

HIGHER-ORDER INCOME RISK OVER THE BUSINESS CYCLE*

BY CHRISTOPHER BUSCH AND ALEXANDER LUDWIG

LMU Munich, CESifo, and CEPR, Germany; Goethe University Frankfurt, ICIR, and CEPR, Germany

We explore the consequences of higher-order risk in a standard incomplete-markets life-cycle model. We calibrate the model using a canonical income process with persistent and transitory risk, extended to feature cyclical shock distributions with left-skewness and excess kurtosis. We estimate this income process for U.S. household data, and find shocks to be highly leptokurtic, with countercyclical variance and procyclical skewness of persistent shocks. In the model, first, higher-order risk has sizable welfare implications; second, it matters quantitatively for the welfare costs of cyclical idiosyncratic risk; third, it has nontrivial implications for self-insurance against shocks.

1. INTRODUCTION

The extent of idiosyncratic income risk matters for many macroeconomic questions. As our main contribution, we systematically evaluate a set of economic consequences of *higher-order risk* in an incomplete-markets model. Higher-order risk refers to features of the distribution of idiosyncratic risk captured by moments beyond the variance. We systematically decompose the role of the variance, skewness, and kurtosis of the shock distributions for economic choices—and link those to properties of preferences (risk attitudes). As our second contribution, we estimate higher-order idiosyncratic income risk and its variation over the business cycle within the canonical transitory-persistent decomposition (dating back to Gottschalk and Moffitt, 1994).

1.1. Implications of Higher-Order Risk. We consider a standard incomplete-markets life-cycle model with ex post heterogeneity, that is, the only source of inequality in the model is the risky idiosyncratic component of household income. Households receive stochastic income throughout their working lives, after which they enter a retirement phase and receive income through a pay-as-you-go pension system. The only means of self-insurance against income risk is private savings in a risk-free asset.

*Manuscript received September 2021; revised November 2023.

We thank Helge Braun for numerous helpful discussions as well as Chris Carroll, Russell Cooper, Johannes Gierlinger, Fatih Guvenen, Daniel Harenberg, Nikolas Hilger, Greg Kaplan, Fatih Karahan, Magne Mogstad, Serdar Ozkan, Luigi Pistaferri, Irina Popova, Luis Rojas, Raül Santaaulàlia-Llopis, Kjetil Storesletten, Peter Zorn, and seminar and conference participants at various places for insightful comments. We also thank four anonymous referees and the editor, Iourii Manovskii, for very helpful comments. We thank Rocio Madera for sharing her code to set up the PSID. Christopher Busch gratefully acknowledges financial support from the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D, and the European Research Council under grant Horizon 2020 GA no. 788547 (APMPAL-HET). Alexander Ludwig gratefully acknowledges financial support by the Research Center SAFE, funded by the State of Hessen initiative for research LOEWE and from NORFACE Dynamics of Inequality across the Life-Course (TRISP) grant: 462-16-120. Please address correspondence to: Christopher Busch, LMU Munich, Department of Economics, Geschwister-Scholl-Platz 1, 80539 Munich, Germany; E-mail: christopher.busch@econ.lmu.de
Alexander Ludwig, Goethe University Frankfurt, RuW PF 65, Theodor-W.-Adorno-Platz 4, 60629 Frankfurt am Main, Germany; E-mail: alexander.ludwig@econ.uni-frankfurt.de

We compare model outcomes under two different calibrations of the income process: (i) with shock distributions that feature (estimated) dispersion, nonzero skewness, and excess kurtosis of shocks, and (ii) with Gaussian shock distributions that feature the same dispersion (but with skewness and excess kurtosis of zero). The goal of the analysis is to establish how higher-order risk of the distribution matters for economic outcomes predicted by the model framework. Households have recursive preferences over consumption (Epstein and Zin, 1989, 1991; Weil, 1989), which allows us to separately control the intertemporal elasticity of substitution and the coefficient of risk aversion. The latter also pins down higher-order risk attitudes, and through this is a crucial determinant of the behavioral reaction to higher-order risk.

Our analysis delivers three main findings. First, evaluated from an ex-ante perspective higher-order risk has sizable negative welfare implications for strong risk attitudes: the consumption equivalent variation (CEV) that makes agents in the economy with lognormal shocks indifferent to the economy with higher-order risk ranges between -0.5% (for a coefficient of risk aversion of 2) and -14.5% (relative risk aversion of 4). The dominant economic mechanism driving this result is a change of the consumption profile over the life cycle: when facing riskier income, risk-sensitive agents form more precautionary savings, and thus consume less at young ages. With weak risk attitudes (like log utility), the welfare effect can be positive (CEV of 0.4%).¹

Second, higher-order risk matters for the welfare costs of business cycles. Since Lucas (1987, 2003) argued that the gains of smoothing cycles beyond what the existing tax and transfer system does would be small, several studies have explored the role of both ex ante and ex post heterogeneity, with Imrohroglu (1989) being the first to emphasize the importance of idiosyncratic risk and incomplete markets. In a model similar to hers, Storesletten et al. (2001) allow for cyclical variance of persistent shocks as estimated in Storesletten et al. (2004). Following a similar strategy, we provide the first systematic assessment of the welfare consequences of cyclical higher-order risk as captured in a continuous distribution function, and thus bridge this approach to papers that explore cyclical downside risk in the form of unemployment (e.g., Krusell and Smith, 1999; Krebs, 2003, 2007; Beaudry and Pages, 2001; Krusell et al., 2009). Under higher-order risk, we find welfare costs—computed as CEV making households in the noncyclical economy indifferent to the cyclical economy—which are 0.3 percentage points (risk aversion of 2) to 7.6 percentage points (risk aversion of 4) larger than under lognormal shocks.

Third, higher-order risk crucially matters for the degree of self-insurance against shocks. We employ a measure of self-insurance introduced in the literature by Blundell et al. (2008), who suggest to evaluate the degree of partial insurance against income shocks by estimating the pass-through of the identified transitory and permanent shocks to consumption changes. In the context of our model based analysis, we follow Kaplan and Violante (2010), who study how much of the empirically estimated partial insurance can be generated in a standard incomplete markets model. We find that incorporating higher-order risk leads to weaker pass-through of income shocks to consumption. However, this does not necessarily represent *better* insurance against negative shocks. In a scenario with higher-order risk agents form more precautionary savings (relative to a scenario in which they face lognormal shocks), which implies a weaker consumption reaction to *positive* shocks. *Negative* shocks can actually translate more strongly into negative consumption changes, if the higher savings do not suffice to smooth out negative shocks that are more pronounced relative to normal shocks. In particular, for transitory shocks, we find this to be the case. Therefore, we caution against using only the insurance coefficient by Blundell et al. (2008) for the analysis of the degree of partial insurance against income risk.

¹ What might appear surprising at first glance, follows mechanically: the introduction of *third-order risk* (left-skewness) reduces *second-order risk* (variance) when characterizing the distribution in levels and holding the variance in logs constant.

1.2. *Estimation of Higher-Order Risk.* We characterize both transitory and persistent shocks by their second to fourth central moments, which in the case of the persistent shocks we allow to be state-contingent. We estimate these distribution moments using the second to fourth cross-sectional central moments and comoments of incomes—whereas similar estimations traditionally are based solely on the variance–covariance matrix. Identification of cyclicity follows from the fact that the accumulated second to fourth central moments systematically differ across cohorts if these cohorts experience different macroeconomic histories and if the moments of shocks differ systematically over the business cycle. This identification idea was introduced in Storesletten et al. (2004) for the second moment, and we extend it to higher-order moments.

Although Storesletten et al. (2004) analyze household-level income including government transfers from the Panel Study of Income Dynamics (PSID) and find *countercyclical variance*² of persistent shocks, later evidence in Guvenen et al. (2014) based on administrative social security data (SSA) for individual males in the United States suggests that individual downside risk is larger in contractions, whereas upside risk is smaller—this is reflected in a more pronounced left-skewness of the distribution of earnings changes in contractions. Busch et al. (2022) conduct a nonparametric analysis and find *procyclical skewness* of individual and household-level annual earnings changes in Germany, Sweden, France, and the United States.³

Our estimation approach draws a richer image of cyclical income changes within the transitory-persistent framework and thus bridges the previous studies. We use the PSID, which allows us to control for a rich set of household-level information and to take into account several relevant public transfer components. We focus on a measure of household net income defined as labor income plus public transfers minus taxes. This reflects the view that this measure represents the amount of risk remaining after insurance through the main insurance mechanisms other than self-insurance—that is, within-household insurance and government taxes and transfers (cf. Blundell et al., 2008)—and as such delivers the necessary ingredient for our model analysis, in which agents insure against this remaining risk using private savings.

We find that both transitory and persistent shocks to household net income feature strong left-skewness, and that persistent shocks are significantly cyclical: in contractions, their distribution is more dispersed and more left-skewed. The magnitude of dispersion is in line with the estimates in Storesletten et al. (2004). Finally, we find strong excess kurtosis of transitory and persistent shocks. One related recent study of cyclical risk is Angelopoulos et al. (2022), who adapt a version of our estimator and document procyclical skewness of persistent shocks in data from the British Household Panel Study.

Our article is part of a growing literature that explicitly analyzes the implications of new insights on cyclical skewness of persistent earnings shocks for macroeconomic questions. Golosov et al. (2016) allow for time-varying skewness of idiosyncratic risk in a study of optimal fiscal policy, Catherine (2021) analyzes the implications of procyclical skewness of idiosyncratic income risk for the equity premium, and McKay (2017) links procyclical skewness to aggregate consumption dynamics. Our analysis is also related to work on the implications of rich earnings dynamics in general (without considering the cyclicity of risk). De Nardi et al. (2020) feed an income process a la Arellano et al. (2017) into an incomplete markets model and study the role of richer earnings dynamics for consumption insurance and the

² This terminology has been introduced in the macroeconomic asset pricing literature, see Mankiw (1986), Constantinides and Duffie (1996), and Storesletten et al. (2007). Building on the framework of Storesletten et al. (2004), Bayer and Juessen (2012) focus on residual hourly wages and based on PSID data estimate countercyclical dispersion of persistent shocks in the United States.

³ More recently, Guvenen et al. (2021) document that, in a given year, most individuals experience small earnings changes, whereas some workers experience very large changes of their earnings. This is summarized by a high kurtosis—relative to what the conventional assumption of log-normality implies. See also De Nardi et al. (2021) for the Netherlands, and Druedahl and Munk-Nielsen (2020) for Denmark.

welfare costs of idiosyncratic risk. Their analysis features nonlinearity, nonnormality, and age-dependence of the income process and corroborates results from Karahan and Ozkan (2013) regarding the role of age-dependent persistence and distributions of shocks. Civale et al. (2017) analyze implications of left-skewed and leptokurtic idiosyncratic shocks for the interest rate and aggregate savings in an otherwise standard Aiyagari economy. Besides the particular outcomes of interest, our contribution to that literature is that we provide a transparent link between moments of the shock distribution and the outcomes of interest, emphasizing the relevant properties of preferences.

The remainder of the article is structured as follows: Section 2 illustrates the implications of higher-order risk for welfare and savings in a simple two-period model, and then sets up the quantitative version of the life-cycle economy. Section 3 presents our empirical approach, discusses identification of the income process, and presents the results of applying the estimator to U.S. household earnings data from the PSID. Section 4 shows the consequences of higher-order risk in the quantitative model, and Section 5 concludes.

2. HIGHER-ORDER RISK IN A LIFE-CYCLE MODEL

2.1. A Two-Period Model. Before we turn to a quantitative life-cycle economy, we consider a simple two-period consumption-savings problem, in which agents face riskless first-period income and risky second-period income. Within this framework, we illustrate that the distribution of risk (i) matters for welfare, (ii) affects precautionary savings, and (iii) affects the marginal propensity to consume (MPC). The utility and behavioral implications depend on risk attitudes of households.

2.1.1. Setting. The household receives riskless income y_0 and risky income y_1 in periods 0 and 1, respectively. Denote the distribution function of y_1 by $\Psi(y)$, which is characterized by its central moments around the mean, μ_k^y .⁴ Households enter period 0 with zero assets and, in the general formulation of the model, have access to a risk-free savings technology with zero interest. The budget constraints are $c_0 + a_1 = y_0$ and $c_1 \leq a_1 + y_1$, where a_1 denotes savings in period 0, and c_0 and c_1 denote consumption in the two periods. Preferences over consumption streams are additively separable and per-period utility is given by $u(c)$. We assume that $u(c)$ is continuously differentiable, that the derivatives feature alternating signs with positive odd derivatives, and denote the k th derivative of the utility function by $u^{(k)}$. We assume no discounting of the future such that expected lifetime utility is $U = u(c_0) + \mathbb{E}[u(c_1)]$. We also assume that $\mathbb{E}[y_1] = y_0$. The assumptions imply that there is no intertemporal savings motive in this simple model.

