

Consumption Dynamics and Welfare under Non-Gaussian Earnings Risk

Fatih Guvenen, Serdar Ozkan, Rocio Madera

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Consumption Dynamics and Welfare under Non-Gaussian Earnings Risk

Abstract

Recent empirical studies document that the distribution of earnings changes displays substantial deviations from lognormality: in particular, earnings changes are negatively skewed with extremely high kurtosis (long and thick tails), and these non-Gaussian features vary substantially both over the life cycle and with the earnings level of individuals. Furthermore, earnings changes display nonlinear (asymmetric) mean reversion. In this paper, we embed a very rich “benchmark earnings process” that captures these non-Gaussian and nonlinear features into a lifecycle consumption-saving model and study its implications for consumption dynamics, consumption insurance, and welfare. We show four main results. First, the benchmark process essentially matches the empirical lifetime earnings inequality—a first-order proxy for consumption inequality—whereas the canonical Gaussian (persistent-plus-transitory) process understates it by a factor of five to ten. Second, the welfare cost of idiosyncratic risk implied by the benchmark process is between two-to-four times higher than the canonical Gaussian one. Third, the standard method in the literature for measuring the pass-through of income shocks to consumption—can significantly overstate the degree of consumption smoothing possible under non-Gaussian shocks. Fourth, the marginal propensity to consume out of transitory income (e.g., from a stimulus check) is higher under non-Gaussian earnings risk.

JEL-Codes: E240, J240, J310.

Keywords: idiosyncratic earnings risk, higher-order earnings risk, non-Gaussian shocks, incomplete markets models, consumption insurance.

Fatih Guvenen
University of Minnesota &
FRB of Minneapolis / MN / USA
guvenen@umn.edu
fatihguvenen.com

Serdar Ozkan
FRB of St. Louis / MO / USA &
University of Toronto / Canada
serdar.ozkan@gmail.com
www.serdarozkan.me

Rocio Madera
Southern Methodist University, University Park / TX / USA
romadera@smu.edu
s2.smu.edu/romadera

May 20, 2024

The order of coauthors has been assigned randomly using AEA’s Author Randomization Tool. We thank Fatih Karahan for contributions to some of the early results in this paper. For helpful comments, we thank Dirk Krueger, Gonzalo Paz-Pardo, Richard Rogerson, Ayşegül Şahin, Rob Shimer, and conference participants at the Gerzensee Conference on Heterogeneity and Macroeconomics of Labor Markets. Ozkan acknowledges financial support from the Social Sciences and Humanities Research Council of Canada.

1 Introduction

This paper studies the implications of non-Gaussian earnings risk for consumption dynamics, consumption insurance, and welfare. A growing number of studies have shown that the most salient features of earnings dynamics cannot be captured by a linear (ARMA(p,q)) process with Gaussian innovations. In particular, these studies have documented several important deviations from log-normality and significant heterogeneity in earnings dynamics across workers. Among these, three features are especially worth highlighting.¹

First, the distribution of earnings shocks is not symmetric but instead displays strong negative skewness, and it is not bell-shaped but instead displays extremely high kurtosis. The negative skewness implies that workers are more likely to experience very large negative changes in their earnings (“disaster shocks”) compared to very large positive shocks (big “upside surprises”). The high excess kurtosis implies that in a given year, most individuals experience very small changes in earnings relative to overall dispersion, while a few experience extremely large shocks. We refer to these non-Gaussian properties as “higher-order risk.” Second, and just as important, there is substantial heterogeneity in these non-Gaussian properties of earnings shocks both over the life cycle and across the earnings distribution. In particular, older and/or higher-income workers experience smaller but more leptokurtic and more negatively skewed earnings shocks.² Third, earnings shocks display important nonlinearities in the form of “asymmetric” mean reversion: for high-income workers, positive earnings shocks are fairly transitory, whereas negative shocks are very long-lasting. The opposite is true for low-income individuals: positive shocks are much more persistent for them than negative shocks are. Furthermore, persistence also depends on the size of shocks.

Given the central importance of earnings risk in shaping individuals’ economic behavior when markets are incomplete, we revisit some key questions about individual consumption-savings behavior in light of these new findings. To this end, we solve and simulate a rich life-cycle consumption-savings model that allows for heterogeneous, nonlinear, non-Gaussian earnings risk and study its implications for consumption dynamics, partial consumption insurance, and the welfare costs of idiosyncratic earnings risk.

¹See, e.g., [Arellano, Blundell and Bonhomme \(2017a\)](#) and [Guvenen, Karahan, Ozkan and Song \(2021\)](#) for the US data, and the November 2022 special issue of *Quantitative Economics* on Global Income Dynamics for evidence on 12 additional countries. [Guvenen, Pistaferri and Violante \(2022b\)](#) summarizes the evidence documented in these articles and show that the significant deviations from lognormality documented for the United States are also observed in these 12 countries.

²We use the terms “earnings” and (labor) “income” interchangeably throughout the paper.

The income process we use is the benchmark model in [Guvenen, Karahan, Ozkan and Song \(2021\)](#) (hereafter, GKOS), which is a very rich stochastic process with 21 parameters and estimated by targeting more than 2000 moments that capture a wide range of nonlinear and non-Gaussian features of earnings dynamics. The model fits the features described above (as well as many others not mentioned above) quite well. This stochastic process features (i) an AR(1) process with normal mixture innovations; (ii) a (potentially full-year) non-employment shock with scarring effects whose incidence varies by age and income level; (iii) a purely transitory shock, and (iv) a heterogeneous income profiles (hereafter, HIP) component. Although the process has many parameters, all dynamics are captured through only one state variable, the same as in the canonical persistent-plus-transitory earnings dynamics model that has been the workhorse in the incomplete markets literature. The nonlinear and non-Gaussian features make discretization methods for the income process inaccurate for some key moments, so we solve the lifecycle model with the full process estimated by GKOS, with shocks drawn from continuous distributions.

We embed this earnings process into a lifecycle consumption-savings model with stochastic lifetimes, retirement, borrowing constraints, and intergenerational persistence of labor earnings, among others. We model several mechanisms that can provide partial consumption insurance and smoothing in the form of accidental and warm-glow bequests that are inherited by newborn offspring, progressive labor income taxation, a redistributive retirement pension that mimics the US Social Security system, and a minimum consumption floor. This rich set of smoothing opportunities reduces the volatility of individual consumption substantially relative to earnings and is therefore important for the soundness of our quantitative analysis. We then solve the same lifecycle model with the canonical Gaussian (persistent-plus-transitory) earnings process, which is widely used in the literature. Comparing the implications of the two models, we reach four main conclusions.

First, as a prelude to our analysis of consumption dynamics, we compare the distributions of *lifetime* earnings generated by the *benchmark* (non-Gaussian) earnings process and the *canonical* Gaussian process (Table IV). This comparison provides a useful starting point because with full insurance (complete markets) and identical preferences, the distribution of consumption would mirror the distribution of lifetime earnings ([Friedman, 1957](#)). While this is a simplified benchmark, it provides a first look at the differences across individuals in their lifetime resources. We find that the canonical earnings process understates the overall lifetime earnings inequality measured by the 90th-to-10th percentile ratio (P90-P10) by a factor of 5 (a ratio of almost 15 in the data versus 2.8), whereas the benchmark process essentially matches it (a ratio of 16). The gap is even larger at the

top, with the canonical Gaussian process understating the P99-P10 ratio by a factor of 10 (a ratio of 44 in the data versus 4.3), whereas the benchmark process generates it almost exactly (41). We should note that GKOS do not discuss the benchmark process’s ability to match the lifetime earnings inequality; thus, this comparison is new to this paper.

The second and most novel contribution of our paper is to quantify the welfare costs of an idiosyncratic earnings risk implied by each earnings process. With a relative risk aversion of 2, we find that the average individual is willing to pay 8.5% of her consumption at every date and state to eliminate idiosyncratic earnings fluctuations generated by the canonical Gaussian process. The analogous welfare cost is almost *four times higher*, at 33.2%, for earnings fluctuations generated by the non-Gaussian benchmark process from GKOS. In this comparison, we chose the parameters for the canonical Gaussian process to reflect typical estimates in previous studies so that the resulting comparison is directly relevant for a large number of studies that are calibrated using those estimates. However, those parameters understate the variance of earnings fluctuations in the data targeted by GKOS. So, as an alternative, we recalculate the welfare cost of a Gaussian process that is estimated to match the same moments as in GKOS, which brings up the welfare cost to 16.8%. Although much higher, this value is still half the welfare cost of the non-Gaussian benchmark process. Furthermore, the welfare cost is significantly more heterogeneous across the income distribution under the benchmark process, for example, ranging from 7.3% to 18.3% for different groups at age 45 compared with a range of 2.5% to 3.6% for the canonical Gaussian process.

Among the different components of the benchmark process, the persistent AR(1) component with normal mixture shocks is responsible for almost 2/3 of the welfare cost (down to 11.4% without AR(1) from 33.2%). The non-employment shock with scarring effects is also very costly, accounting for 1/3 of the welfare cost (22.3%). On the other side, some of the smoothing opportunities built into the model turn out to be very effective: the welfare cost would have been 46.9% without the progressive tax system and 42.2% if the borrowing limit was reduced by half. Similarly, the welfare cost drops to 20.2% if the minimum income floor is raised to \$10,000 from about \$3,000 in the baseline calibration. Overall, these results show that idiosyncratic earnings shocks are very costly for individual welfare, and accurately modeling their nonlinear and non-Gaussian features is essential for properly measuring their welfare costs.

Third, we link the estimated welfare costs of earnings fluctuations to measures of the insurability of idiosyncratic shocks through the “partial insurance” coefficient (Blundell *et al.* (2008), Kaplan and Violante (2010), and others). Although related, the two con-

cepts are somewhat different: the partial insurance coefficient measures the proportion of a typical shock that is transmitted to consumption without reference to the size of the shock, whereas the welfare cost measure incorporates both the size and nature of the underlying fluctuations in addition to their insurability. When the partial insurance coefficients are estimated using the standard moment conditions in the literature, the insurability of persistent shocks appear to be almost twice as high under the benchmark process relative to a Gaussian process (62% vs 32% insured). We show that this moment condition does not isolate persistent shocks under the benchmark process, thereby overstating the extent of partial insurance. We also show that within-cohort consumption inequality over the life cycle increases significantly more—consistent with the high welfare costs noted above—and reveals less insurability of shocks under the benchmark process. Furthermore, the benchmark process implies a left-skewed and leptokurtic distribution of consumption growth, which has some support from the data (e.g., [Constantinides and Ghosh \(2017\)](#), [Toda and Walsh \(2015\)](#), and [Rodriguez Mora *et al.* \(2022\)](#)).

Fourth and finally, we compare the implications for the marginal propensity to consume (MPC) out of various types of earnings shocks. Specifically, the MPC out of a \$500 stimulus check is substantially higher under the non-Gaussian benchmark process: it is about 10% at the 20th percentile of the cash-on-hand distribution versus 4% under the Gaussian process; and it is 7.5% at the median cash-on-hand versus again 4% for the Gaussian process. Furthermore, even though the MPC declines relatively slowly with cash-on-hand for the benchmark process, it is almost completely flat under the Gaussian process. We also investigate the MPCs out of persistent and transitory earnings shocks (which scale with the level of earnings unlike the stimulus check).

A potential limitation of our analysis is that we focus on individuals as our unit of analysis—so as to be consistent with the earnings process we use, which is estimated on individual earnings—which leaves out potential consumption insurance within the household. It is not clear how important this choice is for our results because the empirical evidence on the effectiveness and extent of spousal insurance is mixed, with some studies finding little or no effect against income shocks (e.g., [Hayashi *et al.* \(1996\)](#) and [Busch *et al.* \(2022\)](#)), while others find non-negligible effects ([Blundell *et al.* \(2016\)](#); [Halvorsen *et al.* \(2023\)](#); [Pruitt and Turner \(2020\)](#)). As a simple check, we report some key moments of earnings dynamics, which are quite similar for individuals and households using data from the Panel Study of Income Dynamics (PSID). [De Nardi, Fella, Knoef, Paz-Pardo and Van Ooijen \(2021\)](#) also reach the same conclusion using administrative earnings data from Netherlands.

Literature Review

Although a vast literature that starts with Lucas [Lucas \(1987\)](#) has attempted to measure the welfare costs of aggregate fluctuations, there is little previous work on the welfare costs of idiosyncratic earnings shocks. A recent exception is [Constantinides \(2021\)](#) who separates idiosyncratic and aggregate shocks and finds large welfare costs of the former. The paper takes a clever approach that utilizes an exchange economy model with a no-trade equilibrium (extending [Constantinides and Duffie \(1996\)](#)) whose parameters are estimated by fitting asset prices and cross-sectional moments of household consumption growth. One challenge is that the consumption data used from the Consumer Expenditure Survey is observed for only five quarters, whereas the model only has permanent shocks. This requires assuming that the quarterly consumption changes in the data are completely permanent, when in fact these short-term variations are likely to contain a large transitory component due to measurement error, which can introduce an upward bias the estimated welfare costs. There are also some studies that brought idiosyncratic shocks into the welfare costs of business cycles by allowing aggregate and idiosyncratic shocks to be correlated (see, e.g., [Storesletten *et al.* \(2001\)](#); [Krebs \(2003\)](#); [Krusell *et al.* \(2009\)](#), among others). Our paper differs from these studies by modeling a very rich earnings process with non-Gaussian features and allowing several consumption-smoothing mechanisms and measuring the costs of idiosyncratic risk in this setting.

Our paper is also related to two recent papers who compare the degree of insurance and welfare costs of idiosyncratic income risk under a permanent-plus-transitory standard Gaussian income process and a non-linear non-Gaussian income process: [De Nardi, Fella and Paz-Pardo \(2020\)](#) find a larger degree of partial insurance under the non-Gaussian income process (see our results in [5.1](#)), something we also find in this paper. In contrast, they find the welfare costs of idiosyncratic income risk to be lower for non-Gaussian risk. Our analysis differs from theirs in four main ways: (1) the persistence of earnings shocks in the benchmark process is not solely captured by the AR(1) component but also through nonemployment (disaster) shocks, which accounts for a third of the welfare costs; (2) they estimate an earnings process using survey data (the PSID), in which measurement error possibly attenuates the higher-order moments;³ (3) we compare to a Gaussian process, with persistent rather than permanent shocks; and (4) we introduce an income floor and show it is crucial to capture the degree of insurability against negative shocks. Closer to

³For example, the Crow-Siddiqui kurtosis measure in their study is much lower than the GKOS reports from the SSA data. As we discuss below, skewness is also much less pronounced in the PSID due to possible measurement error.

our exercise, [Busch and Ludwig \(2024\)](#) also find higher partial insurance and welfare costs under non-Gaussian risk, provided a minimum level of risk aversion.⁴ Our analysis differs from theirs along channels (1), (2), and (4) described earlier in this paragraph.

Our paper hence also relates to the literature that studies the role of government insurance in smoothing higher-order income risk. [Busch *et al.* \(2022\)](#) studies the effectiveness of various government social insurance policies in France, Germany, Sweden, and the United States for mitigating the procyclical fluctuations in the skewness of earnings risk. In a different context, [De Nardi *et al.* \(2020\)](#) studies the optimal policy mix in a world with higher-order risk and find income-floor policies to be the best at insuring against this kind of risk, which is precisely the kind of government policy we include in our exercise and find to be quantitatively very important.

