

# How Much Consumption Insurance in the U.S.?<sup>\*</sup>

Dmytro Hryshko<sup>†</sup>  
University of Alberta

Iourii Manovskii<sup>‡</sup>  
University of Pennsylvania

## Abstract

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 Panel Study of Income Dynamics (PSID). We find that there are two sets of households in the PSID that differ dramatically in the dynamics of their income and consumption. Households headed by males from the original PSID sample or by their sons have a highly persistent income process, and permanent shocks to their incomes almost fully pass through to consumption. Households headed by males who marry daughters of the original PSID sample members have a much less persistent income process and a substantially higher degree of consumption insurance. In contrast to the excess insurance puzzle documented on the combined sample, the degree of consumption insurance against permanent income shocks in each subsample is consistent with the prediction of the standard incomplete-markets model.

The sharp difference in income and consumption dynamics between households formed by sons and daughters of the original PSID families is unsettling because the PSID design implies that both sets of households are expected to be representative of the same U.S. population. We contrast PSID data with the U.S. administrative earnings data and with the patterns in other datasets from around the world that were modeled on the PSID. The findings point to an important role of differential attrition based on the dynamic properties of incomes in inducing the differences. The set of households formed by daughters of the original PSID families is less affected by selective attrition and is a better guide to the income dynamics and consumption insurance available to a typical U.S. household.

**KEYWORDS:** Consumption insurance, income processes, incomplete markets models, attrition, Panel Study of Income Dynamics.

**JEL Classification:** D12, D15, D31, E21

---

<sup>\*</sup>We are very 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 participants at the PSID Annual User Conference, NBER Summer Institute Working Group on Consumption, AEA and 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, 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 the National Science Foundation Grants SES-1357903 and SES-1824520, and SSHRC IG grant 435-2018-1275 is gratefully acknowledged.

<sup>†</sup>Department of Economics, University of Alberta, 8-14 HM Tory Building, T6G 2H4, Edmonton, AB, Canada. E-mail: dhryshko@ualberta.ca.

<sup>‡</sup>Department of Economics, University of Pennsylvania, Suite 150, PCPSE Building, 133 South 36th Street, Philadelphia, PA 19104, USA. E-mail: manovski@econ.upenn.edu.

# 1 Introduction

The dynamic properties of earnings and income play a central role in many areas of macro and labor economics. The durability of shocks to income is key for assessing the permanent or transitory nature of inequality. It also affects the degree of adjustment that households must undertake to smooth fluctuations in consumption. The close link between income dynamics and consumption decisions at the household level implies a tight connection between the level and evolution of income and consumption inequality. It is also a key input into the optimal design of tax and transfer policies.<sup>1</sup>

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 Panel Study of Income Dynamics (PSID). For example, Blundell, Pistaferri, and Preston (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 provides a valuable empirical benchmark for assessing the performance of workhorse 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. 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.”

---

<sup>1</sup>See Arellano, Blundell, and Bonhomme (2017) and Daly, Hryshko, and Manovskii (2016) for references to the literature.

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 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 Beckett, Gould, Lillard, and Welch (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-male-headed households is well described by a random-walk model, the households headed by nonsample males 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, 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 (the standard procedure 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 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 comes with a fairly sizable standard error so that it 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.<sup>2</sup> 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 genuine? 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 Social Security Administration (SSA) merged with the survey records by the Health and Retirement Study (HRS). The merged data help us to rule out the dynamics of measurement error in reported earnings by different types of households as an important driver of our findings. More

---

<sup>2</sup>In a related paper, Alan, Browning, and Ejrnæs (2018) allow for pervasive *unobserved* heterogeneity in the income dynamics and consumption insurance in the PSID. They do not consider differences in the income dynamics between sample and nonsample households and do not allow for separate permanent and transitory shocks to incomes.

importantly, however, 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, the degree of consumption insurance estimated on this set likely 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 explores the reasons for different income and consumption dynamics between sample and nonsample families; and Section 7 concludes.

## 2 Data

At the core of our study is the dataset used and made publicly available by Blundell, Pistaferri, and Preston (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 1978–1992. The focus on continuously married couples is to eliminate the potential effects of dramatic family composition change, such as divorce. As we discuss below, the actual implementation of data construction allows for sample females (but generally not sample males) to marry and divorce inside the 1978–1992 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 Differences in Insurance 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. 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,  $\alpha_i$ , 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,  $\tau_{it} = \epsilon_{it} + \theta\epsilon_{it-1}$ . 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:

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

where  $\Delta c_{it}$  is household  $i$ 's consumption growth at time  $t$ ,  $\xi_{it}$  is the permanent shock to household  $i$ 's disposable income,  $\epsilon_{it}$  is the transitory shock,  $\zeta_{it}$  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  $\phi$  and  $\psi$  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.<sup>3</sup>

Following BPP, we estimate  $\phi$  and  $\psi$ , 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,  $\sigma_\zeta^2$ , 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 Insurance Estimates for the Combined Sample and for Households Headed by Sample and Nonsample Males

In column (1) of Table 1, 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

---

<sup>3</sup>For instance, Blundell, Pistaferri, and Saporta-Eksten (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, Blundell, and Bonhomme (2017) study consumption insurance against persistent and transitory shocks to *household earnings* due to assets, and the tax and transfer system.

TABLE 1: CONSUMPTION INSURANCE FOR THE COMBINED SAMPLE AND FOR HOUSEHOLDS HEADED BY SAMPLE AND NONSAMPLE MALES

	Combined (1)	Sample (2)	Nonsample (3)
$\phi$ , transmission of perm. shock	0.6436 (0.0858)	0.9430 (0.1508)	0.4303 (0.0950)
$\psi$ , transmission of trans. shock	0.0291 (0.0436)	-0.0108 (0.0469)	0.1014 (0.1009)

*Notes:* Standard errors in parentheses. p-value for test of equal  $\phi$  ( $\psi$ ) between sample and nonsample families equals 0.4% (31%).

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 The Effects of Marriage and Divorce

As we mentioned above, the dataset construction by BPP treats sample and nonsample households asymmetrically by allowing households headed by nonsample males to be formed through marriage or end in divorce inside the 1978–1992 sample window while this is generally not allowed for the households headed by sample males. The reason for this asymmetry is technical but simple. BPP select the sample based on the following criteria: the head’s marital status does not change, and the head remains married to the same wife. Husbands of PSID sample females are only tracked by the PSID while they are married to PSID sample females. Thus, even if a PSID sample female marries or divorces within the sample window, her nonsample husband (defined by the PSID as the household head) will only be tracked while he is married to her, so his marital status as recorded by the PSID cannot change. In contrast, sample males are continuously tracked before and after the marriage, so all their marital status changes are recorded and, if present, are used to deselect them from the sample. The quantitative importance of this asymmetry is highlighted by Table 2 which indicates that only 8% of sample households and more than 50% of nonsample households are formed inside the sample window (the shares of households with code values for the family composition change variable 2 through 6).<sup>4</sup>

<sup>4</sup>Twenty-three sample males heading the households with code value 2 formed new households precisely in 1978. Fifty-seven sample households with code value 5 are all formed by PSID sample males who became heads of household within the 1978–1992 period for the first time in the PSID – they were never seen in the

TABLE 2: FAMILY COMPOSITION CHANGE IN THE YEAR OF FIRST ENTRANCE INTO THE SAMPLE

	Code value						Total
	0	1	2	4	5	6	
Sample	710 (73.58)	175 (18.13)	23 (2.38)	0 (0.00)	57 (5.91)	0 (0.00)	965 (100)
Nonsample	290 (36.25)	73 (9.13)	0 (0.00)	316 (39.50)	1 (0.13)	120 (15.00)	800 (100)

*Notes:* 0=“No change,” 1=“Change in members other than head or wife,” 2=“Head same but wife left/died and/or head has new wife,” 4=“Female head from previous year got married, husband (nonsample member) now head,” 5=“Some sample member other than head or wife has become head of this family unit,” 6=“Some female in family unit other than the previous-year head got married and nonsample member is now head.” Numbers in parentheses are percentages of the “Total.”

Newlywed or divorcing couples may experience substantial changes in spending or labor supply behavior at the start or end of their marriages which may lead to atypical income and/or consumption dynamics (Altonji and Vidangos 2014), and may potentially affect the consumption insurance estimates. Thus, we now assess whether the differences in insurance between sample and nonsample households documented above are induced by the asymmetric treatment of marriage and divorce between them.

To evaluate the importance of marriage, we group households into those who got married before the start of the sample in 1978 (the code for the family composition change variable in 1978 equals 0 or 1 in Table 2), and those who got married in 1978 or after (the code for the family composition change variable in the first year those households enter the sample equals 2, 4, 5 or 6). As reported in Panels A and B of Table 3, the estimated insurance of permanent shocks for sample households is close to none while nonsample households appear to insure a substantial fraction of permanent shocks regardless of whether we look at the couples who enter the BPP sample being married (columns 2 and 3), or marry into the sample (the results in column 5 are based only on eighty sample families formed in 1978 or later, and therefore are somewhat imprecise).

To assess the impact of keeping nonsample households whose marriage ends in divorce in the estimation sample, we restrict the data to only households that had been surveyed in 1992 – the last sample year. The results reported in Panel C of Table 3 continue to exhibit the same pattern: families headed by sample males appear to be substantially less insured against permanent shocks than families headed by nonsample males.

Finally, we consider the set of households that had been first surveyed in 1978 and last PSID either as single or married heads before 1978.

TABLE 3: CONSUMPTION INSURANCE: THE EFFECTS OF MARRIAGE AND DIVORCE

	Combined (1)	Sample (2)	Nonsample (3)	Combined (4)	Sample (5)	Nonsample (6)
	Panel A. Married before 1978			Panel B. Married in/after 1978		
$\phi$ , transmission perm. shock	0.7138 (0.0981)	0.9265 (0.1313)	0.5310 (0.1073)	0.5937 (0.1815)	1.4665 (0.5491)	0.3283 (0.1153)
$\psi$ , transmission trans. shock	0.0213 (0.0458)	-0.0369 (0.0428)	0.0554 (0.1174)	0.0205 (0.1309)	0.0762 (0.1281)	0.1570 (0.1277)
	Panel C. Surveyed in 1992			Panel D. Surveyed in 1978 & 92		
$\phi$ , transmission perm. shock	0.6293 (0.0914)	0.8920 (0.1584)	0.4802 (0.1044)	0.6844 (0.1787)	0.8986 (0.1602)	0.3665 (0.1032)
$\psi$ , transmission trans. shock	0.0249 (0.0463)	0.0269 (0.0481)	0.0445 (0.0990)	0.1077 (0.0399)	0.0231 (0.0481)	0.0886 (0.1378)

*Notes:* In Panel A, p-value for test of equal  $\phi$  ( $\psi$ ) between sample and nonsample families equals 3% (46%); in Panel B, the respective p-values are 4% and 65%; in Panel C, 3% and 87%; in Panel D, 1% and 65%.

surveyed in 1992, the first and last years in the BPP sample, respectively. This set of households lacks both newlywed couples and the couples divorcing inside the 1978–1992 period. The results are in Panel D. They are based on much smaller samples than the full sample, which is reflected in the precision of the estimates. Remarkably, the point estimates for the transmission coefficient for permanent shocks for sample and nonsample families are fairly similar to those obtained for the full sets of those households – nonsample households achieve a substantial insurance of permanent shocks while sample families achieve virtually none.

