

Estimating household consumption insurance

Arpita Chatterjee¹ | James Morley²  | Aarti Singh²

¹School of Economics, University of New South Wales, Sydney, New South Wales, Australia

²School of Economics, University of Sydney, Sydney, New South Wales, Australia

Correspondence

James Morley, School of Economics, University of Sydney, Sydney, New South Wales, Australia.

Email: james.morley@sydney.edu.au

Funding information

Australian Research Council, Grant/Award Number: DE130100806

Summary

Blundell, Pistaferri, and Preston (*American Economic Review*, 2008, 98(5), 1887–1921) report an estimate of household consumption insurance with respect to permanent income shocks of 36%. In replicating findings for their model and data, we find that this estimate is distorted by a code error and is not robust to weighting scheme for generalized method of moments (GMM) or consideration of quasi maximum likelihood estimation (QMLE), which produces a significantly higher estimate of consumption insurance at 55%. For sub-groups by age and education, the differences between estimates across methods are even more pronounced, and QMLE provides new insights into heterogeneity across households compared to the original study. Monte Carlo experiments using non-normal shocks suggest that consumption insurance estimates for the model are more accurate for QMLE than GMM, including when correcting for bias and especially given a smaller sample such as is only available when looking at sub-groups.

1 | INTRODUCTION

How does idiosyncratic income risk impact consumption when households have only limited access to insurance via formal markets or informal arrangements? In a seminal study, Blundell et al. (2008) (BPP hereafter) constructed a novel panel dataset of income and (imputed) consumption for the Panel Study of Income Dynamics and, by employing generalized method of moments (GMM) minimum distance methods, estimated a low degree of household consumption insurance based on their proposed model for the data. Numerous other influential studies have followed their approach (e.g., Auclert, 2019; Kaplan, Violante & Weidner, 2014) or used their estimate to calibrate structural models (e.g., Kaplan, Mitman, & Violante, 2020). In this replication study, we re-visit the estimation in BPP and consider robustness issues, including to using quasi maximum likelihood estimation (QMLE) instead.

Based on diagonal weights for GMM, BPP reported an estimate of consumption insurance with respect to permanent income shocks of 36%. In replicating findings for their model and data, we find that this estimate is distorted by an error in their original code. In particular, not all sample moments are matched correctly to model-implied moments due to a misplaced transpose. For the BPP data, fixing this error does not alter the GMM estimate based on diagonal weights very much (we find an estimate of 35%), but it could lead to extremely inaccurate estimates in other applications and we found it did so when conducting Monte Carlo experiments, including bootstrap replications used to correct for possible bias. The GMM estimate is also fairly imprecise and highly dependent on weighting scheme. Given optimal weights instead, we find a very different estimate of consumption insurance at 68%. Furthermore, GMM estimates become much less precise and more sensitive to weighting scheme when considering smaller samples of sub-groups by age and education.

For the same data, we find that QMLE leads to a relatively precise and significantly higher estimate of consumption insurance than reported in BPP at 55%. For sub-group analysis, differences in estimates between GMM and QMLE are

even more pronounced. Notably, QMLE produces significantly different estimates across households grouped both by age and by education. For sub-groups based on age, we find that younger households have significantly lower consumption insurance, but BPP did not report estimates in this case due to their apparent imprecision. As in BPP, we find that households with college education have more consumption insurance. But our estimates based on QMLE are quite different than theirs and are again much more precise.

Unlike with the GMM approaches considered in our replication analysis, we note that QMLE avoids the need to estimate a weighting matrix, although it does require specifying a distribution. We directly consider an unobserved components representation of the BPP model and apply the Kalman filter to construct the likelihood under the assumption of normality. However, given the widely noted non-normality of the data (see, e.g., Guvenen, Karahan, Ozcan, & Song, 2015), we interpret this approach as corresponding to QMLE, following White (1982).

