A Tractable Income Process for Business Cycle Analysis*

Fatih Guvenen† Alisdair McKay‡ Conor Ryan§

March 9, 2026

Abstract

We estimate a tractable 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 within-cohort 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 Toronto, FRB of Minneapolis, and NBER; fatih.guvenen@utoronto.ca

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

§Corresponding Author: 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 is 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 enables 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 find 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. Second, our income process features a persistent first-order autoregressive process (AR(1)) with innovations drawn from a normal mixture distribution that can 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 different sensitivities (or exposures) to aggregate fluctu-

[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).

ations. This factor structure captures a systematic component in idiosyncratic income fluctuations, which has empirical support in the data (see Figure 4).[2]

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.

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 7, we discuss a simple, heterogeneous-agent, business cycle model as an example of incorporating this income process into an application. In that application, we find that time-varying idiosyncratic risk makes a substantial contribution to aggregate consumption dynamics and our income processes that incorporate this generate a good fit to the observed consumption dynamics.

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 skewness 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 and 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,

[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.

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 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.

We contribute to the large literature on income dynamics, particularly the branch of this literature that develops statistical models of income processes with non-Gaussian innovations. Geweke and Keane (2000) argue that a normal mixture model is able to capture the leptokurtic nature of earnings changes and investigate the implications for earnings mobility. Arellano et al. (2017) estimate a non-linear model of earnings dynamics with a flexible (non-parametric) specification for the innovation distribution and the persistence of earnings. Guvenen, Karahan, Ozkan, and Song (2021) estimate an income process with normal mixture innovations and non-employment shocks similar to what we do here. Relative to these papers, a key focus of our analysis is to incorporate business cycle variation in the innovation distribution and incorporate unequal incidence of the business cycle across income levels.[3] Structural modeling is an alternative approach to modeling income dynamics. For example, Hubmer (2018) shows that a search-theoretic job ladder model can rationalize negative skewness and excess kurtosis in earnings dynamics. With respect to cyclical effects on earnings risks, Huckfeldt (2022) shows that a search model with skill-intensive and skill-insensitive jobs can account for the cyclicality of post-displacement earnings losses. The structural approach is valuable to gain an understanding of the deeper causes of the earnings dynamics we explore, but within the context of heterogeneous-agent business-cycle models, researchers commonly use more statistical approaches for their computational tractability.

[3]In a very recent paper, Almuzara et al. (2025) extend the Arellano et al. (2017) approach to include aggregate shocks. Their non-parametric approach can incorporate non-linear dynamics, but is perhaps more challenging to adopt in applications as the process is described by non-parametric functions as opposed to a small number of parameters as in our approach.

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 reports the sensitivity of our estimates to changes in the empirical moments. Section 7 discusses incorporating the income process into a dynamic programming problem, similar to those commonly used in heterogeneous agent models of the business cycle and reports simulations of aggregate consumption dynamics.

2 Motivating Facts

In this section, we present the key features of individual income dynamics that we believe are most relevant 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.[4] Moreover, the dispersion in incomes within a cohort also informs us about the dispersion of individual fixed effects. If we allow for heterogeneity 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

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

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

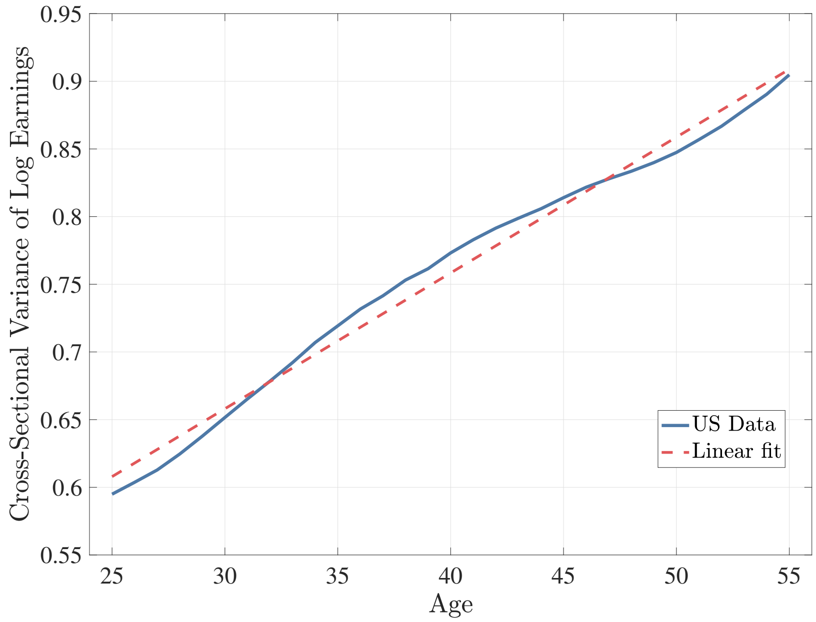

Note: The cross-sectional variance of log income increases almost linearly over the life cycle. The blue line represents the variance of log earning within each age group. The dashed line shows the fit of a linear regression of cross-sectional variance on age.