In summary, all of the panels point to the statistically significant difference in the estimated insurance of permanent shocks among sample and nonsample families. We do not find an important role of the differential selection of nonsample families with respect to marriage and divorce in inducing the differences in the extent of permanent insurance achieved by those families relative to their sample counterparts.

### 3.4 Consumption Insurance among Sample and Nonsample Households Matched on Observables

Sample and nonsample households differ in some observable characteristics which may, in turn, lead to different levels of consumption insurance. We, therefore, next, form pairs of

TABLE 4: CONSUMPTION INSURANCE FOR SAMPLE AND NONSAMPLE HOUSEHOLDS MATCHED ON EDUCATION, RACE, AND YEAR OF BIRTH OF HEAD

	Combined (1)	Sample (2)	Nonsample (3)
$\phi$ , transmission of perm. shock	0.6741 (0.1508)	1.0810 (0.3151)	0.4279 (0.1576)
$\psi$ , transmission of trans. shock	0.0371 (0.0612)	-0.0086 (0.0796)	0.0920 (0.1069)

*Notes:* Standard errors, calculated using a normal approximation to the interquartile range of 1,000 replications, reported in parentheses. p-value for test of equal  $\phi$  ( $\psi$ ) between sample and nonsample families equals 6% (45%).

sample and nonsample households of the same age, schooling, and race.<sup>5</sup> We have 501 of such pairs. These pairs are not unique – for example, there are sixteen households headed by white sample males born in 1920 with no college education and only one such household headed by a nonsample male; these households will create one pair which shares the same year of birth, schooling, and race but there are sixteen different pairs we could form. We, therefore, estimate insurance coefficients for a random set of 501 pairs of sample and nonsample households exactly matched on age, education, and race, repeat this exercise 1,000 times, and average the results. The results reported in Table 4 are somewhat less precise as they are based on a smaller set of households; they are similar, however, in terms of point estimates to the results in Table 1 – nondurable consumption of sample households fully absorbs permanent income shocks while nonsample households insure more than fifty percent of permanent shocks.

### 3.5 Consumption Insurance among Households Formed by PSID Sons and Daughters

Although our previous experiment matched families on head’s year of birth, education, and race, sample and nonsample families may still potentially differ on a variety of other characteristics. By construction, nonsample families do not include couples formed in 1968 but will contain females marrying nonsample males in 1969 and later, keeping their families intact until they are last observed during 1978–1992, or (re-)marrying during 1978–1992. To put selection of sample and nonsample families on an equal footing, we allow PSID sample males to (re-)marry and divorce during 1978–1992, keeping data for each newly-formed couple with

<sup>5</sup>Households are grouped into two education groups – with (some) college and no college education.

the same sample male head in the final dataset.<sup>6</sup> We further split the resulting dataset of sample families into those who had been married in 1968 and stayed married until they were last seen in 1978–1992, 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).<sup>7</sup>

In Appendix Table A-3 we tabulate 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 non-durable 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 things. We further used a LASSO regression for predicting if a family (among the set of households formed by sample sons and daughters) belongs to the group of nonsample families using a wide range of variables.<sup>8</sup> The results are depicted in Appendix Figure A-1 – although some regressors are picked by LASSO as having nonzero predictive power for the nonsample status, their predictive strength is minimal as the average prediction and the range of predicted values for sample and nonsample families are similar, and substantially deviate from their true respective values of zero and one. The analysis corroborates the conclusion based on a simple comparison of means in Table A-3 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 5 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.

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

---

<sup>6</sup>This selection is also recently used in Blundell, Pistaferri, and Saporta-Eksten (2016).

<sup>7</sup>Relative to the original BPP data, additional fourteen nonsample families are added as nonsample males from those families changed their marital status during 1978–1992 and were followed by the PSID after the change (some nonsample individuals were designated as followable since 1990).

<sup>8</sup>We are grateful to Max Sties for proposing this experiment.

TABLE 5: CONSUMPTION INSURANCE FOR THE UPDATED SAMPLE

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

*Notes:* p-value for test of equal  $\phi$  ( $\psi$ ) 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%).

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 5 we restrict the samples even further to 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.

### 3.6 Differences in Consumption Insurance among Sample and Nonsample Households by Gender of the Respondent

The characteristic in Table A-3 that differs the most between sample and nonsample households is that nonsample families have a higher incidence of wives responding to the survey. If male and female responses were different with respect to a persistent component of measurement error, this might induce different estimates of consumption insurance against permanent income shocks between sample and nonsample households. In this section, we present three pieces of evidence that lead us to conclude that the observed differences in gender of the respondent are not driving our findings.

First, in Table 6, we consider whether consumption insurance differs between households where either a male or a female spouse is permanently responding to the survey.<sup>9</sup> Considering a combined set of sample and nonsample households, or nonsample households separately, we observe that the transmission coefficients are not significantly affected by the gender of the

<sup>9</sup>Permanent responding status is less likely to be endogenous to income shocks and insurance.

TABLE 6: CONSUMPTION INSURANCE BY GENDER OF THE RESPONDENT

	Head respondent		Wife respondent	
	All (1)	Nonsamp. (2)	All (3)	Nonsamp. (4)
$\phi$ , transmission perm. shock	0.5920 (0.1030)	0.2149 (0.1233)	0.6080 (0.2630)	0.3812 (0.0913)
$\psi$ , transmission trans. shock	0.0092 (0.0541)	0.3899 (0.1986)	0.1294 (0.1090)	0.0845 (0.2352)

*Notes:* “All” includes sample and nonsample households. Columns (1) and (2) (Columns (3) and (4)) contain the results for households whose male head (female spouse) is recorded as the respondent in all survey waves.

respondent. In particular, we find that nonsample households have much lower transmission coefficients regardless of whether males or females respond to the survey.

In our second experiment, we directly compare the stochastic properties of measurement error defined as the difference between household survey earnings reports and administrative earnings records. Unfortunately, at this time, the PSID responses cannot be matched to administrative records. Instead, we consider the data from the Health and Retirement Study (HRS), which contains both survey responses to questions about household earnings and linked records from the Social Security Administration (SSA) Master Earnings File. The HRS is limited to individuals over the age of 50, so we can only assess the dynamic properties of measurement error for this sample.<sup>10</sup> Specific details of sample construction, estimation procedure, and the results can be found in Appendix III. Here it suffices to report that the autocovariance function for the growth rate in measurement error does not differ between male and female respondents to the survey and it is consistent with measurement error being no more persistent than a moving average process of order one. This implies that measurement error cannot affect our estimates of the permanent component of income and the degree of associated consumption insurance.

Finally, in Section 6, we will document that differences between sample and nonsample households are observed not only in the PSID but also in similar datasets from a number of other countries. Although those datasets maintain the overall design of the PSID, income questions are typically asked separately from each individual in the household, implying

<sup>10</sup>Underlying this experiment is an assumption that measurement error that might be present in administrative earnings data is uncorrelated with whether a man or a woman is designated as a respondent to a household survey. This assumption does not appear very strong, especially taking into account that designated financial respondents are generally fixed across the years in the HRS.

TABLE 7: ROBUSTNESS TO DIFFERENT INCOME AND CONSUMPTION MEASURES

	Sample, orig. (1)	Sample, sons (2)	Nonsamp. (3)	Sample, orig. (4)	Sample, sons (5)	Nonsamp. (6)
	Panel A. Nondur. cons., total earnings			Panel B. Nondur. cons., male earnings		
$\phi$ , transmission perm. shock	0.4544 (0.1087)	0.4986 (0.1149)	0.1519 (0.0604)	0.6606 (0.0704)	0.5690 (0.1170)	0.1443 (0.0448)
$\psi$ , transmission trans. shock	0.0347 (0.0336)	0.0723 (0.0519)	0.1693 (0.0691)	-0.0183 (0.0608)	0.0462 (0.0302)	0.0362 (0.0615)
	Panel C. Food, total earnings			Panel D. Food, male earnings		
$\phi$ , transmission perm. shock	0.3108 (0.0770)	0.3467 (0.0947)	0.1165 (0.0515)	0.4333 (0.1212)	0.3689 (0.0945)	0.0923 (0.0357)
$\psi$ , transmission trans. shock	0.0221 (0.0309)	0.0339 (0.0442)	0.1140 (0.0516)	0.0307 (0.0482)	-0.0176 (0.0266)	0.0363 (0.0479)

*Notes:* In Panel A, p-value for test of equal  $\phi$  ( $\psi$ ) between columns (1) and (3) equals 2% (8%), columns (2) and (3) 1% (26%); in Panel B, the respective p-values are <1% (88%) for columns (4) and (6) and <1% (53%) for columns (5) and (6); in Panel C, the respective p-values are 4% (13%) for columns (1) and (3) and 3% (24%) for columns (2) and (3); in Panel D, the respective p-values are <1% (92%) for columns (4) and (6) and <1% (43%) for columns (5) and (6).

that sample and nonsample households in those datasets do not differ in the identity of the respondent to income questions and yet the dynamic properties of survey earnings are different between sample and nonsample households.

### 3.7 Consumption Insurance among Sample and Nonsample Households for Other Income and Consumption Measures

In Table 7, we report results based on different concepts of income and consumption. To ensure that the results are not driven by potential miscalculation of household taxes or poor measurement of nonlabor income, we consider earnings as a source of risk to household budgets. Specifically, we report results based on the combined head's and wife's earnings in columns (1) and (2), and male earnings in columns (3) and (4). As was the case with net family income, nonsample households are found to be better insured against permanent shocks to earnings than sample households; see Panels A and B.