To help understand our empirical results, we conduct Monte Carlo experiments for the BPP model with two different sample sizes and highly non-normal shocks drawn from skewed and heavy-tailed empirical distributions. The Monte Carlo results support the use of bias-corrected estimates, especially for GMM. For the same large sample size as the full BPP dataset, we find that QMLE is the most accurate in terms of root mean squared error and reported precision. GMM using diagonal weights is far less accurate, whereas GMM using optimal weights is closer in accuracy to QMLE. However, for a smaller sample size, QMLE performs much better than GMM regardless of weighting scheme. These Monte Carlo results generally reconcile the differences in estimates for the BPP data and support our consideration of QMLE as well as GMM when estimating household consumption insurance.

2 | DATA, MODEL, AND ESTIMATION

For our replication analysis, we use the dataset created by BPP available at <https://www.aeaweb.org/articles?id=10.1257/aer.98.5.1887>. Most of the original data come from the Panel Study of Income Dynamics (PSID) sample of continuously married couples headed by a male (with or without children) aged 30 to 65 years. Income is measured by family disposable income, which includes transfers. For consumption, a similar sample selection strategy is adopted 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, annual nondurable consumption for each household is imputed using estimates of food demand from the CEX. The 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, we follow BPP by using residuals from regressions of log household income and consumption on controls for 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.

We consider an unobserved components representation of the BPP model of idiosyncratic household income and consumption with time-varying volatility for income and consumption shocks. Income and consumption for household i at time t are given as follows:

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

$$c_{i,t} = \gamma_{\eta}\tau_{i,t} + \kappa_{i,t} + v_{i,t}, \quad v_{i,t} \sim i.i.d.(0, \sigma_{v,t}^2), \quad (2)$$

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 $v_{i,t}$ is a transitory consumption shock. The trends are specified as random walks:

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

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

while the structure of the time-varying volatility for each shock is assumed to be deterministic and the same as in BPP.¹ The transitory income shock, $\epsilon_{i,t}$, captures events such as a surprise bonus or temporary leave due to illness, whereas the transitory consumption shock, $v_{i,t}$ 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, whereas 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 main parameters we focus on are γ_η and γ_ϵ , which respectively capture the responses of consumption to permanent and transitory income shocks. The implied “consumption insurance” against permanent income shocks is then $1 - \gamma_\eta$.

To estimate the model in (1)–(4), we propose using QMLE. This involves casting the model into state-space form (see the online appendix) and assuming the shocks are normally distributed in order to use the Kalman filter to calculate the likelihood based on the prediction error decomposition. Again, as noted in the introduction, because the actual shocks are likely to be non-normal at the household level, this approach should be thought of as QMLE. We consider how well QMLE performs relative to GMM for this model despite non-normal shocks in our Monte Carlo analysis.

For purposes of comparison, we also consider GMM estimation using the same moment conditions as in BPP.² We employ two methods for 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, whereas OMD allows for covariance between moment conditions in the weighting matrix. See BPP and their online appendix for details of GMM estimation, including for different weighting schemes.

As discussed below in Section 4, preliminary Monte Carlo analysis reported in our online appendix suggests there can be nontrivial bias in the parameter estimates, even in the large-sample case. Related to Newey and Smith (2004), this finding could reflect asymptotic bias, including due to skewness in the distribution of the data and, in the case of GMM, estimation of the weighting matrix. Thus, to address possible bias, we apply a bootstrap correction to the estimates of γ_η and γ_ϵ . In particular, we take initial parameter estimates from QMLE or GMM and calculate smoothed estimates of shocks based on the Kalman filter/smoothen in order to construct artificial samples of data from a semi-parametric bootstrap data generating process (DGP). For each bootstrap replication, we apply QMLE or GMM, exactly as done for the sample data, and record the resulting estimate. The difference between the average estimate across bootstrap replications and the initial sample estimate (i.e., the “true” parameter for the bootstrap DGP) provides our estimated bias used to correct the initial estimate. To save computational time and motivated by the use of a stopping rule in Davidson and MacKinnon (2007), we continue bootstrap replications until bias estimates for all model parameters remain stable up to three decimals.