Finally, our paper also contributes to the literatures on the MPC of earnings and wealth ([Johnson *et al.*, 2006](#); [Parker *et al.*, 2013](#); [Kaplan and Violante, 2014](#)) and the insurability of income shocks ([Blundell *et al.*, 2008](#); [Kaplan and Violante, 2010](#); [Güvenen and Smith, 2014](#); [Blundell *et al.*, 2012](#)) by studying the ramifications of non-Gaussian earnings risk, which has not been considered by these studies. A few notable exceptions are as follows: [Wang, Wang and Yang \(2016\)](#) build a tractable buffer-stock model to study the extent of consumption smoothing with large discrete income jumps and find larger marginal propensity to consume (MPC) than these models usually imply under Gaussian risk—similar to our finding. Recently, [Commault \(2022\)](#) also incorporates the GKOS benchmark process into a Bewley model but calibrates it to match the net liquid wealth (which implies less wealth in the economy relative to our calibration) and finds larger MPCs against negative transitory income shocks. Our analysis complements these papers by embedding a very general non-Gaussian, non-linear stochastic process into a rich life-cycle model and analyzing a range of properties of consumption dynamics. [Ghosh and Theloudis \(2023\)](#) estimate the extent of partial insurance against non-Gaussian shocks using a higher-order log-approximation of optimal consumption decision. They find that the pass-through is larger for negative permanent shocks. [Busch and Ludwig \(2024\)](#) find, exactly like us, that the MPC out of permanent (transitory) shocks is lower (higher) when income risk is non-Gaussian due to a stronger precautionary savings motive. In our paper, we further show this is driven by the lowest decile of the wealth distribution.

⁴A relative risk aversion above 1 is enough to yield higher welfare costs under higher-order risk.

2 Model

The main focus of our analysis is on the implications of non-Gaussian earnings risk for consumption behavior and the welfare costs that such risk entails. However, as economists, much of our intuition about how individuals evaluate risk comes from the famous Arrow-Pratt thought experiment, which implicitly relies on a *thin-tailed* and *symmetric* distribution (such as a Gaussian) to model uncertainty. So, before we delve into our full-blown quantitative model, we begin by showing in a stylized example how allowing for skewness and thick tails (kurtosis) affects the risk premium demanded by individuals.

2.1 Risk Premium with Higher-order Risk

Let us reconsider the Arrow-Pratt thought experiment that confronts a decision maker with a choice between consuming the outcome of a risky bet, $c \times (1 + \tilde{\delta})$, versus the expected payoff of the bet minus a risk premium, $c \times (1 - \pi)$. The question is: how much is the premium, π , that makes the decision maker indifferent between the two options:

$$u(c \times (1 - \pi)) = \mathbb{E} \left[u(c \times (1 + \tilde{\delta})) \right]? \quad (1)$$

The usual derivation proceeds by taking a first-order Taylor approximation to the left-hand side and a second-order approximation to the right-hand side. Why do we stop at second order? Because of two implicit assumptions: (i) that the skewness of the distribution of $\tilde{\delta}$ is not too far from zero, so the third-order term may not be too important, and (ii) the distribution is sufficiently thin tailed (no excess kurtosis) that realizations of $\tilde{\delta}$ far from zero have very low probability, so the fourth-order term may not be too important. In light of the recent empirical evidence discussed above that income shocks have nonzero skewness and excess kurtosis, let us see what happens when we take a fourth-order approximation to the right hand side:

$$\begin{aligned} u(c) - u'(c)c\pi &\approx \mathbb{E} \left(u(c) + u'(c)c\tilde{\delta} + \frac{u''(c)}{2}c^2\tilde{\delta}^2 + \frac{u'''(c)}{6}c^3\tilde{\delta}^3 + \frac{u''''(c)}{24}c^4\tilde{\delta}^4 \right) \\ \pi &\approx \mathbb{E} \left(\frac{1}{2} \frac{u''(c)}{u'(c)} c \tilde{\delta}^2 + \frac{u'''(c)}{6u'(c)} c^2 \tilde{\delta}^3 + \frac{u''''(c)}{24u'(c)} c^3 \tilde{\delta}^4 \right). \end{aligned}$$

Furthermore, let $u(c) = (c^{1-\gamma}/(1-\gamma))$ take the CRRA form and rearrange the terms

TABLE I – Exact Solution to the Risk Premium in Equation (1)

Gambles:	Risk Premium (π)	
	$\tilde{\delta}^{\text{Gaussian}}$	$\tilde{\delta}^{\text{Non-Gaussian}}$
Mean	0.0	0.0
Standard Deviation	0.10	0.10
Skewness	0.0	-2.0
Excess Kurtosis	0.0	27.0
Risk premium	4.9%	22.2%

Note: The first column shows the risk premium under the assumption that the bet has a Gaussian distribution. The second column shows it for a Non-Gaussian distribution whose parameters are taken from [Güvenen *et al.* \(2021\)](#) for a 45-year-old male at the 90th percentile of the income distribution. The standard deviation of 0.1 is about 1/5 of its empirical counterpart (0.51), which implicitly assumes that 80% of the income shock has been insured and is not passed through to consumption.

to get:

$$\pi^* \approx \underbrace{\gamma}_{\text{variance aversion}} \times \frac{\sigma_\delta^2}{2} - \underbrace{\frac{(\gamma+1)\gamma}{6}}_{\text{neg. skew aversion}} \times \sigma_\delta^3 \times s_\delta + \underbrace{\frac{(\gamma+2)(\gamma+1)\gamma}{24}}_{\text{kurtosis aversion}} \times \sigma_\delta^4 \times k_\delta \quad (2)$$

where s_δ and k_δ denote, respectively, the skewness and the kurtosis coefficients (3rd and 4th standardized moments). The first term on the right-hand side is traditionally called the risk premium, and γ is called the relative risk aversion parameter. However, this expanded expression shows that γ only captures the *aversion to variance*, whereas the risk premium also depends on the *aversion to negative skewness* (second term) and *aversion to kurtosis* (third term). These latter terms have two counteracting effects determining their importance. On the one hand, they depend on higher powers of γ , which is typically greater than one; on the other hand, they feature higher powers of σ_δ , which is smaller than one. Therefore, the impact of these terms on the risk premium depends on the empirical values of γ , σ_δ , s_δ , and k_δ . Notice that the sign of the skew aversion term is negative, indicating that a negatively skewed distribution requires a higher premium, and vice versa for a positively skewed distribution.

Table I shows an illustrative example for a relatively high $\gamma = 10$ and a bet with a standard deviation of $\sigma_\delta = 0.10$ under two different assumptions. In one case, the bet is assumed to be Gaussian, so it has zero skewness and no excess kurtosis. In the second case, the bet is non-Gaussian with its skewness (-2) and excess kurtosis (27), set to equal that

for a 45-year-old male US worker in the 90th percentile of the US earnings distribution (from GKOS). We should note that in the US data, the standard deviation of one-year earnings growth for the just-mentioned worker is not 0.10—in fact, it is above 0.50. So, to be conservative, the assumption in this example is that 80% of the idiosyncratic income risk has been insured, and the individual’s consumption is only subject to the remaining 20% of the risk in her earnings.

As seen in Table I, the risk premium under the non-Gaussian risk is 22.2%, more than four times the premium under Gaussian risk (4.9%).⁵ The amplification is especially strong when risk aversion is higher, as could be predicted from the appearance of higher-order polynomials in γ in the formula for (2). In the rest of the paper, we will show that the much higher risk premium demanded by individuals to bear non-Gaussian risk will carry over to a properly calibrated dynamic life cycle model with borrowing and saving, as well as government insurance and transfers.

2.2 A Life-Cycle Consumption-Savings Model

We consider a life cycle consumption-savings model with income uncertainty, retirement, stochastic lifetimes, and imperfect altruism. Individuals face an age-dependent probability of death every period, and the conditional survival probability from age t to $t + 1$ is denoted with δ_t . Individuals can (potentially) work for the first T_W years of their life, retire at age $T_W + 1$, and die with certainty by age $t > T$. An individual who dies is replaced with an offspring, who inherits her parent’s positive assets after paying a flat estate tax. Parents derive utility from leaving a bequest according to a warm-glow bequest function:

$$\phi(\mathbf{b}) = \phi_1 \frac{(\mathbf{b} + \phi_2)^{1-\gamma}}{1-\gamma}$$

as in De Nardi and Yang (2016). This is a flexible functional form that allows us to model bequests either as a necessity or luxury good depending on parameterization.

Individuals have CRRA preferences over consumption and therefore supply labor inelastically. Individuals can borrow or save using a risk-free asset with gross return R , where borrowing is subject to an age-dependent, worker-type-specific limit, denoted by \underline{A}_t^k , described further below. The type of a worker is given by his fixed type, Υ^k , in the earnings

⁵The values reported here are computed exactly from equation ((1)) using numerical integration, so it does not rely on the approximation in ((2)).

equation described later below. The dynamic programming problem of an individual is

$$V_t^i(\mathbf{a}_t^i, z_t^i; \gamma^k) = \frac{(\mathbf{c}_t^i)^{1-\gamma}}{1-\gamma} + \beta \left[(1 - \delta_t) \mathbb{E}_t (V_{t+1}^i(\mathbf{a}_{t+1}^i, z_{t+1}^i; \gamma^k)) + \delta_t \frac{\phi_1 (\mathbf{b}_{t+1}^i + \phi_2)^{1-\gamma}}{(1-\gamma)} \right]$$

s.t.

$$\mathbf{c}_t^i + \mathbf{a}_{t+1}^i = \mathbf{a}_t^i \mathbf{R} + Y_t^{\text{disp},i}, \quad \forall t, \quad (3)$$

$$Y_t^{\text{disp},i} = \lambda \max \left\{ \underline{Y}, \tilde{Y}_t^i \right\}^{1-\tau}, \quad t = 1, \dots, T_W, \quad (4)$$

$$Y_t^{\text{disp},i} = \lambda \left(\tilde{Y}_R^k(z_{T_W}^i) \right)^{1-\tau}, \quad t = T_W + 1, \dots, T,$$

$$\mathbf{a}_{t+1}^i \geq \underline{\mathbf{A}}_t^k, \quad \forall t, \quad (5)$$

$$\text{and equations (7) to (13),} \quad (6)$$

where $\delta_{T+1} \equiv 1$, β is the time discount factor, and γ governs risk aversion. The budget constraint is given as in equation (3) where \mathbf{c}_t^i and \mathbf{a}_t^i denote consumption and asset holdings, respectively, and $Y_t^{\text{disp},i}$ is disposable income, which differs from gross income, \tilde{Y}_t^i , in two ways. First, the government provides social insurance by guaranteeing a minimum level of income, \underline{Y} , to all individuals—more on this in a moment. Second, the government imposes a progressive tax on after-transfer income. Following [Benabou \(2000\)](#) and [Heathcote, Storesletten and Violante \(2014\)](#), we take this tax function to have a power form, with exponent $1 - \tau$.

Furthermore, retirees receive pension income, $\tilde{Y}_R^k(z_{T_W}^i)$, specified to mimic the U.S. Social Security Administration’s OASDI system’s “*primary insurance amount (PIA)*,” which is the benefit a person would receive if she begins receiving retirement benefits at the normal retirement age. The retirement pension is a function of the average lifetime earnings below Social Security maximum taxable earnings. To avoid introducing another state variable (i.e., lifetime earnings), we approximate lifetime earnings by the persistent component of a worker’s earnings in the last year of the working life, $z_{T_W}^i$. In particular, for each worker type k , we regress the simulated average earnings on worker’s $z_{T_W}^i$ with an intercept, which we then use to approximate worker’s lifetime earnings, $\text{LE}^k(z_{T_W}^i)$. The following equation specifies the pension system as a function of $\text{LE}^k(z_{T_W}^i)$:

$$\tilde{Y}_R^k(z_{T_W}^i) = \begin{cases} 0.9\text{LE}^k(z_{T_W}^i) & \frac{\text{LE}^k(z_{T_W}^i)}{\text{AE}} < 0.23 \\ 0.2\text{LE}^k(z_{T_W}^i) + 0.32(\text{LE}^k(z_{T_W}^i) - 0.23\text{AE}) & 0.23 < \frac{\text{LE}^k(z_{T_W}^i)}{\text{AE}} < 1.38 \\ 0.57\text{AE} + 0.15(\text{LE}^k(z_{T_W}^i) - 1.38\text{AE}) & \frac{\text{LE}^k(z_{T_W}^i)}{\text{AE}} > 1.38 \end{cases}$$

for $t = T_W + 1, \dots, T$, where AE is the average earnings in the economy.

2.3 Specifications of Earnings Processes

The specification of the benchmark non-Gaussian and nonlinear earnings process follows [Guvnen, Karahan, Ozkan and Song \(2021\)](#) (hereafter, GKOS) and has the following components: (i) an AR(1) process (z_t^i) with innovations drawn from a mixture of normals; (ii) a nonemployment shock whose incidence probability ($p_v^i(t, z_t)$) can vary with age or z_t or both, and whose duration (v_t^i) is exponentially distributed; (iii) a heterogeneous income profiles component (HIP); and (iv) an i.i.d. normal-mixture transitory shock (ε_t^i):

$$\text{Level of earnings: } \tilde{Y}_t^i = (1 - v_t^i)e^{(g(t) + \alpha^i + \theta^i t + z_t^i + \varepsilon_t^i)} \quad (7)$$

$$\text{Persistent component: } z_t^i = \rho z_{t-1}^i + \eta_t^i, \quad (8)$$

$$\text{Innovations to AR(1): } \eta_t^i \sim \begin{cases} \mathcal{N}(\mu_{\eta,1}, \sigma_{\eta,1}) & \text{with prob. } p_z \\ \mathcal{N}(\mu_{\eta,2}, \sigma_{\eta,2}) & \text{with prob. } 1 - p_z \end{cases} \quad (9)$$

$$\text{Initial condition of } z_t^i: z_0^i \sim \mathcal{N}(0, \sigma_{z_0}) \quad (10)$$

$$\text{Transitory shock: } \varepsilon_t^i \sim \begin{cases} \mathcal{N}(\mu_{\varepsilon,1}, \sigma_{\varepsilon,1}) & \text{with prob. } p_\varepsilon \\ \mathcal{N}(\mu_{\varepsilon,2}, \sigma_{\varepsilon,2}) & \text{with prob. } 1 - p_\varepsilon \end{cases} \quad (11)$$

$$\text{Nonemployment duration: } v_t^i \sim \begin{cases} 0 & \text{with prob. } 1 - p_v(t, z_t^i) \\ \min\{1, F_{\text{exp}}(\varphi)\} & \text{with prob. } p_v(t, z_t^i) \end{cases} \quad (12)$$

$$\text{Prob of Nonemp. shock: } p_v^i(t, z_t) = \frac{e^{\xi_t^i}}{1 + e^{\xi_t^i}}, \text{ where } \xi_t^i \equiv \mathbf{a} + \mathbf{b}t + \mathbf{c}z_t^i + \mathbf{d}z_t^i t. \quad (13)$$

In equation (7), $g(t)$ is a quadratic polynomial—where $t = (\text{age} - 24)/10$ is a normalized age—that captures the lifecycle profile of earnings *common* to all individuals. The random vector (α^i, θ^i) determines ex ante heterogeneity in the level and growth rate of earnings and is drawn from a multivariate normal distribution with zero mean and covariance matrix Σ_α . As we describe in Section 3, we will also allow (α^i, θ^i) to be correlated

across generations to capture the positive intergenerational correlation in earnings observed in the US data.

