Estimating Household Consumption Insurance

By Arpita Chatterjee, James Morley, and Aarti Singh*

Draft: 5 July 2019

Blundell, Pistaferri, and Preston (2008) report an estimate of household consumption insurance with respect to permanent income shocks of 36%. Their estimate is distorted by an error in their code and is not robust to weighting scheme for GMM. We propose instead to use quasi maximum likelihood estimation (QMLE), which produces a more precise and significantly higher estimate of consumption insurance at 55%. For sub-groups by age and education, differences between estimates are even more pronounced. Monte Carlo experiments with non-Normal shocks demonstrate that QMLE is more accurate than GMM, especially given a smaller sample size.

JEL: E21; C13; C33

Keywords: consumption insurance; weighting schemes; quasi maximum likelihood

How does idiosyncratic income risk impact consumption when households have limited access to insurance via formal markets or informal arrangements?[1] In a seminal paper, Blundell, Pistaferri, and Preston (2008) (BPP hereafter) con-

* Chatterjee: UNSW, School of Economics, arpita.chatterjee@unsw.edu.au. Morley: University of Sydney, School of Economics, james.morley@sydney.edu.au. Singh: University of Sydney, School of Economics, aarti.singh@sydney.edu.au. We thank Moshe Buchinsky, Yongsung Chang, Bruce Hansen, James Hansen, Bo Honoré, Greg Kaplan, Jay Lee, Masao Ogaki, Thijs van Rens 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, Frontiers in Econometrics Workshop (Sydney), Continuing Education in Macroeconometrics Workshop (Sydney), IAAE Conference (Sapporo), SNDE Symposium (Tokyo), ANU, Keio University, Monash University, UQ, University of Melbourne, and UTS for helpful comments. Patrick Meyer provided excellent research assistance. Code is available at `https://drive.google.com/open?id=1-CzPYEnMUmX8plDz4eKBOCBSPNOfnu3e`. We are grateful for the financial support from the Australian Research Council grant DE130100806 (Singh). The usual disclaimers apply.

[1]See Jappelli and Pistaferri (2011) for a comprehensive survey of this literature.

struct a novel panel dataset of income and (imputed) consumption for the Panel Study of Income Dynamics and employ GMM to estimate the degree of household consumption insurance based on their proposed model for the data. Numerous studies have followed their approach (e.g., Kaplan, Violante and Weidner 2014; Auclert, 2019). In this paper, we re-visit the estimation in BPP and propose a quasi maximum likelihood estimation (QMLE) approach instead.

Using diagonal weights for GMM, BPP report an estimate of consumption insurance with respect to permanent income shocks of 36%. We find that their estimate is distorted by an error in their original code. In particular, not all sample moments are matched correctly to model-implied moments. For their dataset, this does not alter the estimate very much (33% instead of 36%), but it does have a large effect on the performance of the estimator in repeated samples.[2] Of potentially greater concern, the GMM estimate of consumption insurance is imprecise and highly dependent on the weighting scheme used for GMM. Given optimal weights, we find a very different estimate of 67%. Furthermore, the estimates are even less precise and more sensitive to weighting scheme when considering sub-groups by age and education.

For the same dataset, we find that QMLE leads to a more precise and significantly higher estimate of consumption insurance than BPP at 55%. The QMLE approach avoids having to make a choice about weighting scheme and gains efficiency by implicitly imposing the model structure in fitting sample moments across time. By contrast, GMM using diagonal weights assumes independence of sample moments over time, while GMM using optimal weights allows for potential dependence, but does not take the full structure of the model into account in estimation and appears to perform poorly in small samples.[3] QMLE is made

[2]For the Monte Carlo experiments considered in this paper, preliminary analysis revealed that the error in the code leads to extremely inaccurate estimates for GMM using diagonal weights.

[3]See Altonji and Segal (1996) on small-sample biases for GMM using optimal weights in a cross-sectional setting. In our Monte Carlo analysis, we find similar problems for GMM using optimal weights when considering a smaller sample size in a panel setting. Meanwhile, Altonji and Segal (1996) speculate that QMLE would suffer from the same issues as GMM using optimal weights. However, we do not find this is the case in our Monte Carlo analysis.

feasible in the panel setting by considering an unobserved components version of the BPP model and applying the Kalman filter to construct the likelihood under the assumption of Normality.[4] Given the widely-noted non-Normality of the data (see, for example, Guvenen, Karahan, Ozcan, and Song, 2015), we regard estimation as quasi maximum likelihood (White, 1982) and demonstrate its performance relative to GMM via Monte Carlo experiments with highly non-Normal shocks drawn from heavy-tailed empirical distributions.

For sub-group analysis, the differences in estimates between GMM and QMLE are even more pronounced. QMLE produces significantly different estimates across households grouped both by age and education. As in BPP, we find that college-educated households have more consumption insurance. However, our estimates are quite different than theirs and much more precise. For sub-groups based on age, BPP do not report estimates due to imprecision. Using QMLE, we find that younger households have significantly lower consumption insurance than older households.

Motivated by these empirical results, we conduct Monte Carlo experiments for two different sample sizes. Given the same effective sample size as the full BPP dataset, we find that QMLE is the most accurate in terms of root mean squared error. GMM using diagonal weights is far less accurate, although GMM using optimal weights is reasonably close in accuracy to QMLE. However, given a smaller effective sample size corresponding to the older household sub-group, QMLE performs much better than GMM regardless of weighting scheme. The Monte Carlo results generally reconcile the differences in the estimates for the BPP dataset and suggest that QMLE should be used instead of GMM, especially given a smaller sample size.