superimposed on a Gaussian (normal) distribution with the same mean and variance.[5] 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 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

[5]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)

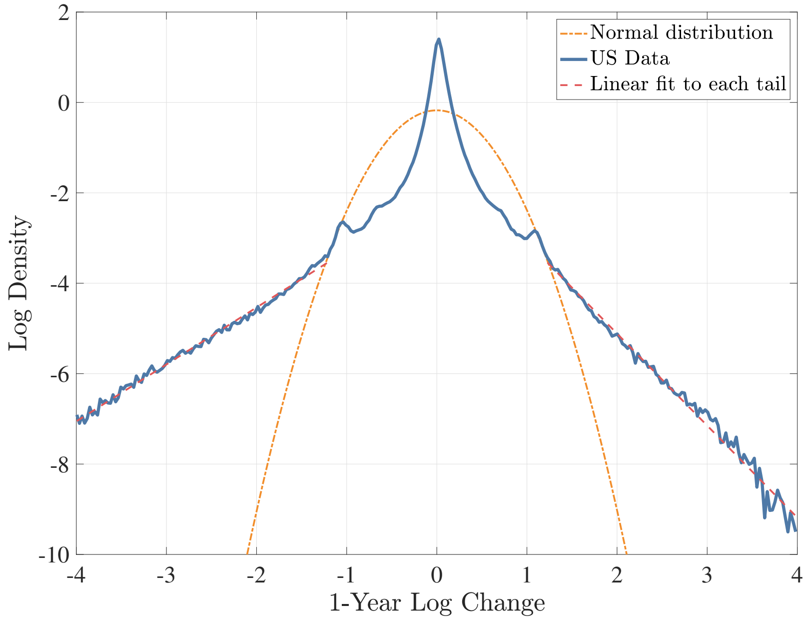

Note: The figure shows the distribution of one-year income growth between 1995 and 1996 (Guvenen et al., 2021). The dashed lines show the fit of a linear regression of the log-density on the one-year log change within the right and left tails. The dash-dot line shows a Gaussian distribution with the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

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 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.[6] 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

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

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

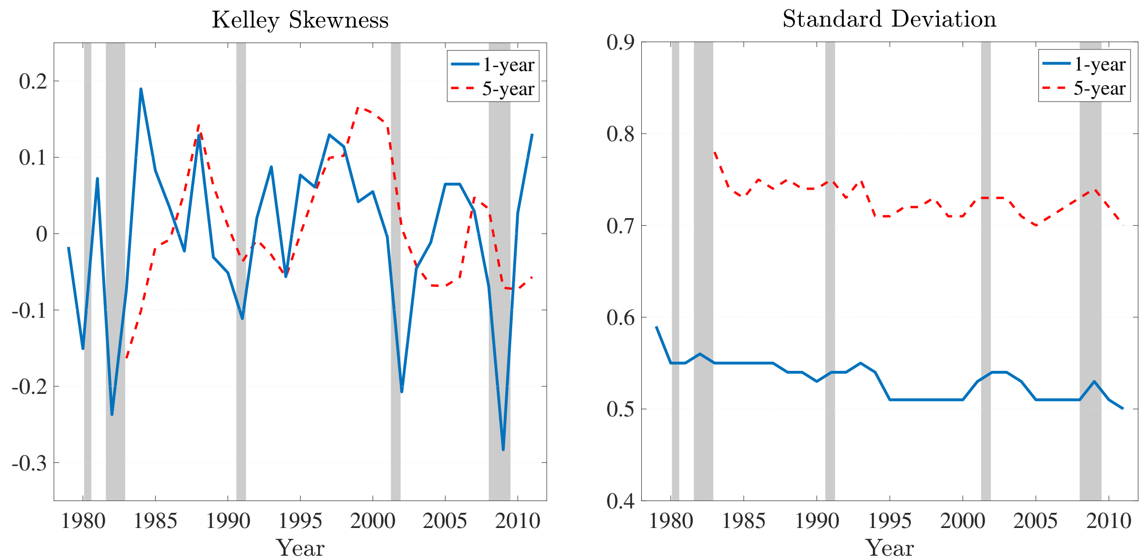

Note: The left panel displays the time series of Kelley skewness of the distribution of one-year (solid line) and five-year (dashed line) income growth. The right panel shows the analogous standard deviations. The plots are aligned with the first year of the time differenced data.