2.1.2. Welfare. Consider hand-to-mouth consumers and shut down the savings technology, that is, $a_1 = 0$. Mechanically, the welfare implication of changing the k th central moment of $\Psi(y_1)$ is pinned down by the k th derivative of $u(c)$. To see this, consider an exact Taylor series expansion of the objective function around the mean of period 1 consumption, $c_1 = \mathbb{E}[y_1]$:⁵

$$(1) \quad U = u(c_0) + \mathbb{E}[u(c_1)] = u(c_0 = y_0) + \sum_{k=0}^{\infty} \frac{u^{(k)}(c_1 = \mathbb{E}[y_1])}{k!} \mu_k^{y_1}.$$

We refer to an increase of even moments and/or a reduction of odd moments as an increase in risk: a higher variance (larger $\mu_2^{y_1}$), higher kurtosis (larger $\mu_4^{y_1}$ for a given $\mu_2^{y_1}$), or more nega-

⁴ $\mu_k^y = \mathbb{E}[(y - \mathbb{E}[y])^k]$ for $k = 1, 2, 3, \dots$

⁵ The expression given for the Taylor series expansion around $c_1 = \mathbb{E}[y_1]$ follows from the linearity of the expectation operator: $\mathbb{E}[u(c_1)] = \mathbb{E}[\sum_{k=0}^{\infty} (u^{(k)}(\mathbb{E}[y_1])/k!)(y_1 - \mathbb{E}[y_1])^k] = \sum_{k=0}^{\infty} (u^{(k)}(\mathbb{E}[y_1])/k!) \mathbb{E}[(y_1 - \mathbb{E}[y_1])^k] = \sum_{k=0}^{\infty} (u^{(k)}(\mathbb{E}[y_1])/k!) \mu_k^{y_1}$. See Appendix A.1 for a derivation of all analytical expressions used in this section.

tive skewness (smaller, or “more negative,” $\mu_3^{y_1}$ for a given $\mu_2^{y_1}$) all imply a riskier distribution. Per Equation (1)—and our assumption of alternating signs of the derivatives $u^{(k)}$ —we see that the household experiences a welfare reduction for each of these increases in risk—and how strong that reduction is depends on the size of the corresponding derivative of the utility function. The relevant risk attitudes that imply that the household cares about second-, third-, and fourth-order risk are referred to as *risk aversion*, *prudence*, and *temperance*, respectively (see, e.g., Eeckhoudt, 2012 for a discussion of these risk attitudes).

2.1.3. *Precautionary savings.* Now, we allow for savings. Given that households face risky income in period 1, they will form precautionary savings to forearm against low-income realizations. The first-order condition characterizing the optimal savings choice is

$$(2) \quad f \equiv u^{(1)}(c_0) - \mathbb{E}[u^{(1)}(c_1)] = u^{(1)}(y_0 - a_1) - \sum_{k=0}^{\infty} \frac{u^{(k+1)}(\mathbb{E}[y_1] + a_1)}{k!} \mu_k^{y_1} = 0,$$

where after plugging in the budget constraints, we express marginal utility tomorrow by a Taylor expansion. As before, this step allows us to see how the change of some central moment of period 1 income translates into the expected value—in this case, of marginal utility, for a given choice a_1 . Equation (2) shows that an increase of risk (higher $\mu_2^{y_1}$, lower $\mu_3^{y_1}$, higher $\mu_4^{y_1}$) implies an increase of expected marginal utility. Thus, the household chooses a higher a_1 such that the FOC holds in the higher risk scenario. The Taylor expansion shows that derivative $u^{(k+1)}(c)$ matters to induce a behavioral reaction of a_1 to a change of the k th moment.⁶ Thus, preferences need to feature a positive third derivative (*prudence*) for $\mu_2^{y_1}$ to matter, a negative fourth derivative (*temperance*) for $\mu_3^{y_1}$ to matter, and a positive fifth derivative (*edginess*) for $\mu_4^{y_1}$ to matter.⁷ If they do, then households reduce the adverse utility consequences of increased risk (second and higher order) by increasing savings a_1 .⁸

2.1.4. *MPC.* From (2), we can derive the marginal propensity to save (MPS) out of a shock to period 0 income by setting the total derivative df/dy_0 to zero. The MPC then follows as $MPC=1-MPS$:

$$(3) \quad MPC \equiv 1 - \frac{\partial a_1}{\partial y_0} = 1 - \left(1 + \frac{\mathbb{E}[u^{(2)}(y_1 + a_1)]}{u^{(2)}(y_0 - a_1)} \right)^{-1}.$$

We are now interested in $dMPC/d\mu_k^{y_1}$ for $k = 2, 3, 4$. There are two effects of an increase of risk: first, the precautionary savings response discussed above increases a_1 , which per Equation (3) decreases the MPC.⁹ This intuitively simply says that due to increased risk, households form more precautionary savings, and therefore, they consume less out of an additional unit of income received today. Second, for a given a_1 , an increase of risk reduces the (negative) expected value $\mathbb{E}[u^{(2)}(y_1 + a_1)]$: again, this becomes apparent when considering a Taylor expansion around the mean, which gives $\mathbb{E}[u^{(2)}(y_1 + a_1)] = \sum_{k=0}^{\infty} (u^{(k+2)}(\mathbb{E}[y_1] + a_1)/k!) \mu_k^{y_1}$. Thus, the expression in parenthesis in (3) is larger for a given a_1 , giving a larger MPC. This reaction of the MPC to an increase of risk is driven by high-order risk attitudes: for example, the fourth derivative—*temperance*—matters for MPC reactions to variance. The overall effect on the MPC of increasing risk is thus ambiguous and a quantitative question.

⁶ We explore precautionary savings in response to higher-order risk in the absence of a liquidity constraint. This complements Carroll et al. (2021), who analyze the savings response to a liquidity constraint.

⁷ The term *edginess* was coined by Lajeri-Chaherli (2004).

⁸ This extends the result on precautionary savings in response to increases in the variance in Kimball (1990) and corroborates findings in Eeckhoudt and Schlesinger (2008) using a slightly different approach.

⁹ To see this, recall that $u^{(2)}(c)$ is negative and $u^{(3)}(c)$ is positive.

2.2. *A Life-Cycle Model.* Per the illustration in the two-period model, we know that deviations from lognormal shocks matter in principle—but does the estimated deviation from the canonical income process matter quantitatively? Put differently, what is the *economic* relevance of the estimated skewness and kurtosis? To answer this question, we use a quantitative version of the simple two-period model of Section 2.1 by extending it to a standard multiperiod life-cycle model.

2.2.1. *Endowments.* Households live from age $j = 0$ to age $j = J$. Before they retire exogenously at age j_r , they receive exogenous earnings at each age $j < j_r$ that consist of an age-specific deterministic component e_j and two stochastic income components: transitory income shock ε , drawn i.i.d. from distribution \tilde{F}_ε , and persistent income z , which is AR(1) with autocorrelation coefficient ρ and shock η . The distribution of the persistent shock, $\tilde{F}_\eta(s)$, depends on the aggregate state s . In retirement, persistent income stays constant, and thus,

$$(4) \quad z'(z; s') = \begin{cases} \rho z + \eta' & \text{for } j < j_r, \text{ where } \eta' \underset{i.i.d.}{\sim} \tilde{F}_\eta(s') \\ z & \text{for } j \geq j_r. \end{cases}$$

For $j \geq j_r$, households earn a fixed pension income contingent on the last persistent income state before retirement, $b(z)$. Net household income is thus

$$(5) \quad y(j, z, \varepsilon) = \begin{cases} e_j \cdot \exp(z + \varepsilon) & \text{for } j < j_r \\ b(z) & \text{for } j \geq j_r. \end{cases}$$

There is a risk-free savings technology with (acyclical) rate of return r , and a zero borrowing constraint. Thus, the dynamic budget constraint is

$$(6) \quad a'(j, z, \varepsilon; s) = a \cdot (1 + r) + y(j, z, \varepsilon) - c(j, z, \varepsilon; s) \geq 0.$$

The aggregate state $s \in \{C, E\}$ follows a first-order Markov process with time-invariant transition matrix Π_s . We abstract from a first-order effect of the cycle on income.

2.2.2. *Preferences and household problem.* Households maximize recursive utility by solving a consumption-savings problem every period. They discount the future at factor $\beta > 0$.¹⁰ The state variables of the household are age j , asset holdings a , persistent income state z , transitory shock ε , and the aggregate state of the economy s . The recursive problem is

$$(7) \quad V_j(a, z, \varepsilon; s) = \max_{c, a'} \begin{cases} \left((1 - \tilde{\beta})c^{1-\frac{1}{\gamma}} + \tilde{\beta}(v(V_{j+1}(a', z', \varepsilon'; s')))^{1-\frac{1}{\gamma}} \right)^{\frac{1}{1-\gamma}} & \gamma \neq 1 \\ \exp \{ (1 - \tilde{\beta}) \ln c + \tilde{\beta} \ln (v(V_{j+1}(a', z', \varepsilon'; s'))) \} & \gamma = 1 \end{cases}$$

s.t. (4), (5), and (6),

where $v(V_{j+1}(\cdot))$ denotes the certainty equivalent of next period's continuation value, $\tilde{\beta} = \beta/(1 + \beta)$ the relative utility weight on this certainty equivalent, and γ the intertemporal elasticity of substitution between current period utility from consumption and $v(V_{j+1}(\cdot))$. Parameter θ pins down risk attitudes of households (see Section 2.1), and thus, determines the certainty equivalent of the risky continuation value:

$$(8) \quad v(V_{j+1}(a', z', \varepsilon'; s')) = \begin{cases} (\mathbb{E}_j[V_{j+1}(a', z', \varepsilon'; s')^{1-\theta}])^{\frac{1}{1-\theta}} & \theta \neq 1 \\ \exp(\mathbb{E}_j[\ln V_{j+1}(a', z', \varepsilon'; s')]) & \text{otherwise,} \end{cases}$$

¹⁰ It is straightforward to show that our intuitive arguments of Section 2.1 extend to recursive preferences, cf. our working paper version, Busch and Ludwig (2020).

where conditional expectations \mathbb{E}_j are defined with respect to the realization of next period's aggregate state of the economy s' , transitory income shock ε' , and persistent income shock η' .

2.2.3. Equilibrium and model solution. We consider a partial equilibrium in which the interest rate is exogenously given and the pension system clears. Given constant aggregate income across different scenarios of idiosyncratic income risk, this ensures that also average pension income is constant across these scenarios—this rules out first-order income difference across scenarios and thus a change of the life-cycle savings incentives. We solve for household policy and value functions using the method of endogenous gridpoints (cf. Carroll, 2005). We aggregate by explicit aggregation iterating forward on the age-specific cross-sectional distribution $\Phi_j(a_j, z_j, \varepsilon; s)$, which follows from the initial distribution $\Phi_0(a_0, z_0, \varepsilon_0; s)$ and the transition function $G_j(a_j, z_j, \varepsilon_j; s)$. The latter is induced by the exogenous laws of motion of s and z , the exogenous distribution of ε , and the endogenous age-specific policy function $a'_j(a_j, z_j, \varepsilon_j; s)$.

2.2.4. Risk scenarios. We explore the consequences of higher-order income risk through the lens of the model by, first, calibrating the distributions of shocks, \tilde{F}_ε and $\tilde{F}_\eta(s)$, such that they feature higher-order risk, and, second, comparing model outcomes to a calibration with lognormal shock distributions. To inform the calibration, we estimate an income process consistent with the one specified above (outside of the model), which we then feed into the model. To this end, we proceed in three steps. First, we estimate the second to fourth central moments of the shock distributions by the generalized method of moments (GMM). Second, we fit parametric distribution functions to the estimated moments. Third, we discretize the distributions using their quantile functions to use them in the model. Thus, our approach directly translates the estimated central moments into the model's shock distributions without the need to simulate. We calibrate the remaining elements of the model in a standard way, with details provided in Section 4.1. We now turn to the estimation of the income process.

3. ESTIMATING HIGHER-ORDER INCOME RISK

3.1. Income Process with Higher-Order Risk. Let log income of household i of age j in year t be

$$(9) \quad \ln(y_{ijt}) = f(\mathbf{X}_{ijt}, Y_t) + \tilde{y}_{ijt},$$

where $f(\mathbf{X}_{ijt}, Y_t)$ is the *deterministic* component of income, that is, the part that can be explained by observable individual and aggregate characteristics, \mathbf{X}_{ijt} and Y_t , respectively, and \tilde{y}_{ijt} is the *residual* part of income, which is assumed to be orthogonal to $f(\mathbf{X}_{ijt}, Y_t)$. The deterministic component $f(\mathbf{X}_{ijt}, Y_t)$ is a linear combination of a cubic in age j , $f_{age}(j)$, the log of household size, year fixed effects, and an education premium $f_{EP}(t)$ for college education, which we allow to vary over years t :

$$(10) \quad f(\mathbf{X}_{ijt}, Y_t) = \beta_{0t} + f_{age}(j) + \mathbb{1}_{edu_i=c} f_{EP}(t) + \beta^{size} \ln(hhsiz_e_{ijt}),$$

where $f_{age}(j) = \beta_1^{age} j + \beta_2^{age} j^2 + \beta_3^{age} j^3$, $f_{EP}(t) = \beta_0^{EP} + \beta_1^{EP} t + \beta_2^{EP} t^2$, and indicator $\mathbb{1}_{edu_i=c}$ takes on value 1 for college-educated households.