3 | EMPIRICAL RESULTS

Table 1 reports bias-corrected estimates for the two parameters of interest related to the responses of consumption to permanent and transitory income shocks when considering the full BPP dataset (see the online appendix for initial estimates of all parameters in the BPP model). The estimates are relatively precise compared to DWMD and lie in between the GMM estimates for both parameters. The DWMD estimate of γ_η is close to what was reported in Table 6 of BPP. However, there are differences in the estimates due to bias correction and an error in the original BPP code.³ The OMD estimates are very different from the DWMD estimates, suggesting a high sensitivity to weights on moments. QMLE

¹Shocks are generally allowed to have different variances in each period, although some variances are set to be the same across certain years to help with econometric identification. See BPP or our online appendix for more details.

²GMM estimation is applied to first differences of income and consumption, whereas QMLE is applied to levels. For QMLE, we assume the initial conditions for the unobserved stochastic trends are exogenous and not drawn from a stationary distribution. However, by placing highly diffuse priors on their values when initializing the Kalman filter and evaluating the likelihood from the second time period of the data in levels, our approach accommodates the lack of a stationary distribution and is equivalent to considering the likelihood for the model in first differences. The benefit of working with levels is that it allows us to incorporate more information in estimation if there are missing observations for households in one of two consecutive time periods, as is often the case in the BPP dataset. Estimates for QMLE based on first differences are similar but less precise due to more missing observations. See the online appendix for more details about estimation in levels and differences, as well as the estimates for QMLE with first differences.

³In the BPP code, `md_AER.prg`, there is a missing transpose on line 289 and an 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 transpose results in a mismatch of some model-implied moments with sample moments and leads to substantially different estimates on average in Monte Carlo and bootstrap analysis.

and OMD imply a significantly higher degree of consumption insurance with implied point estimates between 55–68% compared to 35% for DWMD. Estimates for the effect of a transitory income shock are all insignificantly different than zero, but the estimate for QMLE is the most precise.

Table 2 reports bias-corrected estimates of the consumption response parameters when considering sub-groups by age and education. As with the full dataset, the results vary substantially by estimation method, with the estimates for QMLE being relatively precise compared to DWMD and often taking on values in between the GMM estimates. The DWMD estimates for younger (ages 30–47 years) and older (ages 48–65 years) households are not precise enough to show a significant difference from each other and, indeed, BPP did not report estimates by age, but explicitly noted their imprecision in footnote 31 of their paper. Meanwhile, the DWMD estimate of γ_η by education is again similar to what was reported in Table 6 of BPP, with differences in estimates due to bias correction and the original code error noted above. In terms of the effect of permanent income shocks, only the QMLE results correspond to the intuitive pattern of higher consumption insurance for both older and college-educated households. Contrary to what must be true in population, DWMD implies very low consumption insurance for both younger and older households compared to the estimate for all households in Table 1, with younger households having counter-intuitively higher consumption insurance than older households at 29% versus 13%. OMD suggests the counter-intuitive result that households with no college education have substantially higher consumption insurance of 72% versus 57% for households with college education. Again, the estimates display a high sensitivity to weights on moments, particularly with implied consumption insurance for older households ranging from 13% for DWMD to 81% for OMD and for households with no college education

TABLE 1 Bias-corrected estimates of consumption responses for the full dataset

Parameter	QMLE	DWMD	OMD
γ_η	0.45 (0.04)	0.65 (0.09)	0.32 (0.03)
γ_ϵ	0.03 (0.02)	-0.04 (0.04)	0.05 (0.03)

Note: Point estimates are reported, with standard errors in parentheses. The point estimates are corrected for bias using semi-parametric bootstrap replications based on initial parameter estimates and smoothed estimates of shocks until bias estimates remain stable up to three decimals. The standard errors for QMLE are based on the Huber–White sandwich formula using numerical derivatives, and the standard errors for GMM are calculated as in BPP. There are a total of 1765 households and 15 years of data in levels but with many missing observations.

