THE UNIVERSITY OF SYDNEY

Economics Working Paper Series

2017 – 04

Likelihood-Based Estimates of Household Consumption Insurance

Arpita Chatterjee, James Morley

and Aarti Singh

April 2018

Likelihood-Based Estimates of Household Consumption Insurance

Arpita Chatterjee* James Morley† Aarti Singh‡

April 22, 2018§

Abstract

We propose a panel unobserved components model of household income and consumption. The model allows both idiosyncratic income and consumption to have permanent and transitory components, with possible cointegration and spillovers between components. We use likelihood-based methods for inference and find that (quasi) maximum likelihood estimation for a simple version of the model provides more precise and robust estimates of household consumption insurance against permanent income shocks than generalized method of moments for a widely-used annual panel dataset. Monte Carlo analysis demonstrates that (quasi) maximum likelihood estimation produces more accurate inferences than generalized method of moments in finite samples. Bayesian estimation allows straightforward comparison of different model specifications and favors the simple version of the model that has no cointegration or persistence in the transitory components. In subgroup analysis, we find that consumption insurance is higher for older and more educated households.

Keywords: panel unobserved components; Bayesian model comparison; permanent income; household consumption behavior

JEL codes: E21; C32; C11

*School of Economics, UNSW Sydney; arpita.chatterjee@unsw.edu.au.

†School of Economics, University of Sydney; james.morley@sydney.edu.au.

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

§We thank James Hansen, Greg Kaplan and conference and seminar participants at the Workshop of the Australasian Macroeconomics Society (Brisbane), Annual Conference on Economic Growth and Development (Delhi), Sydney Macro Reading Group Workshop and Internal Brown Bag at University of Sydney, ANU, Monash, University of Queensland, University of Melbourne, University of Technology Sydney, Continuing Education in Macroeconometrics, 2017 IAAE Conference and 2018 SNDE Symposium in Tokyo for comments. We are grateful for the financial support from the Australian Research Council grant DE130100806 (Singh). The usual disclaimers apply.

1 Introduction

"*Measurement, even without understanding of mechanisms, can be of great importance in and of itself – policy change is frequently based on it – and is necessary if not sufficient for any reasoned assessment of policies...*" - Angus Deaton (Measuring and Understanding Behavior, Welfare, and Poverty. Nobel Prize Lecture, December 8, 2015)

How does idiosyncratic income risk impact consumption when households have limited access to insurance via formal markets or informal arrangements? Starting with Aiyagari (1994) and Huggett (1997) and continuing more recently with Kaplan and Violante (2010), the academic literature has keenly focused on this question to inform incomplete-markets modeling and to prescribe policies given an implied extent of market incompleteness.[1] Such analysis requires reliable measures of household income risk and the degree of consumption insurance against this risk. The purpose of this paper is to apply likelihood-based inference, which is often used in the time series literature with aggregate data, to estimate income risk and consumption insurance from household data.

Numerous studies have examined the fraction of income risk that gets transmitted to consumption.[2] For example, Mace (1991), Cochrane (1991), and Townsend (1994) consider cross-sectional regressions, motivated and interpreted in part by theory, to measure household consumption insurance. Deaton (1997) proposes measuring consumption insurance using panel methods that quantify the overall degree of risk sharing in the economy, while remaining agnostic about the exact mechanisms behind it. In this vein, Blundell, Pistaferri, and Preston (2008) (BPP hereafter) construct a novel annual panel dataset of household income and (imputed) consumption for the Panel Study of Income Dynamics (PSID) and employ generalized method of moments (GMM) to estimate income risk and consumption insurance without imposing a particular theory-based structure.

We propose an alternative approach based on likelihood-based methods for estimation of a panel unobserved components (UC) model of household income and consumption in order to measure consumption insurance. To our knowledge, likelihood-based methods have not been previously employed in this context, likely due to the short time dimension of panel datasets frequently used in the literature.[3] We are able to overcome this limiting feature of the data and successfully estimate the panel UC model by assum-

[1] See Blundell, Pistaferri, and Saporta-Eksten (2016) and Heathcote, Storesletten, and Violante (2014), among others, that examine the role of different insurance mechanisms such as labor supply and progressive taxation in providing insurance against wage shocks.

[2] See Jappelli and Pistaferri (2011) for an excellent survey of this literature.

[3] Nakata and Tonetti (2016) propose likelihood-based estimation of labor income processes and show via Monte Carlo analysis that Bayesian estimates of income risk and other parameters have favorable finite-sample properties. However, they do not apply their methods to actual data or consider a multivariate environment necessary to measure consumption insurance.

ing a common distribution and independent idiosyncratic shocks across households.[4]

Our panel UC model allows both idiosyncratic income and consumption to have permanent and transitory components, with possible cointegration and spillovers between components. We find that (quasi) maximum likelihood estimation ((Q)MLE) for a simple version of the model with no cointegration or persistence in the transitory components produces more precise and robust estimates of consumption insurance than GMM for the BPP dataset. Comparing with the results in BPP, our estimates for income risk are similar, but our estimates of consumption insurance are considerably higher. Our preferred estimate implies that 58 percent of an permanent shock to idiosyncratic family disposable labor income do not pass through to household consumption, with 95% confidence bands of 55-61 percent. The estimate in BPP is only 36 percent and is much less precise.[5] Meanwhile, subgroup estimates of consumption insurance based on (Q)MLE are also higher and more precise than those based on GMM in BPP and reveal a highly intuitive pattern of heterogeneity, with higher consumption insurance for older and more educated households that likely have more buffer-stock wealth and possibly also have more margins on which to adjust their labor supply when faced with permanent income shocks.[6]

Monte Carlo analysis demonstrates that (Q)MLE produces more accurate inferences than GMM in finite samples. The performance of GMM depends on the weighting scheme, with equal weighting outperforming optimal weighting or diagonal weighting, the latter of which was used by BPP. Notably, however, (Q)MLE outperforms all versions of GMM in terms of having lower root mean squared error when considering both Normal shocks and empirically-distributed shocks with a similar extreme leptokurtosis as appears to be present in the data. The relative performance of the likelihood-based approach is even stronger in the case of empirical shocks and is particularly evident given a comparatively small effective sample size which occurs when allowing for time variation in some parameters or with subgroup analysis.

A likelihood-based approach also makes it straightforward to compare different specifications of our model using Bayesian methods. In particular, we calculate marginal likelihoods to consider spillovers between different components of consumption and income, persistence in transitory income and consumption, and possible cointegration between income and consumption. Bayesian model comparison strongly supports the simple version of our model used in our (Q)MLE analysis that only has a spillover from permanent

[4]The Kalman filtering used in our approach also conveniently accounts for the many missing observations in the panel dataset.

[5]Our preferred estimate is based on a different model than in BPP. However, our (Q)MLE estimate for the BPP model is 54 percent and is also relatively precise.

[6]Older and more educated households also likely face higher marginal tax rates, which would mitigate the pass-through of permanent income shocks to consumption because of proportionately smaller effects on disposable income. Whether they have greater access to social assistance schemes or to informal means of insurance, such as relying on transfers from relatives, is less obvious.

income to permanent consumption, but no cointegration and no persistence in transitory components. However, we find that consumption insurance estimates are quite robust across the various different specifications under consideration.

Having a flexible empirical model, including for the income process, is particularly important for our analysis because accurate measurement of consumption insurance involves estimating permanent income risk, which has been shown to be sensitive to model misspecification. Following Friedman and Kuznets (1945), household income is typically assumed to have a random walk permanent component, a transitory component that dies away, and zero correlation between movements in the two components.[7] However, it is straightforward to show that, if the zero correlation assumption is incorrect, the model misspecification will bias the estimate of permanent income risk, a key ingredient in heterogeneous agent quantitative macro models. At the same time, Ejrnaes and Browning (2014) show that, without assumptions such as a zero correlation, the decomposition of shocks into persistent and transitory components is indeterminate and illustrate its practical importance in PSID data.[8] Domeij and Floden (2010) show that the estimation of permanent income risk is subject to significant bias if serial correlation in transitory shocks is ignored. Motivated by these studies, our general model specification allows for correlated movements in unobserved components of income and consumption, with random walk permanent components and possibly persistent dynamics for the transitory components. We find that estimation of this general version of the model is straightforward with a Bayesian approach, although many parameters may be weakly identified in practice given the empirical support from Bayesian model comparison for the simple version of the model without all of the additional features.

The rest of this paper is organized as follows: Section 2 presents the panel UC model proposed in this paper. Section 3 describes the data. Section 4 reports empirical results and includes our Monte Carlo analysis. Section 5 concludes.