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 2010. The graph shows a clear 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

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

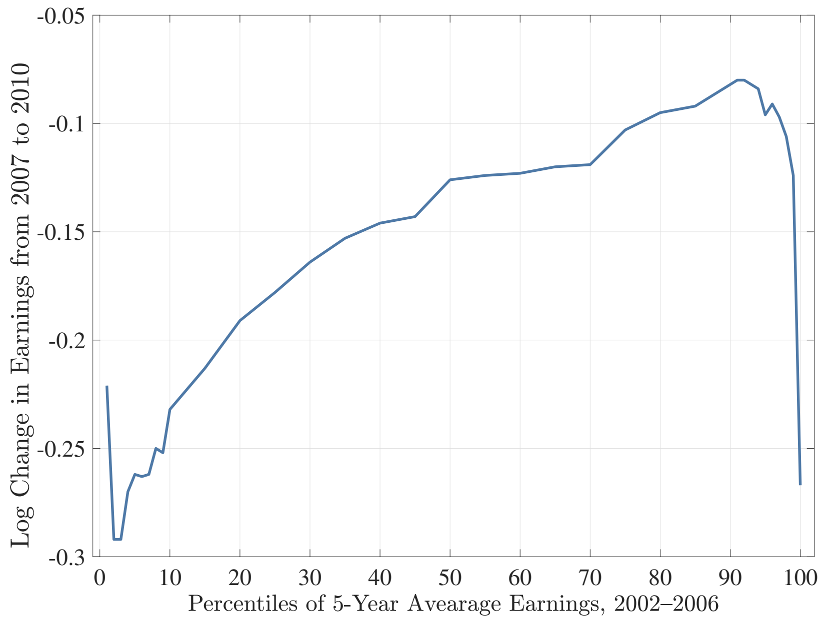

Note: The figure plots the log income change during the Great Recession for prime-age males ranked by their pre-recession average income (2002 to 2006). The income data are averaged within each percentile bin before calculating the log difference between 2007 and 2010. The figure shows income growth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

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

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

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:

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

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 impose $z_{i,h_i} = 0$ as initial condition.

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 $1 - \ell_{i,t}$. As ζ_{it} is the log of this factor, its support is therefore $(-\infty, 0]$. 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

$$\begin{aligned} \mu^{\eta}_{1,t} &= \bar{\mu}^{\eta}, \\ \mu^{\eta}_{2,t} &= \mu^{\eta}_{2} - x_t, \\ \mu^{\eta}_{3,t} &= \mu^{\eta}_{3} - x_t. \end{aligned}$$

The parameter $\bar{\mu}^{\eta}$ is a normalization such that $\bar{\mu}^{\eta} p^{\eta}_{1} + \mu^{\eta}_{2} p^{\eta}_{2} + \mu^{\eta}_{3} p^{\eta}_{3} = 0$. 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. For tractability, we also assume that $p^{\eta}_{2} = p^{\eta}_{3}$ and $\sigma^{\eta}_{2} = \sigma^{\eta}_{3}$. 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.[7] 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.

In what follows, we will consider several special cases of the general process in (1). Many of these cases will assume lifecycle income profiles are homogeneous (κ_i =0) and in these cases we impose that $\rho = 1$. This restriction is without much consequence as ρ would be estimated to be close to 1 even without this restriction, because 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 7.

We also consider versions of several specifications with HIP, in which the linear re-

[7]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.

lationship 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 $z_{i,t} = 0$.[8] 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).[9]

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

[8] 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.

[9] 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.

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.[10] 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.[11] 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).

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.

[10] 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.

[11] 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 ζ	η	ψ	ρ	σ^{κ}	Factor Str.
(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.

