permanent shocks to income.[5]

To calculate the implied consumption insurance based on our UC model, a change in consumption at date t due to the permanent income shock η_t is $\gamma_\eta + \lambda_{c\eta}$. Therefore, the consumption insurance coefficient is

$$\vartheta_c = 1 - (\gamma_\eta + \lambda_{c\eta}). \tag{8}$$

Kaplan and Violante (2010) define the insurance coefficient with respect to permanent income shock as the share of the variance of the shock that does not translate into consumption growth such that

$$\vartheta_c = 1 - \frac{cov(\Delta c_t, \eta_t)}{var(\eta_t)} = 1 - \frac{cov(\gamma_\eta \eta_t + z_t^c, \eta_t)}{\sigma_\eta^2} = 1 - \frac{\gamma_\eta \sigma_\eta^2 + \lambda_{c\eta}\sigma_\eta^2}{\sigma_\eta^2}, \tag{9}$$

which the same as equation (8) in our empirical model. See Appendix B for more details and definition of z_t^c.

2.2 Data, methods, and priors

In this subsection we first briefly describe the data created by BPP and discuss our methods and priors. For full details of the data, we refer the reader to the BPP paper.

BPP use the PSID sample from 1978-1992 of continuously married couples headed by a male (with or without children) age 30 to 65. The income variable is family disposable income which includes transfers. They adopt a similar sample selection adopted in the CEX data. Since CEX data has detailed non-durable consumption data, unlike PSID which primarily has food expenditure data, they impute non-durable consumption for each household per year in using the estimates of the food demand from CEX. The final data is a panel of income and imputed non-durable consumption. To get the residual income (and consumption), they regress income (consumption) of household on a vector of regressors including demographic,

[5] In Kaplan and Violante's (2010) language this would mean that there is *no* "short memory" in our model.

education, ethnic and other income characteristics observable/known by consumers. The residual, unobservable, random and idiosyncratic income and consumption is modeled in section 2.1.

Instead of relying on a few moments to estimate key parameters of the permanent-transitory model, we use the entire likelihood. A clear benefit of this full-information approach is that it addresses the extreme sensitivity of inferences to particular moments, such as those related to levels vs. differences (add citations). Moreover, in our approach we compute the unobserved components using Kalman filtering techniques.

Due to short time series typical of panel data on income and consumption, maximum likelihood estimation of a flexible unobserved components model is not viable in practice. Instead we rely on prior distributions for parameters based on past studies and a priori reasoning and apply Bayesian techniques to explore the posterior distribution. In particular, we estimate the panel UC model using Bayesian posterior simulation based on Markov-chain Monte Carlo (MCMC) methods. We use multi-block random-walk chain version of the Metropolis-Hastings (MH) algorithm with 20,000 draws after a burn-in of 20,000 draws. To check the robustness of our posterior moments, we use different starting values.

The prior distributions are chosen relying on vast empirical literature on modeling income and consumption dynamics. The priors used in our estimation are the following. The priors for the precision (inverse variances) are $\Gamma(2.5, 2.5)$. Since there is no consensus in the literature regarding the estimate of the impact of permanent income shock on consumption, γ_η we chose an uninformative $U(0, 1)$ prior. We truncate it to ensure that it lies between -1 and 1. The impact coefficients, $\lambda_{y\eta}$, $\lambda_{c\eta}$ and $\lambda_{c\epsilon}$ are $N(0, 0.5^2)$ and are truncated to ensure that they lie between -1 and 1; the autoregressive coefficients are $N(0, 0.5^2)$ and are truncated to ensure stationarity.