We model residual income \tilde{y}_{ijt} as the sum of three components: a persistent component z_{ijt} , an i.i.d. transitory shock ε_{ijt} , and an idiosyncratic *fixed effect* χ_i . The idiosyncratic fixed effect is a shock drawn once upon entering the labor market from a distribution that is the same for every cohort.¹¹ The persistent component is modeled as an AR(1) process with innovation η_{ijt} :

¹¹ Thus, from the econometric perspective, we estimate a random effects model.

$$(11a) \quad \tilde{y}_{ijt} = \chi_i + z_{ijt} + \varepsilon_{ijt}, \text{ where } \varepsilon_{ijt} \underset{i.i.d.}{\sim} F_\varepsilon, \chi_i \underset{i.i.d.}{\sim} F_\chi$$

$$(11b) \quad z_{ijt} = \rho z_{ij,t-1} + \eta_{ijt}, \text{ where } \eta_{ijt} \underset{id}{\sim} F_\eta(s(t)),$$

where F_χ , F_ε , and $F_\eta(s(t))$ denote the density functions of χ , ε_{ijt} , and η_{ijt} , respectively. The density function of the persistent shock depends on the aggregate state of the economy in period t , $s(t)$. This income process features the canonical transitory-persistent specification (e.g., Moffitt and Gottschalk, 2011). However, instead of focusing on the variance alone, we are interested in estimating the second to fourth central moments of the shocks, which we denote by μ_2^x , μ_3^x , and μ_4^x , for $x \in \{\chi, \varepsilon, \eta(s)\}$.¹²

As in Storesletten et al. (2004), the economy in period t can be in one of two aggregate states $s(t)$, which we denote by E (expansion) and C (contraction), $s(t) \in \{E, C\}$. Thus, the central moments of the persistent shock $\mu_k^\eta(s(t))$ are equal to $\mu_k^{\eta,E}$ or $\mu_k^{\eta,C}$ if $s(t) = E$ or $s(t) = C$, respectively (for $k = 2, 3, 4$). Both empirical evidence (e.g., Blundell et al., 2008) and model-based analyses (e.g., Kaplan and Violante, 2010) find that households can insure well against transitory shocks. We therefore follow Storesletten et al. (2004) and only consider the cyclicity of persistent income shocks, which have long-lasting effects in the context of a life-cycle decision-making problem. We still do capture skewness and kurtosis of the (acyclical) transitory component and explore its quantitative role. In Appendix B.2, we show robustness with respect to this benchmark choice.

We assume that upon entering the labor market, in addition to drawing the fixed effect χ_i , each worker draws the first realizations of transitory and persistent shocks, ε_{it} and η_{it} , from the distributions F_ε and $F_\eta(s(t))$, respectively. Thus, the moments of the distribution of the persistent component for the cohort entering in year t at age $j = 0$ are $\mu_k(z_{i0t}) = \mu_k^\eta(s(t))$.

3.1.1. *Implied central moments.* Central (co-)moments are given by

$$(12a) \quad \mu_k(\tilde{y}_{ijt}; \theta) = \mathbb{E} \left[(\tilde{y}_{ijt} - \mathbb{E}[\tilde{y}_{ijt}] | s^t) \right]^k,$$

$$(12b) \quad \mu_{kl}(\tilde{y}_{ijt}, \tilde{y}_{ij+1t+1}; \theta) = \mathbb{E} \left[(\tilde{y}_{ijt} - \mathbb{E}[\tilde{y}_{ijt}] | s^t) (\tilde{y}_{ij+1t+1} - \mathbb{E}[\tilde{y}_{ij+1t+1}] | s^t) \right]^l,$$

where $\theta = (\rho, \mu_2^\chi, \mu_2^\varepsilon, \mu_2^{\eta,E}, \mu_2^{\eta,C}, \mu_3^\chi, \mu_3^\varepsilon, \mu_3^{\eta,E}, \mu_3^{\eta,C}, \mu_4^\chi, \mu_4^\varepsilon, \mu_4^{\eta,E}, \mu_4^{\eta,C})$ is a vector of second-stage parameters, and s^t summarizes the history of aggregate states up to year t .¹³ The process implies the following cross-sectional moments of the distribution of residual income at age j in year t :

$$(13a) \quad \mu_2(\tilde{y}_{ijt}; \theta) = \mu_2^\chi + \mu_2^\varepsilon + \mu_2(z_{ijt}),$$

$$(13b) \quad \mu_3(\tilde{y}_{ijt}; \theta) = \mu_3^\chi + \mu_3^\varepsilon + \mu_3(z_{ijt}),$$

$$(13c) \quad \mu_4(\tilde{y}_{ijt}; \theta) = \mu_4^\chi + \mu_4^\varepsilon + \mu_4(z_{ijt}) + 6(\mu_2^\chi \mu_2^\varepsilon + (\mu_2^\chi + \mu_2^\varepsilon) \mu_2(z_{ijt})),$$

and the following (auto) comoments:

$$(14a) \quad \mu_{11}(\tilde{y}_{ijt}, \tilde{y}_{ij+1t+1}; \theta) = \mu_2^\chi + \rho \mu_2(z_{ijt}),$$

¹² One potential disadvantage of using central moments to characterize the shocks in the income process is that they are hard to interpret by themselves. However, in the data, the central moments of the cross-sectional income distribution are strongly correlated with percentile-based counterparts to those moments. We are thus confident that the estimated central moments—and the implied standardized moments *skewness* and *kurtosis*—do capture salient features of the distribution.

¹³ Note that we need to condition only on s^t , not on s^{t+1} , because period $t + 1$ shocks are uncorrelated with all shocks accumulated up to period t .

$$(14b) \quad \mu_{21}(\tilde{y}_{ijt}, \tilde{y}_{ij+1t+1}; \theta) = \mu_3^x + \rho\mu_3(z_{ijt}),$$

$$(14c) \quad \mu_{31}(\tilde{y}_{ijt}, \tilde{y}_{ij+1t+1}; \theta) = \mu_4^x + \rho\mu_4(z_{ijt}) + 3(\mu_2^x\mu_2^\varepsilon + (\mu_2^x + \rho(\mu_2^x + \mu_2^\varepsilon))\mu_2(z_{ijt})),$$

where $\mu_k(z_{ijt})$, for $k = 2, 3, 4$ are given recursively by $\mu_2(z_{ijt}) = \rho^2\mu_2(z_{ij-1t-1}) + \mu_2^\eta(s(t))$, $\mu_3(z_{ijt}) = \rho^3\mu_3(z_{ij-1t-1}) + \mu_3^\eta(s(t))$, and $\mu_4(z_{ijt}) = \rho^4\mu_4(z_{ij-1t-1}) + 6\rho^2\mu_2(z_{ij-1t-1})\mu_2^\eta(s(t)) + \mu_4^\eta(s(t))$, respectively.

3.1.2. *Identification.* All parameters of (11) are identified using the moments (13a)–(14c) and their empirical counterparts. As seen in (13a), the sum $(\mu_2^x + \mu_2^\varepsilon)$ is identified as the intercept of the variance profile over age. Next, the difference between the two expressions for (13a) and (14a) identifies μ_2^x separately from μ_2^ε . The persistence parameter ρ is identified from the curvature of the variance age profile. To see this, ignore t for the moment, and consider the change of variance from age j to $j + 1$: $\Delta\mu_2(\tilde{y}_{ij}) = \mu_2(\tilde{y}_{ij+1}) - \mu_2(\tilde{y}_{ij}) = (\rho^2 - 1)\mu_2(z_{ij}) + \mu_2^\eta = \dots = \mu_2^\eta\rho^{2j}$. Thus, the relative slope at two different ages j and $j' > j$ identifies ρ : $\Delta\mu_2(\tilde{y}_{ij'})/\Delta\mu_2(\tilde{y}_{ij}) = \rho^{2j'}/\rho^{2j} = \rho^{2(j'-j)}$. A concave variance profile over age implies $\rho < 1$. The overall magnitude of the increase of the cross-sectional variance over age identifies the variance of persistent shocks, μ_2^η .

The above logic translates almost one-for-one to the third moment: $(\mu_3^x + \mu_3^\varepsilon)$ is identified via the intercept of the age profile of the third central moment, as seen in (13b). The difference between the expressions for the third central moment and comoment, Equations (13b) and (14b), identifies μ_3^x separately from μ_3^ε .¹⁴ Given ρ and the variance parameters μ_2^x for $x \in \{\chi, \varepsilon, \eta(s)\}$, Equations (13c) and (14c) identify the fourth central moments μ_4^x for $x \in \{\chi, \varepsilon, \eta(s)\}$ in the same way as for the second and third central moments.

Let us now turn to cyclicity. The difference between $\mu_2^{\eta,C}$ and $\mu_2^{\eta,E}$ is identified by the difference between the cross-sectional variances of different cohorts at the same age (for cohorts which by that age lived through different histories of contractions and expansions). The use of cross-sectional moments for identification allows us to exploit macroeconomic information that predates the micro panel, thereby incorporating more business cycles in the analysis than covered by the sample, as pointed out by Storesletten et al. (2004). Consider the persistent component of the income process in Equation (11b). The variance of the innovations accumulates as a cohort ages, as can be seen in Equation (13a). If the innovation variance is higher in contractionary years, then a cohort that lived through more contractions will have a higher income variance at a given age than a cohort at the same age that lived through fewer contractions, if the persistence is high.

Our extension of Storesletten et al. (2004) is based on the insight that other central moments accumulate in a similar fashion, as seen in Equations (13b) and (13c). Consider the third central moment and compare again two cohorts when they reach a certain age. If the third central moment of the shock was smaller (more negative) in a contraction than in an expansion, that is, $\mu_3^{\eta,C} < \mu_3^{\eta,E}$, then this would imply a more negative cross-sectional third central moment for the cohort that worked through more contractions.¹⁵ For a given dispersion, this implies a reduction of skewness (a more left-skewed distribution).

Note that by restricting the transitory shocks to not vary over the business cycle, we do not bias the estimated cyclicity of persistent shocks, which is identified via accumulated shock distributions. Note that Huggett and Kaplan (2016) use a similar strategy based on second and

¹⁴ Also, the curvature of the age profile of the third moment gives additional overidentifying restrictions for the persistence parameter ρ .

¹⁵ Let us emphasize that the cross-sectional distribution of \tilde{y}_{ijt} does not converge to a Normal distribution: the third and fourth central moments of the shocks accumulate over age. This allows identification of higher-order moments of the shock distributions based on cross-sectional moments. Of course, if the shock distributions are symmetric in logs, then this is identified as well, as $\mu_3(\tilde{y}_{ijt}) = 0$ in this case.

third central moments and comoments, without resorting to presample aggregate information in the spirit of Storesletten et al. (2004) as we do.

3.2. Data and Sample Selection. We use data from the PSID, which interviews households in the United States annually from 1968 to 1997 and every other year since then. The representative core sample consists of about 2,000 households in each wave, and we use data from 1977–2012.¹⁶ We estimate the income process at the household level. De Nardi et al. (2020) show that at the individual level, the PSID sample captures well the salient features of earnings dynamics documented in administrative social security data by Guvenen et al. (2021), who also resort to it for the analysis of wage and hours dynamics. Similarly, Arellano et al. (2017) estimate a rich earnings process using the PSID. Busch et al. (2022) document that the cyclical changes of the distribution of annual earnings changes in the PSID reflect the dynamics in social security data documented by Guvenen et al. (2014).¹⁷

Household net income is defined as household labor income plus public transfers minus taxes. As a measure of labor income, we use annual total labor income that includes income from wages and salaries, bonuses, and the labor part of self-employment income. We then sum up head and spouse annual labor income. We impute taxes using Taxsim, and add 50% of the estimated payroll taxes to the sum of head and spouse labor incomes to obtain pregovernment income. We aggregate public transfers to the household level and include measures of unemployment benefits, workers' compensation, combined old-age social security and disability insurance (OASI), supplemental security income, aid to families with dependent children (AFDC), food stamps, and other welfare.

We deflate all nominal values with the annual consumer price index, and select households of at least two adults if the household head is between 25 and 60 years of age. The minimum of our measure of household income and a pregovernment measure, which does not include the tax and transfer component, needs to be above a constant threshold, which is defined as the income from working 520 hours at half the minimum wage. Central moments (especially of higher order) are imprecisely estimated in small samples. We therefore estimate the moments for a given year and age group based on a sample from a five-year window over age within the year, which also smoothes the age profiles of these moments.