The innovations, η_t^i , to the AR(1) component are drawn from a mixture of two normals. An individual draws a shock from $\mathcal{N}(\mu_{\eta,1}, \sigma_{\eta,1})$ with probability p_z and otherwise from $\mathcal{N}(\mu_{\eta,2}, \sigma_{\eta,2})$. Without loss of generality, we normalize η to have zero mean (i.e., $\mu_{\eta,1}p_z + \mu_{\eta,2}(1 - p_z) = 0$) and assume $\mu_{\eta,1} < 0$ (without loss of generality) for identification. Heterogeneity in the initial value of z_t is captured by $z_0^i \sim \mathcal{N}(0, \sigma_{z_0})$. Transitory shocks, ε_t^i , are also drawn from a mixture of two normals (eq. 11), with analogous identifying assumptions (zero mean and $\mu_{\varepsilon,1} < 0$). Solving a dynamic programming problem with normal mixture shocks requires modest adjustments to the computational methods commonly used with Gaussian shocks.⁶

The last component of the earnings process is a nonemployment shock (eq. 12) that is intended to primarily capture movements in the extensive margin. Specifically, a worker is hit with a nonemployment shock with probability p_v whose duration $v_t > 0$ is drawn from an exponential distribution, F_{exp} with mean $1/\varphi$, and is truncated at 1 (corresponding to full-year nonemployment with zero annual income). This shock differs from z_t and ε_t by scaling the *level* of annual income—not its logarithm—which allows the process to capture the sizable fraction of workers who transition into and out of long-term nonemployment every year.⁷ Furthermore, the nonemployment incidence p_v depends on age t and z_t through the logistic function shown in equation 13. The dependence of p_v on z_t —which GKOS refer to as “state dependence”—turns out to be especially important as it *induces* persistence in nonemployment from one year to the next (despite v_t itself being independent over time).

Although this process has many parameters, all dynamics are captured through a single state variable (z_t^i), which makes it relatively straightforward to embed it into a dynamic programming problem. We construct a two-dimensional discrete grid over continuous ex-ante type variables (α^i, θ^i), where each grid point corresponds to a worker type Υ^k . Therefore, some aspects of an individual’s problem depends on his (discrete) ex ante type

⁶The main change is that the expectation on the right hand side of the Bellman equation becomes the weighted sum of two integrals (e.g., with weights p_z and $1 - p_z$), each using the density of the respective Gaussian distribution that makes up the mixture. These integrals are most conveniently computed using continuous integration methods (e.g. based on Gaussian quadrature, etc.) rather than discretizing the shock spaces, which is difficult because the optimal location of grid points needs to be different for each Gaussian component.

⁷It takes a -350 log point shock to z_t or ε_t for a worker earning \$50,000 to drop below Y_{min} . Generating such large shocks with sufficiently high frequency to match worker exit rates makes it challenging to simultaneously match the high frequency of smaller shocks.

TABLE II – Parameters Calibrated outside of the Model

PARAMETER		VALUE
Curvature of utility function	γ	2
Retirement (model) age	T_W	36
Maximum (model) age	T	60
Borrowing limit tightness	ψ	0.58
Tax progressivity	τ	0.185
Tax level parameter	λ	See Table VI
Risk free rate	$R - 1$	3%
Intergen. corr. labor fixed effect	ρ_α	0.5

Notes: In addition to these parameters, survival probabilities, δ_t , are taken from [Bell and Miller \(2002\)](#) (omitted from the table). We recalibrate λ for each income process to generate an average tax rate of 20%. See Table VI.

k , whereas others—including the dynamics of income—are drawn from continuous distributions that are also fully individual specific. This problem can be solved using standard numerical techniques; see computational Appendix A.

3 Model Parameterization

3.1 Demographics, Preferences, and Smoothing Opportunities

Households enter the labor market at age 25, retire at 60 ($T_W = 36$), and die with certainty by 85 ($T = 60$). We set the coefficient of relative risk aversion, γ , to 2 but also consider a value of 5 in the robustness analysis. The net interest rate, $R - 1$ is set to 3%. For the bequest function, we follow [De Nardi and Yang \(2016\)](#) and calibrate the parameters, ϕ_1 and ϕ_2 , to match a bequest–wealth ratio of 0.88% ([Gale and Scholz \(1994\)](#)) and the 90th percentile of bequest distribution normalized by income (4.53) ([Hurd and Smith \(2002\)](#)).

For the borrowing limit, we follow [Guvenen and Smith \(2014\)](#) and assume that banks use a potentially higher interest rate to discount households’ future labor income during working years (to account for income uncertainty) in calculating the borrowing limit, but simply apply the risk-free rate for discounting retirement income, according to this formula:

$$\underline{A}_t^k \equiv \underline{Y} \left[\sum_{j=1}^{R-t} \left(\frac{\psi}{R}\right)^j + \psi^{R-t+1} \sum_{j=R-t+1}^{T-t} \left(\frac{1}{R}\right)^j \right], \quad (14)$$

where $\psi \in [0, 1]$ measures the tightness of the borrowing limit. When $\psi = 0$, no borrowing is allowed against future labor income; when $\psi = 1$, households can borrow up to the natural limit. An appealing feature of the formulation in (14) is that for different values of $\psi \in (0, 1)$, it can generate a pattern of a borrowing limit that increases, decreases, or is U-shaped over the life cycle.⁸

In our baseline calibration, we set $\psi = 0.58$ which is Guvenen and Smith (2014)’s benchmark estimate, but we also experiment with lower values of ψ as well as with a more standard ad hoc borrowing limit later. Given these parameter values, β is calibrated to generate a wealth/income ratio of 4.⁹ Although values lower than this can also be justified on empirical grounds (and have often been used in the literature), a lower ratio implies less wealth available for smoothing income shocks and hence larger welfare costs of idiosyncratic risk. Therefore, a value of 4 represents a more conservative choice for the purposes of this paper.

To sum up, among the model parameters, four of them— $\beta, \lambda, \phi_1,$ and ϕ_2 —are calibrated inside the model, so their values are different for each income process we consider. The values are reported below in Table VI.

The income floor, \underline{Y} , is a key parameter as it limits the severity of large negative shocks to income. Guner *et al.* (2022) provide a comprehensive evaluation of the US welfare system and provide estimates of the magnitude of welfare assistance programs for different demographic groups. Their estimates for total government non-medical transfers for a married household with no children is 6.55% of mean household income. The same figure is similar for a single male with no children, at 6.84%. Therefore, regardless of which interpretation one adopts for the appropriate unit of analysis for the lifecycle model that we analyze in this paper, we set the income floor, \underline{Y} , to 6.75% of the average earnings, which seems like a reasonable, middle-ground estimate. Given the importance of this parameter for our results, below we will also report results for an income floor of \$10,000, which corresponds to 22% of average earnings (which is \$45,000 in 2010 dollars in our sample).¹⁰

⁸Because our model features a nonzero borrowing limit and stochastic survival, households may die with outstanding debt. For simplicity, we assume that both borrowing and saving interest rates are equal to the risk-free rate. However, in a robustness check, we modify the borrowing rate to incorporate the probability of death in addition to the risk-free rate and find very similar results.

⁹This value is based on the Flow of Funds Z1 tables. The total wealth-to-income ratio is defined to be total asset holdings in the population relative to the sum of total before-tax labor income and capital income.

¹⁰In an earlier paper, Hubbard *et al.* (1995) estimate the total value of various government programs (food stamps, AFDC, housing subsidies, etc.) for a female-headed family with two children and no outside earnings and assets. They obtain a value of about \$7,000 in 1984 dollars. Using the OECD equivalence

TABLE III – Parameters of Stochastic Processes for Earnings

Stochastic Process	Canonical Gaussian	Gaussian+	Benchmark-R	Benchmark
	(1)	(2)	(3)	(4)
Parameters				
σ_α	0.12	0.42	0.472	0.300
$\sigma_\theta \times 10$	—	—	—	0.196
$\text{corr}_{\alpha\theta}$	—	—	—	0.768
ρ	0.980	0.980	0.991	0.959
p_z	—	—	17.6%	40.7%
$\mu_{\eta,1}$	—	—	-0.524	-0.085
$\sigma_{\eta,1}$	0.11	0.19	0.113	0.364
$\sigma_{\eta,2}$	—	—	0.046	0.069
σ_{z_0}	0.278	0.001	0.450	0.714
φ	—	—	0.016	0.0001
p_ε	—	—	4.4%	13.0%
$\mu_{\varepsilon,1}$	—	—	0.134	0.271
$\sigma_{\varepsilon,1}$	0.30	0.49	0.762	0.285
$\sigma_{\varepsilon,2}$	—	—	0.055	0.037

Notes: The table reports the parameter values of the earnings processes used to calibrate the lifecycle model. The parameter estimates for the Canonical Gaussian model (column 1) are taken from [Karahan and Ozkan \(2013\)](#); the estimates for the remaining three columns are taken from [Güvenen *et al.* \(2021\)](#).

Finally, the curvature parameter τ , which controls the progressivity of the tax system, is set to 0.185, following [Heathcote *et al.* \(2014\)](#), and λ is calibrated such that average after-tax, after-transfer income over the working life is 80% of average before-tax, before-transfer labor income. Last but not least, flat estate tax is calibrated as 40% on positive inheritances. Table II summarizes the parameter values set exogenously outside the model.

3.2 Idiosyncratic Earnings Processes

We consider four specifications for the earnings process, shown in Table III. The first one (column 1) is the most widely used process in the literature, which features an individual fixed effect, a persistent shock (AR(1)), and a transitory shock ($y_t^i = \alpha^i + z_t^i + \varepsilon_t^i$, using the notation above), with all Gaussian innovations. We will refer to this as the “canonical Gaussian” process. We take the parameter values for this process from [Karahan and Ozkan](#)

scale for one adult plus two children, this comes to \$6,400 per adult person in 2010 dollars.

(2013) (hereafter KO), reported in column (1) of Table III.¹¹ As we will see below, however, this process fails to match some key features of the data, such as the (very high) level of lifetime income inequality, which are important for our analysis. Therefore, we consider a second process that keeps the same specification as the canonical one but targets a broader set of data moments used to estimate the benchmark process, except for the distribution of *lifetime* employment rates. This process has been estimated by GKOS, from whom we borrow the parameter values from. We call this the “Gaussian+” process (column 2 of Table III).

GKOS estimate the canonical process by also matching the distribution of *lifetime* employment rates, but the estimates of key parameters for this process are extremely large (column (1) of GKOS’s Table IV). To match the large fraction of persistently nonemployed individuals, the estimation requires a large variance of fixed effects. However, this process does not offer a good fit to the data in most of the other dimensions (Güvönen *et al.* (2021)). Therefore, we report the parameters of this process and its welfare costs in Table D.2 under column Gaussian++ in Appendix C.

The next two specifications are slightly different versions of the full specification described in Section 2.3. The main difference is that one version restricts the heterogeneity in earnings growth rates by setting $\theta^i \equiv 0$ for all i in (7), whereas the other one does not.¹² We refer to these as the Benchmark-R (-R for restricted, column 3 of Table III) and Benchmark (column 4) processes, respectively.¹³

Intergenerational Correlation of Labor Productivity. Finally, a well-documented empirical fact in the US data is that parents’ and children’s labor earnings are positively correlated, with an estimated correlation around 0.5 (see, e.g., Halvorsen *et al.* (2022); Haider and Solon (2006) and the references therein). We capture this correlation by modeling (α^i, θ^i) as imperfectly transmitted from parent to children according to an AR(1) process:

$$\begin{bmatrix} \alpha_{\text{child}}^i \\ \theta_{\text{child}}^i \end{bmatrix} = \rho_\alpha \begin{bmatrix} \alpha_{\text{child}}^i \\ \theta_{\text{child}}^i \end{bmatrix} + \sqrt{(1 - \rho_\alpha^2) \Sigma_\alpha} \begin{bmatrix} \varepsilon_{\alpha_i} \\ \varepsilon_{\theta_i} \end{bmatrix}, \quad (15)$$

and set $\rho_\alpha = 0.5$. The innovations $(\varepsilon_{\alpha_i}, \varepsilon_{\theta_i})$ are drawn from standard normal distributions and Σ_α denotes the covariance matrix for (α^i, θ^i) , as we define in Section 2.3.

¹¹Although these parameter values are fairly representative of values typically used in the literature, Table D.2 shows results under different sets of parameter values estimated in other papers.

¹²The former version corresponds what is also called a “restricted income profiles” or RIP process whereas the latter corresponds to a “heterogenous income profiles” or HIP process (e.g., Güvönen (2009)).

¹³The parameter values are taken from columns (5) and (6) of GKOS’s Table IV.

Individual vs. Household Earnings Dynamics

A potential concern is that GKOS estimate their income process using earnings data for men, which leaves out the spousal labor income as a potential insurance within the household. Given the lack of household links in the data used to estimate our benchmark process, here we use the PSID data to discuss whether the moments of individual (before-tax and transfer) earnings changes—which GKOS use in their estimation—differ significantly from those for households. To compute these moments, we use the recent biennial waves of the PSID (1999-2017). These contain more detailed information on consumption expenditures, whose moments we will discuss in section 5.1. We include all households whose head is between ages 25 and 65, regardless of the gender. Appendix B contains the details of the data construction and sample selection.¹⁴

In Table D.4 we include a comparison of moments of individual and household earnings changes. The first three columns report the statistics at different ages for individual earnings; the next three columns report the same statistics for household earnings using a sample of individuals who could be single or married. The last three columns further restrict the sample to only married couples. Importantly, adding both head and spousal income does not diminish the importance of higher-order risk. Second- to fourth-order standardized moments as well as their percentile-based counterparts are very similar across all three samples and all income measures of different age groups. The similarities are particularly striking for the volatility, but also for Kelley Skewness and Crow-Siddiqui Kurtosis.¹⁵ Even for the most restricted sample, married couples, household earnings changes display similar non-Gaussian features (e.g., left skewness for those above 40 and excess kurtosis for all).¹⁶ We conclude that while our process is posed on individual earnings before taxes and transfers, it captures the household earnings dynamics well too. Section 5.1 discusses the distribution of consumption growth for the same sample from the PSID.

¹⁴We do not adjust household income or consumption by household size. Given the focus on log-changes and that household structure is quite stable for our age sample the equalisation does not make a substantial difference.

¹⁵This point is also apparent in Figure 8 in De Nardi, Fella, Knoef, Paz-Pardo and Van Ooijen (2021), which shows the second to fourth moments of individual male labor income, household labor income, and household after-tax-and-transfer income for different ages and income percentiles. Like in our sample of the PSID, the standard deviation and Kelley Skewness of labor earnings changes are remarkably similar across individuals and households. Some differences arise in Crow-Siddiqui kurtosis in the lowest deciles of the income distribution, but mostly for households whose head is above 55.