2 A Panel Unobserved Components Model

In this section, we present the details of our panel UC model of household income and consumption measured as residuals from regressions of household income and consump-

[7]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.

[8]In the time series literature using aggregate U.S. quarterly real GDP data, Morley, Nelson, and Zivot (2003) establish that the assumption of zero correlation between permanent and transitory movements can be rejected in the univariate case, while Morley (2007) finds evidence in favor of correlated movements using U.S. quarterly real GDP and aggregate consumption data in a multivariate unobserved components model. Note, however, Morley (2007) considers total income for the aggregate data, not just idiosyncratic labor income, as is considered for household data in this paper.

tion on standard observed factors.[9] We also represent the empirical model used in BPP in a similar form in order to better understand how it compares with our proposed model.

2.1 General model specification

Our panel UC model decomposes idiosyncratic income and consumption for household i into permanent components and transitory deviations from the permanent components:

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

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

The permanent components are specified as random walks with possible drift:

$$\tau_{i,t} = \mu_{\tau,i} + \tau_{i,t-1} + \eta_{i,t}, \qquad \eta_{i,t} \sim i.i.d.N(0, \sigma_\eta^2), \tag{3}$$

$$\kappa_{i,t} = \mu_{\kappa,i} + \kappa_{i,t-1} + u_{i,t}, \qquad u_{i,t} \sim i.i.d.N(0, \sigma_u^2), \tag{4}$$

where for household i, the common stochastic trend of income and consumption is $\tau_{i,t}$ and $\kappa_{i,t}$ is an additional trend for consumption. In our specification, γ_η captures the impact of permanent income shocks on permanent consumption.

The transitory components are specified as ARMA(p,q) processes:

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

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

where $\phi_j(L) = (1 - \phi_{j,1}L - \phi_{j,2}L^2 - ... - \phi_{j,p}L^p)^{-1}$ and $\theta_j(L) = (1 - \theta_{j,1}L - \theta_{j,2}L^2 - ... - \theta_{j,q}L^q)^{-1}$ for $j = \{y, c\}$ are lag polynomials that satisfy stationarity and invertibility constraints, respectively.

The permanent income shock, $\eta_{i,t}$, can be interpreted as reflecting severe health shocks, promotion, or other idiosyncratic factors that result in a change in idiosyncratic permanent income. Other permanent shocks to consumption, $u_{i,t}$, beyond permanent shocks to income could be taste and preference shocks or other shocks to non-labor income, such as wealth shocks. The transitory income shock is $\epsilon_{i,t} \sim i.i.d.N(0, \sigma_\epsilon^2)$, which could capture events such as a surprise bonus or temporary leave due illness or other events, while the transitory consumption shock is $v_{i,t} \sim i.i.d.N(0, \sigma_v^2)$, which could capture measurement error due to the imputation of nondurable consumption. In some specifications, we will allow for time variation in the variance of shocks, as was done in BPP.

[9] In particular, in the BPP dataset described in Section 3, idiosyncratic income and consumption for households are calculated by 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, and presence of income recipients other than husband and wife.

Instead of directly specifying shocks to be correlated across equations, as in Morley, Nelson, and Zivot (2003) and Morley (2007), we assume that shocks are orthogonal, but permanent shocks can affect the transitory components according to impact coefficients $\lambda_{y\eta}$ and $\lambda_{c\eta}$. Thus, permanent and transitory movements can still be correlated, as in Morley, Nelson, and Zivot (2003) and Morley (2007). However, following Morley and Singh (2016), we explicitly model the source of this correlation as due to the effects of permanent shocks on transitory components. For example, we might expect $\lambda_{y\eta}$ to be positive if there are (even partially) delayed labor supply responses to permanent income shocks, meaning some of the initial movement in income following a shock is offset at longer horizons. Meanwhile, $\lambda_{c\epsilon}$ captures the response of consumption to transitory income shocks. For simplicity, we assume no corresponding effect of transitory consumption shocks on income.

We assume that income and consumption shocks are drawn from a Normal distribution in order to make it straightforward to derive the likelihood function used for estimation via MLE.[10] Notably, instead of relying only on matching particular moments (e.g., autocovariances, but not higher moments) to estimate parameters, the likelihood-based approach implicitly makes use of all moments in the data implied by the parametric model. A clear benefit of this approach is that it addresses possible extreme sensitivity of inferences to particular moments. For example, in the idiosyncratic income/wage risk literature Heathcote, Perri, and Violante (2010) find that the estimates of the variance of the wage shocks are different whether one uses moment conditions based on log residual wage growth or moment conditions based on log residual wage level.[11] It also allows us to avoid having to specify a weighting matrix for different moments. However, for completeness, we compare our MLE estimates to GMM estimates based on the full set of second moments, as in BPP, and different weighting schemes.[12] Meanwhile, to the extent that the actual shocks are non-Normal, as suggested in some recent literature (see for example Guvenen, Karahan, Ozcan, and Song, 2015), our estimation approach can be thought of as QMLE.[13] We will consider how well QMLE works in practice relative to GMM in the Monte Carlo analysis in Section 4.

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

[10]See Appendix A for more details on estimation.

[11]This inconsistency between the estimates has been reported by Brzozowski, Gervais, Klien and Suzuki (2010) for Canada; Fuchs-Schundeln, Krueger, and Sommer (2010) for Germany, Domeij, and Floden (2010) for Sweden and Chatterjee, Singh, and Stone (2016) for Australia. See Daly, Hryshko, and Manovskii (2016) who study this issue more closely.

[12]We follow BPP's estimation approach and details can be found in their paper and appendix.

[13]Some recent studies, for example Madera (2017), address possible non-Normality in studying the joint distribution of income and durable and non-durable consumption by examining tails of the distributions.

as follows:

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

which suggests that changes in consumption depend on the full history of permanent shocks to income and transitory shocks to consumption and income.[14] Then to calculate the implied consumption insurance based on our 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, denoted ϑ_c, is

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

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

$$\vartheta_c = 1 - \frac{cov(\Delta c_{i,t}, \eta_{i,t})}{var(\eta_{i,t})}, \tag{9}$$

which can be shown to be equal to $1 - (\gamma_\eta + \lambda_{c\eta})$ for our panel UC model.[15]

2.2 BPP model

We put the BPP model in a similar form to our panel UC model for the purpose of comparison.[16] In particular, the BPP model has an implicit UC representation for income:

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

The permanent component of income is specified as follows:

$$\tau_{i,t} = \mu_i + \tau_{i,t-1} + \eta_{i,t}, \qquad \eta_{i,t} \sim i.i.d.N(0, \sigma_\eta^2). \tag{11}$$

The transitory component has a moving average, in particular, an $MA(1)$, specification as follows:

$$(y_{i,t} - \tau_{i,t}) = \epsilon_{i,t} + \theta\epsilon_{i,t-1}, \qquad \epsilon_{i,t} \sim i.i.d.N(0, \sigma_\epsilon^2). \tag{12}$$

[14] Commault (2017) examines why the estimates of pass-through from transitory income shocks to consumption using structural models are smaller than estimates using natural experiments. She finds that the main reason is that structural estimation often ignores the impact of past shocks on consumption growth. However, through the possibility of persistent transitory dynamics for household consumption in the general specification of our panel UC model, we explicitly allow for the possibility of persistent effects of past shocks. Following the terminology in Kaplan and Violante (2010), this means that there is no "short memory" in our model. However, in practice, we do not find empirical support for such dynamics beyond a one-year horizon.

[15] See Appendix A for the state-space representation of the panel UC-AR(2) model and the implied variances that allow us to see that the analytical expression for equation (9) is the same as equation (8).

[16] The state-space representation and implied variances for the BPP model are also given in Appendix A.

Meanwhile, consumption growth is given by the following process:

$$\Delta c_{i,t} = \gamma_\eta \eta_{i,t} + \gamma_\epsilon \epsilon_{i,t} + u_{i,t} + \Delta v_{i,t}, \qquad u_{i,t} \sim i.i.d.N(0, \sigma_u^2), \tag{13}$$

where $\eta_{i,t}$ and $\epsilon_{i,t}$ are the permanent and transitory income shocks, $u_{i,t}$ is the permanent shock to consumption, and $v_{i,t} \sim i.i.d.N(0, \sigma_v^2)$ is measurement error for consumption.

