A Tractable Income Process for Business Cycle Analysis[*]

Fatih Guvenen[†] Alisdair McKay[‡] Conor Ryan[§]

February 3, 2023

Abstract

We estimate an income process that is consistent with key facts on individual income risk and its variation over the business cycle. In particular, the estimated process generates income fluctuations that display (i) flat and acyclical variance, (ii) volatile and procyclical skewness, (iii) very high kurtosis, and (iv) a moderate rise in cross-sectional inequality over the life cycle, all consistent with the US data. Furthermore, the income process captures the predictable nature of business cycle income risk: income changes during a business cycle episode are partly predicted by income levels before that episode. The estimated process features a time-varying distribution of innovations as well as a factor structure for business cycle exposure. Incorporating the estimated process into a business cycle model adds only one state variable—as in the workhorse persistent-plus-transitory income process—making it a tractable option for modelers.

JEL Codes: D31, E24, E32, H31

Keywords: Idiosyncratic income risk, business cycle models, higher-order risk, non-Gaussian risk, skewness, kurtosis, factor structure.

[*]For helpful comments, we thank participants at the 20th Macro-Finance Society Workshop in Athens, the Bank of Canada Annual Research Conference, and the UCL-FCA interdisciplinary seminar.

[†]University of Minnesota, FRB of Minneapolis, and NBER; guvenen@umn.edu

[‡]Federal Reserve Bank of Minneapolis; alisdair.mckay@mpls.frb.org

[§]Pennsylvania State University; conor.ryan@psu.edu

1 Introduction

Recent years have seen a surge of interest in heterogeneous-agent models for business cycle analysis. Many of these models build on the Bewley-Huggett-Aiyagari tradition, in which uninsurable idiosyncratic income risk is a key driver of ex post heterogeneity across individuals. Traditionally, the calibration of the income process in these models has been based on the second moment properties of income dynamics estimated from survey-based panel data. In recent years, the increasing availability of large administrative panel data sets enabled a more precise estimation of the properties of the higher-order moments of income growth, including how they vary over the business cycle. These recent studies have found that the distribution of income growth rates has mostly flat and acyclical variance, procyclical skewness, very high kurtosis, and unequal exposure to the business cycle across the income distribution (a factor structure).[1] In this paper, we estimate an income process that can capture these features of the data, yet is tractable enough to be incorporated in macroeconomic models.

The income process we consider allows for three key departures from the workhorse persistent-plus-transitory Gaussian model of income dynamics. First, our modeling allows for fat-tailed *declines* in income that partially revert after a period, leaving behind a potentially long-lasting "scarring" effect. While our modeling is purely statistical in nature, these shocks can capture the types of patterns commonly associated with nonemployment: a transitory income loss during nonemployment and a partial recovery with reemployment, followed by a persistent scarring effect. The scarring effect allows the model to generate a left tail of the income growth distribution that is thicker than the right tail, as seen in the data (see Figure 2). Second, our income process features a persistent first-order autoregressive process (AR(1)) with innovations drawn from a normal mixture distribution that allows the flexibility to generate nonnormalities (skewness and excess kurtosis) in income dynamics. We find that introducing time variation by allowing the mean of the normal distributions in the mixture to depend on average wage income growth each year provides a good fit to business cycle variation in the data moments, including the procyclical fluctuations in skewness. Third, the model includes a factor structure, whereby workers in different parts of the income distribution can exhibit

[1] See, for example, Guvenen, Ozkan, and Song (2014); Arellano, Blundell, and Bonhomme (2017); Harmenberg (2021); Kramarz, Nimier-David, and Delemotte (2021); Guvenen, Pistaferri, and Violante (2022).

different sensitivities (or exposures) to aggregate fluctuations. This factor structure captures a systematic component in idiosyncratic income fluctuations, which has empirical support in the data (see Figure 4).[2]

An important advantage of our specification is that it introduces only one state variable to a dynamic programming problem—just as the workhorse persistent-plus-transitory model does—while capturing significantly more complex dynamics. In Section 6, we discuss the steps involved in incorporating this income process into dynamic models.

The features of the income data that we emphasize in this paper have important implications for a number of applied problems. For example, leptokurtic idiosyncratic income risk has an important effect on the value of social insurance (Saez, 2001; Golosov, Troshkin, and Tsyvinski, 2016) and interacts with borrowing and saving decisions, with consequences for the distribution of wealth and marginal propensities to consume (Kaplan, Moll, and Violante, 2018). Another example is that the unequal incidence of business cycle fluctuations has important implications for the welfare cost of business cycles (e.g., Storesletten, Telmer, and Yaron, 2001; Krebs, 2003, 2007), the conduct of stabilization policy (e.g., McKay and Reis, 2021; Bhandari, Evans, Golosov, and Sargent, 2021), and asset pricing (e.g., Mankiw, 1986; Constantinides and Duffie, 1996; Schmidt, 2014; Constantinides and Ghosh, 2016), among others.

We estimate the econometric model by targeting a list of data moments that capture the levels of the first four moments of the individual income growth distribution as well as the time variation in the first three moments over three decades, starting in the early 1980s. These moments have been estimated from panel data on individual income histories from Social Security Administration (SSA) records for male workers, reported by Guvenen, Ozkan, and Song (2014) and Guvenen, Karahan, Ozkan, and Song (2021). We present results for model specifications that build in complexity, starting from the persistent-plus-transitory Gaussian model and culminating with our full benchmark model. In doing so, we show how the model elements we add allow us to fit particular moments in the data. Some features of the data, and therefore some aspects of our income process, may be more or less critical in a particular application.

While idiosyncratic risk in heterogeneous-agent business cycle models has traditionally been modeled as a linear-Gaussian income process that is estimated to match the

[2]See, e.g., Guvenen, Ozkan, and Song (2014); Guvenen, Schulhofer-Wohl, Song, and Yogo (2017) for the US and Bell, Bloom, and Blundell (2021) for the UK.

second moments of income dynamics, some recent studies have started to incorporate higher-order moments. The income process in Kaplan et al. (2018) captures the leptokurtic nature of income growth rates but does not feature any business cycle variation apart from the level of income. McKay (2017), McKay and Reis (2021), and Catherine (2021) incorporate income processes that allow for procyclical skewness of income growth, but do not target the high kurtosis of income growth rates and do not allow for the factor structure in the exposure to the business cycle. Bhandari et al. (2021) allow for a factor structure but do not target higher-moment properties of income risk.

The paper is organized as follows. Section 2 presents the features of the data that we seek to match with our income process. Section 3 presents the parametric specification of the process. Section 4 gives the details of the moments we seek to match and describes our estimation procedure. Section 5 presents the estimation results and describes how the components of the process relate to particular aspects of the data. Section 6 provides guidance on how the income process can be incorporated into the dynamic programming problems commonly used in heterogenous agent models of the business cycle.

2 Motivating Facts

In this section, we present the key features of individual income dynamics that we believe are of most relevance to heterogeneous-agent models of business cycles. We take these empirical moments from Guvenen et al. (2014) and Guvenen et al. (2021), who estimated them from panel data on the income histories of US (male) workers from US Social Security records covering the period from 1978 through 2011.

The first feature is the age profile of within-cohort income inequality, shown in Figure 1. This age-inequality profile has been a key target for heterogeneous-agent models to match since it was first documented (e.g., Deaton and Paxson (1994); Storesletten et al. (2004)). The profile is informative about both the persistence and size of (persistent) income shocks. Figure 1 shows that the cross-sectional variance of log income grows nearly linearly with age. Under the assumption that all individuals share the same deterministic lifecycle profile for income, a linear age-variance profile implies random walk behavior for persistent shocks.[3] Moreover, the dispersion in incomes within a cohort also informs us about the dispersion of individual fixed effects. If we allow for heterogeneity

[3] If $z_t = z_{t-1} + \eta_{t-1}$, the cross-sectional variance at age t is $\text{var}(z_t) = \text{var}(z_{t-1}) + \sigma^2_\eta$, which grows linearly with t.

Figure 1 – Within-Cohort Variance of Log Income over the Life Cycle

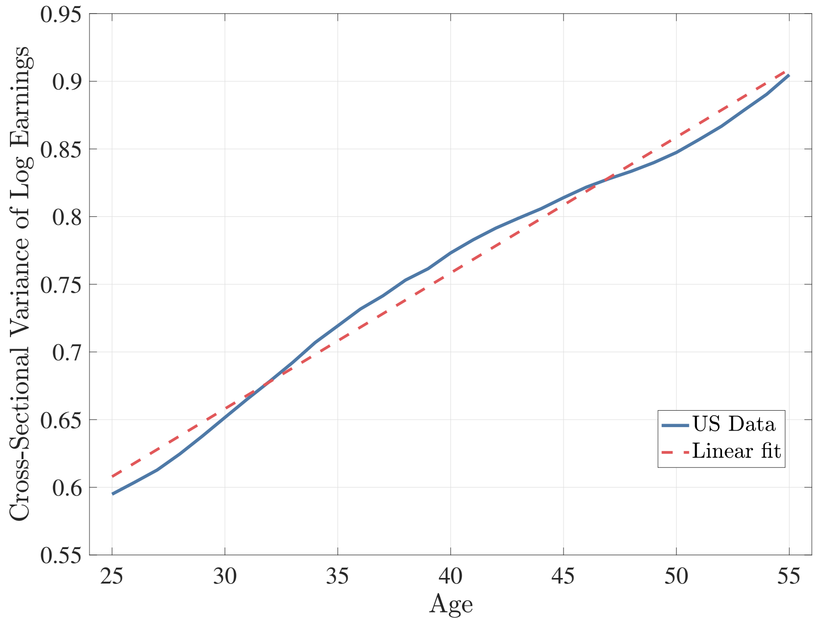

Note: The cross-sectional variance of log income increases almost linearly over the life cycle. The blue line represents the variance of log earning within each age group. The dashed line shows the fit of a linear regression of cross-sectional variance on age.

in income growth rates (or heterogeneous income profiles, HIP), the shape of the age-variance profile informs us about the combined effects of dispersed deterministic income profiles and the accumulated persistent shocks.

While the lifecycle variance profile speaks to the persistence and variance of income innovations, it does not characterize the full shape of the distribution from which they are drawn. Figure 2 shows the empirical log-density of annual growth in log income superimposed on a Gaussian (normal) distribution with the same mean and variance.[4] The distribution has high kurtosis, as demonstrated by a dramatic peak at the center and long, thick tails—a stark contrast to the Gaussian distribution. Moreover, both tails are approximately linear, which corresponds to a double-Pareto distribution in the tails. In fact, this linearity is present over a very wide range, between annual log growth rates of 1 and 4 on the right side (roughly corresponding to 3-fold to 55-fold increases in income) and between −1 and −4 on the left (corresponding to 68% to 98% declines). The figure also makes clear how much a Gaussian density understates the likelihood of tail

[4]The empirical density is for the 1995–1996 change and is taken from Guvenen et al. (2021), who argue that the figure is qualitatively the same in other years in their sample.

FIGURE 2 – Log Density of Annual Income Growth (with Double-Pareto Tails)

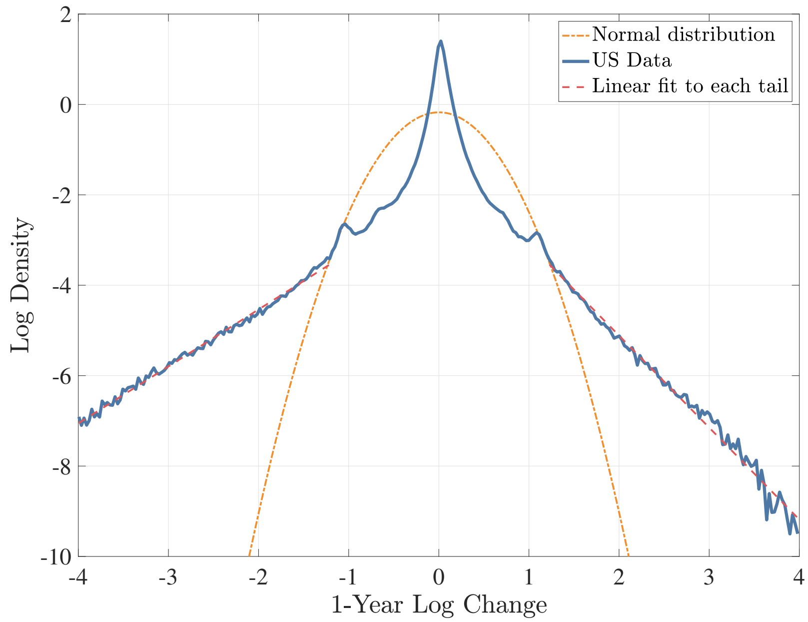

Note: The figure shows the distribution of one-year income growth between 1995 and 1996 ((Guvenen et al., 2021)). The dashed lines show the fit of a linear regression of the log-density on the one-year log change within the right and left tails. The dash-dot line shows a Gaussian distribution with the same mean and variance as the empirical distribution.

shocks. For example, a log income decline of –2 (i.e., a decline of 86%) is 100 times more likely (log likelihood ratio 4.6) in the data than what would be predicted by a Gaussian distribution with the same variance as in the data.