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 kurtosis of income growth observed in the data. It also 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

TABLE II – Estimated Parameters

Parameters		Model Specifications (1)	(2)	(3)	(4)	(5)	(6)
σ_γ	St. dev. of fixed effects	0.789	0.623	0.609	0.611	0.616	0.634
σ_1^ζ	St. dev. for transitory shock	0.204	–	–	–	–	–
p^ζ	1 – Probability of nonempl. shock	–	0.560	0.613	0.612	0.706	0.705
λ	Transitory exponential parameter	–	3.453	3.144	3.150	2.701	2.932
ψ	Scarring effect of transitory shock	–	0.064	0.172	0.162	0.376	0.293
p_2^η	Mix. probab. for persist. innov 2	–	–	0.045	0.030	0.075	0.082
p_3^η	Mix. probab. for persist. innov 3	–	–	0.045	0.030	0.075	0.082
σ_1^η	Std. dev. for persistent innov. 1	0.063	0.115	0.037	0.027	0.096	0.089
σ_2^η	Std. dev. for persistent innov. 2	–	–	0.081	0.138	0.262	0.320
σ_3^η	Std. dev. for persistent innov. 3	–	–	0.081	0.138	0.262	0.320
μ_1^η	Center for persistent component 1	–	–	0.011	0.009	0.025	0.018
μ_2^η	Center for persistent component 2	–	–	-0.082	-0.110	-0.150	-0.127
μ_3^η	Center for persistent component 3	–	–	0.309	0.395	0.432	0.314
β	Loading on aggregate wage	–	–	-8.193	-9.256	-6.783	-6.775
α_1	Factor struct. slope, low income	–	–	–	-0.390	–	-0.768
α_2	Factor struct. slope, high income	–	–	–	0.595	–	1.264
$\bar{q}$	Factor structure threshold	–	–	–	0.760	–	0.660
ρ	AR(1) coefficient	–	–	–	–	0.842	0.824

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. The center for persistent component 1 is the implied normalization, $\bar{\mu}$.

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

Table III – Targeted and Fitted Moments

Moments	US Data	Model Specifications (1)	(2)	(3)	(4)	(5)	(6)
P10, one-year change	–0.434	–0.367	–0.438	-0.418	-0.420	-0.416	-0.415
P10, three-year change	–0.585	–0.379	–0.485	-0.476	-0.475	-0.606	-0.579
P10, five-year change	–0.631	–0.380	–0.524	-0.501	-0.498	-0.689	-0.641
P50, one-year change	0.020	0.002	0.011	0.019	0.019	0.023	0.018
P50, three-year change	0.061	0.007	0.015	0.026	0.027	0.023	0.019
P50, five-year change	0.103	0.023	0.030	0.046	0.044	0.044	0.043
P90, one-year change	0.474	0.370	0.431	0.427	0.447	0.453	0.442
P90, three-year change	0.705	0.393	0.493	0.514	0.525	0.651	0.620
P90, five-year change	0.848	0.426	0.566	0.605	0.620	0.778	0.738
Kurtosis, one-year change	20.00	3.000	23.40	24.247	24.233	22.89	23.19
Kurtosis, five-year change	12.00	3.000	17.61	17.135	17.306	11.07	12.29
Cross-sectional var., age 25	0.595	0.668	0.555	0.545	0.562	0.578	0.597
Cross-sectional var., age 35	0.719	0.708	0.692	0.695	0.712	0.744	0.727
Cross-sectional var., age 45	0.814	0.745	0.834	0.838	0.850	0.814	0.800
Cross-sectional var., age 55	0.905	0.787	0.969	0.958	0.981	0.927	0.912
Left-tail mass	0.024	0.000	0.023	0.023	0.023	0.023	0.021
Left-tail slope	1.260	—	1.292	1.299	1.308	1.277	1.274
Right-tail mass	0.015	0.000	0.020	0.018	0.018	0.015	0.016
Right-tail slope	–2.035	—	–1.538	-1.599	-1.592	-1.928	-1.805
Objective value		6.696	2.333	1.544	1.176	0.977	0.582
Quantiles		0.814	0.279	0.252	0.235	0.034	0.047
Kurtosis		1.284	0.247	0.228	0.240	0.027	0.026
Cross-sectional var. profile		0.395	0.116	0.125	0.122	0.026	0.002
Histogram		2.995	0.444	0.277	0.305	0.020	0.099
Skewness time series		0.698	0.749	0.160	0.154	0.176	0.154
Factor Structure		0.499	0.498	0.502	0.119	0.693	0.251

