Consumption Partial Insurance in the Presence of Tail Income Risk

Ghosh, Anisha; Theloudis, Alexandros

Publication date:
2023

Document Version
Early version, also known as pre-print

CentER Discussion Paper

No. 2023-024

CONSUMPTION PARTIAL INSURANCE IN THE PRESENCE OF TAIL INCOME RISK

By

Anisha Ghosh, Alexandros Theloudis

14 September 2023

ISSN 0924-7815
ISSN 2213-9532

Tilburg University

Consumption Partial Insurance in the Presence of Tail Income Risk*

Anisha Ghosh† Alexandros Theloudis‡

September 12, 2023

Abstract

We measure the extent of consumption insurance to income shocks accounting for high-order moments of the income distribution. We derive a nonlinear consumption function, in which the extent of insurance varies with the sign and magnitude of income shocks. Using PSID data, we estimate an asymmetric pass-through of bad versus good permanent shocks – 17% of a 3σ negative shock transmits to consumption compared to 9% of an equal-sized positive shock – and the pass-through increases as the shock worsens. Our results are consistent with surveys of consumption responses to hypothetical events and suggest that tail income risk matters substantially for consumption.

Keywords: Income risk, skewness, kurtosis, consumption, partial insurance, PSID.
JEL Classification: D12, D15, D31, E21

*We thank Manuel Arellano, Richard Blundell, Jeff Campbell, George Constantinides, Mariacristina De Nardi, Dmytro Hryshko, Christian Julliard, Chris Telmer, and seminar participants in the Structural Econometrics Group at Tilburg University and the BSE Summer Forum 2023 (Macroeconomics and Social Insurance) for many helpful comments. We also thank Ties Schalij for excellent research assistance.

†Desautels Faculty of Management, McGill University; email: anisha.ghosh@mcgill.ca.

‡Department of Econometrics & Operations Research, Tilburg University; email: a.theloudis@gmail.com.

1 Introduction

How does consumption respond to income shocks? The answer to this question is important for determining the welfare effects of shifts in the income distribution, i.e., the link between income and consumption inequality, for understanding how families cope after adverse events, for designing government insurance, and for the dynamics of business cycles. A consistent empirical finding is that consumption is partially insured against income shocks, so income fluctuations translate into consumption fluctuations less than one-to-one (e.g., Blundell, Pistaferri, and Preston, 2008; Heathcote, Storesletten, and Violante, 2014).[1] With a few exceptions discussed subsequently, most measures of consumption insurance do not distinguish between good and bad shocks, or between small and large ones. However, recent evidence in multiple countries shows that the distribution of income risk has a *long and fat left* tail; that is, there is an asymmetrically large probability that households face some really bad income shocks. This paper measures the consumption response to income shocks precisely when the distribution of said shocks may exhibit long and fat tails.

In a seminal paper, Blundell, Pistaferri, and Preston (2008) –henceforth, BPP– introduce a methodology to measure the pass-through of income shocks of varying persistence into consumption. Using Panel Study of Income Dynamics (PSID) and Consumer Expenditure Survey (CEX) data over 1980–1992, BPP find evidence of partial insurance to permanent shocks and full insurance to transitory ones. BPP rely on a workhorse covariance between income and consumption growth, namely a *second* moment of their joint distribution. The idea is that the extent to which income growth (driven by income shocks) varies with consumption growth reflects the strength of the transmission of income shocks into consumption.

In a separate strand of literature, Guvenen, Karahan, Ozkan, and Song (2021) use data from the U.S. Social Security Administration to characterize the distribution of unexplained income growth, namely the distribution of income shocks. They establish two important facts. First, income growth is very negatively skewed, i.e., it has a long left tail. This means there are far more people in the data who experience large negative than large positive shocks. Second, income growth exhibits excess kurtosis, i.e., it has fat tails. This means there are far more people in the data who experience either small or large income shocks, than people who experience moderate ones, relative to a Gaussian density. These are now benchmark results in the income dynamics literature, replicated in many countries (e.g., De Nardi, Fella, Knoef, Paz-Pardo, and Van Ooijen, 2021, compare the U.S. and the Netherlands).

[1]The first empirical studies of consumption insurance date back to tests of the permanent income hypothesis (e.g., Deaton and Paxson, 1994) or tests (and rejections) of complete markets (e.g., Cochrane, 1991; Attanasio and Davis, 1996; Hayashi, Altonji, and Kotlikoff, 1996; Brav, Constantinides, and Geczy, 2002). Meghir and Pistaferri (2011) offer a review of this literature.

If most shocks away from the mean are really bad shocks, their impact on consumption would, intuitively, differ from that of small/moderate shocks. In other words, left skewness and excess kurtosis, which are statements about the third and fourth moments of the income distribution, should matter for the pass-through of shocks into consumption. Yet, by revolving around the second moments of the joint income-consumption distribution, the second moment being a poor summary of the tails, the benchmark methodology in BPP ignores these salient features of income dynamics. This is precisely what we address in this paper.

The paper measures the degree of consumption partial insurance to income shocks of varying persistence, accounting for higher-order moments of the distribution of income and consumption. We offer three main contributions.

First, we replicate BPP (1980–1992) using recent PSID data (1999–2019). The replication targets the original income and consumption second moments. However, while BPP imputed consumption from the CEX (the PSID previously lacked comprehensive consumption data), such data are now internally available in the PSID. The replication is thus interesting not only because of the different time period, but also because we no longer need to rely on imputed consumption from the CEX.

Second, we introduce into BPP the third and fourth moments of the joint distribution of income and consumption, in addition to the second moments that BPP originally used. We thus go beyond the workhorse income-consumption covariance to measure consumption insurance targeting additional information available in the tails of the distribution.

Third, we generalize the method of measuring consumption partial insurance to explicitly permit a role for tail income risk. BPP log-linearize the optimality conditions and intertemporal budget constraint of a lifecycle consumption/savings problem to obtain a consumption function that depends *linearly* on income shocks. We argue (see also Carroll, 2001) that this linear function is a poor approximation to behavior when income is subject to tail shocks. We then derive a higher-order (quadratic) consumption function that nests BPP's linear specification. We characterize the consumption pass-through of income shocks in the generalized model and we establish identification of its parameters. In this environment, the degree of partial insurance depends on the magnitude and sign of shocks, so bad or extreme shocks can potentially shift consumption differently than good or small shocks do.

Our baseline empirical implementation uses income and consumption panel data from the PSID over the years 1999–2019. We focus on a representative sample of stable married couples that is similar to the sample in BPP. We show that unexplained income growth in this sample features large negative skewness and excess kurtosis, qualitatively and numerically close to the skewness and kurtosis in the administrative data of Guvenen, Karahan, Ozkan, and Song (2021). We then obtain four main findings.

First, permanent and transitory shocks to disposable household income are negatively skewed, with skewness coefficients at −0.95 and −1.31 respectively, and highly leptokurtic, with kurtosis coefficients at 44.39 and 49.23 respectively. Left skewness implies that a negative shock is typically more unsettling, more hurtful than a positive shock, because it is further away from the zero mean. Excess kurtosis implies that more shocks are concentrated either in the middle of the distribution or far out in the tails, relative to the shoulders. Taken together, large left skewness and excess kurtosis suggest that the realizations of income shocks tend to be either close to their zero mean or far out in the left tail; that is, income shocks for most households are either small/average or quite bad.

Second, in replicating BPP over the recent years, we find full insurance to transitory shocks and a very small pass-through rate of permanent shocks at 0.15 versus 0.64 in BPP. A 10% permanent income cut thus reduces consumption by only 1.5%, vis-à-vis 6.4% in BPP. In other words, we find much larger partial insurance to permanent shocks than originally found in BPP.[2] We attribute one third of the difference from BPP to the consumption imputation from the CEX that BPP use. The imputation brings in large amounts of measurement error, inflating the covariance between contemporary income and consumption growth. Another third of the difference arises because the modern PSID provides data biennially (versus annually previously), which tends to mute higher frequency shocks and their transmission into consumption. The last third reflects the different calendar times.

Third, introducing the third and fourth moments of income and consumption growth into BPP's linear consumption setting does not change our estimates of partial insurance. This result is counterintuitive because a big literature highlights the large welfare implications of tail risk (e.g., De Nardi, Fella, and Paz-Pardo, 2019; Busch and Ludwig, 2021). One would thus expect that targeting higher-order moments would affect the degree of partial insurance, which is not what we find. This result is suggestive of a potential misspecification in the linear consumption function, which motivates our generalization of the methodology through the introduction of a higher-order consumption specification.

Fourth, in estimating our quadratic consumption function, we find an asymmetry in the pass-through of bad versus good shocks, and of small versus large ones. With a quadratic consumption function, the pass-through depends on the magnitude and sign of the shocks; and targeting higher-order income moments becomes essential for the identification of the generalized transmission parameters. We draw three main conclusions. First, bad *permanent* shocks are more hurtful than good ones because they have larger transmission rates into consumption – a three standard deviations negative shock (a 50% permanent income

[2]Using the recent PSID and an alternative model, Arellano, Blundell, Bonhomme, and Light (2023) find a similarly low pass-through rate of permanent shocks, precisely as we do.

cut) has a pass-through of 0.17 (17% of the shock transmits into consumption), whereas a positive shock of similar magnitude and persistence has a pass-through of 0.09. Second, the pass-through of bad shocks increases with the severity of the shock. So, not only are bad permanent shocks more extreme than good ones (they lie further away from the zero mean in the left tail), they also have increasingly larger transmission rates as their magnitude worsens. Third, *transitory* shocks provide mixed evidence. They are fully insured in the full sample, at least from a statistical point of view, even though the point estimates of their transmission parameters suggest that large transitory windfalls shift consumption considerably, consistent with quasi-experimental evidence of large responses to transitory gains (e.g., Fagereng, Holm, and Natvik, 2021). Among low wealth households, however, transitory shocks shift consumption statistically significantly, with negative shocks passing through at an increasingly larger rate than positive ones. Overall, the results are consistent with an environment in which households save a good income shock but they dissave in the presence of a bad one. As there is clearly a limit to how much one can dissave or borrow, a large negative shock would typically affect consumption by relatively more.

The asymmetric transmission of bad versus good permanent shocks intensifies with the age of the household. While older households can typically insure average permanent shocks better than younger households, the former are associated with larger transmission rates of extreme shocks than the latter. In other words, severely bad permanent shocks cause more harm among older households than among younger ones. This agrees with Guvenen, Karahan, Ozkan, and Song (2021), who show that income skewness becomes more negative with the age of the earner – there is more room for income to fall as older workers have higher incomes; we find a similar pattern in our sample too. We further document that there is also more room for consumption to fall, so bad income draws are doubly hurtful among the elderly. Overall, the results from our quadratic specification are indicative of limited insurance to left tail income risk. Moreover, if the econometrician assumed a linear consumption function, she would overestimate the true extent of insurance to it, i.e., she would infer excessive insurance.

The paper is primarily related to the extensive literature that studies the transmission of income shocks into consumption, i.e., the link between income and consumption inequality. There are two distinct strands in this literature: a quasi-reduced form branch that *measures* the degree of partial insurance in the data without fully specifying the underlying insurance mechanisms, and a structural branch that fully *models* specific insurance channels.

The seminal paper in the first branch, namely BPP, sparked a large volume of subsequent research.[3] Blundell, Pistaferri, and Saporta-Eksten (2016) extend BPP to an environment

[3] Hall and Mishkin (1982), Deaton and Paxson (1994) and Blundell and Preston (1998) characterize the

with wage risk to measure the contribution of family labor supply to partial insurance. Theloudis (2021) extends this further to allow for household heterogeneity, which he estimates using information in the third income-consumption moments. Commault (2022) allows consumption to respond to past transitory shocks and finds a strong response to them, similar to what quasi-experimental studies find. Hryshko and Manovskii (2022) carry out an interesting subsample analysis of BPP, identifying sets of households in the PSID with different degrees of insurance based on certain features of their incomes. Crawley and Kuchler (2023) offer another empirical extension, addressing neglected time aggregation in income and consumption in the PSID. Arellano, Blundell, and Bonhomme (2017) and Arellano, Blundell, Bonhomme, and Light (2023) generalize the income process by letting shocks of different size/sign to feature varying persistence.[4] Abstracting from income skewness and kurtosis, they measure the consumption response to income in the modern PSID at about 0.2 on average (specification with household heterogeneity in the latter paper), though the response varies over the distribution of shocks.

In the present paper, by contrast, we generalize both the income process, by allowing shocks to be non-Gaussian, and the consumption function, by allowing for non-linearities in the transmission of income shocks. The value added is our finding of asymmetric pass-through of shocks depending on their size and sign – a recurrent finding in recent surveys of responses to hypothetical scenarios (e.g., Fuster, Kaplan, and Zafar, 2020). To measure this asymmetry, information in higher-order income moments becomes indispensable. We thus characterize carefully how our measures of partial insurance change and what can be learned as we gradually introduce higher-order information.

The second branch studies specific insurance mechanisms through calibrating lifecycle models of consumption/savings.[5] Much of the early research focused on self-insurance, providing mixed evidence for how much insurance there is. Kaplan and Violante (2010) find less insurance to permanent shocks in the model compared to the data, while Guvenen and Smith (2014) estimate income risk and the degree of self-insurance jointly and find that about one half of income shocks are smoothed, i.e., they find more insurance than BPP. More recent works study the insurance consequences of labor supply (Heathcote, Storesletten, and Violante, 2014), durable purchases (Madera, 2019), health (Blundell, Borella, Commault, and Nardi, 2020), or added workers (Wu and Krueger, 2021). Except Madera (2019), these papers abstract from higher-order moments of income, partly due to the computational difficulties in modelling those. De Nardi, Fella, and Paz-Pardo (2019) embed a non-Gaussian

empirical distribution of consumption and income, so they are early precursors in this literature.

[4]The latter paper advances the former by allowing for flexible heterogeneity and unbalanced panels.

[5]Gourinchas and Parker (2002) and Krueger and Perri (2006) are early precursors in this strand.

income process into a model of consumption/savings and show that tail income risk increases the degree of partial insurance to permanent shocks due to stronger precautionary motives. Busch and Ludwig (2021) estimate a similar model of self-insurance targeting income skewness and kurtosis. They confirm that higher-order moments generate stronger precautionary savings; they distinguish, however, between good and bad shocks and they show that the latter are in fact worse insured than the former. Unlike this vast quantitative literature, we do not take a stance on the insurance mechanisms available to households; this enables us to *measure* the overall degree of insurance from all mechanisms together, keeping our model partially unspecified while also accounting for the higher-order income moments.

The paper is also related to the modern income dynamics literature that investigates the non-Gaussian features of the income distribution. Geweke and Keane (2000) document that earnings growth exhibits left skewness, while Bonhomme and Robin (2010) document its excess kurtosis, both using PSID data. Guvenen, Karahan, Ozkan, and Song (2021) provide evidence of left skewness and excess kurtosis in earnings growth in the population of working males in the U.S., and show that these features get accentuated with age and the level of earnings. Guvenen, Ozkan, and Song (2014) and Busch, Domeij, Guvenen, and Madera (2018) study the cyclical nature of income skewness in several countries. While these papers are mostly descriptive of the higher-order features of the income distribution, we go a step further to investigate what these features imply for consumption partial insurance. We thus contribute to the growing literature in economics and finance that studies the various implications of tail risk (e.g., Constantinides and Ghosh, 2017; McKay, 2017; De Nardi, Fella, and Paz-Pardo, 2019; Ai and Bhandari, 2021, for the equity premium and asset prices, business cycles, consumer welfare, and labor market dynamics, respectively).

The asymmetric response to bad versus good shocks has previously been seen in the literature that uses surveys to elicit responses to various hypothetical scenarios. For example, Bunn, Le Roux, Reinold, and Surico (2018) find that the marginal propensity to consume from bad shocks is larger than from good shocks; Christelis, Georgarakos, Jappelli, Pistaferri, and van Rooij (2019) and Fuster, Kaplan, and Zafar (2020) find that responses to losses are larger than to gains with magnitudes increasing with the severity of the loss.[6] We show that these patterns are also present in survey data, such as the PSID.

In the remainder of the paper, section 2 presents the theoretical framework through which we study the pass-through of income shocks into consumption and 3 establishes identification of the relevant parameters. Section 4 discusses the empirical implementation, section 5 presents the results, and section 6 concludes.

[6]Given our finding of heterogeneous responses to shocks, our paper is also broadly related to the literature that studies MPC heterogeneity (e.g., Crawley and Kuchler, 2023; Lewis, Melcangi, and Pilossoph, 2022).

2 Lifecycle model with tail income risk

The framework through which we study the extent of partial insurance to income shocks is a consumption lifecycle model with idiosyncratic income risk.

Household i chooses lifetime consumption $\{C_{it}\}_{t=0}^{T}$ and savings $\{A_{it+1}\}_{t=0}^{T}$ to maximize its expected discounted utility over its lifecycle. The problem is formulated as

$$\max_{\{C_{it},A_{it+1}\}_{t=0}^{T}} \mathbb{E}_0 \sum_{t=0}^{T} \beta^t U(C_{it}; \mathbf{Z}_{it}), \tag{1}$$

subject to the lifetime budget constraint

$$A_{i0} + \mathbb{E}_0 \sum_{t=0}^{T} \frac{Y_{it}}{(1+r)^t} = \mathbb{E}_0 \sum_{t=0}^{T} \frac{C_{it}}{(1+r)^t}, \tag{2}$$

where $U(\cdot; \mathbf{Z}_{it})$ is utility over consumption, $\mathbf{Z}_{it}$ is a vector of observed taste shifters (e.g., age, education), β is the discount factor, r is the deterministic interest rate, A_{i0} is beginning-of-time financial wealth, $\mathbb{E}_0$ denotes expectations over uncertain future income Y, and T is the end of the horizon. We assume that utility satisfies standard regularity conditions, in particular it is twice differentiable, but we do not otherwise parameterize it.

2.1 Income process

Income is the only source of (idiosyncratic) uncertainty that the household faces. Its logarithm is given by

$$\ln Y_{it} = \mathbf{X}_{it}'\boldsymbol{\delta} + P_{it} + v_{it},$$

where $\mathbf{X}_{it}$ is a vector of characteristics in the $t = 0$ information set of the household, such as age, education, race, that drive the deterministic profile of income over the lifecycle; $\boldsymbol{\delta}$ is the loading factor of those variables on the logarithm of income. The remaining terms make up the stochastic component of income, consisting of permanent income P_{it} and transitory shock v_{it}. We assume that permanent income has a unit root, $P_{it} = P_{it-1} + \zeta_{it}$, where ζ_{it} is the permanent income shock. These modeling choices give rise to the familiar permanent-transitory formulation[7]

$$\Delta y_{it} = \zeta_{it} + \Delta v_{it}, \tag{3}$$

where $\Delta y_{it} = \Delta \ln Y_{it} - \Delta \mathbf{X}_{it}'\boldsymbol{\delta}$ is income growth net of the deterministic profile, i.e., unexplained income growth, and Δ is the first difference operator.

[7]See Meghir and Pistaferri (2011) for a comprehensive review of modern income processes.

A first point of departure from BPP lies in the specification of the permanent and transitory income shocks, ζ_{it} and v_{it}, respectively. Unlike BPP, we allow these shocks to be non-Gaussian, with first four central moments given by

$$\mathbb{E}(\zeta_{it}^m) = \begin{cases} 0 & \text{for } m = 1 \\ \sigma_{\zeta_t}^2 & \text{for } m = 2 \\ \gamma_{\zeta_t} & \text{for } m = 3 \\ \kappa_{\zeta_t} & \text{for } m = 4 \end{cases} \quad \text{and} \quad \mathbb{E}(v_{it}^m) = \begin{cases} 0 & \text{for } m = 1 \\ \sigma_{v_t}^2 & \text{for } m = 2 \\ \gamma_{v_t} & \text{for } m = 3 \\ \kappa_{v_t} & \text{for } m = 4. \end{cases}$$

$\sigma_{\zeta_t}^2$ and $\sigma_{v_t}^2$ describe the dispersion of the distributions of, respectively, the permanent and transitory shocks about their mean; the mean is zero by construction. γ_{ζ_t} and γ_{v_t} denote the skewness/asymmetry of the distributions, i.e., the relative lengths of the upper and lower tails. Finally, κ_{ζ_t} and κ_{v_t} denote the kurtosis, i.e., the tendency of the distributions to amass away from the middle $[-\sigma, \sigma]$, and therefore characterize the thickness of the tails.

We allow the central moments of shocks to depend on t to reflect that different stages of the lifecycle may exhibit different amounts of dispersion, skewness, or thickness of the tails of income risk; Guvenen, Karahan, Ozkan, and Song (2021) provide some evidence for that. Otherwise, shocks across households are identically distributed within t. Permanent and transitory shocks are mutually independent and independent over time.

As discussed in the introduction, several recent studies document that the cross-sectional distribution of income growth exhibits left skewness and excess kurtosis, i.e., a long and fat left tail. Our specification of the moments of income shocks enables our income process to be consistent with these facts, allowing us to subsequently decompose its second- and higher-order moments into the relative contributions of permanent and transitory shocks.

2.2 Linear consumption function

In line with BPP, we first consider a linear specification of the household consumption function, given by

$$\Delta c_{it} = \xi_{it} + \phi_{it}^{(1)} \zeta_{it} + \psi_{it}^{(1)} v_{it}, \tag{4}$$

where Δc_{it} is the residual from a regression of real consumption growth, $\Delta \ln C_{it}$, on the taste observables $\mathbf{Z}_{it}$ (i.e., consumption growth net of its predictable profile). The linear function (4) stems from a log-linearization (i.e., a first-order Taylor series expansion) of the problem's first-order conditions and the household budget constraint, as we show in appendix A. These approximations, which are independent of the distributions of income shocks, mimic similar derivations in BPP.

Permanent and transitory income shocks drive unexplained consumption growth in (4). The transmission parameters $\phi_{it}^{(1)} = \partial \Delta c_{it} / \partial \zeta_{it}$ and $\psi_{it}^{(1)} = \partial \Delta c_{it} / \partial v_{it}$ reflect the pass-through of permanent and transitory shocks, respectively, into consumption. Adopting BPP's terminology, $1 - \phi_{it}^{(1)}$ is then the extent of insurance to permanent shocks (i.e., the fraction of a permanent shock that does *not* translate into a consumption fluctuation), while $1 - \psi_{it}^{(1)}$ is the extent of insurance to transitory shocks. Therefore, $\phi_{it}^{(1)} = \psi_{it}^{(1)} = 0$ implies that consumption is *fully* insured to income shocks, while $\phi_{it}^{(1)} = \psi_{it}^{(1)} = 1$ implies *no insurance*. By contrast, values $\phi_{it}^{(1)}, \psi_{it}^{(1)} \in (0, 1)$ reflect *partial* insurance: a given income shock translates into a consumption fluctuation, albeit less than one-to-one.

The transmission parameters vary, in principle, with i (household) and t (age), reflecting heterogeneity in financial wealth, among other reasons. The idea is that a household can use wealth to smooth consumption when income shocks hit; and the more wealth one holds the better they can self-insure.[8] We establish this link from wealth to $\phi_{it}^{(1)}$, $\psi_{it}^{(1)}$ in appendix A. So while the pass-through of shocks depends on wealth, it does not depend on the magnitude or sign of the shock itself. Function (4) thus implies that a household can insure big or small, good or bad shocks similarly, a point to which we return subsequently.

In addition to income shocks, unexplained consumption growth is also driven by ξ_{it}, which reflects unobserved (but non-stochastic) consumption taste heterogeneity across households, independent of the income shocks. We let the distribution of ξ_{it} across households to have zero mean and variance $\sigma_{\xi_t}^2$, as in BPP. We also let its third and fourth central moments to be γ_{ξ_t} and κ_{ξ_t} respectively; these will appear subsequently in expressions for the skewness and kurtosis of consumption growth.

2.3 Quadratic consumption function

The linear consumption function offers an approximation to the true consumption response to income shocks when such shocks do not move the household too far away from its consumption 'steady-state' – i.e., its deterministic consumption profile. Function (4) relies on a linearization of marginal utility (the problem's first-order conditions) around such 'steady-state'; but such linearization will be inaccurate in the presence of tail income risk, e.g., when extreme shocks may induce large shifts in consumption as a consequence of concave utility.[9]

Our second and main point of departure from BPP is the generalization of the consump-

[8] In addition, households in Blundell, Pistaferri, and Saporta-Eksten (2016) have access to insurance through family labor supply, while households in Theloudis (2021) differ in their preferences; both are reasons why the transmission parameters may vary with t or i.

