How Much Consumption Insurance in the U.S.?*

Dmytro Hryshko[a], Iourii Manovskii[b]

[a] *University of Alberta, Department of Economics, 8–14 HM Tory Building, Edmonton, AB T6G 2H4*
[b] *University of Pennsylvania, Department of Economics, Suite 150, PCPSE Building, 133 South 36th Street, Philadelphia, PA 19104*

Abstract

We identify two sets of households in the Panel Study of Income Dynamics (PSID) differing dramatically in their income and consumption dynamics, although both should be equally representative. The degree of consumption insurance in each subsample is consistent with the standard incomplete-markets model's prediction. We contrast PSID and administrative earnings data and study the patterns in international datasets modeled on the PSID. We find an important role of differential attrition based on the dynamic properties of incomes in inducing the differences and identify PSID households providing a better guide to income dynamics and consumption insurance in the U.S.

JEL: D12, D15, D31, E21

Keywords:

Consumption insurance, income processes, incomplete markets models, attrition, Panel Study of Income Dynamics

*We are grateful to David Johnson and Frank Stafford from the PSID for a thorough reading of the manuscript and detailed comments. We also benefited from the comments of Yuriy Gorodnichenko, an anonymous referee, and participants at the AEA Annual Meetings session "Celebrating 50 years of the PSID," PSID Annual User Conference, NBER Summer Institute Working Group on Consumption, Barcelona GSE Summer Forum, CEA Annual Meetings, Society for Economic Dynamics Meetings, Asian Meeting of the Econometric Society, CEPR ESSIM Conference, Workshop on Inequality at St. Louis FED, SAET Meetings, 2nd European Midwest Micro/Macro Conference, Econometric Society Virtual World Congress, REDg Workshop at Universidad Carlos III de Madrid, Workshop on Household Finance and Retirement Savings at Collegio Carlo Alberto, Deutsche Bank's International Conference on Household Finance and seminar participants at the Universities of Alberta, Delaware, Saskatchewan, Concordia, and Copenhagen Business School. Support from NSF Grants SES–1357903 and SES–1824520, and SSHRC IG grant 435-2018-1275 is gratefully acknowledged.

Email addresses: dhryshko@ualberta.ca (Dmytro Hryshko), manovski@econ.upenn.edu (Iourii Manovskii)

1. Introduction

Economic theory tightly links the evolution of households' consumption to the dynamic properties of their incomes. Empirical measures of joint income and consumption dynamics are thus essential for understanding households' behavior, for developing and disciplining the economic theory, and for the evaluation of policy changes that affect households' budgets. Indeed, consumption represents the dominant share of GDP, and knowledge of the marginal propensities of consume from persistent and transitory shocks to households' incomes are crucial for understanding the macroeconomic impact of changes in tax and transfer policies, labor or credit market reforms, and for designing stabilization and income-maintenance policies.[1]

Much of our knowledge of the joint income and consumption dynamics at the household level in the U.S. is based on the data from the PSID. For example, Blundell et al. (2008) (BPP hereafter), use these data to estimate consumption insurance for permanent and transitory idiosyncratic income shocks, i.e., the fraction of those shocks that does not translate into movements in consumption. This direct evidence has become a central empirical benchmark for calibrating or for assessing the performance of quantitative models of household consumption and saving choices. It indicates that household consumption is excessively insured against permanent shocks to net household incomes relative to the prediction of the standard incomplete-markets model. This finding spurred active ongoing research on the ways of modifying the canonical model to bring its predictions in line with the degree of insurance measured in the PSID data.

In contrast, we provide evidence that the degree of insurance and income dynamics vary quite dramatically and systematically across two sets of households in the PSID. Conditional on income dynamics, the estimated insurance against permanent shocks for both types of households is in line with the prediction of the standard incomplete-markets model.

To understand the distinction between these two types of households, it is necessary to briefly describe PSID data. The PSID started in 1968 with a representative cross-section of U.S. households. These households, as well as their children, grandchildren, etc., are followed over time and form the PSID sample. The idea is to learn about the population at large by following this branch of the U.S. family tree. Note that individuals who become married to the core or "sample" PSID members are not considered to be part of the branch, and are labeled as "nonsample" individuals by the PSID. The information on these individuals is collected while they are attached to a core PSID member, but they are not followed either before or after this period of attachment. An analogy might be helpful in highlighting the distinction. Imagine all individuals who were originally interviewed by the PSID in 1968 were endowed with the "PSID gene." All individuals born to or adopted by somebody with the "PSID gene" acquire the gene themselves and are followed by the PSID. The "gene" is not passed to the spouse. Thus, "sample" PSID members are the ones with the "gene" and "nonsample" PSID members are the ones without the "gene."

We find that households headed by sample males (who have the "PSID gene") are characterized by a virtually complete pass-through of permanent shocks to net family incomes to consumption. In contrast, the households headed by nonsample males (who do not have the "PSID gene") show a dramatically higher degree of insurance against permanent shocks.

The large discrepancy in the degree of insurance is not explained by observed cross-sectional

[1]See Arellano et al. (2017) and Daly et al. (2021) for references to the literature.

differences among sample and nonsample households. (We refer to households headed by sample males as "sample households" and households headed by nonsample males as "nonsample households.") In one of the most comparable groupings, we consider the sample males and females (who all have the "PSID gene") who marry after 1968. One can roughly describe the two groups as consisting of sons and daughters of the original PSID sample, with their spouses being nonsample females and nonsample males, respectively. As can be expected, these groups are virtually identical with respect to all cross-sectional observables. Yet, nearly 90% of permanent income shocks are passed through to consumption of households headed by PSID sons, while only 46% of permanent income shocks are passed through to consumption of households headed by nonsample PSID sons-in-law, married to PSID sample daughters.

While our finding on the dramatic difference in the degree of insurance is novel in the literature, the finding that there is little cross-sectional difference among comparable sample and nonsample individuals in the PSID is consistent with early studies by Becketti et al. (1988) and Lillard (1989). However, to our knowledge, the literature has never compared the dynamic properties of income or earnings among sample and nonsample PSID individuals or households. This is a significant omission as the dynamic properties of incomes are the crucial ingredients in the analysis of consumption insurance and, indeed, in any model with incomplete insurance markets. We present evidence of substantial differences. Specifically, while the permanent component of the income process among sample households is well described by a random-walk model, the nonsample households have a far less persistent permanent component of income. Although the literature traditionally considers the pooled sample, we argue that it might be essential to recognize the heterogeneity in income dynamics between the two groups.

We show that assuming a common income process, and in particular, that the persistent income component follows a random walk, contributes to the well-known discrepancy between the estimates of the household income process targeting the moments in levels and differences. Specifically, using a random-walk process common to the two groups results in inflated estimates of the variance of permanent shocks when estimation targets the income moments in differences (as is standard in this literature). As these shocks are not truly permanent, consumption responds relatively little to them, as predicted by the standard theory. This results in some (but not very large) overestimation of the degree of insurance of permanent shocks.

Perhaps more importantly, as pointed out by Kaplan and Violante (2010) and Blundell (2014), correctly measuring the persistence of income innovations is key for interpretation of the resulting insurance coefficients. For example, the findings of BPP, who considered only the combined sample and assumed a random walk process, suggest a considerably higher degree of insurance against permanent income shocks relative to the predictions of the standard models of imperfect consumption risk-sharing via self-insurance through saving and borrowing. Our estimates, based separately on sample and nonsample households, point to a different conclusion. We show that the amount of insurance achieved by nonsample households is roughly in line with the prediction of the standard model given that "permanent" shocks to their incomes have only limited persistence. On the other hand, the point estimate of almost no insurance against truly permanent income shocks achieved by sample households suggests lower insurance than implied by the theory. However, this point estimate in the data has a fairly sizable standard error and is not statistically different from the prediction of the standard model.

While the difference in persistence of income shocks can rationalize the differential degree of insurance between sample and nonsample households in the PSID, the existence of highly systematic differences in their income processes appears quite unexpected. A random sample of U.S. households in 1968 would be expected to have random samples of sons and daughters. But then, one would expect the two sets of households formed by them to be also random and similar not only with respect to their cross-sectional characteristics but also with respect to their income dynamics. Yet, the differences we document are so large and consistent, that they appear highly unlikely to be induced by sampling noise. But what could be behind these differences? Are they a manifestation of a systematic measurement error? Are they induced by the PSID study design? These are important questions because the PSID is the foundation of our knowledge of household income and consumption dynamics. Most quantitative incomplete markets models in the literature are either estimated using the PSID data or take PSID estimates of the income process as the key input. The PSID is also widely used in many other areas of social sciences, and it has served as a model for designing the datasets in numerous other countries.

To address these questions, we first use confidential earnings data from the U.S. Social Security Administration (SSA) merged with the survey records by the Health and Retirement Study (HRS). We compare the earnings dynamics of the original cohorts first surveyed in 1968 and followed over time by the PSID to the earnings dynamics in the administrative data retrospectively extracted for the corresponding cohorts of workers present in the HRS in 1998. We find that male and family earnings in the PSID have a considerably higher persistence than earnings in the administrative data for comparable cohorts.