¹⁶Notice that the skewness measures from the PSID are less pronounced compared to those from the administrative data, possibly due to added noise in the PSID to the extent that measurement error is symmetric. In addition, changes in the PSID are biennial, which can smooth out earnings changes.

3.3 Implications for Non-Employment Risk

Before delving into the implications of each earnings process for consumption and welfare, we first examine their implications for nonemployment risk and lifetime earnings inequality, which are both intimately related to lifetime welfare.¹⁷

Starting with full-year non-employment risk, Figure 1 shows significant heterogeneity across workers in the data: the fraction of individuals who are non-employed next year increases sharply as earnings fall below the median of the recent earnings distribution (see also Ozkan *et al.* (2023)).¹⁸ The benchmark and benchmark-R processes capture this highly nonlinear relationship quite well, especially for the bottom 90% of the earnings distribution.¹⁹ In contrast, both versions of the Gaussian process generate almost no full-year non-employment, so the graph is completely flat, except for a little blip at the very low end of the recent earnings distribution. Because full-year non-employment spells cause substantial earnings losses today and in the future, the inability of Gaussian processes to capture this extensive margin is a crucial shortcoming of these specifications for welfare analyses of earnings risk.

3.4 Implications for Lifetime Earnings Inequality


We next ask how much lifetime earnings inequality is generated by each earnings process. This is of interest for two reasons. First, lifetime earnings inequality is intimately related to consumption inequality we study in the next section, so getting a sense about the former helps us anticipate some of the upcoming results. Second, it is not common to target lifetime earnings inequality as a moment when earnings processes are estimated in the literature, and this is also true for the four processes used in this paper. So, as a validation step, it is useful to check the extent to which each earnings process generates a plausible level of lifetime inequality compared to what we see in the data.

¹⁷As shown by GKOS, the benchmark process fits a wide range of empirical moments of individual earnings dynamics for US workers, including the levels of the higher-order moments (skewness and kurtosis) of earnings growth, their variation with age and with recent earnings levels; the impulse responses of earnings (i.e., their mean reversion patterns) and their variation by recent earnings levels as well as with the size and sign of the earnings shocks; and finally, the heterogeneity in the lifecycle earnings growth rates by lifetime earnings levels, among others. In interest of space, we do not reproduce those results here and just focus on two features of the data not discussed in detail in GKOS.

¹⁸“Full-year non-employed” is defined in GKOS as an individual whose annual earnings is below a threshold that corresponds to the pay for one quarter of half-time work (260 hours total) at the legal minimum wage in that year, \bar{Y}_t . For example, in 2010, this amount corresponded to \$1,885.

¹⁹It is worth noting that the Benchmark-R process matches the data even better than the benchmark process for the bottom 70 percent of the population.

FIGURE 1 – Full-Year Nonemployment in t by \bar{Y}_{t-1}



Notes: This figure shows the nonemployment risk between t and $t+1$ conditional on recent earnings computed as the average earnings between $t-1$ and $t-5$. Our sample includes individuals who have income above the minimum threshold \bar{Y}_t in $t-1$ and in at least two more years between $t-5$ and $t-2$.

Table IV reports various measures of dispersion for lifetime earnings. The statistics for the US data in the first column are taken from Guvenen *et al.* (2022a), who calculate lifetime earnings as the sum of annual earnings between ages 25 and 55 without discounting. They report the statistics for the sample of individuals who make at least \$50,000 in lifetime earnings in 2012 dollars (therefore including individuals with zero earnings in some years). The canonical Gaussian process severely understates lifetime inequality for all measures considered. For example, the standard deviation of log lifetime earnings is 0.4 compared with 1.3 in the data, the 90th- to 10th-percentile ratio (P90-P10) is only 2.8 compared with 14.9 in the data. The gap is even larger at the top, with a P99-P10 ratio that is an *order of magnitude* smaller than in the data (4.3 vs. 43.8).

In contrast, the benchmark process generates a much more plausible distribution of lifetime earnings, with the standard deviation of log measure slightly lower than in the data and the P90-P10 ratio close to its empirical counterpart (16.1 vs. 14.9). The benchmark process also does well at the very top, generating a P99-P10 ratio of 40.9 compared with 43.8 in the data. The only notable discrepancy from the data is that the benchmark process somewhat overstates the inequality above the median (P90-P50 ratio of 3.6 vs 2.5 in the data) and understates the inequality below the median (P50-P10 ratio of 4.5 vs 6.1 in the data). Overall, however, the benchmark process is much more closely aligned with the distribution of lifetime earnings seen in the US data compared with the canonical

TABLE IV – Lifetime Earnings Inequality: US Data vs. Canonical and Benchmark Earnings Processes

<i>Inequality measure:</i>	US Data	Canonical Gaussian	Benchmark
Std dev of log	1.32	0.40	1.08
P90/P10	14.88	2.77	16.12
P90/P50	2.46	1.68	3.56
P50/P10	6.05	1.67	4.53
P99/P10	43.82	4.31	40.85

Notes: The statistics for the US data in the first column are taken from the online appendix F of [Guvenen et al. \(2022a\)](#), who compute them from the US Social Security Administration’s Master Earnings File of individual earnings histories for the US population between 1978 and 2013. For comparability, we compute lifetime earnings in the model the same way as done in that paper, by summing annual earnings between ages 25 and 55 without discounting for the sample of individuals who earn at least \$50,000 in labor income between those ages.

TABLE V – Lifetime Earnings Inequality: Decomposing the Benchmark Income Processes

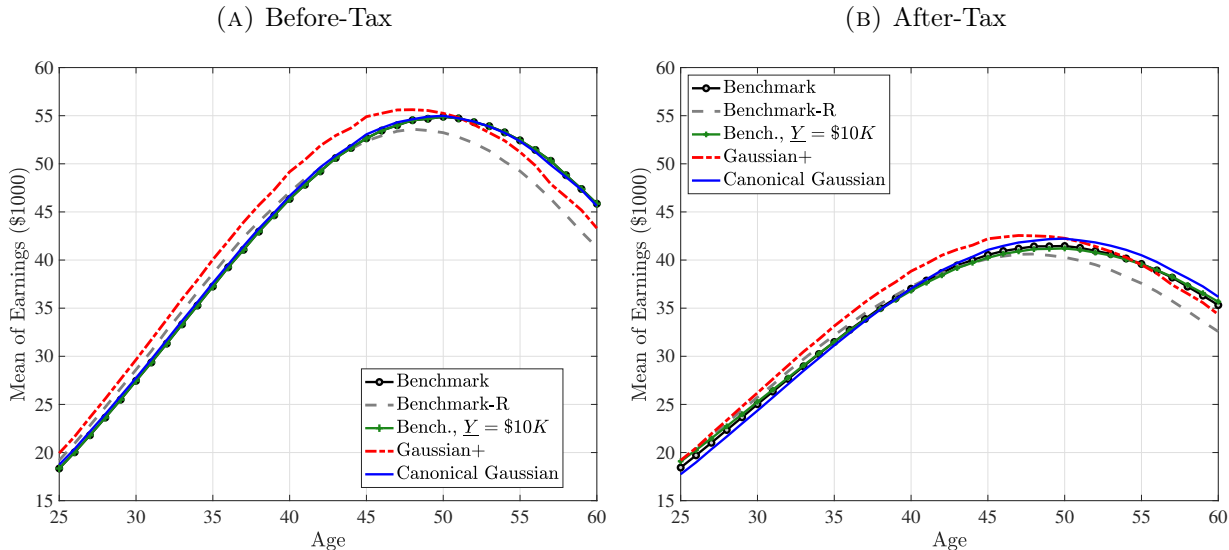
Decomposing Lifetime Inequality Generated by the Benchmark Process						
	Benchmark	No fixed initial heterogeneity	No initial heterogeneity	No nonemp. shocks	No persist. shocks	No transit. shocks
		$\sigma_\alpha, \sigma_\theta \equiv 0$	$\sigma_\alpha, \sigma_\theta, \sigma_{z_0} \equiv 0$	$\mathbf{v}_t \equiv 0$	$\mathbf{z}_t \equiv 0$	$\boldsymbol{\varepsilon}_t \equiv 0$
Std dev log	1.07	0.91	0.82	0.90	0.81	1.07
ln(90/10)	16.12	9.97	7.77	10.38	8.17	16.12
ln(90/50)	3.56	2.44	2.16	3.25	2.64	3.56
ln(50/10)	4.53	4.14	3.60	3.19	3.10	4.53
ln(99/10)	40.85	19.30	14.30	26.31	12.81	40.85

This table reports various inequality measures for lifetime income when each component of the benchmark earnings process indicated in the column header is turned off (one at a time). In the next-to-last column where we turn off persistent shocks, we set $\mathbf{z}_t \equiv 0$ in equation (7) but not in (13) so that the nonemployment probability is unaffected by this change.

Gaussian process, which understates lifetime inequality up to an order of magnitude.

A natural question to ask is how much each component of the benchmark process contributes to lifetime inequality. Table V provides the answer: Each column (after the first) shuts down one component of the benchmark process at a time and reports the resulting lifetime inequality. There are several takeaways. First, with the exception of transitory shocks (last column), all components contribute nontrivially to lifetime inequality. Second, initial conditions and persistent shocks have the largest impact on overall lifetime

FIGURE 2 – Before- and After-tax Average Income




inequality (first two rows) and top-end inequality (last row). For example, shutting down the former ($(\sigma_\alpha, \sigma_\theta, \sigma_{z_0}) \equiv 0$) lowers the P90-P10 ratio from 16.1 to 7.8, and shutting down the latter ($z_t \equiv 0$) lowers it to 8.2. Similarly, shutting down these components reduces the P99-P10 ratio to 14.3 and 12.8, respectively, from 40.9 in the full process. Third, nonemployment shocks also have a significant effect on overall inequality (reducing P90-P10 ratio from 16.1 to 10.38), and most of this effect comes from lowering inequality below the median, with the P50-P10 ratio falling from 4.5 to 3.2 while inequality above the median remains less affected (falling slightly from 3.6 to 3.3). The bottom line is that all three main components of the benchmark process matter for lifetime earnings inequality. In the next section, we will revisit this question from a slightly different angle and ask how much each contributes to the welfare costs of idiosyncratic earnings risk.

3.5 Implications for Lifecycle Profiles

Figures 2 and 3 plot the average lifecycle profiles of before-tax and after-tax income, both in levels (former) and in logs (latter). Notice that the average lifecycle profiles of both before- and after-tax earnings are quite similar across earnings processes. This is because we normalize the parameters of the deterministic lifecycle profile across earnings processes so as to generate the same average earnings (*levels*) profile from age 25 to 55 (Figure 2). However, the mean and variance profiles of *log* after-tax, after-transfer earnings look quite different across earnings processes (Figure 3) because of a Jensen inequality effect.

FIGURE 3 – After-tax, After-transfer Log Income



Namely, each process generates different variance profiles of earnings over the life cycle (e.g., benchmark and Gaussian+ generate a steeper increase in within-cohort inequality than the canonical Gaussian). Conditional on having the same (*levels* of) average earnings profiles, those that generate a higher dispersion mechanically imply a lower *log* earnings profile in Figure 3. So, for example, the benchmark process has the lowest *log* earnings profile in Figure 3 but displays the same *level* as the canonical Gaussian in Figure 2.

4 Welfare Costs of Idiosyncratic Earnings Risk

In this section, we quantify the welfare costs of non-Gaussian idiosyncratic income risk and compare them to their Gaussian counterpart. Specifically, we ask: *What fraction of consumption at every date and state would an individual in the benchmark model be willing to give up to live in a hypothetical world with no income uncertainty?* This hypothetical world is defined as one with the income process set to its average value at each age. We conduct two versions of this experiment. In the first and main exercise, we calculate the welfare effects for an individual with the average type, $\alpha^i \equiv 0$ (and $\theta^i \equiv 0$ when applicable), which abstracts from the uncertainty inherent in drawing different values of the fixed type and simply focuses on the uncertainty coming from the *stochastic evolution of earnings* over the life cycle. We believe that this calculation better captures what we have in mind by the welfare costs of idiosyncratic earnings risk.

TABLE VI – Welfare Costs of Idiosyncratic Earnings Risk

	Canonical Gaussian	Gaussian+	Benchmark-R	Benchmark
	(1)	(2)	(3)	(4)
Parameters				
β	0.985	0.975	0.963	0.960
λ	1.65	1.73	1.74	1.763
ϕ_1	0.648	0.031	1.138	1.473
ϕ_2	9.11	8.56	6.74	5.95
Welfare cost				
Fluctuations	8.48%	16.79%	36.81%	33.21%
Fluct. + Type	9.20%	24.03%	42.02%	41.51%

Note: The next to last row (“Fluctuations”) reports the risk from idiosyncratic fluctuations for an individual with the average type: $\alpha^i = 0$ (and $\theta^i = 0$ for the benchmark process). The last row reports the total risk as viewed from behind the veil of ignorance—i.e., it includes the additional risk from the uncertainty of drawing α^i (and θ^i for benchmark). The parameter τ is not included in the table because we keep it fixed at 0.185 in all calibrations.

The second exercise is conducted behind the veil of ignorance, that is, before the individual learns his type Υ^k , so it combines the risk of idiosyncratic fluctuations with that of drawing an undesirable type (low α^i). Because this is a well-understood experiment, we relegate the equations to Appendix A. In most of our discussions, we will focus on the first measure of welfare and refer to the latter only when relevant. We also compare the welfare costs of idiosyncratic risk across different income processes. Each time, we recalibrate the model (the parameters shown in Table VI) to match the wealth-to-income ratio, the distribution of bequests, and the average taxes paid in the economy.

In the benchmark model (Table VI, column 4), the individual is willing to give up about 33% of consumption at every date and state, which indicates very large welfare costs of idiosyncratic income fluctuations. The first column reports the corresponding number for the canonical Gaussian model, which is 8.5%, or a quarter of the welfare cost for the benchmark model. Column 2 shows that the Gaussian+ process generates about twice the welfare cost compared with the canonical model (16.8% vs 8.5%), which is not very surprising considering that the standard deviation of its persistent shocks are almost twice as high as the canonical model (standard deviation of 0.19 vs 0.11—see Table III).²⁰

²⁰Recall that we also consider a canonical process that can match the distribution of employment rates in Table D.2 under column Gaussian++. It should not be surprising that for this process the welfare cost behind the veil of ignorance is extremely high at 52% (because of its extremely large variance of fixed effects). Despite its large transitory and nonstationary persistent shocks, the welfare effects for an

Recall that the benchmark process includes a HIP component in addition to higher-order risk. To isolate from the effects of the former and focus on the more novel latter component, we present in Table VI column (3) the welfare cost of the benchmark-R process ($\theta^i \equiv 0$), which is even somewhat higher than the benchmark process at 36.8%. The reason for the higher welfare cost can be seen from the parameter estimates in Table III. Without the flexibility of the HIP component, the Benchmark-R process estimates a persistence parameter close to a unit root ($\rho = 0.991$) significantly higher than in the Benchmark case ($\rho = 0.959$), which is harder to self-insure.²¹ That said, both welfare figures are substantial and not materially different from each other, showing that the large welfare costs are not sensitive to whether or not a HIP process is included. (We investigate the welfare costs of idiosyncratic risk from other widely used income processes in the literature in Table D.2.)