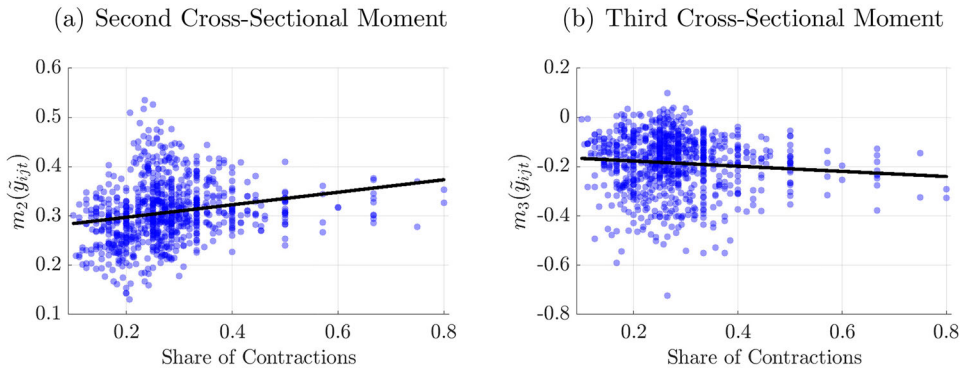
3.2.1. Defining business cycles. We classify years (from 1942 to 2012) as contractions or expansions. Starting from the NBER dating of peaks and troughs, we classify a year as a contraction if (i) it is completely in a contractionary period, which is defined as the time from peak to trough, (ii) if the peak is in the first half of the year and the contraction continues into the next year, and (iii) if a contraction started before the year and the trough is in the second half of the year. Given the sluggish synchronization of labor market outcomes with the macroeconomic indicators that the NBER takes into account (cf. Guvenen et al., 2014, Huggett and Kaplan, 2016), we expand the dating based on mean earnings of males in the PSID. We classify the following years as contractions: 1945, 1949, 1953, 1957, 1960, 1970, 1974, 1980–83, 1990–91, 2001–02, 2008–10. All years that are not classified as contractions are classified as expansions.

3.3. Estimation Results: Cyclical Idiosyncratic Risk.

3.3.1. Illustration of identification. In Figure 1, we plot the cross-sectional second and third central moments of residual household net income against the share of years classified as contractions out of all years a cohort went through since age 25 once reaching the given year. Each marker denotes a moment for households of some age j in some year t , $m_k(\tilde{y}_{ijt})$ for $k = 2, 3$. The pattern that emerges is that a higher share of contractionary years correlates

¹⁶ We do not use earlier waves because of poor coverage of income transfers before 1977. Although the change to a biannual survey frequency implies that households are observed only every other year after 1997, the estimated income process continues to be at the annual level. Thus, starting in 1997, two shocks realize between two observation periods.

¹⁷ Hryshko and Manovskii (2018) discuss some heterogeneity of income dynamics across PSID samples.



NOTES: Cross-sectional moments of residual income are net of age effects. Share of contractions for a given moment is the fraction of years classified as contraction since age 25. The slopes of the fitted lines are 0.13 and -0.11 for m_2 and m_3 , respectively. Moments for shares of 0 or 1 are not displayed for visualization reasons (they are used in the GMM estimation).

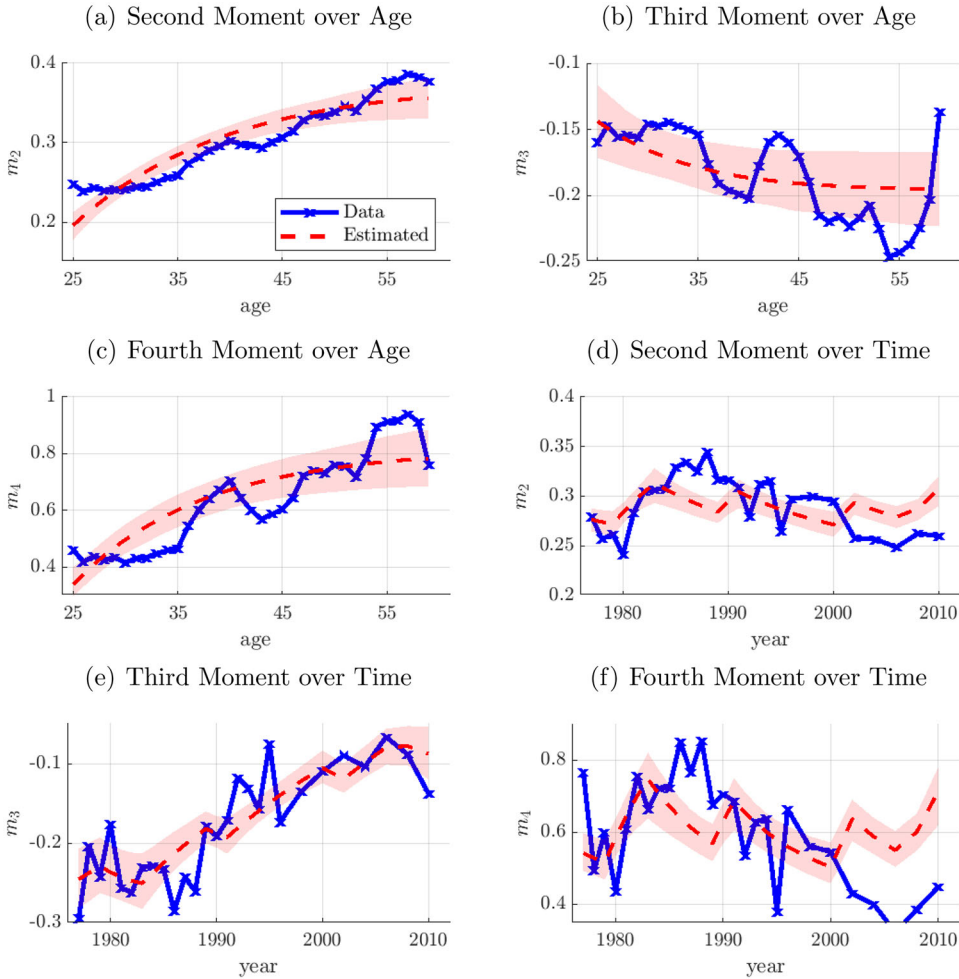
FIGURE 1

CROSS-SECTIONAL MOMENTS BY AGGREGATE HISTORY

positively with the cross-sectional second moments, and negatively with the cross-sectional third moments. These correlations identify the cyclicity of the moments of the shocks.

3.3.2. Estimation. We use the number of observations that contribute to an empirical moment as weights for the moment conditions, and this way assign more weight to those moments that are themselves estimated more reliably in the data. Daly et al. (2022) discuss the importance of the panel composition for consistency of estimates based on income changes versus cross-sectional moments. In the context of our quantitative life-cycle evaluation, it is crucial for us to match cross-sectional distribution moments over the life cycle. Further, the accumulated shocks captured by those cross-sectional moments identify our cyclicity estimates. Still, partly reflecting the insights from Daly et al. (2022), we include as additional moment conditions the averages over years of the second to fourth central moments of one- to five-year income changes. This ensures that the estimated income process is both, consistent with moments of the cross-sectional distribution and broadly consistent with moments of income changes. We give a collective weight of 10% to the latter.

In addition to the structure imposed so far, we hold the kurtosis of η fixed over the business cycle. Note that this does not mean that we restrict the fourth central moment to be acyclical: Let α_i denote the i th standardized moment, that is, $\alpha_i = \mu_i / \mu_2^{i/2}$. Assuming $\alpha_4^\eta(s(t)) = \alpha_4^\eta$ implies $\mu_4^{\eta,C} = \alpha_4^\eta (\mu_2^{\eta,C})^2$ and $\mu_4^{\eta,E} = \alpha_4^\eta (\mu_2^{\eta,E})^2$. In Appendix B.2, we show robustness with respect to this restriction. This leaves us with 12 parameters to estimate. We use moments (13a) and (14a) to estimate the variance parameters and the persistence ρ . Given an estimate for ρ , we then use moments (13b) and (14b) to estimate the third central moments. Likewise, given estimates for ρ and the variance parameters, we use moments (13c) and (14c) to estimate the fourth central moments. The third central moment of the cross-sectional distribution features a low-frequency change (see Panel (e) of Figure 2). To accommodate this in the estimation, and to not confound the estimated cyclicity, we add a linear trend to the third central moment of transitory shocks. This serves to detrend the moment, as we fit a stationary process. We report the time average of the implied moment. For inference, we apply a block bootstrap procedure and resample households to preserve the autocorrelation structure of the data. Table 1 shows the estimates, and Figure 2 illustrates the fit over age and time (further, Appendix B shows standardized moments).



NOTES: Moments are cross-sectional central moments. For each moment, age and year profiles are based on a regression of the moment on a set of age and year dummies. Blue lines: empirical moments; red dashed lines: theoretical moments implied by point estimates; shaded area denotes a 90% confidence band based on the bootstrap iterations.

FIGURE 2

FIT OF ESTIMATED PROCESS FOR HOUSEHOLD NET INCOME

TABLE 1
ESTIMATION RESULTS FOR HOUSEHOLD NET INCOME

Estimated Central Moments					
ρ	0.9683 [0.9463;0.9841]				
μ_2^X	0.1076 [0.0897;0.1237]	μ_3^X	-0.0508 [-0.0780; -0.0253]	μ_4^X	0.0173 [0.0000;0.0741]
μ_2^E	0.0752 [0.0677;0.0816]	μ_3^E	-0.0866 [-0.0935; -0.0771]	μ_4^E	0.2300 [0.1927;0.2664]
$\mu_2^{\eta,C}$	0.0223 [0.0152;0.0291]	$\mu_3^{\eta,C}$	-0.0167 [-0.0266; -0.0062]	$\mu_4^{\eta,C}$	0.0666 [0.0363;0.0847]
$\mu_2^{\eta,E}$	0.0085 [0.0044;0.0153]	$\mu_3^{\eta,E}$	-0.0013 [-0.0073; 0.0040]	$\mu_4^{\eta,E*}$	0.0098 [0.0022;0.0272]

NOTE: Estimated central moments for household income after taxes and transfers. Brackets show 5th and 95th percentiles of 1,000 bootstrap estimates (998 of the bootstrap iterations converge). * $\mu_4^{\eta,E}$ not separately estimated.

3.3.3. *Cyclical dispersion.* The first column of Table 1 reports the persistence of the AR(1) component of income along with the estimates of the variances of the components of the income process estimated jointly. The estimated persistence parameter (ρ) is 0.97. The estimates imply a countercyclical variance of persistent shocks: in aggregate downturns, the cross-sectional distribution of shocks is more dispersed. Our estimate of countercyclicality is quantitatively similar to the one estimated by Storesletten et al. (2004): the estimated standard deviation of persistent shocks is 61% higher in aggregate contractions.

3.3.4. *Cyclical skewness.* The second column of Table 1 reports the third central moments. We find that all shock components have negative third central moments, implying negative skewness of shocks. The third central moment of persistent shocks is significantly negative in contractions; the point estimate of the third central moment of persistent shocks in expansions is also negative, however, not statistically different from zero. The second and third central moments together translate into the third standardized moment, the coefficient of skewness, which is informative about the shape of the distribution. The cyclicity of the third central moment is stronger relative to the cyclicity of the second moment, which translates into the standardized moment displaying procyclicality. Thus, aggregate contractions are periods in which negative persistent shocks become relatively more pronounced.

3.3.5. *Excess kurtosis.* The third column of Table 1 reports the fourth central moments. The fixed effects are very imprecisely estimated; the point estimates imply relatively flat distributions (compared to a Normal distribution, which has a kurtosis of 3): the implied kurtosis coefficient at the point estimates is 1.5. The transitory and persistent shocks are estimated to display very pronounced excess kurtosis of about 41 and 134. Note that while these estimates of kurtosis seem very high at first glance, they imply a good fit of the cross-sectional distribution over age and over years, as shown in Figure 2. Furthermore, the estimated income process is in line with the average kurtosis of income changes.¹⁸

3.3.6. *Robustness.* In Appendix B.2, we present a range of alternative empirical specifications, such as estimating parameters jointly (instead of step-wise), varying the weight given to moments of changes, allowing for cyclicity of the transitory component, among others. The key insight is robustness of the core estimates, that is, those informing the model calibration—the persistence and distributional moments of the persistent component, and the moments of transitory shocks.

4. QUANTITATIVE ROLE OF HIGHER-ORDER RISK


4.1. Calibration of the Life-Cycle Model.

4.1.1. *Idiosyncratic income.* A model period is one year. Households start working at age 25 ($j = 0$) and retire at age 60 ($j = 35$). We calibrate the deterministic age profile e_j by the fitted age polynomial $f_{age}(j)$ of the first-stage estimation, where we abstract from heterogeneity by education, labor market experience, or household size by taking the weighted average of college and noncollege age profiles that display the usual hump-shaped pattern, cf. Appendix C.1. We normalize average productivity to one, that is, $\frac{1}{j_r} \sum_{j=0}^{j_r-1} e_j = 1$.

Next, for each stochastic component (i.e., $\varepsilon, \eta(E), \eta(C)$), we use a flexible generalized lambda distribution (FGLD). The FGLD is developed in Freimer et al. (1988), and can be characterized by its quantile function

$$(15) \quad Q(p; \lambda) = F^{-1}(p; \lambda) = x = \lambda_1 + \frac{1}{\lambda_2} \left(\frac{p^{\lambda_3} - 1}{\lambda_3} - \frac{(1-p)^{\lambda_4} - 1}{\lambda_4} \right),$$

¹⁸ We also estimate the process on pregovernment income, cf. Appendix B.1.



NOTES: Discretized log distribution functions for the persistent shock η . LKSW: FGLD with estimated variance, skewness, and kurtosis. Markers denote the grid points used in the discretized distribution. Normal: Normal distribution with estimated variance discretized using Gaussian quadrature method. Log density is the base 10 logarithm of the PDF.