Our panel UC model differs from the BPP specification in one key way. In particular, following UC models for the aggregate data, our model assumes that transitory income shocks are small enough to only significantly impact transitory consumption, while in the BPP model, transitory income shocks are assumed to have only a permanent impact on consumption. To see this, we can rewrite the level of consumption, after suppressing the individual specific subscript for simplicity, as

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

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

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

To the extent transitory income shocks are quite large, such as might be thought about winning a lottery if it were classified as income, there may be nontrivial permanent effects on consumption. However, we treat it as an empirical issue as to how large transitory income shocks are in practice and which model specification is preferred. We note in advance that our estimates suggest transitory income shocks are similar in magnitude to permanent income shocks and do not persist much beyond a one-year horizon, suggesting a household that follows the permanent income hypothesis and has nontrivial life expectancy would barely adjust its consumption to such a shock.

3 Data

In this section, we briefly describe the novel dataset constructed by BPP and look at sample autocorrelations in idiosyncratic income and consumption growth to help motivate our model specification. For full details of the dataset, we refer readers to BPP.

3.1 BPP dataset

BPP use the Panel Study of Income Dynamics (PSID) sample from 1978-1992 of continuously married couples headed by a male (with or without children) aged 30 to 65. The income variable is family disposable income, which includes transfers. They adopt a similar sample selection in the Consumer Expenditure Survey (CEX). Since CEX has detailed nondurable consumption data, unlike PSID which primarily has food expenditure data,

TABLE 1. SAMPLE ACF AND PACF

$\hat{\rho}_1$	$\hat{\rho}_2$	$\hat{\rho}_3$	$\hat{\phi}_1^{p=1}$	$\hat{\phi}_2^{p=2}$	$\hat{\phi}_3^{p=3}$
I. Idiosyncratic Income Growth					
−0.31	−0.05	−0.01	−0.31	−0.16	−0.08
II. Idiosyncratic Consumption Growth					
−0.45	0.00	0.03	−0.45	−0.25	−0.09

Notes: Sample autocorrelations ($\hat{\rho}_j$'s) and partial autocorrelations ($\hat{\phi}_j$'s) are calculated by pooling growth rate data for up to 14 years for 1,765 households.

they impute nondurable consumption for each household per year by using the estimates of the food demand from CEX. The constructed dataset is a panel of income and imputed nondurable consumption. To get idiosyncratic (residual) income and consumption, BPP regress income and consumption for households on a vector of regressors including demographic and ethnic factors and other income characteristics observable/known by consumers. It is the residual idiosyncratic income and consumption that we model.

3.2 Sample autocorrelations

To help motivate our model specification, we compute the sample autocorrelation function (ACF) and partial autocorrelation function (PACF) for idiosyncratic income and consumption growth from the dataset by pooling individuals of all ages and over all years. Table 1 reports the results.

Based on the sample ACFs and PACFs, we can see that the autocorrelations completely die off after 1 lag, but the partial autocorrelations die off more gradually for both income growth and consumption growth. This pattern is consistent with MA(1) processes, not an MA(2) process, as would be implied for income growth by the BPP model. Moreover, this pattern is suggestive of a basic specification for the general panel UC model in Section 2. In particular, it is consistent with a simple version of the model in which both income and consumption follow random walk permanent components plus noise for their respective transitory components. We start with this specification as a benchmark, but also conduct formal model comparison to determine the preferred specification in practice.

4 Empirical Results

In this section, we present our empirical results. First, we compare our (Q)MLE estimates to those for GMM for a simple version of the panel UC model. Second, we investigate the finite-sample properties of the competing estimators using Monte Carlo analysis, including with empirically-distributed shocks, to help evaluate the differences in estimates across different methods. Third, we use Bayesian methods to consider the robustness of findings about consumption insurance to different model specifications and to determine the preferred specification. Fourth, we report (Q)MLE estimates for subgroups based on age and education.

4.1 Estimates for a simple version of the model

Motivated by the sample ACFs and PACFs, we estimate a simple version of the model in Section 2 using maximum likelihood (or quasi maximum likelihood if we assume Normality does not actually hold in the data). For the simple version of the model, which we refer to as the UC-WN model hereafter, the transitory components are assumed to have no persistence (i.e., the ϕ's and θ's in the general model in Section 2 are set to zero) and, for simplicity, we only consider a spillover from permanent income to permanent consumption, as captured by γ_η (i.e., the λ's in the general model are set to zero). This model structure is consistent with the apparent autocorrelation structure for income and consumption growth noted in Section 3. Meanwhile, despite having only a short time series for each individual in the sample, with a maximum time dimension of T=15, (Q)MLE is feasible for this model because each two-year block of income and consumption growth for a given household is essentially treated as an independent draw from the same joint distribution, making the effective sample size of independent observations approximately $(T-1)N/2$.

Table 2 reports estimates for the two key parameters of interest that determine income risk and consumption insurance for the UC-WN model: σ_η and γ_η.[17] Based on (Q)MLE, the variance of the permanent income shock, σ_η^2 is 0.02 (=0.13^2). This is similar to the estimates in BPP and is also close to what one finds in the related idiosyncratic income risk literature. However, what is notable is that, using the same dataset as BPP, but based on (Q)MLE, our implied estimate of consumption insurance, $\vartheta_c = 1 - \gamma_\eta$, is 58 percent, while the highlighted estimate in BPP using their preferred model specification is only 36 percent.

To address why our estimate of consumption insurance is so different, we also consider GMM estimation of the UC-WN model to see what role estimation method plays

[17] The estimates for the full set of model parameters are provided in Appendix B.

TABLE 2. ESTIMATES FOR THE UC-WN MODEL

Parameter	(Q)MLE	GMM		
		DWMD	EWMD	OMD
σ_η	0.125	0.179	0.177	0.123
	(0.002)	(0.002)	(0.002)	(0.001)
	[0.122, 0.129]	[0.175, 0.183]	[0.173, 0.181]	[0.121, 0.125]
γ_η	0.421	0.376	0.416	0.316
	(0.017)	(0.052)	(0.058)	(0.029)
	[0.389, 0.449]	[0.272, 0.480]	[0.300, 0.532]	[0.258, 0.374]

Notes: The table reports point estimates, with standard deviations in parentheses and 95% confidence intervals in square brackets. (Q)MLE confidence sets based on inverted likelihood ratio tests end up as intervals. GMM confidence intervals are based on inverted t tests.

versus model specification. We report GMM estimates for the UC-WN model for three different weighting schemes for the second moments: diagonal (as used in BPP), equal, and optimal. There are a few things to notice about the GMM results. First, estimates of income risk are somewhat higher for GMM with diagonal and equal weighting, but is very similar to (Q)MLE for GMM with optimal weighting. Second, estimates for consumption insurance are similar to (Q)MLE, especially for diagonal and equal weighting, and much higher than in BPP. Thus, it appears that model specification could be playing an important role in our different estimates. Third, it is notable that the (Q)MLE estimate of consumption insurance is much more precise than the GMM estimates, especially when comparing to diagonal or equal weighting.

One major difference in BPP from our basic UC-WN model is that they allow for time-varying variances for permanent income shocks, transitory income shocks, and transitory consumption shocks. To see what role this plays, we estimate a UC-WN-TVV model that allows for time-varying variances, similar to BPP. Table 3 reports the estimates for the two key parameters for this UC-WN-TVV model (in this case, reporting the average $\bar{\sigma}_\eta$ for easy comparison with the results in Table 2). The estimates of the average level of income risk are essentially unchanged. However, what is striking is that the implied estimates of consumption insurance are basically unaffected for (Q)MLE and GMM with optimal weighting, but are noticeably lower for GMM with diagonal or equal weighting. In the case of diagonal weighting, which was used in BPP, the implied consumption insurance is $\hat{\vartheta}_c = 49$ percent, which gets us about half the way down to their estimate of 36 percent. That is, allowing for time-varying variances appears to be important for the moments-based estimators of consumption insurance, except when using optimal

TABLE 3. ESTIMATES FOR THE UC-WN-TVV MODEL

Parameter	(Q)MLE	GMM		
		DWMD	EWMD	OMD
$\overline{\sigma}_\eta$	0.127	0.175	0.178	0.131
γ_η	0.418	0.513	0.465	0.314
	(0.016)	(0.067)	(0.068)	(0.029)
	[0.387, 0.447]	[0.379, 0.647]	[0.329, 0.601]	[0.256, 0.372]

Notes: The table reports point estimates, with standard deviations in parentheses and 95% confidence intervals in square brackets. (Q)MLE confidence sets based on inverted likelihood ratio tests end up as intervals. GMM confidence intervals are based on inverted t tests.