TABLE 1. ESTIMATES OF PANEL UC MODELS

	UC-AR(2)			UC-WN
	full	$\sigma_u = 0$	$\lambda = 0$	full
		INCOME		
ϕ_{y1}	-0.02 (0.01)	-0.10 (0.02)	-0.08 (0.01)	
ϕ_{y2}	-0.05 (0.01)	-0.10 (0.02)	-0.11 (0.01)	
σ_η	1.41 (0.01)	1.42 (0.01)	1.41 (0.01)	1.39 (0.01)
σ_ϵ	1.61 (0.02)	1.58 (0.02)	1.64 (0.02)	1.69 (0.01)
$\lambda_{y\eta}$	0.11 (0.01)	0.10 (0.01)		
		CONSUMPTION		
ϕ_{c1}	-0.12 (0.01)	-0.37 (0.01)	-0.30 (0.01)	
ϕ_{c2}	-0.07 (0.01)	-0.32 (0.01)	-0.20 (0.01)	
σ_u	1.32 (0.01)		1.36 (0.01)	1.30 (0.01)
σ_v	2.03(0.02)	2.87 (0.02)	1.90 (0.02)	2.14 (0.02)
$\lambda_{c\eta}$	0.05 (0.01)	0.02 (0.01)		
$\lambda_{c\epsilon}$	-0.03 (0.01)	-0.02 (0.01)		
γ_η	0.47 (0.02)	0.43 (0.02)	0.47 (0.02)	0.47 (0.02)
		IMPLIED VARIANCE		
Δy	0.08 (0.00)	0.08 (0.00)	0.08 (0.00)	0.08 (0.00)
Δc	0.16 (0.00)	0.12 (0.00)	0.17 (0.00)	0.15 (0.00)
		MARGINAL LIKELIHOOD (IN LOGS)		
	-89595	-110416	-90295	-89041

Notes: The table reports posterior means of panel UC model parameters with posterior standard deviations reported in parentheses. The third panel reports the variance of residual income and residual consumption growth implied by the model and the marginal likelihood is in the bottom panel. The total number of observations are 1765.

2.3 Bayesian estimates and model comparison

The literature on earnings has in recent years moved away from a simple model where the permanent component is a random walk while the transitory component is white noise, which we refer to as the UC-WN model. It is believed that the earnings dynamics are more complex with serially correlated transitory components, see see Atkinson et al., 1992, for a survey. Our two main specifications of interest, UC-AR(2) and UC-WN, therefore encompass the views held in the literature. As the literature is large, we list here a few other specifications to highlight that a wide variety of specifications have been used before. MaCurdy (1982) and Abowd and Card (1982) find that the covariance matrix of earnings differences fits an MA(2), Gottschalk and Moffitt (1993) fit random walk plus ARMA(1,1) in levels which is an ARMA(1,2) in first differences, Heathcote, Storesletten and Violante (2010, 2014) employ a very persistent component and a white noise transitory component. Motivated by the aggregate time series literature, we also investigate whether dynamics play an important role in panel data and if income and consumption share a common trend like they do in the aggregate data. The results along are in Table 1.

Estimates of the UC-AR(2) model in column 2 show that some of the dynamics, for example the AR dynamics of consumption and the impact of permanent income shocks on transitory income, which is positive, are significantly different from zero. Other studies, e.g. Hyrshko (2010) and Belzil and Bognanno (2008), however find a negative correlation between the permanent and transitory shocks. For both the income and consumption process, transitory shocks are more volatile compared to permanent shocks. The implied variance of income and consumption closely matches their corresponding counterparts in the data, 0.09 and 0.16 (see last column in the bottom panel of Table 1).[6] If we shut down the additional permanent shocks to consumption beyond permanent income shocks, col-

[6]Is the UC-AR(2) model is identified? Morley, Nelson, Zivot (2003) show that as long as p≥q+2, the UC-ARMA (p,q) model is identified and this condition is met in our UC-AR(2) specification.

umn 3 in Table 1, not surprisingly the variance of consumption is much lower than the variance of consumption in the data. In column 4 we set the impact coefficients to zero and finally in column 5, we shut down all the dynamics.

Note that when we estimate different models, the variance of income shocks, transitory shocks to consumption and the pass-though of the permanent income shock to consumption are quite similar. Recall, that the consumption insurance in the full model and the one in which we shut down the permanent shocks to consumption is $1 - \lambda_{c\eta} - \gamma_{\eta}$.

Using Bayesian likelihood methods with a more flexible panel UC model we can compare different models that we have estimated to determine which model fits the data best. We do that by computing the marginal likelihood following Chib and Jeliazkov (2001). The last row in Table 1 clearly shows that the UC-WN model is clearly preferred. Note that UC-WN is a simple canonical model often used in the earnings empirical literature. Another advantage of using Bayesian methods is that we can estimate the variances of residual income and residual consumption growth which are in fact functions of the parameters of the model and report their standard deviation. The implied volatility of income and consumption growth in our preferred model is 0.08 and 0.15. The corresponding counterpart in the BPP dataset are 0.09 and 0.16.