[9] This echoes Carroll (2001)'s critique of the log-linearized Euler equation. Kaplan and Violante (2010) and Blundell, Low, and Preston (2013) assess the approximations in BPP across alternative scenarios, though not in one with respect to the severity of income shocks.

tion function to a nonlinear (quadratic) specification, given by

$$\Delta c_{it} = \xi_{it} + \phi^{(1)}_{it}\zeta_{it} + \psi^{(1)}_{it}v_{it} + \phi^{(2)}_{it}\zeta^2_{it} + \psi^{(2)}_{it}v^2_{it} + \omega^{(22)}_{it}\zeta_{it}v_{it}, \tag{5}$$

where Δc_{it} and ξ_{it} are defined as previously. The quadratic function (5) stems from a second-order Taylor approximation to the problem's first-order conditions and the household budget constraint, as we show in appendix A. These operations and resulting expressions generalize the original approximations in BPP and in several papers thereafter.

The quadratic function is characterized by several parameters. As in the linear case, $\phi^{(1)}_{it}$ and $\psi^{(1)}_{it}$ are the transmission parameters into consumption of the linear part of a polynomial in the income shocks. By extension, $\phi^{(2)}_{it}$, $\psi^{(2)}_{it}$, and $\omega^{(22)}_{it}$ are the transmission parameters of the quadratic part of the same polynomial. In the context of long left and fat tails in income risk, $\phi^{(1)}_{it}$ and $\psi^{(1)}_{it}$ can be thought of as the transmission parameters of the average shock (a shock close to the mean of the distribution), while $\phi^{(2)}_{it}$, $\psi^{(2)}_{it}$, $\omega^{(22)}_{it}$ can be seen as the transmission parameters of the typical shock away from the middle of the distribution. The parameters may again vary with i (household) and t (age), reflecting heterogeneity in financial wealth among other things; we establish this link in appendix A.

In this generalized framework, $\phi^{(1)}_{it}$ no longer represents the pass-through of permanent shocks to consumption. Instead, the pass-through is given by

$$\frac{\partial \Delta c_{it}}{\partial \zeta_{it}} = \phi^{(1)}_{it} + 2\phi^{(2)}_{it}\zeta_{it} + \omega^{(22)}_{it}v_{it}, \tag{6}$$

which depends on multiple parameters *and* on the magnitude and sign of the income shocks. As such, there is an asymmetric pass-through of good versus bad shocks into consumption. Suppose, for example, that $\phi^{(1)} > 0$ and $\phi^{(2)} < 0$, as our results subsequently suggest. At $v_{it} = 0$ (the average transitory shock), a negative *permanent* shock has a higher pass-through to consumption relative to a positive shock of the same magnitude. The pass-through of negative shocks would thus be underestimated (i.e., excessive insurance concluded) if the econometrician assessed the extent of partial insurance on the basis of $\phi^{(1)}$ alone, ignoring the non-linearity in the consumption response. The size of this bias depends on the magnitude of $\phi^{(2)}$ (a more negative $\phi^{(2)}$ implies bigger bias) and on the size of the permanent shock itself (a more severe bad shock implies bigger bias). When income risk is non-Gaussian (e.g., it is left tailed), the realization of shocks can be far from their mean, thereby causing large errors in the inference of the degree of partial insurance based on $\phi^{(1)}$ alone.

A similar asymmetry exists in the transmission of good versus bad transitory income shocks. In this case, the pass-through into consumption is given by

$$\frac{\partial \Delta c_{it}}{\partial v_{it}} = \psi_{it}^{(1)} + 2\psi_{it}^{(2)} v_{it} + \omega_{it}^{(22)} \zeta_{it}, \tag{7}$$

which again depends on multiple parameters *and* on the magnitude and sign of the income shocks. Importantly, this generalized pass-through does not rule out the possibility of full insurance to transitory shocks, a recurrent empirical finding in the literature. Suppose, for example, that $\psi^{(1)} \approx \psi^{(2)} \approx 0$; then $\partial \Delta c_{it}/\partial v_{it} \approx 0$ at $\zeta_{it} = 0$ (the average permanent shock), indicating full insurance to transitory shocks.

3 Identification of income and consumption parameters

Our specification of the income and consumption processes imposes restrictions on the second and higher-order moments of income and consumption that can be used to identify the various parameters. The set of parameters consists of the variance ($\sigma^2_{\zeta_t}$, $\sigma^2_{v_t}$), skewness (γ_{ζ_t}, γ_{v_t}), and kurtosis (κ_{ζ_t}, κ_{v_t}) of the distributions of permanent and transitory shocks,[10] the partial insurance parameters in the linear ($\phi_t^{(1)}$, $\psi_t^{(1)}$) and quadratic consumption functions ($\phi_t^{(1)}$, $\psi_t^{(1)}$, $\phi_t^{(2)}$, $\psi_t^{(2)}$, $\omega_t^{(22)}$), and the moments of the unobserved consumption taste heterogeneity ($\sigma^2_{\xi_t}$, γ_{ξ_t}, κ_{ξ_t}). For simplicity of the illustration, we restrict the partial insurance parameters to not vary across households; this is *not* an identifying restriction, and in fact we explore subsequently their variation across certain demographic groups.

Two important aspects of the data must be addressed before proceeding to identification: the biennial nature of the modern PSID and measurement error.

As we describe in section 4, the PSID provides data only every two years within our time period. This necessitates to recast the income and consumption growth processes of section 2 to this lower frequency, a straightforward operation that we carry out in appendix B. Nevertheless, for consistency with the earlier literature, we provide the identifying statements subsequently *as if* the data come at a yearly frequency. The complete identifying statements at the lower frequency of the modern PSID appear in appendix B.

Measurement error in income and consumption is a more challenging issue. In the permanent-transitory specification of income, the moments of classical income error are not separately identifiable from the moments of the economically relevant transitory shock. This necessitates that we restrict the income error to be Gaussian and retrieve its variance from

[10] With a slight abuse of terminology, we use the terms skewness and kurtosis to refer to the third and fourth central moments of shocks. We will be clear whenever those terms refer to *standardized* moments.

validation studies of the PSID, as in Meghir and Pistaferri (2004) and Blundell, Pistaferri, and Saporta-Eksten (2016). We return to this point in section 4. By contrast, the moments of classical consumption error are identifiable; this introduces three additional parameters for the variance, skewness, and kurtosis of consumption error ($\sigma^2_{u^c_t}$, $\gamma_{u^c_t}$, $\kappa_{u^c_t}$). For simplicity, however, we subsequently show identification *as if* income and consumption errors are not present; we provide our full identifying statements with measurement error in appendix B.

3.1 Income process parameters

Given the income process in (3) and the properties of shocks, the second and higher-order moments of the cross-sectional distribution of income growth are given by

$$\begin{aligned}
\mathbb{E}((\Delta y_{it})^2) &= \sigma^2_{\zeta_t} + \sigma^2_{v_t} + \sigma^2_{v_{t-1}}, \\
\mathbb{E}((\Delta y_{it})^3) &= \gamma_{\zeta_t} + \gamma_{v_t} - \gamma_{v_{t-1}}, \\
\mathbb{E}((\Delta y_{it})^4) &= \kappa_{\zeta_t} + \kappa_{v_t} + \kappa_{v_{t-1}} + 6\sigma^2_{\zeta_t}(\sigma^2_{v_t} + \sigma^2_{v_{t-1}}) + 6\sigma^2_{v_t}\sigma^2_{v_{t-1}}.
\end{aligned}$$

The second moment depends on the variance of permanent shocks at t and the sum of variances of transitory shocks at t and $t-1$; it thus reflects the dispersion of one permanent and two transitory shocks. The fourth moment reflects, similarly, contributions from the dispersion and kurtosis of both types of shocks. By contrast, the third moment is driven primarily by skewness in the permanent shock. Barring strong non stationarities in the distribution of transitory shocks, the difference between γ_{v_t} and $\gamma_{v_{t-1}}$ will generally be small (and precisely zero if the distributions are time-invariant), leaving γ_{ζ_t}, the skewness in permanent shocks, to drive entirely the third moment of income.

While the above expressions are already partly informative of the parameters of the income process, the variance of shocks is formally identified as

$$\begin{aligned}
\sigma^2_{\zeta_t} &= \mathbb{E}(\Delta y_{it} \times (\Delta y_{it-1} + \Delta y_{it} + \Delta y_{it+1})), \\
\sigma^2_{v_t} &= -\mathbb{E}(\Delta y_{it} \times \Delta y_{it+1}).
\end{aligned}$$

The long sum in the first line strips income growth at t of its contemporary transitory shock; its covariance with Δy_{it} thus picks up the variance of the permanent shock. The covariance of two consecutive income growths in the second line identifies the variance of the transitory shocks due to mean reversion in the shock.

In a similar way, the skewness of shocks is identified as[11]

$$
\begin{aligned}
\gamma_{\zeta_t} &= \mathbb{E}((\Delta y_{it})^2 \times (\Delta y_{it-1} + \Delta y_{it} + \Delta y_{it+1})), \\
\gamma_{v_t} &= -\mathbb{E}((\Delta y_{it})^2 \times \Delta y_{it+1}),
\end{aligned}
$$

while their kurtosis is identified from

$$
\begin{aligned}
\kappa_{\zeta_t} &= \mathbb{E}((\Delta y_{it})^4) - \kappa_{v_t} - \kappa_{v_{t-1}} - 6\sigma^2_{\zeta_t}(\sigma^2_{v_t} + \sigma^2_{v_{t-1}}) - 6\sigma^2_{v_t}\sigma^2_{v_{t-1}}, \\
\kappa_{v_t} &= \mathbb{E}((\Delta y_{it})^2 \times (\Delta y_{it+1})^2) - \sigma^2_{\zeta_t}\sigma^2_{\zeta_{t+1}} - \sigma^2_{\zeta_t}\sum_{\tau=0}^{1}\sigma^2_{v_{t+\tau}} - \sigma^2_{\zeta_{t+1}}\sum_{\tau=-1}^{0}\sigma^2_{v_{t+\tau}} \\
&\quad - \sigma^2_{v_{t-1}}\sigma^2_{v_t} - \sigma^2_{v_{t-1}}\sigma^2_{v_{t+1}} - \sigma^2_{v_t}\sigma^2_{v_{t+1}}.
\end{aligned}
$$

There are several over-identifying restrictions for both skewness and kurtosis.[12]

3.2 Linear consumption function parameters

Given the linear consumption function in (4) and the properties of shocks, the second and higher-order moments of the distribution of consumption growth are given by

$$
\begin{aligned}
\mathbb{E}((\Delta c_{it})^2) &= (\phi^{(1)}_t)^2\sigma^2_{\zeta_t} + (\psi^{(1)}_t)^2\sigma^2_{v_t} + \sigma^2_{\xi_t}, \\
\mathbb{E}((\Delta c_{it})^3) &= (\phi^{(1)}_t)^3\gamma_{\zeta_t} + (\psi^{(1)}_t)^3\gamma_{v_t} + \gamma_{\xi_t}, \\
\mathbb{E}((\Delta c_{it})^4) &= (\phi^{(1)}_t)^4\kappa_{\zeta_t} + (\psi^{(1)}_t)^4\kappa_{v_t} + \kappa_{\xi_t} + 6(\phi^{(1)}_t)^2(\psi^{(1)}_t)^2\sigma^2_{\zeta_t}\sigma^2_{v_t} \\
&\quad + 6(\phi^{(1)}_t)^2\sigma^2_{\zeta_t}\sigma^2_{\xi_t} + 6(\psi^{(1)}_t)^2\sigma^2_{v_t}\sigma^2_{\xi_t}.
\end{aligned}
$$

The moments of the distribution of consumption growth reflect the underlying distributions of income shocks, as well as the distribution of unobserved taste heterogeneity. In general, larger dispersion, skewness, or kurtosis in the distributions of income shocks implies larger dispersion, skewness, or kurtosis in the distribution of consumption growth. The link between the two depends on the degree of insurance with respect to these shocks, so these moments provide restrictions that help identify the parameters of partial insurance.

[11] In the presence of income measurement error, the moments that identify the variance and skewness of transitory shocks also pick up the variance and skewness of the error. Retrieving its variance from validation studies (see section 4) while assuming error is Guassian (thus has zero skewness) allows us to identify the parameters of the economically relevant transitory shock. Appendix B reports the details of identification in the presence of measurement error.

[12] Identification of the fourth moments requires prior identification of the second moments; identification of kurtosis of the permanent shock also requires prior identification of kurtosis of the transitory shock.

Formally, the transmission parameter of permanent shocks is identified as

$$\phi_t^{(1)} = \mathbb{E}(\Delta c_{it} \times (\Delta y_{it-1} + \Delta y_{it} + \Delta y_{it+1}))/\sigma_{\zeta_t}^2,$$

while the transmission parameter of transitory shocks is identified as

$$\psi_t^{(1)} = -\mathbb{E}(\Delta c_{it} \times \Delta y_{it+1})/\sigma_{v_t}^2;$$

in both cases, the variance of shocks is given by the earlier expressions.[13] These statements reflect the intuition behind identification in BPP: the *covariance* between consumption and (appropriately defined) income growth captures the strength of the relation between the two series and, consequently, the pass-through of income shocks into consumption.[14]

This workhorse covariance, however, is a *second* moment of the joint distribution of consumption and income and, as such, neglects information in the tails. If most income shocks are either zero or rather bad, as Guvenen, Karahan, Ozkan, and Song (2021) find, identifying partial insurance based on second moments alone may be misleading. However, the partial insurance parameters are heavily (over-)identified by higher-order moments of the marginal and joint distributions of income and consumption; inspect, for example, the expressions for $\mathbb{E}((\Delta c_{it})^3)$ or $\mathbb{E}((\Delta c_{it})^4)$. We will thus subsequently exploit these higher-order moments in the context of the linear consumption function in our first attempt to understand the implications of non-Gaussianity for partial insurance.

3.3 Quadratic consumption function parameters

As per the quadratic specification in (5) and the properties of shocks, the second moment of the distribution of consumption growth is given by

$$\begin{aligned}\mathbb{E}((\Delta c_{it})^2) = {} & (\phi_t^{(1)})^2\sigma_{\zeta_t}^2 + (\psi_t^{(1)})^2\sigma_{v_t}^2 + 2\phi_t^{(1)}\phi_t^{(2)}\gamma_{\zeta_t} + 2\psi_t^{(1)}\psi_t^{(2)}\gamma_{v_t} + (\phi_t^{(2)})^2\kappa_{\zeta_t} + (\psi_t^{(2)})^2\kappa_{v_t} \\ & + (\omega_t^{(22)})^2\sigma_{\zeta_t}^2\sigma_{v_t}^2 + 2\phi_t^{(2)}\psi_t^{(2)}\sigma_{\zeta_t}^2\sigma_{v_t}^2 + \sigma_{\xi_t}^2.\end{aligned}$$

This depends not only on the variance of permanent and transitory income shocks, as in the linear case, but also on their skewness and kurtosis. Contrary to the linear case, however, the higher-order moments of consumption growth depend on higher than fourth-order moments

[13] The variance of consumption taste heterogeneity is further given by $\sigma_{\xi_t}^2 = \mathbb{E}((\Delta c_{it})^2) - (\phi_t^{(1)})^2\sigma_{\zeta_t}^2 - (\psi_t^{(1)})^2\sigma_{v_t}^2$ while its skewness by $\gamma_{\xi_t} = \mathbb{E}((\Delta c_{it})^3) - (\phi_t^{(1)})^3\gamma_{\zeta_t} - (\psi_t^{(1)})^3\gamma_{v_t}$. The kurtosis κ_{ξ_t} is similarly identified from the fourth moment of consumption growth.

[14] An equivalent view is that the partial insurance parameters are identified from a regression of consumption growth on income growth, using permanent or future income as instruments.

of the shocks, which we do not model.

Identification of the partial insurance parameters comes from the joint moments of consumption and income, conceptually similar to identification in the linear case. Given the non-linearity, however, multiple moments contribute jointly to the identification of any single parameter, and the identifying statements are inevitably more involved. Formally, the transmission parameters are given by

$$\begin{pmatrix} \phi_t^{(1)} \\ \psi_t^{(1)} \\ \phi_t^{(2)} \\ \psi_t^{(2)} \\ \omega_t^{(22)} \end{pmatrix} = \mathbf{A}^{-1} \begin{pmatrix} \mathbb{E}\left(\Delta c_{it} \times \Delta y_{it}\right) \\ \mathbb{E}\left(\Delta c_{it} \times (\Delta y_{it})^2\right) \\ \mathbb{E}\left(\Delta c_{it} \times \Delta y_{it+1}\right) \\ \mathbb{E}\left(\Delta c_{it} \times (\Delta y_{it+1})^2\right) \\ \mathbb{E}\left(\Delta c_{it} \times \Delta y_{it} \times \Delta y_{it+1}\right) \end{pmatrix}, \tag{8}$$

where square matrix $\mathbf{A}$ depends exclusively on the second, third, and fourth moments of income shocks, and it is generally full rank and invertible.[15,16]

As shown in section 3.1, the higher-order moments of income shocks are identified through the third and fourth moments of income. Therefore, the higher than second-order moments of income play an essential role in enabling the identification of the transmission parameters of the quadratic consumption function. Unlike the linear case, in other words, the transmission parameters in the generalized specification cannot be identified with second-order income moments alone. By contrast, no higher-order moments of consumption are needed; the identifying moments in (8) are covariances between consumption growth and various powers of income growth. This is an advantage because, while in principle we may estimate higher than second-order moments of *consumption* growth in the PSID, we cannot be certain how good these estimates are. This is in contrast to higher-order *income* moments, for which Guvenen, Karahan, Ozkan, and Song (2021) serves as a natural benchmark given their use of administrative income data. We return to this point in section 4.

What is the bias in the transmission parameters if one assumes a linear specification

[15] The matrix of coefficients is given by

$$\mathbf{A} = \begin{pmatrix} \sigma^2_{\zeta_t} & \sigma^2_{v_t} & \gamma_{\zeta_t} & \gamma_{v_t} & 0 \\ \gamma_{\zeta_t} & \gamma_{v_t} & \kappa_{\zeta_t} + \sigma^2_{\zeta_t}(\sigma^2_{v_t} + \sigma^2_{v_{t-1}}) & \kappa_{v_t} + \sigma^2_{v_t}(\sigma^2_{\zeta_t} + \sigma^2_{v_{t-1}}) & 2\sigma^2_{\zeta_t}\sigma^2_{v_t} \\ 0 & -\sigma^2_{v_t} & 0 & -\gamma_{v_t} & 0 \\ 0 & \gamma_{v_t} & \sigma^2_{\zeta_t}(\sigma^2_{\zeta_{t+1}} + \sigma^2_{v_t} + \sigma^2_{v_{t+1}}) & \kappa_{v_t} + \sigma^2_{v_t}(\sigma^2_{\zeta_{t+1}} + \sigma^2_{v_{t+1}}) & 0 \\ 0 & -\gamma_{v_t} & -\sigma^2_{\zeta_t}\sigma^2_{v_t} & -\kappa_{v_t} & -\sigma^2_{\zeta_t}\sigma^2_{v_t} \end{pmatrix}.$$

There are no obvious linear dependencies across columns so the matrix is generally full rank and invertible. This is true also if the underlying distributions of shocks are Gaussian.

[16] The variance of consumption taste heterogeneity is identified from $\mathbb{E}((\Delta c_{it})^2)$. Its skewness and kurtosis enter the third and fourth moments of consumption growth. Those depend on higher than fourth-order moments of income, which we do not model, so they are generally not identified in the current setting.

when the true consumption function is quadratic? The identifying condition for $\phi_t^{(1)}$ in the linear case identifies

$$\mathbb{E}(\Delta c_{it} \times (\Delta y_{it-1} + \Delta y_{it} + \Delta y_{it+1}))/\sigma_{\zeta_t}^2 = \phi_t^{(1)} + \phi_t^{(2)}(\gamma_{\zeta_t}/\sigma_{\zeta_t}^2),$$

namely a linear combination of $\phi_t^{(1)}$ and $\phi_t^{(2)}$, the latter weighed by the skewness of permanent shocks. If $\gamma_{\zeta_t} < 0$, i.e., permanent shocks are left skewed as we find subsequently, and $\phi_t^{(2)} < 0$, as our empirical estimates suggest, linearity biases $\phi_t^{(1)}$ upwards; the magnitude of this bias increases with the extent of income skewness. Similarly, the identifying condition for $\psi_t^{(1)}$ in the linear case identifies

$$-\mathbb{E}(\Delta c_{it} \times \Delta y_{it+1})/\sigma_{v_t}^2 = \psi_t^{(1)} + \psi_t^{(2)}(\gamma_{v_t}/\sigma_{v_t}^2).$$

If $\gamma_{v_t} < 0$, i.e., transitory shocks are left skewed as we find subsequently, and $\psi_t^{(2)} > 0$, as our baseline estimate suggests, linearity biases $\psi_t^{(1)}$ downwards; the bias is stronger the larger left skewness is. This may thus help explain the discrepancy between BPP (who find little pass-through of transitory shocks) and studies based on natural experiments (who find large consumption responses to said shocks).[17]

4 Empirical implementation

4.1 Data

As is evident from the previous section, panel data in income and consumption are needed to identify the moments of the permanent and transitory shocks and the degrees of insurance to them. The Panel Study of Income Dynamics (PSID) provides such data for a representative and continuously evolving sample of U.S. households since 1968. For most of this period, however, the PSID only collected consumption expenditures on food items. It is well known that food is an imperfect measure of consumption of nondurable goods, representing a declining fraction of the consumer's basket over time. Since 1999, however, the PSID was drastically redesigned in that it started collecting information over a much broader set of consumption items, currently comprising over 70% of consumption reported in the National

[17] The discrepancy is between studies that use covariance restrictions in the joint distribution of income and consumption (they find small pass-through of transitory shocks) and studies that employ natural experiments (e.g., tax refunds or lotteries; they find large pass-through). Commault (2022) explains the discrepancy on the basis of neglected effects from *past* transitory shocks on consumption, while Crawley and Kuchler (2023) explain it on the basis of neglected time aggregation in the data. Both are alternative types of misspecification in the linear consumption function/empirical implementation vis-à-vis our focus on nonlinearities here.

Income and Product Accounts (Blundell, Pistaferri, and Saporta-Eksten, 2016). We, therefore, use the PSID over the period between 1999 and 2019, the last available wave when writing this paper; the data are available only biennially. This is in contrast to BPP, who use PSID data over 1980–1992 and rely on consumption imputations from the Consumer Expenditure Survey (CEX). We return to this point below.

Our baseline sample consists of married couples in the PSID with a male spouse aged 30 to 65. We focus on a sample that is representative of the U.S. population, eliminating the supplementary low-income subsample (SEO). Households are also removed in the event of a break-up or remarriage; our analysis thus focuses on income risk rather than on divorce or other household dissolution factors. Naturally, we require non-missing data on income, expenditure, and basic demographics.[18] Our final sample consists of 20,866 household-year observations; on average, a household is observed for 5.6 periods, which corresponds to about 10 calendar years given the biennial nature of the data.

Table 1 presents summary statistics over the sample period 1999–2019. About 89% of men and 79% of women work in the labor market (but some households have intermittent employment); 68% of men have had some college education, while women fare slightly higher at 71%. Our measure of income is household disposable income, defined as the sum of taxable income (e.g., earnings and asset income), transfers, and social security income of the male and female spouses and of other family unit members, minus taxes paid. Average disposable income is $142,910 in 2018 prices.[19] Our measure of consumption is real expenditure on nondurable goods and services, comprising food (at home and outside), utilities, out-of-pocket health expenses (doctors, prescriptions), public transport and own vehicle expenses (gasoline, parking, insurance), education, and child care. Average real consumption is $26,766, with food at home accounting for about 30% of this. Housing, including rent and home insurance, is the main expenditure we exclude; we do so for comparability with BPP who also exclude it. Including housing does not change any of the subsequent results meaningfully.[20] Finally, table 1 reports summary statistics also for samples sorted by wealth and education; we postpone the discussion of those samples to section 5.4.

Aside from the PSID, we use a similarly selected sample drawn from the CEX to facilitate comparison with BPP. We defer the discussion of this secondary data to section 4.3.

[18]We also remove observations with extreme jumps in income or consumption in consecutive years, where a jump is a $\geq$10-fold increase or a $\geq$90% reduction, or the bottom 0.25% of the distribution of products of growth rates at t and $t-1$, i.e., large increases followed by large drops, likely reflecting measurement error.

[19]We express all monetary figures in 2018 prices because earnings are 1-year retrospective and our latest data are from 2019. We deflate using the Consumer Price Index from the Bureau of Labor Statistics.

[20]Blundell, Pistaferri, and Saporta-Eksten (2016) use consumption data from the modern PSID, including housing, to study family labor supply as an insurance mechanism. House prices, however, fluctuated dramatically in 2007–2009, artificially affecting housing expenditure especially for homeowners.