The rest of this paper is organized as follows: Section I describes the data and provides an unobserved components version of the BPP model used for estimation

[4]Somewhat related to this approach, Primiceri and van Rens (2009) employ likelihood-based Bayesian estimation with flat and uninformative priors for an unobserved components model of household income and consumption in order to examine changes in income inequality across households.

via QMLE. Section II reports empirical results for the full BPP dataset and subgroups by age and education. Section III presents our Monte Carlo analysis. Section IV concludes.

I. Data, Model, and Estimation

In our empirical analysis, we use the dataset created by BPP. They consider the Panel Study of Income Dynamics (PSID) sample 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). Because the CEX has detailed nondurable consumption data, while the PSID primarily has only food expenditure consumption data for the years under consideration, BPP impute annual nondurable consumption for each household using the estimates of food demand from the CEX. Their constructed dataset is then a panel of annual observations for income and imputed nondurable consumption over the sample period of 1978-1992. To calculate idiosyncratic income and consumption, BPP first regress income and consumption for households on demographic and ethnic factors, as well as other income characteristics observable/known by consumers, and then calculate the residuals.[5] Following BPP, it is these residuals that we model to estimate household consumption insurance.

For our analysis, we consider an unobserved components version of the BPP model of idiosyncratic household income and consumption with time-varying volatility for income and consumption shocks. In particular, income and consumption for household i are given as follows:

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

[5] These variables include 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. See BPP for full details.

$$c_{i,t} = \gamma_\eta \tau_{i,t} + \kappa_{i,t} + \upsilon_{i,t}, \qquad \upsilon_{i,t} \sim i.i.d.(0, \sigma^2_{\upsilon,t}), \tag{2}$$

where $\tau_{i,t}$ is a common stochastic trend for income and consumption ("permanent income"), $\epsilon_{i,t}$ is a transitory income shock with moving-average parameter $|\theta| < 1$, $\kappa_{i,t}$ is an additional stochastic trend for consumption, and $\upsilon_{i,t}$ is a transitory consumption shock. The stochastic trends are specified as random walks:

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

$$\kappa_{i,t} = \kappa_{i,t-1} + \gamma_\epsilon \epsilon_{i,t} + u_{i,t}, \qquad u_{i,t} \sim i.i.d.(0, \sigma^2_{u,t}), \tag{4}$$

while the structure of the time-varying volatility for each shock is assumed to be deterministic and the same as in BPP.[6]

In terms of economic interpretation, the transitory income shock, $\epsilon_{i,t}$, captures events such as a surprise bonus or temporary leave due to illness, while the transitory consumption shock, $\upsilon_{i,t}$ could capture measurement error due to the imputation of nondurable consumption. The permanent income shock, $\eta_{i,t}$, reflects severe health shocks, promotion, or other idiosyncratic factors that result in a change in permanent income, while the permanent shock to consumption, $u_{i,t}$, could reflect taste and preference shocks or other shocks to non-labor income, such as wealth shocks.

The key parameters that we focus on in our analysis are γ_η and γ_ϵ, which capture the impacts of permanent and transitory income shocks on permanent consumption. The implied "consumption insurance" against permanent income shocks is then $1 - \gamma_\eta$.

We propose estimating the model in (1)-(4) using QMLE. In particular, we cast the model into state-space form (see the appendix) and assume that the shocks are

[6] BPP generally allow different variances of shocks in each period, although they assume some variances are the same across certain years. See BPP for details.

Normally distributed in order to use the Kalman filter to calculate the likelihood based on the prediction error decomposition. Because the actual shocks are likely to be non-Normal at the household level, as shown in some recent literature (e.g., Guvenen, Karahan, Ozcan, and Song, 2015), our proposed approach should be thought of as quasi maximum likelihood following White (1982). We will consider how well QMLE works relative to GMM despite non-Normal shocks with our Monte Carlo analysis in Section III.

For purposes of comparison, we also consider GMM estimation using the same moment conditions as BPP. We employ two approaches to weighting the moment conditions corresponding to the diagonally weighted minimum distance (DWMD) estimator used in BPP and the optimal minimum distance (OMD) estimator. DWMD generalizes an equally weighted minimum distance approach, but allows for heteroskedasticity, while OMD allows for covariance between moment conditions in the weighting matrix. See BPP and their appendix for details of GMM estimation, including for different weighting schemes.

We note that QMLE implicitly imposes the model structure in fitting sample moments across time. For example, in the extreme case of constant shock variances, the model would imply the exact same autocorrelations and cross-correlations for income and consumption growth over time and QMLE would impose this when fitting the data (i.e., it would implicitly pool the sample moments together across time to match the model-implied moments). Even when allowing for time-varying volatility as in BPP, QMLE still imposes the model structure when considering observations of income and consumption growth over time and, hence, should be more precise in estimating the time-invariant consumption response parameters in particular. By contrast, GMM treats sample moments for income and consumption growth data from each time period separately, although OMD estimates how these moments are related across time when determining weights and should be as efficient as QMLE asymptotically.[7] DWMD treats the

[7]In small samples, however, GMM estimates will be less efficient due to the errors in estimating the

Table 1—Estimates of Consumption Responses for All Households

Parameter	QMLE	GMM	
		DWMD	OMD
γ_η	0.45	0.67	0.33
	(0.02)	(0.09)	(0.03)
	[0.41, 0.50]	[0.49, 0.85]	[0.27, 0.39]
γ_ϵ	0.04	0.03	0.07
	(0.01)	(0.04)	(0.03)
	[0.02, 0.06]	[-0.05, 0.11]	[0.01, 0.13]

Notes: Point estimates are reported with standard errors in parentheses and 95% confidence intervals in square brackets. Confidence intervals are constructed by inverting t tests. There are 1,765 households and 15 years of data in levels, but with many missing observations.

sample moments across time as independent and so is not as efficient, although it is consistent.