2.4 What is the affect of small T on our estimates?

In this section we examine how our method fairs when we use simulated panel data with a small T from a UC-WN model. To be conservative we set N=700 and consider T=10. The results for T=6 are provided in the appendix. We consider three cases where we vary γ_{η}. Table 2 reports our simulation results. Our methods are able to recover the true parameters quite well, while again it is worth noting that MLE would not be viable given a small T. Note that, as T becomes smaller, the estimate of γ_{η} appears to have an upward bias (see appendix) and therefore our estimate

TABLE 2. SIMULATION RESULTS

	DGP	$\gamma_\eta = 0.25$	$\gamma_\eta = 0.45$	$\gamma_\eta = 0.65$
		INCOME		
σ_η	1.40	1.41 (0.02)	1.40 (0.02)	1.38 (0.02)
σ_ϵ	1.7	1.54 (0.02)	1.54 (0.02)	1.56 (0.02)
		CONSUMPTION		
σ_u	1.3	1.49 (0.02)	1.49 (0.02)	1.48 (0.02)
σ_v	2.1	1.70 (0.02)	1.70 (0.02)	1.73 (0.02)
γ_η		0.25 (0.03)	0.47 (0.03)	0.72 (0.03)

Notes: The table reports posterior means of model parameters with posterior standard deviations reported in parentheses. In the simulated data, N=700 and T=10.

of consumption insurance could be seen to provide a lower bound for the true consumption insurance in the data. In our preferred specification, consumption insurance is 53 percent. Heathcote, Storesletten and Violante (2014) find consumption insurance to be close to 60 percent in a structural model on income and consumption dynamics.

3 Comparison with the BPP model

How does our preferred model compare with the BPP model and their estimates? In this section we also evaluate whether using our Bayesian approach, we can recover the BPP estimates.

3.1 BPP model

Here we first re-write the BPP model for completeness (the state space representation is in Appendix A):

$$y_{i,t} = \tau_{i,t} + (y_{i,t} - \tau_{i,t}). \tag{10}$$

The permanent component of income is specified as follows:

$$\tau_{i,t} = \mu_i + \tau_{i,t-1} + \eta_{i,t}, \qquad \eta_{i,t} \sim i.i.d.N(0, \sigma_\eta). \tag{11}$$

The transitory component has an $MA(1)$ specification as follows:

$$(y_{i,t} - \tau_{i,t}) = \epsilon_{i,t} + \theta\epsilon_{i,t-1}, \qquad \epsilon_{i,t} \sim i.i.d.N(0, \sigma_\epsilon), \tag{12}$$

Therefore, consumption growth, or change in logarithm of residual consumption has the following process:

$$\Delta c_{i,t} = \gamma_\eta \eta_{i,t} + \gamma_\epsilon \epsilon_{i,t} + u_{i,t} + \Delta u^\star_{i,t}, \qquad u_{i,t} \sim i.i.d.N(0, \sigma_u), \tag{13}$$

where η_t and ϵ_t are the permanent and transitory income shocks, u_t is the permanent shock to consumption and $u^\star_{i,t} \sim i.i.d.N(0, \sigma_{u^\star})$ is measurement error for consumption. As discussed in BPP, the permanent shock to consumption captures taste and preference shocks, while the measurement error reflects the imputation of consumption.

Note that in this specification, the transitory income shocks has a permanent effect on consumption. To see this, we can rewrite the level of consumption, after suppressing the individual specific subscript for simplicity, as

$$c_t = \gamma_\eta \tau_t + \gamma_\epsilon Z_{\epsilon,t} + Z_{u,t} + u^\star_{i,t}, \tag{14}$$

$$Z_{\epsilon,t} = Z_{\epsilon,t-1} + \epsilon_t, \tag{15}$$

$$Z_{u,t} = Z_{u,t-1} + u_t, \tag{16}$$

Our panel UC model differs from the BPP specification along two key dimensions. First, in our preferred UC-WN specification, there are no persistent dynamics for the transitory income and consumption processes, while in BPP transitory income follows an MA(1) process. Second, and

TABLE 3. SAMPLE ACF AND PACF

a_1	a_2	a_3	p_1	p_2	p_3
		Δy			
-0.29	-0.03	-0.01	-0.29	-0.13	-0.07
		Δc			
-0.34	-0.01	-0.02	-0.34	-0.14	-0.04

Notes: The total number of 12041 observations are used to calculate autocorrelations and partial autocorrelations.

most importantly, a transitory income shock can only impact transitory consumption in our model, while in the BPP model, transitory income shocks have a permanent impact on consumption. In the next section, we examine which specification is supported by the BPP dataset.