Another feature of the empirical density is its asymmetry, which is evident in the shape of the tails. The slope of the log-density is significantly steeper in the right tail (slope of –2.2) than in the left tail (slope of 1.4): negative income shocks have a fatter tail than positive income shocks. Capturing the complex shape of this distribution—its large variance, negative skewness, high kurtosis, and its long, asymmetric double-Pareto tails—is one of our objectives.

We now turn to the business cycle variation in income dynamics. While average income falls in a recession, the impact of business cycles is not felt equally across the population. One manifestation of this unequal incidence is that the distribution of income shocks changes in recessions. Using administrative data on individual income histories, Guvenen et al. (2014) show that the variance of income growth rates is flat and largely

FIGURE 3 – Skewness and Dispersion of Five-year Income Growth

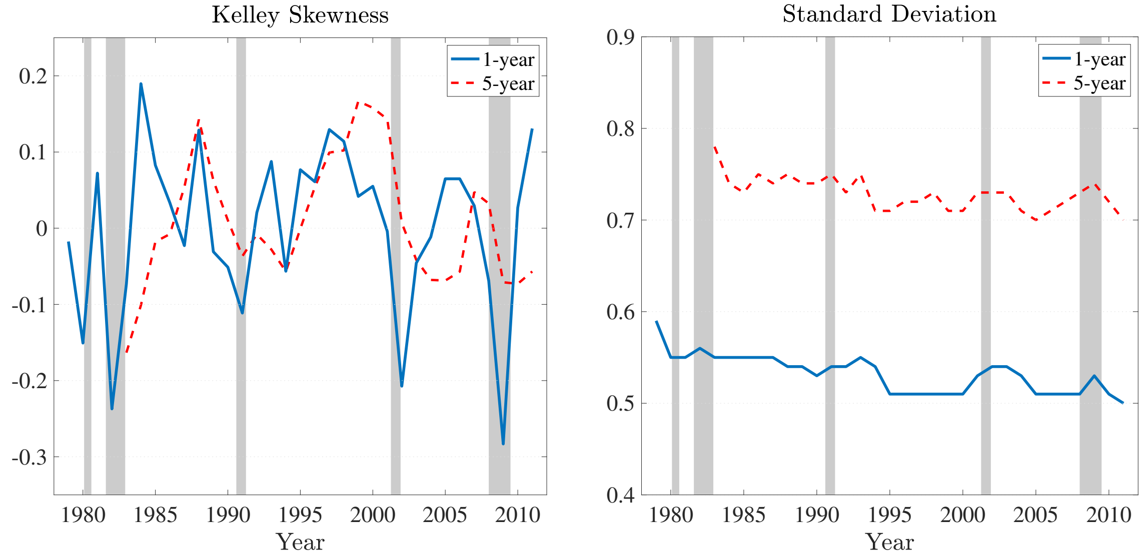

Note: The left panel displays the time series of Kelley skewness of the distribution of one-year (solid line) and five-year (dashed line) income growth. The right panel shows the analogous standard deviations. The plots are aligned with the first year of the time differenced data.

acyclical, whereas the skewness of income growth rates is strongly pro-cyclical. This can be seen in Figure 3, which plots the Kelley skewness and standard deviation of one-year and five-year income growth rates over time.[5] The shaded regions depict NBER recessions. In every recession, the skewness of one-year income growth falls significantly, while the standard deviation shows only limited fluctuations over time and no discernible cyclical pattern. Procyclical skewness is a natural pattern to expect: during recessions, large negative income shocks are more prominent while large income gains are not, and vice versa during expansions.

The second dimension of unequal business cycle incidence is that the sensitivity of a worker's income to aggregate fluctuations depends on the position of that worker in the income distribution. This factor structure can be seen in Figure 4, which shows how workers in different parts of the income distribution fared during the Great Recession. Specifically, workers are sorted into 100 percentile bins based on their five-year average income before the recession started (2002 to 2006), as shown in the x-axis. The y-axis shows the (log) change in the average income of each percentile group from 2007 to

[5] Kelley's skewness is a robust measure of skewness and is calculated as $\mathcal{S}_\kappa = [(P90 - P50) - (P50 - P10)] / (P90 - P10)$.

Figure 4 – Factor Structure: Income Changes during the Great Recession Varied Systematically with Pre-recession Income Rank

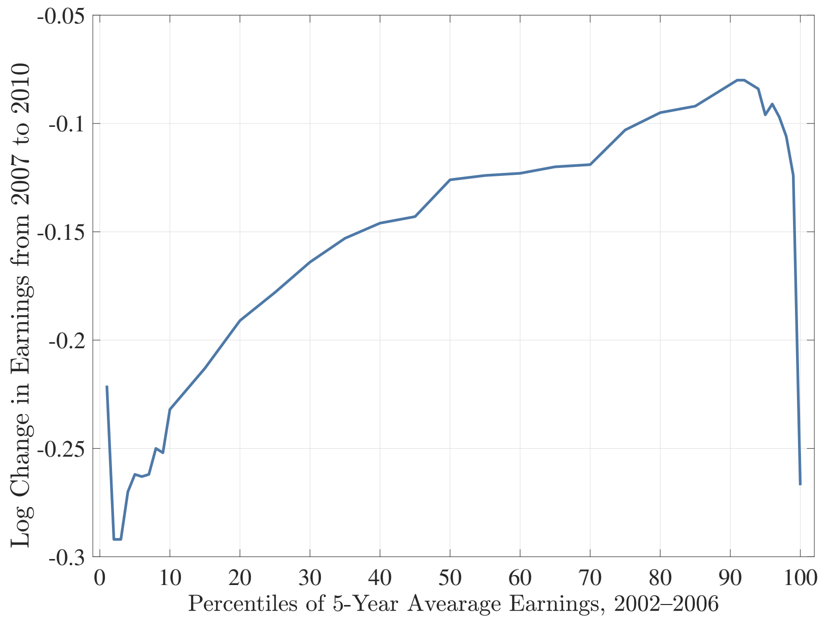

Note: The figure plots the log income change during the Great Recession for prime-age males ranked by their pre-recession average income (2002 to 2006). The income data are averaged within each percentile bin before calculating the log difference between 2007 and 2010. The figure shows income growth after removing age effects, so the average income growth implied by the figure will not match aggregate income growth.

2010. The graph shows a clear systematic pattern. First, for workers below the 90th percentile, the lower a worker's average income was before the recession, the larger was the average income loss experienced during the Great Recession. The magnitude of this systematic variation is large—those who entered the recession at the 10th percentile of income lost 18 percentage points more than those who entered at the 90th percentile. Because these losses are correlated with income levels, the factor structure causes a widening of income inequality during the recession. However, and second, the pattern reverses itself for workers in the top 10 percent: the higher a worker's average income was before the recession the larger were the losses he experienced during the recession. The magnitude of these losses was nearly as large as that of the losses in the bottom 10% of the distribution. As we show below, expansions show the mirror image pattern: income gains are larger at the top and bottom than in the middle. The factor structure we introduce into the income processes will aim to capture these systematic and highly nonlinear variations.

3 Income Process Specification

In this section, we describe a stochastic income process that aims to capture the features of the data described above. Much like the widely used linear-Gaussian processes, the income process requires only a single individual state variable and is thus tractable enough to incorporate into macroeconomic business cycle models. The general process for log income, $y_{i,t}$, is

$$y_{i,t} = \gamma_i + z_{i,t} + \tilde{\zeta}_{i,t} + \left[1 + f(\gamma_i + z_{i,t})\right] w_t + \kappa_i (t - h_i), \tag{1}$$

where γ_i is an individual fixed effect distributed with standard deviation σ^{γ}; $\tilde{\zeta}_{i,t}$ is a transitory shock; $z_{i,t}$ is the persistent idiosyncratic state; and w_t is the aggregate cyclical component. The effect of the aggregate component on an individual's income is mediated by the factor structure, f, a function of the fixed effect and the persistent state, which evolves as an AR(1) process:

$$z_{i,t} = \rho z_{i,t-1} + \tilde{\eta}_{i,t} \tag{2}$$

with $z_{i,0} = 0$ as initial condition. The parameter ρ governs the persistence, and $\tilde{\eta}_{i,t}$ is the innovation. The final term in (1) allows for heterogeneous income profiles (HIP): κ_i is the slope of the individual income profile distributed with standard deviation σ^{κ}, and h_i records the cohort year of the individual (the year he or she turns 25).

We allow for correlation between the transitory and persistent innovations, which we interpret as "scarring" effects of the transitory shock. Specifically, we allow for two independent shocks $\zeta_{i,t}$ and $\eta_{i,t}$ that determine $\tilde{\zeta}_{i,t}$ and $\tilde{\eta}_{i,t}$ according to

$$\begin{aligned} \tilde{\zeta}_{i,t} &= (1 - \psi)\zeta_{i,t} \\ \tilde{\eta}_{i,t} &= \eta_{i,t} + \psi \zeta_{i,t}, \end{aligned} \tag{3}$$

where the parameter ψ determines the extent of the scarring effect.

Each of the elements of the income process plays a role in allowing us to fit the features of the data that we highlighted in Section 2. A natural starting point for a model of labor income dynamics is a simple persistent-transitory specification with normal innovations, which is a special case of our model if we assume that ζ and η are drawn from i.i.d. normal distributions, there is no income scarring effect ($\psi = 0$), aggregate shocks affect all individuals equally ($f(\gamma_i + z_{i,t}) \equiv 0$), and income profiles are restricted ($\kappa_i = 0$).

While such a model can capture the mean and variance of income growth, it fails to match the richer moments of the data described above.

Relative to that starting point, income scarring and a non-Gaussian distribution for ζ_{it} allow the model to match the tails of the income distribution shown in Figure 2 and thereby generate the high kurtosis of the distribution. We let ζ_{it} follow a "nonemployment" process in which ζ_{it} is equal to 0 with probability p^ζ and equal to $\log(1-\ell_{i,t})$ with probability $1-p^\zeta$, and $\ell_{i,t}$ is drawn from an exponential distribution with parameter λ conditional on being in the interval $[0,1]$. Intuitively, this shock represents a risk that an individual's income will be cut by a factor $1-\ell_{i,t}$. This shock generates a long left tail of income risk through its arrival (e.g., job loss) and a long right tail of income risk through its departure (e.g., job gain). Second, we allow the transient income shock to have a scarring effect on income if $\psi > 0$. As a result of this partial persistence, the individuals that experience a negative income shock do not return to the same income level when the negative shock disappears. This asymmetry can generate the steeper right tail of income growth.

We assume a flexible distribution of innovations to the persistent component, η, which allows the model to match the acyclical variance and procyclical skewness of income growth as well as its excess kurtosis. We assume that η is drawn according to

$$\eta_{i,t} \sim \begin{cases} N(\mu^\eta_{1,t}, \sigma^\eta_1) & \text{with prob. } p^\eta_1, \\ N(\mu^\eta_{2,t}, \sigma^\eta_2) & \text{with prob. } p^\eta_2, \\ N(\mu^\eta_{3,t}, \sigma^\eta_3) & \text{with prob. } p^\eta_3, \end{cases}$$

subject to $p^\eta_1 + p^\eta_2 + p^\eta_3 = 1$. The means change over time as driven by the latent variable x_t such that

$$\begin{aligned} \mu^\eta_{1,t} &= \bar{\mu}^\eta_t, \\ \mu^\eta_{2,t} &= \bar{\mu}^\eta_t + \mu^\eta_2 - x_t, \\ \mu^\eta_{3,t} &= \bar{\mu}^\eta_t + \mu^\eta_3. \end{aligned}$$

The parameter $\bar{\mu}^\eta_t$ is a normalization such that $\mathbb{E}_i[exp\{\eta_{i,t}\}] = 1$ in all periods. This normalization implies that x_t has no effect on mean income. When we estimate the income process, we impose restrictions on μ^η_2 and μ^η_3 so that the first component of the mixture will affect the center of the distribution, and the second and third components will affect the left and right tails, respectively. Following Catherine (2021), we posit that $x_t = \beta \Delta w_t$, where β is a parameter that controls the extent of cyclical variation in income risk.