II. Empirical Results

Table 1 reports estimates for the parameters related to the responses of consumption to income shocks for all households in the BPP dataset (see the appendix for the full set of estimates). The QMLE estimates are the most precise and lie in between the GMM estimates for both parameters. The DWMD estimates are close to those reported in Table 6 of BPP. However, there are differences due to an error in the original BPP code.[8] The OMD estimates are very different from the DWMD estimates, suggesting a high sensitivity to weighting of moments. The QMLE estimate for the effect of permanent income shocks is closer to the OMD estimate, with both implying a significantly higher degree of consumption insurance than reported in BPP of between 55-67% rather than 36%.

weights from the data.

[8] In the BPP code, *md_AER.prg*, available from the AER website, there is a missing transpose on line 289 and unnecessary transpose on line 290. If we use their original code, we obtain identical estimates based on DWMD to those reported in Table 6 of BPP. However, the misplaced transposes lead to a mismatching of some model-implied moments to sample moments and result in substantially different estimates on average in Monte Carlo analysis.

The estimates for the effect of a transitory income shock are more similar across estimation methods, however they are only significantly positive for QMLE and OMD given greater precision compared to DWMD.

Table 2 reports estimates of the same consumption response parameters for sub-groups by age and education. As with the full dataset, the results vary substantially by estimation method, with QMLE estimates almost always being the most precise and generally taking on values in between the GMM estimates. The DWMD estimates for education are again similar to those in Table 6 of BPP, with differences due to the original code error noted above. Meanwhile, for age, the DWMD estimates are highly imprecise.[9] In terms of the effect of permanent income shocks, only the QMLE results correspond to higher consumption insurance for both college-educated and older households. GMM based on OMD suggests the counter-intuitive result that households without college education have substantially higher consumption insurance of 75% versus 53% for households with college education. GMM based on DWMD implies very low consumption insurance for both younger and older households compared to the estimate for all households, with younger households having higher consumption insurance than older households at 27% versus 15%. Again, the estimates display a high sensitivity to weighting of moments, particularly with implied consumption insurance for households without college education ranging from 5% for DWMD to 75% for OMD and for older households ranging from 15% for DWMD to 81% for OMD. Meanwhile, the sub-group estimates for QMLE imply relatively high, but moderate levels of consumption insurance based on education (between 44% and 64%) and age (between 48% and 70%). Furthermore, the QMLE estimates for the effect of a transitory income shock are always more precise than the GMM estimates, which are sometimes counter-intuitively negative.

[9]BPP do not report estimates for sub-groups by age, but note their imprecision in footnote 31 of their paper.

Table 2—Estimates of Consumption Responses for Sub-Groups

Parameter	QMLE	GMM	
		DWMD	OMD
		No College	
γ_η	0.56	0.95	0.25
	(0.02)	(0.17)	(0.04)
	[0.52, 0.60]	[0.61, 1.29]	[0.17,0.33]
γ_ϵ	0.05	0.07	0.18
	(0.02)	(0.06)	(0.03)
	[0.01, 0.09]	[-0.05, 0.19]	[0.12,0.24]
		College	
γ_η	0.36	0.47	0.47
	(0.03)	(0.09)	(0.04)
	[0.30, 0.42]	[0.29, 0.65]	[0.39, 0.55]
γ_ϵ	0.04	-0.01	-0.12
	(0.02)	(0.05)	(0.03)
	[0.00, 0.08]	[-0.11, 0.09]	[-0.18,-0.06]
		Younger (30-47)	
γ_η	0.52	0.73	0.52
	(0.03)	(0.11)	(0.05)
	[0.46, 0.58]	[0.51, 0.95]	[0.42, 0.62]
γ_ϵ	0.01	-0.02	0.02
	(0.02)	(0.07)	(0.04)
	[-0.03, 0.05]	[-0.05, 0.19]	[-0.06,0.10]
		Older (48-65)	
γ_η	0.30	0.85	0.19
	(0.04)	(0.22)	(0.03)
	[0.22, 0.38]	[0.44, 1.28]	[0.13, 0.25]
γ_ϵ	0.08	0.06	0.15
	(0.02)	(0.05)	(0.03)
	[0.04, 0.12]	[-0.04, 0.16]	[0.09,0.21]

Notes: Point estimates are reported, with standard errors in parentheses and 95% confidence intervals in square brackets. Confidence intervals are constructed by inverting t tests. For the sub-groups based on education, there are 883 households classified as 'No College' and 882 households as 'College'. For age, there are 1,413 households classified as 'Younger' and 708 households classified as 'Older'. There are 15 years of data in levels, but with many missing observations.

III. Monte Carlo Analysis

Although the QMLE estimates in Tables 1 and 2 are the most precise, it is a reasonable question whether this is a false precision due to the Normality assumption made in constructing the likelihood for the unobserved components version of the BPP model. To address this question, we consider a Monte Carlo experiment where the data generating process corresponds to the BPP model with shocks drawn from their heavy-tailed empirical distributions based on the estimated model in the previous section.[10] We draw, with replacement, from the empirically-distributed shocks for each time period and use the BPP model with stylized parameters ($\gamma_\eta = 0.50$, $\gamma_\epsilon = 0.10$, $\theta = 0.20$) to construct artificial panels of idiosyncratic income and consumption data that have the same dimension and structure in terms of missing observations as the BPP dataset.