TABLE 2 Bias-corrected estimates of consumption responses for sub-groups

Parameter	YOUNGER			OLDER		
	QMLE	DWMD	OMD	QMLE	DWMD	OMD
γ_η	0.55 (0.06)	0.71 (0.11)	0.50 (0.05)	0.25 (0.06)	0.87 (0.22)	0.19 (0.03)
γ_ϵ	-0.01 (0.00)	-0.14 (0.07)	-0.06 (0.04)	0.11 (0.01)	0.07 (0.05)	0.20 (0.03)
Parameter	NO COLLEGE			COLLEGE		
	QMLE	DWMD	OMD	QMLE	DWMD	OMD
γ_η	0.64 (0.08)	0.95 (0.17)	0.28 (0.04)	0.29 (0.04)	0.45 (0.09)	0.43 (0.04)
γ_ϵ	0.02 (0.01)	0.05 (0.06)	0.24 (0.03)	0.05 (0.02)	-0.15 (0.05)	-0.34 (0.04)

Note: Point estimates are reported, with standard errors in parentheses. The point estimates are corrected for bias using semi-parametric bootstrap replications based on initial parameter estimates and smoothed estimates of shocks until bias estimates remain stable up to three decimals. The standard errors for QMLE are based on the Huber–White sandwich formula using numerical derivatives, and the standard errors for GMM are calculated as in BPP. For the sub-groups based on education, there are 883 households classified as “No College” and 882 households as “College.” For age, there are 1413 households classified as “Younger” (ages 30–47 years) and 708 households classified as “Older” (ages 48–65 years). There are 15 years of data in levels but with many missing observations and changes in classification by age.

ranging from 5% for DWMD to 72% for OMD. In comparison, the sub-group estimates for QMLE imply plausible levels of consumption insurance based on age (45% for younger versus 75% for older) and education (36% for no college versus 71% for college) given the estimate of 45% for all households. Finally, the estimates for the effect of a transitory income shock are always the most precise for QMLE.

4 | MONTE CARLO ANALYSIS

Although the QMLE estimates in Tables 1 and 2 are relatively precise, it is a reasonable question whether this might be a false precision due to the normality assumption made when constructing the likelihood for the unobserved components representation of the BPP model. To address this question, we consider Monte Carlo experiments where the DGP corresponds to the BPP model with shocks drawn from their skewed and heavy-tailed empirical distributions based on the estimated model in the previous section. In particular, 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 negative skewness and an extremely high degrees of kurtosis. Simulated data using these as the “true” shocks exhibit similar skewness and kurtosis to the sample data. For our experiments, 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 missing observations as the original BPP data. Then to evaluate the accuracy of a particular estimator, we consider root mean squared error (RMSE). 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 between the estimates across methods in a given sample.

For this particular DGP, preliminary Monte Carlo analysis reported in the online appendix revealed nontrivial bias in the initial estimates for both QMLE and GMM.⁴ Thus, similar to the empirical application, we correct parameter estimates using semi-parametric bootstrap replications based on average initial estimates and smoothed estimates of shocks given the average large-sample QMLE initial estimates. A full bootstrap correction for each simulation would be computationally prohibitive, but the bias corrections based on average initial estimates should at least be indicative of what the corrections would be for each simulation. Also, by using average large-sample QMLE initial estimates to estimate empirical shocks, we are keeping the estimated shocks closer to the “true” shocks used in the Monte Carlo experiments than might occur in a full bootstrap and, therefore, may be somewhat overstating the accuracy of the bias correction in practice, again especially for GMM. Despite this, we find there is generally still some remaining bias for the estimators, especially for DWMD in the case of γ_η . However, the bias corrections are successful to the extent that RMSEs are only 0.01 larger than the standard deviations of the estimators in the DWMD case for γ_η and the same to two decimals in all of the other cases. Notably, these findings imply that the comparative results for the different estimators would hold even if the bias could be perfectly corrected in each case.

Table 3 reports on the accuracy of the estimators for different sample sizes. Focusing first on the results given a large sample size that is the same as for the full BPP dataset, QMLE performs best in terms of RMSE, with OMD and especially DWMD considerably less accurate in terms of γ_η . The estimates also differ substantially across methods in a given sample, with root mean squared differences generally as large or sometimes even larger than the RMSEs. Particularly notable is how much the GMM estimates across weighting schemes differ from each other in a given sample. Given that the correct model is specified in the Monte Carlo experiments, we note that this finding is consistent with Altonji and Segal's (1996) argument that large differences in GMM estimates across weighting schemes do not necessarily provide evidence of model misspecification.