TABLE 4. ESTIMATES FOR THE UC-BPP-TVV MODEL

Parameter	MLE	GMM		
		DWMD	EWMD	OMD
$\overline{\sigma}_\eta$	0.114	0.145	0.149	0.121
γ_η	0.455	0.642	0.539	0.264
	(0.020)	(0.094)	(0.099)	(0.035)
	[0.416, 0.494]	[0.454, 0.830]	[0.341, 0.737]	[0.194, 0.334]

Notes: The table reports point estimates, with standard deviations in parentheses and 95% confidence intervals in square brackets. (Q)MLE confidence sets based on inverted likelihood ratio tests end up as intervals. GMM confidence intervals are based on inverted t tests.

weighting. Again, the (Q)MLE estimates are more robust and much more precise than the GMM estimates.

Table 4 completes the comparison with BPP by reporting estimates of the two key parameters for the UC-BPP-TVV model presented in Section 2 that allows for time-varying variances. The estimates for income risk are all shifted a bit lower compared to the UC-WN and UC-WN-TVV models, but the pattern across estimators is the same as before. The key result is that implied estimates of consumption insurance are fairly similar in level and precision for (Q)MLE, but much more varied and less precise for GMM. In the case of GMM with diagonal weighting, we get the original BPP result of consumption insurance of $\hat{\vartheta}_c = 36$ percent, but with a large standard error. The estimate for GMM with equal weights of $\hat{\vartheta}_c = 46$ percent is closest to the (Q)MLE estimate of $\hat{\vartheta}_c = 54$ percent, but is also much less precise. The estimate for GMM with optimal weights is $\hat{\vartheta}_c = 74$ percent. However, the literature has produced a number of concerns about biases for GMM with optimal weights in small samples (see Altonji and Segal, 1996), concerns that we will see are well justified based on our Monte Carlo analysis.

The main takeaways from these results for the BPP dataset are that (Q)MLE provides precise and reasonably robust estimates of consumption insurance at in the 54-58 percent range, while GMM estimates are much less precise and sensitive to model specification and weighting scheme.[18] A reasonable question, though, is whether the precision of (Q)MLE is in some sense "false", perhaps due to an incorrect assumption of Normality. Another reasonable question is whether these model specifications are supported by the data and whether estimates of consumption insurance are sensitive to alternative specifications. We address these two questions sequentially in the next two subsections.

4.2 Monte Carlo analysis

To investigate whether (Q)MLE estimates are indeed more precise, we consider Monte Carlo analysis. In particular, we simulate data from a UC-WN data generating process (DGP) and compare (Q)MLE and GMM estimators for σ_η and γ_η. We make two assumptions about shocks: i) Normal distributions and ii) empirical distributions based on the (Q)MLE estimated UC-WN-TVV model for the BPP dataset. We also consider two sample sizes: i) "small" (TN=1,050) and ii) "large" (TN=15,000). We use root mean squared error (RMSE) across 5,000 simulations to evaluate the accuracy of an estimator, although we also consider the empirical coverage rate of 95% confidence sets/intervals (CIs) based on

[18]We also verify that these key empirical results are not driven by the imputed non-durable consumption data by estimating these various different specifications on income and food consumption data directly from PSID. All results are available upon request.

TABLE 5. ESTIMATORS IN LARGE SAMPLE WITH NORMAL SHOCKS

Parameter	Estimator Property	MLE	GMM		
			DWMD	EWMD	OMD
$\sigma_\eta = 0.2$	Mean estimate	0.200	0.200	0.200	0.194
	Mean standard error	0.002	0.003	0.003	0.002
	RMSE	0.002	0.002	0.002	0.006
	95 % CI coverage	0.948	0.947	0.950	0.190
$\gamma_\eta = 0.65$	Mean estimate	0.655	0.649	0.650	0.658
	Mean standard error	0.021	0.027	0.027	0.017
	RMSE	0.017	0.022	0.022	0.025
	95 % CI coverage	0.943	0.951	0.953	0.715

Notes: The table reports averages (and frequency of coverage in case of the 95% CI coverage) across 5,000 simulations, where the sample size is T*N=15,000 and shocks are Normally distributed.

inverted likelihood ratio (ILR) tests for (Q)MLE and inverted t tests for GMM.[19]

Table 5 reports the Monte Carlo results for the case of a large sample with Normal shocks. As might be expected given a large sample, the various estimators perform similarly, with very little bias and comparable standard errors. The RMSE for σ_η is the same for MLE and GMM with diagonal or equal weighting. It is somewhat higher for GMM with optimal weighting. Related, the coverage of the CIs is quite accurate for all but GMM with optimal weighting, which severely undercovers the true value of the parameter, despite being less precise. In terms of γ_η, the results are somewhat similar, with MLE and GMM with diagonal or equal weighting performing well in terms of RMSE and CI coverage, but GMM with optimal weighting less so. MLE seems to show a slight upward bias, but it is also more precise in terms of standard errors and the RMSE is slightly lower than for the other cases. GMM with optimal weighting also show a slight upward bias and appears precise, but RMSE reveals a lower accuracy compared to the other estimators. The key point, though, is that the various estimators, which should all behave reasonably well in large samples, do so, with the possible exception of GMM with optimal weighting.

Table 6 reports the Monte Carlo results for the case of a small sample with Normal shocks. Compared to the large sample case, there are more differences across the MLE and GMM estimators. In particular, MLE clearly performs best in terms of RMSE for both

[19]Nakata and Tonetti (2016) conduct Monte Carlo analysis to evaluate the RMSE of their likelihood-based Bayesian estimators of income risk and find that they perform well. In some cases, they also allow for mixtures-of-Normals for shocks to capture leptokurtosis, although doing so complicates estimation considerably. We leave such extensions of model specification for future research.

TABLE 6. ESTIMATORS IN SMALL SAMPLE WITH NORMAL SHOCKS

Parameter	Estimator Property	MLE	GMM		
			DWMD	EWMD	OMD
$\sigma_\eta = 0.2$	Mean estimate	0.200	0.195	0.200	-
	Mean standard error	0.009	0.010	0.010	-
	RMSE	0.007	0.009	0.008	-
	95 % CI coverage	0.950	0.908	0.942	-
$\gamma_\eta = 0.65$	Mean estimate	0.656	0.634	0.654	-
	Mean standard error	0.081	0.104	0.104	-
	RMSE	0.063	0.088	0.083	-
	95 % CI coverage	0.953	0.922	0.948	-

Notes: The table reports averages (and frequency of coverage in case of the 95% CI coverage) across 5,000 simulations, where the sample size is T*N=1,050 and shocks are Normally distributed.

σ_η and γ_η. It is relatively precise in terms of standard errors and also has accurate CI coverage. GMM with equal weighting performs next best with a similar slight upward bias for γ_η and similar CI coverage to MLE, but less precision and higher RMSE. GMM with diagonal weighting, the approach used in BPP, has more noticeable (downward) bias for both parameters and its CIs undercover the true parameters. Meanwhile, GMM with optimal weighting often did not converge given a small sample, so we do not report results for it. In any event, its relatively poor performance even in the large sample case casts doubt on the empirical results for it in Tables 2-4 (again, see Altonji and Segal, 1996, on the shortcomings of GMM with optimal weighting).

A reasonable supposition is that MLE outperforms GMM in Table 6 because it imposes Normality, which is the correct DGP, while GMM is, by its nature, more general from a distributional standpoint, albeit possibly at the cost of small sample efficiency. We address this by considering a DGP with shocks drawn from their empirical distributions. In particular, we use the Kalman filtered estimates of shocks from the (Q)MLE estimated UC-WN-TVV model for the BPP data. The shocks are standardized using the estimated time-varying variances and then multiplied by 0.2 to have a variance of 0.04. We then draw the various shocks from their empirical distributions with replacement. This approach is designed to capture the extreme leptokurtosis of the data (and estimated shocks), which seems immune to controlling for the time-varying variances and remains present even when just considering more homogeneous subgroups. In the BPP dataset, the sample kurtosis is 11.85 and 14.75 (vs. 3 under Normality) for idiosyncratic income and consumption growth, respectively. For a simulated UC-WN process using shocks

TABLE 7. ESTIMATORS IN LARGE SAMPLE WITH EMPIRICAL SHOCKS

Parameter	Estimator Property	QMLE	GMM		
			DWMD	EWMD	OMD
$\sigma_\eta = 0.2$	Mean estimate	0.200	0.197	0.200	0.182
	Mean standard error	0.002	0.003	0.004	0.002
	RMSE	0.003	0.004	0.003	0.018
	95 % CI coverage	0.807	0.871	0.950	0.001
$\gamma_\eta = 0.65$	Mean estimate	0.655	0.649	0.650	0.659
	Mean standard error	0.021	0.029	0.027	0.017
	RMSE	0.017	0.024	0.022	0.026
	95 % CI coverage	0.933	0.938	0.944	0.702

Notes: The table reports averages (and frequency of coverage in case of the 95% CI coverage) across 5,000 simulations, where the sample size is T*N=15,000 and shocks are are drawn from empirical distributions, with replacement.

from their estimated empirical distribution, the population kurtosis based on a large simulation (TN=75,000) is 11.08 and 12.66 for idiosyncratic income and consumption growth, respectively. Thus, our approach of using a DGP with empirically distributed shocks seems to produce simulated data with non-Normalities of the scale found in the actual data. Then, applying MLE based on the incorrect assumption (in this case) of Normal shocks corresponds to QMLE, which often performs well at estimating conditional mean and variance parameters (see White, 1982). Our Monte Carlo analysis investigates how well QMLE performs in the particular case of the UC-WN model.