The choice of a mixture of normal distributions balances richness and flexibility with computational simplicity. In terms of flexibility, any density that satisfies mild regularity conditions can be approximated by mixing a sufficient number of normals.[6] Normal mixtures allow for persistent income shocks to come from a rich distribution with a shape that can vary with the business cycle. And while the additional richness comes at the cost of estimating additional parameters, it can be incorporated into dynamic programming problems with only minor extensions to the computational methods used to integrate and simulate standard Gaussian distributions. Within the class of normal mixture models, there are a number of parameters that could vary with the business cycle. For example, the mixture probabilities or distribution variances could be cyclical, though we have found that neither option works well to match the procyclical volatility in skewness while also maintaining an acyclical variance. In contrast, shifting the means in the manner above is able to generate both patterns observed in the data.

As for the factor structure, we allow the effect of a recession or expansion on an individual's income to depend on his or her position in the persistent component of the income distribution, given by $q \equiv \gamma_i + z_{i,t}$. We model this exposure with the piecewise-linear function $f(q)$,

$$f(q) = \begin{cases} \alpha_1 q & \text{if } q < \bar{q} \\ \alpha_2 (q - \bar{q}) + \alpha_1 \bar{q} & \text{if } q \geq \bar{q}, \end{cases}$$

where $\bar{q}$ is a kink point, α_1 captures the slope of the factor structure below $\bar{q}$, and α_2 captures the slope above $\bar{q}$. The piecewise-linear specification allows flexibility to capture the non-monotonic (V-shape or inverse V-shape) factor structure seen in Figure 4, which we will see again in Figure 10 below.

We estimate each of these specifications while imposing that $\rho = 1$. This restriction is without much consequence if lifecycle income profiles are assumed to be homogeneous (κ_i =0), as ρ would be estimated to be close to 1 even without this restriction, because

[6] See Ferguson (1973) for the classic theorem, which has been generalized in many directions. See, for example, Bacharoglu (2010), who requires only continuity and compact valuedness of the real-valued density, or Frühwirth-Schnatter (2006) for a textbook treatment with alternative conditions.

the lifecycle profile of cross-sectional variance of income within a cohort is nearly linear in age. The restriction that ρ is identical to 1 is convenient for applications with homothetic preferences because it allows for a normalization that can further reduce the number of state variables, as we explain in Section 6.

We also consider versions of several specifications with HIP, in which the linear relationship between cross-sectional inequality and age is consistent with values of $\rho < 1$. Because heterogeneity in κ_i contributes to accelerating dispersion in income within a cohort over time, the linear lifecycle profile of the cross-sectional variance implies less persistent income shocks.

4 Estimation

We fit the model using a simulated method of moments estimation procedure. For a given parameter vector, we simulate a panel with 360,000 individuals per year. While our empirical moments begin only in 1978, we begin the simulation in 1947 to provide for a pre-sample burn-in period, as some of our moments refer to the cross-sectional distribution of income. Our income process does not have age-dependent parameters, but nevertheless we impose a "life cycle" in the simulation by simulating each individual for 36 years starting with $z_{i,t} = 0$.[7] There are two reasons for this life cycle. First, some of our specifications involve random walk shocks to z, and the life cycle structure keeps the cross-sectional distribution of z stationary. Second, one of our data targets is the growth of the cross-sectional variance of income over the life cycle. We assume a uniform age distribution, so there are 10,000 individuals in each of the 36 age groups from ages 25 to 55.

The aggregate component w_t varies over time, with consequences for the innovation distribution and therefore for the income distribution going forward. We normalize w to zero at the beginning of our simulation and then construct a time series by accumulating the demeaned time series for average income growth (one-year changes) reported by Guvenen et al. (2014).[8]

We target several types of moments. First, we target the shape of the income growth

[7] While we assume a deterministic length of life, applications can use a perpetual youth demographic structure to avoid keeping track of age as a state variable.

[8] In our pre-sample period, we use the growth rate of real wages and salary compensation per worker constructed from FRED series A4102C1Q027SBEA, CPIAUCSL, and PAYEMS.

distribution using the 10th, 50th, and 90th percentiles of the distributions for one-year, three-year, and five-year changes. We average these moments across all years, 1979 to 2011, for a total of nine moments. We also target a kurtosis of one-year income growth of 20 and a kurtosis of five-year income growth of 12 (two moments). We target the cross-sectional variance of income at ages 25, 35, 45, and 55 (four moments).

Second, we use two sets of moments to target the tails of the distribution. We target the masses of one-year changes in log income above 1.2 and below –1.2 (two moments). We also target the asymmetric slopes in the tails using an indirect inference procedure. We estimate the density of one-year income growth using a histogram. We then compute the slope of the tails by fitting two lines through the log density on the domains $[-4.0, -1.2]$ and $[1.2, 4.0]$. Figure 2 shows the lines we seek to match. We treat the two slopes of these lines as targets.

Third, to capture cyclical risk, we use the full time series of Kelley's skewness for one-year, three-year, and five-year income changes (93 moments in total). Finally, we target the factor structure of business cycle incidence. For each business cycle episode, Guvenen et al. (2014) construct a figure analogous to Figure 4, using average income over the five years before the business cycle episode to rank individuals in the distribution.[9] For each figure, we fit a line between the 11th and 80th percentiles and another line between the 80th and 100th percentiles. The two slopes of these lines are targets, totaling 14 moments across seven recession and expansion episodes.

The empirical moments are taken from (or in some cases constructed from) data reported by Guvenen et al. (2014) and Guvenen et al. (2021), and the target values are reported in Table III. We compute the squared percentage difference between the simulated and data moments.[10] In general, the moments are weighted equally, but with a few exceptions; notably, the 93 skewness moments are down-weighted to put them collectively on a more equal footing with the other moments. Appendix A provides more details about the construction of the moments and the objective function. We minimize the objective function using the TikTak global optimization algorithm discussed in Arnoud et al. (2019).

[9] We use the data for the recessions of 1979-1983, 1990-1992, 2000-2002, and 2007-2010 and the expansions of 1983-1990, 1992-2000, and 2002-2007.

[10] In a few cases, the data moments take values near zero and those are treated differently. See Appendix A for more details.

Table I – Model Summary

Model	Key Components of Stochastic Process					Factor Str.
	ζ	η	ψ	ρ	σ^{κ}	
(1)	Gaussian	Gaussian	$= 0$	$= 1$	$= 0$	
(2)	Non-emp.	Gaussian	> 0	$= 1$	$= 0$	
(3)	Non-emp.	Mixture	> 0	$= 1$	$= 0$	
(4)	Non-emp.	Mixture	> 0	$= 1$	$= 0$	✓
(5)	Non-emp.	Mixture	> 0	≤ 1	> 0	
(6)	Non-emp.	Mixture	> 0	≤ 1	> 0	✓

Notes: We estimate six models that each maintain different specifications for six different aspects of the model. The transient income innovation, ζ, is either Gaussian distributed or a nonemployment shock. The persistent income innovation, η, is distributed by either a simple Gaussian or a normal mixture distribution. The income scarring parameter, ψ, is either assumed to be zero or estimated as a positive number. The persistence of the AR(1) process, ρ, is either assumed to be 1 or estimated as a number less than or equal to one. Wherever ρ is estimated, we calibrate a HIP process with $\sigma^k > 0$, and wherever ρ is restricted to one, we assume homogeneous income profiles with $\sigma^k = 0$. Finally, we estimate models with and without a factor structure in exposure to the aggregate component of income.

5 Results

We estimate six models of increasing complexity, beginning with the canonical linear-Gaussian model and gradually adding components. In the following subsections, we discuss each model in turn. In Table I, we summarize the features of the six models. Table II presents the parameter estimates for each model. For the most part, parameter values are stable with the addition of more complexity to other aspects of the model. In Table III, we present the targeted and predicted moments, with the exception of the skewness time series and the factor structure. The lower panel of Table III shows the objective function value as well as the contributions from each set of moments. Figures 5 to 11 show key moments for each model.

5.1 Model 1: Canonical (Permanent-Plus-Transitory) Gaussian Model

Under Model 1, log income is normally distributed, as both η and ζ follow Gaussian distributions. This specification does a reasonable job of matching the growth of the cross-sectional variance of income over the life cycle but cannot generate any of the excess

Table II – Estimated Parameters

Parameters		Model Specifications					
		(1)	(2)	(3)	(4)	(5)	(6)
σ_γ	St. dev. of fixed effects	0.776	0.646	0.597	0.604	0.630	0.623
σ_1^ζ	St. dev. for transitory shock	0.196	–	–	–	–	–
p^ζ	1 – Probability of nonempl. shock	–	0.550	0.611	0.618	0.704	0.704
λ	Transitory exponential parameter	–	3.357	3.081	3.054	2.786	2.857
ψ	Scarring effect of transitory shock	–	0.094	0.151	0.166	0.377	0.351
p_2^η	Mix. probab. for persist. innov 2	–	–	0.109	0.084	0.119	0.092
p_3^η	Mix. probab. for persist. innov 3	–	–	0.062	0.018	0.074	0.049
σ_1^η	Std. dev. for persistent innov. 1	0.086	0.103	0.010	0.030	0.056	0.049
σ_2^η	Std. dev. for persistent innov. 2	–	–	0.164	0.190	0.404	0.449
σ_3^η	Std. dev. for persistent innov. 3	–	–	0.192	0.195	0.415	0.225
μ_1^η	Center for persistent component 1	–	–	–0.012	–0.034	0.006	0.024
μ_2^η	Center for persistent component 2	–	–	0.034	0.086	0.086	0.050
μ_3^η	Center for persistent component 3	–	–	0.181	0.304	0.254	0.498
β	Loading on aggregate wage	–	–	–7.948	–7.639	–8.839	–9.954
α_1	Factor struct. slope, low income	–	–	–	–0.678	–	–0.853
α_2	Factor struct. slope, high income	–	–	–	0.931	–	0.911
$\overline{q}$	Factor structure threshold	–	–	–	0.745	–	0.594
ρ	AR(1) coefficient	–	–	–	–	0.810	0.822

Note: This table contains the estimated parameters for each model as specified in Table I. In Model 1, the transitory ζ shock is realized every period and has a Gaussian distribution.

kurtosis of income growth observed in the data. Relatedly, it vastly under-predicts the mass in the tails of the income growth distribution, as can be seen from the histogram shown in panel (a) of Figure 5. This model predicts so few simulated individuals in the tails that we are unable to reliably compute the slopes of those tails. Panel (a) of Figure 8 shows this model predicts a very low standard deviation of income growth. In principle, this model could fit that moment, but with a Gaussian framework, there is a tension between matching the cross-sectional variance of income levels and the dispersion of income growth rates. A similar tension between estimates in levels and differences has been observed by Heathcote et al. (2010). By construction, this model does not generate