FIGURE 3

DISCRETIZED LOG DISTRIBUTION FUNCTIONS: PERSISTENT SHOCK

where p denotes a percentile and λ is a vector of four parameters with location parameter λ_1 , scale parameter λ_2 , and tail index parameters λ_3, λ_4 .¹⁹ The four parameters of the FGLD can be calibrated well to match the first four central moments of a distribution (cf. Lakhany and Mausser, 2000; Su, 2007). Calibration targets for each shock $x \in \{\varepsilon, \eta(E), \eta(C)\}$ are the central moments $\{\hat{\mu}_i^x\}_{i=2}^4$ of distributions F_ε and $F_\eta(s)$, estimated in the second stage; we choose the first moment such that the expected value in levels is normalized to one. We then discretize by spanning equidistant grids for the respective random variable $x \in \{\varepsilon, \eta(s)\}$ and by assigning to each grid point probabilities from the integrated probability density function of the distribution (details in Appendix C.1).

We consider two alternative parameterizations of the FGLD to which we refer as *distribution scenarios*. The first scenario, LKSW, features leptokurtic and left-skewed shock distributions with the estimated second, third, and fourth central moments.²⁰ Panel (a) in Figure 3 shows the log density functions of the persistent shock $\eta(s)$ in contractions and expansions. The second scenario, NORM, features FGLD distributions with the first four central moments of a Gaussian distribution with the estimated variance. Panel (b) shows the corresponding Gaussian distributions for the persistent component.²¹

Appendix C.1 reports the estimated, fitted, and discretized moments of the shocks, as well as the parameter vectors λ for all shocks under the two scenarios NORM and LKSW.²² It also shows central moments 2–4 of cross-sectional income in logs and levels that result from our parameterization.

¹⁹ The parametric constraints are $\lambda_2 > 0$, and $\min\{\lambda_3, \lambda_4\} > -\frac{1}{4}$.

²⁰ We also impose a minimum household net income that remains unchanged across scenarios, which turns out to be nonbinding in either scenario.

²¹ Note that the FGLD distribution NORM is not exactly identical to a Gaussian, but with no quantitative implications: any differences between FGLD distribution NORM and its Gaussian counterpart are captured by the even central moments higher than the fourth (the sixth, eighth, tenth, etc.). Following our analytical analysis in Section 2.1, these mechanically become less and less important with the order. As documented in Appendix D.1, all quantitative results for a Gaussian distribution are numerically almost identical to those obtained for FGLD distribution NORM, which we choose to compare to benchmark LKSW given its nesting in the FGLD distribution.

²² Following Huggett and Kaplan (2016), we assume that one-third of the estimated variance of the transitory shock is measurement error and reduce the targeted variance accordingly. We assume that this measurement error is symmetric and accordingly adjust the third and fourth central moments such that the implied coefficients of skewness and kurtosis are unchanged.

4.1.2. *Pension system.* The model pension system approximates the current U.S. bend point formula, in which we approximate individual average indexed monthly earnings (AIME) by the realization of the persistent income shock before entering into retirement $z_{j,t-1}$. This yields the model equivalent to the primary insurance amount (PIA), denoted by $p(z_{j,t-1})$, which we scale by ϱ to obtain pension benefits as $b(z_{j,t-1}) = \varrho \cdot p(z_{j,t-1})$. Pension contributions obey a linear schedule with the contribution rate calibrated to 11.7%. By our normalization of income, aggregate contributions are constant across scenarios. We calibrate ϱ so that the pension system clears, which implies that average pension income is also constant across scenarios, and thus, we rule out changes in life-cycle savings motives. Appendix C.2 provides the details of the contribution scheme, the pension budget, and the calibrated values of ϱ .

4.1.3. *Aggregate shock process.* Based on our classification of time periods as contractions and expansions for the U.S. economy, we estimate a Markov transition process on this data. We estimate $\pi(s' = E|s = E) = 0.7885$ and $\pi(s' = C|s = C) = 0.3889$, implying the stationary invariant distribution $\Pi_s = [0.2571, 0.7429]'$.

4.1.4. *Initial assets and interest rate.* We assume that all households are born with the same initial assets $a_0 = \bar{a}_0$. We compute those from the average asset to net earnings data at age 25, which we calculate from PSID data as 0.89. We set the annual interest rate of the risk-free asset to $r = 3\%$ as in Kaplan and Violante (2010). This choice is consistent with a risk-free rate of about 4% as in Siegel (2002) and a technological growth rate of 1% (given that there is no technological growth in our model economy).

4.1.5. *Preferences.* We consider risk-sensitive preferences Tallarini (2000)—and accordingly, set the intertemporal elasticity of substitution to $\gamma = 1$ —and four alternative parameterizations for $\theta \in \{1, 2, 3, 4\}$. Based on the estimates in Cooper and Zhu (2016), we take $\theta = 4$ as our benchmark parameterization.²³ For each $\theta \in \{1, 2, 3, 4\}$, we determine the discount rate $\rho = 1/\beta - 1$ by targeting the average life-cycle asset profile scaled by net earnings, which we compute from PSID data for the sample which we use in the income process estimation. Since our model is not designed to match saving patterns in retirement (there is neither survival risk nor a bequest motive), we match assets for ages 25–60, the working period in our model. This calibration is done for distribution scenario LKSW, and we then hold the calibrated discount rate constant when moving to scenario NORM, for each calibration of θ .²⁴

Calibrated discount rates range from 1.86% for $\theta = 1$ to 2.60% for $\theta = 4$, see Table 2. The reason for the positive relationship of the calibrated discount rates and θ is that stronger risk attitudes (larger θ) imply higher precautionary savings that is offset in the calibration to the asset target by a larger ρ so that the life-cycle savings motive is less potent.

4.1.6. *Alternative calibration.* As the main robustness analysis, we consider an alternative calibration of our model economy, following the benchmark calibration in Kaplan and Violante (2010), which we refer to as KV-calibration. There are two differences relative to the baseline calibration described above. First, we set initial assets to zero, and second, we calibrate the discount rate targeting an aggregate capital–output ratio of 2.5—the ratio implicitly targeted in our baseline calibration is about 5.4. As we show in Section 4.5, the welfare implications of higher-order risk are larger in the KV-calibration.

4.2. *Welfare Effect of Higher-Order Income Risk.* First, we assess the welfare effect of higher-order income risk per se—that is, we translate the deviation from Gaussian shock dis-

²³ Cooper and Zhu (2016) estimate a risk aversion of 4.4 and an IES of 0.6. We choose an IES of 1 as a natural benchmark. This is also very convenient when we decompose the welfare effects as described in Appendix A.2.

²⁴ Recalibrating the discount rate under scenario NORM does not alter the results by a relevant margin as reported in Appendix D.3 for completeness.

TABLE 2
CALIBRATED PARAMETERS

Working Period	25 ($j = 0$) to 60 ($j = j_r - 1$)
Maximum Age	80
Elasticity of intertemporal of substitution	$\gamma = 1$
Coefficient of risk aversion	$\theta \in \{1, 2, 3, 4\}$
Discount rate (baseline)	$\rho \in \{0.0186, 0.0205, 0.0229, 0.0260\}$
Discount rate (KV-calibration)	$\rho \in \{0.0361, 0.0462, 0.0633, 0.0919\}$
Interest rate	$r = 0.03$
Pension contribution rate	$\tau^p = 0.117$
Pension benefit level	ϱ shown in Table C.4
Average tax rate	$\tau = 0.168$
Aggregate shocks	$\pi(s' = C s = C) = 0.3889, \pi(s' = E s = E) = 0.7885$
Initial ass. / inc.	$\bar{a}_0 = 0.89$

NOTE: The discount rate ρ is calibrated endogenously to match the average asset-to-income life-cycle profile from the PSID (baseline) or the aggregate capital-output ratio of 2.5 (KV). The pension benefit level is calibrated to clear the pension budget.

tributions toward skewed and leptokurtic distributions into welfare consequences. To this end, we define the social welfare function from an ex ante perspective as the certainty equivalent of being born into the economy under a given distribution scenario. We express the welfare function as a function of the stochastic consumption stream $\mathbf{c}^i(a)$, which denotes optimal consumption at every age for every possible history, given starting out with asset holdings a , when facing shock distribution scenario i : $W^i(\mathbf{c}^i(a)) = v(V_0^i(a, z, \varepsilon; s))$ for $i \in \{\text{NORM}, \text{LKSW}\}$. The certainty equivalent is given by (8), and the expectation operator within $v(\cdot)$ is taken over the idiosyncratic shocks and the aggregate state. In particular, agents draw a shock ε from distribution \tilde{F}_ε , and an initial persistent shock η from conditional distribution $\tilde{F}_\eta(s)$, whereas $s \in \{E, C\}$ is drawn from the stationary distribution of the aggregate state, Π_s .

We evaluate the welfare function for households starting with average assets, that is, $W^i(\mathbf{c}^i(a = \bar{a}_0))$. Given the optimal consumption plans under each distribution scenario, we then calculate the CEV that households need to receive in scenario NORM to be indifferent to scenario LKSW. Homogeneity of the utility function implies that $W^i((1 + g_c)\mathbf{c}^i(a = \bar{a}_0)) = (1 + g_c)W^i(\mathbf{c}^i(a = \bar{a}_0))$, where g_c proportionately varies consumption in every state and period. From $W^{\text{LKSW}}(\mathbf{c}^{\text{LKSW}}(a = \bar{a}_0)) = W^{\text{NORM}}((1 + g_c)\mathbf{c}^{\text{NORM}}(a = \bar{a}_0))$, the CEV follows as $g_c = W^{\text{LKSW}}/W^{\text{NORM}} - 1$, where $W^i = W^i(\mathbf{c}^i(a = \bar{a}_0))$.

Table 3 contains four panels for the different calibrations of risk attitudes. Within each panel, the first row shows the CEV along with a decomposition that we explain below. Focus first on the total CEV. In line with the discussion for the two-period model, welfare losses show up for $\theta > 1$, ranging from about -0.5% for $\theta = 2$ to -14.5% for $\theta = 4$. For weak risk attitudes ($\theta = 1$), higher-order risk leads to welfare gains (CEV of 0.4%). The latter results from the fact that we consider mean-preserving changes of risk: a mean-preserving increase (in absolute terms) of the negative third central moment implies an increase of the mean in logs, and under log utility, this translates into welfare gains. We formalize this in Proposition 1 in Appendix A.3.

To understand the mechanism behind the welfare effect, we next decompose the CEV into three components. When facing different income risk, households make different savings decisions (depending on their risk attitudes). The first consequence of this is potentially different (overall) mean consumption, that is, consumption averaged cross-sectionally and over age. We call the welfare consequence of this change of mean consumption the *mean effect*, g_c^{mean} . The remaining welfare effect stems from a change of the distribution around this mean (and we refer to this welfare component as the *distribution effect*, g_c^{distr}).

The change of the distribution consists of two components. First, a change over age: the profile of average consumption over the life-cycle changes, which gives rise to a *life-cycle distribution effect*, g_c^{lcd} . Second, at a given age, the cross-sectional distribution of consumption

TABLE 3
WELFARE EFFECTS OF (CYCLICAL) IDIOSYNCRATIC RISK

	g_c	g_c^{mean}	g_c^{lcd}	g_c^{csd}	Δg_c^{cr}
		Risk Aversion, $\theta = 1$			
NORM→LKSW	0.384	-0.112	0.482	0.015	-
NORM: cyclical risk	-1.918	0.360	-2.142	-0.136	-
LKSW: cyclical risk	-1.620	0.289	-1.793	-0.116	0.298
		Risk Aversion, $\theta = 2$			
NORM→LKSW	-0.504	-0.106	-0.365	-0.033	-
NORM: cyclical risk	-3.622	0.651	-3.998	-0.274	-
LKSW: cyclical risk	-3.943	0.605	-4.261	-0.287	-0.321
		Risk Aversion, $\theta = 3$			
NORM→LKSW	-5.193	0.270	-5.085	-0.378	-
NORM: cyclical risk	-5.106	0.894	-5.594	-0.406	-
LKSW: cyclical risk	-8.128	1.027	-8.514	-0.641	-3.022
		Risk Aversion, $\theta = 4$			
NORM→LKSW	-14.450	0.957	-14.153	-1.254	-
NORM: cyclical risk	-6.373	1.108	-6.961	-0.520	-
LKSW: cyclical risk	-13.976	1.460	-14.196	-1.240	-7.603

NOTE: Welfare gains (positive numbers) and losses (negative numbers) of higher-order income risk, expressed as a consumption equivalent variation (CEV) in percentages in scenario NORM that makes households indifferent to the higher-order income risk scenario LKSW. Also CEV in the noncyclical scenario that makes households indifferent to the cyclical scenario. g_c : total CEV, g_c^{mean} : CEV from changes of mean consumption, g_c^{lcd} : CEV from changes of consumption over the life cycle, g_c^{csd} : CEV from changes in the cross-sectional distribution, where $g_c = g_c^{mean} + g_c^{lcd} + g_c^{csd}$. $\Delta g_c^{cr} = g_c^{LKSW,cr} - g_c^{NORM,cr}$.

around the age-specific average changes, which gives rise to a *cross-sectional distribution effect*, g_c^{csd} . The overall CEV is the sum of the three components: $g_c = g_c^{mean} + g_c^{lcd} + g_c^{csd}$ (cf. Appendix A.2 for explicit expressions), which are shown in columns 2–4 of Table 3.