Table 1: Descriptive statistics

	Baseline sample		Low wealth	High wealth	No college	Some college
	Mean (1)	Median (2)	Mean (3)	Mean (4)	Mean (5)	Mean (6)
Consumption						
Nondurable+services	26,766	22,895	23,106	29,852	22,137	28,908
Food at home	8,256	7,609	7,861	8,580	7,462	8,630
Utilities	3,515	3,238	3,270	3,721	3,394	3,571
Food out	3,318	2,511	2,683	3,840	2,649	3,634
Health	1,614	863	1,390	1,803	1,429	1,699
Transportation	5,866	4,900	5,398	6,261	5,436	6,065
Education	3,202	0	1,529	4,614	1,262	4,100
Childcare	977	0	937	1,009	543	1,181
Rent (or rent eq.)	18,105	13,821	10,949	23,991	11,647	21,149
Home insurance	894	747	654	1,092	702	985
Income						
Disposable income	142,910	113,581	98,527	179,416	97,675	164,239
Taxable+transfer inc.	137,277	109,003	93,358	173,401	90,810	159,187
Earnings	118,423	97,154	87,005	144,266	78,565	137,217
Male	81,857	62,185	56,817	102,452	50,199	96,783
Female	36,567	29,183	30,187	41,814	28,366	40,433
Total assets	547,761	191,193	69,731	950,823	284,036	669,788
Demographics						
% working (male)	0.89	1.00	0.92	0.88	0.86	0.91
% working (female)	0.79	1.00	0.81	0.78	0.79	0.79
Age (male)	46	45	43	49	46	46
% some college (male)	0.68	1.00	0.55	0.79	0.00	1.00
% some college (female)	0.71	1.00	0.62	0.78	0.44	0.84
Periods per household	5.63	5.00	4.89	6.45	5.21	5.86
Observations	20,866		9,417	11,449	6,686	14,180

Notes: The table reports summary statistics for consumption, income, wealth, and basic demographics in our baseline sample from the PSID, and in subsamples sorted by wealth and education, over 1999–2019. See text for variable definitions. Rent for owners is calculated as 6% of the self-reported value of the house.

4.2 Estimation

Estimation follows a three-step approach. In the first step, we compute income and consumption growth net of their predictable components, $\mathbf{X}_{it}$ and $\mathbf{Z}_{it}$, respectively, in model notation, by regressing real income and consumption growth on a rich set of time effects and household characteristics.[21] The residuals from these regressions constitute the empirical counterparts to unexplained income and consumption growth, Δy_{it} and Δc_{it}, respectively. In the second step, we estimate the parameters of the income process – i.e., the second and higher moments of permanent and transitory shocks – using income data alone. In the third step, we estimate the transmission parameters of shocks and the moments of taste heterogeneity and consumption measurement error; we do so using income and consumption data conditional on the parameters of the income process from the second stage. We estimate the second and third step parameters using GMM and an identity weighting matrix.[22]

We present estimates of the transmission parameters separately for the linear and quadratic specifications of the consumption function. For the linear function, we first present results when only second moments of income and consumption growth are targeted, exactly as in BPP. This purports to replicate BPP over the more recent sample period. We next report results for the linear function when, in addition to the second moments, third and fourth moments are also targeted. This helps assess whether and how, through the lenses of the linear specification, information in the tails of the income distribution affects the measurement of partial insurance. We then move on to the results from the quadratic specification. Our choice to estimate the income parameters in an earlier step enables those parameters to remain unchanged regardless of the specification of the consumption function fitted. Appendix B.4 lists the moments targeted in each case.

Given the multi-step approach, we conduct inference using the block bootstrap jointly over all stages. We thus account for serial correlation and heteroskedasticity of arbitrary form, as well as for our use of pre-estimated parameters in the later stages of the multi-step approach. To avoid the standard errors being affected by extreme draws, we apply a normal approximation to the inter-quartile range of bootstrap replications.[23]

As we established in section 3, we must use an external estimate for the variance of income measurement error. We cannot otherwise separate the moments of the transitory

[21]The full set of controls includes year and year-of-birth dummies, dummies for education, race, region of residence, family size and number of children (level and change from last period), employment status (current and past), presence of outside dependents in the family, and presence of income recipients other than the husband and wife (both current and past). The dummies for education, race, and region are interacted with the year dummies to allow the effect of these characteristics to vary over time.

[22]Our baseline results come from equally weighted GMM. A diagonal weighting matrix, where the diagonal entries are the t-statistics of the sample moments, delivers similar results (not reported here but available upon request), although we did experience some convergence issues in that case.

[23]Instead of estimating standard errors as the standard deviation of the parameters' point estimates across replications, we compute them as the normalized interquartile range of the replications, i.e., the difference between the 75$^{\text{th}}$ and 25$^{\text{th}}$ percentiles divided by the interquartile range of the standard normal cdf. This is equivalent to applying a normal approximation to the distribution of bootstrap replications, thereby shielding the standard errors from extreme resampling draws. See Blundell, Pistaferri, and Saporta-Eksten (2016) and Theloudis (2021) for details. Standard errors without the normal approximation are very similar.

shock from those of the measurement error. Following Meghir and Pistaferri (2004) and Blundell, Pistaferri, and Saporta-Eksten (2016), we retrieve the error variance from Bound, Brown, Duncan, and Rodgers (1994), a popular validation study of the PSID, setting it equal to 4% of the variance of log income in any given year.[24]

4.3 Empirical moments of income and consumption

Table 2 reports the empirical second, third, and fourth moments of the cross-sectional distributions of unexplained income and consumption growth in the sample.

Income moments. Panel A presents the moments of income growth. The variance is 0.153, while the first-order autocovariance is -0.049 and highly significant.[25] The income process of section 2.1 implies that income growth is negatively first-order autocorrelated due to mean reversion in the transitory shock, and this is supported by the data.[26]

The third *standardized* moment of income growth is -0.143. Despite the small size of the cross-section and the inherent difficulties in estimating higher moments, the point estimate misses statistical significance at the 10% level only by a small margin. The point estimate suggests that income growth exhibits left skewness, i.e., a long left tail. In simple terms, the left tail of the distribution of income growth (i.e., of income shocks) is longer than the right tail, so a negative shock is on average more severe, more unsettling, than a positive one. This feature of the data is in line with Guvenen, Karahan, Ozkan, and Song (2021), who document strong left skewness in the distribution of male earnings growth using large administrative data from the U.S. Social Security Administration.

The fourth *standardized* moment of income growth is 10.032 and, as expected, highly significant. Income growth is strongly leptokurtic, i.e., it exhibits very high kurtosis. In simple terms, the distribution of income growth (i.e., of income shocks) has thicker tails but less mass in the shoulders, relative to a Gaussian density. Shocks are thus either close to the center of the distribution, i.e., zero, or out in the tails, i.e., rather extreme. This is again

[24] Bound, Brown, Duncan, and Rodgers (1994) survey workers from a single manufacturing firm in 1983 and 1987. They obtain information on earnings in a way that mimics the PSID questionnaire and coding practices and compare them to administrative data from the firm. There are two caveats in this approach: first, the sample of workers in the validation study comes from two decades prior to the data in this paper; second, we use estimates of error in male earnings to correct for error in household income.

[25] In contrast to the identifying statements of section 3, the data are biennial so we report moments for income/consumption growth at a biennial frequency. We maintain the notation Δx_t for the first difference of variable x over time, noting that this corresponds to a difference over two years. See appendix B for details.

[26] The second-order autocovariance is an order of magnitude smaller and only marginally significant. This feature and the biennial nature of the data, whereby the second-order autocovariance reflects correlations 4 calendar years apart, motivates our white noise specification for the transitory shock.

Table 2: Empirical moments of income and consumption growth

	Baseline sample			
Income data:	PSID		PSID	
Consumption data:	PSID		imputed from CEX	
	(1)		(2)	
Panel A. Moments of income growth				
$\text{Var}(\Delta y_{it})$	0.153	(0.005)	0.153	(0.005)
$\text{Cov}(\Delta y_{it}, \Delta y_{it+1})$	-0.049	(0.002)	-0.049	(0.002)
$\text{Skew}(\Delta y_{it})$	-0.143	(0.098)	-0.143	(0.098)
$\text{Cov}((\Delta y_{it})^2, \Delta y_{it+1})$	0.009	(0.003)	0.009	(0.003)
$\text{Kurt}(\Delta y_{it})$	10.032	(0.509)	10.032	(0.509)
$\text{Cov}((\Delta y_{it})^2, (\Delta y_{it+1})^2)$	0.080	(0.007)	0.080	(0.007)
Panel B. Moments of consumption growth				
$\text{Var}(\Delta c_{it})$	0.126	(0.002)	0.323	(0.015)
$\text{Cov}(\Delta c_{it}, \Delta c_{it+1})$	-0.044	(0.001)	-0.140	(0.009)
$\text{Skew}(\Delta c_{it})$	-0.039	(0.034)	-0.177	(0.083)
$\text{Cov}((\Delta c_{it})^2, \Delta c_{it+1})$	-0.003	(0.001)	0.190	(0.036)
$\text{Kurt}(\Delta c_{it})$	4.445	(0.100)	17.109	(2.111)
$\text{Cov}((\Delta c_{it})^2, (\Delta c_{it+1})^2)$	0.027	(0.001)	0.909	(0.198)
Panel C. Joint moments				
$\text{Cov}(\Delta y_{it}, \Delta c_{it})$	0.008	(0.001)	0.012	(0.002)

Notes: The table presents the second, third, and fourth moments of income and consumption growth in the baseline sample. See text for the definitions of income and consumption. Skewness and kurtosis correspond to the third and fourth standardized moments respectively. Column 1 reports moments of consumption internally in the PSID, while column 2 reports moments of consumption imputed from the CEX into the PSID. See appendix C.2 for details on the imputation. We maintain the notation Δx_t for the first difference of variable x over time, noting that, given the biennial nature of the data, this corresponds to a difference over two calendar years. Block bootstrap standard errors are in parentheses.

in line with Guvenen, Karahan, Ozkan, and Song (2021), who also document strong excess kurtosis in male earnings growth.

Guvenen, Karahan, Ozkan, and Song (2021) focus on male earnings while our model and, therefore, our statistics for income growth herein, revolve around household disposable income. To facilitate the comparison, appendix C.1 reports moments of male *earnings* growth in our baseline sample. The variance is very similar between the PSID and the U.S. Social Security Administration data, at 0.271 and 0.323 respectively. Skewness in our sample is -0.569 (and highly statistically significant) versus -1.039 in the administrative data. While the nature of skewness is similar in both cases, the population data feature substantially more left skewness than the PSID; Guvenen, Karahan, Ozkan, and Song (2021) confirm this. Kurtosis, finally, is identical between the two, at 13.711 and 13.494 respectively.

Left skewness and excess kurtosis are thus not just features of the growth rates in disposable household income (table 2), but they are also observed in alternative measures of income, such as male earnings (table C.1, panel A) and household earnings (table C.1, panel B). Together, left skewness and excess kurtosis suggest that income changes for most households are either close to their mean value of zero or far out in the left tail (large income *reductions*). These results corroborate (and, in fact, are very close numerically with) findings by De Nardi, Fella, and Paz-Pardo (2019) and De Nardi, Fella, Knoef, Paz-Pardo, and Van Ooijen (2021) over the early period of the PSID.

Consumption moments. Panel B of table 2 presents the moments of consumption growth; column 1 specifically reports moments from our baseline sample in the PSID. The variance of consumption growth is very similar to that of income growth (0.126 versus 0.153), reflecting either a large pass-through of income shocks or large taste heterogeneity/measurement error. The first-order autocovariance is -0.044 and highly significant. In the context of our model, this autocovariance reflects the variance of consumption error $\sigma^2_{u^c_t}$ (see appendix B).

The third *standardized* moment of consumption growth is -0.039. Albeit not statistically significant, it is in line with Constantinides and Ghosh (2017), who use the CEX and document negative skewness in (quarterly) consumption growth. Similar to income growth, the distribution of consumption growth in the PSID exhibits left skewness, i.e., it has a long left tail, albeit at a lesser extent than income.[27]

The fourth *standardized* moment of consumption growth is 4.445 and highly significant. There is excess kurtosis and the distribution is thus leptokurtic. In contrast to the distribution of income, however, excess kurtosis is substantially less pronounced here.

The covariance between income and consumption growth is 0.008. As we established in section 3, this moment is crucial for the identification of the pass-through of income shocks into consumption. While we cannot immediately gauge its magnitude, the strong statistical significance reflects that income shocks do transmit into consumption.

Unlike income, there are neither administrative data on consumption, nor has there been a validation study for consumption in the PSID.[28] As a consequence, we cannot benchmark our estimates of the consumption higher moments against the literature. The CEX, however, provides high quality consumption data for a representative sample of the U.S. population,

[27]The covariance between squared consumption growth and one-period-ahead consumption growth identifies, in the context of the linear model, the skewness of consumption error. The covariance is small, so measurement error does not contribute much to the negative skewness in consumption.

[28]Nevertheless, Andreski, Li, Samancioglu, and Schoeni (2014) and Li, Schoeni, Danziger, and Charles (2010) compare consumption in the PSID and in the CEX; Meyer and Sullivan (2023) offer a critical discussion of the CEX data against the PSID.

although it lacks a panel dimension. We can thus use the CEX to select a sample of households as close to our baseline as possible, and estimate similar cross-sectional moments across the surveys. The CEX is useful also for another reason. Since the PSID collected little consumption data prior to 1999, the early literature on consumption insurance, notably BPP, resorted to imputing expenditure from the CEX into the PSID. To facilitate the comparison with BPP, we can use the CEX to replicate the consumption imputation into the PSID. This is not strictly needed in our case (the PSID provides comprehensive data over our sample period) but the imputation is nevertheless informative, as we illustrate later.

Consumption in the CEX. We use CEX interview survey data over 1999-2019 and select a sample of stable married households that mimics closely our baseline selection in the PSID. We report details on the sample selection in appendix C.2.

Appendix table C.2 offers a comparison between the PSID and CEX samples over time. The two samples are very close with respect to age, family size, number of children, education, or labor market participation of the wife and the husband. Household income is higher in the PSID; BPP noted this and argued that it may be due to the more comprehensive definition of income in the PSID. Food expenditure in the PSID is also higher, which is in contrast to BPP who find that food expenditure over the earlier period 1980–1992 is similar across the two surveys. Consequently, overall consumption in the PSID is, on average, higher by about 25% than a similarly selected sample in the CEX.

Appendix figure C.1 plots the variance, skewness, and kurtosis of *log* consumption in the PSID and CEX samples. The patterns for the variance and kurtosis are quite similar in the two surveys over time, though the variance in the PSID is lower and the kurtosis is slightly higher. Skewness, however, is markedly different. Skewness of log consumption is strongly positive in the PSID, while it is mostly negative in the CEX; the patterns also diverge over time. Overall, these statistics suggest that consumption across otherwise similar samples in the PSID and CEX exhibits notable differences over our sample period.

This comparison, so far, is not informative about the moments of consumption *growth*, which is the relevant variable in our case. The CEX lacks the panel dimension needed for our analysis, so it cannot be used to obtain measures of annual (or biennial) consumption growth. This is why BPP imputed consumption from the CEX into the PSID: the CEX lacked the necessary panel dimension while the PSID lacked comprehensive consumption data before 1999. We repeat BPP's imputation over the *recent* years (i.e, in our sample period 1999–2019). The imputation consists of the estimation of food demand in the CEX (food expenditure expressed as a function of prices and nondurable consumption), the inversion of demand under standard assumptions, and the use of the inverted equation to impute nondurable consumption in the PSID based on food that is available in both surveys. This allows us to obtain an alternative, external measure of consumption in the PSID – and thus of consumption growth. Appendix C.2 provides details about the imputation.

The moments of imputed consumption growth are in column 2 of table 2. The variance at 0.323 is two-and-a-half times higher than the variance in the true series (column 1). In fact, the variance of imputed consumption growth is more than double the variance of income growth. The higher-order moments also differ markedly between the imputed and true series. Skewness in imputed consumption growth is -0.177 (versus -0.039 in the true series), while kurtosis is 17.109, almost four times higher than kurtosis in the true series.[29]

The main qualitative features of the distribution remain, however, similar. Consumption growth is non-Gaussian no matter which angle one looks at it. Nevertheless, the imputation likely imparts substantial measurement/imputation error to consumption, which likely explains why the second and higher-order moments are several-fold accentuated relative to the true series. Since the imputation is not strictly needed over our sample period, we use consumption *internally* available in the PSID as our baseline measure of consumption when estimating the model. Yet, the present analysis suggests that we may not want to rely too much on the higher-order moments of consumption, especially given the discrepancy in skewness of *log* consumption between PSID and CEX (this discrepancy is independent of the imputation). Fortunately, identification in the quadratic specification does not use such higher-order consumption moments, as we illustrated in section 3.3.

5 Results

We present our baseline results assuming stationarity of the distributions of income shocks, taste heterogeneity, and consumption measurement error. We also assume constant transmission parameters of income shocks over the lifecycle. We relax the stationarity/lifecycle-invariance assumptions in section 5.4.

5.1 Income process

We present estimates of the income process in three specifications. First, we target only second moments of income growth, which enables us to estimate the second moments of income shocks. This is similar to what BPP do over the earlier period 1980–1992. Second, we target second *and* third moments of income growth, which enables us to estimate also the skewness of income shocks. Finally, we target *up to* the fourth moments of income growth,

[29] Appendix C.2 reports additional summary statistics for imputed consumption.

which allows us to estimate all the parameters of the income process as presented in section 2.1. The results from each specification appear in table 3, columns 1-3, respectively. The full set of moments targeted in each case is shown in appendix B.4.1.

When only the second moments of income growth are targeted, we estimate the variance of permanent and transitory shocks at 0.027 and 0.033 respectively. These are very close to BPP and Meghir and Pistaferri (2004), who estimate these parameters over an earlier period. As the variance of income growth is driven by the variance of the contemporaneous permanent shock and the sum of the variances of transitory shocks in the current and previous periods, the contribution of the transitory shock to the variance of income growth is almost 2.5-fold that of the permanent shock (as in Gottschalk and Moffitt, 2009).

Targeting the third moments of income growth (in addition to the second moments) leaves the variances of shocks unchanged. The third moments convey additional information about the skewness of shocks, which we estimate at -0.947 and -1.310 for the permanent and transitory shock respectively (third standardized moment). Both shocks thus exhibit substantial negative skewness, i.e., their distribution has a much longer left than right tail. This is consistent with Geweke and Keane (2000) who find left skewness of earnings shocks in the early PSID, and Theloudis (2021) who finds left skewness of wage shocks in the recent period. Assuming for now that good and bad shocks transmit similarly into consumption, the long left tail implies that bad shocks are more damaging than good ones because they are typically further away from the zero mean – they are more extreme. This is true for both permanent and transitory shocks. Skewness of income growth, however, is driven mostly by permanent shocks (fully so under stationarity). Therefore, if targeting the third moments of income growth subsequently changes the degree of consumption insurance in the household, this must reflect the implications of permanent rather than of transitory shocks.

Targeting all moments of income growth up to fourth-order leaves the third central moments of permanent and transitory shocks unchanged and modifies the variances only by little (0.030 versus 0.027 and 0.031 versus 0.033, respectively). The fourth moments convey additional information about the kurtosis of shocks, which we estimate at 44.393 and 49.226 respectively (fourth standardized moment). In both cases, the distributions are highly leptokurtic: they are sharp pointed around the mean, have relatively little mass in the shoulders and very fat tails. This is consistent with Busch and Ludwig (2021), who use the PSID until 2012 and estimate numerically similar kurtosis coefficients for both types of shocks. Together with left skewness, large excess kurtosis implies that most households experience either negligible shocks (shocks close to the zero mean) or shocks that are quite bad. Ignoring these stark features of income shocks may lead to erroneous inference of the degree of consumption partial insurance, to which we now turn.

Table 3: Estimates of the income process, baseline sample

Income moments:			2nd	2nd, 3rd	2nd, 3rd, 4th
Income data:			PSID	PSID	PSID
			(1)	(2)	(3)
		Panel A. Permanent shocks			
Variance	σ_ζ^2		0.027	0.027	0.030
			(0.002)	(0.002)	(0.002)
Skewness	γ_ζ	central		-0.004	-0.004
				(0.003)	(0.003)
		standardized		-0.947	-0.831
				(0.649)	(0.571)
Kurtosis	κ_ζ	central			0.039
					(0.007)
		standardized			44.393
					(11.080)
		Panel B. Transitory shocks			
Variance	σ_v^2		0.033	0.033	0.031
			(0.002)	(0.002)	(0.003)
Skewness	γ_v	central		-0.008	-0.008
				(0.003)	(0.003)
		standardized		-1.310	-1.431
				(0.510)	(0.566)
Kurtosis	κ_v	central			0.048
					(0.005)
		standardized			49.226
					(10.119)

Notes: The table presents the estimates of the parameters of the income process, assuming stationarity over time/lifecycle. Column 1 presents parameter estimates when only the second moments of income growth are targeted; column 2 targets also third moments; column 3 targets all moments up to fourth-order. Estimation is via equally weighted GMM; block bootstrap standard errors are in parentheses.

5.2 Linear consumption function

We present estimates of the linear consumption function, first, targeting only second moments of the joint income-consumption distribution and, subsequently, targeting second as well as higher-order moments of the joint distribution. The full set of moments targeted in each case is shown in appendix B.4.2.

Targeting second-order moments only. The estimation of the linear consumption function targeting second moments alone is effectively a replication of BPP over the recent years

(1999-2019) while using consumption data that are internally available in the PSID.[30] This replication is important in its own right. The recent years offer superior consumption data (in the sense that we do not need to rely on a consumption imputation from the CEX), thus reducing the scope of measurement/imputation error. The data, however, come at a biennial rather than annual frequency. Moreover, there have been many changes in taxation and redistribution policies since 1980–1992, the original period studied in BPP, which may matter for the extent of partial insurance available to households today (see Borella, De Nardi, Pak, Russo, and Yang, 2022, for a discussion of such changes).[31]

The results are in column 1 of table 4. We estimate the transmission parameter of permanent shocks at $\phi^{(1)} = 0.152$. A 10% permanent change in income thus results in only 1.52% shift in consumption. This reflects a very low pass-through of permanent shocks and a substantial degree of consumption insurance. While highly significant, the estimate for $\phi^{(1)}$ is strikingly different from BPP, whose baseline estimate of 0.64, four times bigger, indicates a much larger pass-through of permanent shocks and a much lower degree of insurance. By contrast, we find evidence of full insurance to transitory shocks, similarly to BPP. We estimate the pass-through of said shocks at $\psi^{(1)} = -0.006$, statistically indistinguishable from zero. Therefore, transitory shocks do not pass through into consumption.[32]

What explains the discrepancy in $\phi^{(1)}$ with BPP? As suggested earlier, there are a few important differences between our replication and the original exercise in BPP: new consumption data and a different time period. To help decipher their role, we impute consumption from the CEX into the PSID in the same way BPP did for the earlier period, and we repeat the estimation; see appendix C.2 for details on the imputation. The results are in column 2 of table 4. The transmission parameter of permanent shocks doubles to $\phi^{(1)} = 0.288$, while the transmission parameter of transitory shocks remains statistically zero (though at a larger magnitude). Using imputed consumption data thus inflates the pass-through of permanent shocks, substantially lowering the degree of partial insurance. This is most likely the byproduct of measurement/imputation error, evident in the variance $\sigma^2_{u_c}$ of consumption error that almost quadruples between the true series (column 1) and the imputed data from the CEX (column 2).[33] By contrast, the variance of taste heterogeneity σ^2_{ξ} remains similar in

[30] We use the pre-estimated second moments of income shocks obtained while targeting only the second moments of income growth, i.e., column 1 of table 3.

[31] The nature and size of families and family networks may have also changed over the last 5 decades, thus affecting various informal insurance channels available to households.

[32] Commault (2022) generalizes BPP's linear consumption function by allowing *past* transitory shocks to enter the specification in (4). She then finds a strong consumption response to transitory income, similar to the response estimated from quasi-experimental data. We abstract from this generalization, as our focus here is on higher-order income moments and on non-linearities in the consumption response.

[33] The large extent of consumption measurement/imputation error is also evident in the comparison of the consumption moments of the true vs. imputed series in table 2, as well as in appendix figure C.1.