Table III – Targeted and Fitted Moments

Moments	US Data	Model Specifications (1)	(2)	(3)	(4)	(5)	(6)
P10, one-year change	–0.434	–0.361	–0.448	–0.435	–0.426	–0.435	–0.419
P10, three-year change	–0.585	–0.388	–0.486	–0.495	–0.485	–0.629	–0.604
P10, five-year change	–0.631	–0.403	–0.517	–0.529	–0.52	–0.713	–0.681
P50, one-year change	0.020	0.002	0.015	0.009	0.014	0.018	0.018
P50, three-year change	0.061	0.007	0.020	0.012	0.016	0.016	0.019
P50, five-year change	0.103	0.023	0.035	0.021	0.025	0.027	0.032
P90, one-year change	0.474	0.365	0.435	0.434	0.432	0.465	0.451
P90, three-year change	0.705	0.402	0.487	0.513	0.518	0.663	0.635
P90, five-year change	0.848	0.449	0.552	0.591	0.599	0.773	0.746
Kurtosis, one-year change	20.00	3.000	22.73	23.28	23.88	22.00	23.18
Kurtosis, five-year change	12.00	3.000	17.78	16.83	16.97	10.81	11.66
Cross-sectional var., age 25	0.595	0.649	0.592	0.535	0.556	0.592	0.587
Cross-sectional var., age 35	0.719	0.723	0.711	0.692	0.717	0.746	0.736
Cross-sectional var., age 45	0.814	0.795	0.832	0.841	0.853	0.812	0.804
Cross-sectional var., age 55	0.905	0.872	0.953	0.974	0.991	0.922	0.917
Left-tail mass	0.024	0.00	0.024	0.024	0.024	0.023	0.022
Left-tail slope	1.260	—	1.282	1.282	1.294	1.246	1.284
Right-tail mass	0.015	0.00	0.021	0.018	0.018	0.016	0.015
Right-tail slope	–2.035	—	–1.482	–1.561	–1.551	–1.935	–1.850
Objective value		10.505	2.439	1.592	1.195	0.988	0.564
Quantiles		0.748	0.295	0.238	0.236	0.045	0.046
Kurtosis		1.286	0.251	0.189	0.210	0.020	0.026
Cross-sectional var. profile		0.100	0.035	0.187	0.157	0.017	0.011
Histogram		7.175	0.525	0.341	0.347	0.021	0.052
Skewness time series		0.699	0.833	0.143	0.143	0.185	0.136
Factor Structure		0.498	0.498	0.496	0.101	0.699	0.293

Notes: This table shows the model fit for each estimated model. The top panel displays all the individual targeted moments with the exception of the time series for the Kelley skewness of one-year and five-year income growth and the factor structure moments. The first column contains the targeted moments computed from SSA data (Guvenen et al. (2014, 2021)), and subsequent columns show the implied values from the estimated models. The bottom panel shows the weighted contribution of selected sets of moments to the objective function. The top row of the bottom panel shows the total value of the objective function, including factor structure moments.

Figure 5 – Log Density of Annual Income Growth

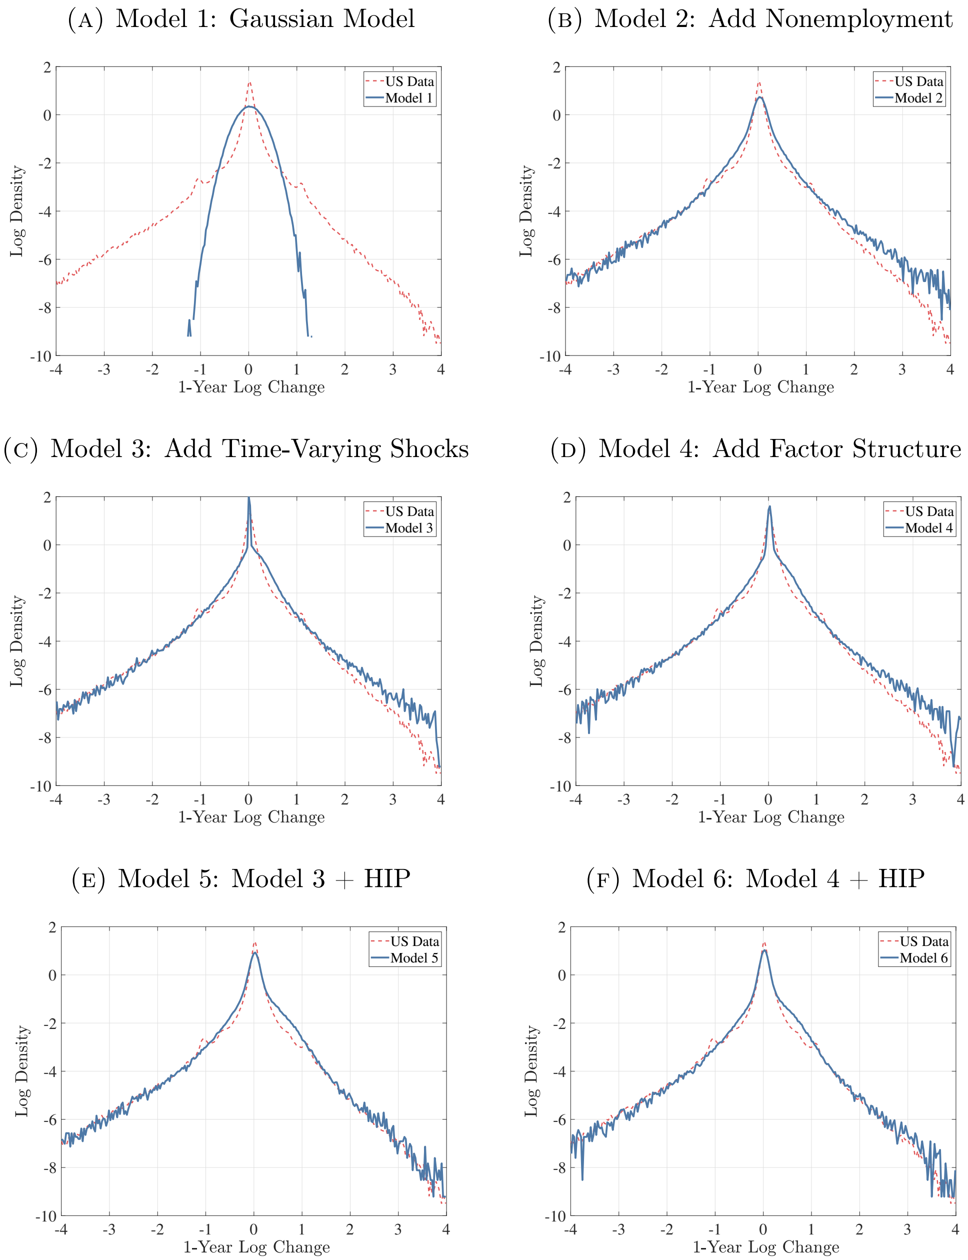

Note: This figure shows the fit of each model as specified in Table I to the log density of one-year income growth from 1995 to 1996. The dashed line is the empirical distribution (Guvenen et al., 2021), and the solid line is the log density of income growth for 360,000 simulated individuals in each model.

any time series variation in the standard deviation or skewness of income growth, nor does it generate a factor structure in business cycle incidence.

5.2 Model 2: Adding Nonemployment Shocks

Model 2 changes the distribution of the transitory shock to nonemployment shocks with scarring effects. The estimates in column 2 of Table II imply that an individual receives a nonemployment shock, ℓ_{it}, with an annual probability of 45% and mean $\lambda^{-1} \approx 0.30$. The passthrough coefficient, ψ, which determines the fraction of ζ that leaves a permanent effect (through equation 3) is about 9.4%. Although this may seem like a small number, keep in mind that $log(1 - \ell_{it})$ can get very negative given the exponential distribution of ℓ_{it}, which has a very long tail, and the log transformation. For example, each year, 8.6% of individuals see their income fall by 50% or more because of the nonemployment shock alone, and 1.8% see their income fall to effectively zero, corresponding to full year nonemployment.

As a result, Model 2 is able to generate thick and long tails and a more peaked center of the income growth distribution, which provides a much closer match to the empirical density than Model 1—compare Figure 5(b) to 5(a). This is also reflected in the kurtosis values of one- and five-year income growth, which rise from the Gaussian benchmark of 3 in Model 1 to 22.7 and 17.8 in Model 2 (Table III), somewhat exceeding their empirical counterparts (of 20 and 12, respectively). While the model fits the left-tail slope almost exactly, it overstates the thickness of the right tail (with a slope of −1.48 versus −2.04 in the data), leading to the higher kurtosis. Income scarring is helpful in making the right tail steeper than the left tail, but the estimation cannot push this mechanism too far without generating too much dispersion in persistent income changes. Below, we will show that allowing for HIP relaxes this tension by making the persistent component of income risk less persistent ($\rho < 1$).

The lower panel of Table III shows a substantial improvement in the objective function from Model 1 to Model 2 that comes from better fitting the histogram, the kurtosis, and the quantiles. Like Model 1, this one does not generate any business cycle dynamics in the higher moments of the distribution of income growth, but it is better able to match the *levels* of the standard deviations of income growth (panel (b) of Figures 8 and 9).[11]

[11] We also estimated a version of Model 2 that features the nonemployment shock but without the scarring component ($\psi \equiv 0$ in equation 3). This model does not perform as well, with the main deterioration in fit coming from the histogram, kurtosis, and the lifecycle variation in inequality. The estimates are reported in Tables IV and V in the appendix.

Figure 6 – Skewness of Annual Income Growth

(a) Model 1: Gaussian Model

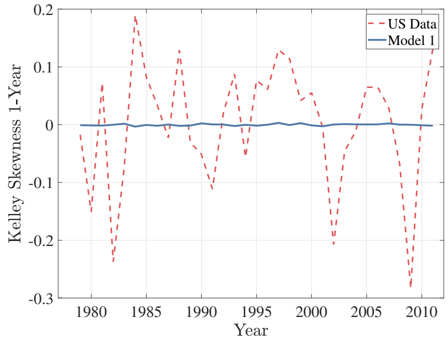

(b) Model 2: Add Nonemployment

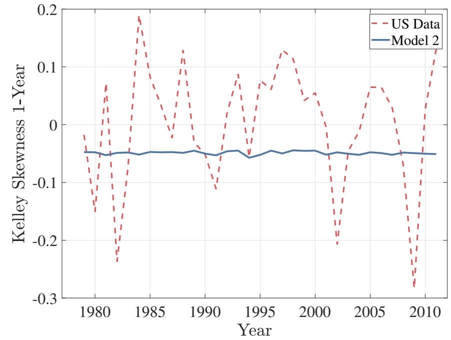

(c) Model 3: Add Time-Varying Shocks

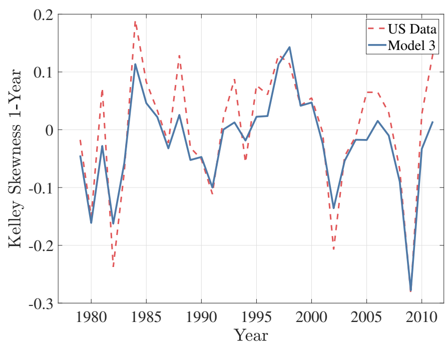

(d) Model 4: Add Factor Structure

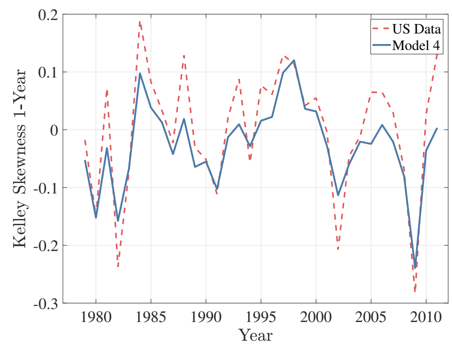

(e) Model 5: Model 3 + HIP

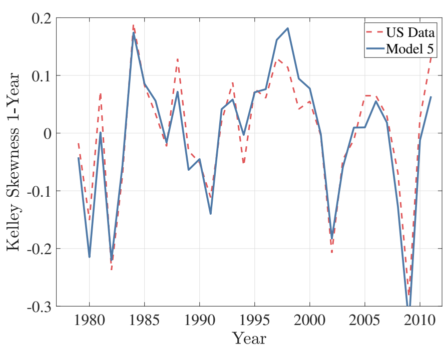

(f) Model 6: Model 4 + HIP

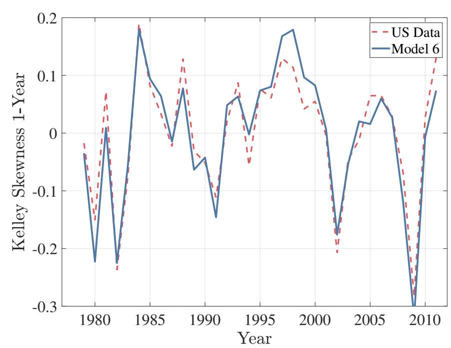

Note: This figure shows the fit to the time series of the Kelley skewness of one-year income growth for each model, as specified in Table I. The dashed line represents the Kelley skewness of the empirical distribution of income growth with respect to one year prior (Guvenen et al. (2014)), and the solid line is the simulated time series of the Kelley skewness of one-year income growth in each model.