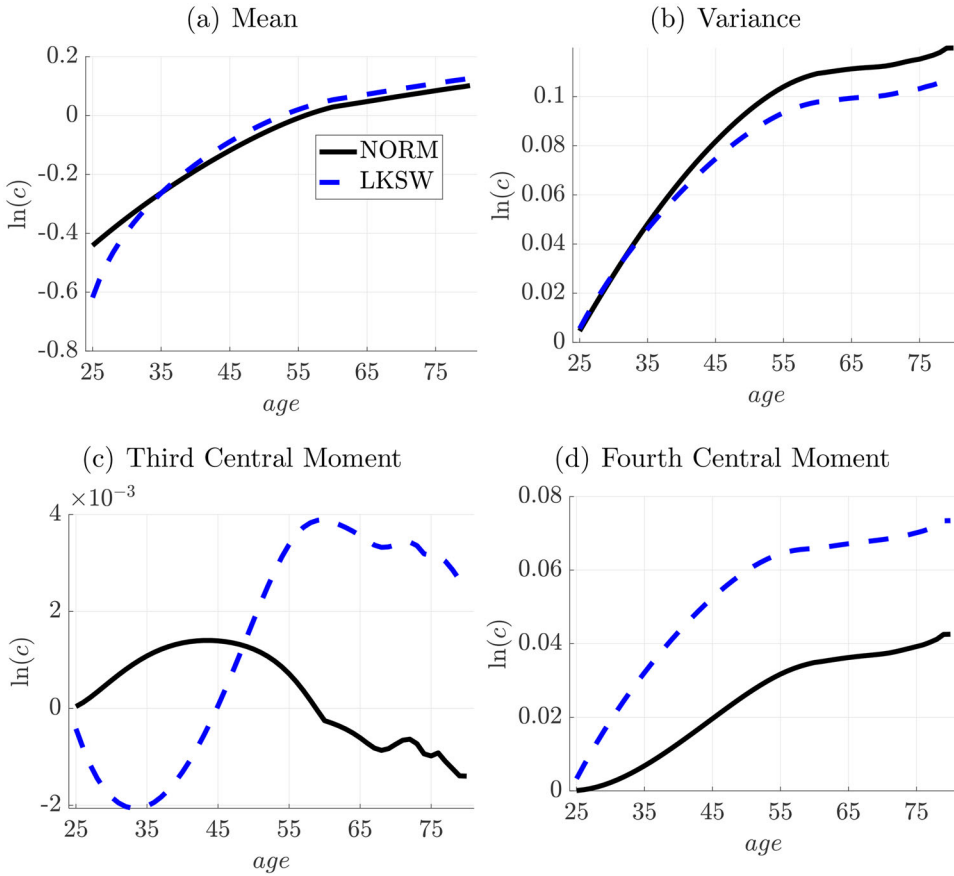
The dominant channel behind the welfare results turns out to be the different life-cycle consumption profile reflected in g_c^{lcd} . Driven by the precautionary savings motive, young households choose lower consumption when facing scenario LKSW compared to scenario NORM. The average life-cycle consumption profiles under the two scenarios are shown in Panel (a) of Figure 4.²⁵ In welfare terms, lower young-age consumption dominates higher old-age consumption due to discounting. For $\theta = 3$ and $\theta = 4$, the mean effect g_c^{mean} is positive because higher old-age consumption dominates.

Panels (b)–(d) of Figure 4 show the second to fourth cross-sectional central moments of the consumption distribution over the life cycle, which are responsible for the cross-sectional distribution effect g_c^{csd} . To interpret it, note that the variance is lower in scenario LKSW than in scenario NORM for most ages, whereas the third central moment is initially negative and the fourth central moment is higher at all ages.²⁶ Lower variance contributes positively to g_c^{csd} , which dominates for weaker risk attitudes, whereas the more negative third and higher fourth central moment contribute negatively, which dominates for stronger risk attitudes.

4.2.1. *Decomposing moments.* In Appendix D.2, we analyze an additional distribution scenario, which features symmetric (in logs) shocks with the excess kurtosis of LKSW. For cali-

²⁵ The corresponding asset profile is shown in Appendix C.3. Qualitatively, effects are the same in the other calibrations of θ . Consumption is monotonically increasing due to a strong life-cycle savings motive; we address this in Section 4.5.

²⁶ The Gini coefficient for assets for a risk aversion of 4 is at 0.36 in scenario NORM, and at 0.34 in scenario LKSW: the introduction of higher-order income risk does not increase the Gini coefficient in a quantitative model such as ours. Also, note that the Gini coefficient in our calibrated model is substantially lower than in the data and also lower than what is typically found in quantitative work; for example, Krueger and Ludwig (2016) compute a Gini coefficient of assets of 0.55 in an overlapping generations model calibrated to the U.S. economy. Main reason for the relatively modest asset inequality lies in our focus on ex post heterogeneity. In our alternative KV-calibration, the Gini coefficient for assets is 0.50 under scenario NORM and 0.38 under scenario LKSW.



NOTES: Moments of cross-sectional distribution of log consumption over the life cycle under scenarios NORM and LKSW.

FIGURE 4
CENTRAL MOMENTS OF LOG CONSUMPTION BY AGE ($\theta = 4$)

bration $\theta = 4$, the welfare costs are 10.9% (see Table D.2), which implies that the kurtosis accounts for about 75% of the welfare costs of higher-order risk, and thus, skewness accounts for about 25%.

4.3. *Welfare Costs of Cyclical Idiosyncratic Risk.* Second, we quantify the utility consequences of *cyclical idiosyncratic risk* under the different scenarios. For each $i \in \{NORM, LKSW\}$, let $W^{i,ncr}$ denote the social welfare function in a (counterfactual) *no cyclical risk scenario*, in which we shut down the cyclical variation of the shock distributions. We then compute the CEV that makes households in this counterfactual scenario ex ante indifferent to being born into the (actual) scenario with cyclical risk, $g_c^{i,cr} = W^i/W^{i,ncr} - 1$. By holding mean wages and interest rates constant over the cycle, the welfare effects of cyclical risk we report constitute a lower bound for each scenario.²⁷

When shutting down cyclical risk, we assume that households always draw from the *expansion-distribution* of the scenario. When using lognormal distributions of shocks, one

²⁷ Note that the direct effect of business cycles is typically found to be small. For example, Storesletten et al. (2001) find the direct effect to be an order of magnitude smaller than the role of cyclical variation in idiosyncratic risk. However, there can be indirect utility “interactions” between aggregate and idiosyncratic risk, which may be large Harenberg and Ludwig (2019), and which we abstract from in order to focus on the role of the idiosyncratic shock distribution.

approach in the literature is to consider an *average* distribution, which features the average of expansion and contraction variances, see, for example, Storesletten et al. (2001). This approach is not applicable in our analysis as it is conceptually not clear what characterizes such an *average* distribution once other moments than the variance are taken into account. To the extent that some average distribution represents a better noncyclical counterfactual scenario, the pure effect of cyclical idiosyncratic risk is overstated in our analysis.²⁸ However, we are mainly interested in the difference of welfare costs of cyclical income risk across scenarios, that is, the *difference-in-difference* comparison between $g_c^{LKSW,cr}$ and $g_c^{NORM,cr}$, that is, $\Delta g_c^{cr} = g_c^{LKSW,cr} - g_c^{NORM,cr}$. Thus, our approach to shutting down cyclical idiosyncratic risk is of second-order importance as it is consistent across scenarios.²⁹

Table 3 reports the resulting overall welfare costs of cyclical idiosyncratic risk in scenarios NORM and LKSW, and the decomposition of the total CEV into its components for each calibration of θ . In both scenarios $i \in \{NORM, LKSW\}$, we observe that, first, the welfare costs of business cycles, $g_c^{i,cr}$, increase monotonically in θ . Second, the effect is dominated by the life-cycle distribution effect $g_c^{i,cr,lcd}$, which reflects the change of the mean consumption profile over age: agents face higher risk to which they respond by increasing savings early in the life cycle. Third, higher savings increase expected consumption in the middle of the life cycle, which pushes up overall mean consumption and is reflected in a positive mean effect $g_c^{i,cr,mean}$.


Consistent with the previous result that with logarithmic utility the total welfare effect from higher-order income risk is positive, we now find that welfare losses from cyclical idiosyncratic risk are about 0.3 percentage points lower in scenario LKSW. Similarly, with moderate risk attitudes (risk aversion of 2), the welfare losses of cyclical income risk in scenario LKSW are only mildly higher than those obtained in scenario NORM. With strong risk attitudes ($\theta = 4$), the welfare losses are significantly (about 7.6 percentage points) higher in scenario LKSW compared to scenario NORM. Thus, the welfare effects of cyclical risk are strongly underestimated in conventional approaches based on Gaussian distributions of innovations if risk attitudes are strong.

4.3.1. Decomposing moments. In the calibration with $\theta = 4$, welfare costs of cyclical risk are about 4.9 percentage points higher in the distribution scenario with excess kurtosis but zero skewness than in scenario NORM (see Table D.2). Combined with the lower part of Table 3, we thus find that the differential welfare losses from higher-order risk approximately 64% ($\approx 4.885/7.608 \cdot 100$ [%]) are due to the excess kurtosis and the remaining 36% are due to the left-skewness of shocks.

4.4. Self-Insurance and the Propensity to Consume. We extensively used the concept of increased precautionary savings—that is, self-insurance—in response to higher-order risk to explain the ex ante welfare effects that largely stem from reduced consumption at early ages. Now, we study how this translates into measures of self-insurance against income shocks $x_j(s) \in \{\varepsilon_j, \eta_j(s)\}$. We start by employing a measure of consumption insurance as introduced in the empirical literature by Blundell et al. (2008), which is based on estimating the pass-through of shocks to consumption adjustments. As all shocks are observed within the

²⁸ In one of our sensitivity checks in Appendix D.3, we consider CRRA preferences with $\theta = 2$, and in scenario NORM, obtain welfare costs of about 2.9%. With the same specification for preferences, Storesletten et al. (2001) find welfare costs of cyclical (Gaussian) risk of about 1.3%. Besides other differences between our model and theirs, one reason for the higher welfare costs in our analysis lies in the different approach to characterizing the noncyclical scenario.

²⁹ One alternative is to follow the “integrating out” principle (see Krusell and Smith, 1999 and Krusell et al., 2009), which first isolates a true idiosyncratic component of the shock, and then integrates over the probability distribution of the aggregate state.



NOTES: Figures show the degree of consumption insurance against transitory and persistent shocks separately by age.

FIGURE 5

INSURANCE COEFFICIENTS: STRONG RISK ATTITUDES, $\theta = 4$

TABLE 4
AGGREGATE PASS-THROUGH AND ITS DECOMPOSITION, $\theta = 4$

	Components of $1 - \phi$:				Conditional $1 - \phi$:	
	$1 - \phi^\varepsilon$	$\frac{\pi^- A^\varepsilon}{1 - \phi^\varepsilon}$	$\frac{\pi^+ B^\varepsilon}{1 - \phi^\varepsilon}$	$\frac{C^\varepsilon}{1 - \phi^\varepsilon}$	$1 - \phi^\varepsilon \varepsilon^-$	$1 - \phi^\varepsilon \varepsilon^+$
<i>Transitory:</i>						
NORM	0.051	-0.014	0.883	0.131	0.041	0.063
LKSW	0.045	0.303	0.525	0.172	0.030	0.090
<i>Persistent:</i>						
	$1 - \phi^\eta$	$\frac{\pi^- A^\eta}{1 - \phi^\eta}$	$\frac{\pi^+ B^\eta}{1 - \phi^\eta}$	$\frac{C^\eta}{1 - \phi^\eta}$	$1 - \phi^\eta \eta^-$	$1 - \phi^\eta \eta^+$
NORM	0.448	0.409	0.576	0.015	0.463	0.494
LKSW	0.402	0.523	0.454	0.023	0.350	0.511

NOTE: Column 1 shows the aggregate consumption pass-through coefficient, and columns 2–4 its decomposition into components according to Equation (17), expressed as shares of total pass-through. π^- and π^+ are short for the probabilities of the shock being negative or positive. Columns 5 and 6 show conditional pass-through coefficients for negative and positive shocks.

model, we calculate the model-consistent insurance coefficient following Kaplan and Violante (2010).³⁰

Conditional on aggregate state s and age j , the pass-through coefficient $1 - \phi_j^x(s)$ is the coefficient of a linear regression of consumption growth on shock x :

$$(16) \quad 1 - \phi_j^x(s) = \frac{\text{cov}(\Delta \ln(c_{j+1}(s' | s)), x_{j+1}(s'))}{\text{var}(x_{j+1}(s'))},$$

where $\Delta \ln(c_j(s' | s)) = \ln(c_{j+1}(s' | s)) - \ln(c_j(s))$, and $\phi_j^x(s)$ denotes the insurance coefficient.