One possible explanation for this finding is that while the original random PSID sample was representative cross-sectionally, it was not representative with respect to the earnings dynamics. It is difficult to definitively rule out this possibility but a strong piece of evidence against it is that we find qualitatively similar differences in income persistence among sample and nonsample households in the data from the German Socio-Economic Panel, the British Household Panel Survey, the Household, Income and Labour Dynamics in Australia, the Korean Labor and Income Panel Study, and the Swiss Household Panel, all of which used the PSID design as a template. As these datasets also started with random cross-sectional samples, it appears unlikely that all these samples would be biased with respect to the earnings dynamics in the same way as the PSID. Instead, we find that the PSID and all those datasets share important properties of sample attrition. In particular, sample males are more likely to attrit than females, and the attrition is related to income dynamics with attritors having lower income persistence. This leads to a stronger selection among sample households, inducing higher persistence of their earnings relative to nonsample households and relative to the retrospective administrative data that do not feature such selection. As the set of nonsample households features lower attrition in survey data it is likely more representative. We indeed find that the income properties of nonsample families line up quite closely with those based on a large and arguably representative sample of families in the Current Population Survey (CPS) as well as U.S. SSA data, while this is not the case for sample PSID families. These findings suggest that the degree of consumption insurance estimated on the set of nonsample PSID families provides a better guide to the extent of insurance available to a representative U.S. household.

The rest of the paper is structured as follows. Section 2 describes the PSID data used; Section 3 documents differences in consumption insurance among sample and nonsample households in the PSID; Section 4 models and estimates income processes for sample and nonsample families; Section 5 compares empirical estimates with the predictions of a standard incomplete markets model; Section 6 investigates the reasons for different income and consumption dynamics between sample and nonsample families; Section 7 explores what can be inferred about income dynamics and the extent of consumption insurance in the U.S.; and Section 8 concludes.

2. Data

At the core of our study is the dataset used and made publicly available by Blundell et al. (2008). We augment these data with additional variables extracted from the PSID, most importantly, the ones that indicate whether a particular individual is a sample or nonsample PSID member. As summarized above, the PSID started in 1968 interviewing about 4,800 families; 2,930 of them were nationally representative (SRC sample), while the rest belonged to income-poor households (SEO sample). Members of these original households, as well as their descendants (children, grandchildren, etc.), are referred to as sample members by the PSID, whereas individuals entering the PSID due to marriage or living arrangements with the original sample members are labeled nonsample (e.g., a male marrying a sample female after 1968 will become a head of household and will be treated as a nonsample PSID member). The major distinction of nonsample persons is that the PSID typically makes no attempts to contact these individuals once they separate from a sample person. While the PSID provides weights for sample individuals, which makes it possible to achieve nationally representative results using individual data, the nonsample members have zero (longitudinal, and cross-sectional up to 1997) weights in the PSID.

Unless explicitly stated otherwise, we maintain all of the sample restrictions made by BPP, and we refer the reader to that paper for the detailed discussion of the motivation behind those restrictions. Briefly, the main objective was to focus on a sample of continuously married couples headed by a male (with or without children). BPP aimed to restrict the sample to households with male heads of ages 30–65 who do not change their marital status and are continuously married to the same spouse during 1979–1993. The focus on continuously married couples is to eliminate the potential effects of dramatic family composition change, such as divorce. As we discuss in Online Appendix I.1, the actual implementation of data construction allows for sample females (but generally not sample males) to marry and divorce inside the 1979–1993 window. However, this aspect of sample construction is not responsible for the data patterns that motivate this paper.

Our initial sample is the same as in BPP. It excludes SEO families and contains 1,765 households, among them 965 families headed by sample males, and 800 families headed by nonsample males. Various modifications to this sample will be considered and explained below.

3. Documenting Insurance Differences among Sample and Nonsample Households

In this section, we document large and robust differences in the measured insurance against permanent income shocks among sample and nonsample households. Due to space constraints, we report numerous additional results in Online Appendix I. We begin by briefly summarizing the empirical measures of insurance proposed and implemented by BPP.

3.1. Methodology

BPP assume that household i's idiosyncratic net family income, y_{it}, is composed of a fixed effect, α_i, a random-walk permanent component, $p_{it} = p_{it-1} + \xi_{it}$, and a transitory component modeled as a moving average process of order one, $\tau_{it} = \epsilon_{it} + \theta\epsilon_{it-1}$.[2] Idiosyncratic income and idiosyncratic consumption are residuals from panel regressions of the logs of net family income, and (imputed) nondurable consumption on a number of observables (listed in BPP).

BPP consider the following equation for residual consumption growth:

$$\Delta c_{it} = \phi\xi_{it} + \psi\epsilon_{it} + \zeta_{it} + \Delta u_{it}, \tag{1}$$

where Δc_{it} is household i's consumption growth at time t, ξ_{it} is the permanent shock to household i's disposable income, ϵ_{it} is the transitory shock, ζ_{it} is an innovation to consumption growth independent of the two income components, and u_{it} is an i.i.d. measurement (and imputation) error in nondurable consumption. All of the shocks are assumed to be independent of each other. Coefficients ϕ and ψ measure the transmission of permanent and transitory shocks to consumption. Conversely, $1-\phi$ and $1-\psi$ measure the extent of household consumption self-insurance against permanent and transitory shocks to net income due to accumulated assets. For other measures of income, $1-\phi$ and $1-\psi$ will have different interpretations.[3]

Following BPP, we estimate ϕ and ψ, the parameters of the income process (the moving-average parameter and the time-varying variances of permanent and transitory shocks), the variance of random growth in consumption, σ^2_ζ, and time-varying variances of measurement (and imputation) error in consumption using the minimum-distance method. The parameters are recovered by minimizing the weighted distance between the full set of autocovariances of income and consumption growth, the full set of their cross-covariances, and their model counterparts. The weights are obtained from the diagonal weighting matrix constructed from the diagonal of the variance-covariance matrix of the data moments.

3.2. Benchmark Consumption Insurance Estimates

In column (1) of Table 1 Panel A, we tabulate the results based on the full sample of 1,765 PSID families. As reported by BPP, consumption is almost perfectly insulated from transitory shocks ($\hat{\psi}$ is close to zero) while about 36% of permanent shocks are insured ($\hat{\phi} = 0.64$).

Next, we consider separately the households headed by sample and nonsample males. The results are in columns (2) and (3): sample families insure only about 6% of permanent shocks while nonsample families insure up to 57% of permanent shocks; the difference in the insurance of permanent shocks between sample and nonsample families is significant at the 1% level whereas the difference in the insurance of transitory income shocks is not statistically significant at any conventional level.

3.3. Consumption Insurance among Households Formed by PSID Sons and Daughters

Online Appendix I.1 documents that the data selection procedure in BPP treats sample and nonsample households differently. In particular, households headed by nonsample males can be formed through marriage or end in divorce inside the 1979–1993 sample window while

[2]We do not consider other alternatives to this income process, such as those allowing for heterogeneous income profiles as in Guvenen (2007) and Guvenen (2009), since we use BPP as an organizing framework in this paper.

[3]For instance, Blundell et al. (2016) measure the extent of consumption insurance against permanent and transitory shocks to *husband's wages* due to changes in own and spousal labor supply, accumulated assets, and the tax and transfer system, whereas Arellano et al. (2017) study consumption insurance against persistent and transitory shocks to *household earnings* due to assets, and the tax and transfer system.

this is generally not allowed for the households headed by sample males. Although we show that this differential selection based on marriage and divorce does not drive the differences in insurance between sample and nonsample families, those families may still potentially differ on a variety of characteristics. To put selection of sample and nonsample families on an equal footing, we allow PSID sample males to (re-)marry and divorce during 1979–1993, keeping data for each newly-formed couple with the same sample male head in the final dataset.[4] We further split the resulting dataset of sample families into those who had been married by 1968 and stayed married until they were last seen in 1979–1993, and those who, similarly to nonsample families, married or re-married in 1969 or later. We label them "Sample orig." and "Sample sons," respectively, because the latter sample is dominated by the sons of original PSID households in addition to a few original sample members who married after 1969. In total, we have 669 original sample families, 854 families formed and headed by sample "sons," and 814 families formed by sample "daughters" and headed by their nonsample husbands (the latter group is the same as the set of nonsample families).[5]

Table A-4 in Appendix I.3 reports means of various observables for the resulting three subsamples. Original sample families are older and thus different from the other two subsamples with respect to many cross-sectional characteristics. In contrast, households formed by sample sons and sample daughters are very similar with respect to age, average nondurable consumption, net family income, head's earnings, assets, head's and wife's hours worked, incidence of unemployment, occupation and industry switching, precision of food and income measurement, immigrant status of the head, incidence of owning a business and homeownership rates, among many other dimensions. Figure A-1 in Appendix I.3 documents that a wide range of such variables used in a LASSO regression does not predict the nonsample status of the family (among the set of households formed by sample sons and daughters). This confirms that the families formed by sample sons and daughters do not significantly differ on a wide range of observable characteristics.

Despite the sets of households formed by sample sons and sample daughters being nearly identical with respect to their cross-sectional characteristics, the results in columns (1) and (2) of Table 1, Panel B indicate that they differ dramatically in the degree of consumption insurance against permanent income shocks, with nonsample households (i.e., the ones formed by PSID sample daughters) being significantly better insured.[6]

In contrast, despite being different on many observable dimensions, original sample families and younger sample families formed mostly by their sons have quite similar insurance against permanent income shocks – columns (3) and (1), respectively. In column (4), therefore, we group them obtaining similar in magnitude but a more precise estimate of the insurance coefficient for permanent income shocks.

In columns (5) and (6) of Table 1 Panel B we restrict the samples even further to households

[4]This selection is also recently used in Blundell et al. (2016).

[5]Relative to the original BPP data, additional fourteen nonsample families are added as nonsample males from those families changed their marital status during 1979–1993 and were followed by the PSID after the change (some nonsample individuals were designated as followable since 1990).

[6]The characteristic that significantly differs between sample and nonsample households is that the female spouse is more likely to be responsible for filling out PSID questionnaires for the nonsample households. Online Appendix I.4 details various empirical exercises that make us conclude that the respondent status is of no importance for our results on the differential consumption insurance among sample and nonsample households.

formed by brothers and sisters who are children of the original PSID sample families and thus share some unobservable characteristics imparted by their common background. Although these sibling samples are smaller, the patterns they reveal remain the same – insurance against permanent income shocks is much higher for the families of sisters.