Figure 7 – Skewness of Five-year Income Growth

(a) Model 1: Gaussian Model

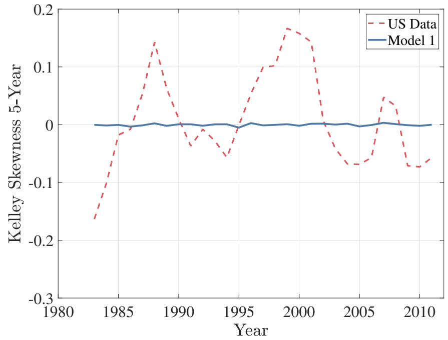

(b) Model 2: Add Nonemployment

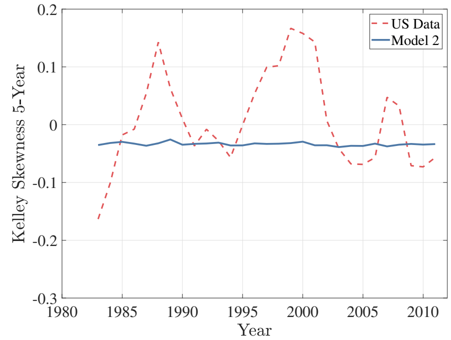

(c) Model 3: Add Time-Varying Shocks

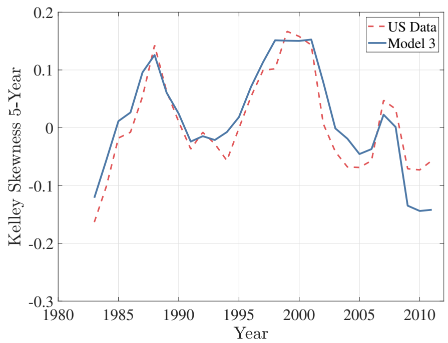

(d) Model 4: Add Factor Structure

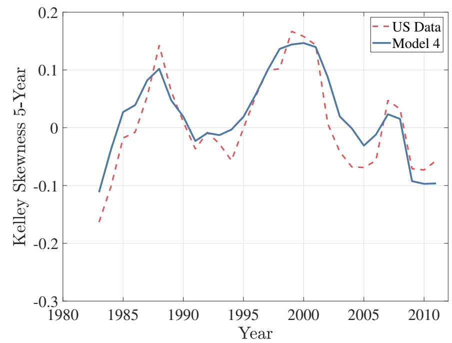

(e) Model 5: Model 3 + HIP

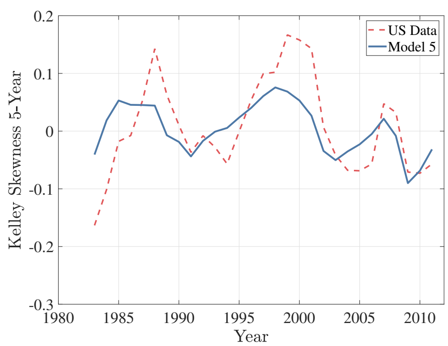

(f) Model 6: Model 4 + HIP

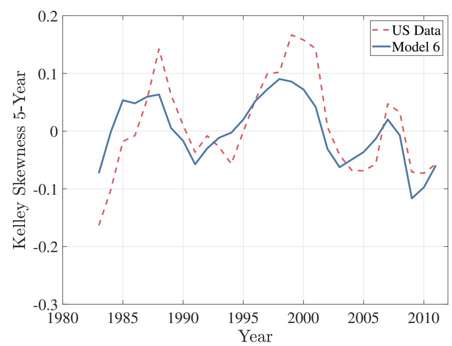

Note: This figure shows the fit to the time series of the Kelley skewness of five-year income growth for each model, as specified in Table I. The dashed line represents the Kelley skewness of the empirical distribution of income growth with respect to 5 years prior (Guvenen et al. (2014)), and the solid line is the simulated time series of the Kelley skewness of five-year income growth in each model.

Variances of Log Income and Change in Log Income: Resolving a Puzzle

A well-known puzzle in the income dynamics literature is that the persistent-plus-transitory Gaussian model cannot simultaneously fit the variances (and covariances) of log income and log income changes (Heathcote et al. (2010)). This can be seen in our results for Model 1 above: while the model does a reasonably good job of fitting the variances of log income *levels* at ages 25, 35, 45, and 55 in Table III, it understates the variances of one- and five-year log income *changes* by 60% to 80% (obtained by squaring the standard deviation lines in Figures 8(a) and 9(a)).[12] The introduction of the non-employment shock in Model 2 goes a long way towards resolving this puzzle: the model improves the fit to the variances of log income levels relative to Model 1 (Table III), while now matching the variance of one-year log income change exactly and understating that of five-year change by only 30%.

5.3 Model 3: Introducing Normal Mixture with Time Variation

Model 3 changes the distribution of the persistent shock to a time-varying mixture of normal distributions. The estimated distribution features one shock, η_1, that is realized very often (with about 83% probability) and two other shocks that are realized with 10.9% and 6.2% probability. The frequent shock is very small, with about 1% standard deviation, whereas the other two shocks are much larger, with 16.4% and 19.2% standard deviations. These estimates are consistent with the plausible idea that in most years, persistent income changes are small, with large jumps happening less frequently. The estimates of the scarring effect component are not affected greatly by the introduction of the normal mixture, with the passthrough coefficient rising somewhat (from 9.4% to 15.1%) and the shock frequency declining slightly (from 45% to 39%).

The main improvement that this model offers is to generate procyclical skewness and acyclical variance of income growth rates. Figures 6(c) and 7(c) show the model is able to closely match the dynamics of Kelley's skewness for both one-year and five-year income growth. The model is able to generate an acyclical standard deviation for both one-year and five-year income growth rates (Figures 8(c) and 9(c)). The levels of standard deviation remain virtually the same as those in Model 2, with the one-year matching the data exactly and five-year understating its empirical counterpart by 10 log points

[12] Alternative explanations for this puzzle have been proposed based on the treatment of job-to-job transitions (Daly et al., 2022) and time aggregation (Crawley et al., 2022).

FIGURE 8 – Standard Deviation of Annual Earnings Growth

(A) Model 1: Gaussian Model

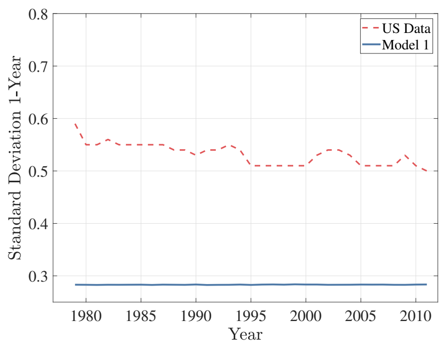

(B) Model 2: Add Nonemployment

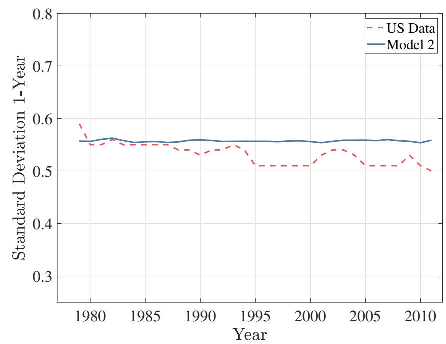

(C) Model 3: Add Time-Varying Shocks

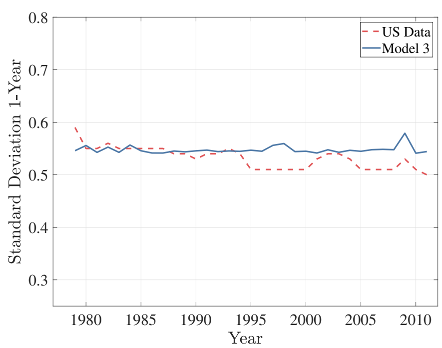

(D) Model 4: Add Factor Structure

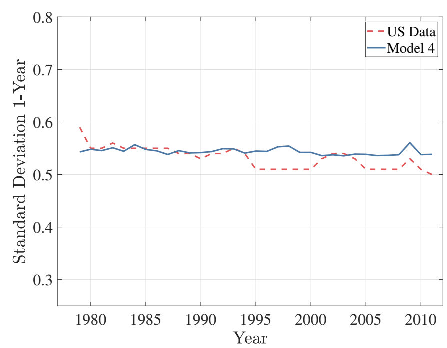

(E) Model 5: Model 3 + HIP

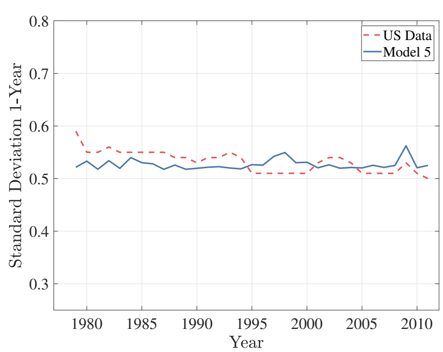

(F) Model 6: Model 4 + HIP

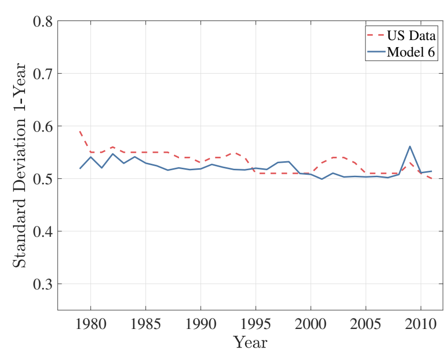

Note: This figure shows the fit to the time series of the standard deviation of one-year income growth for each model, as specified in Table I. The dashed line represents the standard deviation of the empirical distribution of income growth with respect to one year prior (Guvenen et al. (2014)), and the solid line is the simulated time series of the standard deviation of one-year income growth in each model.

Figure 9 – Standard Deviation of Five-year Income Growth

(a) Model 1: Gaussian Model

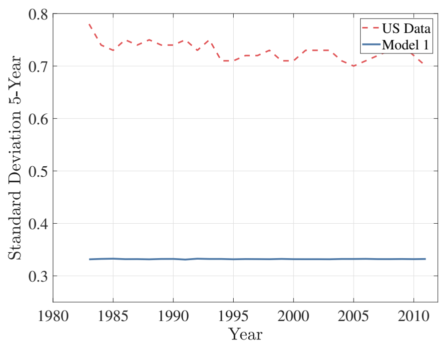

(b) Model 2: Add Nonemployment

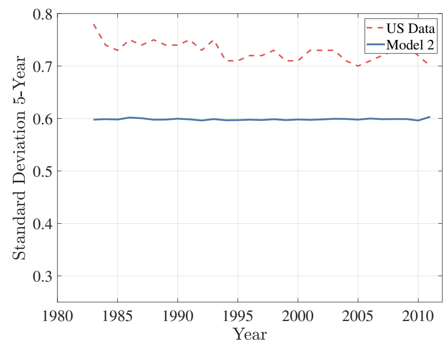

(c) Model 3: Add Time-Varying Shocks

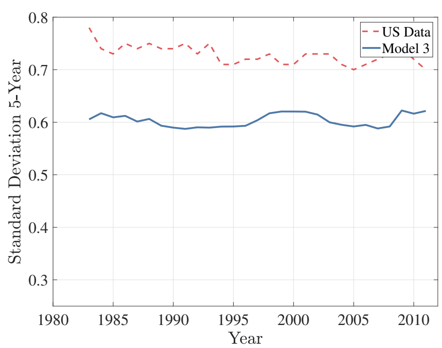

(d) Model 4: Add Factor Structure

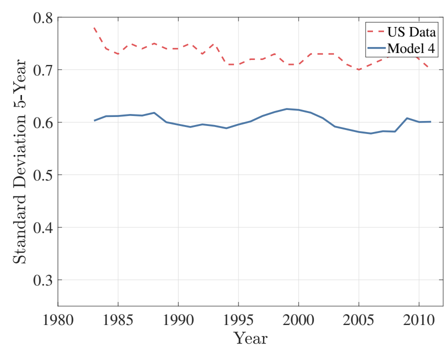

(e) Model 5: Model 3 + HIP

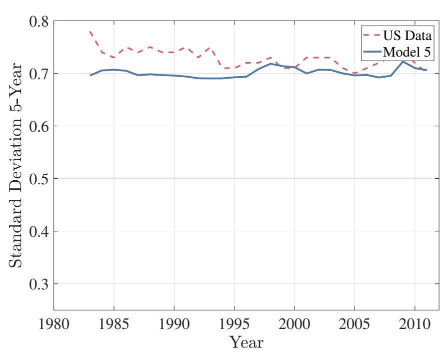

(f) Model 6: Model 4 + HIP

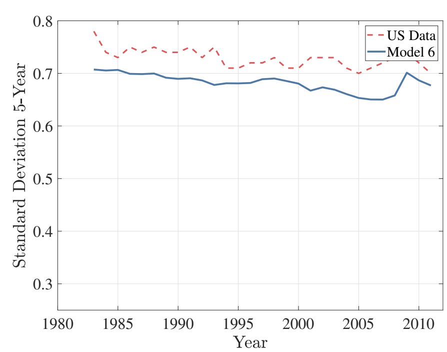

Note: This figure shows the fit to the time series of the standard deviation of five-year income growth for each model as specified in Table I. The dashed line represents the standard deviation of the empirical distribution of income growth with respect to 5 years prior (Guvenen et al. (2014)), and the solid line is the simulated time series of the standard deviation of five-year income growth in each model.