To evaluate the accuracy of a particular estimator, we consider root mean squared error (RMSE).[11] We also report on the underlying sources of the overall RMSE in terms of bias and standard deviation of an estimator, as well as the root mean squared differences of the GMM estimators compared to QMLE. Each statistic is calculated based on averaging across 2,500 simulations. We focus on results for our key parameters of interest, γ_η and γ_ϵ.

Table 3 reports on the accuracy of different estimators given a large sample with the same effective sample size as the BPP dataset. QMLE performs best in terms of RMSE, although OMD is quite similar. DWMD is much less accurate in terms of γ_η. The main reason for the strong performance of QMLE is a much lower standard deviation of the estimator for both parameters, although this is somewhat offset by a bit more bias than OMD. Notably, the GMM estimates also differ considerably from QMLE in a given sample, with root mean squared

[10] Given QMLE parameter estimates, we employ the Kalman filter to extract estimates of the underlying permanent and transitory income and consumption shocks. Notably, these shocks display some negative skewness (ranging from -0.9 to -0.4) and extremely high degrees of kurtosis (ranging from 8 to 13).

[11] Nakata and Tonetti (2015) conduct Monte Carlo analysis to evaluate the RMSE of likelihood-based Bayesian estimators of income risk in a univariate setting and find that they perform better than GMM.

Table 3—Properties of Estimators given a Large Sample

Parameter	Estimator Property	QMLE	GMM	
			DWMD	OMD
$\gamma_\eta = 0.50$	RMSE	0.06	0.12	0.07
	Bias	0.05	0.05	0.03
	Standard Deviation	0.03	0.11	0.06
	Difference from QMLE	-	0.11	0.06
$\gamma_\epsilon = 0.10$	RMSE	0.03	0.03	0.03
	Bias	-0.03	-0.01	0.00
	Standard Deviation	0.00	0.03	0.03
	Difference from QMLE	-	0.04	0.04

Note: RMSE, bias, standard deviation, and (root mean squared) difference from QMLE for different estimators are based on averages across 2,500 simulations with sample size of 32,547 observations in the same location of the panel as for the BPP dataset.

differences of a similar magnitude to their RMSEs.[12]

Looking back at Table 1, the Monte Carlo results in Table 3 help explain some of the key differences across the estimates. In particular, the relative precision of the estimates corresponds closely to the standard deviations of the estimators, with QMLE being most precise and DWMD being least precise. The difference between the QMLE and OMD estimates for γ_η can be partly reconciled by the somewhat higher bias for QMLE and the substantial root mean squared difference from QMLE for OMD. However, the key point is that the QMLE estimator is the most accurate overall based on RMSE and, assuming any bias is similar to the Monte Carlo result, implies consumption insurance between 55-64%.[13] Meanwhile, the QMLE and OMD estimates for γ_ϵ can be reconciled entirely by a difference in bias similar to the Monte Carlo results, with an implied non-trivial impact of transitory income shocks on consumption.

[12]The GMM estimates across weighting schemes also differ considerably from each other in a given sample. Consistent with these Monte Carlo results in which the correct model is specified, Altonji and Segal (1996) argue that finding a large difference in GMM estimates across weighting schemes in practice does not necessarily provide evidence of model misspecification.

[13]It is possible to consider a formal bootstrap correction for bias, but we leave this for future research.

Table 4—Properties of Estimators given a Small Sample

Parameter	Estimator Property	QMLE	GMM	
			DWMD	OMD
$\gamma_\eta = 0.50$	RMSE	0.07	0.26	0.20
	Bias	0.05	0.12	0.03
	Standard Deviation	0.05	0.23	0.20
	Difference from QMLE	-	0.24	0.20
$\gamma_\epsilon = 0.10$	RMSE	0.04	0.06	0.08
	Bias	-0.03	-0.01	0.00
	Standard Deviation	0.03	0.06	0.08
	Difference from QMLE	-	0.06	0.08

Notes: RMSE, bias, standard deviation, and (root mean squared) difference from QMLE for different estimators are based on averages across 2,500 simulations with sample size of 11,437 observations located the same in the same location of the panel as for the sub-group of older households in the BPP dataset.

Table 4 reports on the accuracy of different estimators given a smaller sample with the same effective sample size as the sub-group of older households in the BPP dataset. Given well-known concerns about OMD in small samples for cross-sectional analysis (see Altonji and Segal, 1996), our aim is to understand how the performance of QMLE and GMM compare given a smaller sample size in a panel setting, with the number of observations for older households being smallest amongst the sub-groups.[14] As in the large sample case, QMLE performs best in terms of RMSE, but the improvements over GMM are much more pronounced in this case, especially compared to OMD. The lower RMSE again results from a comparatively low standard deviation of the QMLE estimator for both parameters, with similar biases as before in almost every case. The only change in bias is a large increase for γ_η with DWMD, leading it to being the worst performing estimator again for this key parameter of interest. However, the striking difference from Table 3 is the severe deterioration of OMD, which even performs worse

[14]Out of the total 32,547 observations of data for all of the households, there are 15,735 observations for households without college education versus 16,812 observations for college-educated households and 21,110 observations for younger households versus 11,437 observations for older households.

in terms of RMSE and standard deviation than DWMD for γ_ϵ. By contrast, the accuracy of QMLE is almost as good as it was in the large sample case, suggesting the efficiency gains from imposing the model structure in fitting sample moments across time are substantial. Notably, it is also almost three times more accurate than OMD for γ_η and twice as accurate for γ_ϵ.[15] Meanwhile, the GMM estimators differ from QMLE, with root mean squared differences again of a similar magnitude to their RMSEs.