4. Income Processes of Sample and Nonsample Households

The body of evidence presented so far points to substantial differences in insurance against permanent shocks to net family incomes for sample versus nonsample families. Underlying these findings was the standard maintained assumption that the income process is the same across sample and nonsample households, assumed to consist of a random walk permanent component and an MA(1) transitory component as in BPP. The assumption of common income process appears consistent with the evidence of cross-sectional similarity of sample and nonsample households across observable characteristics. However, the dynamic properties of incomes of sample and nonsample households have never been examined in the literature, to our knowledge. If they differ, the estimates of insurance might be biased. Moreover, the interpretation of insurance coefficients depends on the dynamic properties of shocks to household budgets. For example, the insurance of about 60% of permanent shocks, found for nonsample families, appears excessive for consumption models with incomplete markets when the permanent component is a random walk process, but the value may be reasonable for the income process with low persistence of shocks to the permanent component. In this section, we provide evidence that income processes indeed differ systematically across the two types of households and in the next section we reinterpret the differences in consumption insurance in light of the differences in income processes.

Interpreting the relationship between income dynamics and consumption insurance is only possible within the context of a model. For the narrow task of understanding the degree of consumption insurance, the persistence of permanent and transitory shocks plays a key role while the variances of these shocks are less important (Carroll, 2009). However, to provide a relevant framework for interpreting the data, the model must reflect the correct incentives, and, thus, it has to replicate the wealth and income distributions in the data. To achieve this objective, accurately measuring the variances of income shocks in the data is necessary (Carroll, 1992; Heathcote et al., 2010). Thus, our ultimate goal in this section is to provide an accurate measurement of both persistences and variances of permanent and transitory shocks experienced by sample and nonsample families.

4.1. Descriptive Evidence on Different Income Dynamics between Sample and Nonsample Families

Panel (a) of Figure 1 plots the autocorrelation functions of net family incomes for households headed by sample and nonsample males.[7] As can be seen, income dynamics differs markedly between the two sets of families. Although the income process of nonsample families appears less persistent, this evidence is not conclusive as the figures are also affected by potentially different variances of shocks.

[7] Autocorrelation of order j in year t is calculated as $\frac{E[y_{it}y_{it+j}]}{\sqrt{E[y_{it}y_{it}]}\sqrt{E[y_{it+j}y_{it+j}]}}$. In the figure, for each j, we plot autocorrelations averaged over all t's.

We next examine the moments constructed from residual income growth which were targeted in the minimum-distance estimation above. Panel (b) of Figure 1 suggests noticeable differences in these moments across sample and nonsample families.[8] In particular, for nonsample families, the variance of income growth rates had not experienced any clear trend.[9] The latter fact manifests itself in the correlation of just 27% between the variance of income growth rates for the two subsamples.[10] There are also some differences in the trends for the first- and second-order autocovariances, but the most important is the plot for the third-order autocovariance. Under the null of the income process in BPP – net family income is the sum of a random walk and an MA(1) component – the autocovariances beyond second order should not differ from zero. While the average third-order autocovariance is not significantly different from zero for the sample families, it is statistically different from zero, at less than 1% level, for the families headed by nonsample males. As the minimum-distance estimation targets not only the third-order but all of the higher-order autocovariances of income growth, we next test if all higher-order autocovariances above the second order are jointly equal to zero, as in Abowd and Card (1989). The p-value of the test for sample families is 42%, but less than 2% for nonsample families. These results suggest that the random-walk plus an MA(1) component is an adequate description of the income process for the sample families. However, the permanent component of nonsample households' incomes appears to be less persistent than a random walk.[11]

Finally, in Online Appendix II.2 we generalize the permanent component to an autoregressive process, $p_{it} = \rho p_{it-1} + \xi_{it}$, and estimate the persistence ρ by GMM using a bias-corrected estimator of Chen et al. (2019).[12] We find that the persistence is considerably higher for sample families. We also confirm, in a smaller set of families of siblings of original PSID families, that the persistence of permanent income shocks is noticeably higher for the families of brothers, that is, sample families.

4.2. Different Income Dynamics and the Fit of Minimum Distance Estimation for Sample and Nonsample Families

Our ultimate objective in this section is to provide an improved measurement of income dynamics for sample and nonsample families. To do so, it is instructive to first examine the fit of the standard BPP model estimated above which assumed that the permanent component is a random walk and targeted the moments for income and consumption growth rates.

First, in panel (a) of Figure 2 we consider nonsample families. The bottom row of panels indicates that the fit of the model (long-dashed line) to specifically targeted moments of income growth rates is quite good. (Data moments are plotted using the solid lines.) In contrast, the

[8]Incomes recorded in the PSID in a given year reflect incomes received in the previous year.

[9]See Online Appendix V for additional evidence on the differences in income inequality and income volatility trends for sample and nonsample families.

[10]In a regression of the cross-sectional variances in income growth rates on a constant and trend, the estimated coefficient on trend is not significantly different from zero for nonsample families, but significant at less than 2% level for sample families.

[11]This evidence is also consistent with the possibility that an MA(1) process does not fully capture the dynamics of the transitory component of income for nonsample families. We have found a significantly better fit to the data for the parsimonious AR(1) plus MA(1) process than an alternative of maintaining the random walk assumption for permanent shocks but relaxing the assumption of an MA(1) transitory component instead.

[12]Our conclusions are the same if employ the standard GMM instead of the debiased estimator.

top row of panels indicates that the fit of the model to the moments of income levels is poor. In the data, the variance of log residual incomes rises from about 0.12 to 0.18, while the model predicts a rise to about 0.43. Thus, remarkably, the variance of incomes in levels is overestimated by about 140% in the last sample year.

In Figure 2, panel (b) we consider sample families. The fit of the model (long-dashed line) to the targeted moments of income growth rates is once again quite good. The variance of log income levels is overestimated in 1993 by about 30%, which is substantially lower relative to overestimation for the nonsample families described above.

The descriptive evidence discussed above indicated that income shocks experienced by nonsample families are considerably less persistent than the shocks impacting sample families. Yet, the standard model assumed that permanent shocks are described by the random walk process for both sets of households. Using the identifying moments in Heathcote et al. (2010), in Online Appendix II.3 we show that restricting the permanent component to a random walk when its true persistence is lower may lead to inflated estimates of the variances of permanent (transitory) shocks when targeting the moments in growth rates (levels). Misspecification will lead to negligible biases if the persistence, ρ, is close to one; the biases, however, are expected to be larger for smaller values of ρ. As the results above point to a value of ρ substantially lower than one for nonsample families, the variance of permanent shocks would be significantly overestimated when moments in growth rates are targeted, leading to the dramatic overestimation of the variance of income levels observed above. We also show in the Appendix that one may expect a larger downward bias in the estimated transmission coefficient for permanent shocks using the random-walk assumption for smaller values of the true persistence ρ.

There is another potential source of bias in estimated variances and persistences of income shocks that induces a poor fit to income moments in levels when income growth moments are targeted in estimation. Daly et al. (2021) show that this bias arises if income records in the beginning or end of incomplete income spells are systematically different in their means or variances. This can occur, for example, at the beginning of marriages for the newly-formed couples, or at the end of marriages for the couples which dissolve during 1979–1993. Indeed, in Table A-10 in Online Appendix II.4 we document the presence of these effects in our net family income data. Most prominently, the variance of incomes is high at the start of incomplete income spells relative to income observations from the interior of contiguous income spells.

4.3. Refining the Estimates of Income Dynamics and Consumption Insurance

We now introduce two modifications to the specification of the income process implied by the findings above and re-estimate the income dynamics and consumption insurance of sample and nonsample families. First, we relax the assumption of a random walk in incomes, and model the permanent component as a persistent AR(1) process, $p_{it} = \rho p_{it-1} + \xi_{it}$, estimating, in addition, persistence ρ. Second, we follow Daly et al. (2021) and augment the estimating consumption equation with an additional shock to household incomes to which consumption may react: $\Delta c_{it} = \zeta_{it} + \phi \xi_{it} + \psi \epsilon_{it} + \psi_\nu \nu_{it} + \Delta u_{it}$, where ν_{it} is an i.i.d. shock (with mean and variance estimated from the data), which appears only in the first and last periods of incomplete income spells.[13] The resulting income process is as follows:

$$
\begin{aligned}
y_{it} &= \alpha_i + p_{it} + \tau_{it} + \chi_{it}, \;\; t = t_0, \dots, T \\
p_{it} &= \rho p_{it-1} + \xi_{it} \\
\tau_{it} &= \epsilon_{it} + \theta \epsilon_{it-1} \\
\chi_{it+j} &= \begin{cases} \nu_{it} & \text{if } y_{it-1} \text{ or } y_{it+1} \text{ is missing and } t-1 \geq t_0,\, t+1 \leq T, j = 0 \\ \theta \nu_{it} & j = 1 \\ 0 & \text{otherwise,} \end{cases}
\end{aligned}
\tag{2}
$$

where α_i is individual i's fixed effect, t_0 is the first sample year (1979), and T is the last sample year (1993).

To recover additional parameters, in addition to all of the moments in the original BPP estimation, we also target all the regression coefficients reported in Table A-10. We estimate the model by the method of simulated minimum distance, assuming that persistent, transitory, and ν-shocks are drawn from normal distributions, and using the diagonal weighting matrix calculated by block-bootstrap.[14] In estimations, we assumed that the fixed effect in family incomes is independent of the shocks.