(about 0.6 versus 0.7). In addition to matching the procyclical skewness and acyclical variance of income growth, this model also leads to a slight improvement in the shape of the distribution as measured by the quantile, kurtosis, and histogram components of the objective function.[13]

5.4 Model 4: Adding the Factor Structure

Model 4 adds a factor structure for business cycle incidence. Compared with those in the middle of the income distribution, individuals with either low or high levels of (the persistent component of) income are more exposed to cyclical fluctuations in income. Figure 10 shows the fit of Model 4 for each of the seven business cycle episodes. The top row of the figure shows recessions, and the bottom row shows expansions. In deep recessions (1979–1983 and 2007–2010), there is a clear upward slope through the middle of the income distribution. This indicates that lower-income individuals experienced larger income losses. The model successfully captures this pattern. The slope is less evident in the milder recessions (1990–1992 and 2000–2002), which were also shorter, and the model also captures this. The difference between these episodes is captured by the change in the aggregate component of income, w_t, which falls more in deep recessions, making the uneven incidence more important. At the top end of the income distribution, recessions have different characters. In the 2000–2002 and 2007–2010 recessions, there is a large fall in income at the top of the distribution. The model is able to match this in the Great Recession but not in the 2000–2002 recession. Again, this difference in predictions is due to the larger movement in w_t in the Great Recession.[14] In the 1979–1983 recession, the data do not show a decline in top incomes despite the large drop in w_t.

The pre-2000 expansions show income growth at the bottom and top of the income distribution, and the model is able to match these patterns fairly well. The 2002-2007 expansion does not show a factor structure for the middle of the income distribution but does show strong growth at the top. During this episode, average income grew slowly (w_t rises by only 10 log points), and as a result, the unequal incidence in the model is of

[13] In order to show the extent to which the improvement in fit over Model 2 comes from the time-varying moments, we estimate a specification with a static normal mixture distribution. As seen in Appendix Tables IV and V, this process (Model 3') still leads to a modest improvement in the fit over Model 2—with the objective value falling from 2.44 to 2.26. The main improvements come from a better fit of the tails of the distribution and the average skewness.

[14] Adding other aggregate factors, e.g. stock indices, that may affect very high incomes can help to better fit the cyclical income risk at the top of the distribution.

little importance.

Overall, the model is able to capture the business cycle factor structure and account for some, though not all, of the differences across business cycle episodes. Adding the factor structure does not detract from the model's ability to fit other moments, as can be seen in the lower panel of Table III, which shows that all the other components of the objective function are little changed relative to Model 3.

Factor structure without time-varying moments. Finally, in some applications, researchers may want to model the factor structure without the time variation in skewness. We have estimated a version of Model 4 that corresponds to this case, which shuts down time variation ($\beta \equiv 0$) but keeps the factor structure. Although the fit is slightly worse than Model 4's, even ignoring the moments with time variation, the fit to the factor structure remains remains largely intact. Because this specification may be useful to researchers, we include the parameter estimates in Appendix B (Model 4' in Table IV).

5.5 Model 5: Adding HIP

Persistence in the income process leads to a tension between matching the roughly linear increase in cross-sectional variance over the life cycle and the broader empirical distribution of income changes. In particular, the models that assume $\rho = 1$ have a difficult time matching the right tail of the distribution of annual income growth (Figure 5). This tension is also reflected in the fit of the standard deviation and kurtosis of five-year changes.

Model 5 adds HIP to Model 3 (i.e. which does not feature the factor structure). We set $\sigma^{\kappa} = 0.015$, based on the empirical estimates in Baker (1997) and Guvenen et al. (2021), among others. We relax the restriction of $\rho = 1$ and estimate ρ to be 0.80. Despite the low value of ρ, the model continues to fit the nearly linear growth in the cross-sectional variance of income, as shown in Table III. This is possible because HIP gives rise to a convex variance profile that offsets the concave contribution from the persistent income shocks.

Comparing panels (c) and (e) of Figure 5, we can see this model better captures the shape of the income distribution. In particular, the right-tail slope is much closer to the data. In Table III, we see the fit of P10 and P90 for the three-year and five-year changes is also much improved. Comparing panels (c) and (e) of Figure 6, we see the fit to the time series of Kelley's skewness for one-year income growth improves as well, but the fit for five-year changes (shown in Figure 7(e)) is somewhat worse, as the amplitude of the fluctuations in skewness is somewhat attenuated. On the other hand, the level of the standard deviation of five-year changes rises to match the data more closely (see Figure 9 (e)).

5.6 Model 6: HIP with a Factor Structure

Our last specification introduces the factor structure alongside HIP components. The estimate of ρ is nearly unchanged at 0.794. The improvements to the model fit that HIP brings remain mostly intact (see Table III and Figures 5 through 9). However, the fit to the factor structure is worse compared with that of Model 4 (see Figure 11). In particular, the model struggles to fit the greater exposure of low-income individuals to recessions. With $\rho < 1$, there is a tendency for mean reversion by which low-income individuals experience faster income growth. During a recession, the factor structure must overcome this force. On the other hand, the faster growth for low-income individuals in expansions is naturally generated by mean reversion without a strong factor structure in business cycle incidence, implying a weaker factor structure. The lower panel of Table III shows that the objective function contribution from the factor structure is larger for this specification than for Model 4. Nevertheless, the other components of the objective function improve, and overall this model gives a substantially better fit.

FIGURE 10 – Growth in Log Average Income by Previous Income Percentile (Model 4)

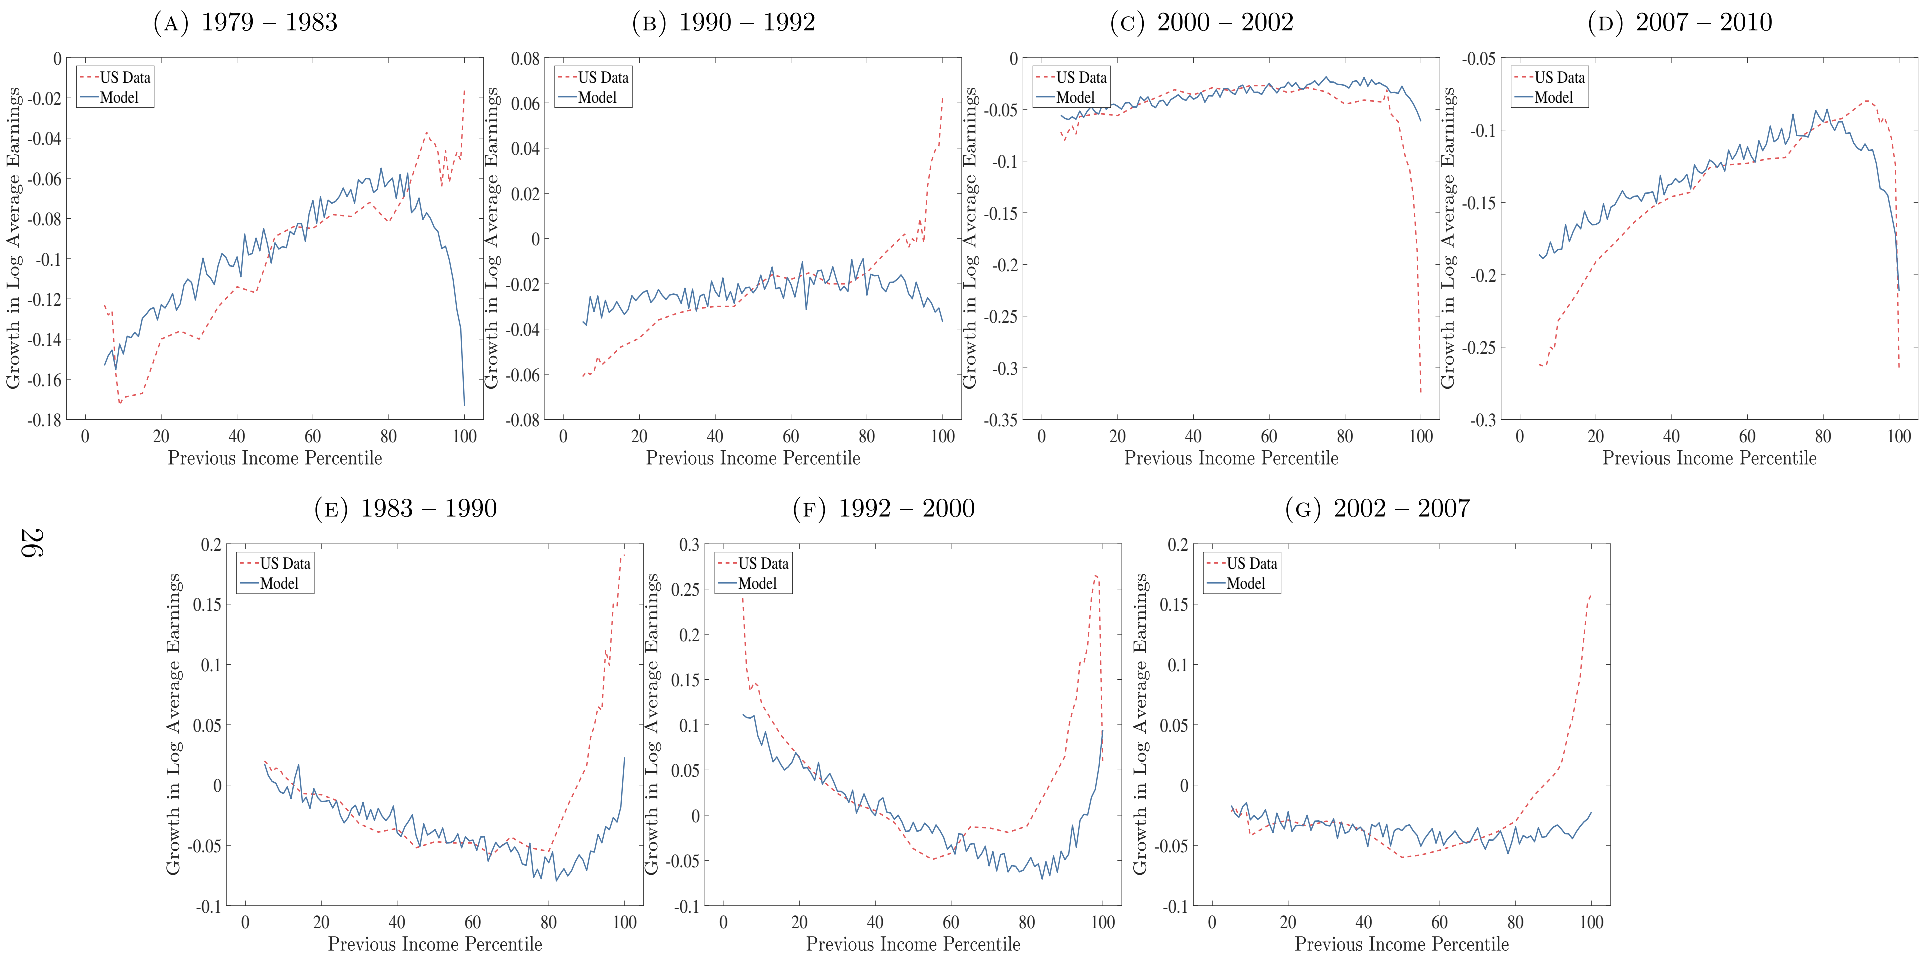

Note: This figure shows the model fit to the factor structure for Model 4 as specified in Table I. The dashed line is the growth in log income over each expansion or recession period for individuals ranked by the income percentile for the 5 years prior to the beginning of the period (Guvenen et al. (2014)). The solid line is the analogous income growth for the individuals simulated by the model and ranked according to income percentile. The figures show income growth net of predicted age effects, which explains the low level of average income especially in expansions, which tend to be longer episodes.