TABLE 8: ROBUSTNESS TO IMPUTATION PROCEDURES

Instruments:	BPP, real			Alt. IV, real		
	Sample, orig. (1)	Sample, sons (2)	Nonsamp. (3)	Sample, orig. (4)	Sample, sons (5)	Non-samp. (6)
$\phi$ , transmission	1.0517	0.8054	0.3895	1.3442	1.0296	0.4621
perm. shock	(0.1940)	(0.1956)	(0.0890)	(0.2530)	(0.2648)	(0.1087)
$\psi$ , transmission	0.0360	0.0258	0.1046	0.0573	0.0413	0.1461
trans. shock	(0.0404)	(0.0753)	(0.0888)	(0.0476)	(0.0957)	(0.1111)

*Notes:* p-value for test of equal  $\phi$  ( $\psi$ ) in columns (1) and (3) equals <1% (48%); in columns (2) and (3) equals 5% (50%); in columns (4) and (6) equals 5% (47%); in columns (5) and (6) equals <1% (46%). “BPP, real” estimation uses the original BPP instruments for imputation – the average, by cohort, year and education, head’s and wife’s hourly wages but in real terms. “Alt. IV, real” estimation uses the alternative instruments for imputation – the average, by cohort and year, real hourly head’s and wife’s wages.

In Panels C and D, we use food as a measure of household consumption, because food is free of potential imputation biases. Not surprisingly, food is better insured than nondurable consumption, but nonsample households are once again found to be substantially better insured against permanent income and earnings shocks than their sample counterparts.

### 3.8 Effects of Consumption Imputation

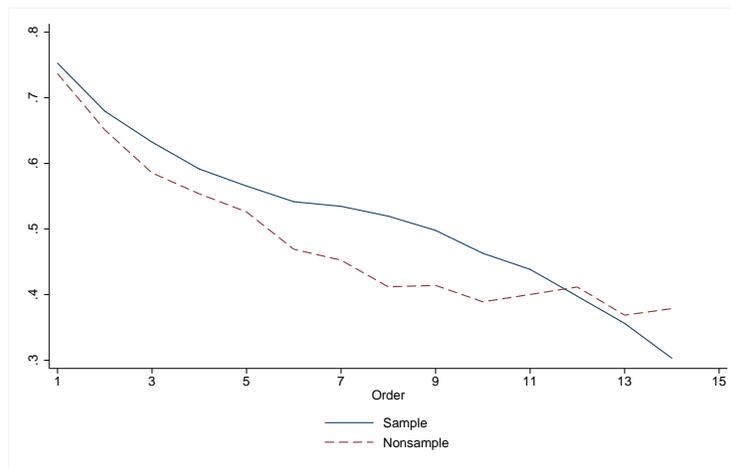
In Table 8 we consider alternative imputation procedures for nondurable consumption to evaluate further whether the procedure adopted in BPP induces the systematic difference in insurance between sample and nonsample households. BPP used the food demand equation estimated on Consumer Expenditure Survey data to impute nominal nondurable consumption to the PSID households. Since nominal nondurable expenditures, a right-hand-side variable of the equation, are potentially measured with error, BPP instrumented nominal nondurable consumption with the average – by cohort, year and education – hourly nominal head’s and wife’s wages. Campos and Reggio (2014) proposed to use instead real nondurable expenditures as a right-hand-side variable in the imputation equation and real hourly head’s and wife’s wages as instruments (“BPP, real” in Table 8); they also suggested alternative instruments – real hourly head’s and wife’s wages averaged only by cohort and year (“Alt. IV, real” in Table 8). The relative difference among sample and nonsample households in the insurance of permanent income shocks is preserved across alternative imputation procedures.

## 4 Income Processes of Sample and Nonsample Households

The body of evidence presented so far points to the 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, Perri, and Violante 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.

FIGURE 1: AUTOCORRELATION FUNCTION OF NET FAMILY INCOMES.



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

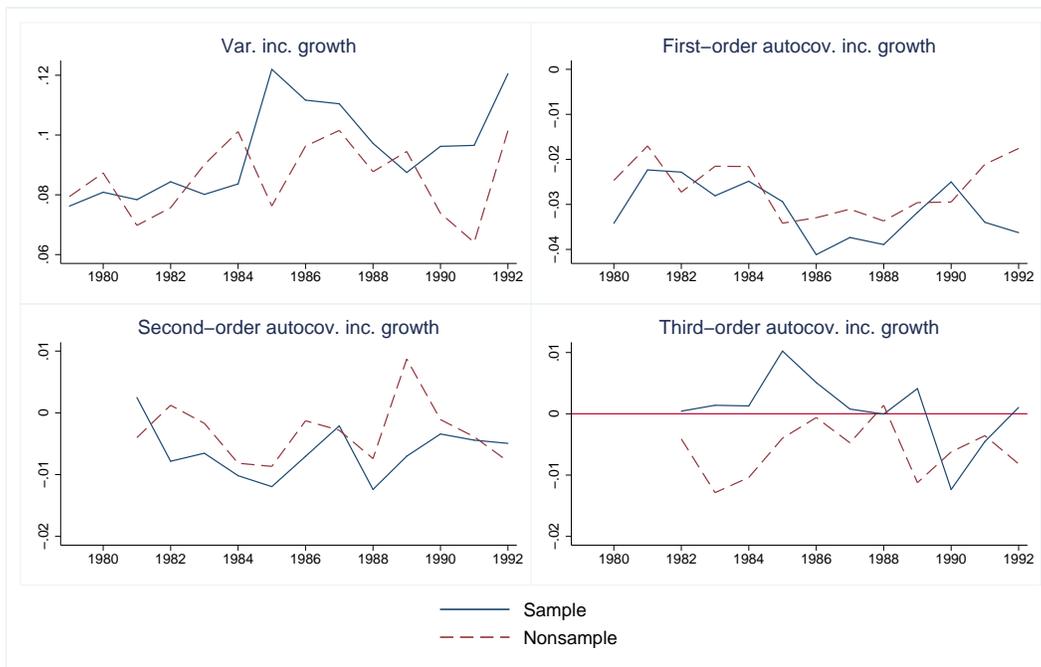
Figure 1 plots the autocorrelation functions of net family incomes for households headed by sample and nonsample males.<sup>11</sup> As can be seen, income dynamics differs markedly between the two sets of families. Although the income process of nonsample families appears less persistent, this evidence is not conclusive as the figures are also affected by potentially different variances of shocks.

We next examine the moments constructed from residual income growth which were targeted in the minimum-distance estimation above. Figure 2 suggests noticeable differences in these moments across sample and nonsample families. In particular, for nonsample families, the variance of income growth rates had not experienced any clear trend. The latter fact manifests itself in the correlation of just 27% between the variance of income growth rates for the two subsamples.<sup>12</sup> 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

<sup>11</sup>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.

<sup>12</sup>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.

FIGURE 2: DATA MOMENTS.



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.<sup>13</sup>

Finally, we assess the potential differences in income persistence between sample and nonsample families using GMM. This complements the preceding analysis because this estimation approach does not rely on the autocovariance function of incomes. We continue to assume that the (residual) income process for household  $i$  in year  $t$  is  $y_{it} = \alpha_i + p_{it} + \tau_{it}$ , but instead of imposing random walk, we now generalize the permanent component to an autoregressive process,  $p_{it} = \rho p_{it-1} + \xi_{it}$ . As before, we assume that permanent and transitory shocks to income are independent,<sup>14</sup> and that the transitory component is an MA(1) process,  $\tau_{it} = \epsilon_{it} + \theta \epsilon_{it-1}$ . We follow a large literature on estimating  $\rho$  in a GMM setting

<sup>13</sup>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.

<sup>14</sup>The assumption is standard. See Ejrnæs and Browning (2014) and Hryshko (2014) for exceptions.

TABLE 9: GMM ESTIMATES OF PERSISTENCE

	Overall data			Sibling pairs	
	Samp. orig. (1)	Samp. sons (2)	Nonsamp. (3)	Sons (4)	Daught. (5)
$\rho$ , persistence	0.94	0.96	0.82	0.92	0.79
perm. shock	(0.03)	(0.04)	(0.05)	(0.04)	(0.07)

*Notes:* The results from a two-step GMM estimation. Bootstrap standard errors in parentheses.

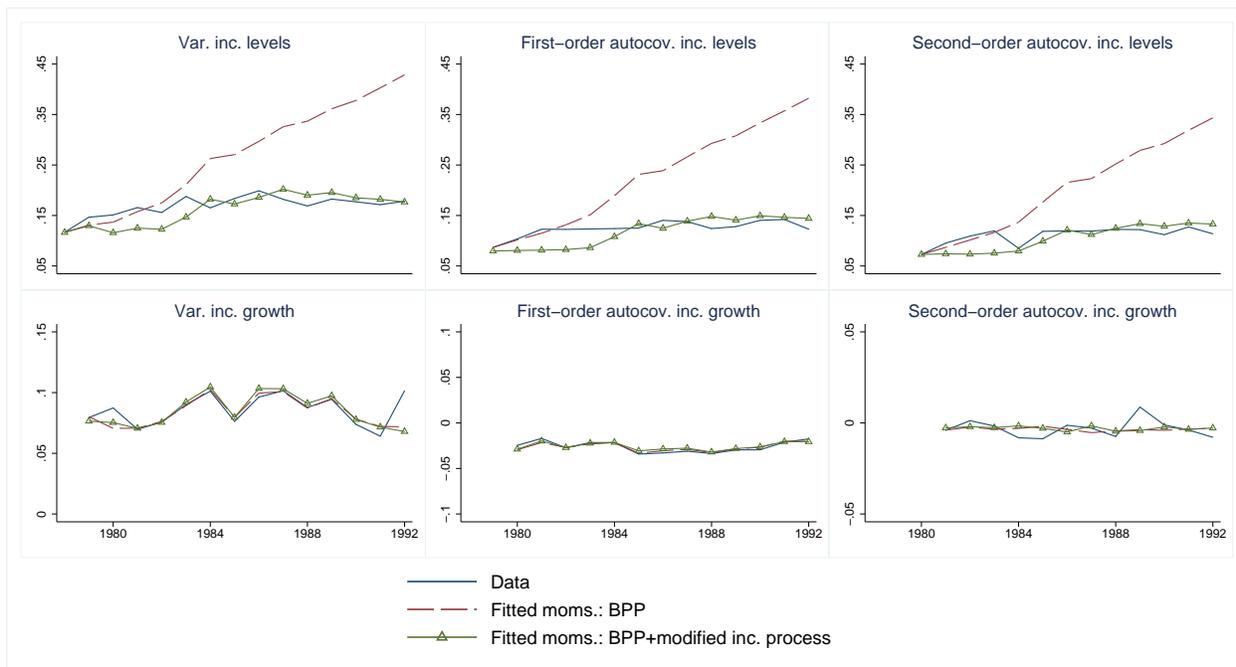
that commonly restricts  $\alpha_i$  to be an i.i.d. component; see, e.g., Arellano and Honoré (2001) for a review. We also assume that  $p_{i0} = m + \xi_{i0}$  so that  $y_{i0} = m + \alpha_i + \xi_{i0} + \epsilon_{i0} + \theta\epsilon_{i,-1}$ . We now rewrite the income process as  $y_{it} = (1 - \rho)\alpha_i + \rho y_{it-1} + \xi_{it} + \epsilon_{it} - (\rho - \theta)\epsilon_{it-1} - \rho\theta\epsilon_{it-2}$ . Given our assumptions, the time- $t$  quasi-difference  $y_{it} - \rho y_{it-1}$  will be uncorrelated with income growth measured at times  $t - j$ ,  $j \geq 3$ . In particular, we can use the following set of orthogonality conditions to identify  $\rho$ :  $E[(y_{it} - \rho y_{it-1})\Delta y_{it-j}] = 0$ ,  $t = 1982, \dots, 1992$ ,  $j \geq 3$ . This is the GMM estimator in levels; see Bun and Sarafidis (2015) for more details.