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

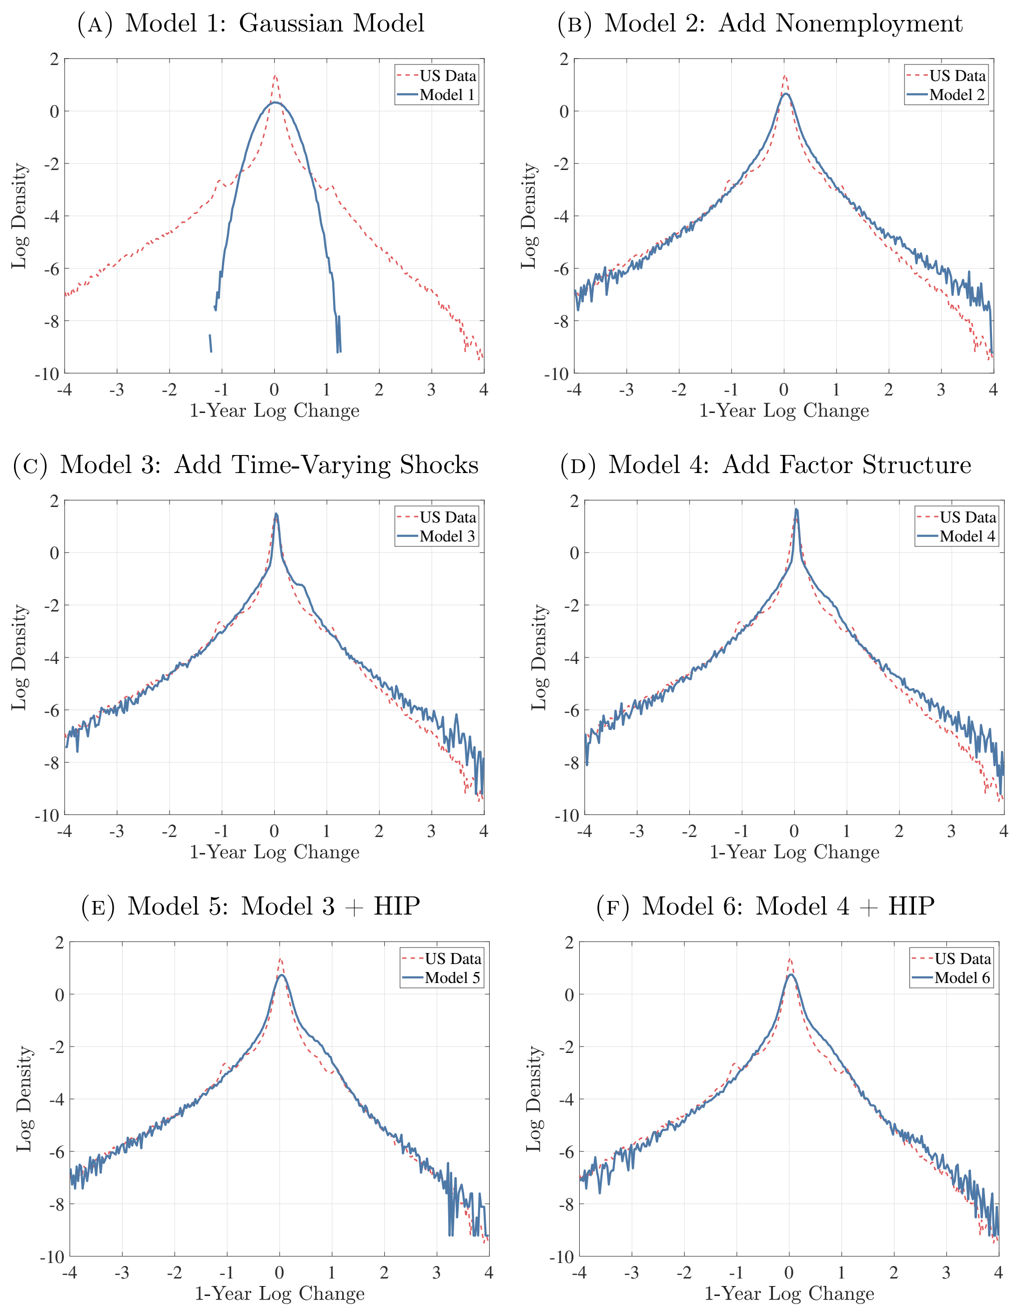

(A) Model 1: Gaussian Model

(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 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.

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 44% and mean 0.26. The passthrough coefficient, ψ, which determines the fraction of ζ that leaves a permanent effect (through equation 3) is about 6.4%. 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, 8.1% of individuals see their income fall by 50% or more because of the nonemployment shock alone, and 1.7% see their income fall by more than 95%.

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 23.4 and 17.6 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 –1.54 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, Model 2 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).[12]

[12] 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. The estimates are reported in Tables IV and V in the appendix.