Figure 5 reports the insurance coefficients ϕ_j^x for all ages $j \in \{0, \dots, J\}$, as a weighted average of the coefficients in contractions and expansions (using the stationary invariant distribution Π_s), for the transitory shock ε in Panel (a) and for the persistent shock $\eta(s)$ in Panel (b). Results are quantitatively similar for different values of risk attitudes, so we focus on $\theta = 4$. In the first column of Table 4, we report the pass-through aggregated over age. All in all, in

³⁰ Thus, we do not consider the extent to which an empirical measure is valid, given that shocks are not observed in observational data and need to be extracted using some identifying assumptions, as discussed in, for example, Blundell et al. (2008), Kaplan and Violante (2010), or Commault (2022).

scenario LKSW, consumption insurance against both transitory and persistent shocks as measured by the ϕ -coefficients is higher than in scenario NORM. This is in line with results presented in De Nardi et al. (2021).

One part of the higher insurance coefficient is explained directly by higher precautionary savings, which imply a lower pass-through of income gains to consumption, as we discuss in Section 2.1 for the two-period setting. To explore whether the higher insurance coefficient represents *better insurance* against income losses, we consider the following decomposition, which captures the contribution of consumption comovement with positive and negative shocks, respectively, to the aggregate pass-through coefficient for shock $x \in \{\eta, \varepsilon\}$:

$$(17) \quad 1 - \phi^x = \frac{E[\Delta \ln(c(\cdot))x] - E[\Delta \ln(c(\cdot))]E[x]}{\text{var}(x)} = \underbrace{\pi^- \cdot \frac{E[\Delta \ln(c(\cdot))x|x < 0]}{\text{var}(x)}}_{A^x} + \underbrace{\pi^+ \cdot \frac{E[\Delta \ln(c(\cdot))x|x > 0]}{\text{var}(x)}}_{B^x} - \underbrace{\frac{E[\Delta \ln(c(\cdot))]E[x]}{\text{var}(x)}}_{C^x}.$$

The terms π^- and π^+ are short for the probabilities of the shock being negative or positive, respectively. The decomposition of the aggregate coefficient expressed as shares is shown in Table 4 for $\theta = 4$ (in Appendix Table D.4, we show the same table for $\theta = 2$; the other calibrations of θ yield the same patterns). In scenario NORM, the comovement of consumption changes with negative transitory shocks plays only a negligible role (-1.4%) for the overall pass-through coefficient. On the other hand, in scenario LKSW, the (negative) consumption reaction to negative transitory shocks accounts for sizable 30.3% of the pass-through coefficient. A possible rationalization is that built-up savings do not suffice to smooth out the negative shocks in scenario LKSW as well as they do in scenario NORM. The nonzero third term in the decomposition follows from the normalization in levels, which implies a nonzero mean of the shocks in logs (which is larger in scenario LKSW than in scenario NORM). Turning to persistent shocks, in scenario NORM, 40.9% of the pass-through is accounted for by consumption reductions in response to negative shocks, whereas 57.6% come from consumption increases with positive shocks. A larger fraction (52.3%) of the overall smaller pass-through is accounted for by consumption reductions in response to negative shocks in scenario LKSW.

The last two columns of Table 4 show pass-through coefficients conditional on the sign of the shocks. For example, $1 - \phi^\varepsilon|\varepsilon^- = \text{cov}(\Delta \ln(c(\cdot)), \varepsilon|\varepsilon < 0)/\text{var}(\varepsilon|\varepsilon < 0)$, and likewise for positive shocks and for η . For both transitory and persistent shocks, under scenario LKSW, the conditional pass-through is larger for positive shocks and smaller for negative shocks, compared to scenario NORM.³¹

To further interpret these findings, note that the decomposition captures three effects. First, the shock distributions differ, and thus, do the probabilities of positive and negative shocks, and the (typical) size of the shocks of either sign. Second, the equilibrium asset distribution is different—and the amount of assets matters for the optimal consumption-savings choice. Third, for a given amount of assets, the reaction to a shock of the same magnitude is different.

For better comparability across distribution scenarios, we now first turn to the MPC, defined as the slope of the consumption function. Table 5 reports the average MPC for different levels of risk attitudes. It turns out that the aggregate marginal consumption response is somewhat larger in scenario LKSW.

Next, in Table 6, we consider explicitly the consumption response to transitory and persistent shocks of different signs, and of three sizes. Small shocks are defined as one standard deviation of the persistent shock (0.11 based on the weighted variances in expansions and con-

³¹ Comparison of the coefficients is not straightforward as the conditional variance of the shocks changes. Consider negative transitory shocks: under scenario NORM, the conditional covariance of shocks with consumption changes is 0.0009, whereas it is an order of magnitude larger (0.0085) under LKSW, which displays a larger conditional variance of negative shocks.

TABLE 5
MARGINAL PROPENSITY TO CONSUME

	Risk Aversion			
	$\theta = 1$	$\theta = 2$	$\theta = 3$	$\theta = 4$
NORM	0.0390	0.0410	0.0433	0.0461
LKSW	0.0390	0.0417	0.0463	0.0529

NOTE: MPC is calculated as the slope of the consumption policy function. Reported value aggregates over the distribution of households (over age, income, and assets).

TABLE 6
PROPENSITY TO CONSUME ($\theta = 4$)

		Negative Shocks			Positive Shocks		
		Small	Medium	Large	Small	Medium	Large
<i>Transitory:</i>	NORM	0.0460	0.0460	0.0460	0.0460	0.0459	0.0459
	LKSW	0.0528	0.0530	0.0533	0.0525	0.0523	0.0519
<i>Persistent:</i>	NORM	0.3991	0.4026	0.4045	0.3914	0.3868	0.3704
	LKSW	0.3847	0.3889	0.3981	0.3800	0.3748	0.3594

NOTE: Propensity to consume out of a given shock. The reported numbers are evaluated at the mean asset profile over age from scenario LKSW.

tractions); medium shocks as one standard deviation of the transitory shock (0.22), and large shocks as twice that (0.44). For both transitory and persistent shocks, we show the consumption reaction to a shock of a given sign and size for households with age-specific average income and the (baseline) asset holdings from the LKSW calibration, aggregated over age.³² Thus, the difference between NORM and LKSW represents changes in the policy function. For transitory shocks, we find that the propensity to consume is larger in scenario LKSW, reflecting the somewhat larger MPC mentioned already. In particular, negative shocks (of the same size) translate somewhat more strongly into consumption reductions than in scenario NORM. For persistent shocks, the consumption responses are weaker in scenario LKSW compared to scenario NORM. The reason for this direction of the effect is that for persistent shocks, the precautionary savings motive is more pronounced. As we discussed in Section 2.1, a stronger savings motive reduces the propensity to consume out of a given shock.³³

Summing up, for both transitory and persistent shocks, we observe a smaller pass-through (increase of the insurance coefficient) in scenario LKSW than in scenario NORM. In the case of the persistent shock, this is driven by a reduced propensity to consume, and in this sense reflects better insurance. In the case of transitory shocks, however, negative shocks actually translate more into consumption, which is also reflected in a higher MPC. The reduction of the pass-through therefore does not reflect better insurance against a same-size shock for an observationally equivalent household. Thus, the simple aggregate insurance coefficient based on comovements of shocks and consumption is an imprecise measure of insurance, if one plausibly has in mind that better insurance means that negative shocks translate less into consumption—and that this is what should be reflected by a larger value of an insurance coefficient.

4.4.1. *Decomposing moments.* Table D.3 shows that the difference between scenarios NORM and LKSW of the relative importance of positive and negative shocks for the over-

³² Ghosh and Theloudis (2023) estimate a nonlinear consumption function on biannual PSID data and find patterns of the propensity to consume over shock sign and size that qualitatively resemble our model patterns.

³³ Considering the propensity to consume out of a persistent shock in the two-period framework, one can show that there is an additional term added to (3) that increases in the persistence of the shock; results are available upon request.

TABLE 7
WELFARE ANALYSIS IN KV-CALIBRATION

	g_c	g_c^{mean}	g_c^{lcd}	g_c^{csd}	Δg_c^{cr}
		Risk Aversion, $\theta = 2$			
NORM→LKSW	-4.008	0.329	-4.075	-0.262	-
NORM: cyclical risk	-2.472	0.898	-3.288	-0.082	-
LKSW: cyclical risk	-4.397	1.088	-5.406	-0.079	-1.925
		Risk Aversion, $\theta = 4$			
NORM→LKSW	-25.553	2.563	-27.540	-0.576	-
NORM: cyclical risk	-3.081	1.315	-4.269	-0.127	-
LKSW: cyclical risk	-16.958	2.170	-18.666	-0.462	-13.877

NOTE: Welfare gains (positive numbers) and losses (negative numbers) of higher-order income risk, expressed as a consumption equivalent variation (CEV) in percentages in scenario NORM that makes households indifferent to the higher-order income risk scenario LKSW. Also, CEV in the noncyclical scenario that makes households indifferent to the cyclical scenario. g_c : total CEV, g_c^{mean} : CEV from changes of mean consumption, g_c^{lcd} : CEV from changes of consumption over the life cycle, g_c^{csd} : CEV from changes in the cross-sectional distribution, where $g_c = g_c^{mean} + g_c^{lcd} + g_c^{csd}$. $\Delta g_c^{cr} = g_c^{LKSW,cr} - g_c^{NORM,cr}$.

all pass-through is driven by the introduction of left-skewness—under scenario LK, the shares are very close to scenario NORM.

4.5. *Alternative Calibrations.* We start by considering the alternative calibration a la Kaplan and Violante (2010) (KV-calibration), where we set initial assets to zero, and calibrate the discount factor targeting an aggregate capital–output ratio of 2.5. Relative to our baseline calibration, the endogenously calibrated discount rate now exceeds the (exogenous) interest rate, cf. Table 2, which in isolation would lead to a consumption profile negatively sloped over age. As households engage in precautionary savings, they save when young to build up a buffer against negative income risk. Together, the two forces lead to a hump-shaped consumption profile (see Figure D.2 and the corresponding asset profile in Figure C.2).

Table 7 summarizes the results of the two welfare analyses for $\theta = 2$ and $\theta = 4$. For each value of θ , the first row shows the welfare costs of higher-order risk, which turn out substantially higher under the KV-calibration. To insure themselves against the higher risk in scenario LKSW, households save relatively much at young ages, which tilts the average consumption profile compared to the one in scenario NORM. The difference between the two profiles is stronger than in our baseline calibration, which is driven by (exogenously set) zero initial assets and an (endogenously determined) higher discount rate. This interpretation is further supported by investigating the fraction of borrowing constrained households. In the baseline calibration, this fraction is basically zero in both risk scenarios with high-risk attitudes of $\theta = 4$. In the KV calibration of scenario NORM with a risk aversion of $\theta = 4$, about 16% of households are initially (at biological age 25) borrowing constrained. In scenario LKSW, this fraction is substantially lower, at about 1% only, because households face higher risk that induces them to save more. The worse outcomes at young ages are crucial given the ex ante perspective on welfare and the higher discount rate.

The remaining rows in Table 7 show the welfare costs of cyclical idiosyncratic risk in scenarios NORM and LKSW under the KV-calibration. Qualitatively, the results mimic the baseline calibration, and differences in magnitude are driven by the same mechanisms as described for the overall welfare costs.

We next analyze the aggregate measure of (self-)insurance against idiosyncratic risk. We again focus on the results for $\theta = 4$, as there is no relevant variation across different values of risk attitudes. The pass-through coefficient, as shown in Table 8, is smaller in scenario LKSW relative to scenario NORM for both transitory and persistent shocks. A smaller fraction of the lower pass-through is accounted for by consumption increases in response to positive shocks, which is again in line with increased precautionary savings.

TABLE 8
KV-CALIBRATION—AGGREGATE PASS-THROUGH AND DECOMPOSITION, $\theta = 4$

		<i>Components of $1 - \phi$:</i>			<i>Conditional $1 - \phi$:</i>	
<i>Transitory:</i>	$1 - \phi^\varepsilon$	$\frac{\pi^- A^\varepsilon}{1 - \phi^\varepsilon}$	$\frac{\pi^+ B^\varepsilon}{1 - \phi^\varepsilon}$	$\frac{C^\varepsilon}{1 - \phi^\varepsilon}$	$1 - \phi^\varepsilon \varepsilon^-$	$1 - \phi^\varepsilon \varepsilon^+$
NORM	0.101	0.562	0.445	-0.007	0.100	0.106
LKSW	0.081	0.558	0.383	0.058	0.058	0.146
<i>Persistent:</i>	$1 - \phi^\eta$	$\frac{\pi^- A^\eta}{1 - \phi^\eta}$	$\frac{\pi^+ B^\eta}{1 - \phi^\eta}$	$\frac{C^\eta}{1 - \phi^\eta}$	$1 - \phi^\eta \eta^-$	$1 - \phi^\eta \eta^+$
NORM	0.620	0.529	0.472	-0.001	0.657	0.666
LKSW	0.499	0.563	0.425	0.011	0.450	0.602

NOTE: Column 1 shows the aggregate consumption pass-through coefficient, and columns 2–4 its decomposition into components according to Equation (17), expressed as shares of total pass-through. π^- and π^+ are short for the probabilities of the shock being negative or positive. Columns 5 and 6 show conditional pass-through coefficients for negative and positive shocks.