Table 4: Estimates of the consumption function, baseline sample

Consumption fn.:	**Linear**				**Quadratic**
Income moments:	2nd	2nd	2nd, 3rd	2nd, 3rd, 4th	2nd, 3rd, 4th
Consumption data:	PSID	CEX	PSID	PSID	PSID
	(1)	(2)	(3)	(4)	(5)
$\phi^{(1)}$	0.152	0.288	0.160	0.112	0.134
	(0.028)	(0.044)	(0.029)	(0.041)	(0.026)
$\psi^{(1)}$	-0.006	-0.109	-0.012	0.007	-0.003
	(0.045)	(0.078)	(0.038)	(0.056)	(0.052)
$\phi^{(2)}$					-0.040
					(0.026)
$\psi^{(2)}$					0.015
					(0.034)
$\omega^{(22)}$					0.320
					(0.770)
σ^2_{ξ}	0.019	0.019	0.019	0.020	0.019
	(0.001)	(0.004)	(0.001)	(0.001)	(0.001)
γ_{ξ}			-0.001	-0.001	
			(0.001)	(0.001)	
κ_{ξ}				0.005	
				(0.001)	
$\sigma^2_{u_c}$	0.044	0.140	0.044	0.044	0.044
	(0.002)	(0.009)	(0.002)	(0.001)	(0.002)
γ_{u_c}			0.002	0.002	
			(0.001)	(0.001)	
κ_{u_c}				0.013	
				(0.001)	

Notes: The table presents the estimates of the parameters of the consumption function, assuming homogeneous transmission parameters over the lifecycle/in the cross-section and stationarity of taste heterogeneity and consumption measurement error. Columns 1-4 present parameter estimates in the linear function, while column 5 presents estimates in the quadratic case; the order of moments targeted in each case is shown at the top of the table. Except for column 2 where we use imputed data from the CEX, all other columns use consumption data internally available in the PSID. Appendix table C.4 shows additional results from the CEX. Estimation is via equally weighted GMM; block bootstrap standard errors are in parentheses.

the two cases, and also very close to BPP. The imputation, therefore, seems to bring in large amounts of measurement error, explaining about 1/3 of the difference between our lower $\phi^{(1)}$ and BPP's higher estimate. We report additional insights in appendix C.2.

Although the imputation inflates $\phi^{(1)}$, the inflated value in column 2 is still substantially lower than BPP. The PSID switched in 1997 from annual to biennial frequency. It is thus possible that at the lower frequency certain higher frequency income shocks are smoothed out or undone by subsequent shocks. To assess this, we would need to re-estimate our model on annual growth rates, which is impossible after 1997. However, we can use the original BPP data over 1980–1992 to estimate the model *annually* versus *biennially*. We do this in appendix C.3. While we obtain $\phi^{(1)} = 0.64$ using annual growth rates (exactly as BPP), this drops to $\phi^{(1)} = 0.4$–0.5 using biennial rates. The biennial nature of the modern data may thus explain another 1/3 of the difference between our lower $\phi^{(1)}$ and BPP's higher estimate.[34] We attribute the remaining gap to calendar time effects, e.g., changes in taxes/benefits or other insurance channels between 1980–1992 and 1999–2019. It is worth noting that we are not the first paper to find such a low pass-through of permanent shocks in the recent years; Arellano, Blundell, and Bonhomme (2017) and Arellano, Blundell, Bonhomme, and Light (2023) use the modern PSID and find (in the specification closest to ours) a pass-through of about 0.2, numerically close to us.

Introducing higher-order moments. The estimation of the linear consumption function targeting higher-order moments is our first attempt to uncover how the tails of income risk affect the pass-through of shocks into consumption. In other words, we want to measure the degree of partial insurance going beyond the workhorse income-consumption covariance to account for the non-Gaussian features of modern income dynamics. We introduce the higher-order moments incrementally; we first target third moments (in addition to second); we then introduce fourth moments (in addition to second and third).[35]

Column 3 of table 4 presents the results from targeting third- in addition to second-order moments of consumption and income. The partial insurance parameters remain largely unchanged upon inclusion of the third moments. We estimate $\phi^{(1)} = 0.160$ (as opposed to 0.152 with second moments only) and $\psi^{(1)} = -0.012$ (again indistinguishable from zero). We thus again find large partial insurance to permanent shocks (a 10% permanent change in income results in only 1.6% shift in consumption) and full insurance to transitory shocks. Skewness in taste heterogeneity and consumption error is small (much smaller than skewness in income shocks), while their variances remain unchanged from the specification with second moments only. Therefore, left skewness in consumption is driven by negatively skewed income shocks rather than by taste heterogeneity or consumption measurement error.

Column 4 of table 4 presents the results from targeting all income and consumption mo-

[34] In both annual and biennial cases, $\phi^{(1)}$ measures the pass-through of the permanent shock at an *annual* rate. This becomes clear in appendix B that shows the estimating equations with biennial data. The difference in the estimates of $\phi^{(1)}$ between the two cases arises because with biennial data we only measure an 'aggregate' pass-through within a given biennial period. Consequently, some higher-frequency shocks will be inevitably smoothed out at the observed lower frequency and their transmission will be muted.

[35] We use pre-estimated income parameters targeting, respectively, up to 3^{rd} moments of income growth (column 2, table 3) and up to 4^{the} moments of income growth (column 3, table 3).

ments up to fourth-order. The earlier findings hold up. We estimate $\phi^{(1)} = 0.112$ (versus 0.152 and 0.160 in the lower-order cases) and $\psi^{(1)} = 0.007$ (indistinguishable from zero). A 10% permanent change in income thus results in 1.12% shift in consumption. Therefore, if anything, we find evidence for slightly *more* insurance to permanent shocks upon inclusion of the fourth-order moments. The dispersion and skewness of taste heterogeneity and consumption error do not change much from the lower-order cases; we find evidence for excess kurtosis in both cases, albeit at a much lesser extent than for income shocks.[36]

Clearly, targeting higher-order moments in the context of the linear consumption function has not altered the degree of partial insurance vis-à-vis targeting second-order moments alone. This seems counterintuitive given that tail income risk has been shown to matter for various phenomena in finance and macro. For example, tail risk explains asset pricing puzzles in Constantinides and Ghosh (2017) and plays a crucial role in business cycles in McKay (2017). Thematically closer to us, Ai and Bhandari (2021) show that tail risk is uninsured in realistically parameterized labor markets, while Busch and Ludwig (2021) find sizable welfare implications of higher-order risk. One would thus expect that tail risk *increases* the pass-through of shocks, as there must be a limit beyond which households cannot further insure consumption against extreme events.[37]

Yet, the linear consumption function leaves no room to such events. BPP rely on a log-linearization of the optimal consumption path around a lifecycle 'steady-state' (appendix A). But tail events move the agent far from such 'steady-state': if utility is sufficiently concave, a large shock would command a large consumption response. A linearized marginal utility, however, will fail to pick up such large response (echoing Carroll, 2001), which in turn likely explains why we do not see much action upon including higher-order moments. In other words, tail risk has, by design, no role in the linear consumption function, which in turn motivates and justifies our higher-order specification subsequently.

To summarize our findings from the linear consumption function: 1.) we replicate BPP over the recent years and find much lower pass-through of permanent shocks (larger degree of partial insurance); the imputation of consumption from the CEX and the biennial nature of

[36]Appendix table C.4 uses imputed consumption from the CEX to target higher-order moments as in columns 3-4, table 4. We do see slightly higher pass-through of permanent shocks upon inclusion of the higher-order moments, but we also find a strong, significant, and negative response to transitory shocks, which is hard to explain. We do not think those results are very credible; the imputation brings in substantial measurement/imputation error, which becomes worse the higher the order of the target moments is.

[37]Similar to our results with fourth moments, De Nardi, Fella, and Paz-Pardo (2019) and Busch and Ludwig (2021) find slightly higher insurance with non-normal vs. normal income shocks, which they attribute to stronger precautionary motives. Busch and Ludwig (2021) caution against BPP's linear framework, showing evidence of bias in the presence of tail risk. Our quadratic specification addresses their concern.

the modern PSID explain the bulk of the difference from BPP; 2.) targeting higher-order moments does not alter the degree of consumption partial insurance, which is not entirely surprising as tail risk has, by design, no role in the linear framework.

5.3 Quadratic consumption function

We now turn to the estimation of the quadratic consumption function. The results appear in column 5 of table 4; the full set of targeted moments is shown in appendix B.4.3.[38,39]

We estimate the transmission parameter of the average *permanent* shock (the coefficient on the *linear* part of a polynomial in the permanent shock) at $\phi^{(1)} = 0.134$, highly significant and close to the values 0.11–0.16 obtained from the linear specification. We estimate the transmission parameter of the average *transitory* shock (the coefficient on the *linear* part of a polynomial in the transitory shock) at $\psi^{(1)} = -0.003$, statistically indistinguishable from zero. As $\phi^{(1)}$ and $\psi^{(1)}$ remain largely unaltered relative to the linear case, interest lies in the transmission parameters of the typical shocks *away* from the middle of the distribution, i.e., the coefficients on the *quadratic* part of a polynomial in the shocks. We estimate a negative $\phi^{(2)} = -0.04$ for permanent shocks (statistically significant at the margin for now), while we estimate a positive $\psi^{(2)} = 0.015$ for transitory shocks (statistically insignificant for now). The coefficient $\omega^{(22)}$ on the product of permanent and transitory shocks is large but imprecisely estimated. The variance of taste heterogeneity σ^2_ξ and consumption measurement error $\sigma^2_{u_c}$ remain effectively unchanged from the linear case.[40]

The values of $\phi^{(2)}$ and $\psi^{(2)}$ have important implications for the pass-through of shocks into consumption. To help interpret the results, recall that the pass-through of permanent shocks is given by (6), that is, it depends on $\phi^{(1)}$, $\phi^{(2)}$, $\omega^{(22)}$, and on the magnitudes and signs of both income shocks. If, for simplicity, we set the transitory shock to its mean value ($v_{it} = 0$) and plug in the parameter estimates from table 4, we get

$$\partial \Delta c_{it} / \partial \zeta_{it} = 0.134 - 0.08 \times \zeta_{it}.$$

[38] Estimation of the quadratic consumption function requires up to 4^{th}-order moments of the income shocks, so we use the pre-estimated income parameters from column 3, table 3.

[39] The consumption-income moments targeted in the quadratic case differ from the moments targeted in the most general linear case, i.e., in column 4, table 4. This is because in the quadratic case certain 3^{rd}- and 4^{th}-order consumption moments require up to 8^{th} moments of income shocks, which we do not model. This raises an issue of comparability between the linear and quadratic specifications, as the target moments differ (appendix table B.2). However, we can estimate the linear function targeting the exact same moments that we use in the quadratic case, as the linear function is identified by a smaller set of lower-order moments anyway (section 3). Our estimates of the linear function in that case do not differ from those in column 4, table 4, so estimation is robust to the choice of moments. These results are available upon request.

[40] Skewness and kurtosis of taste heterogeneity and consumption error are not identified because they require higher-order moments of consumption, which depend on higher than fourth-order moments of income.

Two main results follow.

First, bad (negative) permanent shocks are more hurtful than good (positive) ones because they are associated with a much larger transmission parameter into consumption. Consider, for example, a large permanent shock to household income, $\zeta_{it} = \pm 0.5$, approximately three standard deviations away from the mean. If the shock is negative ($\zeta_{it} = -0.5$), then $\partial \Delta c_{it} / \partial \zeta_{it} = 0.174$; by contrast, if the shock is positive ($\zeta_{it} = 0.5$), then $\partial \Delta c_{it} / \partial \zeta_{it} = 0.094$. We can generalize this to bad and good shocks of any magnitude and conclude that

$$\left.\frac{\partial \Delta c_{it}}{\partial \zeta_{it}}\right|_{\zeta<0} > \left.\frac{\partial \Delta c_{it}}{\partial \zeta_{it}}\right|_{\zeta>0}.$$

Therefore, bad permanent shocks are not just more extreme than good shocks, i.e., further away from the zero mean – see left skewness in table 3, but they also have much larger pass-through rates into consumption. Bad permanent shocks are thus double hurtful.

Second, severe bad shocks are more hurtful than moderate ones because they too are associated with larger transmission parameters into consumption. Contrast, for example, a very large permanent bad shock, $\zeta_{it} = -0.7$, about four standard deviations away from the mean, with a moderate one, $\zeta_{it} = -0.07$, half standard deviation from the mean. In the first case, $\partial \Delta c_{it} / \partial \zeta_{it} = 0.19$ while in the latter case $\partial \Delta c_{it} / \partial \zeta_{it} = 0.139$. We can generalize this to bad shocks of any magnitude and conclude that

$$\left.\frac{\partial \Delta c_{it}}{\partial \zeta_{it}}\right|_{\zeta<<<0} > \left.\frac{\partial \Delta c_{it}}{\partial \zeta_{it}}\right|_{\zeta<0}.$$

This result is not hard to comprehend. When a bad shock hits, households can insure consumption by running down savings or by using other formal or informal insurance arrangements. But there is clearly a limit to this. As the magnitude of the bad shock increases, these insurance arrangements will be eventually exhausted, thus resulting in a larger fraction of the shock eventually transmitting into consumption.[41]

Why would a positive or small shock be better insured? While we cannot assess candidate mechanisms without a structural model, our best explanation is that people save a good income shock (thus consumption does not respond much) but they dissave in the presence of a bad one. As there is clearly a limit to how much households can dissave or borrow, especially in the presence of borrowing constraints, a large negative shock would typically affect consumption by relatively more.

[41] The parameter estimates also imply that the pass-through of a *good* shock decreases with its magnitude. While this seems counterintuitive, recall that *extreme* good shocks are rare in our data, which follows from left skewness and excess kurtosis in table 3. We should thus not extrapolate the implications of the parameter estimates for the case of *extreme* good shocks, as such shocks are effectively not present in our sample.

Our finding of asymmetric consumption response to bad versus good income news has recently been documented in multiple survey questions in which subjects directly report how they would react in different hypothetical scenarios. Based on the 2011-2014 Bank of England/NMG Consulting Survey, Bunn, Le Roux, Reinold, and Surico (2018) find that the marginal propensity to consume from negative income shocks is larger than from positive shocks. Similarly, Christelis, Georgarakos, Jappelli, Pistaferri, and van Rooij (2019) and Fuster, Kaplan, and Zafar (2020) find that responses to losses are much larger than responses to gains, and the magnitude of the response increases with the magnitude of the shock. While their hypothetical shocks are more transitory in nature, they argue that these responses are consistent with lifecycle models with concave utility and borrowing constraints.[42]

Moving on to transitory shocks, their pass-through is given by (7); this depends on $\psi^{(1)}$, $\psi^{(2)}$, $\omega^{(22)}$, and on the magnitudes and signs of both income shocks. If we set the permanent shock to its mean value ($\zeta_{it} = 0$) and plug in the parameter estimates from table 4, we get

$$\partial \Delta c_{it} / \partial v_{it} = -0.003 + 0.03 \times v_{it}.$$

Unlike the transmission parameters of permanent shocks, neither $\psi^{(1)}$ nor $\psi^{(2)}$ are statistically significant in the full sample. This implies that transitory shocks, be they good or bad, moderate or extreme, do not pass through into consumption. This is similar to our findings for the linear case as well as to BPP. By contrast, this is not the case for certain well-defined subsamples, but we postpone the discussion of the subsamples for section 5.4.

Despite the lack of statistical significance, the partial derivative with respect to transitory shocks has, at face value, an interesting implication: large good transitory shocks have a numerically larger pass-through than smaller ones, whose pass-through is economically zero. While we should be cautious in drawing conclusions about *extreme* good transitory shocks (given left skewness and excess kurtosis, large good transitory shocks are not present in our sample), this implication helps reconcile the large pass-through found in quasi-experimental studies with the negligible pass-through found in BPP and other semi-structural studies that use covariance restrictions on survey data. The former studies typically exploit large and infrequent transitory episodes (e.g., tax rebates or lotteries as in Parker, Souleles, Johnson, and McClelland, 2013; Misra and Surico, 2014; Fagereng, Holm, and Natvik, 2021), which are unlikely to be present in the pooled survey data that the latter studies exploit (e.g., BPP, Blundell, Pistaferri, and Saporta-Eksten, 2016; Theloudis, 2021). In an environment

[42] This is precisely the type of model that Busch and Ludwig (2021) simulate; they find larger consumption pass-through of negative permanent shocks versus positive ones. While we do not model borrowing constraints here, a strength of our approach is that it enables us to measure partial insurance in the data without taking a stance on the precise underlying structural model/economic mechanism at play.

where the transmission of shocks depends on their magnitude, these two types of settings will lead to opposite conclusions. Our nonlinear consumption function reconciles the two by permitting transitory shocks to pass through to consumption at a varying rate that depends on the shocks' magnitude and sign.

To summarize our findings from the quadratic consumption function: 1.) bad permanent shocks reduce consumption by more than good permanent shocks of the same magnitude increase it; 2.) severe bad permanent shocks reduce consumption by more than moderate ones do. Tail risk from permanent income shocks thus matters for consumption. For a given value of $\phi^{(1)}$, the econometrician will overestimate the true degree of partial insurance against bad permanent shocks if she assumes a linear consumption function instead of the nonlinear specification herein. By contrast, transitory shocks, be they good or bad, small or large, do to not impact consumption, at least not in a statistically significant way.[43]

5.4 Sensitivity to age, wealth, and education

If bad permanent shocks impact consumption more severely than good ones, one would expect this asymmetry to be influenced by household age, wealth, or education, all of which seem to matter for the insurance means available to households. To assess this claim, we repeat the estimation over subsamples sorted on the basis of these household characteristics.

We first consider how partial insurance changes with age, by repeating the estimation over three, almost equally sized, subsamples of households with the male spouse between 30–40, 41–50, and 51–65 years old respectively.[44] Intuitively, the transmission parameters of income shocks should exhibit some age dependence since wealth accumulates over the life cycle, which in turn influences the household's ability to insure against income fluctuations. In fact, the derivation of the consumption function in appendix A shows that the transmission parameters are linked to the share of financial wealth in the household's total wealth (comprising financial assets and lifetime income); for example, $\phi_{it}^{(1)} = 1 - \pi_{it}$, where π_{it} is the aforementioned share of financial wealth. Therefore, $\phi_{it}^{(1)}$ should decrease monotonically over the lifecycle as households gradually accumulate more wealth.

The results appear in table 5. Column 1 presents results from the linear specification when only second-order income and consumption moments are targeted. The transmission parameter $\phi^{(1)}$ exhibits a non-monotonic pattern over the three age brackets – decreasing from 0.154 over ages 30–40 to 0.126 over 41–50 before increasing to 0.180 over ages 51–65.

[43] Appendix table C.4 shows that similar conclusions are obtained if we use consumption imputed from the CEX. In particular, $\phi^{(2)} < 0$ (albeit not statistically significant), with a magnitude similar to that obtained directly from the PSID (-0.036 versus -0.04), and $\psi^{(2)} > 0$.

[44] We target the same moments as earlier, broken down by the age of the male spouse. The moments of income shocks are also age-specific but we do not report those results for brevity.

A similar non-monotonic pattern is observed in column 2, where second- and third-order moments are targeted, and is further accentuated in column 3 where all moments up to fourth-order are included in the estimation. This is again indicative of potential misspecification in the linear function, in the sense that its theoretical restrictions are not supported by the empirical results.[45] By contrast, the transmission parameter $\psi^{(1)}$ is statistically indistinguishable from zero in all three age brackets.

Column 4 presents the results from the quadratic specification. In this case, $\phi^{(1)}$ decreases monotonically with age, so the average permanent shock (the shock close to the mean of the distribution) impacts consumption gradually *less* as households age. In line with theory, this pattern reflects the age-increasing role of financial wealth for self-insurance. Perhaps more importantly, $\phi^{(2)}$, the transmission parameter of the typical shock *away* from the mean, also decreases monotonically with age. $\phi^{(2)}$ drops from 0.006 to -0.008 as we move from ages 30–40 to ages 41–50 (neither statistically significant) and becomes large and statistically significantly negative at $\phi^{(2)} = -0.105$ in the 51–65 age bracket. As a result, large negative permanent shocks impact consumption more severely at older than at younger ages.

The last finding echoes Guvenen, Karahan, Ozkan, and Song (2021), who show that skewness of earnings growth becomes more negative with the level of earnings and, consequently, with the earner's age on average. There is more room for income to fall sharply in this case than among the low-earners (and therefore the young). In fact, we do find larger left skewness in income growth as households age in the PSID and a more negative coefficient of skewness of permanent shocks among older households.[46] So, not only do permanent shocks become more negative as households age, but they also transmit into consumption at an increasingly larger rate. In other words, not only is there more room for income to fall among older earners, but this also causes even greater harm to consumption as households age. By contrast, despite a similar pattern for transitory shocks, their transmission parameters remain statistically indistinguishable from zero over the entire lifecycle.

We next consider two subsamples formed on the basis of household wealth. We define a household as 'low wealth' ('high wealth') if its average wealth over its entire observation window is less (more) than median real wealth in the sample over 1999–2019. Wealth is defined as the sum of home equity (value of main house minus mortgage), other estates, vehicles, value of any business, investment, pension wealth, savings, and other assets, net of

[45] This conclusion seems to be consistent with BPP who find a decreasing pattern with age only when they impose a linear age trend on the transmission parameters. By contrast, when they generalize this to a quadratic trend, they report a worsening of their results without including further details.

[46] The coefficient of skewness of permanent shocks more than triples in magnitude from -0.65 to -2.39 as we move from ages 30–40 to 51–65. The coefficient of skewness of transitory shocks, by contrast, does not change much with the age – its estimate is -1.43 and -1.45, respectively, over ages 30–40 and 51–65.

Table 5: Estimates of the consumption function, by age

Consumption fn.:		Linear						Quadratic	
Income moments:		2^{nd}		2^{nd}, 3^{rd}		2^{nd}, 3^{rd}, 4^{th}		2^{nd}, 3^{rd}, 4^{th}	
Consumption data:		PSID		PSID		PSID		PSID	
		(1)		(2)		(3)		(4)	
$\phi^{(1)}$	30-40	0.154	(0.045)	0.164	(0.051)	0.089	(0.058)	0.149	(0.050)
	41-50	0.126	(0.051)	0.146	(0.055)	0.079	(0.089)	0.115	(0.049)
	51-65	0.180	(0.049)	0.175	(0.051)	0.136	(0.061)	0.113	(0.057)
$\psi^{(1)}$	30-40	0.094	(0.079)	0.083	(0.073)	0.095	(0.107)	0.105	(0.090)
	41-50	0.010	(0.079)	-0.016	(0.072)	-0.075	(0.061)	0.019	(0.090)
	51-65	-0.046	(0.070)	-0.044	(0.066)	0.037	(0.101)	-0.063	(0.088)
$\phi^{(2)}$	30-40							0.006	(0.040)
	41-50							-0.008	(0.063)
	51-65							-0.105	(0.054)
$\psi^{(2)}$	30-40							0.036	(0.061)
	41-50							0.048	(0.046)
	51-65							-0.052	(0.079)
$\omega^{(22)}$	30-40							-1.090	(1.465)
	41-50							-0.290	(1.195)
	51-65							3.003	(1.306)
σ^2_ξ	30-40	0.014	(0.002)	0.014	(0.002)	0.016	(0.002)	0.013	(0.004)
	41-50	0.022	(0.002)	0.022	(0.002)	0.022	(0.002)	0.022	(0.002)
	51-65	0.021	(0.002)	0.021	(0.002)	0.023	(0.002)	0.010	(0.008)
γ_ξ	30-40			-0.001	(0.001)	-0.001	(0.001)		
	41-50			0.006	(0.001)	0.006	(0.002)		
	51-65			-0.006	(0.002)	-0.006	(0.002)		
κ_ξ	30-40					0.004	(0.002)		
	41-50					0.007	(0.002)		
	51-65					0.006	(0.002)		
$\sigma^2_{u_c}$	30-40	0.046	(0.003)	0.046	(0.003)	0.046	(0.003)	0.046	(0.003)
	41-50	0.040	(0.002)	0.040	(0.002)	0.040	(0.002)	0.040	(0.002)
	51-65	0.043	(0.002)	0.043	(0.002)	0.043	(0.002)	0.043	(0.002)
γ_{u_c}	30-40			0.000	(0.002)	0.000	(0.002)		
	41-50			0.002	(0.002)	0.002	(0.002)		
	51-65			0.004	(0.002)	0.004	(0.002)		
κ_{u_c}	30-40					0.014	(0.003)		
	41-50					0.012	(0.002)		
	51-65					0.012	(0.002)		

Notes: The table presents the estimates of the parameters of the consumption function, allowing them to vary over ages 30–40, 41–50, and 51–65 of the male spouse. Columns 1-3 present parameter estimates in the linear function, while column 4 presents estimates in the quadratic case; the order of moments targeted in each case is shown at the top of the table. All columns use consumption data internally available in the PSID. Estimation is via equally weighted GMM; block bootstrap standard errors are in parentheses.

medical, student, and credit card debt. Table 1 provides summary statistics for the wealth subsamples. Average wealth in the first group is \$69,731 in 2018 prices, while the second group has 14 times more wealth, an average of \$950,823. Income and earnings are almost twice as much in the second group, while consumption is only 30% higher. Labor force participation is slightly higher in the low-wealth sample.