Figure 11 – Growth in Log Average Income by Previous Income Percentile (Model 6)

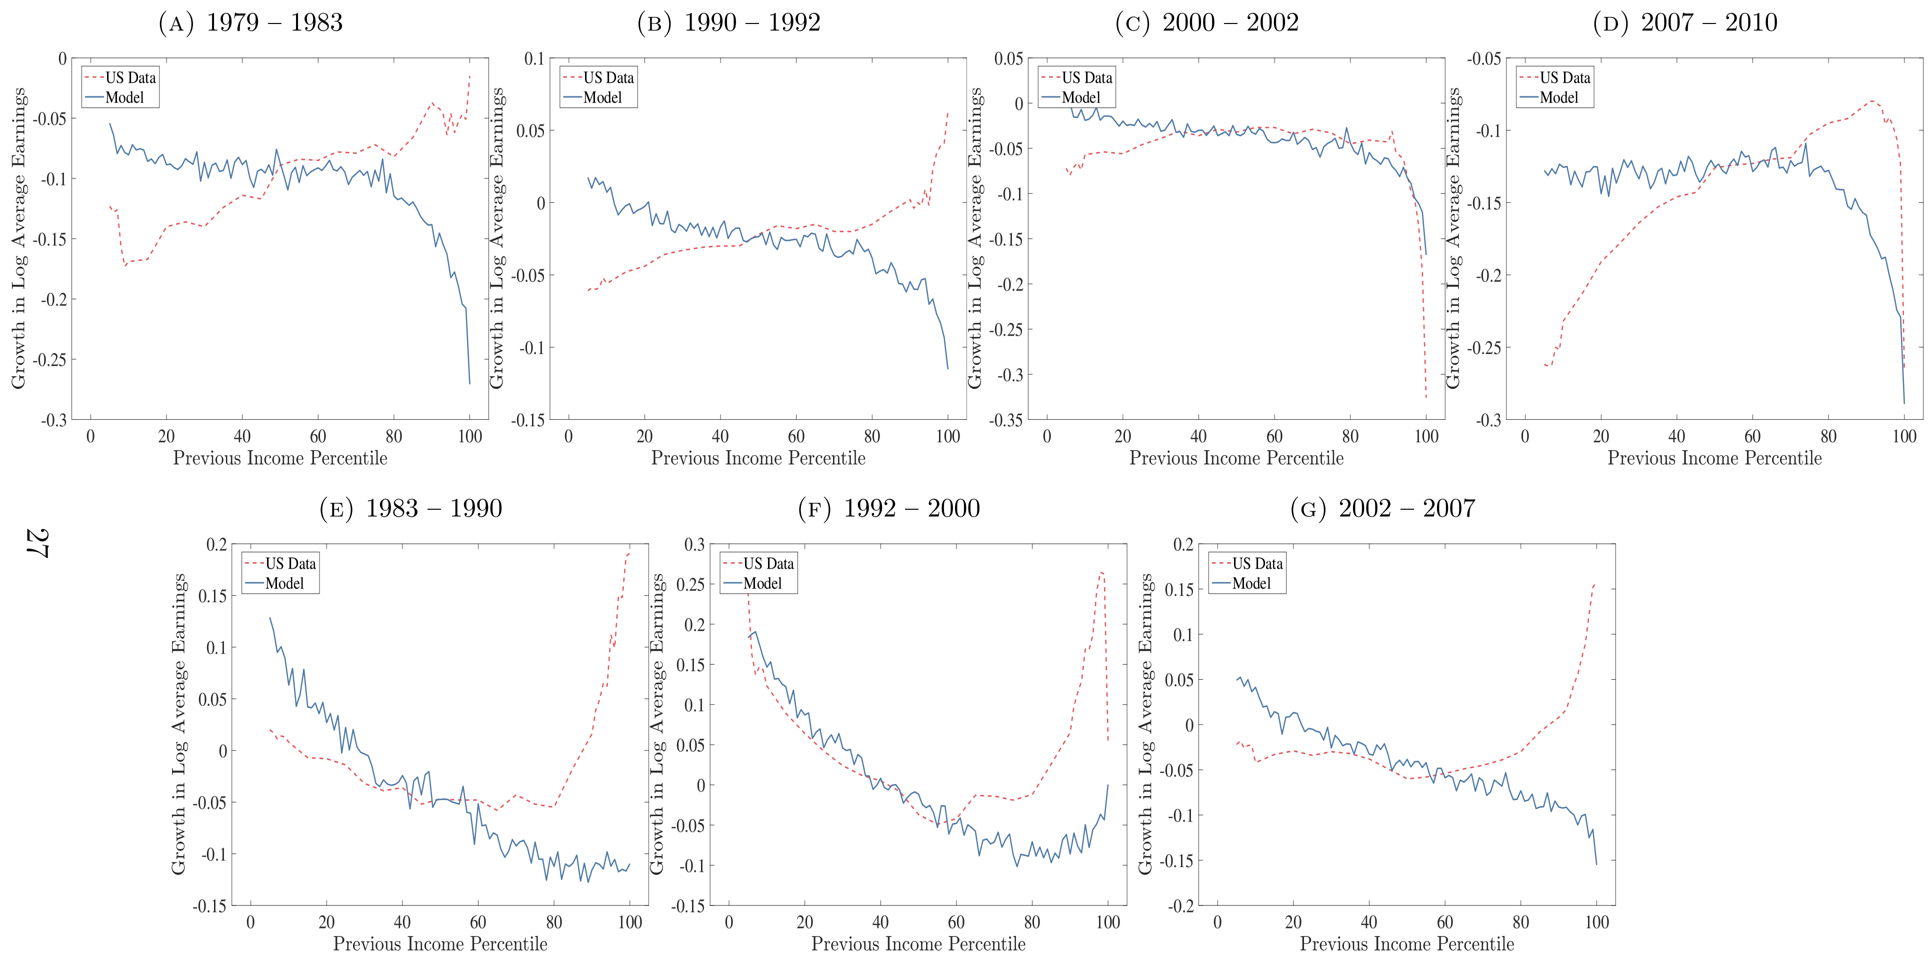

Note: This figure shows the model fit to the factor structure of business cycles for Model 6, as specified in Table I. The dashed line is the growth in log income over each expansion or recession period for individuals ranked by the income percentile for the 5 years before the beginning of the period (Guvenen et al. (2014)). The solid line is the analogous income growth for the individuals simulated by the model and ranked according to income percentile. The figures show income growth net of predicted age effects, which explains the low level of average income especially in expansions, which tend to be longer episodes.

6 Suggestions to Modelers

The income process we have presented is quite rich, and not all of its elements may be relevant for a given application. We recommend Model 5 as a useful benchmark if the factor structure of the business cycle is not a focus of the analysis. If the factor structure is of special importance, we recommend Model 4, as it fits the factor structure better than Model 6. If parsimony is of particular importance, Model 3 affords a normalization that reduces the number of state variables in models with homothetic preferences (see below).

We now give some guidance on how one can incorporate this income process into the dynamic programming problem that is at the core of many heterogeneous-agent business cycle models. The Bellman equation for such a problem could take the form

$$V(m, z, S; \gamma, \kappa) = \max_{c,a'} \{u(c) + \beta E [V(R(S')a' + Y(z', \zeta', \gamma, \kappa, S), z', S'; \gamma, \kappa)]\}$$

subject to

$$\begin{aligned} c + a' &= m \\ a' &\geqslant \underline{a}, \end{aligned}$$

where m is cash on hand, $R(S)$ is a rate of return that may depend on the aggregate state S and $Y(z, \gamma, \kappa, S)$ gives the level of income. This income function is given by $\exp(y)$, where y is determined by Equation 1 with $w_t = w(S)$. The persistent component z evolves according to Equation 2. The parameter $\underline{a}$ is a borrowing limit.

From the structure of the problem, one can see that the income process affects the dynamic program in two places. First, we need to evaluate income, $Y(z, \zeta, \gamma, \kappa, S)$, for a given set of state variables. This just requires evaluating the expression in Equation 1, which is straightforward. The second place the income process affects the dynamic programming problem is in taking the expectation over z' and ζ'. A quadrature method is likely the most practical approach for this. To take the expectation, one can create quadrature nodes and weights for different outcomes of η' and ζ'. With the time-varying mixture of normals for the distribution of η', some of the quadrature nodes for η' will depend on S and S'. See McKay and Reis (2021) for an example of such a quadrature approach.

In a special case with $\rho = 1$, homothetic $u(\cdot)$, and no factor structure $f(z+\gamma) = 0$, the level of income does not affect the decision problem. In this case, one can greatly simplify the problem by normalizing all variables by $\exp(z+\gamma)$ and eliminating z as a state variable of the problem (see, e.g., Carroll et al. 2017).

7 Discussion

Recent analyses of income dynamics have highlighted several features of the data that are not captured by the income processes typically used in heterogeneous agent models of the business cycle. In particular, income growth rates show double-Pareto tails, high kurtosis, procyclical skewness, acyclical variance, and a factor structure in business cycle incidence. We have presented a model for an income process that can capture these aspects of the data while retaining the simplicity of a single state variable as in commonly used processes.

Each of the elements that we add to the income process has its own implications for the analysis of the business cycle. The kurtosis of income risk has been shown to be important to analyses of social insurance policies, the distribution of wealth, and marginal propensities to consume. Cyclical variation in income risk has been shown to be important to the welfare cost of business cycles and the value of stabilization policy. The incidence of the business cycle is relevant for cyclical variation in inequality.

In a particular application, one may find it useful to include only some elements of the income process or to modify the model in other ways. The set of moments we target and the estimation procedure we use can easily be applied to estimate alternative income processes.

References

Arellano, M., R. Blundell, and S. Bonhomme (2017): "Earnings and consumption dynamics: a nonlinear panel data framework," *Econometrica*, 85, 693–734.

Arnoud, A., F. Guvenen, and T. Kleineberg (2019): "Benchmarking Global Optimizers," Tech. rep., National Bureau of Economic Research.

Bacharoglu, A. (2010): "Approximation of Probability Distributions by Convex Mixtures of Gaussian Measures," *Proceedings of the American Mathematical Society*, 138, 2619–2628.

Baker, M. (1997): "Growth-Rate Heterogeneity and the Covariance Structure of Life-Cycle Earnings," *Journal of Labor Economics*, 15, 338–375.

Bell, B., N. Bloom, and J. Blundell (2021): "Income Dynamics in the United Kingdom 1975-2020," Working paper, Stanford University.

Bhandari, A., D. Evans, M. Golosov, and T. J. Sargent (2021): "Inequality, Business Cycles, and Monetary-Fiscal Policy," *Econometrica*, 89, 2559–2599.

Carroll, C., J. Slacalek, K. Tokuoka, and M. N. White (2017): "The distribution of wealth and the marginal propensity to consume," *Quantitative Economics*, 8, 977–1020.

Catherine, S. (2021): "Countercyclical Labor Income Risk and Portfolio Choices over the Life Cycle," *The Review of Financial Studies*, hhab136.

Constantinides, G. M. and D. Duffie (1996): "Asset Pricing with Heterogeneous Consumers," *The Journal of Political Economy*, 104, 219–240.

Constantinides, G. M. and A. Ghosh (2016): "Asset Pricing with Countercyclical Household Consumption Risk," *Journal of Finance*.