The bottom row of Table VI shows the welfare costs including the type risk (α^i and when relevant θ^i), which increases the welfare costs across the board but does not change the substantive conclusions. For example, for the benchmark model the welfare cost behind the veil of ignorance is 41.5% compared to the 33.2% for an individual with the average type ($\alpha^i \equiv 0, \theta^i \equiv 0$). However, the average welfare costs reported in the bottom row mask significant heterogeneity across ex ante types (not reported in the table). For example, ranking all individual types k by the welfare cost they face in the benchmark model, we find that the 90th percentile of this distribution is close to 46.5%, whereas the 10th percentile is 18.5%. The highest welfare costs are associated with types who have high values of (α^i, θ^i) . That is, high-income individuals suffer more from idiosyncratic risk. Although this may seem surprising at first blush, there is a simple reason for this: these individuals are less protected by the social safety net, the magnitude of which is too small to make a difference in their income fluctuations.

Decomposing the Welfare Costs

As we did in the previous section, we seek again to decompose the contribution of each component of the benchmark process to the welfare costs we found. For this purpose, we shut down each component of the model one at a time and calculate the welfare cost again, reported in Table VII. In column (2) we shut down nonemployment shocks by setting $\nu_t \equiv 0$, which has a substantial effect on the welfare cost, reducing it from 33.2%

individual with the average type, $\alpha^i \equiv 0$ is less than 19%, which is only roughly half the welfare costs we find for the benchmark process.

²¹See Guvenen (2009) for an explanation of why restricting growth rate heterogeneity biases the estimate of ρ upward. Also, notice that even though $\rho = 0.959$ and $\rho = 0.991$ may appear close, they are quite different: for example, at 20 year horizon, the impact of a shock falls to about 43% under the former but still keeps about 83% of its initial impact under the latter.

TABLE VII – Decomposing Welfare Costs

	Benchmark	$\nu_t \equiv 0$	$z_t \equiv 0$	$\varepsilon_t \equiv 0$	No Tax	$0.5 \times \underline{Y}$	$\underline{Y} = \$10K$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Welfare costs							
Fluctuations	33.21%	22.25%	11.44%	32.68%	46.86%	42.24%	20.17%
Fluct. + Type	41.51%	33.64%	23.17%	41.34%	56.83%	49.93%	29.67%

Notes: The first row (“Fluctuations”) reports the welfare cost of risk from idiosyncratic fluctuations for an individual with the average type, i.e., $\alpha^i \equiv 0$ (and $\theta^i \equiv 0$ for the benchmark process). The second row (“Fluct. + Type”) reports the welfare cost of “total risk” as viewed from behind the veil of ignorance, by including the risk from α^i draw (and θ^i for benchmark).

to 22.3%. Recall that nonemployment shocks have an induced persistence through their dependence on z . To quantify the role of this persistence, we eliminate the dependence of the probability function, p_ν , on z_t , allowing it to vary only by age. This change has two effects: one, because z_t is very persistent, eliminating the dependence on them makes nonemployment shocks completely transitory; and two, nonemployment ceases to vary with the income level and hits all workers of a given age with the same probability of the worker with $z_t = 0$. Interestingly, the welfare costs in this case are quite close to the case if we were to eliminate nonemployment shocks completely, suggesting that most of the welfare costs of nonemployment shocks are due to their persistence and concentration among already low-income individuals.

In Column (3), we shut down persistent shocks so that they do not directly affect income. Yet, as noted earlier, they still govern the probability of nonemployment shocks in the background. This has an even larger effect compared to eliminating nonemployment shocks (in Column (2)), reducing the welfare cost to 11.4%—almost a third of its benchmark value. Finally, column 4 shows that shutting down the transitory shock also has a trivial effect on the welfare costs—reducing it to 32.7% from 32.9%.

In the next three columns, we examine the effects of two sources of insurance embedded in the model on mitigating the effects of idiosyncratic shocks. First, in column (5), we eliminate the tax system (setting $\tau = 0$ and $\lambda = 1$). This raises the welfare cost to 46.9%, highlighting the important role played by progressive taxation in smoothing idiosyncratic risk. Next, in column (6), we reduce the guaranteed minimum income level by half (from 6.75% of average earnings to 3.375%), and in column (7), we increase it to \$10,000 (22% of average earnings). The welfare costs change from 32.9% to 42.2% and 20.2% when income floor is reduced and increased, respectively. Interestingly, the welfare cost rises only by

a little in the Gaussian model when the income floor is reduced by half. The reason for this asymmetry is that in the benchmark model, individuals occasionally receive very large negative shocks (including full-year nonemployment shocks) and therefore benefit from the insurance provided by the income floor, which is less of a case in the Gaussian model, where tail shocks are much less likely. Therefore, weakening the safety net is more costly when the true income process is as in the benchmark model.

Robustness Analysis

To investigate the sensitivity of these results, we consider several alternative assumptions in Table D.3. In row (1), we replace the baseline borrowing constraint with a limit that is a constant 20% of average earnings over the life cycle. In row (2), we keep our baseline borrowing constraint but make it 50% tighter. Making the borrowing limit tighter for workers increases the welfare costs by a modest amount, from 32.9% to 34.9%. We also find similar very small effects of a tighter borrowing limit on welfare costs for the canonical Gaussian process. Thus, we conclude that our results are robust to the generosity of the borrowing limit. In row (3), we remove the warm-glow bequest motive and recalibrate the model, which turns out to have a minimal effect on welfare costs as well.

In row (4) of Table D.3, we raise relative risk aversion to $\sigma = 5$ for the canonical Gaussian and the benchmark income processes, respectively. Not surprisingly, a higher risk aversion implies a significantly higher welfare costs for both processes. Interestingly, the increase is larger for the canonical Gaussian process, for which it more than doubles from 8.5% in column (1) of Table VI to 19.7% here. The rise is also significant but smaller in percentage terms for the benchmark process, going from 33.2% in the benchmark model to 44.9%.²² This result confirms our intuition we revealed using equation 2. In particular, the higher the risk aversion, the larger the risk premium against the non-Gaussian risk compared to the Gaussian risk.

4.1 Heterogeneity in Welfare Costs of Idiosyncratic Risk

The average welfare costs reported in Table VI mask significant heterogeneity across different types of households. Next, we discuss the between-group heterogeneity in the welfare costs of idiosyncratic risk. Specifically, we investigate how different age and income groups are affected by income risk. For this purpose, we calculate the fraction of lifetime consumption each group would be willing to give up to avoid income risk and instead

²²Notice, however, that the welfare costs behind the veil of ignorance in the bottom row show a larger increase for benchmark process to 61.2 (from 41.6% in column 4 of Table VI).

TABLE VIII – Heterogeneity in Welfare Costs

Welfare Cost with Average Assets For Each Group					
	≤P25	P45–P55	P70–P80	P90–P95	P95–P100
Age	Canonical Gaussian Process				
25	6.23%	6.30%	6.32%	6.36%	6.56%
35	5.60%	5.02%	4.98%	4.91%	5.34%
45	3.61%	2.77%	2.78%	2.52%	2.86%
55	1.18%	0.70%	0.71%	0.64%	0.71%
Age	Benchmark Process				
25	25.13%	24.63%	26.46%	26.45%	27.61%
35	26.15%	16.58%	18.76%	18.39%	19.81%
45	18.27%	7.25%	9.04%	8.72%	11.41%
55	6.01%	1.50%	1.85%	1.48%	3.13%

Notes: This table shows the fraction of lifetime consumption households in each cell would be willing to give up to avoid income risk and instead receive a constant stream of income equal to the group’s average earnings at all ages going forward.

receive a constant stream of income equal to the group’s average earnings at all ages. Furthermore, to focus on the role of idiosyncratic risk individuals face going forward, we assume that each individual starts with a level of assets equal to the average asset holdings of their respective group. We also compare the benchmark process with the canonical Gaussian process. Table VIII summarizes our results.

First, the welfare costs systematically decrease by age for all income groups in both specifications. This is not surprising given that the ratio of risky human wealth relative to financial wealth—the safe asset in our model—declines by age (Huggett and Kaplan (2016)). Furthermore, labor earnings become less risky as individuals approach retirement, after which they receive a steady pension income. The specific functional form we use for the borrowing limit (equation 14), which is estimated from data by Guvenen and Smith (2014), also allows for a more generous borrowing limit for middle-age workers.

Next, looking at heterogeneity by income-group, we first observe that for the youngest workers there is relatively little variation in welfare costs of income risk over the income distribution, and relative differences grow over the working life. For example, at age 25 the welfare costs of the benchmark process vary only very little, from 26.5% to 27.6%, across

the income distribution. In contrast, among 55-year-olds, the bottom income quartile workers are willing to give up, as a fraction of their lifetime consumption, more than three times as much (6.01%) as those in the upper-middle income group (1.85%). We also find that there is relatively little variation in welfare costs of the canonical Gaussian process among the youngest workers.

Second, and interestingly, welfare costs are not monotone over the income distribution but broadly exhibit a U-shaped pattern. For example, among 45-year-olds, the welfare cost of the benchmark process declines from 18.3% for the bottom quartile workers to 7.25% for those around the median, and then it increases to 11.4% for the workers in the top 5% of the earnings distribution. This pattern can be explained by a combination of different factors. On the one hand, lower-income workers face larger idiosyncratic risk because of their higher nonemployment risk (see Figure 1), and they have less assets to insure themselves. On the other hand, they are also more protected by the social safety net, the magnitude of which is too small to make a difference in income fluctuations of high-income workers. As a result of these opposing forces, the highest welfare costs are associated with those who have the highest and, especially, lowest incomes. As for the canonical Gaussian, the increase in welfare cost is less pronounced in the high end of the income distribution.

5 Consumption Dynamics, Insurance, and MPC

In this section, we study three questions that can be included under the broad umbrella of the response of consumption to earnings changes under a non-Gaussian process. In particular, we ask: (i) How accurate is the standard approach to measuring partial insurance when shocks are non-Gaussian? (ii) How effective is self-insurance in Bewley-Aiyagari models in response to non-Gaussian risk? and (iii) How does the marginal propensity to consume (MPC) with respect to transitory (e.g., a one-time stimulus check) and persistent earnings shocks differ under different earnings processes? We study the first two questions in the following subsection.

5.1 Measuring the Insurability of Income Shocks

The transmission rate of earnings shocks to consumption has received significant attention in the literature, because it is used to measure the extent of partial insurance—that is, insurance above and beyond self-insurance (see, e.g., [Blundell *et al.* \(2008\)](#), [Primiceri and van Rens \(2009\)](#), and [Kaplan and Violante \(2010\)](#)). In an important paper, [Blundell](#)

TABLE IX – Measures of Insurability of Shocks

BPP Insurance Coefficients						
Model:	Canonical Gaussian			Benchmark		
Age:	30	40	50	30	40	50
Permanent	0.34	0.32	0.56	0.50	0.62	0.72
Transitory	0.95	0.95	0.95	0.90	0.89	0.90

True Insurance Coefficients for Persistent Shocks						
Model:	Canonical Gaussian			Benchmark		
Age:	30	40	50	30	40	50
Positive (+)	0.46	0.48	0.60	0.44	0.47	0.57
Negative (−)	0.47	0.50	0.62	0.54	0.60	0.71

Notes: The top panel reports the BPP partial insurance coefficients estimated using equations (16) and (17) following Blundell *et al.* (2008). The bottom panel reports what we call the “true” partial insurance coefficients in the model with respect to persistent shocks using equation (18). A coefficient of 1 (alternatively, 0) indicates no (full) consumption insurance.

et al. (2008) (hereafter BPP) proposed using two simple moment conditions to estimate the insurance coefficients in response to permanent shocks (η) and transitory shocks (ε), respectively. The formulas for the BPP insurance coefficients are:

$$\phi^\eta = 1 - \frac{\text{cov}(\Delta c_t^i, y_{t+1}^{\text{disp},i} - y_{t-2}^{\text{disp},i})}{\text{cov}(\Delta y_t^{\text{disp},i}, y_{t+1}^{\text{disp},i} - y_{t-2}^{\text{disp},i})} \quad (16)$$

for permanent shocks and

$$\phi^\varepsilon = 1 - \frac{\text{cov}(\Delta c_t^i, \Delta y_{t+1}^{\text{disp},i})}{\text{cov}(\Delta y_t^{\text{disp},i}, \Delta y_{t+1}^{\text{disp},i})} \quad (17)$$

for transitory shocks. The insurance coefficients range from 0 to 1, where 0 means no partial insurance above self-insurance (or full transmission) and 1 means perfect consumption insurance (no transmission). Here, we consider the following experiment. Suppose we give a panel dataset of income and consumption simulated from the benchmark model to an econometrician and ask her to estimate the insurance coefficients using BPP’s moments. What would the econometrician conclude about the extent of the insurability of income shocks?

The top panel of Table IX reports the results at different ages. When the true data-generating process is the benchmark model, the BPP procedure estimates that 62% of permanent shocks are insured and the remaining 38% is transmitted to consumption at age 40. The corresponding insurance coefficient for the Gaussian process is much lower—almost half—at 32% (so 68% transmitted). Taken at face value, these findings would indicate that persistent income shocks in the benchmark model are more insurable relative to those in the Gaussian model, which seems surprising given that both the long tails and the negative skewness of the benchmark process would be expected to be harder (at least, not easier) to insure such shocks. To understand this result, note that equations (16) and (17) are derived under the assumption that the earnings process is given by the persistent-plus-transitory process such as the canonical Gaussian model, but unlike the benchmark process. The discrepancy could happen if the benchmark *process* features less persistent shocks (which is partly true, at least judging crudely based on the value of ρ —0.98 vs. 0.959—and the presence of the more transitory non-employment shock process) or if the benchmark *lifecycle model* features more insurance opportunities relative to the Gaussian model (which is not true—they both have self-insurance plus an income floor only).

To avoid these problems, a more direct measure of the partial insurance coefficient with respect to persistent shocks can be obtained from simulated data by directly looking at the covariation between consumption growth and the innovation, η_t , to the persistent component, z_t . Using this covariation, we can also go one step further and measure what we call the “true” partial insurance coefficient with respect to positive and negative innovations separately using this equation:

$$\chi_{\text{insur.}}^{\eta} = 1 - \frac{\text{cov}(\Delta c_t^i, \eta_t \mid \eta_t > 0)}{\text{var}(\eta_t \mid \eta_t > 0)}. \quad (18)$$

The bottom panel of Table IX reports the true coefficients for both processes, which reveal two results. First, under the Gaussian process, the insurance against positive and negative persistent shocks are similar, whereas under the benchmark process, negative shocks are better insured than positive shocks. Second, the insurability of positive persistent shocks are very similar between the two processes, whereas negative shocks are slightly more insurable under the benchmark process. This asymmetry is because the typical negative shock in the benchmark process is drawn from a thicker tail (due to both the negative skewness and excess kurtosis) and therefore reduces the earnings to a level at which the guaranteed minimum income threshold kicks in to smooth those shocks.