Table 6, columns 1-2, presents the results in the quadratic model.[47] Three points emerge. First, $\phi^{(1)}$ is more than twice as large among the low vis-à-vis high wealth households (0.212 versus 0.091). As $\phi^{(1)}$ measures the transmission of the *average* permanent shock, low wealth households exhibit a larger pass-through of average shocks, i.e., they are less able to insure against them, compared to high wealth households. Second, $\phi^{(2)} < 0$ for both groups, albeit statistically insignificant. With the caveat of insignificance, the magnitude of $\phi^{(2)}$ is bigger among the wealthy. Large negative permanent shocks thus impact consumption of the wealthy more severely than they do among the less wealthy. As in the case of older households previously, there is more room for consumption to fall sharply among the wealthier. Third, $\psi^{(2)}$, the transmission parameter of the typical *transitory* shock *away* from the mean, is negative, large, and statistically significant among the less wealthy. Transitory shocks thus have an asymmetric pass-through among lower wealth households, with large bad shocks reducing consumption substantially more than good shocks increase it (expression (7)). So while *average* transitory shocks are fully insured among low wealth households, larger negative transitory shocks are *not*: they reduce consumption quite substantially as low wealth households lack the self-insurance means to smooth out even transitory fluctuations.

We finally turn to the role of education. We define a household as 'no college' ('some college') if the male spouse has not gone to college (has some college education). Table 1 provides summary statistics. Those with some college education earn 70% more income, they have more than twice as much wealth, they consume about 30% more, while they also work slightly more than those without college education.

Table 6, columns 3-4, presents the results in the quadratic model.[48] Similar to the patterns observed by wealth, we find that $\phi^{(1)}$ is higher (double) for households without college education (0.219 versus 0.105). The pass-through of the *average* permanent shock is therefore larger among the lower educated, presumably because of the lower wealth that these households hold. $\phi^{(2)}$ is negative in both groups, suggesting that large bad permanent shocks

[47] Appendix table C.5, columns 1-4, presents the parameter estimates in the linear model. $\phi^{(1)}$ is consistently smaller among the high wealth households, regardless of which set of moments are targeted. $\psi^{(1)}$ is statistically indistinguishable from zero for both groups.

[48] Appendix table C.5, columns 5-8, presents the parameter estimates in the linear model. The results do not present a clear pattern across education groups. When only second moments of income and consumption are targeted, $\phi^{(1)}$ is larger for households without college education – 0.225 versus 0.123. However, when all moments up to fourth-order are included, $\phi^{(1)}$ is almost identical in the two groups.

Table 6: Estimates of the consumption function, by wealth and education

Consumption fn.:	**Quadratic**			
	By wealth		By education	
Subsample:	Low wealth	High wealth	No college	Some college
Income moments:	2^{nd}, 3^{rd}, 4^{th}	2^{nd}, 3^{rd}, 4^{th}	2^{nd}, 3^{rd}, 4^{th}	2^{nd}, 3^{rd}, 4^{th}
Consumption data:	PSID	PSID	PSID	PSID
	(1)	(2)	(3)	(4)
$\phi^{(1)}$	0.212	0.091	0.219	0.105
	(0.041)	(0.033)	(0.057)	(0.032)
$\psi^{(1)}$	-0.081	0.017	-0.071	0.030
	(0.090)	(0.059)	(0.079)	(0.063)
$\phi^{(2)}$	-0.018	-0.039	-0.095	-0.023
	(0.052)	(0.035)	(0.048)	(0.035)
$\psi^{(2)}$	-0.108	0.031	0.050	0.012
	(0.065)	(0.041)	(0.063)	(0.040)
$\omega^{(22)}$	1.116	-0.049	0.668	0.095
	(0.874)	(0.897)	(1.067)	(0.971)
σ^2_{ξ}	0.013	0.022	0.013	0.021
	(0.002)	(0.002)	(0.003)	(0.002)
$\sigma^2_{u_c}$	0.048	0.042	0.043	0.044
	(0.002)	(0.002)	(0.003)	(0.002)

Notes: The table presents the estimates of the parameters of the quadratic consumption function, allowing them to vary across wealth and education groups (linear specification estimates appear in appendix table C.5). Low (high) wealth is defined on the basis of household wealth being less (more) than median real wealth in the sample over 1999–2019. No versus some college is defined on the basis of the highest level of education attained by the male spouse. All columns use consumption data internally available in the PSID. Estimation is via equally weighted GMM; block bootstrap standard errors are in parentheses.

transmit into consumption more than equal-sized good shocks. However, the magnitude of the estimate in the lower education group is over four-fold larger (and statistically significant) compared to the higher education group. Permanent shocks *away* from the mean are therefore way more hurtful among the lower educated. By contrast, the transmission parameters of all transitory shocks are statistically indistinguishable from zero in both groups.

6 Conclusion

We measure the degree of consumption partial insurance to income shocks, accounting for higher-order moments of the distribution of income. Such moments, in particular left skewness and excess kurtosis, are important features of unexplained income growth, highlighted in several recent papers. Our focus is on *measuring* partial insurance without taking a stance on specific mechanisms at the nexus of income and consumption; as such, we relate closely to the seminal work of Blundell, Pistaferri, and Preston (2008) – BPP as we called it throughout – who introduced a methodology to measure partial insurance directly in the data.

Tail income risk, however, plays little role in their framework because they linearize the consumption rule (a poor approximation when extreme shocks command a large consumption response due to concave utility) and they ignore moments higher than second-order. These are all features that the present paper improves on. We use recent income and consumption data from the PSID (1999–2019) and we establish three main new insights.

First, in replicating BPP over the recent years, we find a much lower pass-through of permanent shocks (larger partial insurance) and full insurance to transitory ones. We attribute one third of the difference from BPP to the consumption imputation from the CEX, another third to the biennial nature of our data, and the remainder to time effects.

Second, in introducing third- and fourth-order income and consumption moments into the linear consumption model does not alter the above results.

Third, in estimating a generalized higher-order consumption function, we find that bad permanent shocks transmit at larger rates than good shocks, and their pass-through increases with the severity of the shock. By contrast, transitory shocks are fully insured in the full sample, at least from a statistical point-of-view, but they are not insured among the less wealthy. Among them, in particular, bad transitory shocks transmit into consumption at increasingly larger rates than good ones, exactly as permanent shocks do.

In summary, we contribute an analytical framework to measure the transmission of income shocks into consumption, allowing such transmission to vary with the sign and severity of the shock. To measure the transmission, we exploit information in higher-order income moments and we characterize how our measures of partial insurance change as we gradually introduce those moments. Our results help reconcile the traditional partial insurance measurement literature, in which BPP is the default methodology, with both the quantitative literature (e.g., De Nardi, Fella, and Paz-Pardo, 2019) and the more empirical one based on surveys or quasi-natural experiments (e.g., Fuster, Kaplan, and Zafar, 2020), which document asymmetries in the pass-through of good and bad, or small and large shocks. Overall, our results suggest that tail income risk matters substantially for consumption.

References

AI, H., AND A. BHANDARI (2021): "Asset Pricing With Endogenously Uninsurable Tail Risk," *Econometrica*, 89(3), 1471–1505.

Andreski, P., G. Li, M. Z. Samancioglu, and R. Schoeni (2014): "Estimates of Annual Consumption Expenditures and Its Major Components in the PSID in Comparison to the CE," *American Economic Review*, 104(5), 132–35.

Arellano, M., R. Blundell, and S. Bonhomme (2017): "Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework," *Econometrica*, 85(3), 693–734.

Arellano, M., R. Blundell, S. Bonhomme, and J. Light (2023): "Heterogeneity of Consumption Responses to Income Shocks in the Presence of Nonlinear Persistence," *Journal of Econometrics*.

Attanasio, O., and S. J. Davis (1996): "Relative Wage Movements and the Distribution of Consumption," *Journal of Political Economy*, 104(6), 1227–1262.

Blundell, R., M. Borella, J. Commault, and M. D. Nardi (2020): "Why Does Consumption Fluctuate in Old Age and How Should the Government Insure It?," NBER Working Paper 27348.

Blundell, R., H. Low, and I. Preston (2013): "Decomposing Changes in Income Risk Using Consumption Data," *Quantitative Economics*, 4(1), 1–37.

Blundell, R., L. Pistaferri, and I. Preston (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98(5), 1887–1921.

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016): "Consumption Inequality and Family Labor Supply," *American Economic Review*, 106(2), 387–435.

Blundell, R., and I. Preston (1998): "Consumption Inequality and Income Uncertainty," *The Quarterly Journal of Economics*, 113(2), 603–640.

Bonhomme, S., and J.-M. Robin (2010): "Generalized Non-Parametric Deconvolution with an Application to Earnings Dynamics," *The Review of Economic Studies*, 77(2), 491–533.

Borella, M., M. De Nardi, M. Pak, N. Russo, and F. Yang (2022): "The Importance of Modeling Income Taxes Over time. U.S. Reforms and Outcomes," NBER Working Papers 30725, National Bureau of Economic Research, Inc.

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(3), 345–368.

Brav, A., G. M. Constantinides, and C. C. Geczy (2002): "Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence," *Journal of Political Economy*, 110(4), 793–824.

Bunn, P., J. Le Roux, K. Reinold, and P. Surico (2018): "The consumption response to positive and negative income shocks," *Journal of Monetary Economics*, 96, 1–15.

Busch, C., D. Domeij, F. Guvenen, and R. Madera (2018): "Asymmetric Business-Cycle Risk and Social Insurance," NBER Working Papers 24569, National Bureau of Economic Research, Inc.

Busch, C., and A. Ludwig (2021): "Higher-order Income Risk Over the Business Cycle," Unpublished manuscript.

Carroll, C. D. (2001): "Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation)," *Advances in Macroeconomics*, 1(1).

Christelis, D., D. Georgarakos, T. Jappelli, L. Pistaferri, and M. van Rooij (2019): "Asymmetric Consumption Effects of Transitory Income Shocks," *The Economic Journal*, 129(622), 2322–2341.

Cochrane, J. H. (1991): "A Simple Test of Consumption Insurance," *Journal of Political Economy*, 99(5), 957–976.

Commault, J. (2022): "Does Consumption Respond to Transitory Shocks? Reconciling Natural Experiments and Semistructural Methods," *American Economic Journal: Macroeconomics*, 14(2), 96–122.

Constantinides, G. M., and A. Ghosh (2017): "Asset Pricing with Countercyclical Household Consumption Risk," *The Journal of Finance*, 72(1), 415–460.

Crawley, E., and A. Kuchler (2023): "Consumption Heterogeneity: Micro Drivers and Macro Implications," *American Economic Journal: Macroeconomics*, 15(1), 314–41.

De Nardi, M., G. Fella, M. Knoef, G. Paz-Pardo, and R. Van Ooijen (2021): "Family and Government Insurance: Wage, Earnings, and Income Risks in the Netherlands and the U.S.," *Journal of Public Economics*, 193, 104327.

De Nardi, M., G. Fella, and G. Paz-Pardo (2019): "Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare," *Journal of the European Economic Association*, 18(2), 890–926.

DEATON, A., AND C. PAXSON (1994): "Intertemporal Choice and Inequality," *Journal of Political Economy*, 102(3), 437–467.

FAGERENG, A., M. B. HOLM, AND G. J. NATVIK (2021): "MPC Heterogeneity and Household Balance Sheets," *American Economic Journal: Macroeconomics*, 13(4), 1–54.

FUSTER, A., G. KAPLAN, AND B. ZAFAR (2020): "What Would You Do with $500? Spending Responses to Gains, Losses, News, and Loans," *The Review of Economic Studies*, 88(4), 1760–1795.

GEWEKE, J., AND M. KEANE (2000): "An Empirical Analysis of Earnings Dynamics among Men in the PSID: 1968–1989," *Journal of Econometrics*, 96(2), 293–356.

GOTTSCHALK, P., AND R. MOFFITT (2009): "The Rising Instability of U.S. Earnings," *Journal of Economic Perspectives*, 23(4), 3–24.

GOURINCHAS, P.-O., AND J. A. PARKER (2002): "Consumption Over the Life Cycle," *Econometrica*, 70(1), 47–89.

GUVENEN, F., F. KARAHAN, S. OZKAN, AND J. SONG (2021): "What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?," *Econometrica*, 89(5), 2303–2339.

GUVENEN, F., S. OZKAN, AND J. SONG (2014): "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, 122(3), 621–660.

GUVENEN, F., AND A. A. SMITH (2014): "Inferring Labor Income Risk and Partial Insurance From Economic Choices," *Econometrica*, 82(6), 2085–2129.

HALL, R. E., AND F. S. MISHKIN (1982): "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," *Econometrica*, 50(2), 461–481.

HAYASHI, F., J. ALTONJI, AND L. KOTLIKOFF (1996): "Risk-Sharing between and within Families," *Econometrica*, 64(2), 261–294.

HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2014): "Consumption and Labor Supply with Partial Insurance: An Analytical Framework," *American Economic Review*, 104(7), 2075–2126.

HRYSHKO, D., AND I. MANOVSKII (2022): "How much Consumption Insurance in the U.S.?," *Journal of Monetary Economics*, 130, 17–33.

Kaplan, G., and G. L. Violante (2010): "How Much Consumption Insurance beyond Self-Insurance?," *American Economic Journal: Macroeconomics*, 2(4), 53–87.

Krueger, D., and F. Perri (2006): "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory," *The Review of Economic Studies*, 73(1), 163–193.

Lewis, D., D. Melcangi, and L. Pilossoph (2022): "Latent Heterogeneity in the Marginal Propensity to Consume," Unpublished manuscript.

Li, G., R. Schoeni, S. Danziger, and K. Charles (2010): "New Expenditure Data in the PSID: Comparisons with the CE," *Monthly Labor Review*, 133.

Madera, R. (2019): "The Consumption Response to Tail Earnings Shocks," Unpublished manuscript.

McKay, A. (2017): "Time-varying Idiosyncratic Risk and Aggregate Consumption Dynamics," *Journal of Monetary Economics*, 88, 1–14.

Meghir, C., and L. Pistaferri (2004): "Income Variance Dynamics and Heterogeneity," *Econometrica*, 72(1), 1–32.

——— (2011): "Earnings, Consumption and Life Cycle Choices," vol. 4 of *Handbook of Labor Economics*, pp. 773–854.

Meyer, B. D., and J. X. Sullivan (2023): "Consumption and Income Inequality in the United States since the 1960s," *Journal of Political Economy*, 131(2), 247–284.

Misra, K., and P. Surico (2014): "Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs," *American Economic Journal: Macroeconomics*, 6(4), 84–106.

Parker, J. A., N. S. Souleles, D. S. Johnson, and R. McClelland (2013): "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, 103(6), 2530–53.

Theloudis, A. (2021): "Consumption Inequality across Heterogeneous Families," *European Economic Review*, 136.

Wu, C., and D. Krueger (2021): "Consumption Insurance against Wage Risk: Family Labor Supply and Optimal Progressive Income Taxation," *American Economic Journal: Macroeconomics*, 13(1), 79–113.

Appendices – for online publication

A Derivation of the consumption function

In this appendix we derive the consumption function under a permanent-transitory specification for log income. We follow similar derivations in Blundell, Pistaferri, and Preston (2008), BPP hereafter, though our approach is a refinement of theirs given the second-order approximation we carry out herein. We show below how to derive the most general (quadratic) consumption function (5); the linear function (4) is a special case of this. We let a time unit in the model correspond to one calendar year. However, data in the PSID in recent years are available only every second year, so we adapt the derivations to reflect this.

Let the utility function take the form $U(C_{it}; \mathbf{Z}_{it}) = \widetilde{U}(C_{it} \exp(-\mathbf{Z}'_{it}\boldsymbol{\alpha}))$, which simplifies the subsequent statements. The first-order conditions of the household problem (1) *s.t.* (2) in a generic period $t+s$ are

$$\begin{aligned} [C_{it+s}] : \quad & \widetilde{U}_C(\widetilde{C}_{it+s}) \exp(-\mathbf{Z}'_{it+s}\boldsymbol{\alpha}) = \lambda_{it+s} \\ [A_{it+s+1}] : \quad & \beta(1+r)\mathbb{E}_{t+s}\lambda_{it+s+1} = \lambda_{it+s}, \end{aligned}$$

where $\widetilde{U}_C$ is the marginal utility of consumption and $\widetilde{C}_{it} = C_{it} \exp(-\mathbf{Z}'_{it}\boldsymbol{\alpha})$. λ_{it} is the Lagrange multiplier on the sequential budget constraint (the marginal utility of wealth).

A.1 Taylor approximation to optimality conditions

Approximation to static optimality condition. Applying logs to the static condition and taking a first difference over time yields $\Delta \ln \widetilde{U}_C(\widetilde{C}_{it+s}) - \Delta(\mathbf{Z}'_{it+s}\boldsymbol{\alpha}) = \Delta \ln \lambda_{it+s}$. A first-order Taylor expansion of $\ln \widetilde{U}_C(\widetilde{C}_{it+s})$ around $\widetilde{C}_{it+s-1}$ yields

$$\begin{aligned} \Delta \ln \widetilde{U}_C(\widetilde{C}_{it+s}) &\approx \frac{\widetilde{U}_{CC}(\widetilde{C}_{it+s-1})}{\widetilde{U}_C(\widetilde{C}_{it+s-1})} \exp(-\mathbf{Z}'_{it+s-1}\boldsymbol{\alpha})(C_{it+s} - C_{it+s-1}) \\ &\approx \frac{\widetilde{U}_{CC}(\widetilde{C}_{it+s-1})}{\widetilde{U}_C(\widetilde{C}_{it+s-1})} C_{it+s-1} \exp(-\mathbf{Z}'_{it+s-1}\boldsymbol{\alpha}) \Delta \ln C_{it+s} = \theta^{-1}_{it+s-1} \Delta \ln C_{it+s}, \end{aligned}$$

where $\widetilde{U}_{CC}$ is the second derivative of the utility function and $\theta_{it} = \widetilde{U}^{-1}_{CC}(\widetilde{C}_{it})\widetilde{C}^{-1}_{it}\widetilde{U}_C(\widetilde{C}_{it})$ is the inverse of the consumption substitution elasticity. Plugging the last expression in the log-linearized first-order condition yields $\Delta \ln C_{it+s} - \theta_{it+s-1}\Delta(\mathbf{Z}'_{it+s}\boldsymbol{\alpha}) = \theta_{it+s-1}\Delta \ln \lambda_{it+s}$.

Approximation to Euler equation. The approximation to the intertemporal condition involves future expectations. Let $\exp(\Gamma) = \beta^{-1}(1+r)^{-1}$ for some Γ. A second-order Taylor expansion of $\exp(\ln \lambda_{it+s+1})$ around $\ln \lambda_{it+s} + \Gamma$ yields

$$\lambda_{it+s+1} \approx \frac{\lambda_{it+s}}{\beta(1+r)} \left(1 + (\Delta \ln \lambda_{it+s+1} - \Gamma) + \frac{1}{2}(\Delta \ln \lambda_{it+s+1} - \Gamma)^2\right).$$

Taking expectations at time $t+s$ and noting that $\mathbb{E}_{t+s}\lambda_{it+s+1} = \lambda_{it+s}\beta^{-1}(1+r)^{-1}$, we obtain

$$\mathbb{E}_{t+s}\Delta \ln \lambda_{it+s+1} \approx \Gamma - \frac{1}{2}\mathbb{E}_{t+s}(\Delta \ln \lambda_{it+s+1} - \Gamma)^2 \Rightarrow \Delta \ln \lambda_{it+s+1} \approx \omega_{it+s+1} + \epsilon_{it+s+1}.$$

The first term $\omega_{it+s+1} = \Gamma - \frac{1}{2}\mathbb{E}_t(\Delta \ln \lambda_{it+s+1} - \Gamma)^2$ captures the impact of interest rates, impatience, and precautionary motives on the growth of the marginal utility of wealth. To maintain tractability, we assume that ω is non-stochastic but possibly heterogeneous across households. The second term is an expectation error with $\mathbb{E}_t(\epsilon_{it+1}) = 0$; it captures idiosyncratic revisions to the marginal utility of wealth upon arrival of income shocks.

Combining the approximations to the optimality conditions, we obtain

$$\begin{aligned} \Delta \ln C_{it+s} - \theta_{it+s-1}\Delta(\mathbf{Z}'_{it+s}\boldsymbol{\alpha}) &= \theta_{it+s-1}(\omega_{it+s} + \epsilon_{it+s}) \\ \Delta \ln C_{it+s} - \theta_{it+s-1}\Delta(\mathbf{Z}'_{it+s}\boldsymbol{\alpha}) - \Xi_{it+s} &= \xi_{it+s} + \theta_{it+s-1}\epsilon_{it+s} \\ \Delta c_{it+s} &= \xi_{it+s} + \theta_{it+s-1}\epsilon_{it+s}, \end{aligned} \tag{A.1}$$

where $\Delta c_{it+s} = \Delta \ln C_{it+s} - \theta_{it+s-1}\Delta(\mathbf{Z}'_{it+s}\boldsymbol{\alpha}) - \Xi_{it+s}$. We have split $\theta_{it+s-1}\omega_{it+s}$ into Ξ_{it+s} and ξ_{it+s}; the first term reflects the gradient of consumption due to interest rates and impatience (discounting), the second term reflects unobserved consumption taste heterogeneity. We let demographics and time fixed effects pick up the effect of taste shifters and the gradient in the consumption path. Therefore Δc_{it+s} can be obtained empirically from a regression of log consumption on appropriate observables.

A.2 Taylor approximation to lifetime budget constraint

Let $G(\boldsymbol{\xi}) = \ln \sum_{s=0}^{N} \exp \xi_s$ for $\boldsymbol{\xi} = (\xi_0, \xi_1, \ldots, \xi_N)'$. Applying a second-order Taylor expansion of $G(\boldsymbol{\xi})$ around a deterministic $\boldsymbol{\xi}^0$, and taking expectations given information $\mathcal{I}$, yields

$$\begin{aligned} \mathbb{E}_{\mathcal{I}}G(\boldsymbol{\xi}) \approx G(\boldsymbol{\xi}^0) &+ \sum_{s=0}^{N} \frac{\exp \xi_s^0}{\sum_{\kappa=0}^{N} \exp \xi_\kappa^0} \mathbb{E}_{\mathcal{I}}(\xi_s - \xi_s^0) \\ &+ \frac{1}{2}\sum_{s=0}^{N}\sum_{\ell=0}^{N} \frac{\exp \xi_s^0}{\sum_{\kappa=0}^{N} \exp \xi_\kappa^0} \left(\delta_{s\ell} - \frac{\exp \xi_\ell^0}{\sum_{\kappa=0}^{N} \exp \xi_\kappa^0}\right) \mathbb{E}_{\mathcal{I}}(\xi_s - \xi_s^0)(\xi_\ell - \xi_\ell^0), \end{aligned} \tag{A.2}$$

where $\delta_{s\ell}$ is the Kronecker delta ($\delta_{s\ell} = 1$ if $s = \ell$, $\delta_{s\ell} = 0$ otherwise). We apply this approximation to both sides of the intertemporal budget constraint.