Table 7 reports the Monte Carlo results for the case of a large sample with empirically-distributed shocks. In this case, we can see more differences across estimators even given the large sample. Intuitively, the strong deviations from Normality imply that asymptotic approximations require larger samples to fully apply. In terms of σ_η, QMLE and GMM with equal weighting are unbiased, with QMLE appearing twice as precise. However, the RMSEs are the same and GMM with equal weighting performs best in terms of CI coverage. GMM with diagonal weighting performs less well in terms of accuracy, while GMM with optimal weighting does terribly, with much higher RMSE and CIs that almost never correctly cover the true parameter value. In terms of γ_η, QMLE is best in terms of accuracy based on RMSE. It has a slight upward bias and its CIs slightly undercover. Meanwhile, GMM with equal weighting is next best in terms of RMSE and again performs best in terms of CI coverage. GMM with diagonal weighting is almost as good, while GMM with optimal weighting again performs the worst in terms of accuracy and CI coverage. Notably, though, the estimators of the slope coefficient γ_η that is directly

TABLE 8. ESTIMATORS IN SMALL SAMPLE WITH EMPIRICAL SHOCKS

Parameter	Estimator Property	QMLE	GMM		
			DWMD	EWMD	OMD
$\sigma_\eta = 0.2$	Mean estimate	0.199	0.181	0.199	-
	Mean standard error	0.009	0.011	0.013	-
	RMSE	0.010	0.020	0.011	-
	95 % CI coverage	0.822	0.598	0.933	-
$\gamma_\eta = 0.65$	Mean estimate	0.657	0.651	0.656	-
	Mean standard error	0.081	0.117	0.104	-
	RMSE	0.065	0.102	0.084	-
	95 % CI coverage	0.944	0.923	0.945	-

Notes: The table reports averages (and frequency of coverage in case of the 95% CI coverage) across 5,000 simulations, where the sample size is T*N=1,050 and shocks are are drawn from empirical distributions, with replacement.

informative about consumption insurance generally seem more robust to the strong deviations from Normality than the estimators of the volatility parameter σ_η. This is reasonably intuitive since a strong deviation from Normality, with a direct misspecification in the case of QMLE, is likely to make it harder to estimate the variance of a distribution, but presumably only affects the efficiency of estimates of a slope coefficient for its conditional mean.[20]

Table 8 reports the Monte Carlo results for the case of a small sample with empirically-distributed shocks. As with the Normal shocks, the benefits of QMLE in terms of estimating γ_η in a small sample are more evident than in the large sample case. In particular, the RMSE is considerably lower for QMLE than GMM with diagonal or equal weighting. Again, GMM with optimal weighting often did not converge given a small sample, so we do not report results for it. The precision in terms of standard errors is also better for QMLE, while the CI coverage is as accurate as for GMM with equal weighting and more accurate than GMM with diagonal weighting. All three estimators have a slight upward bias. Meanwhile, similar to the large sample results with empirically-distributed shocks, QMLE and GMM with equal weighting perform similarly in terms of RMSE for σ_η, while GMM with diagonal weighting performs much worse. The precision in terms of standard errors for QMLE is best, although the CI coverage for QMLE is low, but not as low as for

[20]One possibility would be to consider bootstrapping the CIs for QMLE when there are apparent strong deviations from Normality. In particular, it is possible to bootstrap critical values of likelihood ratio test for the estimated parameter values, which would likely make the ILR confidence sets a bit more conservative and have coverage closer to the 95% nominal level. However, given the computational burden involved in running a bootstrap within a Monte Carlo, we leave this for future research.

GMM with diagonal weighting.

Reflecting back on the estimates for the BPP dataset in Tables 2-4 given the Monte Carlo results, we take the higher precision of the (Q)MLE estimates of consumption insurance as corresponding to a desirable property of the estimator that is robust to deviations from Normality. Indeed, we find that the relative performance of QMLE versus GMM in small samples is even better given empirically-distributed shocks, which is likely the relevant case for our application to the BPP dataset. Also, while we might worry that the (Q)MLE CI for income risk in Tables 2-4 undercover the true parameter value, we note that (Q)MLE still has the lowest RMSE in terms of σ_η across all of the Monte Carlo experiments, so the (Q)MLE estimates of income risk are still likely to be most accurate, even if we need to apply caution when looking at the CIs for income risk.

4.3 Bayesian analysis

Some of the literature on earnings has moved away from a simple model in which the permanent component is a random walk and the transitory component is white noise (i.e., the basic UC-WN model). It is often believed that the earnings dynamics are more complicated. For example, MaCurdy (1982) and Abowd and Card (1982) find that the covariance matrix of earnings differences fits an MA(2), Gottschalk and Moffitt (1994) fit random walk plus ARMA(1,1) in levels which is an ARMA(1,2) in first differences, and Heathcote, Storesletten, and Violante (2010) employ a very persistent "permanent" component and a white noise transitory component. We focus on two main specifications of our general panel UC model discussed in Section 2, a UC-AR(2) model and the basic UC-WN model to encompass the main differences in views held in the literature.[21] Motivated by the findings for persistent autoregressive dynamics in the aggregate data found in the time series literature, we investigate whether dynamics play an important role in the household data and whether income and consumption are cointegrated, as they appear to be in the aggregate data. We also consider the BPP model presented in Section 2.

For estimation of more general specifications and model comparison, we consider Bayesian methods. In particular, we estimate the panel UC models using Bayesian posterior simulation based on Markov-chain Monte Carlo (MCMC) methods. We use a 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,

[21] Following Morley, Nelson, and Zivot (2003), a UC-AR(2) model with correlated permanent and transitory movements is identified because $p = q + 2$ for an implied ARMA(p,q) process in first differences. Trenkler and Weber (2016) show that a multivariate UC model, such as ours, is actually identified given $p = 1$ and $q = 0$. However, we found a UC-AR(1) model was more weakly identified, thus we stick with $p = 2$ and $q = 0$.

we use different starting values.

Our prior distributions are loosely motivated by the vast empirical literature on modeling income and consumption dynamics. First, the priors for the precisions (inverse variances) are $\Gamma(2.5, 2.5)$. Meanwhile, because there is no consensus in the literature regarding the estimate of the impact of permanent income shock on consumption, we choose an uninformative $U(0, 1)$ prior for γ_η. The priors for the impact coefficients, $\lambda_{y\eta}$, $\lambda_{c\eta}$, and $\lambda_{c\epsilon}$ are $TN_{[-1,1]}(0, 0.5^2)$ – i.e., they are truncated to ensure that they lie between -1 and 1. The priors for autoregressive and moving-average coefficients are $TN_{|z|>1,\phi(z)=0}(0, 0.5^2)$ and $TN_{|z|>1,\theta(z)=0}(0, 0.5^2)$ – i.e., they are truncated to ensure stationarity or invertibility. For the BPP model, the MA coefficient θ and the impact of a transitory income shock on consumption γ_ϵ are $TN_{[-1,1]}(0, 0.5^2)$. As we will see, the priors are sufficiently uninformative that the implied variances from the models generally match sample variances.

Table 9 reports Bayesian estimates for the different model specifications. First, estimates for the full UC-AR(2) model in the second column suggest no persistent transitory dynamics for income. However, permanent income shocks have an immediate positive impact on transitory income, as would be consistent with a delayed labor supply response to partially offset the initial change in income. This stands in contrast to some other studies, such as Hyrshko (2010) and Belzil and Bognanno (2008), which find a negative correlation between the permanent and transitory income shocks. Also, for both the income and consumption processes, transitory shocks are somewhat more volatile compared to permanent shocks, although of similar magnitude. The implied variances of income and consumption closely match their corresponding sample counterparts in the data, 0.09 and 0.16. Second, when we shut down the additional permanent shocks to consumption beyond permanent income shocks in the third column, the implied variance of consumption is much lower than the variance of consumption in the data. Third, when we set the impact coefficients to zero in the fourth column, we find relatively similar estimates for the other parameters, as for the full UC-AR(2) model in the second column. Fourth, when we shut down all dynamics and set all impact coefficients to zero in the fifth column, which corresponds to the UC-WN model, we find, again, that the estimates of the remaining parameters remain similar to the full UC-AR(2) model. In particular, the variance of income shocks and transitory shocks to consumption, as well as the pass-through of the permanent income shock to consumption, are quite similar across the different specifications. Notably, these results also hold for the BPP model reported in the sixth column.