Looking back at Table 2, the results in Table 4 help explain why the estimates are so different across estimation methods in most cases and raise strong concerns about the accuracy of the GMM estimates given smaller sample sizes. Furthermore, unlike with OMD, the Monte Carlo results provide confidence in the precision of the QMLE estimates. The main takeaway is that consumption insurance appears to be significantly higher for college-educated and older households, with about 20 percentage point higher insurance than the comparison sub-groups in both cases. Again, assuming any bias is similar to the Monte Carlo result, implied consumption insurance is between 45-53% for households without college versus 63-75% for households with college and 47-59% for younger households versus 67-83% for older households. For the effects of transitory income shocks, the QMLE estimates for γ_ϵ are, as in the case of all households, consistent with a non-trivial impact on consumption for all sub-groups, at least when taking the apparent downward bias into account.

IV. Conclusions

In this paper, we have examined the robustness of BPP's low estimated degree of consumption insurance with respect to permanent income shocks. We find that, even for the same dataset, their result is not robust to different estimation methods, with QMLE implying a higher and notably more precise estimate of

[15]As noted previously, GMM using optimal weights estimates how moments should be related across time and can, therefore, achieve the same efficiency as QMLE asymptotically. However, from the Monte Carlo results, sampling uncertainty about weights is clearly relevant given the smaller effective sample size.

consumption insurance. Estimates for sub-groups are also sensitive to estimation method, with QMLE again being the most precise and suggesting intuitive heterogeneity across households grouped by age and education. Monte Carlo analysis assuming non-Normal shocks drawn from heavy-tailed distributions supports the greater accuracy of QMLE versus GMM.

We believe our paper makes two significant contributions to the literature on household consumption insurance. First, we have provided evidence that consumption insurance is considerably higher than previously reported for BPP's dataset, at least based on their model. Kaplan and Violante (2010) have previously argued that BPP's estimate for consumption insurance is downward biased due to model misspecification. Thus, our finding of higher consumption insurance for the dataset is consistent with the implication of Kaplan and Violante (2010), albeit for the different reason of imprecision of the GMM estimator using diagonal weights rather than model misspecification.[16] Second, we have demonstrated the feasibility and accuracy of QMLE in a panel setting with highly non-Normal shocks and a relatively small sample size. A widely-claimed reason why heterogenous agent quantitative models can improve our understanding of the macroeconomy is due to aggregation bias arising out of heterogeneous consumption insurance for various sub-groups of agents. Thus, improved performance in estimating this key parameter in a small sample setting is particularly important and we, therefore, recommend using QMLE in practice.

[16]We leave analysis of the performance of QMLE for different model specifications in terms of permanent and transitory components of income and consumption to future research.

REFERENCES

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.

Auclert, A., 2019, "Monetary Policy and the Redistribution Channel," *American Economic Review* 109, 2333-67.

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

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.

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

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

Kaplan, G, G. L. Violante, and J. Weidner, 2014, "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity* Spring, 77-138.

Nakata, T., and C. Tonetti, 2015, "Small Sample Properties of Bayesian Estimators of Labor Income Processes," *Journal of Applied Economics* 18, 121-148.

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

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

Appendix A. State-Space Representation

In this appendix, we present the state-space representation of the unobserved components version of the BPP model.

Suppressing household-specific subscripts for simplicity, the observation equation is

$$\mathbf{y}_t = \mathbf{H}\mathbf{X_t},$$

where

$$\mathbf{y}_t = \begin{bmatrix} y_t \\ c_t \end{bmatrix}, \ \mathbf{H} = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & \gamma_\eta & 1 \end{bmatrix}, \text{ and } \mathbf{X_t} = \begin{bmatrix} y_t - \tau_t \\ c_t - \gamma_\eta \tau_t - \kappa_t \\ \tau_t \\ \kappa_t \end{bmatrix}.$$

The state equation is

$$\mathbf{X}_t = \mathbf{F}\mathbf{X_t} + \mathbf{v_t},$$

where

$$\mathbf{F} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \ \mathbf{v_t} = \begin{bmatrix} \epsilon_t + \theta \epsilon_{t-1} \\ \upsilon_t \\ \eta_t \\ u_t + \gamma_\epsilon \epsilon_t \end{bmatrix},$$

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

$$\mathbf{Q} = \begin{pmatrix} \sigma_\epsilon^2(1+\theta^2) & 0 & 0 & \gamma_\epsilon \sigma_\epsilon^2 \\ 0 & \sigma_\upsilon^2 & 0 & 0 \\ 0 & 0 & \sigma_\eta^2 & 0 \\ \gamma_\epsilon \sigma_\epsilon^2 & 0 & 0 & \sigma_u^2 + \gamma_\epsilon^2 \sigma_\epsilon^2 \end{pmatrix}.$$

Given the state-space representation and an assumption of Normality, we can use the Kalman filter to calculate the likelihood for the BPP model based on the prediction error decomposition and an assumption of independence of idiosyncratic income and consumption across households. In addition, the Kalman filter can be easily adapted to handle missing observations, which are prevalent in the BPP dataset.

Appendix B. Full Sets of Estimates

In this appendix, we report tables with full sets of point estimates and standard errors in parentheses for the BPP model. The tables are reported on the following pages.