Table 2 contains estimation results.[15] For the families headed by sample males, the persistence of permanent shocks is estimated to be very close to one while for nonsample families, the AR(1) coefficient is estimated to be significantly lower, at only 0.90. The transmission of these shocks to consumption is also estimated to be very different, at 0.99 for sample and 0.48 for nonsample families, respectively.[16] In panels (a) and (b) of Figure 2 we show the fit of the models (lines with triangles) in Table 2 to the data moments. The models with a modified

[13] Daly et al. (2021) showed that it is sufficient to account for the mean and variance of the first and last records of incomplete income spells to eliminate the biases they induce. Hryshko and Manovskii (2019) present evidence that ν_{it} shocks are transitory. This interpretation is also consistent with the underlying thought experiment in BPP. The unit of analysis in BPP and in this paper is a continuous marriage over time and empirically it is only the observations close to the start or end of incomplete family income spells that have systematically different means or variances. The approach in BPP that is based on a sample of continuously married couples yields estimates that are representative of the population at large if marriage and divorce are orthogonal to income shocks. Thus, relative to the income history of a given family, shocks that systematically induce unusual mean and variance of incomes only at the very start or the end of income history are best thought of as being transitory.

[14] We verified that the assumption of normal permanent and transitory shocks (which does not allow for skewness and excess kurtosis) is inconsequential for recovering the variances of income shocks and the transmission coefficients in the baseline BPP model through two experiments: (1) adding the third moments of income and consumption growth to the second moments used for fitting in BPP; and (2) estimating the model with the shocks drawn from a fat-tailed Student t-distribution, the degrees of freedom of which were estimated by matching kurtosis of residual consumption and income growth observed in the data.

[15] As the transmission coefficient for ν-shocks was estimated with a large standard error both for sample and nonsample families, we restricted it to equal the transmission coefficient for transitory shocks. The estimated persistence of permanent shocks is invariant to this assumption.

[16] The finding that the transmission of permanent shocks is higher than the value estimated under the assumption of a random-walk permanent component is consistent with the results in Hryshko and Manovskii (2019) who allow for ν-shocks in estimation and the theoretical prediction in Appendix II.3. The estimated persistence is higher relative to Table A-9 because Table 2 also uses consumption information to identify the parameters of the income process. Moreover, Han and Phillips (2010) show that system-GMM estimates of the persistence may be downward-biased when the true persistence is close to one.

income process match the income growth moments as well as the standard model with the random walk restriction on permanent shocks. The fit to the moments of income levels is, however, improved dramatically.

In panel (c) of Figure 2, we show the fit of the estimated models to the autocorrelation functions of income levels for sample and nonsample families (that were plotted together in Figure 1). Solid lines depict the autocorrelation function in the data, long-dashed lines plot the autocorrelation function implied by the estimates of the BPP model assuming that incomes contain a random-walk permanent component and lines with triangles refer to the autocorrelation function implied by the estimates of the BPP model with a modified income process in Table 2. The estimation with a modified income process shows a much tighter fit to the data moments. For nonsample families, the tighter fit is mainly achieved by a lower estimate of the persistence of longer-lasting shocks, whereas for sample families it is achieved by allowing for additional income variance at the extremes of contiguous income spells. Noteworthy, none of the moments in Figure 2 panel (c) had been targeted in any of the estimations.

5. Quantitative Theory Benchmark

An important objective of measuring income dynamics and consumption insurance in the data is to compare the estimated insurance coefficients with those implied by quantitative models of household consumption and saving choices. To provide such a benchmark, we now describe and calibrate the standard incomplete markets life-cycle model using the estimated income process parameters for sample and nonsample families in Table 2. The simulated model also helps illustrate how sample and nonsample households can be quite similar cross-sectionally despite having very different dynamic properties of income.

5.1. Model

Households start working life at age t_0, retire at age t_R, spend time in retirement until age T, and die at age T with certainty. Households' life spans are uncertain with unconditional probability of being alive at age t equal to s_t. Households supply labor inelastically, value consumption using a CRRA utility function, and discount future with a discount factor β. Household i's problem is:

$$\max_{\{C_{it}\}_{t=t_0}^{T}} E_{i,t_0} \sum_{t=t_0}^{T} \beta^{t-t_0} s_t \frac{C_{it}^{1-\gamma}-1}{1-\gamma},$$

subject to

$$\begin{aligned}
W_{it+1} &= (1+r)(W_{it}+Y_{it}-C_{it}), \\
Y_{it} &= \mu_t P_{it} V_{it}, \quad t=t_0,\ldots,t_R \\
P_{it} &= P_{it-1}^{\rho}\exp(\xi_{it}) \\
V_{it} &= \begin{cases} \exp(\epsilon_{it}), & \text{with prob. } 1-\pi \\ 0, & \text{with prob. } \pi \end{cases} \\
Y_{it} &= \kappa P_{it_R}, \quad t=t_R+1,\ldots,T \\
W_{it} &\geq 0, \ \ t=t_0,\ldots,T.
\end{aligned}$$

Y_{it} is household i's income at age t, stochastic until retirement age t_R, and deterministic afterwards. In the empirical analysis we followed BPP who assumed that the only source of idiosyncratic uncertainty faced by the consumer is net family income excluding capital income. Y_{it} in the model corresponds to this empirical income measure. μ_t is the common lifecycle component of income; P_{it} is the permanent component of income, the log of which follows an AR(1) process with persistence ρ; ξ_{it} is an i.i.d. permanent shock; V_{it} is the transitory component of income which, following Carroll (1997), takes the value of zero with probability π and is positive otherwise, with the values determined by an i.i.d. transitory shock ϵ_{it}.[17] After retirement, household i's income is proportional to the permanent component at age t_R with a replacement rate κ, as in, e.g., Demyanyk et al. (2017). W_{it} is household i's wealth at age t, C_{it} is household i's consumption at age t, and E_{i,t_0} stands for household i's expectation about future resources based on the information available at age t_0. Households cannot borrow but can save into a riskfree asset yielding a net interest rate r.

5.2. Calibration

We calibrate the model to match the data targets for nonsample households (households formed by "daughters" of the original PSID sample members). In one of the quantitative experiments below we will assess the consequences of fixing all parameters calibrated on the sample of daughters, except for changing the parameters governing the income dynamics to those measured on the sample of households formed by the sons of the original PSID sample members.[18]

We assume that households start their life at age 26 with zero assets, retire at age 65, and die at age 90, that is, $t_0 = 26$, $t_R = 65$, and $T = 90$. Before retirement, the unconditional probability of survival, s_t, is set to one; the conditional probabilities of surviving for ages 66 to 90 are taken from Table A.1 in Hubbard et al. (1994). The age-dependent deterministic income profile, μ_t, is taken from Kaplan and Violante (2010). The replacement rate κ is set to 0.70, and the interest rate r is set to 4%, as in Carroll (2009).

We take as given the income process for families of daughters estimated above. Specifically, the variances of permanent and transitory shocks are taken from Column (2) of Table 2; we assume that the shocks ξ_{it} and ϵ_{it} are normally distributed. We then calibrate the CRRA coefficient γ, the probability of a transitory zero-income state, π, and the time discount factor β by matching selected percentiles of the wealth distribution for the families of daughters calculated using the PSID wealth supplements for years 1984, 1989, and 1994 (the years around the period 1979–1993 used for estimation of the income process and consumption insurance, when wealth supplements are available in the PSID). We choose the three parameters by solving the minimization problem

$$\min_{\gamma,\beta,\pi} \sum_{j=\{10,25,50,75,90\}} \left| 100 \cdot \left(\frac{p_j^d(\gamma,\beta,\pi) - p_j^m(\gamma,\beta,\pi)}{p_j^d(\gamma,\beta,\pi)} \right) \right|,$$

[17] Since the estimated moving average parameters in Table 2 are small, for simplicity, we assume that transitory shocks are i.i.d.

[18] Following Hryshko et al. (2011), for the households formed by sons and daughters of original PSID sample members, we found no differences in risk attitudes as revealed by their choices of hypothetical risky gambles in the 1996 wave of the PSID.

where p_j^d and p_j^m are the data and model j's percentiles of the wealth distribution, $j = \{10, 25, 50, 75, 90\}$. In our calibration (and then simulations for the families of sons and daughters), we replicate the age distributions observed in the data. As in BPP PSID sample selection, we drop income growth outliers in the simulated data. The values of internally calibrated parameters are shown in the three bottom rows of Table 3.[19]

5.3. Quantitative Findings

First, observe that the income distribution in the model, driven by the estimated income process, matches well the distribution of income for PSID daughters in the data (a comparison of columns (3) and (4) in the top panel of Table 3). Moreover, the model also matches well the targeted moments of the wealth distribution of families of daughters.

We next fix the model parameters but change the income process to that of PSID sons. Specifically, we now use the values for the persistence of permanent shocks, and variances of permanent and transitory shocks, from column (1) of Table 2. We find that the model matches quite well the income distribution for the sample of sons and does a decent job in matching their wealth distribution.

As we have already seen in Section 3.3, income and wealth are quite similar on average in the data among households formed by the PSID sons and daughters, despite very different estimated income dynamics. Similar patterns are also replicated in the model.

Next, we measure the degree of consumption insurance by applying the same BPP measurement approach to the model generated data as we did when measuring insurance in PSID data in Section 3. To do so, we simulate multiple samples with 814 families of daughters and 854 families of sons. We also allow for measurement error in consumption, assuming that it is distributed normally with the variance of 0.07, as was estimated in Table 2.