Using Bayesian methods, we also compare different models to determine which is supported by the data. We do so by computing the marginal likelihood following the method in Chib and Jeliazkov (2001). The last row in Table 9 shows that the basic UC-WN model is clearly preferred, with the highest marginal likelihood by a huge margin. That is, given similar priors across the models, the UC-WN model would have produced

	UC-AR(2)	$\sigma_u = 0$	$\lambda = 0$	UC-WN	BPP
			I. Income Parameters		
$\phi_{y,1}$	-0.02 (0.01)	-0.06 (0.02)	-0.08 (0.01)		
$\phi_{y,2}$	-0.05 (0.01)	-0.10 (0.02)	-0.11 (0.01)		
θ					-0.04 (0.02)
σ_η	0.14 (0.01)	0.14 (0.01)	0.14 (0.01)	0.14 (0.01)	0.15 (0.02)
σ_ϵ	0.16 (0.02)	0.16 (0.02)	0.16 (0.02)	0.17 (0.01)	0.17 (0.02)
$\lambda_{y\eta}$	0.11 (0.01)	0.10 (0.01)			
			II. Consumption Parameters		
$\phi_{c,1}$	-0.12 (0.01)	-0.37 (0.01)	-0.30 (0.01)		
$\phi_{c,2}$	-0.07 (0.01)	-0.32 (0.01)	-0.20 (0.01)		
σ_u	0.13 (0.01)		0.14 (0.01)	0.13 (0.01)	0.13 (0.01)
σ_v	0.20 (0.02)	0.29 (0.02)	0.19 (0.02)	0.21 (0.02)	0.21 (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)	0.46 (0.02)
γ_ϵ					-0.01 (0.01)
			III. Implied Variances		
Δy	0.08 (0.00)	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)	0.11 (0.00)
			IV. Marginal Likelihoods		
$ln(f(y,c))$	-89595	-110416	-90295	-89041	-89946

Table 9. Bayesian Estimates and Marginal Likelihoods

Notes: The table reports posterior means, with posterior standard deviations in parentheses. The marginal likelihoods are reported in the bottom row. The total number of households are 1,765.

a much better density forecast of the data than the other models.[22] However, we note that the estimates of the key parameters of interest, the variance of the permanent shock to income and the pass-through to consumption, are robust and similar to the (Q)MLE estimates reported in Tables 2-4. The estimates for the variance of the permanent income shock, σ_{η}^2, are always 0.02, while the implied estimates for consumption insurance, $\vartheta_c = 1 - (\gamma_{\eta} + \lambda_{c\eta})$, are between 48-55 percent. Furthermore, all of the underlying parameter values are reasonably precisely estimated.

Based on the Bayesian analysis, we conclude that the basic UC-WN specification used in our (Q)MLE analysis is preferred, but that likelihood-based inferences about the key parameters of interest in terms of income risk and consumption insurance are largely robust to different model specifications, including the BPP model. Summarising our key conclusions, some of which are in contrast with aggregate time series data based model specifications, we find that i) idiosyncratic household consumption is subject to permanent shocks beyond those to permanent income, ii) permanent and transitory movements are generally uncorrelated with each other, except in terms of the key parameter of interest, γ_{η}, which captures the effects of permanent income on permanent consumption, and iii) transitory income and consumption dynamics are not important at an annual frequency. With these results in mind, we can be more confident in the (Q)MLE results based on a UC-WN model presented earlier were not the result of a model specification that is particularly at odds with the data.

4.4 Subgroup analysis

In this subsection, we estimate income risk and consumption insurance for different subgroups based on age and education. Following the Monte Carlo analysis and Bayesian model comparison results, we report (Q)MLE estimates for the UC-WN model only.

Looking at Table 10, permanent income risk, as captured by σ_{η}, is fairly similar across different subgroups, but there is clear evidence of heterogeneity with respect to the implied estimates of consumption insurance, ϑ_c.[23] In particular, for households with no college education, the estimated consumption insurance is 47 percent, while it is 67 percent for households with college. Qualitatively, these results are similar to BPP, although the magnitudes are different. Meanwhile, based on age, we find that consumption insurance is 52 percent for young households (head of household is aged between 30-47

[22] The finding in favour of a simple UC model with no dynamics and no correlation between permanent and transitory movements stand in contrast to what was found for the aggregate data in Morley (2007), 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. Also, we are using annual data, while Morley (2007) considers quarterly data, so we would expect less persistent transitory dynamics.

[23] See Appendix B for estimates of all of the parameter estimates for the subgroups.

TABLE 10. (Q)MLE ESTIMATES FOR SUBGROUPS

	Education		Age	
Parameter	No College	College	Young	Older
σ_η	0.12 (0.00)	0.13 (0.00)	0.11 (0.00)	0.12 (0.00)
γ_η	0.53 (0.04)	0.33 (0.02)	0.48 (0.02)	0.33 (0.03)

Notes: The table reports (quasi) maximum likelihood point estimates, with standard deviations in parentheses. For Education, there are 883 households classified as No College and 882 households as College. For age, there are 1,413 households classified as Young and 708 households classified as Older.

years), while it is 67 percent for older households. This suggests that older households are better able to insulate consumption against fluctuations in income relative to younger households. An obvious source of this higher consumption insurance could be higher buffer-stock savings and/or more access to formal income insurance arrangements, in addition to possibly greater flexibility in adjusting labor supply in response to shocks. BPP mention that they find some evidence of an age profile in their estimates of the consumption insurance parameter, although their estimates are imprecise. By contrast, our (Q)MLE estimates are reasonably precise, in addition to being intuitive. They also suggest higher levels of consumption insurance than reported in BPP.

5 Conclusion

In this paper, we have proposed estimating household consumption insurance using a panel unobserved components model and quasi maximum likelihood rather than generalized method of moments. We show that a likelihood-based approach leads to higher and more precise estimates of consumption insurance than previously reported by Blundell, Pistaferri, and Preston (2008) for the same panel dataset. Monte Carlo analysis suggests this precision reflects a greater accuracy of quasi maximum likelihood, even when shocks are highly non-Normal. Bayesian analysis supports a simple version of our proposed unobserved components model that is straightforward to estimate via quasi maximum likelihood. Subgroup estimates are also precise and suggest intuitive heterogeneity across households grouped by age or education.

Our proposed likelihood-based approach to estimation addresses the apparent sensitivity of results to the choice of which moments to consider and what weighting scheme to apply given a set of moments. Also, much like the use of Bayesian estimation in estimation of structural DSGE models in aggregate macroeconomics, we note that our likelihood-based approach could be potentially useful for estimating structural incomplete-markets models of household income and consumption that would still be consistent with the data. However, our admittedly reduced-form approach has the benefit of being robust to different risk-sharing mechanisms in line with Deaton (1997)'s argument that "Saving is only one of the ways people can protect their consumption against fluctuations in their incomes... Although it is also possible to examine the mechanisms, the insurance contracts, tithes, and transfers, their multiplicity makes it attractive to look directly at the magnitude that is supposed to be smoothed, namely consumption."

References

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

[2] Aiyagari, S.R. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Quarterly Journal of Economics*, 109, 659-684.

[3] Altonji, J. G., and L. M. Segal. 1996. "Small-Sample Bias in GMM Estimation of Covariance Structures." *Journal of Business & Economic Statistics*, 14, 353-366.

[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, R., L. Pistaferri, and I. Preston. 2008. "Consumption Inequality and Partial Insurance." *American Economic Review*, 98, 1887-1921.

[6] Blundell, R., L. Pistaferri, and I. Saporta-Eksten. 2016. "Consumption Smoothing and Family Labor Supply." *American Economic Review*, 106, 387-435.

[7] Brzozowski, M., M. Gervais, P. Klein, and M. Suzuki. 2010. "Consumption, Income, and Wealth Inequality in Canada." *Review of Economic Dynamics*, 13, 52-75.

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

[9] Chib, S., and I. Jeloazkov. 2001. "Marginal Likelihood from the Metropolis-Hastings Output." *Journal of the American Statistical Association*, 96, 270-281.

[10] Chatterjee, A., A. Singh, and T. Stone. 2016. "Understanding Wage Inequality in Australia." *Economic Record*, 92, 348-60.