Table B1—Estimates for All Households

Parameter		QMLE	GMM	
			DWMD	OMD
σ_η	1979-81	0.14 (0.00)	0.10 (0.00)	0.12 (0.00)
.	1982	0.10 (0.01)	0.14 (0.00)	0.14 (0.00)
.	1983	0.13 (0.01)	0.17 (0.01)	0.13 (0.00)
.	1984	0.12 (0.01)	0.17 (0.01)	0.08 (0.00)
.	1985	0.13 (0.01)	0.17 (0.01)	0.14 (0.01)
.	1986	0.10 (0.01)	0.15 (0.01)	0.12 (0.00)
.	1987	0.13 (0.01)	0.17 (0.01)	0.12 (0.00)
.	1988	0.06 (0.02)	0.13 (0.01)	0.14 (0.01)
.	1989	0.12 (0.01)	0.13 (0.01)	0.12 (0.00)
.	1990-92	0.12 (0.01)	0.12 (0.00)	0.10 (0.00)
σ_ϵ	1979	0.18 (0.01)	0.19 (0.01)	0.14 (0.00)
.	1980	0.19 (0.01)	0.17 (0.00)	0.12 (0.00)
.	1981	0.19 (0.01)	0.17 (0.00)	0.14 (0.00)
.	1982	0.20 (0.01)	0.17 (0.00)	0.16 (0.00)
.	1983	0.19 (0.01)	0.16 (0.00)	0.14 (0.00)
.	1984	0.20 (0.01)	0.19 (0.00)	0.16 (0.00)
.	1985	0.25 (0.01)	0.21 (0.01)	0.17 (0.00)
.	1986	0.24 (0.01)	0.21 (0.01)	0.18 (0.00)
.	1987	0.24 (0.01)	0.21 (0.01)	0.18 (0.00)
.	1988	0.22 (0.01)	0.20 (0.01)	0.16 (0.00)
.	1989	0.21 (0.01)	0.20 (0.01)	0.17 (0.00)
.	1990-92	0.23 (0.00)	0.21 (0.00)	0.16 (0.00)
σ_u		0.08 (0.00)	0.10 (0.00)	0.09 (0.00)
σ_v	1979	0.26 (0.01)	0.25 (0.01)	0.23 (0.00)
.	1980	0.24 (0.01)	0.23 (0.01)	0.20 (0.00)
.	1981	0.24 (0.01)	0.23 (0.01)	0.20 (0.00)
.	1982	0.28 (0.01)	0.25 (0.01)	0.21 (0.01)
.	1983	0.26 (0.01)	0.26 (0.01)	0.20 (0.00)
.	1984	0.34 (0.01)	0.31 (0.02)	0.25 (0.01)
.	1985	0.30 (0.01)	0.28 (0.02)	0.25 (0.01)
.	1986	0.27 (0.01)	0.26 (0.01)	0.23 (0.01)
.	1989	0.31 (0.01)	-	-
.	1990-92	0.28 (0.00)	0.26 (0.01)	0.21 (0.00)
θ		0.19 (0.01)	0.11 (0.03)	0.12 (0.02)
γ_ϵ		0.04 (0.01)	0.03 (0.04)	0.07 (0.03)
γ_η		0.45 (0.02)	0.67 (0.09)	0.33 (0.03)

Note: GMM estimation is based on growth rates and, therefore, cannot estimate a variance for transitory consumption in 1989 given missing consumption data in 1988.

Table B2—Estimates for Households without College Education

Parameter		QMLE	GMM	
			DWMD	OMD
σ_η	1979-81	0.13 (0.01)	0.08 (0.00)	0.12 (0.00)
.	1982	0.08 (0.02)	0.12 (0.01)	0.11 (0.00)
.	1983	0.13 (0.01)	0.18 (0.01)	0.12 (0.00)
.	1984	0.13 (0.01)	0.18 (0.01)	0.10 (0.00)
.	1985	0.11 (0.02)	0.17 (0.01)	0.08 (0.00)
.	1986	0.10 (0.02)	0.13 (0.01)	0.14 (0.00)
.	1987	0.11 (0.02)	0.14 (0.01)	0.12 (0.00)
.	1988	0.05 (0.04)	0.11 (0.01)	0.09 (0.00)
.	1989	0.10 (0.02)	0.10 (0.01)	0.10 (0.00)
.	1990-92	0.11 (0.01)	0.10 (0.01)	0.06 (0.00)
σ_ϵ	1979	0.21 (0.01)	0.21(0.01)	0.13 (0.00)
.	1980	0.20 (0.01)	0.18 (0.01)	0.15 (0.00)
.	1981	0.20 (0.01)	0.19 (0.01)	0.15 (0.00)
.	1982	0.22 (0.01)	0.20 (0.01)	0.15 (0.00)
.	1983	0.22 (0.01)	0.19 (0.01)	0.17 (0.00)
.	1984	0.22 (0.01)	0.20 (0.01)	0.18 (0.00)
.	1985	0.22 (0.01)	0.19 (0.01)	0.19 (0.00)
.	1986	0.23 (0.01)	0.22 (0.01)	0.17 (0.00)
.	1987	0.25 (0.01)	0.23 (0.01)	0.19 (0.00)
.	1988	0.23 (0.01)	0.22 (0.01)	0.15 (0.00)
.	1989	0.24 (0.01)	0.23 (0.01)	0.19 (0.00)
.	1990-92	0.24 (0.01)	0.23 (0.01)	0.18 (0.00)
σ_u		0.08 (0.00)	0.08 (0.01)	0.09 (0.00)
σ_v	1979	0.30 (0.01)	0.28 (0.01)	0.20 (0.01)
.	1980	0.27 (0.01)	0.27 (0.01)	0.24 (0.01)
.	1981	0.25 (0.01)	0.25 (0.01)	0.20 (0.00)
.	1982	0.32 (0.01)	0.28 (0.01)	0.22 (0.00)
.	1983	0.29 (0.01)	0.28 (0.02)	0.21 (0.00)
.	1984	0.39 (0.01)	0.35 (0.03)	0.22 (0.01)
.	1985	0.35 (0.01)	0.33 (0.03)	0.23 (0.01)
.	1986	0.31 (0.01)	0.29 (0.01)	0.24 (0.01)
.	1989	0.35 (0.01)	-	-
.	1990-92	0.32 (0.01)	0.29 (0.01)	0.22 (0.00)
θ		0.19 (0.02)	0.13 (0.03)	0.19 (0.02)
γ_ϵ		0.05 (0.02)	0.07 (0.06)	0.18 (0.03)
γ_η		0.56 (0.02)	0.95 (0.17)	0.25 (0.04)