Figure 6 – Skewness of Annual Income Growth

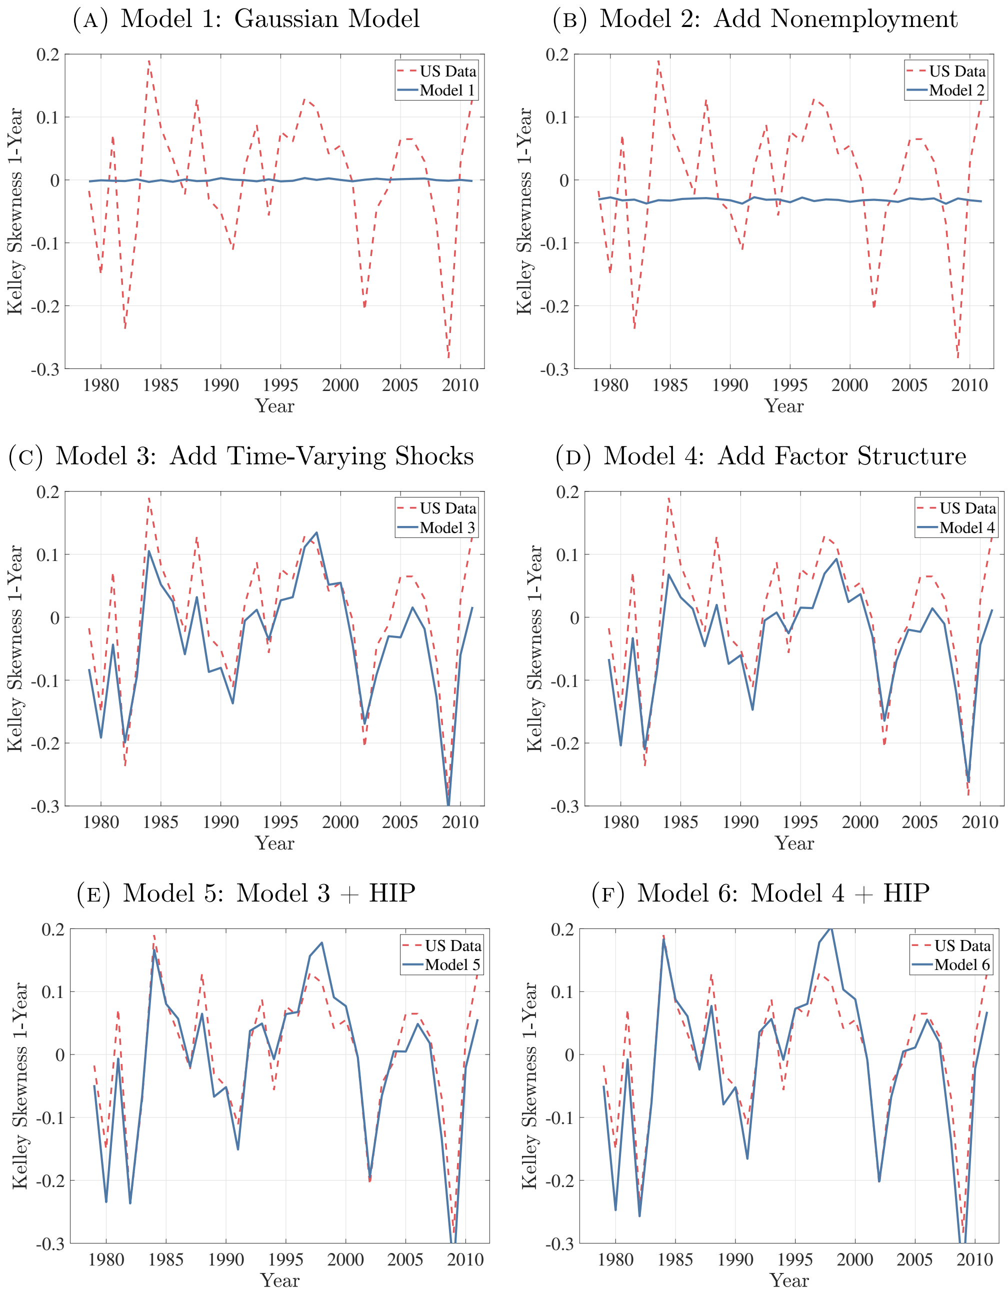

Note: This figure shows the fit to the time series of the Kelley skewness of one-year income growth for each model, as specified in Table I. The dashed line represents the Kelley skewness of the empirical distribution of income growth with respect to one year prior (Guvenen et al. (2014)), and the solid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