3.2 Income and consumption dynamics

We compute sample autocorrelation function and partial autocorrelation functions for residual income and residual consumption from BPP dataset by pooling individuals of all ages and over all years.

In Table 3, we can see that the ACF cuts off after 1 lag but the PACF tails off more gradually for both income growth and consumption growth. This is consistent with MA(1) process, not an MA(2) process, as implied by the BPP model. By contrast, our preferred model based on marginal likelihood analysis implies MA(1) processes for both income growth and consumption growth, consistent with the ACF and PACF results for the BPP dataset.

3.3 Bayesian estimates for BPP model

In estimating the BPP model, we assume that θ and γ_ϵ are truncated $TN_{[-1,1]}(0, 0.5^2)$. That is, they are truncated to lie between -1 and 1. Then, to examine the

TABLE 4. BPP ESTIMATES USING BAYESIAN APPROACH

Prior γ_η	$TN_{[-1,1]}(0.65, 0.25^2)$	$TN_{[-1,1]}(0.65, 1^2)$	$U(0,1)$	BPP estimate/data
		INCOME		
θ	0.01(0.01)	0.001(0.01)	0.04(0.02)	0.11
σ_η	1.41(0.01)	1.43(0.013)	1.45(0.015)	1.41
σ_ϵ	1.71(0.01)	1.69(0.015)	1.66(0.019)	1.73
		CONSUMPTION		
γ_η	0.64(0.01)	0.54(0.02)	0.45(0.02)	0.64
γ_ϵ	0.002(0.00)	0.003(0.003)	0.004(0.004)	0.05
σ_u	1.31(0.01)	1.30(0.01)	1.30(0.01)	1.10
$\sigma_{u^\star}$	2.11(0.02)	2.11(0.02)	2.15(0.02)	NA
		IMPLIED VERSUS ACTUAL VARIANCE		
Δy	0.08(0.001)	0.08(0.001)	0.08(0.001)	0.09
Δc	0.11(0.001)	0.11(0.001)	0.11(0.001)	0.16

Notes: The table reports posterior means of model parameters with posterior standard deviations reported in parentheses. The bottom panel reports the variance of residual income and residual consumption growth implied by the model versus the corresponding averages in the BPP data.

role of the prior on γ_η, we also vary it.[7]

Table 4 reports the estimates of the BPP model using Bayesian methods.It is quite clear our estimation method can recover the volatility of income shocks and the consumption insurance parameter when the prior on γ_η is tight around the BPP estimate.[8] However, when the prior is less informative with $TN_{[-1,1]}(0.65, 1^2)$ the estimate moves towards higher consumption insurance and what we find with a uniform prior for our preferred UC-0 model. Note that in the last row of Table 4, the implied variance of residual consumption growth is 0.11 while the variance of residual consumption growth in the BPP sample is 0.16. This seems plausible as Figure 5 in BPP suggests that the process of consumption growth implied by their baseline model does not match the data that well in the latter part of the sample.

Therefore, using the full information approach, we can estimate the

[7] All the other priors are the same as for our panel UC model.

[8] The estimate of the standard deviation of the measurement error is not reported in Table 6 of BPP.

BPP model and more importantly also determine if some parameters are not well identified. In this case we find that consumption insurance is not well identified. Moreover, when the prior is uninformative, we get consumption insurance of 55 percent.

3.4 Related literature on consumption insurance

Are there formal markets or informal arrangements that insure households against idiosyncratic and unexpected movements in their income or wealth? How does this vary across households and over their life-cycle even when borrowing and lending opportunities change? The literature has keenly focused on these issues, both to inform the incomplete-markets literature and prescribe policies after understanding the extent of market incompleteness.

From the simple test of consumption insurance of Cochrane (1991) and Townsend (1994)'s consumption insurance in village India to the measurement of consumption insurance a panel data on household income and (imputed) non-durable consumption and quasi structural approach suggested by Deaton (1997) and implemented by BPP(2008), to the structural approach of Heathcote, Storesletten and Violante (2014), the literature has made important advances that enhance our understanding of the extent of consumption insurance. It has clearly been established that for the U.S. economy at least, there is no evidence for the two extremes of full insurance or zero insurance. However, quantitative estimates differ significantly and so do the methods with respect to the identification of the permanent and transitory shocks to income. For example, Kruger and Perri (2011) simply compute the change in non-durable consumption because of a change in income. Using Italian Survey of Household Income and Wealth from 1987 to 2008, they find that in the short-run non-durable consumption increases by 23 cents when after-tax income increases by 1 Euro in the households without any business wealth or real estate and the consumption response is stronger for a longer run income changes. Others have proxied permanent and transitory income shocks, for example Souleles (1999). Recent paper by HSV(2014) takes a structural approach.