Note: GMM estimation is based on growth rates and, therefore, cannot estimate a variance for transitory consumption in 1989 given missing consumption data in 1988.

Table B3—Estimates for Households with College Education

Parameter		QMLE	GMM	
			DWMD	OMD
σ_η	1979-81	0.15 (0.01)	0.10 (0.01)	0.09 (0.00)
.	1982	0.12 (0.01)	0.16 (0.01)	0.16 (0.00)
.	1983	0.11 (0.01)	0.15 (0.01)	0.11 (0.00)
.	1984	0.10 (0.01)	0.13 (0.01)	0.06 (0.00)
.	1985	0.15 (0.02)	0.15 (0.01)	-
.	1986	0.09 (0.02)	0.18 (0.01)	0.14 (0.00)
.	1987	0.15 (0.01)	0.19 (0.01)	0.12 (0.00)
.	1988	0.07 (0.02)	0.14 (0.01)	0.18 (0.00)
.	1989	0.13 (0.01)	0.17 (0.01)	0.11 (0.00)
.	1990-92	0.13 (0.01)	0.15 (0.01)	0.09 (0.00)
σ_ϵ	1979	0.15 (0.01)	0.17 (0.01)	0.13 (0.00)
.	1980	0.17 (0.01)	0.17 (0.01)	0.12 (0.00)
.	1981	0.17 (0.01)	0.16 (0.00)	0.09 (0.00)
.	1982	0.18 (0.01)	0.15 (0.00)	0.10 (0.00)
.	1983	0.17 (0.01)	0.14 (0.00)	0.12 (0.00)
.	1984	0.18 (0.01)	0.17 (0.01)	0.15 (0.00)
.	1985	0.27 (0.01)	0.22 (0.01)	0.13 (0.00)
.	1986	0.24 (0.01)	0.21 (0.01)	0.11 (0.00)
.	1987	0.23 (0.01)	0.21 (0.01)	0.12 (0.00)
.	1988	0.21 (0.01)	0.19 (0.01)	0.15 (0.00)
.	1989	0.18 (0.01)	0.15 (0.01)	0.15 (0.00)
.	1990-92	0.21 (0.01)	0.19 (0.00)	0.15 (0.00)
σ_u		0.08 (0.00)	0.11 (0.00)	0.09 (0.00)
σ_v	1979	0.22 (0.01)	0.21 (0.01)	0.19 (0.00)
.	1980	0.20 (0.01)	0.18 (0.01)	0.16 (0.00)
.	1981	0.22 (0.01)	0.22 (0.01)	0.17 (0.00)
.	1982	0.23 (0.01)	0.21 (0.01)	0.19 (0.00)
.	1983	0.24 (0.01)	0.24 (0.01)	0.18 (0.00)
.	1984	0.29 (0.01)	0.27 (0.02)	0.23 (0.01)
.	1985	0.25 (0.01)	0.23 (0.01)	0.22 (0.01)
.	1986	0.24 (0.01)	0.22 (0.01)	0.19 (0.00)
.	1989	0.27 (0.01)	-	-
.	1990-92	0.24 (0.01)	0.24 (0.01)	0.20 (0.00)
θ		0.19 (0.02)	0.11 (0.03)	0.13 (0.03)
γ_ϵ		0.04 (0.02)	-0.01 (0.05)	-0.12 (0.03)
γ_η		0.36 (0.03)	0.47 (0.09)	0.47 (0.04)

Notes: GMM estimation is based on growth rates and, therefore, cannot estimate a variance for transitory consumption in 1989 given missing consumption data in 1988. OMD estimates a negative variance (not reported) for the permanent income shock in 1985.