The results are in Table 4. Columns (1) and (3) reproduce our empirical results from Table 1, columns (1) and (2) of Panel B, respectively, and column (5) tabulates the transmission coefficients for the combined sample of sons and daughters estimated in the data. The model replicates remarkably well the transmission coefficients for permanent and transitory shocks observed in the data both for the families of sons and daughters (the model values reported in columns (2) and (4) respectively) and for the combined sample (the model value reported in column (6)). As in the PSID data, despite having cross-sectionally similar distributions of income and wealth, the model households of sons and daughters differ substantially in the consumption insurance. It is also noteworthy that uncertainty in the point estimates in the PSID data is matched reasonably well in the model.

Finally, we replicate the so-called excess insurance puzzle. To this end, we ignore the observed differences in the income processes across the families of sons and daughters and, instead, make the standard assumption that they share the same income process, which is the sum of a random walk permanent component and a transitory shock. We estimate this process in the data and then recalibrate the same three parameters, γ, π, and β, by following the procedure described above. Using the same minimum-distance method on the simulated data, we end up with substantial underestimation of insurance at about 0.10 in column (7) for the model relative to 0.43 in column (5) for the data.

[19]The calibrated CRRA coefficient is somewhat low, at 0.4, but is consistent with the results in Gourinchas and Parker (2002). The calibrated value of the time discount factor is standard, and the value of the probability of a zero-income state is consistent with Carroll (1997).

The main substantive takeaway from these findings is on the relationship between the degree of insurance and the dynamic properties of income shocks. Specifically, after accounting for heterogeneity in income processes between sample and nonsample households measured in the PSID data, we do not find evidence of excess insurance in the data relative to the standard self-insurance model.

6. How Could Sample and Nonsample Families Have Different Income Dynamics?

If the PSID originally interviewed a representative sample of the U.S. population of households, the survey design is expected to ensure that their sons and daughters also form representative samples of the same underlying population so that one would not expect to observe a large difference in income dynamics and consumption insurance between families formed by PSID sons and daughters that we find.

A large literature has compared the cross-sectional PSID samples to other surveys which are expected to be cross-sectionally representative, such as the CPS, and found relatively small discrepancies between them. However, this finding does not necessarily imply that the original PSID sample was representative of the U.S. population with respect to, e.g., income dynamics. This has never been studied, and for a good reason. The PSID is a unique dataset that tracks people over a long time. There are no other comparable surveys that could be employed for cross-validation of the dynamic properties of PSID data.

In an attempt to provide some first evidence on this issue, we take the following approach. We exploit the fact that for respondents to the 1998 HRS, we can obtain administrative individual earnings data going back to 1978. HRS 1998 is a representative sample of the U.S. population over the age of 50.[20] Thus, we select individuals in the HRS born before 1948 and obtain historical administrative earnings data for these individuals extracted from the IRS Master Earnings File. In the PSID, we select members of the original households born before 1948 who survive in the sample up to 1999 so that they would correspond to the HRS sample in 1998. Assuming that both the PSID and HRS are representative, we now have two independent sets of earnings histories since 1978, allowing us to compare the dynamic properties of earnings. We estimate earnings processes by GMM. The results reported in Table 5 indicate that the original cohort of males and families in the PSID have noticeably more persistent earnings than the set of corresponding individuals or families in the HRS.

One potential explanation for a higher persistence of survey-based earnings in the PSID relative to administrative earnings data from the HRS could be that survey data are generally more persistent due to systematic mismeasurement (if, e.g., measurement error is more persistent than true earnings). However, Abowd and Stinson (2013) using Survey of Income and Program Participation data matched with administrative earnings records find that survey earnings tend to be less persistent than the administrative ones as measured by their autocorrelation functions. Our comparison of overlapping survey-based and administrative records for the same individuals in the HRS reveals a similar pattern. Lower autocorrelation of survey earnings is consistent with mean-reverting measurement error in survey earnings data, as documented in, e.g., Bound and Krueger (1991), and in our Table A-6.

Another possibility is that the original PSID sample was not representative, and biased

[20] See `https://hrs.isr.umich.edu/documentation/survey-design`.

toward individuals with more persistent earnings dynamics.[21] If true, this could potentially rationalize the patterns we find. For example, if sons of the original PSID heads partially inherit their high earnings persistence, but daughters tend to marry a more representative cross-section of males who have lower persistence, it might induce the differences in earnings persistence among sample and nonsample households. Our exploration of the data did not yield convincing evidence in support of this logic. One piece of evidence against the notion that the original PSID sample was not representative is somewhat indirect but powerful. The PSID was based on a pioneering and ingenious design copied later on by the corresponding surveys in other countries, for example by the German Socio-Economic Panel (GSOEP), the British Household Panel Survey (BHPS), the Household, Income and Labour Dynamics in Australia (HILDA) survey, the Korean Labor and Income Panel Study (KLIPS), and the Swiss Household Panel (SHP).[22] These datasets also started from a random cross-section of households and followed them and their descendants just as the PSID does. As their original sampling was also random, there does not appear to be a reason to oversample households with high income persistence in all of them. Yet, as Table 6 illustrates, nonsample households in all those datasets have lower income persistence than sample households, replicating qualitatively the pattern that we have documented for the PSID. Those surveys also have an attractive feature for our purposes, that they interview not one family member about incomes of all other family members, as the PSID typically does, but each adult family member individually.[23] Thus, the discrepancy cannot be caused by the differential tendency of males and females to respond to the income questions between sample and nonsample households.

This leads us to a potentially more plausible explanation based on the selective sample attrition. It has been documented that sample attrition in the PSID is not random. First, Fitzgerald et al. (1998b) and Fitzgerald (2011) note that PSID males are more likely to attrit than PSID females. We confirm this pattern in Table 8 Panel A, where we follow individuals age 0–50 in the 1968 survey from the year they are first observed as heads or wives and check whether they attrit in the subsequent fifteen years. We confirm that males are indeed more likely to attrit and that the same pattern is observed in the other datasets. Note that when a PSID sample male or a PSID sample female attrits, so does his or her entire household. Thus, selective attrition based on gender implies that there is more attrition among sample households than among nonsample ones.

The second feature of attrition that we document is more novel. Inspired by Fitzgerald et al. (1998a), who find that attrition is related to the transitory volatility of past individual earnings, we study the relationship between attrition and persistence properties of family incomes. Specifically, we first estimate the persistence of household income in the PSID for the households whose sample spouse (the one with the "PSID gene") is present in the data for at least twelve years, but separately for the households whose sample spouse does and does

[21] Of course, it is also possible that the HRS is not representative.

[22] A description of these datasets and the details of sample construction are provided in Appendix IV.

[23] The BHPS, SHP, and KLIPS allow for proxy interviews. In the BHPS and SHP, proxy interviews are very rare in our estimation samples, at less than two percent for male and female spouses. In the KLIPS, the extent of proxy interviews does not vary substantively across sample and nonsample families – about ten percent of male and four percent of female interviews are done by some other family member in sample families, while the corresponding numbers in nonsample families are thirteen and one percent.

not exit the data in the subsequent eight years.[24] The results, summarized in Table 7, imply that eventually attriting PSID households have a notably less persistent income process. The same pattern holds among sample and nosample families.

The non-U.S. datasets are of a considerably shorter duration than the PSID. To provide a comparable analysis across all datasets, we consider the set of couples present in the first survey year of each respective dataset. We follow these couples for twelve years and then check whether they attrit in the subsequent several years (Appendix IV details sample selections). In Table 8 Panel B, we report separate estimates of persistence for households that will and will not eventually attrit in each of the datasets. A clear pattern is evident: the families that will eventually attrit have lower income persistence than families that will remain in the sample.

The two sources of selection in attrition – based on gender and income dynamics – imply that the set of sample households is more severely selected in favor of households with higher income persistence than the set of nonsample households. This is consistent with the patterns we document where sample households have a more persistent income process. It is also consistent with our finding of a higher persistence of male and family earnings in the PSID than in the HRS linked to administrative data because the PSID sample is restricted to households who have not attrited between 1968 and 1999. In contrast, the 1998 HRS can be thought of as a less selected cross-sectional sample for which complete retrospective earnings histories are obtained from administrative data.

7. How Much Consumption Insurance in the U.S.?

We now return to the key question motivating this paper: How much consumption insurance in the U.S.? Unfortunately, our conclusion is that this question cannot be answered precisely with the currently available data. Yet, our analysis reveals some evidence that the answer is quite different from what has become the conventional wisdom.

To obtain unbiased estimates of income dynamics and ultimately of consumption insurance it appears necessary to correct for the effects of selective attrition in the PSID.[25] One approach would involve measuring the dynamic properties of incomes of individual attritors and non-attritors, identifying for each attritor a matching non-attritor and then increasing the weight of matched non-attritors in the estimation sample. Practical suitability of this approach is limited by two issues. First, attrition rates are quite large and for most attritors income histories are either short or not observed at all. For example, nonresponse in the initial, 1968 wave of the PSID constituted about 24%. Even for families that were interviewed in 1968, about 30% (25%) of their sons (daughters) aged 0–18 in 1968 never join the PSID as heads or wives by 2015. As no information on incomes of these households is available, correcting for selection

[24] Specifically, we select all families observed during 1968–1997 with the head born in 1920–1959, the same cohorts as in the main analysis. We then drop the families whose head is more than forty-five years of age in the year the family is first observed in the PSID (to define attrition by the time the head is aged sixty-five during a twenty-year window), and select observations with the head's age range 25–65.