Looking back at Table 1, these Monte Carlo results can help explain some of the empirical results. In particular, the reasonably precise standard errors for the QMLE estimates in Table 1, which are based on the Huber–White sandwich formula ($\sqrt{\text{diag}(H^{-1}GG'H^{-1})}$ for Hessian H and Jacobian G) given misspecification of the distribution when constructing the likelihood, are supported by the standard deviations of the estimators for the large-sample case in Table 3. By contrast, the standard errors for GMM in Table 1 appear to understate sampling uncertainty, even in the relatively imprecise case of DWMD. Also, the differences between the QMLE and GMM estimates in Table 1 are generally less

⁴Based on our bootstrap analysis and further Monte Carlo analysis, biases appear to vary considerably with particular DGPs, especially for GMM. This sensitivity for GMM is likely due to small sample and asymptotic biases related to estimation of the weighting matrix, as discussed in Altonji and Segal (1996) and Newey and Smith (2004), respectively.

TABLE 3 Properties of bias-corrected estimators

Parameter	Estimator property	Large sample			Small Sample		
		QMLE	DWMD	OMD	QMLE	DWMD	OMD
$\gamma_\eta = 0.50$	RMSE	0.04	0.18	0.08	0.07	0.39	0.36
	Bias	0.01	0.06	-0.01	0.01	0.11	-0.01
	Standard deviation	0.04	0.17	0.08	0.07	0.38	0.36
	Difference from DWMD	0.17	-	0.18	0.38	-	0.53
	Difference from OMD	0.08	0.18	-	0.36	0.53	-
$\gamma_\epsilon = 0.10$	RMSE	0.01	0.03	0.04	0.02	0.05	0.07
	Bias	0.00	0.00	-0.02	-0.01	-0.01	0.01
	Standard deviation	0.01	0.03	0.04	0.02	0.05	0.07
	Difference from DWMD	0.03	-	0.04	0.05	-	0.08
	Difference from OMD	0.03	0.04	-	0.07	0.08	-

Notes: RMSE, bias, standard deviation, and (root mean squared) differences for estimators are calculated across 500 simulations. Parameter estimates were corrected for bias using 100 semi-parametric bootstrap replications based on initial average parameter estimates and smoothed estimates of shocks given the average large-sample QMLE estimates. The large sample has 32,547 observations and the same missing observations as the full BPP dataset. The small sample has 11,437 observations and the same missing observations as the sub-group of older households in the BPP dataset.

than twice the average differences across the Monte Carlo simulations, although the differences are more than twice as large, but still within simulated ranges, in the case of QMLE and DWMD for γ_ϵ .

Table 3 also reports on the accuracy of different estimators given a small sample size that is the same as for the sub-group of older households in the BPP dataset. In light of the well-known concerns about OMD in small samples highlighted in Altonji and Segal (1996), our aim with this experiment is to understand how the performance of QMLE and GMM compare given a smaller sample size, with the number of observations for older households being the smallest among the sub-groups.⁵ As in the large-sample case, QMLE performs best in terms of RMSE, but the improvements over GMM are even more pronounced, especially compared to OMD and notably in terms of γ_ϵ as well as γ_η .⁶ Meanwhile, the estimators again differ substantially from each other in a given sample, especially in the case of γ_η across weighting schemes for GMM.

Looking back at Table 2, the Monte Carlo results for the small sample case in Table 3 help explain why the estimates are so different across estimation methods given very large average differences across the Monte Carlo simulations and raise strong concerns about the reported precision of the GMM estimates given the large standard deviations with a small sample size. The puzzle of lower implied consumption insurance for both younger and older households based on DWMD in Table 2 compared to what was found for all households can be resolved by the substantial deterioration in accuracy of DWMD with the smaller sample size, as well as, at least in part, by the apparent remaining large positive bias for even the bias-corrected DWMD estimator of γ_η . By contrast, the Monte Carlo results provide some reassurance about the reported precision of the QMLE estimates in smaller samples, notably so given their standard errors always being smaller than for DWMD.

⁵Out of the total 32,547 observations of data for all of the households in the dataset, there are 15,735 observations for households with no college education versus 16,812 observations for college-educated households and 21,110 observations for younger households versus 11,437 observations for older households.