Assuming expectations away, the logarithm of the *left* hand side of the budget constraint (2) in a generic period t is $\ln\left(A_{it} + \sum_{s=0}^{T-t} \exp\left(\ln Y_{it+s} - s\ln(1+r)\right)\right)$. Letting

$$\xi_s = \begin{cases} \ln A_{it+s} & \text{for } s = 0 \\ \ln Y_{it+s-1} - (s-1)\ln(1+r) & \text{for } s = 1, \ldots, T-t+1 \end{cases}$$
$$\xi_s^0 = \begin{cases} \mathbb{E}_{t-2} \ln A_{it+s} & \text{for } s = 0 \\ \mathbb{E}_{t-2} \ln Y_{it+s-1} - (s-1)\ln(1+r) & \text{for } s = 1, \ldots, T-t+1, \end{cases}$$

and declaring the information set to contain time t information, we follow (A.2) to write

$$\begin{aligned}
&\mathbb{E}_t \ln\left(A_{it} + \sum_{s=0}^{T-t} \frac{Y_{it+s}}{(1+r)^s}\right) \approx \ln\left(\exp(\mathbb{E}_{t-2}\ln A_{it}) + \sum_{s=0}^{T-t} \exp\left(\mathbb{E}_{t-2}\ln Y_{it+s} - s\ln(1+r)\right)\right) \\
&+ \pi_{it}\mathbb{E}_t(\ln A_{it} - \mathbb{E}_{t-2}\ln A_{it}) \\
&+ (1-\pi_{it})\sum_{s=0}^{T-t} \vartheta_{it+s}^Y \mathbb{E}_t(\ln Y_{it+s} - \mathbb{E}_{t-2}\ln Y_{it+s}) \\
&+ \frac{1}{2}\pi_{it}(1-\pi_{it})\mathbb{E}_t(\ln A_{it} - \mathbb{E}_{t-2}\ln A_{it})^2 \\
&+ \frac{1}{2}\sum_{s=0}^{T-t}\sum_{\ell=0}^{T-t}(1-\pi_{it})\vartheta_{it+s}^Y(\delta_{s\ell} - (1-\pi_{it})\vartheta_{it+\ell}^Y)\mathbb{E}_t(\ln Y_{it+s} - \mathbb{E}_{t-2}\ln Y_{it+s})(\ln Y_{it+\ell} - \mathbb{E}_{t-2}\ln Y_{it+\ell}) \\
&- \pi_{it}(1-\pi_{it})\sum_{s=0}^{T-t} \vartheta_{it+s}^Y \mathbb{E}_t(\ln A_{it} - \mathbb{E}_{t-2}\ln A_{it})(\ln Y_{it+s} - \mathbb{E}_{t-2}\ln Y_{it+s}).
\end{aligned}$$

The notation is: $\pi_{it} = \frac{Q_{1t}}{Q_{1t}+Q_{2t}}$ is equal to the expected share of financial wealth $Q_{1t} = \exp(\mathbb{E}_{t-2}\ln A_{it})$ in the household's total financial and human wealth at t, the latter defined as $Q_{2t} = \sum_{\kappa=0}^{T-t} \exp\left(\mathbb{E}_{t-2}\ln Y_{it+\kappa} - \kappa\ln(1+r)\right)$, i.e., the household's expected lifetime income for the remaining of life; $\vartheta_{it+s}^Y = \exp\left(\mathbb{E}_{t-2}\ln Y_{it+s} - s\ln(1+r)\right)/Q_{2t}$ is an annuitization factor equal to the expected share of time $t+s$ household income in total lifetime income. π_{it} and ϑ_{it+s}^Y pertain to expectations at $t-2$ so they are both known at $t-2$ and later.

Using the income process and assuming the deterministic observables away (they cancel out in a first difference in expectations at a later stage), we reach $\ln Y_{it+s} = \ln Y_{it-2} + \sum_{\tau=-1}^{s} \zeta_{it+\tau} + v_{it+s} - v_{it-2}$, therefore $\ln Y_{it+s} - \mathbb{E}_{t-2}\ln Y_{it-2} = \sum_{\tau=-1}^{s}\zeta_{it+\tau} + v_{it+s}$. It follows that $\sum_{s=0}^{T-t}\vartheta_{it+s}^Y\mathbb{E}_t(\ln Y_{it+s} - \mathbb{E}_{t-2}\ln Y_{it+s}) = \zeta_{it} + \zeta_{it-1} + \vartheta_{it}^Y v_{it}$ because $\mathbb{E}_t\zeta_{it+\kappa} = \mathbb{E}_t v_{it+\kappa} = 0$ for $\kappa > 0$, and $\sum_{s=0}^{T-t}\vartheta_{it+s}^Y = 1$ by construction. The second-order term in income is more involved. This is given by

$$
\begin{aligned}
&\sum_{s=0}^{T-t}\sum_{\ell=0}^{T-t}(1-\pi_{it})\vartheta^Y_{it+s}(\delta_{s\ell}-(1-\pi_{it})\vartheta^Y_{it+\ell})\mathbb{E}_t(\ln Y_{it+s}-\mathbb{E}_{t-2}\ln Y_{it+s})(\ln Y_{it+\ell}-\mathbb{E}_{t-2}\ln Y_{it+\ell})\\
&=\sum_{s=0}^{T-t}\sum_{\ell=0}^{T-t}(1-\pi_{it})\vartheta^Y_{it+s}(\delta_{s\ell}-(1-\pi_{it})\vartheta^Y_{it+\ell})\mathbb{E}_t\left((\zeta_{it}+\zeta_{it-1})^2+(\zeta_{it}+\zeta_{it-1})(v_{it+s}+v_{it+\ell})+v_{it+s}v_{it+\ell}\right)\\
&=(1-\pi_{it})\pi_{it}(\zeta_{it}+\zeta_{it-1})^2+2\pi_{it}(1-\pi_{it})\vartheta^Y_{it}(\zeta_{it}+\zeta_{it-1})v_{it}+(1-\pi_{it})\vartheta^Y_{it}(1-(1-\pi_{it})\vartheta^Y_{it})v^2_{it}.
\end{aligned}
$$

We derive this making use of $\sum_{s=0}^{T-t}\vartheta^Y_{it+s}=1$, $\sum_{s=0}^{T-t}\sum_{\ell=0}^{T-t}\vartheta^Y_{it+s}\vartheta^Y_{it+\ell}=(\sum_{s=0}^{T-t}\vartheta^Y_{it+s})(\sum_{\ell=0}^{T-t}\vartheta^Y_{it+\ell})=1$, and the expectation of future shocks being zero by construction.

Plugging these results in the approximation of the previous page, noting that assets at t are 'beginning-of-period' and thus independent of the time t income shocks, and taking a first difference in expectations between $t-2$ and t yields

$$
\begin{aligned}
\mathbb{E}_t\ln\left(A_{it}+\sum_{s=0}^{T-t}\frac{Y_{it+s}}{(1+r)^s}\right)-\mathbb{E}_{t-2}\ln\left(A_{it}+\sum_{s=0}^{T-t}\frac{Y_{it+s}}{(1+r)^s}\right)&\approx(1-\pi_{it})(\zeta_{it}+\zeta_{it-1})+(1-\pi_{it})\vartheta^Y_{it}v_{it}\\
&+\frac{1}{2}(1-\pi_{it})\pi_{it}(\zeta_{it}+\zeta_{it-1})^2\\
&+\frac{1}{2}(1-\pi_{it})\vartheta^Y_{it}(1-(1-\pi_{it})\vartheta^Y_{it})v^2_{it}\\
&+(1-\pi_{it})\vartheta^Y_{it}\pi_{it}(\zeta_{it}+\zeta_{it-1})v_{it}.
\end{aligned}
$$

The linear terms in the first line are the first-order approximation of BPP; the quadratic terms in the next lines are the refinement from the second-order approximation.

Assuming expectations away, the logarithm of the *right* hand side of the budget constraint (2) is $\ln\sum_{s=0}^{T-t}\exp(\ln C_{it+s}-s\ln(1+r))$. Adopting the notation of (A.2), we let $\xi_s=\ln C_{it+s}-s\ln(1+r)$ and $\xi^0_s=\mathbb{E}_{t-2}\ln C_{it+s}-s\ln(1+r)$. Letting the information set contain time t information, it follows that

$$
\begin{aligned}
\mathbb{E}_t\ln\sum_{s=0}^{T-t}\frac{C_{it+s}}{(1+r)^s}&\approx\ln\sum_{s=0}^{T-t}\exp(\mathbb{E}_{t-2}\ln C_{it+s}-s\ln(1+r))\\
&+\sum_{s=0}^{T-t}\vartheta^C_{it+s}\mathbb{E}_t(\ln C_{it+s}-\mathbb{E}_{t-2}\ln C_{it+s})\\
&+\frac{1}{2}\sum_{s=0}^{T-t}\sum_{\ell=0}^{T-t}\vartheta^C_{it+s}(\delta_{s\ell}-\vartheta^C_{it+\ell})\mathbb{E}_t(\ln C_{it+s}-\mathbb{E}_{t-2}\ln C_{it+s})(\ln C_{it+\ell}-\mathbb{E}_{t-2}\ln C_{it+\ell}),
\end{aligned}
$$

where $\vartheta^C_{it+s}=\exp\left(\mathbb{E}_{t-2}\ln C_{it+s}-s\ln(1+r)\right)/\sum_{\kappa=0}^{T-t}\exp\left(\mathbb{E}_{t-2}\ln C_{it+\kappa}-\kappa\ln(1+r)\right)$ is an annuitization factor equal to the expected share of time $t+s$ consumption in total lifetime consumption. ϑ^C_{it+s} pertains to expectations at $t-2$, so it is known at $t-2$ and later.

Using the linearized consumption function in (A.1) to replace $\ln C_{it+s}$ recursively, we reach $\ln C_{it+s} = \ln C_{it-2} + \sum_{\tau=-1}^{s} \Psi_{it+\tau} + \sum_{\tau=-1}^{s} \xi_{it+\tau} + \sum_{\tau=-1}^{s} \theta_{it+\tau-1}\epsilon_{it+\tau}$, where $\Psi_{it+\tau} = \theta_{it+\tau-1}\Delta(\mathbf{Z}'_{it+\tau}\boldsymbol{\alpha}) + \Xi_{it+\tau}$, therefore $\ln C_{it+s} - \mathbb{E}_{t-2}\ln C_{it+s} = \sum_{\tau=-1}^{s}\theta_{it+\tau-1}\epsilon_{it+\tau}$ because Ψ_{it} and ξ_{it} are non-stochastic by assumption and they cancel out in the first difference. It follows that $\sum_{s=0}^{T-t}\vartheta^C_{it+s}\mathbb{E}_t(\ln C_{it+s} - \mathbb{E}_{t-2}\ln C_{it+s}) = \theta_{it-1}\epsilon_{it} + \theta_{it-2}\epsilon_{it-1}$ because $\mathbb{E}_t\epsilon_{it+\kappa} = 0$ for $\kappa > 0$ and $\sum_{s=0}^{T-t}\vartheta^C_{it+s} = 1$ by construction. The second-order term is more involved; it is given by

$$\begin{aligned}&\sum_{s=0}^{T-t}\sum_{\ell=0}^{T-t}\vartheta^C_{it+s}(\delta_{s\ell} - \vartheta^C_{it+\ell})\mathbb{E}_t(\ln C_{it+s} - \mathbb{E}_{t-2}\ln C_{it+s})(\ln C_{it+\ell} - \mathbb{E}_{t-2}\ln C_{it+\ell})\\ &= \sum_{s=0}^{T-t}\sum_{\ell=0}^{T-t}\vartheta^C_{it+s}(\delta_{s\ell} - \vartheta^C_{it+\ell})(\theta_{it-1}\epsilon_{it} + \theta_{it-2}\epsilon_{it-1})^2 = 0\end{aligned}$$

We derive this by noting that future expectations errors are zero by assumption, and that $\sum_{s=0}^{T-t}\sum_{\ell=0}^{T-t}\vartheta^C_{it+s}(\delta_{s\ell} - \vartheta^C_{it+\ell}) = \sum_{s=0}^{T-t}\vartheta^C_{it+s}\delta_{ss} + \sum_{s=0}^{T-t}\sum_{\ell\neq s}\vartheta^C_{it+s}\delta_{s\ell} - (\sum_{s=0}^{T-t}\vartheta^C_{it+s})(\sum_{\ell=0}^{T-t}\vartheta^C_{it+\ell}) = 0$. Plugging these results in the approximation above and taking a first difference in expectations between $t-2$ and t yields

$$\mathbb{E}_t \ln\sum_{s=0}^{T-t}\frac{C_{it+s}}{(1+r)^s} - \mathbb{E}_{t-2}\ln\sum_{s=0}^{T-t}\frac{C_{it+s}}{(1+r)^s} \approx \theta_{it-1}\epsilon_{it} + \theta_{it-2}\epsilon_{it-1}.$$

A.3 Biennial consumption growth

Given the biennial nature of our data, it follows from (A.1) that *observed* consumption growth is $\Delta^2 c_{it+s} = \xi_{it+s} + \xi_{it+s-1} + \theta_{it+s-1}\epsilon_{it+s} + \theta_{it+s-2}\epsilon_{it+s-1}$. The notation $\Delta^2 x_t = \Delta x_t + \Delta x_{t-1} = x_t - x_{t-2}$ indicates a difference between t and $t-2$. Dropping the index s and replacing the expectations errors with the quadratic expression in income shocks (the budget must balance so we bring the two sides of the budget constraint together), we obtain the final expression for consumption growth with biennial survey data, given by

$$\begin{aligned}\Delta^2 c_{it} \approx\ &\xi_{it} + \xi_{it-1} + (1-\pi_{it})(\zeta_{it} + \zeta_{it-1}) + (1-\pi_{it})\vartheta^Y_{it}v_{it}\\ &+ \frac{1}{2}(1-\pi_{it})\pi_{it}(\zeta_{it}+\zeta_{it-1})^2 + \frac{1}{2}(1-\pi_{it})\vartheta^Y_{it}(1-(1-\pi_{it})\vartheta^Y_{it})v^2_{it}\\ &+ (1-\pi_{it})\vartheta^Y_{it}\pi_{it}(\zeta_{it}+\zeta_{it-1})v_{it}.\end{aligned} \tag{A.3}$$

The first line of (A.3) corresponds to the linear consumption function (4), adapted here for the biennial nature of the data and given by

$$\Delta^2 c_{it} = \xi_{it} + \xi_{it-1} + \phi_{it}^{(1)} \left(\zeta_{it} + \zeta_{it-1}\right) + \psi_{it}^{(1)} v_{it}.$$

The transmission parameters of income shocks are given by $\phi_{it}^{(1)} = 1 - \pi_{it}$ and $\psi_{it}^{(1)} = (1 - \pi_{it})\vartheta_{it}^{Y}$, where π_{it} and ϑ_{it}^{Y} were defined previously.

The quadratic terms in the next two lines of (A.3) reflect the refinement from the higher-order approximation. This higher-order approximation leads to the quadratic consumption function (5), adapted here for the biennial nature of the data and given by

$$\Delta^2 c_{it} = \xi_{it} + \xi_{it-1} + \phi_{it}^{(1)} \left(\zeta_{it} + \zeta_{it-1}\right) + \psi_{it}^{(1)} v_{it} + \phi_{it}^{(2)} \left(\zeta_{it} + \zeta_{it-1}\right)^2 + \psi_{it}^{(2)} v_{it}^2 + \omega_{it}^{(22)} \left(\zeta_{it} + \zeta_{it-1}\right) v_{it}.$$

The additional transmission parameters of income shocks are given by $\phi_{it}^{(2)} = \frac{1}{2}(1 - \pi_{it})\pi_{it}$, $\psi_{it}^{(2)} = \frac{1}{2}(1 - \pi_{it})\vartheta_{it}^{Y}(1 - (1 - \pi_{it})\vartheta_{it}^{Y})$, and $\omega_{it}^{(22)} = (1 - \pi_{it})\vartheta_{it}^{Y}\pi_{it}$.

B Identification details

This appendix provides detailed identification statements based on our use of the recent biennial PSID data. There are two points of departure from section 3 in the text.

First, we recast the income and consumption processes to reflect the biennial frequency of the modern PSID. Unexplained income growth is thus given by $\Delta^2 y_{it} = \Delta y_{it} + \Delta y_{it-1} = \zeta_{it} + \zeta_{it-1} + \Delta^2 v_{it}$, with Δy_{it} given by (3). The notation $\Delta^2 x_t = \Delta x_t + \Delta x_{t-1} = x_t - x_{t-2}$ indicates a first difference in variable x between t and $t-2$, i.e., the observed frequency.

Second, we allow for classical measurement error. We let observed income y_{it}^* and consumption c_{it}^* be the sum of their corresponding true value and measurement error, namely

$$y_{it}^* = y_{it} + u_{it}^y \qquad \text{and} \qquad c_{it}^* = c_{it} + u_{it}^c.$$

The moments of income measurement error are not separately identifiable from the moments of the transitory shock. This necessitates that we restrict the income error to be Gaussian (but we let the consumption error unrestricted), so we specify

$$\mathbb{E}((u_{it}^y)^m) = \begin{cases} 0 & \text{for } m = 1 \\ \sigma_{u_t^y}^2 & \text{for } m = 2 \\ 0 & \text{for } m = 3 \\ 3(\sigma_{u_t^y}^2)^2 & \text{for } m = 4 \end{cases} \qquad \text{and} \qquad \mathbb{E}((u_{i,t}^c)^m) = \begin{cases} 0 & \text{for } m = 1 \\ \sigma_{u_t^c}^2 & \text{for } m = 2 \\ \gamma_{u_t^c} & \text{for } m = 3 \\ \kappa_{u_t^c} & \text{for } m = 4. \end{cases}$$

We retrieve the variance of income error from the validation study of the PSID in Bound, Brown, Duncan, and Rodgers (1994), as we explain in section 4. We do not need similar distributional assumptions for the consumption error, whose moments we identify below. We further assume that income and consumption errors are mutually independent, independent over time, and independent of the income shocks and taste heterogeneity ξ.

These choices lead to our processes for observed income growth, consumption growth in the linear specification, and consumption growth in the quadratic specification, given by

$$\begin{aligned}
\Delta^2 y_{it} &= \zeta_{it} + \zeta_{it-1} + \Delta^2 v_{it} + \Delta^2 u^y_{it},\\
\Delta^2 c_{it} &= \xi_{it} + \xi_{it-1} + \phi^{(1)}_{it}(\zeta_{it} + \zeta_{it-1}) + \psi^{(1)}_{it} v_{it} + \Delta^2 u^c_{it},\\
\Delta^2 c_{it} &= \xi_{it} + \xi_{it-1} + \phi^{(1)}_{it}(\zeta_{it} + \zeta_{it-1}) + \psi^{(1)}_{it} v_{it}\\
&\quad + \phi^{(2)}_{it}(\zeta_{it} + \zeta_{it-1})^2 + \psi^{(2)}_{it} v^2_{it} + \omega^{(22)}_{it}(\zeta_{it} + \zeta_{it-1}) v_{it} + \Delta^2 u^c_{it},
\end{aligned}$$

respectively. These are the theoretical expressions we bring to the data.

B.1 Income process parameters

Due to the biennial nature of the data, we cannot separately identify the moments of the *yearly* permanent shocks that comprise a two-year period. We thus assume $\sigma^2_{\zeta_t} \approx \sigma^2_{\zeta_{t-1}}$ and $\gamma_{\zeta_t} \approx \gamma_{\zeta_{t-1}}$, which is exactly true if the underlying distributions are stationary. The second and third moments of shocks are then given by

$$\sigma^2_{\zeta_t} = \frac{1}{2}\mathbb{E}(\Delta^2 y_{it} \times \sum_{\kappa=\{-2,0,2\}} \Delta^2 y_{it+\kappa}) \qquad \text{and} \qquad \sigma^2_{v_t} = -\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it+2}) - \sigma^2_{u^y_t},$$

$$\gamma_{\zeta_t} = \frac{1}{2}\mathbb{E}((\Delta^2 y_{it})^2 \times \sum_{\kappa=\{-2,0,2\}} \Delta^2 y_{it+\kappa}) \qquad \text{and} \qquad \gamma_{v_t} = -\mathbb{E}((\Delta^2 y_{it})^2 \times \Delta^2 y_{it+2}).$$

Identification of the fourth moments is more involved; to simplify the statements, we further assume that $\sigma^2_{\zeta_t} \approx \sigma^2_{\zeta_{t+1}} \approx \sigma^2_{\zeta_{t+2}}$, $\sigma^2_{v_t} \approx \sigma^2_{v_{t-2}} \approx \sigma^2_{v_{t+2}}$, and $\sigma^2_{u^y_t} \approx \sigma^2_{u^y_{t-2}} \approx \sigma^2_{u^y_{t+2}}$. These restrictions are *not* needed for identification but they simplify the illustration. The fourth moments are then given by

$$\begin{aligned}
\kappa_{\zeta_t} &= \frac{1}{2}\mathbb{E}((\Delta^2 y_{it})^4) - \kappa_{v_t} - 3(\sigma^2_{\zeta_t})^2 - 3(\sigma^2_{v_t})^2 - 6(\sigma^2_{u^y_t})^2\\
&\quad - 12\sigma^2_{\zeta_t}\sigma^2_{v_t} - 12\sigma^2_{\zeta_t}\sigma^2_{u^y_t} - 12\sigma^2_{v_t}\sigma^2_{u^y_t},\\
\kappa_{v_t} &= \mathbb{E}((\Delta^2 y_{it})^2 \times (\Delta^2 y_{it+2})^2) - 4(\sigma^2_{\zeta_t})^2 - 3(\sigma^2_{v_t})^2 - 6(\sigma^2_{u^y_t})^2\\
&\quad - 8\sigma^2_{\zeta_t}\sigma^2_{v_t} - 8\sigma^2_{\zeta_t}\sigma^2_{u^y_t} - 12\sigma^2_{v_t}\sigma^2_{u^y_t}.
\end{aligned}$$

B.2 Linear consumption function parameters

Given the biennial data, we must assume $\sigma^2_{\zeta_t} \approx \sigma^2_{\zeta_{t-1}}$ and $\sigma^2_{\xi_t} \approx \sigma^2_{\xi_{t-1}}$; both are exactly true if the distributions are stationary. Given the income process, the linear consumption function, and the properties of shocks, the transmission parameters are given by

$$\phi^{(1)}_t = \frac{1}{2}\mathbb{E}(\Delta^2 c_{it} \times \sum_{\kappa=\{-2,0,2\}} \Delta^2 y_{it+\kappa})/\sigma^2_{\zeta_t} \qquad \text{and} \qquad \psi^{(1)}_t = -\mathbb{E}(\Delta^2 c_{it} \times \Delta^2 y_{it+2})/\sigma^2_{v_t}.$$

The variance of taste heterogeneity and consumption measurement error are identified as

$$\sigma^2_{\xi_t} = \frac{1}{2}\mathbb{E}(\Delta^2 c_{it} \times \sum_{\kappa=\{-2,0,2\}} \Delta^2 c_{it+\kappa}) - (\phi^{(1)}_t)^2\sigma^2_{\zeta_t} - \frac{1}{2}(\psi^{(1)}_t)^2\sigma^2_{v_t} \quad \text{and} \quad \sigma^2_{u^c_t} = -\mathbb{E}(\Delta^2 c_{it} \times \Delta^2 c_{it+2}).$$

The third and fourth moments of taste heterogeneity and consumption error are identified by higher-order consumption moments through expressions analogous to those above.

B.3 Quadratic consumption function parameters

Given the biennial data, we must assume $\sigma^2_{\zeta_t} \approx \sigma^2_{\zeta_{t-1}}$, $\gamma_{\zeta_t} \approx \gamma_{\zeta_{t-1}}$, $\kappa_{\zeta_t} \approx \kappa_{\zeta_{t-1}}$, and $\sigma^2_{\xi_t} \approx \sigma^2_{\xi_{t-1}}$. For simplicity, we also assume that $\sigma^2_{\zeta_t} \approx \sigma^2_{\zeta_{t+1}} \approx \sigma^2_{\zeta_{t+2}}$, $\sigma^2_{v_t} \approx \sigma^2_{v_{t-2}}$, and $\sigma^2_{u^y_t} \approx \sigma^2_{u^y_{t-2}}$; the latter restrictions are *not* needed for identification. Given the income process, the quadratic consumption function, and the properties of income shocks, the transmission parameters are given by

$$\begin{pmatrix} \phi^{(1)}_t \\ \psi^{(1)}_t \\ \phi^{(2)}_t \\ \psi^{(2)}_t \\ \omega^{(22)}_t \end{pmatrix} = \mathbf{A}^{-1} \begin{pmatrix} \mathbb{E}(\Delta^2 c_{it} \times \Delta^2 y_{it}) \\ \mathbb{E}(\Delta^2 c_{it} \times (\Delta^2 y_{it})^2) \\ \mathbb{E}(\Delta^2 c_{it} \times \Delta^2 y_{it+2}) \\ \mathbb{E}(\Delta^2 c_{it} \times (\Delta^2 y_{it+2})^2) \\ \mathbb{E}(\Delta^2 c_{it} \times \Delta^2 y_{it} \times \Delta^2 y_{it+2}) \end{pmatrix},$$

where the matrix of coefficients $\mathbf{A}$ depends exclusively on the second, third, and fourth moments of income shocks, and it is generally full rank and invertible.[1] This can be easily

[1]The matrix of coefficients with *biennial* data (different from matrix $\mathbf{A}$ in the main text) is given by

$$\mathbf{A} = \begin{pmatrix} 2\sigma^2_{\zeta_t} & \sigma^2_{v_t} & 2\gamma_{\zeta_t} & \gamma_{v_t} & 0 \\ 2\gamma_{\zeta_t} & \gamma_{v_t} & 2(\kappa_{\zeta_t} + \sigma^2_{\zeta_t}(2\bar{\sigma}_t + \sigma^2_{\zeta_t})) & \kappa_{v_t} + \sigma^2_{v_t}(2\bar{\sigma}_t - \sigma^2_{v_t}) & 4\sigma^2_{\zeta_t}\sigma^2_{v_t} \\ 0 & -\sigma^2_{v_t} & 0 & -\gamma_{v_t} & 0 \\ 0 & \gamma_{v_t} & 4\sigma^2_{\zeta_t}\bar{\sigma}_t & \kappa_{v_t} + \sigma^2_{v_t}(2\bar{\sigma}_t - \sigma^2_{v_t}) & 0 \\ 0 & -\gamma_{v_t} & -2\sigma^2_{\zeta_t}(\bar{\sigma}_t - \sigma^2_{\zeta_t}) & -\kappa_{v_t} - \sigma^2_{v_t}\sigma^2_{u^y_t} & -2\sigma^2_{\zeta_t}\sigma^2_{v_t} \end{pmatrix},$$

where $\bar{\sigma}_t = \sigma^2_{\zeta_t} + \sigma^2_{v_t} + \sigma^2_{u^y_t}$. There are no obvious linear dependencies across the columns of $\mathbf{A}$, which is true also if income shocks are Gaussian, so the matrix is generally full rank and invertible.