To complement this analysis, we consider another way to quantify the degree of partial


insurance commonly used in the literature— by studying how within-cohort consumption inequality evolves over the life cycle. The idea is that if shocks are easily insurable (because either they are not persistent or the economy features rich smoothing opportunities) then consumption inequality should not rise much with age. To investigate this, Figure 4 plots the cross-sectional variance of log consumption in the benchmark and Gaussian models. In the former, consumption inequality rises by about 45 log points from age 25 to age 60, which is about four times as large as the rise in the Gaussian model (about 12 log points). Therefore, this second way of looking at the degree of partial insurance leads to the opposite conclusion: shocks in the benchmark model are less insurable, which causes consumption inequality to rise more than in the Gaussian model. This latter evidence is also more consistent with the higher welfare costs of idiosyncratic shocks that we found for the benchmark model relative to the Gaussian model above. That said, this result depends somewhat on the availability of public insurance. In the presence of a more generous income floor (i.e., $\underline{Y} = \$10,000$), consumption inequality rises by around 30 log points in the benchmark model, which is still much higher than in the canonical Gaussian model.

Recall that the benchmark process generates a steeper increase in within-cohort, after-tax, after-transfer income inequality than the canonical Gaussian specification even with a \$10,000 income floor (see Figure 3). The total effect of earnings shocks on consumption responses also depends on the size of shocks, which depends on the shape of the distribution it is drawn from. The benchmark process with its long tails and left skewness can produce shocks that generate a larger income inequality and stronger *total* consumption response.

Taken together, we conclude from these two pieces of evidence that once we move beyond Gaussian shocks and linear models, extra care is needed to properly measure the extent of insurability of shocks. Nonlinear dynamics and higher-order moments generate interesting new patterns. For example, why does the transmission parameters indicate a low rate of transmission in the benchmark model? An important reason is that the partial insurance coefficient measures the *average* response of consumption growth to income shocks. But it is plausible to expect that the consumption response varies by the sign of the shock, by its size (Beraja and Zorzi, 2024), and by many of its other properties established in the previous sections. So, the average response coefficient could provide an incomplete picture of the transmission of income shocks to consumption.

A natural question that this discussion raises concerns the properties of the entire distribution of consumption growth rates implied by a model, including its higher-order moments. We turn to this point next.

FIGURE 4 – Within-Cohort Variance of Log Consumption Over the Life Cycle



Notes: The green line with “+” markers is simulated using the benchmark income process but by raising the minimum income floor to \$10,000.

Higher-Order Moments of Consumption Growth

Table (X) reports the standard deviation, skewness and kurtosis of consumption growth in the US data (from the PSID), the Gaussian model, and both versions of the benchmark model. For the empirical distribution, we use a total household consumption measure for the sample described in Subsection 3.2; see Appendix B for details.

There are several takeaways from Table (X). First, the standard deviation of consumption growth is twice as high in the benchmark model than for the Gaussian specification, partly reflecting the larger magnitude of shocks in the former. Second, consumption growth is about three times more volatile in the PSID compared with the benchmark process. Although this may suggest that, despite the large welfare costs we found above, consumption smoothing may still be a bit too effective in the model relative to the data, it may also very well be a reflection of the high measurement error found in survey data. The discrepancy may also reflect other shocks not captured in our model (such as marriage, divorce, children, moving, health, and so on).

Third, consumption growth is negatively skewed in the benchmark model, and much more so in the Benchmark-R model, but has slightly positive skewness in the Gaussian model. The minimum income floor induces slight positive skewness in both models—without it, the benchmark model delivers even more negatively skewed consumption growth. Interestingly, consumption growth becomes *less* negatively skewed with age, even

TABLE X – Cross-Sectional Moments of Consumption Growth in the Model

<i>Age</i>	US Data (PSID)			Canonical Gaussian			Benchmark-R			Benchmark		
	30	40	50	30	40	50	30	40	50	30	40	50
Coefficient on:												
Standard deviation	0.46	0.41	0.40	0.06	0.06	0.05	0.14	0.12	0.09	0.13	0.13	0.11
Kelley skewness	0.01	-0.07	-0.01	0.01	0.02	0.02	-0.64	-0.65	-0.62	-0.19	-0.11	-0.10
Crow-Siddiqui kurtosis	3.69	3.63	3.57	2.97	2.95	2.97	6.00	5.73	6.65	6.30	6.82	7.67

Notes: Table X shows the moments of consumption changes. refers to households income after taxes and transfers. The statistics are computed from the PSID waves 2005–2021. In the data, age groups are defined including plus and minus three years.

though income shocks become *more* negatively skewed, a divergence that seems to be because precautionary wealth allows better smoothing at older ages. Skewness is only slightly negative in the PSID, but measurement error may bias it toward zero to the extent that it is symmetric. Indeed, [Constantinides and Ghosh \(2017\)](#) show that household consumption growth is more robustly negatively skewed in the Consumer Expenditure Survey (CE) data, which may suffer less from measurement error (at least over shorter horizons) due to its design that is squarely focused on consumption expenditures.

Finally, consumption growth in the benchmark model has high excess kurtosis, as measured by the Crow-Siddiqui measure. The Gaussian model delivers almost no excess kurtosis in consumption growth, with a Crow measure of 2.95–2.97 compared with 2.91 for a standard Gaussian distribution. The empirical distribution of consumption growth does have excess kurtosis, ranging between 3.57 and 3.69; however, this is significantly lower than what is implied by both versions of the non-Gaussian benchmark model.²³

Deviations from lognormality of the consumption growth distribution have been studied in some recent papers. In addition to [Constantinides and Ghosh \(2017\)](#) just noted above, [Toda and Walsh \(2015\)](#) use CE data to document that the household consumption growth distribution has long tails in the form of a double-Pareto distribution, consistent with excess kurtosis in the benchmark model studied here.²⁴ In very recent work, [Rodriguez Mora et al. \(2022\)](#) make use of millions of transactions from bank records to show that indeed the full distribution of consumption growth deviates from a normal distribution differently across

²³The third and fourth standardized moment measures of skewness and kurtosis—which we do not report here for brevity—confirm the same finding. Further, Figure C.1 in Appendix B confirms that the difference in kurtosis in the empirical distribution of consumption growth compared to the Gaussian reference (3.63 compared to 2.95 at age 40, for example) is indeed significant.

²⁴Separately, [Brav et al. \(2002\)](#) argued that accounting for the skewness of consumption growth distribution helps improve the performance of asset pricing models for explaining the equity premium observed in the data.

ages and income groups.

Overall, these results suggest that while the non-Gaussian earnings process provides a step toward better capturing the properties of the consumption growth distribution in the data, important discrepancies remain that invite further research on this topic. There are channels of adjustment missing in this paper that can also help reconcile theory and data at the tails, such as durable (lumpy) adjustments, as discussed in [Beraja and Zorzi \(2024\)](#).

5.2 The Marginal Propensity to Consume (MPC)

The marginal propensity to consume (MPC) out of earnings (or wealth) is a key ingredient in designing effective fiscal and monetary policies (e.g., [Kaplan *et al.* \(2018\)](#)). The substantial concentration of wealth and earnings in the United States and in many modern economies implies that the effectiveness of such policies can be significantly improved if they are targeted to certain segments of the population *if* MPCs vary in a meaningful way across the population.²⁵ In this section, we investigate the implications of non-Gaussian earnings risk for the level and distribution of MPCs in the population.


MPC Response to \$500 Stimulus Check

We start with a standard exercise and measure the MPC in response to a one-time \$500 stimulus check, similar to the analyses conducted before (e.g., [Johnson *et al.*, 2006](#); [Parker *et al.*, 2013](#); [Kaplan and Violante, 2014](#)). Figure (5) reports the results for different earnings processes. Starting with the benchmark process, we see that the MPC distribution is declining with the cash-on-hand level of individuals and does so in a very convex fashion at low cash-on-hand levels: it falls from about 35% for the bottom 1% group to 15% at the 10th percentile and to about 7.5% at the 30th percentile. It then falls gradually to 5% for individuals in the top 1% of the cash-on-hand distribution. The benchmark-R process shows a similar pattern with a slightly lower level at the low end.

In contrast, the MPC is much lower and very flat under the Gaussian process: it ranges between 4% and 5% from the 5th percentile all the way up to the top 1 percent group (with the exception of the very bottom of the distribution). What is also interesting is that rather than declining with cash-on-hand, it bottoms out around the 40th percentile and increases slightly thereafter. The pattern is essentially the same, with a slightly higher level at the low end, for the Gaussian+ process. The figure also plots the MPCs in the benchmark model when the minimum income floor is raised to \$10,000, which significantly lowers the

²⁵See [Kaplan and Violante \(2022\)](#) for a recent review of the extensive recent literature that studies heterogeneity in MPCs.

FIGURE 5 – MPC out of \$500 Stimulus Check



Notes: The figure plots the MPCs in response to a one-time \$500 stimulus check.


level of the MPC and flattens its shape across the distribution. Although \$10,000 seems too high relative to the empirical estimate of welfare and social safety net programs we use in our baseline calibration (taken from [Guner *et al.* \(2023\)](#)), this result still underscores the importance of the minimum income floor level used in calibration for the quantitative conclusions one reaches in models with idiosyncratic earnings risk. This is all the more true when the non-Gaussian features of earnings dynamics, such as negative skewness and long tails, are modeled.

MPC Response to 1% Persistent or Transitory Earnings Shocks

The appeal of the stimulus check exercise is that it is informative for actual stimulus policies regularly implemented by governments. But a different type of exercise is also informative for other contexts: the MPCs in response to a fixed percentage increase in persistent or transitory earnings. The difference of course is that these increases are anchored to the earnings level of the individual rather than being a fixed dollar amount. These MPCs would be relevant, for example, for a fixed reduction or increase in the labor income tax rate that can be permanent or transitory, or for an increase or decrease in inflation, among others. Analyzing the MPC out of a persistent increase in earnings is also useful because the persistence of earnings shocks in the benchmark process is not solely captured by ρ of the AR(1) process but also through the nonemployment shocks (see [Guvenen *et al.* \(2021\)](#) for a discussion).

The left panel of Figure 6 plots the MPC out of a 1% transitory earnings shock. The overall patterns are similar to the stimulus check exercise above, with some notable dif-

FIGURE 6 – MPC out of a 1% Earnings Shock



Notes: The figures plot the MPCs in response to a transitory increase in earnings (ε) (left panel) and persistent increase (right panel) that is equal to 1% of each individual’s earnings level.

ferences: the MPC is higher for the benchmark process—close to 5% for the median cash-on-hand individual compared with about 3% for the Gaussian model; it declines with cash-on-hand levels, albeit at a slower (less convex) fashion than the fixed stimulus payment; the Gaussian model displays little heterogeneity across earnings levels; and raising the minimum income floor substantially reduces the MPC in the benchmark model.

There are more interesting differences and similarities in the right panel, which plots the MPCs in response to a 1% innovation to the persistent component (z_t). (Notice that the likelihood of the nonemployment shock depends on z_t , so it rises in response to the reduction in z_t , which our calculations take into account.) In this case, both the level and shape of the MPC distribution are fairly similar between the Gaussian and benchmark processes. The MPC at the median cash-on-hand level is about 40% for the benchmark model and 36% for the Gaussian process. The benchmark-R process has a higher MPC at 44% at the median, owing to the higher estimated ρ for that process. As before, increasing the minimum income floor to \$10,000 reduces the MPC for the benchmark model to less than 30% at the median.

6 Conclusions

The nature of income risk has important consequences for many economic decisions and distributional phenomena that involve risk and influence the design of optimal social

insurance and taxation. Recent contributions to the literature have shown significant nonlinearities and nonnormalities in income dynamics (e.g., [Guvenen *et al.* \(2021\)](#); [Arellano *et al.* \(2017b\)](#)). There is a growing literature incorporating nonlinear, non-Gaussian earnings dynamics into quantitative models. In this paper we have studied how higher-order income risk can have quite different implications for consumption dynamics, consumption insurance, and welfare compared with the commonly used Gaussian processes. The overarching finding from our analysis is that the benchmark process implies more severe income risk—e.g., larger lifetime earnings inequality and welfare costs; negatively skewed and leptokurtic consumption growth; larger MPCs out of transitory earnings shocks—than previously modeled through canonical Gaussian processes.

This finding could have implications for a wide range of economic questions where earnings inequality and earnings risk play a central role. For example, an important point of our exercise is the crucial role of public insurance mechanisms against (negative) tail risk in earnings. While an even more generous income floor does not fully insure against idiosyncratic negative changes, our analysis suggests that this channel is quantitatively important to dampen the welfare cost of higher-order earnings risk. While we capture all transfers with an income floor (and progressive taxation), different instruments can be mapped into our catch-all parameter. For example, in the context of idiosyncratic risk correlated with aggregate shocks, [Busch *et al.* \(2022\)](#) find that unemployment insurance is the public instrument that smoothes variation in higher-order earnings risk the most across different countries, more than progressive taxation, welfare, pensions, and SSI. Our result thus has implications for the design of public insurance with varying idiosyncratic income risk and the potential welfare gains of automatic stabilizers.

Our last finding is that a non-Gaussian earnings risk also implies a stronger consumption response to earnings shocks. The cross-sectional distribution of consumption growth implied by the benchmark model displays strong higher-order moments, as well as higher volatility—closer to the distribution of consumption growth from the PSID—than a Gaussian earnings process. Overall, these results suggest that the non-Gaussian earnings process provides a step toward better capturing the properties of the consumption growth distribution in the data. Yet, there are still important discrepancies: incorporating new channels of consumption adjustment, such as durable (lumpy) adjustments ([Beraja and Zorzi \(2024\)](#)), or advancing consumption measurement beyond and as a complement to survey data ([Rodriguez Mora *et al.* \(2022\)](#)) can help better reconcile the theoretical and empirical distribution of consumption growth.