[25] The selection problem we face is different from the often encountered one which would correspond in our setting to the situation where income persistence is different between sample and nonsample families but constant within each group of families. Such selection on observables can be corrected by placing a higher weight on the persistence and insurance estimates for sample families as relatively more of them drop out of the data. This approach is not suitable, however, if attrition is correlated with individual income dynamics, as is the case in our setting, where attritors are more likely to have lower persistence.

among these households appears infeasible (there also do not appear to be good predictors for future income dynamics based on scant family background information available for these individuals). Second, for families that do enter the PSID but attrit eventually, family income histories are typically observed for a short period, making it difficult to estimate their dynamic properties with sufficient precision. For example, we have at most eleven income observations (and typically many fewer) for families that attrit prior to the start of our estimation sample in 1979. To find matching non-attriting families, we need to identify families in the estimation sample that already existed at the same time as the attriting families and shared the same age and other determinants of individual income dynamics. The sets of such families in the PSID are very small. This makes it impossible to verify that the supports of the distributions of, e.g., income persistence coincide for attriting and non-attriting families and, assuming that they do, find reliable matches.

A more promising approach for correcting for selective attrition in estimating income dynamics would be to link the PSID to administrative income records. If it were possible to obtain long administrative income histories for individuals and families who did and did not attrit, we could potentially match attritors to relevant non-attritors based on administrative records and then reweight the non-attritors appropriately when estimating the extent of consumption insurance using the PSID data. Unfortunately, linking of the PSID to administrative income data is not currently possible.

What does this leave the researchers with? The evidence suggests that both sets of sample and nonsample households are affected by attrition and do not provide an unbiased measure of income dynamics and consumption insurance. But given the stark differences between them, the question is then which set of households provides a better guide to the income dynamics and consumption insurance available to U.S. families. We found above that the set of nonsample households is less affected by attrition and is thus more representative. This suggests that measures of income and consumption dynamics on this sample are more informative about the corresponding measures for a representative U.S. household.[26]

To further assess the validity of this interpretation, in Online Appendix V we compare trends in net family income inequality for sample and nonsample families in the PSID to families in the CPS documented in Heathcote et al. (2010) and trends in the volatility of individual earnings growth documented using U.S. Social Security Administration data by Bloom et al. (2017). Both sets of comparisons reveal that nonsample PSID households feature similar trends to the nationally representative CPS and administrative samples, while patterns documented for sample PSID households deviate considerably.

[26] While our focus in this paper is on the level of consumption insurance, our findings also have an interesting implication about its trend. Specifically, using their full PSID sample, BPP show that the slowdown in consumption inequality in the mid-1980s was due to the reduced importance of the variance of permanent shocks in the overall variance of income growth during that period while the degree of insurance against permanent and transitory shocks did not change significantly between the early vs. later parts of their sample. We find that that this finding is driven by nonsample households, whose insurance indeed remains constant. In contrast, the degree of consumption insurance against permanent income shocks among sample households features a statistically significant increase.

8. Conclusion

The dynamic properties of income play a central role in modern macro and labor economics. They are key for understanding the variation in consumption, the permanent and transitory nature of income and consumption inequality, and for the optimal design of tax and transfer policies. Most of what the profession knows about joint income and consumption dynamics at the household level in the U.S. is based on the data from the PSID. Standard measures of consumption insurance obtained using these data imply excess insurance of permanent income shocks relative to the prediction of the workhorse incomplete-markets model.

In this paper, we document that the PSID consists of two sets of households that systematically differ in the dynamics of their income and consumption. Specifically, the PSID comprises the original sample members interviewed in 1968 and their offspring, and nonsample members, who marry PSID sample males or females. We find a nearly complete pass-through of permanent income shocks to consumption for households headed by PSID sample males. In contrast, families headed by nonsample males show a dramatically higher degree of insurance against permanent income shocks. Moreover, income shocks of households headed by nonsample males are considerably less persistent. Conditional on income dynamics, the estimated degree of insurance in each subsample is consistent with the prediction of the standard incomplete-markets model. In particular, we find no evidence of excess consumption insurance beyond that provided by self-insurance due to accumulated household wealth.

While the patterns documented in the paper are highly robust, the existence of large differences in the stochastic properties of income and consumption between households formed by sons and daughters of the original PSID families is unexpected. It raises both the issues of the interpretation as both sets of households are expected to be equally representative of the same U.S. population, and the concern that the differences might be due to systematic mismeasurement. Given the absolutely central role of the PSID in research in economics and other social sciences, understanding these issues is of first order importance. We contribute to building this understanding.

We compared the dynamics of earnings of the original PSID cohorts to administrative earnings records of comparable cohorts in the HRS and found that earnings are more persistent in the PSID. While it is well known from multiple validation studies that the PSID is close to being cross-sectionally nationally representative, it is not known whether it is representative with respect to, e.g., earnings dynamics. A comparison with the administrative data suggests that it might not be. This leads to the thorny question of whether the bias was induced by the original sampling or by the evolution of the sample over time due, in part, to selective attrition. If it were possible to match PSID individuals to administrative data to obtain complete earnings histories of PSID individuals not affected by selective survey attrition, we could compare these complete administrative earnings histories to those of the corresponding U.S. population to infer the representativeness of the original PSID sample. Unfortunately, this is not currently possible.

Instead, we compare the PSID to other datasets, from Germany, U.K., Australia, Korea, and Switzerland, which were modeled on the PSID. We find that those datasets feature differences in the income dynamics between sample and nonsample households that are qualitatively similar to those in the PSID. Their original random samples were also cross-sectionally representative suggesting that they are unlikely to feature a similar bias in sampling households with high income persistence. This indicates that selective attrition might play an important role. Indeed, we find that both in the PSID and the other datasets, males are more likely to attrit than females and that attritors tend to have lower income persistence. Thus, both sample and nonsample families represent selected subsets, but the selection effect is stronger on the set of sample families, explaining a higher persistence of their incomes relative to nonsample families and to administrative records that are not subject to such attrition. This implies that the set of nonsample families is less selected and we indeed find that the cross-sectional properties of family incomes computed on this set of PSID families line up much better with the corresponding statistics computed using the nationally representative sample from the CPS or U.S. Social Security Administration data. This suggests that the set of nonsample PSID families provides a better guide to the income dynamics and the degree of consumption insurance in the U.S.

References

Abowd, J., Card, D., 1989. On the Covariance Structure of Earnings and Hours Changes. Econometrica 57, 411–445.

Abowd, J.M., Stinson, M.H., 2013. Estimating Measurement Error in Annual Job Earnings: A Comparison of Survey and Administrative Data. The Review of Economics and Statistics 95, 1451–1467.

Arellano, M., Blundell, R., Bonhomme, S., 2017. Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework. Econometrica 85, 693–734.

Becketti, S., Gould, W., Lillard, L., Welch, F., 1988. The Panel Study of Income Dynamics after Fourteen Years: An Evaluation. Journal of Labor Economics 6, 472–492.

Bloom, N., Guvenen, F., Pistaferri, L., Sabelhaus, J., Selgado, S., Song, J., 2017. The Great Micro Moderation. Working Paper.

Blundell, R., 2014. Income Dynamics and Life-Cycle Inequality: Mechanisms and Controversies. Economic Journal 124, 289–318.

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

Blundell, R., Pistaferri, L., Saporta-Eksten, I., 2016. Consumption Inequality and Family Labor Supply. American Economic Review 106, 387–435.

Bound, J., Krueger, A.B., 1991. The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right? Journal of Labor Economics 9, 1–24.

Carroll, C.D., 1992. The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence. Brookings Papers on Economic Activity 2, 61–156.

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

Carroll, C.D., 2009. Precautionary Saving and the Marginal Propensity to Consume Out of Permanent Income. Journal of Monetary Economics 56, 780–790.

Chen, S., Chernozhukov, V., Fernández-Val, I., 2019. Mastering Panel Metrics: Causal Impact of Democracy on Growth. AEA Papers and Proceedings 109, 77–82.

Daly, M., Hryshko, D., Manovskii, I., 2021. Improving the Measurement of Earnings Dynamics. International Economic Review (forthcoming).

Demyanyk, Y., Hryshko, D., Luengo-Prado, M.J., Sørensen, B.E., 2017. Moving to a Job: The Role of Home Equity, Debt, and Access to Credit. American Economic Journal: Macroeconomics 9, 149–181.

Fitzgerald, J., Gottschalk, P., Moffitt, R., 1998a. An Analysis of Sample Attrition in Panel Data: The Michigan Panel Study of Income Dynamics. Journal of Human Resources 33, 251–299.

Fitzgerald, J., Gottschalk, P., Moffitt, R., 1998b. An Analysis of the Impact of Sample Attrition on the Second Generation of Respondents in the Michigan Panel Study of Income Dynamics. Journal of Human Resources 33, 300–344.

Fitzgerald, J.M., 2011. Attrition in Models of Intergenerational Links Using the PSID with Extensions to Health and to Sibling Models. The B.E. Journal of Economic Analysis & Policy 11, 1–63.

Gourinchas, P.O., Parker, J., 2002. Consumption Over the Life Cycle. Econometrica 70, 47–89.

Guvenen, F., 2007. Learning Your Earning: Are Labor Income Shocks Really Very Persistent? American Economic Review 97, 687–712.

Guvenen, F., 2009. An Empirical Investigation of Labor Income Processes. Review of Economic Dynamics 12, 58–79.

Han, C., Phillips, P.C.B., 2010. GMM Estimation for Dynamic Panels with Fixed Effects and Strong Instruments at Unity. Econometric Theory 26, 119–151.

Heathcote, J., Perri, F., Violante, G.L., 2010. Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 1967–2006. Review of Economic Dynamics 13, 15–51.

Hryshko, D., Luengo-Prado, M.J., Sørensen, B.E., 2011. Childhood Determinants of Risk Aversion: The Long Shadow of Compulsory Education. Quantitative Economics 2, 37–72.

Hryshko, D., Manovskii, I., 2019. Income Dynamics and Consumption Insurance. Working Paper. University of Pennsylvania.