Table B.1: Targeted moments of income

Moments:	2^{nd}	2^{nd}, 3^{rd}	2^{nd}, 3^{rd}, 4^{th}
1	$\mathbb{E}((\Delta^2 y_{it})^2)$	$\mathbb{E}((\Delta^2 y_{it})^2)$	$\mathbb{E}((\Delta^2 y_{it})^2)$
2	$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it-2})$	$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it-2})$	$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it-2})$
3		$\mathbb{E}((\Delta^2 y_{it})^3)$	$\mathbb{E}((\Delta^2 y_{it})^3)$
4		$\mathbb{E}((\Delta^2 y_{it})^2 \times \Delta^2 y_{it-2})$	$\mathbb{E}((\Delta^2 y_{it})^2 \times \Delta^2 y_{it-2})$
5		$\mathbb{E}(\Delta^2 y_{it} \times (\Delta^2 y_{it-2})^2)$	$\mathbb{E}(\Delta^2 y_{it} \times (\Delta^2 y_{it-2})^2)$
6			$\mathbb{E}((\Delta^2 y_{it})^4)$
7			$\mathbb{E}((\Delta^2 y_{it})^2 \times (\Delta^2 y_{it-2})^2)$
8			$\mathbb{E}((\Delta^2 y_{it})^3 \times \Delta^2 y_{it-2})$
9			$\mathbb{E}(\Delta^2 y_{it} \times (\Delta^2 y_{it-2})^3)$

Notes: The table lists the moments targeted in the GMM estimation of the various specifications of the income process. The notation $\Delta^2 y_t = \Delta y_t + \Delta y_{t-1} = y_t - y_{t-2}$ indicates the first difference of residual log income y between periods t and $t-2$.

checked as soon as the income process is estimated. Upon identification of the transmission parameters, $\mathbb{E}(\Delta^2 c_{it} \times \Delta^2 c_{it+2})$ identifies the variance of consumption error, while $\mathbb{E}((\Delta^2 c_{it})^2)$ identifies the variance of ξ. Higher moments of the consumption error or of ξ require higher than fourth-order moments of income, which we do not model.

B.4 Targeted moments

There is a large number of over-identifying moments that we target in the estimation of the various specifications of the model. Table B.1 lists the moments targeted in the various specifications of the income process while table B.2 lists the moments targeted in the different specifications of the consumption model. For simplicity in deriving these moments, we assume stationarity in the distributions of income shocks, taste heterogeneity, and measurement error, and time-invariance of the partial insurance parameters. Stationarity and time-invariance, while not needed for identification, simplify the following statements considerably as they allow us to remove the time subscript and bundle common terms together.

B.4.1 Moments of income

$$\begin{aligned}
\mathbb{E}((\Delta^2 y_{it})^2) &= 2\sigma_\zeta^2 + 2\sigma_v^2 + 2\sigma_{u^y}^2 \\
\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it-2}) &= -\sigma_v^2 - \sigma_{u^y}^2 \\
\mathbb{E}((\Delta^2 y_{it})^3) &= 2\gamma_\zeta \\
\mathbb{E}((\Delta^2 y_{it})^2 \times \Delta^2 y_{it-2}) &= \gamma_v
\end{aligned}$$

$$\mathbb{E}(\Delta^2 y_{it} \times (\Delta^2 y_{it-2})^2) = -\gamma_v$$
$$\mathbb{E}((\Delta^2 y_{it})^4) = 2\kappa_\zeta + 2\kappa_v + 6(\sigma_\zeta^2)^2 + 6(\sigma_v^2)^2 + 12(\sigma_{u^y}^2)^2 + 24\sigma_\zeta^2\sigma_v^2 + 24\sigma_v^2\sigma_{u^y}^2 + 24\sigma_\zeta^2\sigma_{u^y}^2$$
$$\mathbb{E}((\Delta^2 y_{it})^2 \times (\Delta^2 y_{it-2})^2) = \kappa_v + 4(\sigma_\zeta^2)^2 + 3(\sigma_v^2)^2 + 6(\sigma_{u^y}^2)^2 + 8\sigma_\zeta^2\sigma_v^2 + 8\sigma_\zeta^2\sigma_{u^y}^2 + 12\sigma_v^2\sigma_{u^y}^2$$
$$\mathbb{E}((\Delta^2 y_{it})^3 \times \Delta^2 y_{it-2}) = -\kappa_v - 3(\sigma_v^2)^2 - 6(\sigma_{u^y}^2)^2 - 6\sigma_\zeta^2\sigma_v^2 - 6\sigma_\zeta^2\sigma_{u^y}^2 - 12\sigma_v^2\sigma_{u^y}^2$$
$$\mathbb{E}(\Delta^2 y_{it} \times (\Delta^2 y_{it-2})^3) = -\kappa_v - 3(\sigma_v^2)^2 - 6(\sigma_{u^y}^2)^2 - 6\sigma_\zeta^2\sigma_v^2 - 6\sigma_\zeta^2\sigma_{u^y}^2 - 12\sigma_v^2\sigma_{u^y}^2$$

B.4.2 Moments of consumption – linear function

$$\mathbb{E}((\Delta^2 c_{it})^2) = 2\sigma_\xi^2 + 2(\phi^{(1)})^2\sigma_\zeta^2 + (\psi^{(1)})^2\sigma_v^2 + 2\sigma_{u^c}^2$$
$$\mathbb{E}(\Delta^2 c_{it} \times \Delta^2 c_{it-2}) = -\sigma_{u^c}^2$$
$$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 c_{it}) = 2\phi^{(1)}\sigma_\zeta^2 + \psi^{(1)}\sigma_v^2$$
$$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 c_{it-2}) = -\psi^{(1)}\sigma_v^2$$
$$\mathbb{E}((\Delta^2 c_{it})^3) = 2\gamma_\xi + 2(\phi^{(1)})^3\gamma_\zeta + (\psi^{(1)})^3\gamma_v$$
$$\mathbb{E}((\Delta^2 c_{it})^2 \times \Delta^2 c_{it-2}) = \gamma_{u^c}$$
$$\mathbb{E}(\Delta^2 c_{it} \times (\Delta^2 c_{it-2})^2) = -\gamma_{u^c}$$
$$\mathbb{E}((\Delta^2 y_{it})^2 \times \Delta^2 c_{it}) = 2\phi^{(1)}\gamma_\zeta + \psi^{(1)}\gamma_v$$
$$\mathbb{E}(\Delta^2 y_{it} \times (\Delta^2 c_{it})^2) = 2(\phi^{(1)})^2\gamma_\zeta + (\psi^{(1)})^2\gamma_v$$
$$\mathbb{E}((\Delta^2 y_{it})^2 \times \Delta^2 c_{it-2}) = \psi^{(1)}\gamma_v$$
$$\mathbb{E}(\Delta^2 y_{it} \times (\Delta^2 c_{it-2})^2) = -(\psi^{(1)})^2\gamma_v$$
$$\mathbb{E}((\Delta^2 y_{it-2})^2 \times \Delta^2 c_{it}) = 0$$
$$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it+2} \times \Delta^2 c_{it}) = -\psi^{(1)}\gamma_v$$
$$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it+2} \times \Delta^2 c_{it-2}) = 0$$
$$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it+2} \times \Delta^2 c_{it+2}) = 0$$
$$\begin{aligned}\mathbb{E}((\Delta^2 c_{it})^4) = {} & 2\kappa_\xi + 2(\phi^{(1)})^4\kappa_\zeta + (\psi^{(1)})^4\kappa_v + 2\kappa_{u^c} + 6(\sigma_\xi^2)^2 + 6(\phi^{(1)})^4(\sigma_\zeta^2)^2 + 6(\sigma_{u^c}^2)^2 \\ & + 24(\phi^{(1)})^2\sigma_\xi^2\sigma_\zeta^2 + 12(\psi^{(1)})^2\sigma_\xi^2\sigma_v^2 + 24\sigma_\xi^2\sigma_{u^c}^2 \\ & + 12(\phi^{(1)})^2(\psi^{(1)})^2\sigma_\zeta^2\sigma_v^2 + 24(\phi^{(1)})^2\sigma_\zeta^2\sigma_{u^c}^2 + 12(\psi^{(1)})^2\sigma_v^2\sigma_{u^c}^2\end{aligned}$$
$$\begin{aligned}\mathbb{E}((\Delta^2 c_{it})^2 \times (\Delta^2 c_{it-2})^2) = {} & 4(\sigma_\xi^2)^2 + 8(\phi^{(1)})^2\sigma_\xi^2\sigma_\zeta^2 + 4(\psi^{(1)})^2\sigma_\xi^2\sigma_v^2 + 8\sigma_\xi^2\sigma_{u^c}^2 + 4(\phi^{(1)})^4(\sigma_\zeta^2)^2 \\ & + 4(\phi^{(1)})^2(\psi^{(1)})^2\sigma_\zeta^2\sigma_v^2 + 8(\phi^{(1)})^2\sigma_\zeta^2\sigma_{u^c}^2 + (\psi^{(1)})^4(\sigma_v^2)^2 + 4(\psi^{(1)})^2\sigma_v^2\sigma_{u^c}^2 \\ & + \kappa_{u^c} + 3(\sigma_{u^c}^2)^2\end{aligned}$$
$$\mathbb{E}((\Delta^2 c_{it})^3 \times \Delta^2 c_{it-2}) = -\kappa_{u^c} - 3(\sigma_{u^c}^2)^2 - 6\sigma_\xi^2\sigma_{u^c}^2 - 6(\phi^{(1)})^2\sigma_\zeta^2\sigma_{u^c}^2 - 3(\psi^{(1)})^2\sigma_v^2\sigma_{u^c}^2$$
$$\mathbb{E}(\Delta^2 c_{it} \times (\Delta^2 c_{it-2})^3) = -\kappa_{u^c} - 3(\sigma_{u^c}^2)^2 - 6\sigma_\xi^2\sigma_{u^c}^2 - 6(\phi^{(1)})^2\sigma_\zeta^2\sigma_{u^c}^2 - 3(\psi^{(1)})^2\sigma_v^2\sigma_{u^c}^2$$

Table B.2: Targeted joint moments of consumption and income

Cons. fn.:	**Linear**			**Quadratic**
Moments:	2nd	2nd, 3rd	2nd, 3rd, 4th	2nd, 3rd, 4th (income shocks)
1	$\mathbb{E}((\Delta^2 c_{it})^2)$	$\mathbb{E}((\Delta^2 c_{it})^2)$	$\mathbb{E}((\Delta^2 c_{it})^2)$	$\mathbb{E}((\Delta^2 c_{it})^2)$
2	$\mathbb{E}(\Delta^2 c_{it} \times \Delta^2 c_{it-2})$	$\mathbb{E}(\Delta^2 c_{it} \times \Delta^2 c_{it-2})$	$\mathbb{E}(\Delta^2 c_{it} \times \Delta^2 c_{it-2})$	$\mathbb{E}(\Delta^2 c_{it} \times \Delta^2 c_{it-2})$
3	$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 c_{it})$	$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 c_{it})$	$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 c_{it})$	$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 c_{it})$
4	$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 c_{it-2})$	$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 c_{it-2})$	$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 c_{it-2})$	$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 c_{it-2})$
5		$\mathbb{E}((\Delta^2 c_{it})^3)$	$\mathbb{E}((\Delta^2 c_{it})^3)$	
6		$\mathbb{E}((\Delta^2 c_{it})^2 \times \Delta^2 c_{it-2})$	$\mathbb{E}((\Delta^2 c_{it})^2 \times \Delta^2 c_{it-2})$	
7		$\mathbb{E}(\Delta^2 c_{it} \times (\Delta^2 c_{it-2})^2)$	$\mathbb{E}(\Delta^2 c_{it} \times (\Delta^2 c_{it-2})^2)$	
8		$\mathbb{E}((\Delta^2 y_{it})^2 \times \Delta^2 c_{it})$	$\mathbb{E}((\Delta^2 y_{it})^2 \times \Delta^2 c_{it})$	$\mathbb{E}((\Delta^2 y_{it})^2 \times \Delta^2 c_{it})$
9		$\mathbb{E}(\Delta^2 y_{it} \times (\Delta^2 c_{it})^2)$	$\mathbb{E}(\Delta^2 y_{it} \times (\Delta^2 c_{it})^2)$	
10		$\mathbb{E}((\Delta^2 y_{it})^2 \times \Delta^2 c_{it-2})$	$\mathbb{E}((\Delta^2 y_{it})^2 \times \Delta^2 c_{it-2})$	$\mathbb{E}((\Delta^2 y_{it})^2 \times \Delta^2 c_{it-2})$
11		$\mathbb{E}(\Delta^2 y_{it} \times (\Delta c_{it-2})^2)$	$\mathbb{E}(\Delta^2 y_{it} \times (\Delta^2 c_{it-2})^2)$	
12		$\mathbb{E}((\Delta^2 y_{it-2})^2 \times \Delta^2 c_{it})$	$\mathbb{E}((\Delta^2 y_{it-2})^2 \times \Delta^2 c_{it})$	$\mathbb{E}((\Delta^2 y_{it-2})^2 \times \Delta^2 c_{it})$
13		$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it+2} \times \Delta^2 c_{it})$	$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it+2} \times \Delta^2 c_{it})$	$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it+2} \times \Delta^2 c_{it})$
14		$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it+2} \times \Delta^2 c_{it-2})$	$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it+2} \times \Delta^2 c_{it-2})$	$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it+2} \times \Delta^2 c_{it-2})$
15		$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it+2} \times \Delta^2 c_{it+2})$	$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it+2} \times \Delta^2 c_{it+2})$	$\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it+2} \times \Delta^2 c_{it+2})$
16			$\mathbb{E}((\Delta^2 c_{it})^4)$	
17			$\mathbb{E}((\Delta^2 c_{it})^2 \times (\Delta^2 c_{it-2})^2)$	
18			$\mathbb{E}((\Delta^2 c_{it})^3 \times \Delta^2 c_{it-2})$	
19			$\mathbb{E}(\Delta^2 c_{it} \times (\Delta^2 c_{it-2})^3)$	
20			$\mathbb{E}((\Delta^2 y_{it})^2 \times (\Delta^2 c_{it})^2)$	
21			$\mathbb{E}((\Delta^2 y_{it})^3 \times \Delta^2 c_{it})$	
22			$\mathbb{E}(\Delta^2 y_{it} \times (\Delta^2 c_{it})^3)$	
23			$\mathbb{E}((\Delta^2 y_{it})^2 \times (\Delta^2 c_{it-2})^2)$	
24			$\mathbb{E}((\Delta^2 y_{it-2})^2 \times (\Delta^2 c_{it})^2)$	
25			$\mathbb{E}((\Delta^2 y_{it})^3 \times \Delta^2 c_{it-2})$	
26			$\mathbb{E}(\Delta^2 y_{it} \times (\Delta^2 c_{it-2})^3)$	

Notes: The table lists the moments targeted in the GMM estimation of the different specifications of the consumption model. In all cases, the notation $\Delta^2 x_t = \Delta x_t + \Delta x_{t-1} = x_t - x_{t-2}$ indicates the first difference of the generic variable x between periods t and $t-2$.

$$
\begin{aligned}
\mathbb{E}((\Delta^2 y_{it})^2 \times (\Delta^2 c_{it})^2) &= 4\sigma_\zeta^2\sigma_\xi^2 + 2(\phi^{(1)})^2\kappa_\zeta + 6(\phi^{(1)})^2(\sigma_\zeta^2)^2 + 2(\psi^{(1)})^2\sigma_\zeta^2\sigma_v^2 + 4\sigma_\zeta^2\sigma_{u^c}^2 \\
&+ 4\sigma_v^2\sigma_\xi^2 + 4(\phi^{(1)})^2\sigma_v^2\sigma_\zeta^2 + (\psi^{(1)})^2\kappa_v + (\psi^{(1)})^2(\sigma_v^2)^2 + 4\sigma_v^2\sigma_{u^c}^2 \\
&+ 4\sigma_{u^y}^2\sigma_\xi^2 + 4(\phi^{(1)})^2\sigma_{u^y}^2\sigma_\zeta^2 + 2(\psi^{(1)})^2\sigma_{u^y}^2\sigma_v^2 + 4\sigma_{u^y}^2\sigma_{u^c}^2 + 8\phi^{(1)}\psi^{(1)}\sigma_\zeta^2\sigma_v^2 \\
\mathbb{E}((\Delta^2 y_{it})^3 \times \Delta^2 c_{it}) &= 2\phi^{(1)}\kappa_\zeta + 6\phi^{(1)}(\sigma_\zeta^2)^2 + \psi^{(1)}\kappa_v + 3\psi^{(1)}(\sigma_v^2)^2 \\
&+ 6\psi^{(1)}\sigma_\zeta^2\sigma_v^2 + 12\phi^{(1)}\sigma_\zeta^2\sigma_v^2 + 12\phi^{(1)}\sigma_\zeta^2\sigma_{u^y}^2 + 6\psi^{(1)}\sigma_v^2\sigma_{u^y}^2 \\
\mathbb{E}(\Delta^2 y_{it} \times (\Delta^2 c_{it})^3) &= 2(\phi^{(1)})^3\kappa_\zeta + 6(\phi^{(1)})^3(\sigma_\zeta^2)^2 + (\psi^{(1)})^3\kappa_v + 12\phi^{(1)}\sigma_\xi^2\sigma_\zeta^2 + 6\psi^{(1)}\sigma_\xi^2\sigma_v^2 \\
&+ 6\phi^{(1)}(\psi^{(1)})^2\sigma_\zeta^2\sigma_v^2 + 6(\phi^{(1)})^2\psi^{(1)}\sigma_\zeta^2\sigma_v^2 + 12\phi^{(1)}\sigma_\zeta^2\sigma_{u^c}^2 + 6\psi^{(1)}\sigma_v^2\sigma_{u^c}^2 \\
\mathbb{E}((\Delta^2 y_{it})^2 \times (\Delta^2 c_{it-2})^2) &= 4\sigma_\zeta^2\sigma_\xi^2 + 4(\phi^{(1)})^2(\sigma_\zeta^2)^2 + 2(\psi^{(1)})^2\sigma_\zeta^2\sigma_v^2 + 4\sigma_\zeta^2\sigma_{u^c}^2 \\
&+ 4\sigma_v^2\sigma_\xi^2 + 4(\phi^{(1)})^2\sigma_v^2\sigma_\zeta^2 + (\psi^{(1)})^2\kappa_v + (\psi^{(1)})^2(\sigma_v^2)^2 + 4\sigma_v^2\sigma_{u^c}^2 \\
&+ 4\sigma_{u^y}^2\sigma_\xi^2 + 4(\phi^{(1)})^2\sigma_{u^y}^2\sigma_\zeta^2 + 2(\psi^{(1)})^2\sigma_{u^y}^2\sigma_v^2 + 4\sigma_{u^y}^2\sigma_{u^c}^2 \\
\mathbb{E}((\Delta^2 y_{it-2})^2 \times (\Delta^2 c_{it})^2) &= 4\sigma_\zeta^2\sigma_\xi^2 + 4(\phi^{(1)})^2(\sigma_\zeta^2)^2 + 2(\psi^{(1)})^2\sigma_\zeta^2\sigma_v^2 + 4\sigma_\zeta^2\sigma_{u^c}^2 \\
&+ 4\sigma_v^2\sigma_\xi^2 + 4(\phi^{(1)})^2\sigma_v^2\sigma_\zeta^2 + 2(\psi^{(1)})^2(\sigma_v^2)^2 + 4\sigma_v^2\sigma_{u^c}^2 \\
&+ 4\sigma_{u^y}^2\sigma_\xi^2 + 4(\phi^{(1)})^2\sigma_{u^y}^2\sigma_\zeta^2 + 2(\psi^{(1)})^2\sigma_{u^y}^2\sigma_v^2 + 4\sigma_{u^y}^2\sigma_{u^c}^2 \\
\mathbb{E}((\Delta^2 y_{it})^3 \times \Delta^2 c_{it-2}) &= -\psi^{(1)}\kappa_v - 3\psi^{(1)}(\sigma_v^2)^2 - 6\psi^{(1)}\sigma_\zeta^2\sigma_v^2 - 6\psi^{(1)}\sigma_v^2\sigma_{u^y}^2 \\
\mathbb{E}(\Delta^2 y_{it} \times (\Delta^2 c_{it-2})^3) &= -(\psi^{(1)})^3\kappa_v - 6\psi^{(1)}\sigma_v^2\sigma_\xi^2 - 6(\phi^{(1)})^2\psi^{(1)}\sigma_v^2\sigma_\zeta^2 - 6\psi^{(1)}\sigma_v^2\sigma_{u^c}^2
\end{aligned}
$$

B.4.3 Moments of consumption – quadratic function

$$
\begin{aligned}
\mathbb{E}((\Delta^2 c_{it})^2) &= 2\sigma_\xi^2 + 2(\phi^{(1)})^2\sigma_\zeta^2 + (\psi^{(1)})^2\sigma_v^2 \\
&+ 2(\phi^{(2)})^2(\kappa_\zeta + 3(\sigma_\zeta^2)^2) + (\psi^{(2)})^2\kappa_v + 2(\omega^{(22)})^2\sigma_\zeta^2\sigma_v^2 + 2\sigma_{u^c}^2 \\
&+ 4\phi^{(1)}\phi^{(2)}\gamma_\zeta + 2\psi^{(1)}\psi^{(2)}\gamma_v + 4\phi^{(2)}\psi^{(2)}\sigma_\zeta^2\sigma_v^2 \\
\mathbb{E}(\Delta^2 c_{it} \times \Delta^2 c_{it-2}) &= 4(\phi^{(2)})^2(\sigma_\zeta^2)^2 + 4\phi^{(2)}\psi^{(2)}\sigma_\zeta^2\sigma_v^2 + (\psi^{(2)})^2(\sigma_v^2)^2 - \sigma_{u^c}^2 \\
\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 c_{it}) &= 2\phi^{(1)}\sigma_\zeta^2 + \psi^{(1)}\sigma_v^2 + 2\phi^{(2)}\gamma_\zeta + \psi^{(2)}\gamma_v \\
\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 c_{it-2}) &= -\psi^{(1)}\sigma_v^2 - \psi^{(2)}\gamma_v \\
\mathbb{E}((\Delta^2 y_{it})^2 \times \Delta^2 c_{it}) &= 2\phi^{(1)}\gamma_\zeta + \psi^{(1)}\gamma_v + 2\phi^{(2)}(\kappa_\zeta + 3(\sigma_\zeta^2)^2 + 2\sigma_\zeta^2\sigma_v^2 + 2\sigma_\zeta^2\sigma_{u^y}^2) \\
&+ \psi^{(2)}(\kappa_v + (\sigma_v^2)^2 + 2\sigma_\zeta^2\sigma_v^2 + 2\sigma_v^2\sigma_{u^y}^2) + 4\omega^{(22)}\sigma_\zeta^2\sigma_v^2 \\
\mathbb{E}((\Delta^2 y_{it})^2 \times \Delta^2 c_{it-2}) &= \psi^{(1)}\gamma_v + 4\phi^{(2)}\sigma_\zeta^2(\sigma_\zeta^2 + \sigma_v^2 + \sigma_{u^y}^2) + \psi^{(2)}(\kappa_v + \sigma_v^2(2\sigma_\zeta^2 + \sigma_v^2 + 2\sigma_{u^y}^2)) \\
\mathbb{E}((\Delta^2 y_{it})^2 \times \Delta^2 c_{it+2}) &= 4\phi^{(2)}\sigma_\zeta^2(\sigma_\zeta^2 + \sigma_v^2 + \sigma_{u^y}^2) + 2\psi^{(2)}\sigma_v^2(\sigma_\zeta^2 + \sigma_v^2 + \sigma_{u^y}^2) \\
\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it+2} \times \Delta^2 c_{it}) &= -\psi^{(1)}\gamma_v - 2\phi^{(2)}\sigma_\zeta^2(\sigma_v^2 + \sigma_{u^y}^2) - \psi^{(2)}(\kappa_v + \sigma_v^2\sigma_{u^y}^2) - 2\omega^{(22)}\sigma_\zeta^2\sigma_v^2 \\
\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it+2} \times \Delta^2 c_{it-2}) &= -2\phi^{(2)}\sigma_\zeta^2(\sigma_v^2 + \sigma_{u^y}^2) - \psi^{(2)}\sigma_v^2(\sigma_v^2 + \sigma_{u^y}^2) \\
\mathbb{E}(\Delta^2 y_{it} \times \Delta^2 y_{it+2} \times \Delta^2 c_{it+2}) &= -2\phi^{(2)}\sigma_\zeta^2(\sigma_v^2 + \sigma_{u^y}^2) - \psi^{(2)}\sigma_v^2(\sigma_v^2 + \sigma_{u^y}^2)
\end{aligned}
$$

C Empirical details and additional results

C.1 Comparison with administrative moments

Table C.1 reports the empirical second, third, and fourth moments of the cross-sectional distribution of unexplained growth in male earnings (panel A) and household earnings (panel B), as opposed to the moments of household disposable income in table 2 in the main text. Column 1 presents moments from our baseline sample in the PSID, while column 2 presents the corresponding statistics, wherever available, from the U.S. Social Security Administration data reported in Guvenen, Karahan, Ozkan, and Song (2021).