Table B4—Estimates for Younger Households

Parameter		QMLE	GMM	
			DWMD	OMD
σ_η	1979-81	0.13 (0.00)	0.10 (0.00)	0.09 (0.00)
.	1982	0.10 (0.01)	0.15 (0.00)	0.13 (0.00)
.	1983	0.12 (0.01)	0.16 (0.01)	0.14 (0.00)
.	1984	0.11 (0.01)	0.16 (0.01)	0.10 (0.00)
.	1985	0.07 (0.02)	0.13 (0.01)	0.08 (0.00)
.	1986	0.11 (0.01)	0.17 (0.01)	0.13 (0.00)
.	1987	0.13 (0.01)	0.17 (0.01)	0.12 (0.00)
.	1988	0.08 (0.02)	0.13 (0.01)	0.14 (0.01)
.	1989	0.10 (0.01)	0.14 (0.01)	0.09 (0.00)
.	1990-92	0.11 (0.01)	0.13 (0.00)	0.09 (0.00)
σ_ϵ	1979	0.17 (0.01)	0.18 (0.00)	0.15 (0.00)
.	1980	0.17 (0.01)	0.16 (0.00)	0.13 (0.00)
.	1981	0.18 (0.01)	0.16 (0.00)	0.12 (0.00)
.	1982	0.19 (0.01)	0.16 (0.00)	0.14 (0.00)
.	1983	0.18 (0.01)	0.16 (0.00)	0.14 (0.00)
.	1984	0.18 (0.01)	0.15 (0.00)	0.13 (0.00)
.	1985	0.20 (0.01)	0.18 (0.01)	0.15 (0.00)
.	1986	0.21 (0.01)	0.17 (0.01)	0.15 (0.00)
.	1987	0.23 (0.01)	0.21 (0.01)	0.17 (0.00)
.	1988	0.20 (0.01)	0.18 (0.00)	0.15 (0.00)
.	1989	0.19 (0.01)	0.16 (0.00)	0.14 (0.00)
.	1990-92	0.20 (0.00)	0.18 (0.00)	0.14 (0.00)
σ_u		0.07 (0.00)	0.10 (0.01)	0.09 (0.00)
σ_v	1979	0.28 (0.01)	0.27 (0.01)	0.20 (0.01)
.	1980	0.25 (0.01)	0.23 (0.01)	0.18 (0.00)
.	1981	0.24 (0.01)	0.24 (0.01)	0.19 (0.01)
.	1982	0.28 (0.01)	0.25 (0.01)	0.21 (0.01)
.	1983	0.26 (0.01)	0.25 (0.01)	0.20 (0.00)
.	1984	0.36 (0.01)	0.33 (0.02)	0.23 (0.01)
.	1985	0.33 (0.01)	0.32 (0.03)	0.26 (0.01)
.	1986	0.27 (0.01)	0.24 (0.01)	0.21 (0.00)
.	1989	0.34 (0.01)	-	-
.	1990-92	0.27 (0.01)	0.25 (0.01)	0.21 (0.00)
θ		0.19 (0.02)	0.11 (0.04)	0.17 (0.02)
γ_ϵ		0.01 (0.01)	-0.02 (0.07)	0.02 (0.04)
γ_η		0.52 (0.03)	0.73 (0.11)	0.52 (0.05)

Note: GMM estimation is based on growth rates and, therefore, cannot estimate a variance for transitory consumption in 1989 given missing consumption data in 1988.

Table B5—Estimates for Older Households

Parameter		QMLE	GMM	
			DWMD	OMD
σ_η	1979-81	0.14 (0.01)	0.08 (0.00)	0.14 (0.00)
.	1982	0.07 (0.03)	0.12 (0.01)	0.10 (0.00)
.	1983	0.13 (0.02)	0.18 (0.01)	0.12 (0.00)
.	1984	0.22 (0.02)	0.17 (0.01)	0.12 (0.00)
.	1985	0.05 (0.05)	0.16 (0.01)	0.16 (0.01)
.	1986	0.12 (0.02)	0.09 (0.01)	-
.	1987	0.00 (0.00)	0.17 (0.01)	0.12 (0.00)
.	1988	0.00 (0.00)	0.11 (0.01)	0.04 (0.00)
.	1989	0.11 (0.02)	0.13 (0.01)	0.16 (0.01)
.	1990-92	0.15 (0.01)	0.08 (0.01)	0.15 (0.00)
σ_ϵ	1979	0.22 (0.01)	0.22 (0.01)	0.09 (0.00)
.	1980	0.23 (0.01)	0.20 (0.01)	0.10 (0.00)
.	1981	0.20 (0.01)	0.19 (0.01)	0.15 (0.00)
.	1982	0.22 (0.01)	0.19 (0.01)	0.15 (0.00)
.	1983	0.23 (0.01)	0.17 (0.01)	0.11 (0.00)
.	1984	0.25 (0.01)	0.23 (0.01)	0.18 (0.00)
.	1985	0.31 (0.01)	0.27 (0.02)	0.18 (0.00)
.	1986	0.28 (0.01)	0.27 (0.01)	0.21 (0.00)
.	1987	0.26 (0.01)	0.22 (0.01)	0.18 (0.00)
.	1988	0.25 (0.01)	0.23 (0.01)	0.18 (0.00)
.	1989	0.26 (0.01)	0.25 (0.02)	0.18 (0.00)
.	1990-92	0.26 (0.00)	0.26 (0.01)	0.19 (0.00)
σ_u		0.08 (0.00)	0.06 (0.01)	0.10 (0.00)
σ_v	1979	0.24 (0.01)	0.22 (0.01)	0.16 (0.00)
.	1980	0.22 (0.01)	0.23 (0.01)	0.16 (0.00)
.	1981	0.23 (0.01)	0.22 (0.01)	0.18 (0.00)
.	1982	0.27 (0.01)	0.25 (0.01)	0.13 (0.00)
.	1983	0.27 (0.01)	0.27 (0.02)	0.21 (0.00)
.	1984	0.31 (0.01)	0.27 (0.01)	0.20 (0.01)
.	1985	0.24 (0.01)	0.23 (0.01)	0.21 (0.00)
.	1986	0.26 (0.01)	0.25 (0.01)	0.22 (0.00)
.	1989	0.27 (0.01)	-	-
.	1990-92	0.29 (0.01)	0.29 (0.01)	0.20 (0.01)
θ		0.20 (0.02)	0.11 (0.03)	0.17 (0.02)
γ_ϵ		0.08 (0.02)	0.06 (0.05)	0.15 (0.03)
γ_η		0.30 (0.04)	0.85 (0.22)	0.19 (0.03)

Notes: GMM estimation is based on growth rates and, therefore, cannot estimate a variance for transitory consumption in 1989 given missing consumption data in 1988. OMD estimates a negative variance (not reported) for the permanent income shock in 1986.