GMM is known to be biased when the number of moment conditions (that increase with the time dimension,  $T$ ) is large relative to the sample size ( $N$ ), which is the case for the estimation samples we consider; see, e.g., Alvarez and Arellano (2003). Following Chen, Chernozhukov, and Fernández-Val (2019), we use an estimator that corrects for such a bias by sample splitting along the cross-sectional dimension, and report the estimates that are averages of five random equal splits.<sup>15</sup> The two-step GMM estimates of the persistence of permanent shocks are reported in Table 9. A comparison of the estimates in column (3) to those in columns (1) and (2) reveals that the estimated persistence for nonsample families is lower than for sample families.<sup>16</sup> In columns (4) and (5) we report the results for sibling children of original PSID households who are observed in the data between 1978–1992 as heads or wives of their own households. Even for these smaller sets of families formed by brothers and sisters, the persistence of permanent income shocks is noticeably higher for the families of brothers, that is, sample families.

<sup>15</sup>Our conclusions are the same if we do not employ the debiased estimator and use the standard GMM instead.

<sup>16</sup>Adding first-differenced moments to the levels moments – the estimation known as system-GMM; see Blundell and Bond (1998) – delivers qualitatively similar results in that the estimated values of persistence for the sample families are higher than the corresponding values obtained for the nonsample families. We also estimated the income process using the methodology proposed by Browning, Ejrnæs, and Alvarez (2010) that relies neither on the autocovariance function of income levels and/or growth rates nor on the orthogonality conditions exploited here. The results once again revealed lower persistence of income shocks for nonsample households than for their sample counterparts. See also the results in Hryshko and Manovskii (2018b).

FIGURE 3: MODEL FIT. NONSAMPLE FAMILIES.



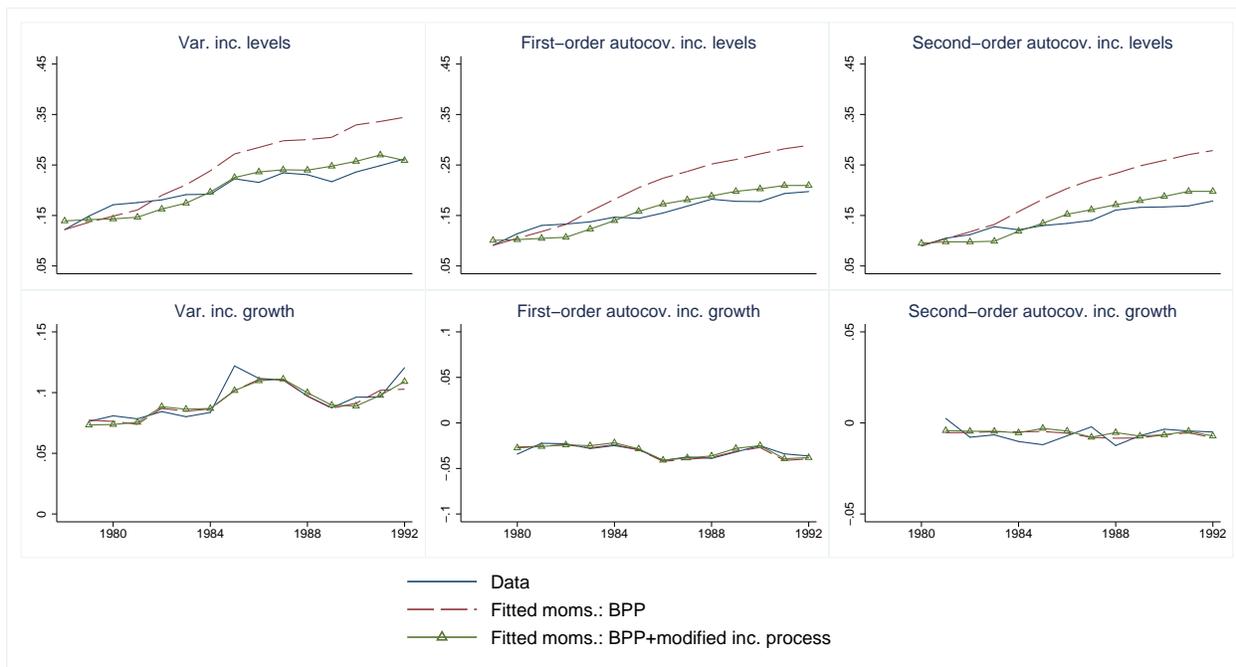
## 4.2 Different Income Dynamics and the Fit of Minimum Distance Estimation for Sample and Nonsample Families

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 Figure 3 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 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.

Figure 4 similarly plots the fit of the standard BPP model to the income moments in levels and growth rates for 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 1992 by about 30%, which is substantially lower relative to overestimation for the nonsample families described above.

FIGURE 4: MODEL FIT. SAMPLE FAMILIES.



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. We now use the identifying moments in Heathcote, Perri, and Violante (2010) to understand how this leads to a relatively poorer fit to the income moments in levels for nonsample households that we just described.

Specifically, consider the income process, which is the sum of the permanent random-walk component and an i.i.d. transitory shock.<sup>17</sup> The following moments identify the variances of permanent and transitory shocks,  $\sigma_{\xi,t}^2$  and  $\sigma_{\epsilon,t}^2$ , when targeting the moments of incomes in

<sup>17</sup>The assumption of an i.i.d. transitory component is reasonable given small estimates of the moving average parameter.

levels and growth rates, respectively:<sup>18</sup>

Differences :

$$\sigma_{\xi,t,\text{diffs}}^2 = E[\Delta y_{it}\Delta y_{it-1}] + E[\Delta y_{it}\Delta y_{it}] + E[\Delta y_{it}\Delta y_{it+1}] \quad (2)$$

$$\sigma_{\epsilon,t,\text{diffs}}^2 = -E[\Delta y_{it}\Delta y_{it+1}]. \quad (3)$$

Levels :

$$\sigma_{\xi,t,\text{levs}}^2 = E[y_{it}y_{it+1}] - E[y_{it}y_{it-1}] \quad (4)$$

$$\sigma_{\epsilon,t,\text{levs}}^2 = E[y_{it}y_{it}] - E[y_{it}y_{it+1}]. \quad (5)$$

The estimated variance of permanent and transitory shocks using these identifying moments in levels and differences should be identical if the true and estimated income processes are identical. If, however, the permanent component is an AR(1) process,  $p_{it} = \rho p_{it-1} + \xi_{it}$ , we show in Appendix I that the difference in the estimated moments (2) and (4), and (5) and (3) will equal, respectively:

$$\sigma_{\xi,t,\text{diffs}}^2 - \sigma_{\xi,t,\text{levs}}^2 = (1 - \rho)(\rho + \rho^3)\text{var}(p_{it-2}) > 0 \quad (6)$$

$$\sigma_{\epsilon,t,\text{levs}}^2 - \sigma_{\epsilon,t,\text{diffs}}^2 = \rho(1 - \rho)\text{var}(p_{it-1}) > 0. \quad (7)$$

Thus, restricting the permanent component to a random walk when its true persistence is lower will lead to inflated estimates of the variances of permanent (transitory) shocks when targeting the moments in growth rates (levels), to the extent the minimum-distance estimation relies on the identifying moments (2)–(5). Misspecification will lead to negligible biases if the persistence,  $\rho$ , is close to one; the biases, however, are expected to be larger for smaller values of  $\rho$ . As the results above point to a value of  $\rho$  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  $\rho$ .

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. As pointed out by Daly, Hryshko, and Manovskii (2016), 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

---

<sup>18</sup>Heathcote, Perri, and Violante (2010) used the moments identifying the variances of the shocks from biennial data but these moments are analogous if one instead relies on the annual data for identification of the variances in each year.

marriages for the newly-formed couples, or at the end of marriages for the couples which dissolve during 1978–1992. To assess the presence of these effects in our net family income data, in Appendix Table A-4, in a regression setting, we compare first and last income records within incomplete income spells with other income observations. The most prominent feature of the results is that the variance of incomes is high at the start of incomplete income spells relative to income observations from the interior of contiguous income spells.<sup>19</sup>

### 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  $\rho$ . Second, we follow Daly, Hryshko, and Manovskii (2016) 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  $\nu_{it}$  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.<sup>20</sup> The resulting income process is as follows:

$$\begin{aligned}
 y_{it} &= \alpha_i + p_{it} + \tau_{it} + \chi_{it}, \quad t = t_0, \dots, T \\
 p_{it} &= \rho p_{it-1} + \xi_{it} \\
 \tau_{it} &= \epsilon_{it} + \theta \epsilon_{it-1} \\
 \chi_{it+j} &= \begin{cases} \nu_{it} & \text{if } y_{it-1} \text{ or } y_{it+1} \text{ is missing and } t-1 \geq t_0, t+1 \leq T, j=0 \\ \theta \nu_{it} & j=1 \\ 0 & \text{otherwise,} \end{cases}
 \end{aligned} \tag{8}$$

where  $\alpha_i$  is individual  $i$ 's fixed effect,  $t_0$  is the first sample year (1978), and  $T$  is the last sample year (1992).