Hubbard, R.G., Skinner, J., Zeldes, S.P., 1994. The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving. Carnegie-Rochester Conference Series on Public Policy 40, 59–125.

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

Lillard, L.A., 1989. Sample Dynamics: Some Behavioral Issues, in: Kasprzyk, D., Duncan, G., Kalton, G., Singh, M.P. (Eds.), Panel Surveys. Wiley Series in Probability and Statistics, pp. 497–511.

Table 1: CONSUMPTION INSURANCE FOR VARIOUS SAMPLES

Panel A: BPP sample

	Combined (1)	Sample (2)	Nonsample (3)
ϕ, transmission of perm. shock	0.6436 (0.0858)	0.9430 (0.1508)	0.4303 (0.0950)
ψ, transmission of trans. shock	0.0291 (0.0436)	–0.0108 (0.0469)	0.1014 (0.1009)

Panel B: Updated sample

	Sample sons (1)	Nonsample (2)	Sample orig. (3)	Sample all (4)	Sibling pairs Sons (5)	Sibling pairs Daught. (6)
ϕ, transmission perm. shock	0.8656 (0.1939)	0.4563 (0.1007)	1.0869 (0.2008)	0.9038 (0.1462)	1.0736 (0.3392)	0.3244 (0.1604)
ψ, transmission trans. shock	0.0650 (0.0800)	0.1204 (0.0886)	0.0356 (0.0412)	0.0494 (0.0391)	0.1512 (0.1308)	–0.1173 (0.1818)

Notes: The table shows the transmission of permanent and transitory shocks to household net incomes to household consumption estimated by minimum distance. Standard errors in parentheses. "Sample" comprise families headed by PSID males, "Nonsample" comprise families formed by PSID females, "Sample orig." comprise families married by 1968 and who stay married until they were last seen in 1979–1993, "Sample sons" comprise families married in 1969 or later and headed by PSID sample males. Panel A uses the original BPP sample, Panel B uses an updated sample. See Sections 2 and 3.3 for details. In panel A, p-value for test of equal ϕ (ψ) between sample and nonsample families equals 0.4% (31%). In panel B, p-value for test of equal ϕ (ψ) in columns (1) and (2) equals 6% (64%); in columns (2) and (3) equals 1% (39%); in columns (2) and (4) equals 1% (46%); in columns (5) and (6) equals 5% (23%).

Table 2: Consumption insurance. AR(1) permanent component

	Sample (1)	Nonsample (2)
ρ, AR coeff.	0.9956	0.9003
	(0.0095)	(0.0314)
σ_{ξ}^2, var.	0.0151	0.0285
perm. shock (avg.)	(0.0029)	(0.0052)
θ, MA coeff.	0.1315	0.0456
	(0.0252)	(0.0653)
σ_{ϵ}^2, var.	0.0407	0.0274
trans. shock (avg.)	(0.0027)	(0.0042)
σ_{u}^2, var.	0.0704	0.0740
cons. meas. err (avg.)	(0.0033)	(0.0082)
ϕ, transmission	0.9903	0.4798
perm. shock	(0.1770)	(0.1142)
ψ, transmission	0.0922	0.0997
trans. shock	(0.0393)	(0.1112)

Notes: The table shows the estimated income process parameters and the transmission of permanent and transitory shocks to household net incomes to household consumption. "Sample" comprise families headed by PSID males, "Nonsample" comprise families formed by PSID females. Simulated minimum distance estimates. Standard errors in parentheses.

Table 3: Model calibration

	Sons		Daughters	
	Data (1)	Model (2)	Data (3)	Model (4)
Various income percentiles, in '000s				
P10	16.9	20.4	16.7	22.0
P25	25.3	28.6	25.4	27.9
P50	35.9	35.5	36.3	37.1
P75	49.3	44.4	49.6	48.8
P90	66.6	64.5	66.5	63.3
Various wealth percentiles, in '000s				
P10	4.7	11.9	4.3*	3.9
P25	19.7	26.4	18*	14.9
P50	54.2	57.7	48*	47.9
P75	119.7	118.2	125.4*	129.8
P90	218.4	220	254.2*	265.4
Internally calibrated parameter values				
Time disc. factor, β			0.969	
Coeff. RRA, γ			0.405	
Prob. of zero inc. state, π			0.006	

Notes: Top two panels of the table show various income and wealth percentiles in the data and in the model for the households formed by sons and daughters of the original PSID families. The bottom panel shows the calibrated parameters. See Section 5 for details. * indicates calibration targets.

Table 4: Consumption insurance in simulated and PSID data

	Sons		Daughters		Combined Sons & Daughters		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Data	Model	Data	Model	Data	Model	RW Model
ϕ, transmission	0.865	0.849	0.456	0.446	0.570	0.618	0.929
perm. shock	(0.193)	(0.139)	(0.101)	(0.066)	(0.090)	(0.053)	(0.114)
ψ, transmission	0.065	0.058	0.120	0.139	0.085	0.082	0.066
trans. shock	(0.080)	(0.047)	(0.089)	(0.063)	(0.062)	(0.040)	(0.063)

Notes: Minimum distance estimates of the coefficients for permanent and transitory income shocks for the model and PSID data. "Sons" and "Daughters" are the households formed by sons and daughters of the original PSID families, respectively. Standard errors (calculated by bootstrap for the model) in parentheses. In columns (2), (4), and (6) income processes are different for the families of sons and daughters, whereas in column (7) sons and daughters share the same income process with the permanent component being a random walk. See Section 5 for details.

Table 5: GMM estimates of persistence: PSID vs. HRS-SSA

	Family earnings		Male earnings	
	PSID	HRS-SSA	PSID	HRS-SSA
	(1)	(2)	(3)	(4)
ρ, persistence	0.98	0.93	1.02	0.96
perm. shock	(0.04)	(0.02)	(0.02)	(0.02)
No. ind./fam.	508	1822	520	2628

Notes: The results from a two-step debiased GMM estimation. Bootstrap standard errors in parentheses. Online Appendix II.2 describes the methodology. Online Appendix III describes sample selections. "HRS-SSA" sample uses data on male and family earnings from administrative tax records linked to the households in the Health and Retirement Study.

Table 6: GMM estimates of persistence. Various datasets

Dataset: (Country):	PSID (U.S.A.)		GSOEP (Germany)		BHPS (U.K.)		HILDA (Australia)		KLIPS (Korea)		SHP (Switz.)	
	S (1)	NS (2)	S (3)	NS (4)	S (5)	NS (6)	S (7)	NS (8)	S (9)	NS (10)	S (11)	NS (12)
ρ, persist.	0.95	0.83	0.93	0.89	0.85	0.80	0.96	0.88	0.92	0.79	0.91	0.82
perm. shock	(0.02)	(0.04)	(0.01)	(0.03)	(0.02)	(0.04)	(0.03)	(0.06)	(0.02)	(0.05)	(0.04)	(0.06)
No. families	1593	889	2044	423	2467	554	3286	949	2181	516	1625	258

Notes: Labels "NS" and "S" stand for nonsample and sample families, respectively. The results from a two-step debiased GMM estimation. Bootstrap standard errors in parentheses. See Online Appendix II.2 for the methodology. See Online Appendix IV for the data description and sample selections.

Table 7: GMM estimates of persistence by attrition. PSID

	Samp.		Nonsamp.	
	Non-attr. (1)	Attr. (2)	Non-attr. (3)	Attr. (4)
ρ, persistence	0.96	0.78	0.89	0.72
perm. shock	(0.02)	(0.05)	(0.04)	(0.07)
No. families	1156	174	585	74

Notes: The results from a two-step debiased GMM estimation. Bootstrap standard errors in parentheses. See Online Appendix II.2 for the methodology.

Table 8: Attrition rates and GMM estimates of persistence by attrition. Various datasets

Dataset: (Country):	PSID (U.S.A.)		GSOEP (Germany)		BHPS (U.K.)		HILDA (Australia)		KLIPS (Korea)		SHP (Switz.)	
	Panel A: Attrition rates											
Men	51.1		63.2		56.3		58.4		67.2		77.8	
Women	44.5		58.5		48.1		56.2		61.9		75.7	
	Panel B: GMM estimates of persistence											
	NA (1)	A (2)	NA (3)	A (4)	NA (5)	A (6)	NA (7)	A (8)	NA (9)	A (10)	NA (11)	A (12)
ρ, persist. perm. shock	0.91 (0.04)	0.81 (0.09)	0.99 (0.04)	0.75 (0.08)	0.88 (0.03)	0.83 (0.09)	0.87 (0.04)	0.66 (0.11)	0.87 (0.06)	0.56 (0.11)	0.96 (0.07)	0.66 (0.08)
No. fam.	573	56	627	101	724	133	836	67	654	86	449	53

Notes: Labels "NA" and "A" stand for families of non-attritors and attritors, respectively. The results from a two-step debiased GMM estimation in Panel B. Bootstrap standard errors in parentheses in Panel B. See Online Appendix II.2 for the methodology. See Appendix IV for the data description and sample selections.