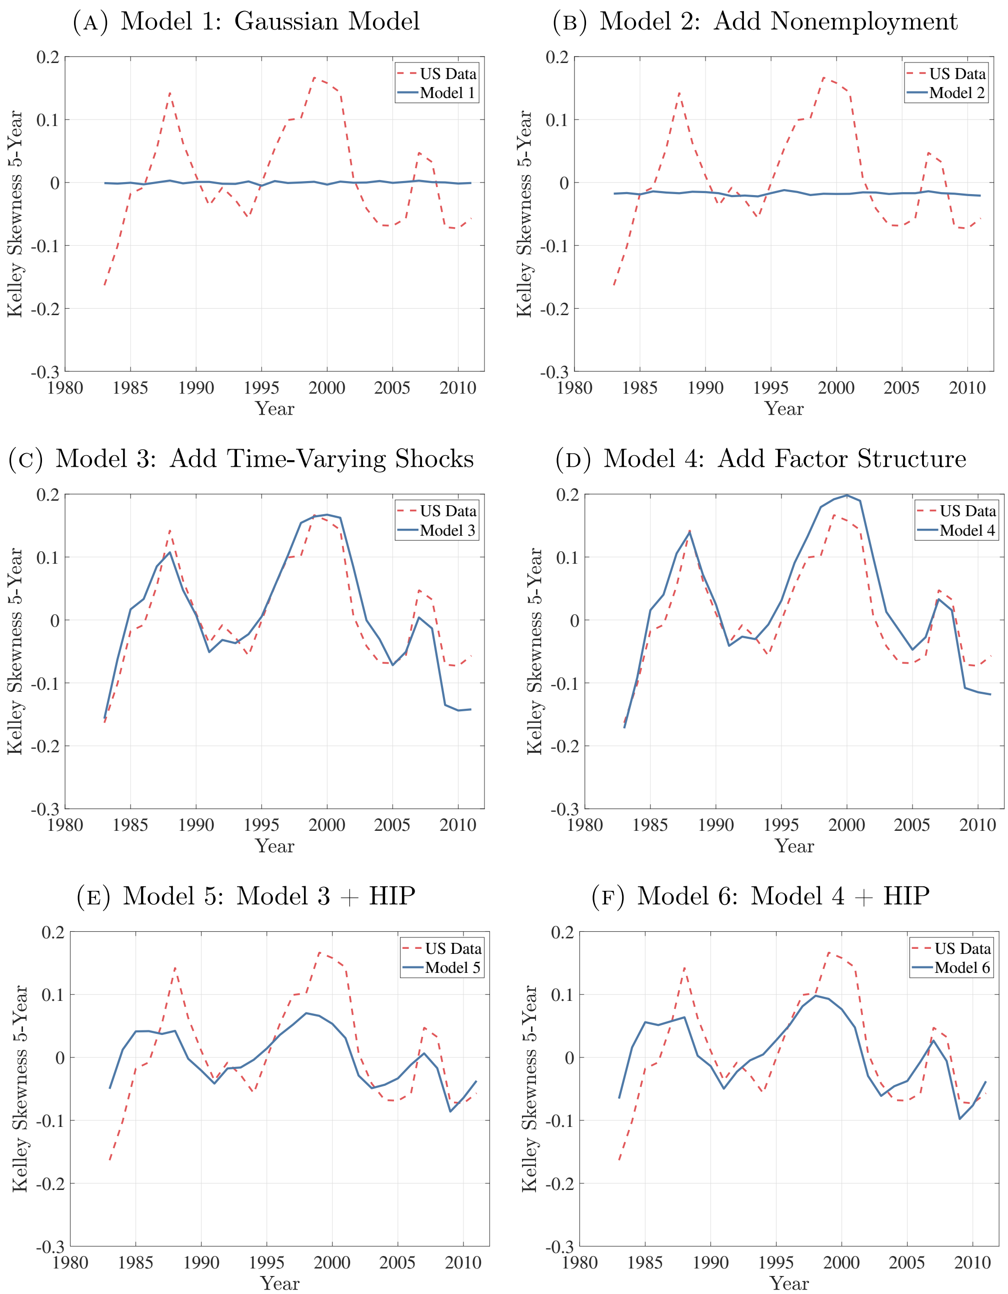

Note: This figure shows the fit to the time series of the Kelley skewness of five-year income growth for each model, as specified in Table I. The dashed line represents the Kelley skewness of the empirical distribution of income growth with respect to 5 years prior (Guvenen et al. (2014)), and the solid 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)).[13] 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 91% probability) and two tail shocks that are realized with 4.5% probability each. The frequent shock is small, with a mean of 1.1% and a standard deviation of 3.7%, whereas the other two shocks are are larger with means of -8.2% and 30.9% and standard deviations of 8.1%. These estimates are consistent with the idea that in most years, persistent income changes are small, with large jumps happening less frequently. The estimates of the scarring effect component show some interaction with the introduction of the normal mixture, with the passthrough coefficient rising (from 6.4% to 17.2%) and the shock frequency declining slightly (from 44% 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