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-4. We estimate the model by the method of simulated minimum distance, assuming that persistent, transitory, and  $\nu$ -shocks are drawn from normal distributions, and using the diagonal weighting

<sup>19</sup>The finding that the levels of family incomes are not significantly different at the start and end of incomplete income spells is in contrast to the pattern documented for male earnings data in Daly, Hryshko, and Manovskii (2016). This is likely due to the smoothing of shocks to male earnings provided by the tax and transfer system and family labor supply.

<sup>20</sup>Daly, Hryshko, and Manovskii (2016) 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.

TABLE 10: CONSUMPTION INSURANCE. AR(1) PERMANENT COMPONENT

	Sample (1)	Nonsample (2)
$\rho$ , AR coeff.	0.9956 (0.0095)	0.9003 (0.0314)
$\sigma_{\xi}^2$ , var. perm. shock (avg.)	0.0151 (0.0029)	0.0285 (0.0052)
$\theta$ , MA coeff.	0.1315 (0.0252)	0.0456 (0.0653)
$\sigma_{\epsilon}^2$ , var. trans. shock (avg.)	0.0407 (0.0027)	0.0274 (0.0042)
$\sigma_u^2$ , var. cons. meas. err (avg.)	0.0704 (0.0033)	0.0740 (0.0082)
$\phi$ , transmission perm. shock	0.9903 (0.1770)	0.4798 (0.1142)
$\psi$ , transmission trans. shock	0.0922 (0.0393)	0.0997 (0.1112)

*Notes:* Simulated minimum distance estimates. Standard errors in parentheses.

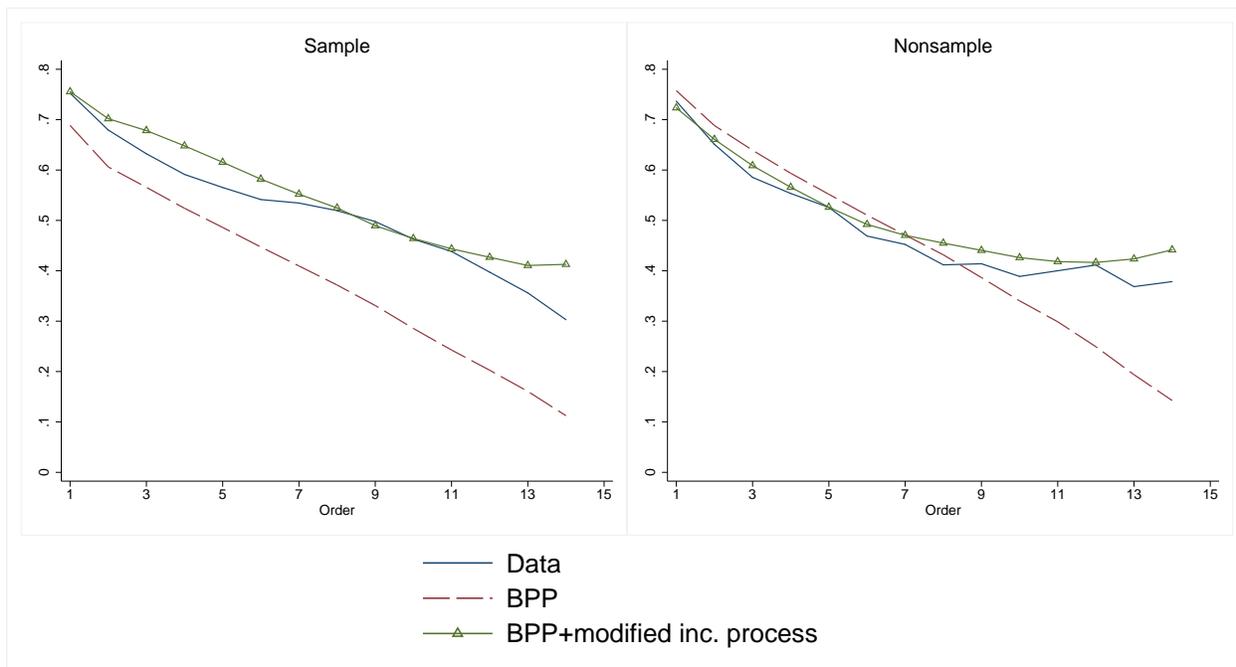
matrix calculated by block-bootstrap.<sup>21</sup> In estimations, we assumed that the fixed effect in family incomes is independent of the shocks.

Table 10 contains estimation results.<sup>22</sup> 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

<sup>21</sup>We verified that the simulated method of moments with the assumption of normal permanent and transitory shocks delivers virtually the same parameter estimates as the estimations which assume that the shocks are drawn from a fat-tailed Student t-distribution, the degrees of freedom of which are estimated by matching kurtosis of residual consumption and income growth observed in the data. We also experimented with updating the BPP estimation to allow for fitting the third moments of income and consumption growth, but we did not find any substantial differences in the estimated insurance against permanent income shocks when fitting both the second and third moments. Importantly, the differences in the estimated permanent insurance for sample versus nonsample families remain a robust feature of the data.

<sup>22</sup>As the transmission coefficient for  $\nu$ -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.

FIGURE 5: MODEL FIT. AUTOCORRELATION FUNCTION OF INCOME LEVELS.



sample and 0.48 for nonsample families, respectively.<sup>23,24</sup> In Figures 3 and 4 we show the fit of the models (lines with triangles) in Table 10 to the data moments. The models with a modified 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 Figure 5, 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 10. 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 allow-

<sup>23</sup>The transmission of permanent shocks is higher than the value estimated under the assumption of a random-walk permanent component. This is consistent with the results in Hryshko and Manovskii (2018a) that allow for infrequent shocks  $\nu$  in estimation and the theoretical prediction in Appendix I.

<sup>24</sup>Note that the values of the estimated persistence are higher than the values reported in Table 9. This is not surprising as the methods use different information; in particular, Table 10 uses, in addition, consumption information to identify the parameters of the income process. Han and Phillips (2010) show that system-GMM estimates of the persistence may be downward-biased when the true persistence is close to unity.

ing for additional income variance at the extremes of contiguous income spells. Noteworthy, none of the moments in Figure 5 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 10. 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  $\beta$ . 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

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

$Y_{it}$  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.  $\mu_t$  is the common

life cycle component of income;  $P_{it}$  is the permanent component of income, the log of which follows an AR(1) process with persistence  $\rho$ ;  $\xi_{it}$  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  $\pi$  and is positive otherwise, with the values determined by an i.i.d. transitory shock  $\epsilon_{it}$ .<sup>25</sup> After retirement, household  $i$ 's income is proportional to the permanent component at age  $t_R$  with a replacement rate  $\kappa$ , as in, e.g., Demyanyk, Hryshko, Luengo-Prado, and Sørensen (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.<sup>26</sup>

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, Skinner, and Zeldes (1994). The age-dependent deterministic income profile,  $\mu_t$ , is taken from Kaplan and Violante (2010). The replacement rate  $\kappa$  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 10; we assume that the shocks  $\xi_{it}$  and  $\epsilon_{it}$  are normally distributed. We then calibrate the CRRA coefficient  $\gamma$ , the probability of a transitory zero-income state,  $\pi$ , and the time discount factor  $\beta$  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 1978–1992 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

---

<sup>25</sup>Since the estimated moving average parameters in Table 10 are small, for simplicity, we assume that transitory shocks are i.i.d.

<sup>26</sup>Following Hryshko, Luengo-Prado, and Sørensen (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.

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

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 the PSID data, we drop income growth outliers in the simulated data. The values of internally calibrated parameters are shown in the three bottom rows of Table 11.<sup>27</sup>

### 5.3 Quantitative Findings

First, observe that 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 11). 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 10. 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 3.5, 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 10.

The results are in Table 12. Columns (1) and (3) reproduce our empirical results from Table 5, columns (1) and (2), 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

---

<sup>27</sup>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).

TABLE 11: 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	27.9	28.6	27.7	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 parameters				
Time disc. factor, $\beta$			0.969**	
Coeff. RRA, $\gamma$			0.405**	
Prob. of zero inc. state, $\pi$			0.006**	

*Notes:* \* indicates calibration targets; \*\* marks calibrated parameter values.

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 then recalibrate the same three parameters,  $\gamma$ ,  $\pi$ , and  $\beta$ , 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.

The main substantive takeaway from these findings is on the relationship between the de-

TABLE 12: CONSUMPTION INSURANCE IN SIMULATED AND PSID DATA

	Sons		Daught.		Combined Sons & Daught.		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Data	Model	Data	Model	Data	Model	RW Model
$\phi$ , transmission	0.865	0.839	0.456	0.446	0.570	0.615	0.918
perm. shock	(0.193)	(0.134)	(0.101)	(0.062)	(0.090)	(0.061)	(0.116)
$\psi$ , transmission	0.065	0.055	0.120	0.139	0.085	0.083	0.057
trans. shock	(0.080)	(0.046)	(0.089)	(0.062)	(0.062)	(0.037)	(0.064)

*Notes:* Minimum distance estimates for the model and PSID data. 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.

gree 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 sampled and 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 Current Population Survey, 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, to our knowledge, 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.

TABLE 13: GMM ESTIMATES OF PERSISTENCE: PSID vs. HRS-SSA

	Family earnings		Male earnings	
	PSID (1)	HRS-SSA (2)	PSID (3)	HRS-SSA (4)
$\rho$ , 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 GMM estimation. Bootstrap standard errors in parentheses. Appendix II describes sample selections.

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 Health and Retirement Study (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.<sup>28</sup> 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 1998 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 13 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-1.

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

<sup>28</sup>See <https://hrs.isr.umich.edu/documentation/survey-design>.

TABLE 14: 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)
$\rho$ , 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 GMM estimation. Bootstrap standard errors in parentheses. See Appendix IV for the data description and sample selections.

toward individuals with more persistent earnings dynamics.<sup>29</sup> 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).<sup>30</sup> 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 14 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.<sup>31</sup> Thus, the discrepancy cannot be caused by the differential tendency

<sup>29</sup>Of course, it is also possible that the HRS is not representative.

<sup>30</sup>A description of these datasets and the details of sample construction are provided in Appendix IV.

<sup>31</sup>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, albeit being somewhat higher for males, 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.