References

- ANDRESKI, P., LI, G., SAMANCIOGLU, M. Z. and SCHOENI, R. (2014). Estimates of annual consumption expenditures and its major components in the psid in comparison to the ce. *American Economic Review*, **104** (5), 132–35.
- ARELLANO, M., BLUNDELL, R. and BONHOMME, S. (2017a). Earnings and consumption dynamics: a nonlinear panel data framework. *Econometrica*, **85** (3), 693–734.
- , — and — (2017b). Earnings and consumption dynamics: a nonlinear panel data framework. *Econometrica*, **85** (3), 693–734.
- BELL, F. C. and MILLER, M. L. (2002). *Life Tables for the United States Social Security Area: 1900-2100*". Actuarial Study 116, Office of the Actuary, Social Security Administration.
- BENABOU, R. (2000). Unequal societies: Income distribution and the social contract. *American Economic Review*, **90** (1), 96–129.
- BERAJA, M. and ZORZI, N. (2024). *Durables and Size-Dependence in the Marginal Propensity to Spend*. Working Paper 32080, National Bureau of Economic Research.
- BLUNDELL, R., PISTAFERRI, L. and PRESTON, I. (2008). Consumption inequality and partial insurance. *American Economic Review*, **98** (5), 1887–1921.
- , — and SAPORTA-EKSTEN, I. (2012). *Consumption Inequality and Family Labor Supply*. Working Paper Series 18445, National Bureau of Economic Research.
- , — and — (2016). Consumption inequality and family labor supply. *American Economic Review*, **106** (2), 387–435.
- BRAV, A., CONSTANTINIDES, G. and GECZY, C. (2002). Asset pricing with heterogeneous consumers and limited participation: Empirical evidence. *Journal of Political Economy*, **110** (4), 793–824.
- BUSCH, C., DOMEIJ, D., GUVENEN, F. and MADERA, R. (2022). Skewed Idiosyncratic Income Risk over the Business Cycle: Sources and Insurance. *American Economic Journal: Macroeconomics*, **14** (2), 207–242.
- and LUDWIG, A. (2024). Higher-order income risk over the business cycle. *International Economic Review*, **forthcoming**.
- COMMAULT, J. (2022). How do persistent earnings affect the response of consumption to transitory shocks?
- CONSTANTINIDES, G. M. (2021). *Welfare Costs of Idiosyncratic and Aggregate Consumption Shocks*. Working Paper 29009, National Bureau of Economic Research.
- and DUFFIE, D. (1996). Asset pricing with heterogeneous consumers. *The Journal of Political Economy*, **104** (2), 219–240.
- and GHOSH, A. (2017). Asset pricing with countercyclical household consumption risk. *The Journal of Finance*, **72** (1), 415–460.
- DE NARDI, M., FELLA, G., KNOEF, M., PAZ-PARDO, G. and VAN OOIJEN, R. (2021). Family

- and government insurance: Wage, earnings, and income risks in the netherlands and the u.s. *Journal of Public Economics*, **193**, 104327.
- DE NARDI, M., FELLA, G. and PAZ-PARDO, G. (2020). Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare. *Journal of the European Economic Association*, **18** (2), 890–926.
- DE NARDI, M., FELLA, G. and PAZ-PARDO, G. (2020). *Wage Risk and Government and Spousal Insurance*. NBER Working Papers 28294, National Bureau of Economic Research, Inc.
- and YANG, F. (2016). Wealth inequality, family background, and estate taxation. *Journal of Monetary Economics*, **77**, 130 – 145.
- FRIEDMAN, M. (1957). *A Theory of the Consumption Function*. Princeton, NJ: Princeton University Press.
- GALE, W. G. and SCHOLZ, J. K. (1994). Intergenerational transfers and the accumulation of wealth. *Journal of Economic Perspectives*, **8** (4), 145–160.
- GHOSH, A. and THELOUDIS, A. (2023). *Consumption Partial Insurance in the Presence of Tail Income Risk*. Papers 2306.13208, arXiv.org.
- GUNER, N., KAYGUSUZ, R. and VENTURA, G. (2023). Rethinking the welfare state. *Econometrica*, **91** (6), 2261–2294.
- , RAUH, C. and VENTURA, G. (2022). *Means-Tested Transfers in the US: Facts and Parametric Estimates*. Working paper, Arizona State University.
- GUVENEN, F. (2009). An Empirical Investigation of Labor Income Processes. *Review of Economic Dynamics*, **12** (1), 58–79.
- , KAPLAN, G., SONG, J. and WEIDNER, J. (2022a). Lifetime incomes in the united states over six decades. *American Economic Journal: Macroeconomics*, **14** (4), 446–479.
- , KARAHAN, F., OZKAN, S. and SONG, J. (2021). What Do Data on Millions of U.S. Workers Say About Labor Income Risk? *Econometrica*, **89** (5), 2303–2339.
- , PISTAFERRI, L. and VIOLANTE, G. (2022b). Global trends in income inequality and income dynamics: New insights from grid. *Quantitative Economics*, **13**, 1321–1360.
- and SMITH, A. A. (2014). Inferring Labor Income Risk and Partial Insurance from Economic Choices. *Econometrica*, **82** (6), 2085–2129.
- HAIDER, S. J. and SOLON, G. (2006). "life-cycle variation in the association between current and lifetime earnings. *American Economic Review*, **96** (4), 1308–1320.
- HALVORSEN, E., HOLTER, H. A., OZKAN, S. and STORESLETTEN, K. (2023). Dissecting Idiosyncratic Earnings Risk. *Journal of the European Economic Association*, **22** (2), 617–668.
- , OZKAN, S. and SALGADO, S. (2022). Earnings dynamics and its intergenerational transmission: Evidence from norway. *Quantitative Economics*, **13** (4), 1707–1746.
- HAYASHI, F., ALTONJI, J. and KOTLIKOFF, L. (1996). Risk-sharing between and within families. *Econometrica*, **64** (2), 261–94.
- HEATHCOTE, J., STORESLETTEN, K. and VIOLANTE, G. L. (2014). Consumption and labor

- supply with partial insurance: An analytical framework. *American Economic Review*.
- HUBBARD, R. G., SKINNER, J. and ZELDES, S. P. (1995). Precautionary saving and social insurance. *The Journal of Political Economy*, **103** (2), 360–399.
- HUGGETT, M. and KAPLAN, G. (2016). How large is the stock component of human capital? *Review of Economic Dynamics*, **22**, 21–51.
- HURD, M. D. and SMITH, J. P. (2002). Expected bequests and their distribution.
- JOHNSON, D. S., PARKER, J. A. and SOULELES, N. S. (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review*, **96** (5), 1589–1610.
- KAPLAN, G., MOLL, B. and VIOLANTE, G. L. (2018). Monetary policy according to hank. *American Economic Review*, **108** (3), 697–743.
- and VIOLANTE, G. L. (2010). How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics*, **2** (4), 53–87.
- and — (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, **82** (4), 1199–1239.
- and — (2022). The marginal propensity to consume in heterogeneous agent models. *Annual Review of Economics*, **14**, 747–775.
- KARAHAN, F. and OZKAN, S. (2013). On the persistence of income shocks over the life cycle: Evidence, theory, and implications. *Review of Economic Dynamics*, **16** (3), 452–476.
- KREBS, T. (2003). Growth and Welfare Effects of Business Cycles in Economies with Idiosyncratic Human Capital Risk. *Review of Economic Dynamics*, **6** (4), 846–868.
- KRUSELL, P., MUKOYAMA, T., SAHIN, A. and ANTHONY A. SMITH, J. (2009). Revisiting the Welfare Effects of Eliminating Business Cycles. *Review of Economic Dynamics*, **12** (3), 393–402.
- LUCAS, R. E. (1987). *Models of Business Cycles*. New York: Basil Blackwell.
- OZKAN, S., SONG, J. and KARAHAN, F. (2023). Anatomy of lifetime earnings inequality: Heterogeneity in job-ladder risk versus human capital. *Journal of Political Economy Macroeconomics*, **1** (3), 506–550.
- PARKER, J. A., SOULELES, N. S., JOHNSON, D. S. and MCCLELLAND, R. (2013). Consumer spending and the economic stimulus payments of 2008. *American Economic Review*, **103** (6), 2530–53.
- PRIMICERI, G. E. and VAN RENS, T. (2009). Heterogeneous life-cycle profiles, income risk and consumption inequality. *Journal of Monetary Economics*, **56** (1), 20 – 39, carnegie-Rochester Conference Series on Public Policy: The Causes and Consequences of Rising Economic Inequality April 25-26, 2008.
- PRUITT, S. and TURNER, N. (2020). Earnings risk in the household: Evidence from millions of us tax returns. *American Economic Review: Insights*, **2** (2), 237–54.
- RODRIGUEZ MORA, J. V., BUDA, G., CARVALHO, V., HANSEN, S., ORTIZ, A. and RODRIGO, T. (2022). *National Accounts in a World of Naturally Occurring Data: A Proof of Concept for Consumption*. CEPR Discussion Papers 17519, C.E.P.R. Discussion Papers.

- STORESLETTEN, K., TELMER, C. I. and YARON, A. (2001). The welfare cost of business cycles revisited: Finite lives and cyclical variation in idiosyncratic risk. *European Economic Review*, **45** (7), 1311–1339.
- , — and — (2004). Cyclical dynamics in idiosyncratic labor market risk. *Journal of political Economy*, **112** (3), 695–717.
- TODA, A. A. and WALSH, K. (2015). The double power law in consumption and implications for testing euler equations,. *Journal of Political Economy*, **123** (5), 1177–1200.
- WANG, C., WANG, N. and YANG, J. (2016). Optimal consumption and savings with stochastic income and recursive utility. *Journal of Economic Theory*.

Supplemental Online Appendix

NOT FOR PUBLICATION

A Details of Consumption Model

A.1 Numerical Solution of the Model

In order to be able to fully capture the features of the rich earnings dynamics we estimated in this paper, we chose not to discretize the income process when solving the consumption model. Thus it is worth discussing our numerical solution methodology to solve an otherwise standard Bewley model. We employ value function iteration to solve for the policy function for the consumption decision and then use the policy function to simulate consumption-savings decisions of individuals. We now discuss the details below.

A.1.1 Bounds and Grids

For a given (α^k, θ^k) – type worker, the value function has three continuous variables for each age t , asset holdings \mathbf{a}_t , and two persistent components, $z_{1,t}, z_{2,t}$, one or both of which may be fully transitory. We start simulating the continuous income process which is used in the simulation of the consumption path of individuals. We take the minimum and maximum values of simulated persistent components to be the bounds of the $z_{1,t}$ and $z_{2,t}$ spaces. We choose grid points in the $z_{1,t}$ and $z_{2,t}$ spaces to be equally spaced.²⁶ We have 11 and 61 grid points for $z_{1,t}$ and $z_{2,t}$ spaces, respectively.

As for the asset space, the borrowing limit is set by equation 14. The upper bound for the asset spaced is chosen such that in simulations no individual’s asset holdings come close to this upper bound, A_t . We choose grids in the \mathbf{a}_t dimension to be polynomial spaced between $\mathbf{a}_t = \bar{A}_k$ and $\mathbf{a}_t = A_t$, with an explicit point at $\mathbf{a}_t = 0$. We have 100 grid points for the asset space.

A.1.2 Integration and Interpolation

The value function of a type- k individual with (α^k, θ^k) is given by:²⁷

$$V_t^{i,k}(\mathbf{a}_t^i, z_{1,t}^i, z_{2,t}^i) = \max_{c_t^i, \mathbf{a}_{t+1}^i} u(c_t^i) + \beta \mathbb{E}_t \left[V_{t+1}^{i,k}(\mathbf{a}_{t+1}^i, z_{1,t+1}^i, z_{2,t+1}^i) \mid z_{1,t}^i, z_{2,t}^i \right].$$

We define the conditional expectation function:

$$\tilde{V}_{t+1}^{i,k}(\mathbf{a}_{t+1}^i, z_{1,t}^i, z_{2,t}^i) = \mathbb{E}_t \left[V_{t+1}^{i,k}(\mathbf{a}_{t+1}^i, z_{1,t+1}^i, z_{2,t+1}^i) \mid z_{1,t}^i, z_{2,t}^i \right].$$

Then the value function becomes:

$$V_t^{i,k}(\mathbf{a}_t^i, z_{1,t}^i, z_{2,t}^i) = \max_{c_t^i, \mathbf{a}_{t+1}^i} u(c_t) + \beta \tilde{V}_{t+1}^{i,k}(\mathbf{a}_{t+1}^i, z_{1,t}^i, z_{2,t}^i),$$

which reduces the interpolation to only along the \mathbf{a}_{t+1} dimension, for which we use one-dimensional spline interpolation.

²⁶Since $z_{1,t}$ and $z_{2,t}$ are in log values polynomially spaced grid points along these dimensions performed worse.

²⁷The budget constraint and the other equations that properly define the value function are omitted here for simplicity.

The conditional expectation, $\tilde{V}_{t+1}^{i,k} \left(\mathbf{a}_{t+1}^i, z_{1,t}^i, z_{2,t}^i \right) = \mathbb{E}_t \left[\mathbf{V}_{t+1}^{i,k} \left(\mathbf{a}_{t+1}^i, z_{1,t+1}^i, z_{2,t+1}^i \right) \mid z_{1,t}^i, z_{2,t}^i \right]$ involves integration over the innovations to both of the persistent components, $\eta_{1,t}^i$ and $\eta_{2,t}^i$ and a non-employment shock, v_t^i . Probability of non-employment shocks is a function of $\left(z_{1,t+1}^i, z_{2,t+1}^i \right)$ given by eq. 13. Then for given $\left(z_{1,t+1}^i, z_{2,t+1}^i \right)$ we first take the expectation over v_t^i . For this purpose, we approximate the non-employment shocks using the Gauss-Laguerre abscissas and weights which is the quadrature that is used for exponential distribution. In particular, let $\left\{ x_{v,t}^1, x_{v,t}^2, \dots, x_{v,t}^{n_v}, x_{v,t}^{n_v+1} = 0 \right\}$ and $\left\{ w_{v,t}^1, w_{v,t}^2, \dots, w_{v,t}^{n_v}, w_{v,t}^{n_v+1} = 1 - p_{v,t} \right\}$ be the Gauss-Laguerre abscissas and weights for the v_t^i , including the case of full employment, $x_{v,t}^{n_v+1} = 0$. Note that the weights depend on $\left(z_{1,t+1}^i, z_{2,t+1}^i \right)$. Then, we define the conditional expectation function over non-employment shocks as the following:

$$\begin{aligned} \mathbb{V}_{t+1}^{i,k} \left(\mathbf{a}_{t+1}^i, z_{1,t+1}^i, z_{2,t+1}^i \right) &= \mathbb{E}_{t,v} \left[\mathbf{V}_{v,t+1}^{i,k} \left(v_t^i; \mathbf{a}_{t+1}^i, z_{1,t+1}^i, z_{2,t+1}^i \right) \mid z_{1,t+1}^i, z_{2,t+1}^i \right] \\ &= \sum_j^{n_v+1} w_{v,t}^j \mathbb{V}_{v,t+1}^{i,k} \left(x_{v,t}^j; \mathbf{a}_{t+1}^i, z_{1,t+1}^i, z_{2,t+1}^i \right). \end{aligned}$$

Innovations to the persistent components, $\eta_{1,t}$ and $\eta_{2,t}$ are drawn from a mixture of two normal distributions. Let p_{z_1} and p_{z_2} be probability to draw an innovation from the first normal distribution for persistent components z_1 and z_2 , respectively. The normal distributions are approximated using the Gauss-Hermite abscissas and weights. Let $\left\{ x_{j,t}^{m,1}, x_{j,t}^{m,2}, \dots, x_{j,t}^{m,n_j^m} \right\}$ and $\left\{ w_{j,t}^{m,1}, w_{j,t}^{m,2}, \dots, w_{j,t}^{m,n_j^m} \right\}$ be the Gauss-Hermite abscissas and weights for the innovations from normal distribution $m = 1, 2$ to the persistent component $j = 1, 2$. Then, the conditional expectation function over the shocks to the persistent components is defined as:

$$\begin{aligned} \tilde{V}_{t+1}^{i,k} \left(\mathbf{a}_{t+1}^i, z_{1,t}^i, z_{2,t}^i \right) &= \mathbb{E}_t \left[\mathbb{V}_{t+1}^{i,k} \left(\mathbf{a}_{t+1}^i, z_{1,t+1}^i, z_{2,t+1}^i \right) \mid z_{1,t}^i, z_{2,t}^i \right] \\ &= p_{z_1} p_{z_2} \sum_{j_1}^{n_1^1} \sum_{j_2}^{n_2^1} w_{1,t}^{1,j_1} w_{2,t}^{1,j_2} \mathbb{V} \left(\mathbf{a}_{t+1}^i, \rho_1 z_{1,t}^i + x_{1,t}^{1,j_1}, \rho_2 z_{2,t}^i + x_{2,t}^{1,j_2} \right) \\ &+ (1 - p_{z_1}) p_{z_2} \sum_{j_1}^{n_1^2} \sum_{j_2}^{n_2^2} w_{1,t}^{2,j_1} w_{2,t}^{1,j_2} \mathbb{V} \left(\mathbf{a}_{t+1}^i, \rho_1 z_{1,t}^i + x_{1,t}^{2,j_1}, \rho_2 z_{2,t}^i + x_{2,t}^{1,j_2} \right) \\ &+ p_{z_1} (1 - p_{z_2}) \sum_{j_1}^{n_1^1} \sum_{j_2}^{n_2^2} w_{1,t}^{1,j_1} w_{2,t}^{2,j_2} \mathbb{V} \left(\mathbf{a}_{t+1}^i, \rho_1 z_{1,t}^i + x_{1,t}^{1,j_1}, \rho_2 z_{2,t}^i + x_{2,t}^{2,j_2} \right) \\ &+ (1 - p_{z_1}) (1 - p_{z_2}) \sum_{j_1}^{n_1^2} \sum_{j_2}^{n_2^2} w_{1,t}^{2,j_1} w_{2,t}^{2,j_2} \mathbb{V} \left(\mathbf{a}_{t+1}^i, \rho_1 z_{1,t}^i + x_{1,t}^{2,j_1}, \rho_2 z_{2,t}^i + x_{2,t}^{2,j_2} \right) \end{aligned}$$

The above part requires interpolation over $\left(z_{1,t+1}^i, z_{2,t+1}^i \right)$ for which we use 2-dimensional nested cubic spline. We use 7 abscissas for innovations to the each of the persistent component.