Figure 1: Data moments

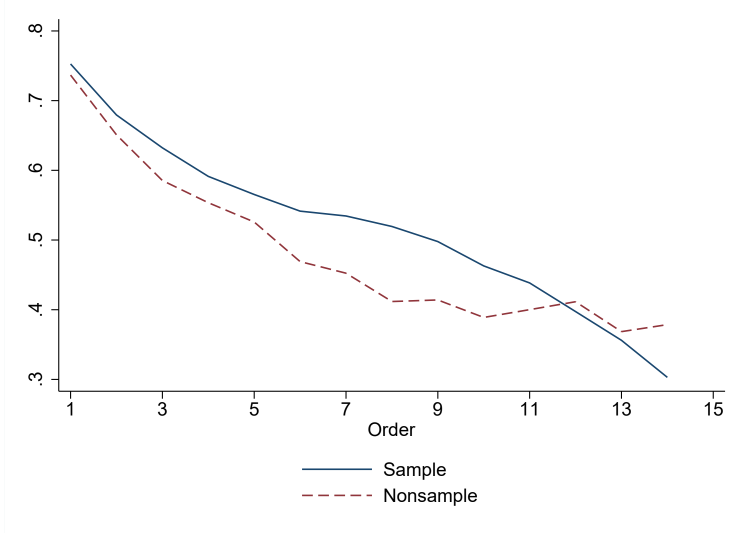

(a) Autocorrelation function of net family incomes

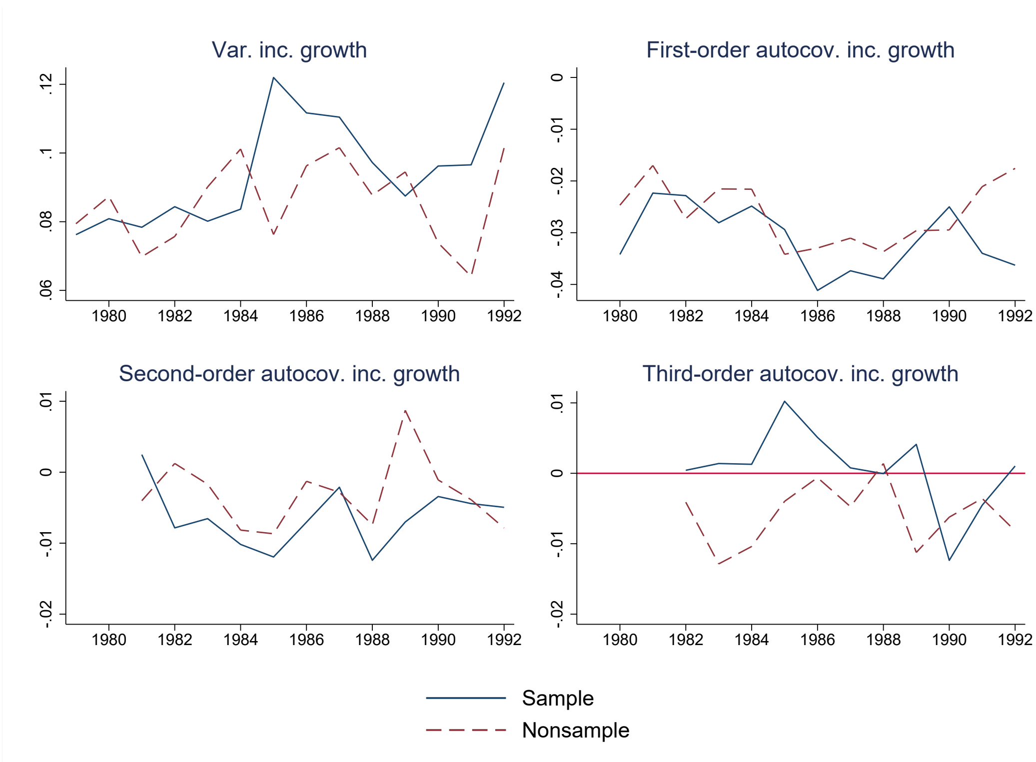

(b) Autocovariances of income growth rates

Notes: The figure shows the autocorrelation function for net family incomes and the autocovariance function, up to the third order, for the growth rate in net family incomes separately for sample and nonsample families. "Sample" comprise families headed by PSID males, "Nonsample" comprise families formed by PSID females.

Figure 2: Model fit to various data moments

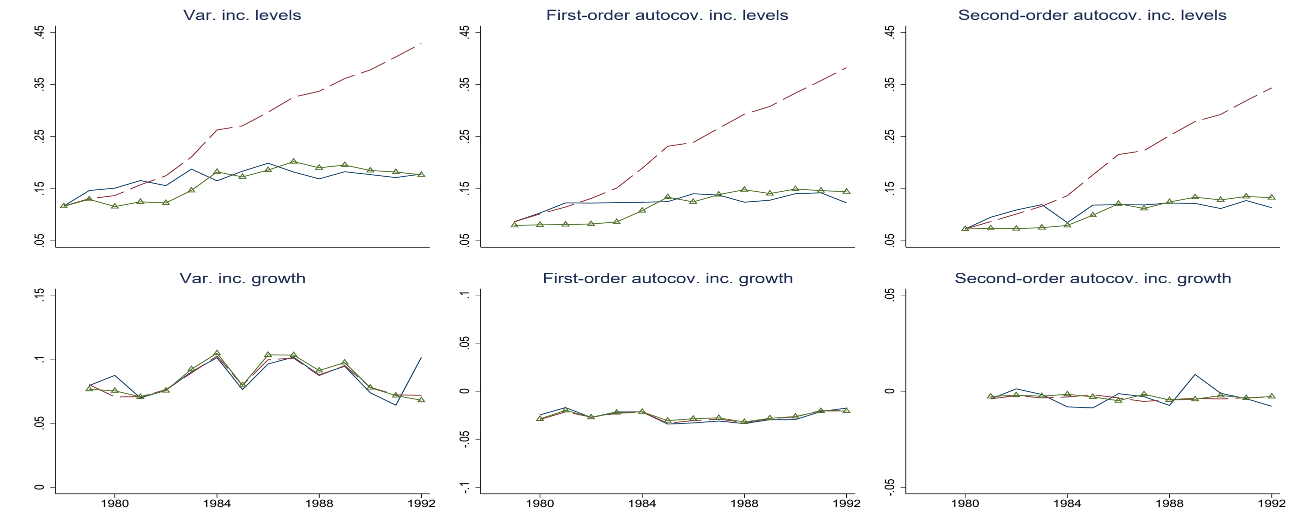

(a) Nonsample Families

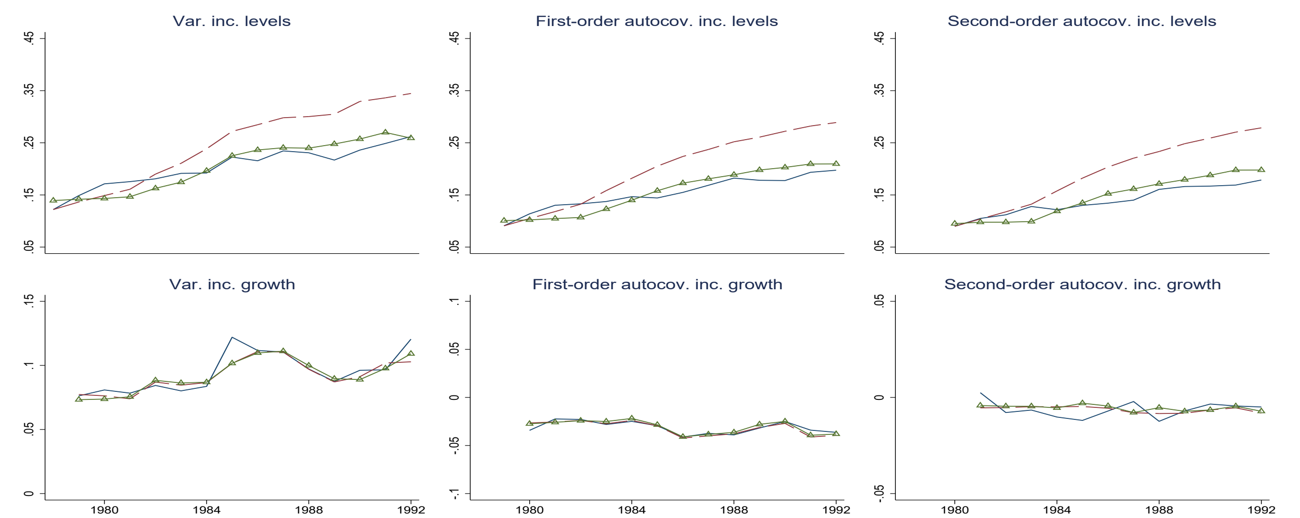

(b) Sample Families

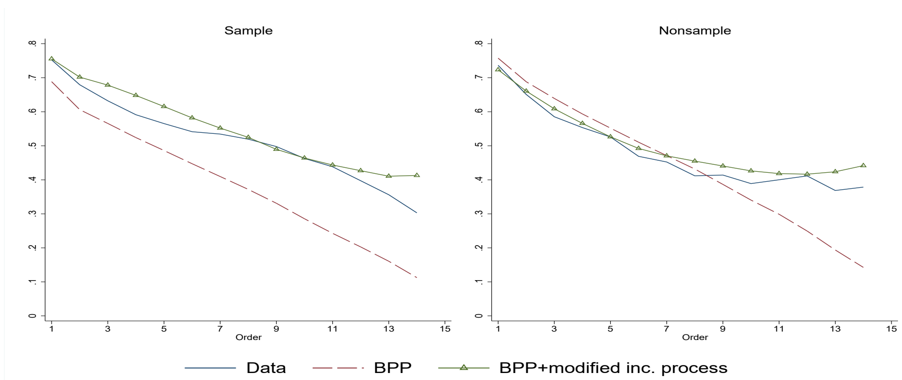

(c) Autocorrelation function of income levels

Notes: The figure shows the data moments (solid lines) and the fit to those moments of 1) the BPP model which assumes that the permanent component is a random walk (dash lines) and 2) the BPP model which assumes that the permanent component is an autoregressive process as in Eq. 2 (solid lines with triangles). "Sample" comprise families headed by PSID males, "Nonsample" comprise families formed by PSID females.