TABLE 15: ATTRITION RATES, IN PERCENT

Dataset: (Country):	PSID (U.S.A.)	GSOEP (Germany)	BHPS (U.K.)	HILDA (Australia)	KLIPS (Korea)	SHP (Switzerland)
	(1)	(2)	(3)	(4)	(5)	(6)
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

*Notes:* See Appendix IV for the details on sample selections.

of males and females to respond to the income questions between sample and nonsample households, confirming our conclusions in Section 3.6.

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, Gottschalk, and Moffitt (1998b) and Fitzgerald (2011) note that PSID males are more likely to attrit than PSID females. We confirm this pattern in Table 15, where we follow individuals present in the 1968 survey after their entry into the sample (as heads, wives, or children) and check whether they attrit in the following 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, Gottschalk, and Moffitt (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 that do and do not exit the data in the subsequent eight years. The results, summarized in Table 16, imply that 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 17, we report separate estimates of persistence for households that will and will not eventually attrit in each of the datasets we consider. A clear pattern is evident:

TABLE 16: GMM ESTIMATES OF PERSISTENCE BY ATTRITION. PSID

	Samp.		Nonsamp.	
	Non-attr. (1)	Attr. (2)	Non-attr. (3)	Attr. (4)
$\rho$ , 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 GMM estimation. Bootstrap standard errors in parentheses.

TABLE 17: 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.)	
	NA (1)	A (2)	NA (3)	A (4)	NA (5)	A (6)	NA (7)	A (8)	NA (9)	A (10)	NA (11)	A (12)
$\rho$ , persist.	0.91	0.81	0.99	0.75	0.88	0.83	0.87	0.66	0.87	0.56	0.96	0.66
perm. shock	(0.04)	(0.09)	(0.04)	(0.08)	(0.03)	(0.09)	(0.04)	(0.11)	(0.06)	(0.11)	(0.07)	(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-attriters and attriters, respectively. The results from a two-step GMM estimation. Bootstrap standard errors in parentheses. See Appendix IV for the data description and sample selections.

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 1998. 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.

This evidence suggests that both sets of sample and nonsample households are affected by attrition and do not provide a perfect measure of income dynamics and consumption insurance. However, the set of nonsample households is more representative and thus provides a better guide to the relevant measures for the U.S. population.

## 7 Conclusion

The dynamic properties of earnings and 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 standard 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 non-sample 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 workhorse 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 headed by sample and nonsample males 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 we consider, 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. More research on this is called for, but as a practical matter, as the set of nonsample families is less selected, it likely provides a better guide to the income dynamics and the degree of consumption insurance in the U.S.

## References

- ABOWD, J., AND D. CARD (1989): “On the Covariance Structure of Earnings and Hours Changes,” *Econometrica*, 57, 411–445.
- ABOWD, J. M., AND M. H. STINSON (2013): “Estimating Measurement Error in Annual Job Earnings: A Comparison of Survey and Administrative Data,” *The Review of Economics and Statistics*, 95, 1451–1467.
- ALAN, S., M. BROWNING, AND M. EJRNÆS (2018): “Income and Consumption: A Micro Semistructural Analysis with Pervasive Heterogeneity,” *Journal of Political Economy*, 126, 1827–1864.
- ALTONJI, J. G., AND I. VIDANGOS (2014): “Marriage Dynamics, Earnings Dynamics, and Lifetime Family Income,” Society for Economic Dynamics Meeting Papers 1230.
- ALVAREZ, J., AND M. ARELLANO (2003): “The Time Series and Cross-Section Asymptotics of Dynamic Panel Data Estimators,” *Econometrica*, 71, 1121–1159.
- ARELLANO, M., R. BLUNDELL, AND S. BONHOMME (2017): “Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework,” *Econometrica*, 85, 693–734.
- ARELLANO, M., AND B. HONORÉ (2001): “Panel Data Models: Some Recent Developments,” in *Handbook of Econometrics*, Vol. 5, ed. by J. J. Heckman, and E. E. Leamer, chap. 53, pp. 3229–3296. Elsevier, Amsterdam.
- BECKETTI, S., W. GOULD, L. LILLARD, AND F. WELCH (1988): “The Panel Study of Income Dynamics after Fourteen Years: An Evaluation,” *Journal of Labor Economics*, 6, 472–492.
- BLUNDELL, R. (2014): “Income Dynamics and Life-Cycle Inequality: Mechanisms and Controversies,” *Economic Journal*, 124, 289–318.
- BLUNDELL, R., AND S. BOND (1998): “Initial Conditions and Moment Restrictions in Dynamic Panel Data Models,” *Journal of Econometrics*, 87, 115–143.
- BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): “Consumption Inequality and Partial Insurance,” *American Economic Review*, 98, 1887–1921.
- BLUNDELL, R., L. PISTAFERRI, AND I. SAPORTA-EKSTEN (2016): “Consumption Inequality and Family Labor Supply,” *American Economic Review*, 106, 387–435.
- BOUND, J., C. BROWN, G. J. DUNCAN, AND W. L. RODGERS (1994): “Evidence on the Validity of Cross-Sectional and Longitudinal Labor Market Data,” *Journal of Labor Economics*, 12, 345–368.
- BOUND, J., AND A. B. KRUEGER (1991): “The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?,” *Journal of Labor Economics*, 9, 1–24.

- BROWNING, M., M. EJRNÆS, AND J. ALVAREZ (2010): “Modelling Income Processes with Lots of Heterogeneity,” *Review of Economic Studies*, 77, 1353–1381.
- BUN, M. J., AND V. SARAFIDIS (2015): “Dynamic Panel Data Models,” in *The Oxford Handbook of Panel Data*, ed. by B. H. Baltagi, chap. 3, pp. 76–110. Oxford University Press, Oxford.
- CAMPOS, R. G., AND I. REGGIO (2014): “Measurement Error in Imputation Procedures,” *Economics Letters*, 122, 197–202.
- CARROLL, C. D. (1992): “The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence,” *Brookings Papers on Economic Activity*, 2, 61–156.
- (1997): “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *Quarterly Journal of Economics*, 112, 1–55.
- (2009): “Precautionary Saving and the Marginal Propensity to Consume Out of Permanent Income,” *Journal of Monetary Economics*, 56, 780–790.
- CHEN, S., V. CHERNOZHUKOV, AND I. FERNÁNDEZ-VAL (2019): “Mastering Panel Metrics: Causal Impact of Democracy on Growth,” *AEA Papers and Proceedings*, 109, 77–82.
- DALY, M., D. HRYSHKO, AND I. MANOVSKII (2016): “Improving the Measurement of Earnings Dynamics,” NBER Working Paper # 22938.
- DEMYANYK, Y., D. HRYSHKO, M. J. LUENGO-PRADO, AND B. E. SØRENSEN (2017): “Moving to a Job: The Role of Home Equity, Debt, and Access to Credit,” *American Economic Journal: Macroeconomics*, 9, 149–181.
- EJRNÆS, M., AND M. BROWNING (2014): “The Persistent-Transitory Representation for Earnings Processes,” *Quantitative Economics*, 5, 555–581.
- FITZGERALD, J., P. GOTTSCHALK, AND R. MOFFITT (1998a): “An Analysis of Sample Attrition in Panel Data: The Michigan Panel Study of Income Dynamics,” *Journal of Human Resources*, 33, 251–299.
- (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., AND J. PARKER (2002): “Consumption Over the Life Cycle,” *Econometrica*, 70, 47–89.
- HAN, C., AND P. C. B. PHILLIPS (2010): “GMM Estimation for Dynamic Panels with Fixed Effects and Strong Instruments at Unity,” *Econometric Theory*, 26, 119–151.

- HEATHCOTE, J., F. PERRI, AND G. L. VIOLANTE (2010): “Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 1967–2006,” *Review of Economic Dynamics*, 13, 15–51.
- HRYSHKO, D. (2014): “Correlated Income Shocks and Excess Smoothness of Consumption,” *Journal of Economic Dynamics and Control*, 48, 41–62.
- HRYSHKO, D., M. J. LUENGO-PRADO, AND B. E. SØRENSEN (2011): “Childhood Determinants of Risk Aversion: The Long Shadow of Compulsory Education,” *Quantitative Economics*, 2, 37–72.
- HRYSHKO, D., AND I. MANOVSKII (2018a): “Income Dynamics and Consumption Insurance,” Working paper, University of Pennsylvania.
- (2018b): “On the Heterogeneity in Family Earnings and Income Dynamics in the PSID,” *AEA Papers and Proceedings*, 108, 292–296.
- HUBBARD, R. G., J. SKINNER, AND S. P. ZELDES (1994): “The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving,” *Carnegie-Rochester Conference Series on Public Policy*, 40, 59–125.
- KAPLAN, G., AND G. L. VIOLANTE (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 *Panel Surveys*, ed. by D. Kasprzyk, G. Duncan, G. Kalton, and M. P. Singh, pp. 497–511. Wiley Series in Probability and Statistics.
- PATTISON, D., AND H. WALDRON (2018): “Trends in Elective Deferrals of Earnings from 1990-2001 in Social Security Administrative Data,” Social Security Administration, Research and Statistics Note No. 2008-03.

# Appendix

## I Derivations

If the true permanent component is an AR(1) process with persistence  $\rho < 1$ , imposing instead  $\rho = 1$  yields the misspecified variance of permanent shocks using the moments in levels:

$$\sigma_{\xi,t,\text{levs}}^2 = E[y_{it}y_{it+1}] - E[y_{it}y_{it-1}] = (\rho^3 - \rho)\text{var}(p_{t-1}) + \rho\sigma_{\xi_t}^2.$$

Using the moments in differences instead, the measured variance will equal

$$\sigma_{\xi,t,\text{diffs}}^2 = E\left[\Delta y_{it} \sum_{j=-1}^1 \Delta y_{it+j}\right] = (\rho - 1)(\rho^4 - \rho)\text{var}(p_{t-2}) + (\rho^3 - \rho^2)\sigma_{\xi_{t-1}}^2 + \rho\sigma_{\xi_t}^2.$$

Since  $\text{var}(p_{t-1}) = \rho^2\text{var}(p_{t-2}) + \sigma_{\xi_{t-1}}^2$ ,

$$\sigma_{\xi,t,\text{diffs}}^2 - \sigma_{\xi,t,\text{levs}}^2 = \text{var}(p_{t-2}) [(\rho - 1)(\rho^4 - \rho) - \rho^2(\rho^3 - \rho)] = (1 - \rho)(\rho + \rho^3)\text{var}(p_{t-2}),$$

which is greater than zero for  $0 < \rho < 1$ .