TABLE 9
KV-CALIBRATION—PROPENSITY TO CONSUME ($\theta = 4$)

		<i>Negative Shocks</i>			<i>Positive Shocks</i>		
		Small	Medium	Large	Small	Medium	Large
<i>Transitory:</i>	NORM	0.0782	0.0796	0.0885	0.0753	0.0743	0.0727
	LKSW	0.0990	0.1002	0.1022	0.0970	0.0959	0.0936
<i>Persistent:</i>	NORM	0.5722	0.5773	0.5809	0.5607	0.5540	0.5298
	LKSW	0.5042	0.5094	0.5204	0.4974	0.4913	0.4735

NOTE: Propensity to consume out of a given shock. The numbers are evaluated at the mean asset profile over age from scenario LKSW in the KV calibration.

Table 9 reports the propensity to consume out of transitory and persistent shocks of different sign and magnitude (11, 22, and 44 log points). We evaluate the consumption responses at the age-specific mean asset holdings under the KV calibration in scenario LKSW. We find that the consumption response to transitory shocks is stronger in scenario LKSW than in scenario NORM, whereas the aggregate consumption response to persistent shocks is weaker. These differences are more pronounced than in the baseline calibration.

In Appendix D.3, we further decompose differences between the baseline and the KV-calibrations. We also explore robustness with respect to the preference specification by considering CRRA utility—stronger risk attitudes θ now simultaneously imply a lower IES γ . Our calibration then determines a lower discount rate because that second effect turns out to be the dominant force for asset accumulation. Consequently, the future is valued more in both scenarios NORM and LKSW, and therefore, for $\theta > 1$, the welfare effects of higher-order risk are lower. This finding underscores the importance of disentangling risk attitudes from intertemporal preferences. Also, we consider a version of the baseline where we recalibrate the discount rate for scenario NORM and find that our results are very little affected.

5. CONCLUSION

We estimate an income process that extends the canonical one by *higher-order risk* of transitory and persistent shocks. Our estimates on PSID household income imply that persistent shocks exhibit countercyclical variance and a procyclical third central moment. All shocks exhibit excess kurtosis.

Within an otherwise standard partial equilibrium life-cycle model with incomplete markets, first, higher-order risk has important welfare consequences relative to a world with lognormal shocks. Second, the presence of higher-order risk matters for the welfare costs of business

cycles. Third, higher-order risk affects the measured degree of consumption self-insurance and the (marginal) propensity to consume. A natural follow-up analysis pertains to the empirical exploration of consumption moments over the life cycle and how a structural model such as ours matches those.

ACKNOWLEDGMENTS

Open access funding enabled and organized by Projekt DEAL.

SUPPORTING INFORMATION

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

- Table B.1: Estimation Results for Household Net Income-Standardized Moments
- Figure B.1: Fit of Estimated Process for Household Net Income: Standardized Moments
- Table B.2: Moments of Income Changes
- Table B.3: Estimation Results for Household Pregovernment Income
- Table B.4: Estimation Results for Household Pregovernment Income-Standardized Moments
- Table B.5: Estimation Results for Robustness Specifications
- Table C.1: Moments in Three Distribution Scenarios
- Table C.2: Fitted Parameters of FGLD
- Table C.3: Moments of the Earnings Distribution in Logs and Levels
- Figure C.1: Moments of Life-Cycle Earnings by Age: Logs
- Figure C.2: Moments of Life-Cycle Earnings by Age: Levels
- Table C.4: Pension Indexation Factor ρ
- Figure C.3: Asset to Income Ratio over the Life Cycle
- Figure D.1: Discretized Log Distribution Functions: Persistent Shock
- Table D.1: Welfare Effects of Cyclical Idiosyncratic Risk: FGLD(NORM) versus Normal Distribution
- Table D.2: The Welfare Effects of Cyclical Idiosyncratic Risk for Distribution Scenario LK
- Table D.3: Aggregate Pass-Through and Its Decomposition for Scenario LK, $\theta = 4$
- Table D.4: Aggregate Pass-Through and Its Decomposition, $\theta = 2$
- Table D.5: Welfare Analysis: Sensitivity Analyses
- Figure D.2: Age Profile of Consumption: Baseline vs. KV-Calibration ($\theta = 4$)

REFERENCES

- ANGELOPOULOS, K., S. LAZARAKIS, and J. MALLEY, "Cyclical Labour Income Risk in Great Britain," *Journal of Applied Econometrics* 37 (2022), 116–30.
- ARELLANO, M., R. BLUNDELL, and S. BONHOMME, "Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework," *Econometrica* 85 (2017), 693–734.
- BAYER, C., and F. JUESSEN, "The Life-Cycle and the Business-Cycle of Wage Risk. Cross-Country Comparisons," *Economics Letters* 117 (2012), 831–33.
- BEAUDRY, P., and C. PAGES, "The Cost of Business Cycles and the Stabilization Value of Unemployment Insurance," *European Economic Review* 45 (2001), 1545–72.
- BLUNDELL, R., L. PISTAFERRI, and I. PRESTON, "Consumption Inequality and Partial Insurance," *American Economic Review* 98 (2008), 1887–921.
- BUSCH, C., D. DOMEIJ, F. GUVENEN, and R. MADERA, "Skewed Idiosyncratic Income Risk over the Business Cycle: Sources and Insurance," *American Economic Journal: Macroeconomics* 14 (2022), 207–42.
- , and A. LUDWIG, "Higher-Order Income Risk over the Business Cycle," CEPR Discussion Paper 14538, 2020.
- CARROLL, C., "The Methods of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economics Letters* 91 (2005), 312–20.

- CARROLL, C. D., M. HOLM, and M. S. KIMBALL, "Liquidity Constraints and Precautionary Savings," *Journal of Economic Theory* 195 (2021), 1–21.
- CATHERINE, S., "Countercyclical Labor Income Risk and Portfolio Choices over the Life Cycle," *Review of Financial Studies* 35 (2021), 4016–54.
- CIVALE, S., L. DIÉZ-CATALÁN, and F. FAZILET, "Discretizing a Process with Non-Zero Skewness and High Kurtosis," Technical Report, University of Minnesota, 2017.
- COMMAULT, J., "Does Consumption Respond to Transitory Shocks? Reconciling Natural Experiments and Semistructural Methods," *American Economic Journal: Macroeconomics* 14 (2022, April), 96–122.
- CONSTANTINIDES, G. M., and D. DUFFIE, "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy* 104 (1996), 219–40.
- COOPER, R., and G. ZHU, "Household Finance over the Life-Cycle: What Does Education Contribute?" *Review of Economic Dynamics* 20 (2016), 63–89.
- DALY, M., D. HRYSHKO, and I. MANOVSKII, "Improving the Measurement of Earnings Dynamics," *International Economic Review* 63 (2022), 95–124.
- DE NARDI, M., G. FELLA, M. KNOEF, G. PAZ-PARDO, and R. VAN OOIJEN, "Family and Government Insurance: Wage, Earnings, and Income Risks in the Netherlands and the U.S.," *Journal of Public Economics* 193 (2021).
- , ———, and G. PAZ-PARDO, "Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare," *Journal of the European Economic Association* 18 (2020), 890–926.
- DRUEDAHL, J., and A. MUNK-NIELSEN, "Higher-Order Income Dynamics with Linked Regression Trees," *Econometrics Journal* 23 (2020), S25–S58.
- EECKHOUDT, L., "Beyond Risk Aversion: Why, How and What's Next," *The Geneva Risk and Insurance Review* 37 (2012), 141–55.
- , and H. SCHLESINGER, "Changes in Risk and the Demand for Saving," *Journal of Monetary Economics* 55 (2008), 1329–36.
- EPSTEIN, L. G., and S. ZIN, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica* 57 (1989), 937–69.
- EPSTEIN, L. G. L., and ———, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis," *Journal of Political Economy* 99 (1991), 263–86.
- FREIMER, M., C. T. LIN, and G. S. MUDHOLKAR, "A Study of the Generalized Tukey Lambda Family," *Communications in Statistics - Theory and Methods* 17 (1988), 3547–67.
- GHOSH, A., and A. THELOUDIS, "Consumption Partial Insurance in the Presence of Tail Income Risk," Technical Report, McGill University, 2023.
- GOLOSOV, M., M. TROSHKIN, and A. TSYVINSKI, "Redistribution and Social Insurance," *American Economic Review* 106 (2016), 359–86.
- GOTTSCHALK, P., and R. MOFFITT, "The Growth of Earnings Instability in the U.S. Labor Market," *Brookings Papers on Economic Activity* 25 (1994), 217–72.
- GUVENEN, F., F. KARAHAN, S. OZKAN, and J. SONG, "What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Dynamics?" *Econometrica* 89 (2021), 2303–39.
- , S. OZKAN, and J. SONG, "The Nature of Countercyclical Income Risk," *Journal of Political Economy* 122 (2014), 1–59.
- HARENBERG, D., and A. LUDWIG, "Idiosyncratic Risk, Aggregate Risk, and the Welfare Effects of Social Security," *International Economic Review* 60 (2019), 661–92.
- HRYSHKO, D., and I. MANOVSKII, "On the Heterogeneity in Family Earnings and Income Dynamics in the PSID," *American Economic Association Papers and Proceedings* 108 (2018), 292–96.
- HUGGETT, M., and G. KAPLAN, "How Large is the Stock Component of Human Capital?," *Review of Economic Dynamics* 22 (2016), 21–51.
- IMROHOROGLU, A., "Cost of Business Cycles with Indivisibilities and Liquidity Constraints," *Journal of Political Economy* 97 (1989), 1364–83.
- KAPLAN, G., and G. L. VIOLANTE, "How Much Consumption Insurance Beyond Self-Insurance?," *American Economic Journal: Macroeconomics* 2 (2010), 53–87.
- KARAHAN, F., and S. OZKAN, "On the Persistence of Income Shocks over the Life-Cycle: Evidence, Theory, and Implications," *Review of Economic Dynamics* 16 (2013), 452–76.
- KIMBALL, M., "Precautionary Savings in the Small and in the Large," *Econometrica* 58 (1990), 53–73.
- KREBS, T., "Growth and Welfare Effects of Business Cycles in Economies with Idiosyncratic Human Capital Risk," *Review of Economic Dynamics* 6 (2003), 846–68.
- , "Job Displacement Risk and the Cost of Business Cycles," *The American Economic Review* 97 (2007), 664–86.
- KRUEGER, D., and A. LUDWIG, "On the Optimal Provision of Social Insurance: Progressive Taxation Versus Education Subsidies in General Equilibrium," *Journal of Monetary Economics* 77 (2016), 72–98.

- KRUSELL, P., T. MUKOYAMA, A. SAHIN, and A. A. SMITH, Jr., "Revisiting the Welfare Effects of Eliminating Business Cycles," *Review of Economic Dynamics* 12 (2009), 393–402.
- , and A. A. SMITH, Jr., "On the Welfare Effects of Eliminating Business Cycles," *Review of Economic Dynamics* 2 (1999), 245–72.
- LAJERI-CHAHERLI, F., "Proper Prudence, Standard Prudence and Precautionary Vulnerability," *Economics Letters* 82 (2004), 29–34.
- LAKHANY, A., and H. MAUSSER, "Estimating the Parameters of the Generalized Lambda Distribution," *ALGO Research Quarterly* 3 (2000), 47–58.
- LUCAS, R. E., *Models of Business Cycles* (New York: Basil Blackwell, 1987).
- , "Macroeconomic Priorities," *American Economic Review* 93 (2003), 1–14.
- MANKIW, N. G., "The Equity Premium and the Concentration of Aggregate Shocks," *Journal of Financial Economics* 17 (1986), 211–19.
- MCKAY, A., "Time-Varying Idiosyncratic Risk and Aggregate Consumption Dynamics," *Journal of Monetary Economics* 88 (2017), 1–14.
- MOFFIT, R., and P. GOTTSCHALK, "Trends in the Transitory Variance of Male Earnings in the U.S., 1970–2004," Working Paper 16833, NBER, 2011.
- Panel Study of Income Dynamics (PSID), Public Use Dataset, Produced and Distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI, 2015.
- SIEGEL, J. J., *Stocks for the Long Run: The Definitive Guide to Financial Market Returns and Long-Term Investment Strategies* (New York: McGraw-Hill, 2002).
- STORELLETTEN, K., C. I. TELMER, and A. YARON, "Asset Pricing with Ideoyncratic Risk and Overlapping Generations," *Review of Economic Dynamics* 10 (2007), 519–48.
- , ———, and ———, "The Welfare Cost of Business Cycles Revisited: Finite Lives and Cyclical Variation in Idiosyncratic Risk," *European Economic Review* 45 (2001), 1311–39.
- , ———, and ———, "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *The Journal of Political Economy* 112 (2004), 695–717.
- SU, S., "Numerical Maximum Log Likelihood Estimation for Generalized Lambda Distributions," *Computational Statistics and Data Analysis* 51 (2007), 3983–98.
- TALLARINI, T. D. J., "Risk-Sensitive Real Business Cycles," *Journal of Monetary Economics* 45 (2000), 507–32.
- WEIL, P., "The Equity Premium Puzzle and the Risk-Free Rate Puzzle," *Journal of Monetary Economics* 24 (1989), 401–21.