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

[12] Commault, J. 2017. "How Does Consumption Respond to a Transitory Income Shock? Reconciling Natural Experiments and Structural Estimations."" Mimeo.

[13] Daly, M., D. Hryshko, and I. Manovskii. 2016. "Improving the Measurement of Earnings Dynamics." NBER Working Paper no. 22938.

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

[15] Domeij, D., and M. Floden. 2010. "Inequality Trends in Sweden 1978-2004. " *Review of Economic Dynamics*, 13, 179-208.

[16] Ejrnaes, M., and M. Browning. 2014. "The Persistent-Transitory Representation for Earnings Processes." *Quantitative Economics*, 5, 555-581.

[17] Friedman, M., and S. Kuznets. 1945. "Income from Independent Professional Practice." *New York: National Bureau of Economic Research.*

[18] Fuchs-Schundeln, N., D. Krueger, and M. Sommer. 2010. "Inequality Trends for Germany in the Last Two Decades: A Tale of Two Countries." *Review of Economic Dynamics*, 13, 103-32.

[19] Gottschalk, P., and R. Moffitt. 1994. "The Growth of Earnings Instability in the U.S. Labor Market." *Brookings Papers on Economic Activity*, 2, 217-72.

[20] Guvenen, F. 2007. "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?" *American Economic Review*, 97, 687-712.

[21] Guvenen, F., F. Karahan, S. Ozcan, and J. Song. 2015. "What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk?" *Staff Reports 710*, Federal Reserve Bank of New York.

[22] Heathcote, J., F. Perri, and G. L. Violante. 2010. "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006." *Review of Economic Dynamics*, 13, 15-51.

[23] Heathcote, J., K. Storesletten, and G. L. Violante. 2010. "The Macroeconomic Implications of Rising Wage Inequality in the United States." *Journal of Political Economy*, 118, 681-722.

[24] Heathcote, J., K. Storesletten, and G. L. Violante. 2014. "Consumption and Labor Supply with Partial Insurance: An Analytical Framework." *American Economic Review*, 104, 2075-2126.

[25] Huggett, M., 1997. "The One-sector Growth Model with Idiosyncratic Shocks: Steady States and Dynamics." *Journal of Monetary Economics*, 39, 385-403.

[26] Jappelli, T., and L. Pistaferri. 2011. "The Consumption Response to Income Changes." *Annual Review of Economics*, 2, 479-506.

[27] Kaplan, G., and G. L. Violante. 2010. "How Much Consumption Insurance Beyond Self-Insurance?" *American Economic Journal: Macroeconomics*, 2, 53-87.

[28] Low, H., C. Meghir, and L. Pistaferri. 2010. "Wage Risk and Employment Risk over the Life Cycle." *American Economic Review*, 100, 1432-67.

[29] Mace, B. 1991. "Full Insurance in the Presence of Aggregate Uncertainty." *Journal of Political Economy*, 99, 928-56.

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

[31] Madera, R. 2017. "How Shocking are Shocks?" Working paper, University of Minnesota.

[32] Moffitt, R. A., and P. Gottschalk. 2002. "Trends in the Transitory Variance of Earnings in the United States." *Economic Journal*, 112, C68-C73.

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

[34] Morley, J. C., C. R. Nelson, and E. Zivot. 2003. "Why are the Beveridge-Nelson and Unobserved-Component Decompositions of GDP so Different?" *Review of Economics and Statistics*, 85, 235-243.

[35] Morley, J. C., and A. Singh. 2016. "Inventory Shocks and the Great Moderation." *Journal of Money, Credit and Banking*, 48, 699-728.

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

[37] Storesletten, K., C. Telmer, and A. Yaron. 2004. "Consumption and Risk Sharing over the Life Cycle." *Journal of Monetary Economics*, 51, 609-33.

[38] Townsend, R. 1994. "Risk and Insurance in Village India. " *Econometrica*, 62, 539-92.

[39] Trenkler, C., and E. Weber. 2016. "On the Identification of Multivariate Correlated Unobserved Components Models. " *Economics Letters*, 138, 15-18.

[40] White, H. 1982. "Maximum Likelihood Estimation of Misspecified Models. " *Econometrica*, 50, 1-25.

A State-Space Representations and Implied Variances

In this appendix, we present the state-space representations and implied variances for a panel UC-AR(2) model and for the BPP model, respectively. We also briefly discuss estimation.

Suppressing the individual specific subscript throughout the appendix for simplicity, the observation equation for the panel UC model is

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

where

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

The state equation is

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

where

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

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

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

For the BPP model, the observation equation is

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

where

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

The state equation is

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

where

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

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

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

Turning next to the computation of implied variances, income and consumption growth for the panel UC-AR(2) model are given as follows:

$$\Delta y_t = \eta_t + z_t^y, \tag{A.1}$$

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

$$\Delta c_t = \gamma_c \eta_t + u_t + z_t^c, \tag{A.2}$$

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: $\mathbf{W}[1,1] = \mathbf{W}[1,5] = \mathbf{W}[5,1] = \mathbf{W}[5,5] = \lambda_{y\eta}^2\sigma_\eta^2 + \sigma_\epsilon^2$, $\mathbf{W}[1,3] = \mathbf{W}[3,1] = \mathbf{W}[1,6] = \mathbf{W}[6,1] = \mathbf{W}[3,5] = \mathbf{W}[5,3] = \mathbf{W}[5,6] = \mathbf{W}[6,5] = \lambda_{y\eta}\lambda_{c\eta}\sigma_\eta^2 + \lambda_{c\epsilon}\sigma_\epsilon^2$, and $\mathbf{W}[3,3] = \mathbf{W}[3,6] = \mathbf{W}[6,3] = \mathbf{W}[6,6] = \lambda_{c\eta}^2\sigma_\eta^2 + \lambda_{c\epsilon}^2\sigma_\epsilon^2 + \sigma_v^2$.

Because 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

$$\begin{aligned} var(\Delta c_t) &= var(\gamma_\eta\eta_t + u_t + z_t^c) \\ &= \gamma_\eta^2\sigma_\eta^2 + \sigma_u^2 + var(z_t^c) + 2cov(\gamma_\eta\eta_t, z_t^c) \\ &= \gamma_\eta^2\sigma_\eta^2 + \sigma_u^2 + var(z_t^c) + 2\gamma_\eta\lambda_{c,\eta}\sigma_\eta^2, \end{aligned}$$

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

For the BPP model, computing the implied variances is relatively simple. They are given as follows:

$$var(\Delta y_t) = \sigma_\eta^2 + 2\sigma_\epsilon^2(1 + \theta^2 - \theta), \tag{A.3}$$

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

Similarly,

$$var(\Delta c_t) = \gamma_\eta^2 \sigma_\eta^2 + \gamma_\epsilon^2 \sigma_\epsilon^2 + \sigma_u^2 + 2\sigma_v^2, \tag{A.4}$$

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

Given the state-space representations and an assumption of Normality, we can use the Kalman filter to construct the likelihood for the panel UC model or the BPP model based on the prediction error decomposition. In addition, the Kalman filter can be easily adapted to handle missing observations, which are prevalent in the BPP dataset. Given the likelihood, we use numerical optimization to conduct (Q)MLE and MCMC methods to conduct Bayesian inference. Details of the posterior simulation and the priors are presented in the main text. Meanwhile, for GMM estimation, we follow the approach in BPP. See their paper and appendix for details.

B Full Sets of Estimates

TABLE A1. FULL SET OF ESTIMATES FOR THE UC-WN MODEL

Parameter	MLE	GMM		
		DWMD	EWMD	OMD
σ_η	0.13 (0.002)	0.18 (0.002)	0.17 (0.002)	0.12 (0.001)
σ_ϵ	0.20 (0.002)	0.16 (0.001)	0.17 (0.001)	0.14 (0.001)
σ_u	0.09 (0.002)	0.12 (0.003)	0.12 (0.005)	0.09 (0.003)
σ_v	0.28 (0.002)	0.25 (0.003)	0.26 (0.004)	0.21 (0.001)
γ_η	0.42 (0.017)	0.37 (0.050)	0.41 (0.050)	0.31 (0.030)

Note: The table reports point estimates, with standard deviations in parentheses.