Our approach is similar in spirit to BPP. However, we differ from them by taking a full information approach to estimating a panel unobserved components model. We see that as the main contribution relative to the large literature that uses partial information approach, primarily because of data limitations. We show that our approach can estimate underlying parameters with a relatively short time dimension, determine which parameters are not well identified, and compare models with different specifications, especially in terms of the transitory income and consumption processes.

4 Subgroup estimates

In this section, we examine how estimates vary across different groups based on age and education and check the robustness of our model specification.

Comparing different models, as in Table 1 for the whole sample, we find that the UC-WN model is preferred by the data even for subgroups. Therefore, in Tables 5 and 6, we report results for this specification only.

From Table 5 it is clear that while the group with college education encounter permanent shocks that are more variable, the pass-through of these shocks to consumption is 34 percent which is approximately half of the pass through of income shocks to consumption for the households with no college education. Qualitatively, these results are similar to BPP however the magnitudes are different. These results also seem consistent with the reasons outlined in KV where they argue that the downward bias in consumption insurance using BPP estimator is much more pronounced for households that are borrowing constrained. In particular, it is likely that households without college education are generally more borrowing-constrained and we find that while consumption insurance for no college group is 6 times higher than BPP (.35 in ours versus .06 in BPP), the down-

TABLE 4. EDUCATION HETEROGENEITY

	No college	College
	INCOME	
σ_η	1.37 (0.02)	1.41 (0.02)
σ_ϵ	1.76 (0.02)	1.62 (0.02)
	CONSUMPTION	
σ_u	1.31 (0.02)	1.28 (0.02)
σ_v	2.38(0.03)	1.89 (0.02)
γ_η	0.65 (0.03)	0.34 (0.02)

Notes: The table reports posterior means of model parameters with posterior standard deviations reported in parentheses. There are 883 observations in the no college group and 882 in the college group.

TABLE 5. AGE HETEROGENEITY

	Young (30-47)	Old (48-65)
	INCOME	
σ_η	1.32 (0.01)	1.42 (0.02)
σ_ϵ	1.50 (0.01)	1.92 (0.02)
	CONSUMPTION	
σ_u	1.30 (0.01)	1.30 (0.02)
σ_v	2.11(0.02)	1.96 (0.02)
γ_η	0.55 (0.03)	0.40 (0.03)

Notes: The table reports posterior means of model parameters with posterior standard deviations reported in parentheses. There are 1413 observations for the young while the number of observations for the old is 708.

ward bias is not so large for households with college education (.66 in ours versus .58 in BPP).

From Table 6, we again see that while income is more variable for the old, both permanent and transitory shocks are more volatile, γ_η for the old is 0.40 while it is 0.55 for the young. According to BPP, there is no clear evidence of an age profile with respect to consumption insurance.

These results once again highlight that using our likelihood approach we can estimate models with smaller N and T and find clear evidence of heterogeneity in consumption insurance with respect to education and age. In particular, college educated or people above the age of 47 have more consumption insurance relative to their counterparts, i.e., those with no college or less than 47 years in age in the BPP dataset. Both these results are consistent with the idea in Krueger and Perri (2006) that higher idiosyncratic income fluctuations can cause a change in the development of financial markets in a way that allows individual households to insure against these shocks better such that there is a lesser impact of these shocks on their consumption.

5 Conclusions

We have provided an econometric explanation for BPP's low estimates, which is primarily due to the partial information estimation approach, although model specification is relevant as well. Our approach produces higher estimates of consumption insurance, but with the caveat that the parameter is not particularly well identified. Our approach can go part way to address KV critique, although our model is not structural and does not have mechanisms KV discuss. According to KV the bias in BPP is higher for more borrowing constrained households and we see that for no college. Consumption insurance in BPP for this group is just 6 percent while we find that to be 35 percent, almost 6 times more. For college, BPP's estimate is 58 percent while ours is 66 percent. But, like BPP, our model provides structural facts to be replicated by more structural models. It also makes full use of data, is consistent with the dynamic properties of the data, and makes clear the role of priors and data in determining estimates.