The misspecified variance of transitory shocks using the moments in levels equals

$$\sigma_{\epsilon,t,\text{levs}}^2 = E[y_{it}y_{it}] - E[y_{it}y_{it+1}] = \rho^2(1 - \rho)\text{var}(p_{t-1}) + (1 - \rho)\sigma_{\xi_t}^2 + \sigma_{\epsilon,t}^2,$$

while using instead the moment in differences it equals

$$\sigma_{\epsilon,t,\text{diffs}}^2 = -E[\Delta y_{it}\Delta y_{it+1}] = -\rho(1 - \rho)^2\text{var}(p_{t-1}) + (1 - \rho)\sigma_{\xi_t}^2 + \sigma_{\epsilon,t}^2.$$

This implies that

$$\sigma_{\epsilon,t,\text{levs}}^2 - \sigma_{\epsilon,t,\text{diffs}}^2 = \rho(1 - \rho)\text{var}(p_{t-1}),$$

which, again, is greater than zero for  $0 < \rho < 1$ .

The identifying moment for the transmission of permanent shocks is

$$\hat{\phi}_t = \frac{E[\Delta c_{it} \sum_{j=-1}^1 \Delta y_{it+j}]}{E[\Delta y_{it} \sum_{j=-1}^1 \Delta y_{it+j}]} = \frac{\phi_t \rho \sigma_{\xi_t}^2}{(\rho - 1)(\rho^4 - \rho)\text{var}(p_{t-2}) + (\rho^3 - \rho^2)\sigma_{\xi_{t-1}}^2 + \rho\sigma_{\xi_t}^2}.$$

Assuming that the variance of persistent shocks does not change over time,  $\text{var}(p_{t-2}) = \frac{1 - \rho^{2(t-2)}}{1 - \rho^2}\sigma_{\xi}^2$ , it follows that

$$\begin{aligned} \hat{\phi}_t &= \frac{E[\Delta c_{it} \sum_{j=-1}^1 \Delta y_{it+j}]}{E[\Delta y_{it} \sum_{j=-1}^1 \Delta y_{it+j}]} = \frac{\rho\phi_t}{(\rho - 1)(\rho^4 - \rho)\frac{1 - \rho^{2(t-2)}}{1 - \rho^2} + \rho^3 - \rho^2 + \rho} \\ &= \frac{\phi_t}{(1 - \rho^3)\frac{1 - \rho^{2(t-2)}}{1 + \rho} - \rho(1 - \rho) + 1}. \end{aligned}$$

## II Comparison of HRS-SSA and PSID earnings data

For our comparison with the PSID, we use the 1998 wave of the Health and Retirement Study (HRS) that contains a comprehensive coverage of individuals age fifty-one and older – representative of the corresponding U.S. population – and their spouses. In 1998, eighty-three percent of the HRS families had at least one member who granted permission to the HRS to access their historical earnings records from the Social Security Administration (SSA).<sup>32</sup> For males (and their spouses) present in the 1998 wave, we use confidential earnings records extracted from W2 tax forms for the period 1978–1992, which coincides with our PSID sample period.<sup>33</sup> SSA earnings between 1 and 49, 250,000 and 300,000, 300,000 and 500,000 dollars are bracketed, and above 500,000 are set to missing: we assign midpoints to the observations falling within those brackets both in the HRS and in the PSID and set the PSID income records above 500,000 dollars to missing.<sup>34</sup> Since SSA earnings for self-employed individuals are top coded until 1994, we drop annual records for self-employed individuals in both datasets. To correct for a nonrandom selection of individuals and couples with matches to the administrative earnings data, we utilize individual-level and family-level social security weights. As in the HRS, in the PSID we select sample males and females (and their families) born before 1948 and present in the 1999 wave (there was no 1998 wave in the PSID; the results based on the 1997 PSID wave are very similar), and use their earnings data for 1978–1992, with the same sample restrictions as in the HRS. Our first-stage regressions for the HRS and the PSID contain the same variables: the full set of year of birth dummies, and interactions of race, and education dummies with the full set of year dummies.

## III Comparison of HRS-SSA and HRS survey earnings data

To get an idea about the extent of measurement error in family earnings in survey data, we use nonimputed earnings data from the HRS, and administrative data from the SSA matched to the HRS survey years 1992–2012. Our measure of administrative earnings includes deferred compensations which gained importance in the early 1990s. We select males and their spouses present in the 2010 HRS wave, born before 1960, and use the social security weights to correct for selective matching.<sup>35</sup>

Since sampling in the HRS is biennial, we drop administrative earnings records for those years that are missing in the HRS. We further select earnings records for couples whose male is aged 30–65, and drop the data if there are earnings growth outliers in the survey earnings reports. The HRS designates one spouse as a financial respondent who is, in particular, answering individual earnings questions for both spouses – this allows us to explore the differential dynamics of measurement error in family earnings separately for male and female

---

<sup>32</sup>These data are not public but can be accessed by entering a contractual agreement with the HRS.

<sup>33</sup>Our measure of earnings is mostly based on the information from W2 Box 1 (“wages, tips, other compensation”) but, due to concerns that some Box 1 records were zeroed out in the process of record updating, we also utilize information from Box 3 (“Social security wages”) and Box 5 (“Medicare wages”), as in Pattison and Waldron (2018).

<sup>34</sup>The numbers of such observations both in the HRS and PSID are small.

<sup>35</sup>Many of these households have both survey and administrative data available for 2012, which we also utilize.

TABLE A-1: AUTOCOVARIANCE STRUCTURE OF MEASUREMENT ERROR  
IN FAMILY EARNINGS IN HRS

	Autocovariance of order:			
	0	1	2	3
Male resp. (s.e.)	0.25 (0.04)	-0.10 (0.03)	-0.01 (0.01)	-0.003 (0.009)
Female resp. (s.e.)	0.22 (0.03)	-0.06 (0.01)	-0.002 (0.006)	0.008 (0.009)

*Notes:* Standard errors, clustered by household, in parentheses.

financial respondents. Our data contain 1,254 (689) couples, with at least two observations on family earnings, whose male (female) is the financial respondent in all waves.<sup>36</sup>

Following Bound and Krueger (1991) and Bound, Brown, Duncan, and Rodgers (1994), we define (log) measurement error as the difference between log family earnings reported in the HRS and log family earnings from SSA data. We then use residuals from a regression of the error on the full set of year of birth dummies, and interactions of race, and education dummies with the full set of year dummies.<sup>37</sup> Both for female and male respondents, the median log error is zero, the mean is 0.09 for females and 0.08 for males (driven by large positive outliers), and the standard deviation is at 0.62 for men and 0.58 for women.

In Table A-1, we show the autocovariance function for the growth rate in measurement error, separately for male and female respondents. Clearly, the dynamics of measurement error does not depend on who provides the survey earnings reports. Both for male and female respondents, the autocovariance function is significant up to the first order – since our data are biennial, this is consistent with measurement error being no more persistent than a moving average process of order one.

<sup>36</sup>There are only forty-eight couples with switching financial respondents in different waves.

<sup>37</sup>The results are nearly the same if use the raw data instead.

## IV Non-U.S. datasets

There are several datasets designed similarly to the PSID, which can be used to compare the income dynamics of sample versus nonsample families, attrition rates for males versus females, and the persistence of permanent income shocks for attritors versus non-attritors. The datasets we use are the German Socio-economic Panel (GSOEP), the British Household Panel Survey (BHPS), the Household, Income and Labour Dynamics in Australia (HILDA), the Korean Labor and Income Panel Study (KLIPS), and the Swiss Household Panel (SHP).<sup>38</sup>

The GSOEP started in 1984, representing the entire population of the Federal Republic of Germany. We use data for the period 1984–2017 for the representative “West German Sample” comprised of 4,528 households whose head did not belong to a group of foreign “guest workers.”

The BHPS started in 1991 with a nationally representative sample of more than 5,000 households, adding various subsamples later on, and was discontinued in 2008. We use data for individuals from the original sample for the period 1991–2008.

HILDA started in 2001 by interviewing about 20,000 individuals from a nationally representative sample of nearly 8,000 Australian households. We use data for Australian-born individuals from the original sample interviewed during the period 2001–2017.

The KLIPS started in 1998 with an equal probability sample of 5,000 households and their members from seven metropolitan cities and urban areas in eight provinces. We use data for the original sample of individuals interviewed during the period 1998–2016.

The SHP started in 1999 with a random sample of 5,074 households and their members and later added two refreshment samples in 2004 and 2013. We use data for the original sample of individuals interviewed during the period 1999–2017.

***Persistence of income shocks for sample and nonsample families.*** For each dataset, we select married couples, observed in the data consecutively, with at least one sample member,<sup>39</sup> and a male spouse aged 25–65. As in the PSID, we define sample families as those with a male spouse being a sample member and nonsample families as those with a nonsample male spouse. Since we do not have a measure of taxes for each dataset, our income measure is gross family income. As our main analysis for the PSID is for net family income, we also report PSID results for gross family income in Table 14. We drop income growth outliers in each dataset, and use residuals from a first-stage regression of gross family income on the same controls as in Section 3. We assume the same income process – a sum of an AR(1) and MA(1) processes – everywhere.

***Attrition rates.*** Attrition rates – reported in Table 15 – are calculated for a subset of the sample individuals who can be potentially observed in each dataset for at least fifteen years from age twenty-five until they turn age sixty-five. In the PSID, we first select heads, spouses, and children ages 0–50 in 1968 – the first survey year.<sup>40</sup> In the GSOEP, we select

---

<sup>38</sup>Access to HILDA and GSOEP requires a formal application process while the other datasets are publicly available.

<sup>39</sup>They are called “original sample members” in the BHPS, KLIPS, and SHP, and “continuing sample members” in HILDA.

<sup>40</sup>Those born in 1968 turn thirty in 1998 and can be potentially observed for seventeen years, until 2015, our last survey year, when they are forty-seven.

TABLE A-2: ATTRITION RATES, IN PERCENT

Dataset: (Country):	PSID (U.S.A.)	GSOEP (Germany)	BHPS (U.K.)	HILDA (Australia)	KLIPS (Korea)	SHP (Switzerland)
	(1)	(2)	(3)	(4)	(5)	(6)
Men	41.4	51.0	55.7	58.4	67.2	74.6
Women	36.7	48.9	47.5	56.2	61.5	72.8

*Notes:* See the text in Appendix IV for the details on sample selections.

heads, spouses, and children ages 7–50 in 1984.<sup>41</sup> In the BHPS, we select heads, spouses, and children ages 23–50 in 1991. In HILDA, we select heads, spouses, and children age 24–50 in 2001. In the KLIPS, we select heads, spouses, and children age 22–50 in 1998. In the SHP, we select heads, spouses, and children age 21–50 in 1999.