Crawley, E., M. B. Holm, and H. Tretvoll (2022): "A Parsimonious Model of Idiosyncratic Income," Finance and Economics Discussion Series 2022-026. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2022.026.

Daly, M., D. Hryshko, and I. Manovskii (2022): "Improving the Measurement of Earnings Dynamics," *International Economic Review*, 63, 95–124.

Deaton, A. and C. Paxson (1994): "Intertemporal Choice and Inequality," *Journal of Political Economy*, 102, 437–67.

Ferguson, T. (1973): "A Bayesian Analysis of Some Nonparametric Problems." *The Annals of Statistics*, 1, 209–230.

Frühwirth-Schnatter, S. (2006): *Finite Mixture and Markov Switching Models*, Springer.

Golosov, M., M. Troshkin, and A. Tsyvinski (2016): "Redistribution and Social Insurance," *American Economic Review*, 106, 359–386.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2015): "What Do Data on Millions of U.S. Workers Say About Labor Income Risk?" Working Paper 20913, National Bureau of Economic Research.

——— (2021): "What Do Data on Millions of U.S. Workers Reveal about Lifecycle Earnings Dynamics?" *Econometrica*, 89, 2303–2339.

Guvenen, F., S. Ozkan, and J. Song (2014): "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, 122, 621–660.

Guvenen, F., L. Pistaferri, and G. Violante (2022): "Global Trends in Income Inequality and Income Dynamics: New Insights from GRID," *Quantitative Economics*, Forthcoming.

Guvenen, F., S. Schulhofer-Wohl, J. Song, and M. Yogo (2017): "Worker Betas: Five Facts about Systematic Income Risk," *American Economic Review P&P*, 107, 397–403.

Harmenberg, K. (2021): "The Labor-Market Origins of Cyclical Skewness," Tech. rep., University of Copenhagen.

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

Kaplan, G., B. Moll, and G. L. Violante (2018): "Monetary Policy According to HANK," *American Economic Review*, 108, 697–743.

Kramarz, F., E. Nimier-David, and T. Delemotte (2021): "Inequality and Earnings Dynamics in France: National Policies and Local Consequences," Working paper, Institut Polytechnique de Paris.

Krebs, T. (2003): "Human Capital Risk and Economic Growth*," *Quarterly Journal of Economics*, 118, 709–744.

——— (2007): "Job Displacement Risk and the Cost of Business Cycles," *American Economic Review*, 97, 664–686.

Mankiw, N. G. (1986): "The equity premium and the concentration of aggregate shocks," *Journal of Financial Economics*, 17, 211–219.

McKay, A. (2017): "Time-varying idiosyncratic risk and aggregate consumption dynamics," *Journal of Monetary Economics*, 88, 1–14.

McKay, A. and R. Reis (2021): "Optimal Automatic Stabilizers," *The Review of Economic Studies*, 88, 2375–2406.

Saez, E. (2001): "Using Elasticities to Derive Optimal Income Tax Rates," *Review of Economic Studies*, 68, 205–229.

Schmidt, L. (2014): "Climbing and Falling Off the Ladder: Asset Pricing Implications of Labor Market Event Risk," Working paper, University of California at San Diego.

Storesletten, K., C. I. Telmer, and A. Yaron (2001): "The welfare cost of business cycles revisited: Finite lives and cyclical variation in idiosyncratic risk," *European Economic Review*, 45, 1311–1339.

——— (2004): "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy*, 112, 695–717.

A Details of the Objective Function

The objective function is the sum of several components (as listed in the lower panel of Table III), and we discuss each component in turn. Starting with the "quantiles" component, let $P_{90,5}$ be the average across years of the 90th percentile of the distribution of five-year income growth. Let $P_{90,5}^{Mod}$ be the model-implied value. We then form

$$\sum_{d\in\{1,3,5\}}\left[\left(\frac{P_{10,d}^{Mod}-P_{10,d}}{P_{10,d}}\right)^2+\left(\frac{P_{50,d}^{Mod}-P_{50,d}}{P_{90,d}}\right)^2+\left(\frac{P_{90,d}^{Mod}-P_{90,d}}{P_{90,d}}\right)^2\right].$$

Notice that we use the 90th percentile in the denominator of the difference between the medians to avoid dividing by a value near zero. Turning to the "kurtosis" component, we add

$$\sum_{d\in\{1,5\}} \left(\frac{K_d^{Mod} - K_d}{K_d}\right)^2,$$

where K_d is the kurtosis of the distribution of d-year changes in income averaged across years. Turning to the cross sectional variance profile, we add

$$10 \times \sum_{a\in\{25,35,45,55\}} \left(\frac{V_a^{Mod} - V_a}{V_a}\right)^2,$$

where V_a is the cross-sectional variance of income among individuals of age a. Turning to the histogram component, we use a histogram to approximate the distribution of income growth between 1995 and 1996. The histogram has 279 equally spaced bins on a domain of log income changes that runs from -4.0 to 4.0. The tail mass moments are the CDF at -1.2 and one minus the CDF at 1.2. The left-tail slope is calculated by regressing the log density on the domain $[-4.0, -1.2]$ on the midpoints of the bins in that domain. The right-tail slope is formed analogously on $[1.2, 4.0]$. We then add

$$\left(\frac{M_{Left}^{Mod} - M_{Left}}{M_{Left}}\right)^2 + 5 \times \left(\frac{S_{Left}^{Mod} - S_{Left}}{S_{Left}}\right)^2,$$

where M_{Left} is the mass in the left tail and S_{Left} is the slope of the left tail. We then add the same components for the right tail. Turning to the skewness time series component, we add

$$\sum_{d\in\{1,3,5\}} \frac{1}{T} \sum_{t=1}^{T} \left(\frac{N_{d,t}^{Mod} - N_{d,t}}{0.2}\right)^2,$$

where $N_{d,t}$ is Kelley skewness of the distribution of d-year changes in income in year t. Skewness can take values near zero, so we divide by 0.2 instead, which is about two standard deviations of the fluctuations in skewness. Finally, we have the factor structure. We have seven business cycle episodes in all, but for exposition let us focus on an episode that starts in year t and ends in year τ. For this calculation, we winsorize the simulated data at the 10th and 90th percentiles of income growth between t and τ to reduce simulation noise. We then compute average income for each individual for years $t-5$ to $t-1$, and we bin the individuals by their percentile in the distribution of past income. With the groups held fixed, we then compute the average income for each group in t and in τ; e.g., $\bar{y}_{p,t} \equiv E[\exp(y_{i,t})|i \in p]$, where p is a group or percentile. We then take the log change in the averages $s_p = \log \bar{y}_{p,\tau} - \log \bar{y}_{p,t}$. For values of $p \in [11, 80]$, we regress s_p on p. Let $L_{[11,80],[t,\tau]}$ be the slope coefficient of this regression. We then regress s_p on p on the domain $[81, 100]$ to form $L_{[81,100],[t,\tau]}$. We add to the objective function

$$\frac{1}{100}\left(\frac{L^{Mod}_{[11,80],[t,\tau]} - L_{[11,80],[t,\tau]}}{L_{[11,80],[t,\tau]}}\right)^2 + \frac{1}{100}\left(\frac{L^{Mod}_{[81,100],[t,\tau]} - L_{[81,100],[t,\tau]}}{L_{[81,100],[t,\tau]}}\right)^2.$$

We repeat these steps for the business cycle episodes: $[1979, 1983]$, $[1983, 1990]$, $[1990, 1992]$, $[1992, 2000]$, $[2000, 2002]$, $[2002, 2007]$, $[2007, 2010]$.

The data moments are taken from the data appendixes provided by Guvenen et al. (2014) (GOS) and Guvenen et al. (2015) (GKOS). The percentiles are taken from Table C1 in the Excel data file that accompanies GOS. The age profile of the cross-sectional variance of log income is taken from Figure A.2 in the data appendix of GKOS. The factor structure of business cycle incidence is taken from Figures 13 and 14 in the data appendix for GOS. Kurtosis is taken from Figure 10 in GKOS, and the log density of one-year income growth comes from Figure 11. The standard deviations of one-year and five-year changes in income are taken from GOS Appendix Table A8.

B Results from Additional Specifications

Table IV – Estimated Parameters

Parameters		(2')	(3')	(4')
σ_γ	St. dev. of fixed effects	0.628	0.608	0.590
p^ζ	Probability of full-year empl.	0.585	0.607	0.620
λ	Transitory exponential parameter	3.355	3.102	3.024
ψ	Scarring effect of transitory shock	–	0.143	0.138
p_2^η	Mix. probab. for persist. innov 2	–	0.134	0.112
p_3^η	Mix. probab. for persist. innov 3	–	0.074	0.082
σ_1^η	Std. dev. for persistent innov. 1	0.114	0.056	0.021
σ_2^η	Std. dev. for persistent innov. 2	–	0.129	0.214
σ_3^η	Std. dev. for persistent innov. 3	–	0.184	0.225
μ_1^η	Center for persistent component 1	–	–0.006	-0.019
μ_2^η	Center for persistent component 2	–	0.045	0.003
μ_3^η	Center for persistent component 3	–	0.236	0.136
α_1	Factor struct. slope, low income	–	–	–0.720
α_2	Factor struct. slope, high income	–	–	1.045
$\bar{q}$	Factor structure threshold	–	–	0.876

Table V – Targeted and Fitted Moments

	US Data	Model Specifications					
Moments		(2)	(2')	(3)	(3')	(4)	(4')
P10, one-year change	–0.434	–0.448	-0.438	–0.435	-0.429	–0.426	-0.433
P10, three-year change	–0.585	–0.486	-0.482	–0.495	-0.478	–0.485	-0.485
P10, five-year change	–0.631	–0.517	-0.519	–0.529	-0.516	–0.52	-0.521
P50, one-year change	0.020	0.015	0.002	0.009	0.014	0.014	0.012
P50, three-year change	0.061	0.020	0.007	0.012	0.015	0.016	0.014
P50, five-year change	0.103	0.035	0.023	0.021	0.028	0.025	0.027
P90, one-year change	0.474	0.435	0.441	0.434	0.433	0.432	0.435
P90, three-year change	0.705	0.487	0.497	0.513	0.509	0.518	0.513
P90, five-year change	0.848	0.552	0.566	0.591	0.584	0.599	0.589
Kurtosis, one-year change	20.00	22.730	23.770	23.280	23.487	23.880	23.76
Kurtosis, five-year change	12.00	17.78	18.298	16.83	17.215	16.97	17.515
Cross-sectional var., age 25	0.595	0.592	0.562	0.535	0.547	0.556	0.542
Cross-sectional var., age 35	0.719	0.711	0.693	0.692	0.694	0.717	0.695
Cross-sectional var., age 45	0.814	0.832	0.821	0.841	0.845	0.853	0.839
Cross-sectional var., age 55	0.905	0.953	0.95	0.974	0.994	0.991	0.988
Left-tail mass	0.024	0.024	0.023	0.024	0.024	0.024	0.025
Left-tail slope	1.260	1.282	1.297	1.282	1.268	1.294	1.271
Right-tail mass	0.015	0.021	0.022	0.018	0.019	0.018	0.02
Right-tail slope	–2.035	–1.482	-1.385	–1.561	-1.516	–1.551	-1.493
Objective value		2.439	2.633	1.592	2.275	1.195	1.943
Quantiles		0.295	0.281	0.238	0.261	0.236	0.247
Kurtosis		0.251	0.311	0.189	0.219	0.210	0.246
Cross-sectional var. profile		0.035	0.070	0.187	0.190	0.157	0.185
Histogram		0.525	0.788	0.341	0.403	0.347	0.467
Skewness time series		0.833	0.688	0.143	0.704	0.143	0.657
Factor Structure		0.498	0.494	0.496	0.496	0.101	0.141

Notes: This table shows the model fit for each estimated model. The top panel displays all the individual targeted moments with the exception of the time series for the Kelley skewness of one-year and five-year income growth and the factor structure moments. The first column contains the targeted moments computed from SSA data (Guvenen et al. (2014, 2021)) and subsequent columns show the implied values from the estimated models. The bottom panel shows the weighted contribution of selected sets of moments to the objective function. The top row of the bottom panel shows the total value of the objective function, including factor structure moments.