A STATE-SPACE REPRESENTATION

In this section we first present the state-space representation of the BPP model and then our general baseline model

For the BPP model, the observation equation is

$$\tilde{\mathbf{y}}_t = \mathbf{H}\,\boldsymbol{\beta}_t$$

where

$$\tilde{\mathbf{y}}_t = \begin{bmatrix} y_t \\ c_t \end{bmatrix},\ \mathbf{H} = \begin{bmatrix} 1 & \theta & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & \gamma_\epsilon & \gamma_\eta & 1 \end{bmatrix} \text{ and } \boldsymbol{\beta}_t = \begin{bmatrix} \epsilon_t \\ \epsilon_{t-1} \\ v_t^\star \\ Z_{\epsilon,t} \\ \tau_t \\ Z_{u,t} \end{bmatrix}$$

The state equation is

$$\boldsymbol{\beta}_t = \boldsymbol{F}\boldsymbol{\beta}_{t-1} + \tilde{\boldsymbol{v}}_t$$

where

$$\mathbf{F} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix},\ \tilde{\boldsymbol{v}}_t = \begin{bmatrix} \epsilon_t \\ 0 \\ \epsilon_t \\ \eta_t \\ u_t \end{bmatrix}$$

and the covariance matrix of $\tilde{\boldsymbol{v}}_t$, $\mathbf{Q}$, is given by

$$\mathbf{Q} = \begin{pmatrix} \sigma_\epsilon^2 & 0 & 0 & \sigma_\epsilon^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma_v^2 & 0 & 0 & 0 \\ \sigma_\epsilon^2 & 0 & 0 & \sigma_\epsilon^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & \sigma_\eta 2 & 0 \\ 0 & 0 & 0 & 0 & 0 & \sigma_u^2 \end{pmatrix}$$

The state-space representation of our model is standard. The observation equation is

$$\tilde{\mathbf{y}}_t = \mathbf{H}\,\boldsymbol{\beta}_t$$

where

$$\tilde{\mathbf{y}}_t = \begin{bmatrix} y_t \\ c_t \end{bmatrix}, \mathbf{H} = \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & \gamma_\eta & 1 \end{bmatrix} \text{ and } \boldsymbol{\beta}_t = \begin{bmatrix} y_t - \tau_t \\ y_{t-1} - \tau_{t-1} \\ \tau_t \\ c_t - \tau_t \\ c_{t-1} - \tau_{t-1} \\ \kappa_t \end{bmatrix}$$

The state equation is

$$\boldsymbol{\beta}_t = \boldsymbol{F}\boldsymbol{\beta}_{t-1} + \tilde{\boldsymbol{v}}_t$$

where

$$\mathbf{F} = \begin{bmatrix} \phi_{y,1} & \phi_{y,1} & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \phi_{c,1} & \phi_{c,2} & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \tilde{\boldsymbol{v}}_t = \begin{bmatrix} \lambda_{y\eta}\eta_t + \epsilon_t \\ 0 \\ \lambda_{c\eta}\eta_t + \lambda_{c\epsilon}\epsilon_t + v_t \\ 0 \\ \eta_t \\ u_t \end{bmatrix}$$

and the covariance matrix of $\tilde{\boldsymbol{v}}_t$, $\mathbf{Q}$, is given by

$$\mathbf{Q} = \begin{pmatrix} \lambda_{y\eta}^2\sigma_\eta^2 + \sigma_\epsilon^2 & 0 & \lambda_{y\eta}\lambda_{c\eta}\sigma_\eta^2 + \lambda_{c\epsilon}\sigma_\epsilon^2 & 0 & \lambda_{y\eta}\sigma_\eta^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \lambda_{y\eta}\lambda_{c\eta}\sigma_\eta^2 + \lambda_{c\epsilon}\sigma_\epsilon^2 & 0 & \lambda_{c\eta}^2\sigma_\eta^2 + \lambda_{c\epsilon}\sigma_\epsilon^2 + \sigma_v^2 & 0 & \lambda_{c\eta}\sigma_\eta^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \lambda_{y\eta}\sigma_\eta^2 & 0 & \lambda_{c\eta}\sigma_\eta^2 & 0 & \sigma_\eta^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & \sigma_u^2 \end{pmatrix}$$

B IMPLIED VARIANCES

In this appendix, we first compute the variance of income and consumption growth in the BPP model and then do the same for our UC model.