⁶The deterioration in the performance of GMM given a smaller effective sample size can also be seen by considering what happens when allowing for a possible structural break in γ_η and γ_ϵ near the middle of the estimation sample. This corresponds to smaller effective T for estimating each parameter. As in the case of the smaller N considered in Table 3, the RMSEs increase substantially, especially for γ_η and for OMD in particular. Meanwhile, consistent with QMLE implicitly linking moments over time without requiring estimation of a weighting matrix, there appears to be a relative loss of efficiency when allowing for a possible structural break in estimating time-invariant parameters in the sense that the RMSE for QMLE also increases and, unlike most of the cases for GMM, as much as it does when reducing N in Table 3. Full results for this Monte Carlo experiment are reported in the online appendix.

5 | CONCLUSIONS

We have examined the replicability of BPP's low estimated degree of consumption insurance in both a "narrow" and "wide" sense. In a narrow sense, there is an error in BPP's code that we have corrected to avoid extremely inaccurate estimates in Monte Carlo analysis and other settings. In a wider sense, we find that the result of a low estimated degree of consumption insurance using GMM based on diagonal weights is not replicable across different estimation methods, even given the same data, with both QMLE and GMM based on optimal weights implying higher and more precise estimates of consumption insurance. Estimates for sub-groups are also sensitive to estimation method, with QMLE in particular suggesting both intuitive and significant heterogeneity across households grouped by age and education. Monte Carlo analysis assuming non-normal shocks suggests greater accuracy of QMLE versus GMM in estimating consumption insurance using BPP's model.

We believe our analysis makes two contributions to the literature on household consumption insurance. First, we have provided evidence that consumption insurance is significantly higher than previously reported for BPP's dataset. Thus, our finding of higher consumption insurance is consistent with a key implication of Kaplan and Violante (2010), albeit for the different reason of imprecision of the GMM estimator using diagonal weights, rather than model misspecification. Second, we have demonstrated the feasibility and apparent accuracy of QMLE in a panel setting with highly non-normal shocks and a relatively small sample size. A widely claimed reason why heterogeneous 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. Thus, improved performance in estimating this key parameter in a small sample setting is particularly important and we plan further analysis of the comparative accuracy of QMLE, including when allowing for different model specifications to address Kaplan and Violante's (2010) concerns, in future research.

ACKNOWLEDGMENTS

We thank the Co-Editor, Thierry Magnac, as well as 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 and Michaela Haderer provided excellent research assistance. We are grateful for the financial support from the Australian Research Council grant DE130100806 (Singh). The usual disclaimers apply.

OPEN RESEARCH BADGES



This article has been awarded Open Data Badge for making publicly available the digitally-shareable data necessary to reproduce the reported results. Data is available at [<http://qed.econ.queensu.ca/jae/datasets/chatterjee001/>].

ORCID

James Morley  <https://orcid.org/0000-0003-2380-6747>

REFERENCES

- Altonji, J. G., & Segal, L. M. (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, 2333–2367.
- Blundell, R., Pistaferri, L., & Preston, I. (2008). Consumption inequality and partial insurance. *American Economic Review*, 98, 1887–1921.
- Davidson, R., & MacKinnon, J. G. (2007). Bootstrap tests: How many bootstraps? *Econometric Reviews*, 19, 55–68.
- Guvenen, F., Karahan, F., Ozcan, S., & Song, J. (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.
- Kaplan, G., Mitman, K., & Violante, G. L. (2020). The housing boom and bust: Model meets evidence. *Journal of Political Economy*, 128, 3285–3345.
- Kaplan, G., & Violante, G. L. (2010). How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics*, 2, 53–87.

- Kaplan, G., Violante, G. L., & Weidner, J. (2014). The wealthy hand-to-mouth. *Brookings Papers on Economic Activity*, 1, 77–138. Spring.
- Newey, W. K., & Smith, R. J. (2004). Higher order properties of GMM and generalized empirical likelihood estimators. *Econometrica*, 72, 219–255.
- White, H. (1982). Maximum likelihood estimation of misspecified models. *Econometrica*, 50, 1–25.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

How to cite this article: Chatterjee A, Morley J, Singh A. Estimating household consumption insurance. *J Appl Econ*. 2021;36:628–635. <https://doi.org/10.1002/jae.2820>