TABLE A2. FULL SET OF ESTIMATES FOR THE UC-WN-TVV MODEL

Parameter		MLE	GMM		
			DWMD	EWMD	OMD
σ_η	1979-81	0.11 (0.004)	0.13 (0.004)	0.13 (0.004)	0.13 (0.002)
.	1982	0.11 (0.009)	0.17 (0.004)	0.17 (0.004)	0.14 (0.002)
.	1983	0.14 (0.008)	0.20 (0.006)	0.20 (0.006)	0.13 (0.002)
.	1984	0.13 (0.009)	0.19 (0.004)	0.19 (0.005)	0.09 (0.002)
.	1985	0.15 (0.009)	0.20 (0.011)	0.21 (0.012)	0.15 (0.005)
.	1986	0.11 (0.011)	0.18 (0.006)	0.17 (0.006)	0.13 (0.002)
.	1987	0.15 (0.009)	0.20 (0.005)	0.20 (0.005)	0.14 (0.003)
.	1988	0.09 (0.013)	0.15 (0.007)	0.15 (0.007)	0.14 (0.004)
.	1989	0.14 (0.009)	0.16 (0.005)	0.17 (0.006)	0.13 (0.003)
.	1990-92	0.14 (0.006)	0.15 (0.005)	0.16 (0.005)	0.11 (0.002)
σ_ϵ	1979	0.20 (0.006)	0.18 (0.005)	0.17(0.005)	0.13 (0.002)
.	1980	0.17 (0.006)	0.15 (0.003)	0.17 (0.004)	0.11 (0.001)
.	1981	0.18 (0.006)	0.15 (0.003)	0.15 (0.003)	0.13 (0.001)
.	1982	0.18 (0.006)	0.14 (0.003)	0.14 (0.003)	0.13 (0.002)
.	1983	0.18 (0.006)	0.14 (0.003)	0.13 (0.003)	0.13 (0.002)
.	1984	0.19 (0.006)	0.16 (0.003)	0.16 (0.003)	0.14 (0.002)
.	1985	0.23 (0.007)	0.19 (0.007)	0.19 (0.007)	0.16 (0.003)
.	1986	0.22 (0.007)	0.18 (0.005)	0.19 (0.005)	0.17 (0.003)
.	1987	0.22 (0.006)	0.19 (0.005)	0.19 (0.005)	0.16 (0.003)
.	1988	0.20 (0.006)	0.17 (0.004)	0.17 (0.004)	0.14 (0.002)
.	1989	0.19 (0.006)	0.17 (0.006)	0.17 (0.006)	0.16 (0.003)
.	1990-92	0.20 (0.005)	0.19 (0.003)	0.19 (0.003)	0.14 (0.001)
σ_u		0.08 (0.002)	0.10 (0.004)	0.11 (0.005)	0.09 (0.001)
σ_v	1979	0.26 (0.006)	0.25 (0.007)	0.27 (0.010)	0.22 (0.004)
.	1980	0.24 (0.007)	0.25 (0.007)	0.27 (0.010)	0.22 (0.004)
.	1981	0.24 (0.007)	0.23 (0.007)	0.22 (0.008)	0.20 (0.003)
.	1982	0.28 (0.007)	0.23 (0.008)	0.23 (0.008)	0.20 (0.004)
.	1983	0.26 (0.007)	0.25 (0.008)	0.25 (0.008)	0.21 (0.004)
.	1984	0.34 (0.008)	0.26 (0.009)	0.26 (0.009)	0.20 (0.004)
.	1985	0.30 (0.007)	0.31 (0.015)	0.31 (0.015)	0.25 (0.008)
.	1986	0.27 (0.007)	0.28 (0.016)	0.28 (0.017)	0.24 (0.008)
.	1989	0.31 (0.008)	0.25 (0.006)	0.27 (0.009)	0.23 (0.004)
.	1990-92	0.28 (0.004)	0.26 (0.007)	0.25 (0.007)	0.21 (0.003)
γ_η		0.42 (0.016)	0.51 (0.060)	0.46 (0.070)	0.31 (0.030)

Note: The table reports point estimates, with standard deviations in parentheses.

Table A3. Full Set of Estimates for the UC-BPP-TVV Model

Parameter		MLE	GMM		
			DWMD	EWMD	OMD
σ_η	1979-81	0.14 (0.004)	0.10 (0.003)	0.10 (0.004)	0.12 (0.002)
.	1982	0.10 (0.012)	0.14 (0.004)	0.14 (0.004)	0.13 (0.003)
.	1983	0.13 (0.009)	0.17 (0.005)	0.18 (0.006)	0.13 (0.003)
.	1984	0.12 (0.010)	0.17 (0.004)	0.17 (0.006)	0.08 (0.002)
.	1985	0.13 (0.011)	0.17 (0.009)	0.18 (0.012)	0.15 (0.006)
.	1986	0.10 (0.013)	0.15 (0.006)	0.14 (0.006)	0.12 (0.003)
.	1987	0.13 (0.010)	0.17(0.006)	0.17 (0.006)	0.12 (0.004)
.	1988	0.06 (0.021)	0.13 (0.006)	0.12 (0.006)	0.14 (0.004)
.	1989	0.20 (0.010)	0.14 (0.005)	0.14 (0.007)	0.12 (0.003)
.	1990-92	0.12 (0.006)	0.12 (0.004)	0.14 (0.005)	0.09 (0.002)
σ_ϵ	1979	0.18 (0.005)	0.20 (0.005)	0.20 (0.006)	0.15 (0.003)
.	1980	0.19 (0.007)	0.17 (0.004)	0.19 (0.005)	0.12 (0.002)
.	1981	0.19 (0.006)	0.17 (0.003)	0.18 (0.004)	0.14 (0.002)
.	1982	0.20 (0.006)	0.17 (0.004)	0.17 (0.004)	0.16 (0.002)
.	1983	0.20 (0.006)	0.16 (0.003)	0.16 (0.004)	0.14 (0.002)
.	1984	0.20 (0.006)	0.19 (0.003)	0.19 (0.004)	0.16 (0.002)
.	1985	0.25 (0.007)	0.21 (0.007)	0.22 (0.008)	0.17 (0.003)
.	1986	0.24 (0.007)	0.21 (0.006)	0.21 (0.006)	0.18 (0.004)
.	1987	0.24 (0.006)	0.22 (0.005)	0.22 (0.005)	0.18 (0.003)
.	1988	0.22 (0.006)	0.20 (0.004)	0.20 (0.004)	0.15 (0.002)
.	1989	0.21 (0.006)	0.19 (.006)	0.19 (0.007)	0.16 (0.003)
.	1990-92	0.23 (0.005)	0.21 (0.004)	0.21 (0.004)	0.16 (0.002)
σ_u		0.08 (0.002)	0.10 (0.004)	0.12 (0.005)	0.10 (0.001)
σ_v	1979	0.26 (0.006)	0.25 (0.008)	0.27 (0.010)	0.23 (0.004)
.	1980	0.24 (0.007)	0.25 (0.008)	0.27 (0.010)	0.23 (0.004)
.	1981	0.24 (0.007)	0.23 (0.007)	0.22 (0.008)	0.21 (0.004)
.	1982	0.28 (0.007)	0.23 (0.008)	0.23 (0.008)	0.20 (0.004)
.	1983	0.26 (0.007)	0.25 (0.008)	0.25 (0.008)	0.22 (0.005)
.	1984	0.34 (0.008)	0.26 (0.009)	0.26 (0.009)	0.21 (0.004)
.	1985	0.30 (0.007)	0.31 (0.015)	0.31 (0.015)	0.27 (0.009)
.	1986	0.27 (0.007)	0.29 (0.016)	0.28 (0.017)	0.25 (0.008)
.	1989	0.31 (0.008)	0.26 (0.006)	0.28 (0.009)	0.24 (0.004)
.	1990-92	0.28 (0.004)	0.26 (0.007)	0.26 (0.007)	0.22 (0.003)
θ		0.19 (0.013)	0.11 (0.024)	0.11 (0.026)	0.13 (0.019)
γ_ϵ		0.04 (0.013)	0.05 (0.043)	0.06 (0.046)	0.17 (0.029)
γ_η		0.46 (0.020)	0.64 (0.094)	0.54 (0.099)	0.26 (0.035)

Note: The table reports point estimates, with standard deviations in parentheses.

Table A4. Full Set of (Q)MLE Estimates for Subgroups

	No college	College	Younger	Older
			I. Income	
σ_η	0.12 (0.00)	0.13 (0.00)	0.11 (0.00)	0.12 (0.00)
σ_ϵ	0.21 (0.00)	0.19 (0.00)	0.18 (0.00)	0.24 (0.00)
			II. Consumption	
σ_u	0.08 (0.00)	0.08 (0.00)	0.07 (0.00)	0.09 (0.00)
σ_v	0.32 (0.00)	0.24 (0.00)	0.29 (0.00)	0.26 (0.00)
γ_η	0.53 (0.04)	0.33 (0.02)	0.48 (0.02)	0.33 (0.03)

Note: The table reports point estimates, with standard deviations in parentheses.