A.1.3 Simulating the Consumption Path

The policy function for the consumption decision has three continuous variables, asset holdings, and two persistent components, $\mathbf{a}_t, z_{1,t}, z_{2,t}$. To interpolate the consumption decision we employ the 3-dimensional nested cubic spline interpolation. We do not need an additional state for the non-employment shocks since they are transitory in nature. Namely, we store the consumption decision in the value function iteration step for only the case of full-year non-employed shocks. In simulating the non-employment shocks less than a year, we just add the difference in labor earnings due to the shorter non-employment spell to the asset holdings of the individuals, which can be thought as cash-on-hand. This allows us to interpolate the consumption decision only over three dimensions, $(\mathbf{a}_t, z_{1,t}, z_{2,t})$ rather than 4-dimensional interpolation.

A.1.4 Other Computational Details

We use MPI for the parallel programming. In particular, we use one core for each worker type- k (with a specific (α^k, θ^k)). The values of (α^k, θ^k) and the number of individuals to be simulated for each type k type are determined using the Gauss-Hermite abscissas and weights, respectively. We have five α^k -types, three θ^k -types. In total, we have 15 ex-ante types of workers and used 15 cores to compute for 15 different value functions. Once each core finishes simulating consumption path for its own type of worker, the main core gathers all the simulated data and compute statistics for the entire economy. It takes 4-5 mins to solve and simulate the model for a given discount factor.

B Empirical Appendix

B.1 Data and Sample

The PSID is a longitudinal study of a representative sample of U.S. households, including income and an extensive list of consumption categories from 2005 to 2021 at a biennial frequency.²⁸ For those waves, a new set of questions targeted to understanding spending dynamics allow us to track consumption on almost all expenditures measured in the cross-sectional Consumer Expenditure Survey (CEX) [Andreski et al. \(2014\)](#). We focus on a sample of the original SRC households that participate in the labor market and are between 25 and 65 years old. We further focus on those households in which the head of household is a male and his labor earnings are at least \$1500 in the given year. This sample is comparable to that in [Blundell et al. \(2016\)](#). We end up with 29,644 observations, spanning 16 years at biennial frequency. We define *household consumption* as the sum of spending in food (at home, away, and delivery), gasoline, health, transport, utilities, clothing, and recreation. [Table C.1](#) summarizes the characteristics of the sample and the subcomponents of consumption.

B.2 Cross-Sectional Distribution of Consumption Growth: Densities

²⁸The PSID has been running since 1968, but the expansion to consumption beyond food and housing only happened in the 1999 wave, with information on 1998 spending for a few consumption categories. In 2005, the PSID finalized the expansion with a comprehensive list of consumption categories.

	(1) 2005-2020
<i>Demographics</i>	
Age	43.27
Labor Income, Head	67,008.74
Labor Income, Spouse	29,022.54
Household Gross Labor Income	96,031.28
Household Net Labor Income	76,259.84
Expenditures	17,181.38
Foot at Home	6,712.55
Foot Away from Home	2,806.67
Food dDelivery	182.29
Gasoline	2,729.81
Clothing	1,639.76
Utilities	2,969.63
Telecommunications	2,649.15
Auto Insurance	1,787.80
Parking	70.68
Bus and Train	77.11
Taxi	53.97
Other Transportation	118.18
Home Insurance	747.15
Education	1,992.09
Childcare	870.33
OOP Medical	520.32
Doctor	870.28
Prescription	413.40
Health Insurance	2,560.80
Trips	2,168.39
Other Recreation	1,107.93
N	29,644

TABLE C.1 – Detailed Summary Statistics

FIGURE C.1 – The Distribution of Earnings and Consumption Changes in the Data: Log Densities

(A) Household Net Labor Earnings ($\mu = 0.04, \sigma = 0.53$) (B) Household Consumption ($\mu = 0.01, \sigma = 0.42$)



Notes: Figure C.1 shows the log density income and consumption changes in our sample. Income refers to households income after taxes and transfers. The sample includes households whose heads are males, between 25-65, and earnings at least \$1500 in a given year from the PSID waves 2005-2021. All distributions are truncated at -2 and 2, so the extremes include all changes below -2 and above 2.

FIGURE D.1 – Before-tax, Before-transfer Log Income


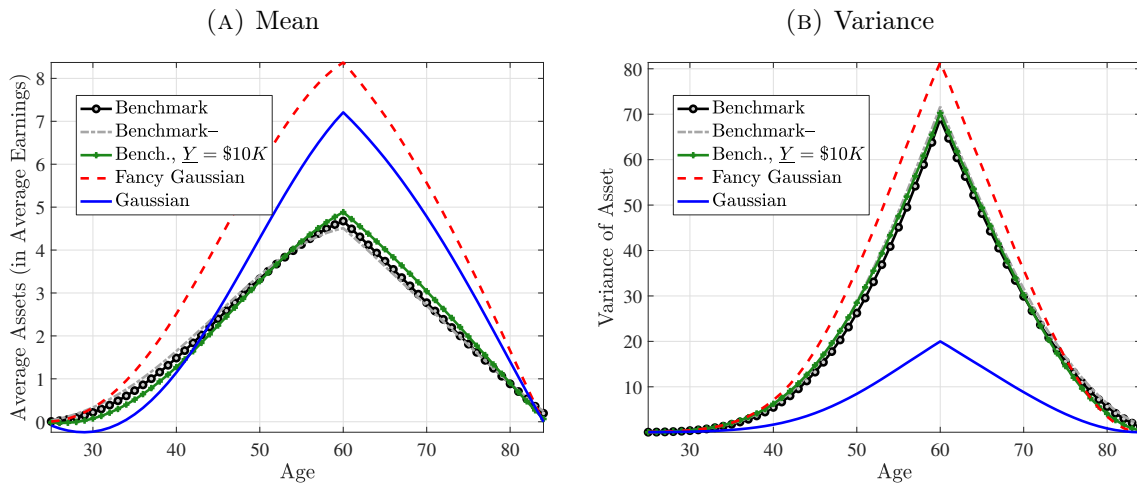



FIGURE D.2 – Asset profile



Notes:


C Additional Figures and Tables

FIGURE D.3 – Consumption Insurance-Persistent shocks



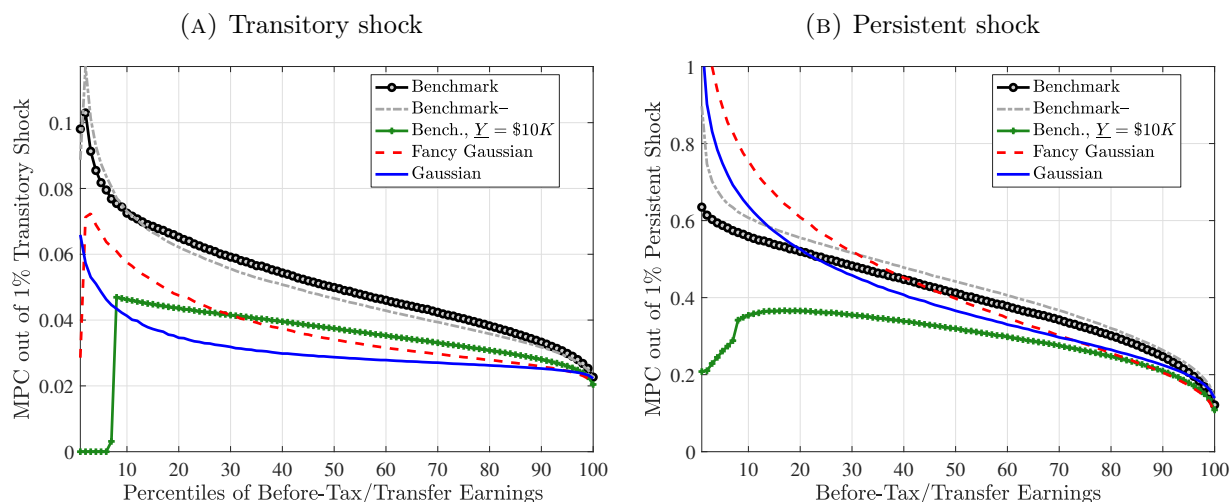
Notes: This figure shows the nonemployment risk between t and $t + 1$ conditional on recent earnings between $t - 1$ and $t - 5$. Benchmark model (left panel) is the non-Gaussian, nonlinear income process. Gaussian (right panel) is a random walk income process with Gaussian shocks, which is estimated so as to generate similar inequality to the benchmark process.

FIGURE D.4 – Consumption Insurance - Transitory shocks



Notes: This figure shows the nonemployment risk between t and $t + 1$ conditional on recent earnings between $t - 1$ and $t - 5$. Benchmark model (left panel) is the non-Gaussian, nonlinear income process. Gaussian (right panel) is a random walk income process with Gaussian shocks, which is estimated so as to generate similar inequality to the benchmark process.

FIGURE D.5 – MPC out of a 1% Shock: Variation By Before-Tax Earnings Percentiles



Notes: The figures plot the MPCs in response to a transitory increase in earnings (ϵ) (left panel) and persistent increase (right panel) that is equal to 1% of each individual's earnings level. The difference from Figure (6) is that here the MPCs are plotted as a function of before-tax/transfer earnings percentiles.

TABLE D.1 – Heterogeneity in Welfare Costs

Welfare Cost with <i>Cohort Average Assets</i>					
	≤P25	P45–P55	P70–P80	P90–P95	P95–P100
Age	Canonical Gaussian				
25	7.06%	7.03%	7.03%	7.05%	7.26%
35	5.89%	5.43%	5.49%	5.55%	6.17%
45	3.17%	2.90%	3.19%	3.21%	3.99%
55	0.81%	0.72%	0.90%	1.01%	1.45%
Age	Benchmark Process				
25	27.37%	27.43%	29.40%	28.36%	28.41%
35	25.33%	16.71%	20.75%	20.65%	22.11%
45	11.22%	6.45%	10.39%	11.81%	14.91%
55	2.53%	1.00%	2.19%	2.92%	6.32%

Note: This table is analogous to Table VIII except that we give each individual the average wealth of the entire cohort at that age rather than the average wealth of the age/income cell she belongs to.

TABLE D.2 – Welfare Costs of Gaussian Processes

	Storesletten <i>et al.</i> 2004	Kaplan and Violante 2010	Karahan and Ozkan 2013	GKOS	GKOS
	(1)	(3)	Gaussian	Gaussian+	Gaussian++
Parameters					
σ_α	0.448	0.387	0.120	0.420	1.182
ρ	0.957	1.0	0.980	0.981	1.005
σ_η	0.188	0.10	0.110	0.190	0.134
σ_{z_0}	0.0	0.0	0.28	0.01	0.343
σ_ε	0.35	0.25	0.30	0.49	0.696
β	0.987	0.993	0.985	0.975	0.969
τ	0.185	0.185	0.185	0.185	0.185
λ	1.75	1.68	1.65	1.79	2.22
ϕ_1	0.001	0.001	0.648	0.031	0.264
ϕ_2	9.77	7.17	9.11	8.516	9.71
Fluctuations	11.3%	8.25%	8.48%	16.79%	52.00%
Fluct. + Type	20.00%	15.17%	9.20%	24.03%	18.77%

TABLE D.3 – Robustness Analysis

		Canonical Gaussian	Benchmark
		(1)	(4)
Specification			
(1)	20%AE Borr.	8.82%	34.90%
(2)	50%Fancy	8.29%	34.62%
(3)	No Beq	8.40%	32.67%
(4)	$\sigma = 5$	19.70%	44.86%

Note: The next to last row (“Fluctuations”) reports the risk from idiosyncratic fluctuations for an individual with the average type: $\alpha^i = 0$ (and $\theta^i = 0$ for the benchmark process). The last row reports the total risk as viewed from behind the veil of ignorance—i.e., it includes the additional risk from the uncertainty of drawing α^i (and θ^i for benchmark).

TABLE D.4 – Cross-Sectional Moments of Earnings Growth, Panel Study of Income Dynamics

Age	Individual Earnings			Household Earnings			Household Earnings		
	<i>All households</i>			<i>All households</i>			<i>Married households</i>		
	30	40	50	30	40	50	30	40	50
Standard deviation	0.63	0.53	0.51	0.58	0.50	0.46	0.50	0.46	0.42
Skewness coeff	0.05	-0.42	-0.17	0.14	-0.49	-0.36	0.13	-0.26	-0.47
Kurtosis coeff.	8.72	13.22	13.68	8.61	14.68	15.59	9.48	17.40	18.07
P90-P10	1.22	0.88	0.83	1.17	0.87	0.78	1.03	0.78	0.72
Kelley skewness	0.07	0.03	-0.06	0.08	-0.02	-0.09	0.10	0.04	-0.09
Crow-Siddiqui kurtosis	6.53	7.91	8.32	5.64	6.80	7.15	5.06	6.33	6.37

Notes: Table D.4 shows the moments of 2-year earnings changes. Earnings refers to individual or households labor income after taxes and transfers. The sample for the first column includes all households including singles with male or female heads; the sample for the second column is the same but the earnings measure is household earnings in the case of married couples. The sample for the last column includes married households only. The statistics are computed from the PSID waves 2005–2021. Age groups are defined including plus and minus three years.