The non-attritors are individuals from the initial selection who became age twenty-five by year 1999 in the PSID, by 2002 in the GSOEP, by 1993 in BHPS, 2001 in KLIPS, and by 2002 in HILDA and SHP and are consecutively observed as heads or spouses for at least fifteen years by age sixty-five.

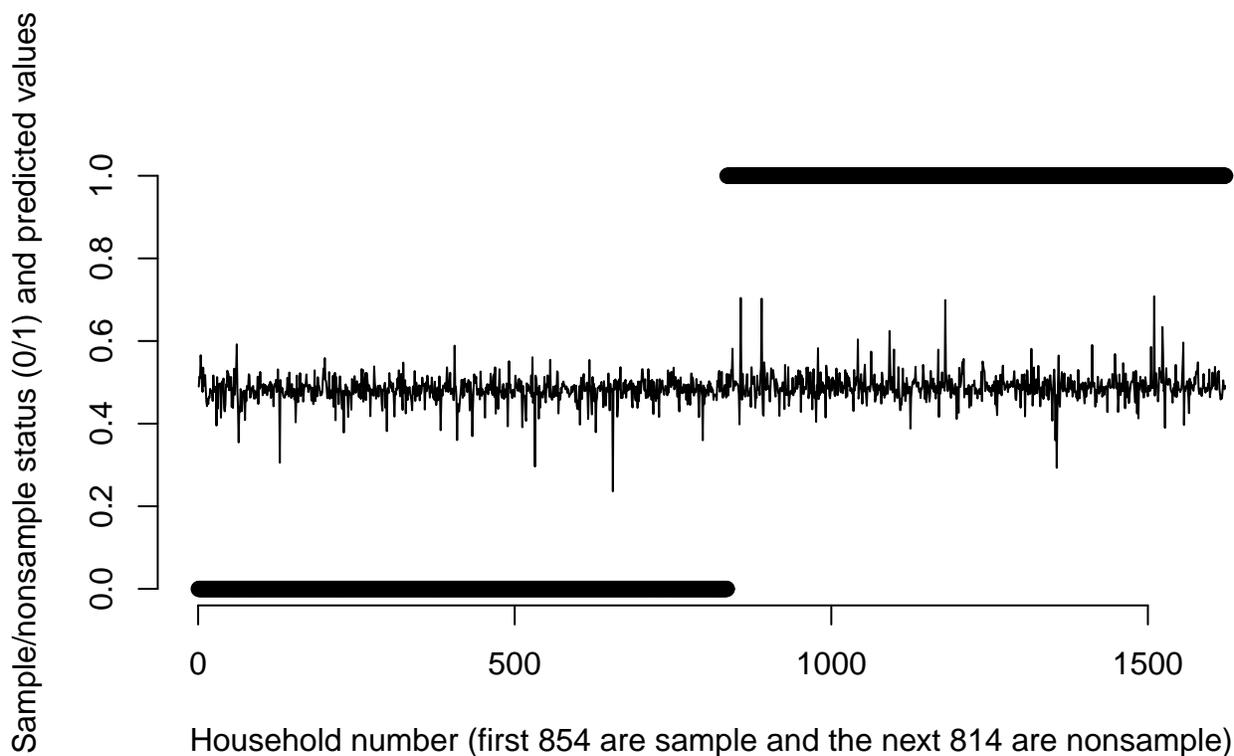
Although these selections are relevant for understanding the extent of attrition in our estimation samples in Table 14, they do not yield fully comparable attrition rates across the datasets due to a selection of different age cohorts in the first survey year, with a broader cohort coverage for longer-running datasets. In Table A-2 we show such comparable attrition rates, for sample individuals ages twenty-four to fifty in the first year of each dataset who are potentially observed for at least fifteen years as heads or spouses by age sixty-five. Clearly, the PSID shows the lowest attrition rates both for males and females. A comparison of Tables 15 and A-2 also reveals that in the GSOEP men attrit relatively more than women in the second generation.

***Persistence of income shocks for attritors and non-attritors.*** For comparability across the datasets, we first select original married couples, with a male spouse aged twenty-five or older, that we can follow for at least fifteen years by the time the male spouse reaches age sixty-five. We further limit our samples to the households observed consecutively in each dataset, with both spouses present for at least twelve years. Attritors are defined as the families with any of the spouses attriting from each respective dataset in the next three years. We drop income growth outliers and use residuals from a first-stage regression of gross family income on the same controls as in Section 3. We assume the same income process – a sum of an AR(1) and MA(1) processes – in each dataset, both for attritors and non-attritors, and report the persistence of permanent shocks estimated by GMM in Table 17.

<sup>41</sup>Those who are age seven in 1984 turn twenty-five in 2002 and we can potentially follow them for fifteen more years until 2017, the last survey year.

## V Additional Figures and Tables

FIGURE A-1: LASSO PREDICTIONS OF THE SAMPLE/NONSAMPLE STATUS.



*Notes:* The following regressors and their squared values (90 in total) are used for predicting the sample/nonsample status: annual hours worked of head and wife, labor income of head and wife, net family income, family size, number of children, combined transfer income of head and wife, transfer income of other family members, house value, head's year of birth, wife's year of birth, help from others, net assets; employment, unemployment and temporary layoff dummies, dummies for disability, residence in a large city, outside dependents, presence of extra earners other than head and wife; net family income and nondurable consumption residuals, food at home, and away from home, food stamps; dummies for a working wife, displacement, business ownership, college and non-college degrees, region of growing up, homeownership; family weight, number of observations on net family incomes and (imputed) nondurable consumption, race dummies, nondurable consumption, standard deviation of family nondurable consumption, and standard deviation of net family income. All variables that vary over time are averaged.

TABLE A-3: MEANS OF SELECTED VARIABLES FOR VARIOUS PSID SAMPLES

	Sample orig. (1)	Sample sons (2)	Nonsample (3)	p-value test: (2)=(3)
<i>A. Demographics</i>				
Head's age	51.149	38.54	38.834	43%
Wife's age	48.775	35.758	35.431	37%
Number of children	0.701	1.544	1.535	87%
Family size	3.181	3.667	3.653	82%
White	0.912	0.936	0.916	16%
Black	0.06	0.044	0.057	29%
Midwest	0.31	0.306	0.306	99%
South	0.309	0.316	0.296	43%
West	0.148	0.191	0.192	94%
MSA: largest city more than 100,000	0.397	0.397	0.423	28%
Years of education, avg.	12.762	13.861	13.722	16%
Region grew: foreign country	0.035	0.012	0.013	41%
<i>B. Income flows, consumption</i>				
Nondurable consumption	44015	24734	25650	43%
Net family income	41845	40389	40498	93%
Head's earnings	27547	27952	28038	94%
Wife's earnings	7074	10254	10042	68%
Transfers, family	2746	1449	1560	52%
Amount of help from relatives	22.004	74.48	98.252	23%
If income of other members>0	0.47	0.236	0.227	51%
If provided money support to others	0.166	0.161	0.167	72%
Number of tax exemptions, head and wife	3.066	3.701	3.668	57%
<i>C. Assets</i>				
Assets	132032	89005	90656	83%
If family owns business	0.198	0.212	0.235	22%
If family owns house	0.923	0.823	0.810	43%
House value/avg. income	3.301	2.675	2.566	31%
Mortgage/house value	0.175	0.509	0.513	86%
<i>D. Labor market</i>				
If head employed	0.814	0.921	0.917	67%
If head unemployed	0.019	0.024	0.03	12%
Head's hours	1933	2162	2181	55%
Wife's hours	949	1205	1200	90%
Head changed occupation	0.331	0.352	0.34	44%
Head changed industry	0.285	0.301	0.308	67%
<i>E. Survey info</i>				
Food at home, minor assignment	0.004	0.003	0.003	75%
Food at home, major assignment	0.005	0.002	0.003	73%
Food away, minor assignment	0.003	0.001	0.002	42%
Food away, major assignment	0.006	0.002	0.004	24%
Percent total family income, major assign.	5.33	2.629	2.759	82%
Percent total family income, minor assign.	6.303	3.323	3.32	99%
If respondent head	0.771	0.831	0.584	0%
If respondent wife	0.227	0.167	0.414	0%
Number of obs. in income spell	11.637	8.097	7.969	27%

TABLE A-4: NET FAMILY INCOME RESIDUALS

	Sample		Nonsample	
	Means	Vars.	Means	Vars.
	(1)	(2)	(3)	(4)
Year observed: first, year = 1978	0.00 (0.33)	-0.08*** (-8.06)	-0.02 (-0.91)	-0.05*** (-3.42)
Year observed: first, year $\neq$ 1978	0.00 (0.09)	0.02 (1.12)	-0.04** (-2.24)	0.03* (1.89)
Year observed: last, year = 1992	-0.02 (-1.12)	0.07*** (3.42)	0.02 (1.07)	0.00 (0.23)
Year observed: last, year $\neq$ 1992	-0.02 (-0.94)	0.05* (1.89)	-0.04 (-1.43)	0.04** (2.36)
1 year before inc. miss.	0.04 (0.23)	0.63*** (2.80)	0.26 (0.93)	0.32 (1.32)
1 year after inc. miss.	-0.13 (-0.72)	0.96*** (2.73)	0.44** (2.25)	0.05 (0.39)
Constant	-0.00 (-0.31)	0.20*** (23.73)	0.01 (0.93)	0.18*** (18.19)
Adj. R sq.	0.000	0.018	0.002	0.002
No. obs.	14672	14672	6483	6483
No. indiv.	1523	1523	814	814

*Notes:* Income data span the period 1979–1993. Income recorded in year  $t$  reflects income received in year  $t - 1$ . In columns (1) and (3), the dependent variables are residual incomes; in columns (2) and (4), the dependent variables are squared residual incomes. The dummies “Year observed: first, year = 1978” (“Year observed: first, year  $\neq$  1978”) is equal to one if an individual’s first income record is in 1979 (after 1979), and is equal to zero otherwise; “Year observed: last, year = 1992” (Year observed: last, year  $\neq$  1992) is equal to one if an individual’s last income record is in 1993 (before 1993), and is equal to zero otherwise. Standard errors are clustered by individual; t-statistics are in parentheses. \*\*\* significant at the 1% level, \*\* significant at the 5% level, \* significant at the 10% level.