Table C.1: Empirical moments, male earnings and household earnings

	Baseline sample		Guvenen et al. (2021)
Income data:	PSID		Social Security Admin.
	(1)		(2)
Panel A. Moments of male earnings growth			
$\text{Var}(\Delta y_{it})$	0.271	(0.010)	0.323
$\text{Cov}(\Delta y_{it}, \Delta y_{it+1})$	-0.065	(0.004)	–
$\text{Skew}(\Delta y_{it})$	-0.569	(0.128)	-1.039
$\text{Cov}((\Delta y_{it})^2, \Delta y_{it+1})$	0.087	(0.010)	–
$\text{Kurt}(\Delta y_{it})$	13.711	(0.551)	13.494
$\text{Cov}((\Delta y_{it})^2, (\Delta y_{it+1})^2)$	0.253	(0.022)	–
Panel B. Moments of household earnings growth			
$\text{Var}(\Delta y_{it})$	0.233	(0.009)	–
$\text{Cov}(\Delta y_{it}, \Delta y_{it+1})$	-0.056	(0.004)	–
$\text{Skew}(\Delta y_{it})$	-0.728	(0.142)	–
$\text{Cov}((\Delta y_{it})^2, \Delta y_{it+1})$	0.045	(0.008)	–
$\text{Kurt}(\Delta y_{it})$	14.728	(0.822)	–
$\text{Cov}((\Delta y_{it})^2, (\Delta y_{it+1})^2)$	0.176	(0.017)	–

Notes: The table presents the second, third, and fourth moments of earnings growth. Skewness and kurtosis correspond to the third and fourth standardized moments respectively. Column 1 reports moments of biennial earnings growth in our baseline sample in the PSID (we maintain the notation Δx_t for the first difference of variable x over time, noting that, given the biennial nature of the PSID, this corresponds to a difference over two calendar years in this case), while column 2 reports moments of male earnings growth from the U.S. Social Security Administration data. The latter moments concern the one-year male earnings growth, unconditional on past earnings, and correspond to averages over ages 30-55 of, respectively, dispersion squared, skewness, and kurtosis, as reported in online appendix tables C23-C25 in Guvenen, Karahan, Ozkan, and Song (2021). Block bootstrap standard errors are in parentheses.

C.2 CEX data and imputation in the PSID

The Consumer Expenditure Survey accounts for about 95% of all household expenditures from a highly disaggregated list of consumption goods and services (see Meyer and Sullivan, 2023, for a recent overview). The CEX is run by the Census Bureau and the Bureau of Labor Statistics, and it forms the basis for the calculation of the consumer's price index.

Sample selection and variables. To parallel the design of the PSID sample, we use CEX interview data between 1999 and 2019.[2] We select a sample of households that mimics closely our baseline selection in the PSID. Specifically, we select continuously married couples with the male spouse aged 30 to 65. We require non-missing data on expenditure and basic demographics, and we drop those with zero food expenditure given the nature of the consumption imputation subsequently. Similar to the consumption measure in the PSID, we define consumption as the sum of real expenditure on nondurable goods and services, namely food (at home and outside), utilities, out-of-pocket health expenses, public transport, vehicle expenses, education, and daycare. As the CEX collects quarterly data over four rolling quarters, we generate annual consumption as the sum of consumption over four quarters. We assign the final observation to the calendar year to which the underlying quarterly data mostly correspond to. Our final sample consists of 31,751 observations.

Table C.2 presents a comparison of means between the PSID and CEX samples for a host of household characteristics over the period 1999–2019. The two samples are very close with respect to age, family size, number of children, race, education, region of residence, or labor market participation. Disposable household income is higher in the PSID than in the CEX; BPP noted this also for the earlier period (1980–1992) and argued that it is due to the more comprehensive definition of income in the PSID. Food expenditure in the PSID is higher, on average, by about 60% than food expenditure in the CEX. This is in contrast to BPP who find that food expenditure over the earlier period 1980–1992 is similar across the two surveys. This results in consumption being about 25% higher in the PSID than in the CEX. Overall, the comparison suggests that the PSID and CEX samples exhibit some differences, with lower income households likely overrepresented in the latter. This may have implications for the degree of consumption partial insurance estimated using the internal consumption data in the PSID versus the data imputed from the CEX.

Figure C.1 plots the variance, skewness, and kurtosis of log consumption in the PSID (blue solid line) and CEX samples (black dotted line). Although the patterns for the variance and kurtosis are quite similar in the two surveys over time, consumption in the PSID is less

[2] To improve the quality of the food demand estimation subsequently, we keep *all* calendar years between 1999 and 2019, i.e., not just the odd ones as in the PSID.

Table C.2: Comparison of sample means, PSID and CEX

	1999		2009		2019	
	PSID	CEX	PSID	CEX	PSID	CEX
Earnings	72,829	59,011	109,047	84,383	130,488	106,316
Consumption	16,115	10,631	22,317	15,454	27,315	18,592
Food	7,434	4,461	9,306	6,090	12,575	7,636
Age	44.70	45.76	46.40	47.59	46.69	48.40
Family size	3.39	3.54	3.24	3.48	3.42	3.44
# children	1.13	1.21	1.01	1.11	1.19	1.07
% white	0.93	0.88	0.91	0.86	0.91	0.84
% high school dropout	0.09	0.12	0.06	0.12	0.07	0.10
% high school graduate	0.28	0.27	0.25	0.25	0.22	0.21
% college dropout	0.62	0.62	0.69	0.63	0.71	0.69
% Northeast	0.17	0.18	0.18	0.19	0.16	0.17
% Midwest	0.32	0.24	0.30	0.25	0.31	0.21
% South	0.33	0.32	0.32	0.34	0.33	0.33
% West	0.18	0.25	0.20	0.22	0.20	0.28
% working (male)	0.89	0.88	0.91	0.93	0.90	0.92
% working (female)	0.81	0.77	0.79	0.81	0.79	0.79

Notes: The table presents a comparison of means in the PSID and CEX samples in 1999, 2009, and 2019.

volatile relative to the CEX but slightly more leptokurtic. For skewness, there is disagreement both in the trends over time and, importantly, with respect to the sign, with the PSID and CEX samples featuring positive and weakly negative skewness, respectively. Overall, these statistics suggest that consumption across otherwise similarly selected samples in the PSID and CEX exhibits notable differences; this casts a caveat over the comparison of the partial insurance parameters across samples, especially in cases when higher-order moments of consumption are used in the estimation.

Consumption imputation from CEX into PSID. The previous comparison, however, is not directly informative about the moments of consumption growth, which is the relevant variable in our case. The CEX lacks the panel dimension needed for our analysis, so it cannot be used to obtain measures of annual (or biennial) consumption growth.[3] We thus follow BPP and impute nondurable consumption from the CEX into the PSID.

The imputation involves two distinct steps. In the first step, we estimate food demand in the CEX. Specifically, we let real food expenditure be a log-linear function of the log of

[3]The CEX consists of repeated cross sections, in which households are interviewed over only 4 quarters. There is also an initial interview, in principle a fifth quarter, but this is regarded as a training quarter and it is not used in empirical analyses.

Figure C.1: Moments of log consumption in the PSID and CEX

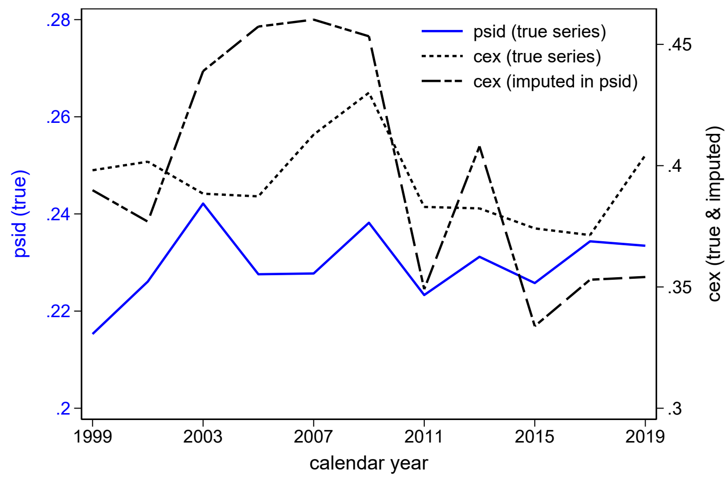

(a) Variance of log consumption

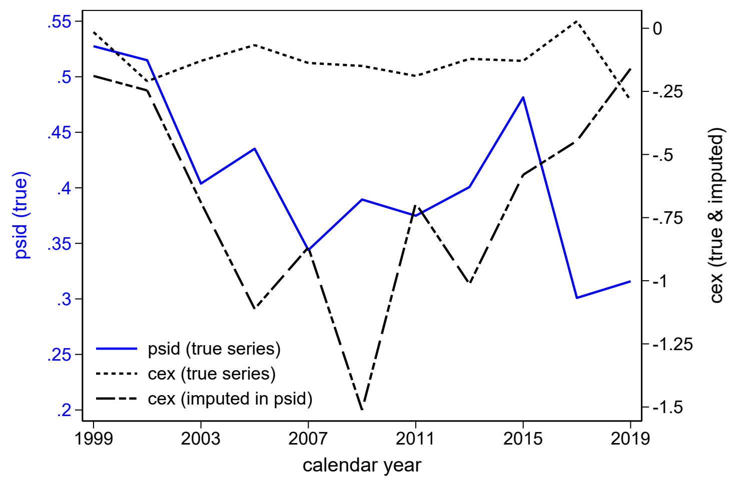

(b) Skewness of log consumption

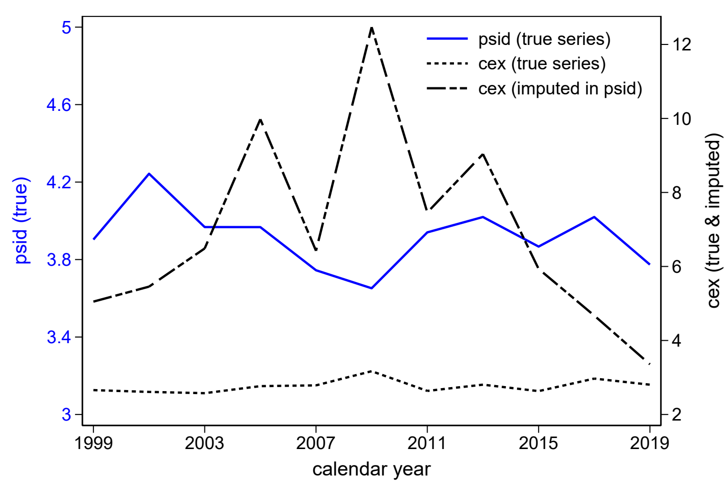

(c) Kurtosis of log consumption

Notes: The figure plots the moments of log consumption in the PSID (blue solid line; left axis) and the CEX (black dotted line for the true series and black dashed-dotted line for the imputed series; right axis).

nondurable expenditure c, relative prices $\mathbf{p}$, and demographics $\mathbf{W}$, namely

$$f_{it} = \beta(D_{it})c_{it} + \mathbf{p}_t'\boldsymbol{\gamma} + \mathbf{W}_{it}'\boldsymbol{\mu} + e_{it}, \tag{C.1}$$

where f_{it} is the log of real annual food expenditure for household i in the CEX in year t, β is the budget elasticity of food demand, which we allow to shift with time and household characteristics D_{it}, and e_{it} captures unobserved heterogeneity in food demand.[4] To address possible endogeneity of and measurement error in total expenditure, we follow an instrumental variables approach using as instruments the mean and standard deviation of household disposable income and the standard deviation of spousal wages by cohort, education, year, the former interacted with time, education, and number of children dummies.[5]

In the second step, we invert (C.1) under the assumption that food demand is monotonic in total expenditure. We can then use the inverted equation to predict nondurable expenditure c for a given level of food expenditure, given prices and demographics. Since the CEX and PSID samples are representative of the same underlying population, and since food expenditure, prices, and demographics are available in the PSID, we can use the inverted (C.1) to impute consumption into the PSID. This allows us to obtain an alternative, external measure of consumption in the PSID – and thus of consumption growth.

Figure C.1 plots the variance, skewness, and kurtosis of log consumption imputed from the CEX (black dashed-dotted line). All moments trend similarly over time between the true PSID series and the imputed series, although the magnitude is quite different across the two. Imputed consumption is more volatile than the true PSID, skewness is negative (as opposed to positive in the PSID) and peaks down in 2007–2011, kurtosis is often more than twice as high as kurtosis in the PSID and peaks upwards in 2007–2011.

Table 2 in the text reports the second and higher-moments of consumption growth in the imputed series, residualized through similar first-stage regressions as in the true data. The main qualitative features of the distributions remain similar across the true PSID data and the imputed series, although the moments of imputed consumption are several-fold accentuated relative to the true data. The imputation likely imparts substantial measurement/imputation error to consumption, which casts doubt on the reliability of the partial insurance estimates derived from imputed consumption data from the CEX.

[4]Prices $\mathbf{p}$ include the relative prices of food, public transportation, utilities, medical services, and childcare. Demographics $\mathbf{W}$ include a quadratic polynomial in age and dummy variables for the number of children, family size, education, region of residence, race, and cohort. D_{it} includes year dummies and dummies for education and the number of children.

[5]These instruments are slightly different from those used in BPP, namely the cohort-education-year specific averages of male and female hourly wages, interacted with time, education, and kids dummies. The use of their instruments over our sample period often resulted in convergence problems.

C.3 Comparison with BPP

Table C.3 uses the original BPP data over 1980-1992 (taken from the replication package of Blundell, Pistaferri, and Preston, 2008, available online) and does three things.

First, we replicate BPP over the original data/period of time, using BPP's *annual* growth rates of income and consumption. This is column 1 in table C.3. We obtain numerically similar results to BPP, namely $\phi^{(1)} = 0.64$ while $\psi^{(1)}$ is statistically zero.

Second, we calculate *biennial* growth rates of income and consumption and re-estimate the linear model. This is column 2 in table C.3; this estimation is analogous to the linear function estimation in column 1 of table 4, albeit using the original BPP sample. Note that despite the use of biennial data, $\phi^{(1)}$ and $\psi^{(1)}$ still express the pass-through at annual rates (see appendix B). $\phi^{(1)}$ drops by 20% to 0.518, while $\psi^{(1)}$ remains qualitatively unaffected. The use of biennial data thus reduces the pass-through of permanent shocks. With biennial data, we only identify the pass-though of an 'aggregate' permanent shock over two years. Some higher frequency shocks will be muted at the observed lower frequency, e.g., a second-year shock may partly undo a large first-year shock, so we cannot observe their transmission. In other words, we only measure an 'aggregate' pass-through within a given biennial period, which means that some higher-frequency shocks will be inevitably smoothed out at the observed lower frequency.

Third, we re-estimate the linear model assuming that the target second-order income and consumption moments are time-invariant. This differs from the original BPP exercise where the moments vary with time, but it is otherwise similar to our main exercise in this paper given that we cannot precisely estimate higher-order moments on a year-by-year basis. The results are in columns 3-4 in table C.3. Time-invariance further reduces the transmission parameter of permanent shocks ($\phi^{(1)}$ drops by another 20% to 0.404), bringing the pass-through closer to the value we estimate in the modern data. Taking averages of moments over time attenuates the largest income-consumption co-movements, and the model can fit the data requiring, ceteris paribus, smaller transmission parameters.

There are some caveats comparing the exercise in table C.3 and our results with the most recent data in table 4. First, our main exercise uses consumption data internally available in the PSID, while table C.3 uses consumption imputed from the CEX. We showed earlier that the imputation imparts substantial error. Second, while the estimation in table C.3 does not correct the income moments for measurement error (in line with BPP), our main exercise in table C.3 does so (in line with Blundell, Pistaferri, and Saporta-Eksten, 2016). Finally, the use of biennial growth rates does not allow identification of θ, the moving average parameter of the transitory shock, so the transitory shock is $MA(0)$ in our baseline ($\theta = 0$), while it is $MA(1)$ in the original BPP exercise and in table C.3.

Table C.3: Replication of BPP, original BPP data 1980-1992

Consumption fn.:	**Linear**			
Income moments:	2^{nd}	2^{nd}	2^{nd}	2^{nd}
Consumption data:	CEX	CEX	CEX	CEX
Time variability:	**Time-varying moments**		**Time-invariant moments**	
Growth rates:	annual (original BPP)	biennial	annual	biennial
	(1)	(2)	(3)	(4)
$\phi^{(1)}$	0.644	0.518	0.504	0.404
	(0.079)	(0.068)	(0.076)	(0.057)
$\psi^{(1)}$	0.024	0.117	0.043	0.158
	(0.044)	(0.050)	(0.044)	(0.051)
$\theta^{\#}$	0.112	0	0.115	0
	(0.025)	(not identified)	(0.026)	(not identified)

Notes: The table presents the estimates of the parameters of the linear consumption function, using the original BPP data over 1980-1992 (PSID income data, consumption data imputed from the CEX). The data are available through Blundell, Pistaferri, and Preston (2008)'s online replication package. Columns 1-2 allow the target second-order income and consumption moments to vary with calendar time, while columns 3-4 estimate and target said moments assuming time-invariance. Columns 1 and 3 estimate the model using annual growth rates of income and consumption, while columns 2 and 4 use biennial rates. In line with BPP, estimation is done via diagonally weighted GMM and asymptotic standard errors are in parentheses.
[#] θ is the moving average parameter of the transitory shock. BPP assume that the transitory component of income is given by $v_{it} = \epsilon_{it} + \theta\epsilon_{it-1}$, where ϵ_{it} is the transitory shock. Given their use of annual growth rates, θ is readily identified by the first-order income autocovariance. Our use of biennial data in this paper prevents identification of θ, which we thus set to 0.

C.4 Additional results

Table C.4 presents the estimates of the parameters of the consumption function using imputed consumption data from the CEX. As such, this table accompanies table 4 in the main text, which reports results using consumption internally available in the PSID.

In the linear specification, the transmission parameter of permanent shocks $\phi^{(1)}$ is consistently higher than that obtained from the consumption data directly available in the PSID, regardless of whether we target only second- or also higher-order moments. This echoes our earlier discussion about the imputation inflating the pass-through of shocks. We also find that the transmission parameter of transitory shocks $\psi^{(1)}$ is negative in all cases. Moreover, when up to fourth-order moments are targeted (column 3), the point estimate is particularly large in magnitude (-0.418) and statistically significant. This is counterintuitive and raises further concerns about the reliability of the imputation procedure. We draw similar conclusions about these two parameters in the quadratic specification. However, $\phi^{(2)} < 0$ in this case with a magnitude very similar to our baseline results from the PSID (-0.036 versus -0.04; albeit not significant). This is, once again, indicative of the asymmetric pass-through of negative versus positive permanent shocks. Similar to our baseline, we also find $\psi^{(2)} > 0$.

Table C.5 presents estimation results for the linear consumption function in subsamples formed by wealth and education. As such, this table accompanies table 6 in the main text, which reported wealth- and education-specific results in the quadratic specification.

Columns 1-4 present parameter estimates among low and high wealth households. $\phi^{(1)}$ is consistently smaller among the wealthier, regardless of which set of moments are targeted, indicating the self-insurance role of assets in that group. $\psi^{(1)}$ is statistically indistinguishable from zero in both groups. Columns 5-8 present parameter estimates among households with and without college education. The results do not present a clear pattern across groups. When only second moments of income and consumption are targeted, $\phi^{(1)}$ is larger for households without college education – 0.225 versus 0.123. However, when all moments up to fourth-order are included, $\phi^{(1)}$ is almost identical in the two groups.

Table C.4: Estimates of the consumption function, baseline sample, imputed consumption

Consumption fn.:	**Linear**			**Quadratic**
Income moments:	2^{nd}	2^{nd}, 3^{rd}	2^{nd}, 3^{rd}, 4^{th}	2^{nd}, 3^{rd}, 4^{th}
Consumption data:	CEX	CEX	CEX	CEX
	(1)	(2)	(3)	(4)
$\phi^{(1)}$	0.288	0.339	0.404	0.260
	(0.044)	(0.074)	(0.118)	(0.043)
$\psi^{(1)}$	-0.109	-0.155	-0.418	-0.108
	(0.078)	(0.098)	(0.247)	(0.086)
$\phi^{(2)}$				-0.036
				(0.038)
$\psi^{(2)}$				0.025
				(0.045)
$\omega^{(22)}$				-0.157
				(1.195)
σ^2_ξ	0.019	0.018	0.000	0.019
	(0.004)	(0.004)	(0.004)	(0.004)
γ_ξ		-0.016	-0.016	
		(0.007)	(0.007)	
κ_ξ			0.000	
			(0.000)	
$\sigma^2_{u_c}$	0.140	0.140	0.152	0.140
	(0.009)	(0.009)	(0.012)	(0.009)
γ_{u_c}		-0.185	-0.185	
		(0.037)	(0.037)	
κ_{u_c}			0.816	
			(0.183)	

Notes: The table presents the estimates of the parameters of the consumption function, assuming homogeneous transmission parameters over the lifecycle/in the cross-section and stationarity of taste heterogeneity and consumption measurement error. Columns 1-3 present parameter estimates in the linear function, while column 4 presents estimates in the quadratic case; the order of moments targeted in each case is shown at the top of the table. All columns use consumption imputed from the CEX (with details of the imputation reported in appendix C.2). Main text table 4 shows the baseline results using consumption internally available from the PSID. Estimation is via equally weighted GMM; block bootstrap standard errors are in parentheses.

Table C.5: Estimates of the linear consumption function, by wealth and education

Consumption fn.:	**Linear**							
	By wealth				By education			
Subsample:	Low wealth		High wealth		No college		Some college	
Income moments:	2^{nd}	2^{nd}, 3^{rd}, 4^{th}	2^{nd}	2^{nd}, 3^{rd}, 4^{th}	2^{nd}	2^{nd}, 3^{rd}, 4^{th}	2^{nd}	2^{nd}, 3^{rd}, 4^{th}
Consumption data:	PSID	PSID	PSID	PSID	PSID	PSID	PSID	PSID
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\phi^{(1)}$	0.223	0.148	0.111	0.099	0.225	0.107	0.123	0.113
	(0.040)	(0.045)	(0.037)	(0.055)	(0.055)	(0.086)	(0.033)	(0.042)
$\psi^{(1)}$	-0.021	0.047	0.011	0.018	-0.070	-0.041	0.024	0.031
	(0.066)	(0.062)	(0.052)	(0.076)	(0.074)	(0.086)	(0.050)	(0.067)
σ^2_ξ	0.014	0.015	0.023	0.023	0.014	0.016	0.021	0.022
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)
γ_ξ		-0.002		0.000		-0.002		0.000
		(0.001)		(0.001)		(0.001)		(0.001)
κ_ξ		0.003		0.007		0.003		0.006
		(0.001)		(0.001)		(0.002)		(0.001)
$\sigma^2_{u_c}$	0.048	0.048	0.042	0.042	0.043	0.043	0.044	0.044
	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)	(0.002)	(0.002)
γ_{u_c}		0.003		0.003		0.003		0.002
		(0.002)		(0.001)		(0.002)		(0.001)
κ_{u_c}		0.015		0.012		0.015		0.012
		(0.002)		(0.002)		(0.003)		(0.001)

Notes: The table presents the estimates of the parameters of the linear consumption function, allowing them to vary across wealth and education groups. This table accompanies the quadratic specification estimates in main text table 6. For reasons of brevity, we do not report results where second- and third-order moments of income and consumption are targeted in the estimation; these are available from the authors upon request. Low (high) wealth is defined on the basis of household wealth being less (more) than median real wealth in the sample over 1999–2019. No versus some college is defined on the basis of the highest level of education attained by the male spouse. All columns use consumption data internally available in the PSID. Estimation is via equally weighted GMM; block bootstrap standard errors are in parentheses.