In BPP, computing these variances is rather simple. They are as follows:

$$var(\Delta y_t) = \sigma_\eta^2 + \sigma_\epsilon^2(1+\theta^2-\theta) \tag{B.1}$$

since $\Delta y_t = \epsilon_t - \epsilon_{t-1} + \theta\epsilon_{t-1} - \theta\epsilon_{t-2} + \eta_t$.

Similarly,

$$var(\Delta c_t) = \gamma_\eta\sigma_\eta^2 + \gamma_\epsilon\sigma_\epsilon^2 + \sigma_u^2 + +2\sigma_v^2 \tag{B.2}$$

since $\Delta c_t = \gamma_\eta\eta_t + \gamma_\epsilon\epsilon_t + u_t + \Delta v_t$.

In our UC model income and consumption growth are given as follows:

$$\Delta y_t = \mu + \eta_t + z_t^y, \tag{B.3}$$

where $(1-\phi_{y,1}L - \phi_{y,2}L^2)z_t^y = (1-L)x_t^y$ and $x_t^y = \lambda_{y\eta}\eta_t + \epsilon_t$ and

$$\Delta c_t = \mu + \gamma_c\eta_t + z_t^c, \tag{B.4}$$

where $(1-\phi_{c,1}L - \phi_{c,2}L^2)z_t^c = (1-L)x_t^c$ and $x_t^c = \lambda_{c\eta}\eta_t + \lambda_{c\epsilon}\epsilon_t + v_t$.

We can then write a vector representation for z_t^y and z_t^c as

$$\mathbf{z}_t = \mathbf{K}\mathbf{z}_{t-1} + \mathbf{w}_t,$$

where

$$\mathbf{z}_t = \begin{bmatrix} z_t^y \\ z_{t-1}^y \\ z_t^c \\ z_{t-1}^c \\ x_t^y \\ x_t^c \end{bmatrix}, \mathbf{K} = \begin{bmatrix} \phi_{y,1} & \phi_{y,2} & 0 & 0 & -1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \phi_{c,1} & \phi_{c,2} & 0 & -1 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \mathbf{w}_t = \begin{bmatrix} x_t^y \\ 0 \\ x_t^c \\ 0 \\ x_t^y \\ x_t^c \end{bmatrix}.$$

Let $\mathbf{W}$ be the covariance matrix of $\mathbf{w}_t$, with the following non-zero entries: $\mathbf{W}[1,1] = \mathbf{W}[1,5] = \mathbf{W}[5,1] = \mathbf{W}[5,5] = \lambda_{y\eta}^2\sigma_\eta^2 + \sigma_\epsilon^2$, $\mathbf{W}[1,3] = \mathbf{W}[3,1] = \mathbf{W}[1,6] = \mathbf{W}[6,1] = \mathbf{W}[3,5] = \mathbf{W}[5,3] = \mathbf{W}[5,6] = \mathbf{W}[6,5] = \lambda_{y\eta}\lambda_{c\eta}\sigma_\eta^2$, and $\mathbf{W}[3,3] = \mathbf{W}[3,6] = \mathbf{W}[6,3] = \mathbf{W}[6,6] = \lambda_{c\eta}^2\sigma_\eta^2 + \sigma_v^2$.

Since the $vec(var(\mathbf{z}_t)) = (I - \mathbf{K} \otimes \mathbf{K})^{-1} vec(\mathbf{W})$, the unconditional variance of output growth is given by

$$var(\Delta y_t) = var(\eta_t + z_t^y)$$
$$= \sigma_\eta^2 + var(z_t^y) + 2cov(\eta_t, z_t^y)$$
$$= \sigma_\eta^2 + var(z_t^y) + 2\lambda_{y,\eta}\sigma_\eta^2$$

where $var(z_t^y)$ is the $[1,1]$ element of $var(\mathbf{z}_t)$. Similarly, unconditional variance of consumption growth is given by

$$var(\Delta c_t) = var(\gamma_c \eta_t + z_t^c)$$
$$= \gamma_\eta^2 \sigma_\eta^2 + var(z_t^c) + 2cov(\eta_t, z_t^c)$$
$$= \sigma_\eta^2 + var(z_t^c) + 2\lambda_{c,\eta}\sigma_\eta^2$$

where $var(z_t^c)$ is the $[3,3]$ element of $var(\mathbf{z}_t)$.

C Simulation results

TABLE 7. SIMULATION RESULTS

	DGP	$\gamma_\eta = 0.25$	$\gamma_\eta = 0.45$	$\gamma_\eta = 0.65$
		INCOME		
σ_η	1.40	1.42 (0.02)	1.41 (0.02)	1.39 (0.02)
σ_ϵ	1.70	1.46 (0.02)	1.50 (0.02)	1.51 (0.02)
		CONSUMPTION		
σ_u	1.30	1.51 (0.02)	1.49 (0.02)	1.48 (0.02)
σ_v	2.10	1.57 (0.02)	1.61 (0.02)	1.56 (0.02)
γ_η		0.27 (0.03)	0.48 (0.03)	0.83 (0.04)

Notes: The table reports posterior means of model parameters with posterior standard deviations reported in parentheses. In the simulated data, N=700 and T=6.

References

[1] Abowd, John M., and David Card. (1989). " On the Covariance Structure of Earnings and Hours Changes. " *Econometrica*, 57, 411-45.

[2] Altonji, Joseph G., and Aloysius Siow. (1987). " Testing the Response of Consumption to Income Changes with (Noisy) Panel Data. " *Quarterly Journal of Economics*, 102, 293-328.

[3] Attanasio, Orazio, and Steven J. Davis. (1996). " Relative Wage Movements and the Distribution of Consumption. " *Journal of Political Economy*, 104, 1227-62.

[4] Belzil, C., and M. Bognanno (2008) " Promotions, Demotions, Halo Effects, and the Earnings Dynamics of American Executives " *Journal of Labor Economics*, 26, 287-310.

[5] Blundell, Richard, and Ian Preston. (1998). " Consumption Inequality and Income Uncertainty. " *Quarterly Journal of Economics*, 113, 603-40.

[6] Blundell, Richard, Luigi Pistaferri, and Ian Preston. (2008). " Consumption Inequality and Partial Insurance. " *American Economic Review*, 98, 1887-1921.

[7] Carroll, Christopher D. (1997). " Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. " *Quarterly Journal of Economics*, 112, 1-55.

[8] Cochrane, John H. 1991. "A Simple Test of Consumption Insurance. " *Journal of Political Economy*, 99, 957-76.

[9] Deaton, Angus, and Christina Paxson. 1994. " Intertemporal Choice and Inequality. " *Journal of Political Economy*, 102, 437-67.

[10] Deaton, Angus. 1997. " The Analysis of Household Surveys: A Microeconomic Approach to Development Policy. " *Baltimore: Johns Hopkins University Press for the World Bank.*

[11] Ejrnaes, Mette, and Martin Browning. 2014. "The persistent-transitory representation for earnings processes. " *Quantitative Economics*, 5, 555-581.

[12] Jonathan Heathcote, Kjetil Storesletten, and Giovanni L. Violante. 2010. " Consumption Inequality and Partial Insurance. " *Journal of Political Economy*, 118, 681-722

[13] Jonathan Heathcote, Kjetil Storesletten, and Giovanni L. Violante. 2014. " Consumption and Labor Supply with Partial Insurance: An Analytical Framework. " *American Economic Review*, 104, 1-52

[14] Hryshko, D. (2013). " Excess smoothness of consumption in an Estimated Life-Cycle Model. " *Manuscript*

[15] Krueger, Dirk, and Fabrizio Perri. 2006. "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory." *Review of Economic Studies*, 73: 163-193.

[16] MaCurdy, Thomas E. 1982. " The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis. " *Journal of Econometrics*, 18: 83-114.

[17] Moffitt, R. A. and P. Gottschalk (2002). "Trends in the transitory variance of earnings in the United States. " *Economic Journal* 112(478):C68-C73.

[18] Morley, James C. (2007) "The Slow Adjustment of Aggregate Consumption to Permanent Income." *Journal of Money, Credit, and Banking*, 39, 615-638.

[19] Morley, James C., Charles R. Nelson and Eric Zivot. (2003) "Why are the Beveridge-Nelson and Unobserved-Component Decompositions of GDP so Different?" *Review of Economics and Statistics*, 85, 235-243.

[20] Primiceri, Giorgio E., and Thijs van Rens. (2009). " Heterogeneous Life-Cycle Profiles, Income Risk and Consumption Inequality. " *Journal of Monetary Economics*, 56, 20-39.

[21] Skinner, Jonathan. 1987.. "A Superior Measure of Consumption from the Panel Study of Income Dynamics. " *Economics Letters*, 23, 213-16.

[22] Townsend, R.(1994). "Risk and Insurance in Village India. " *Econometrica*, 62, 539-592.