[13] 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

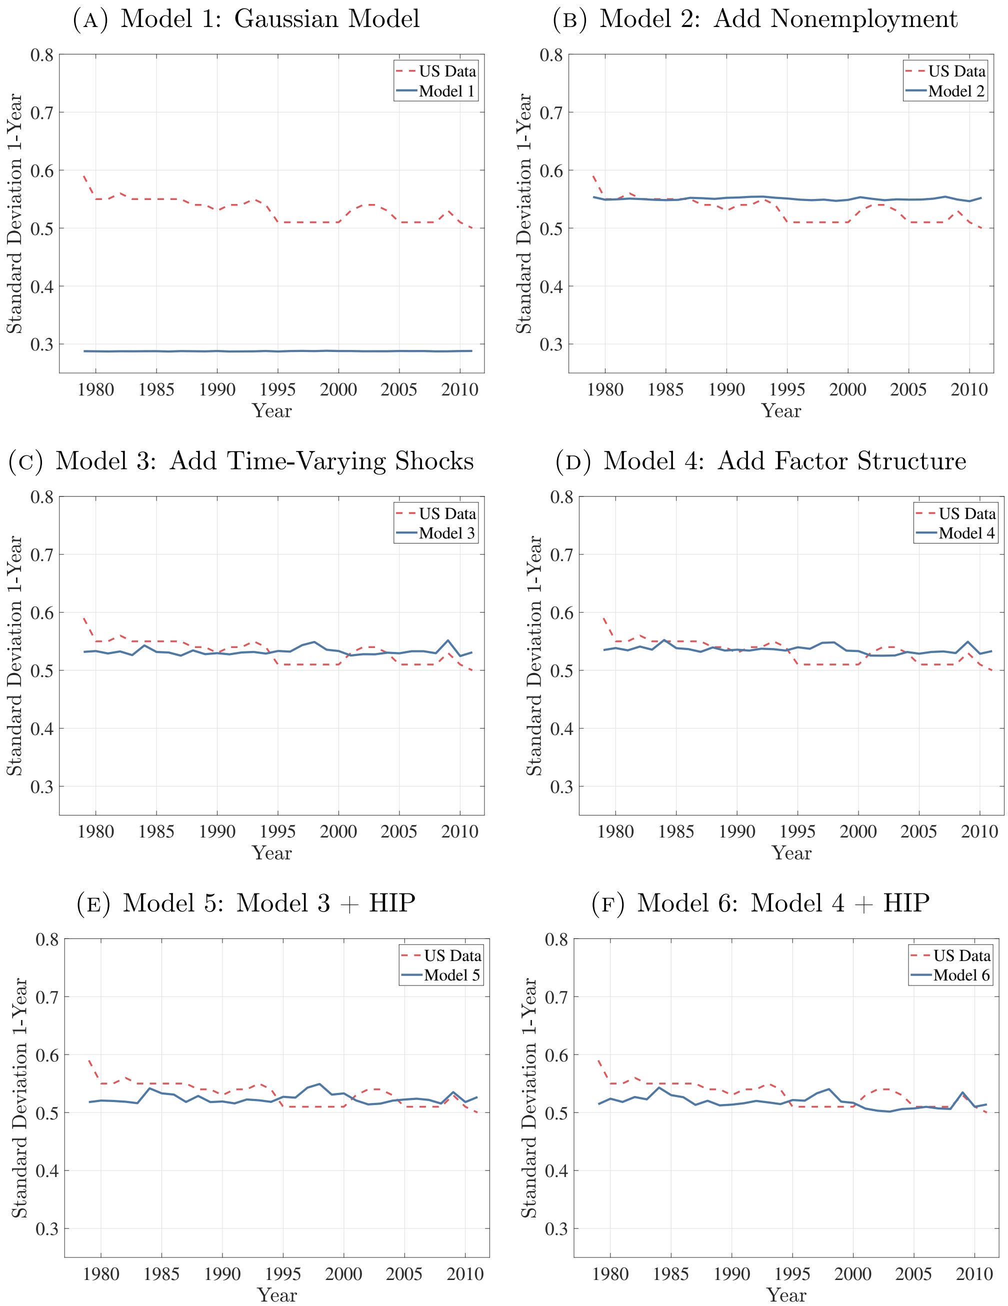

Note: This figure shows the fit to the time series of the standard deviation of one-year income growth for each model, as specified in Table I. The dashed line represents the standard deviation of the empirical distribution of income growth with respect to one year prior (Guvenen et al. (2014)), and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

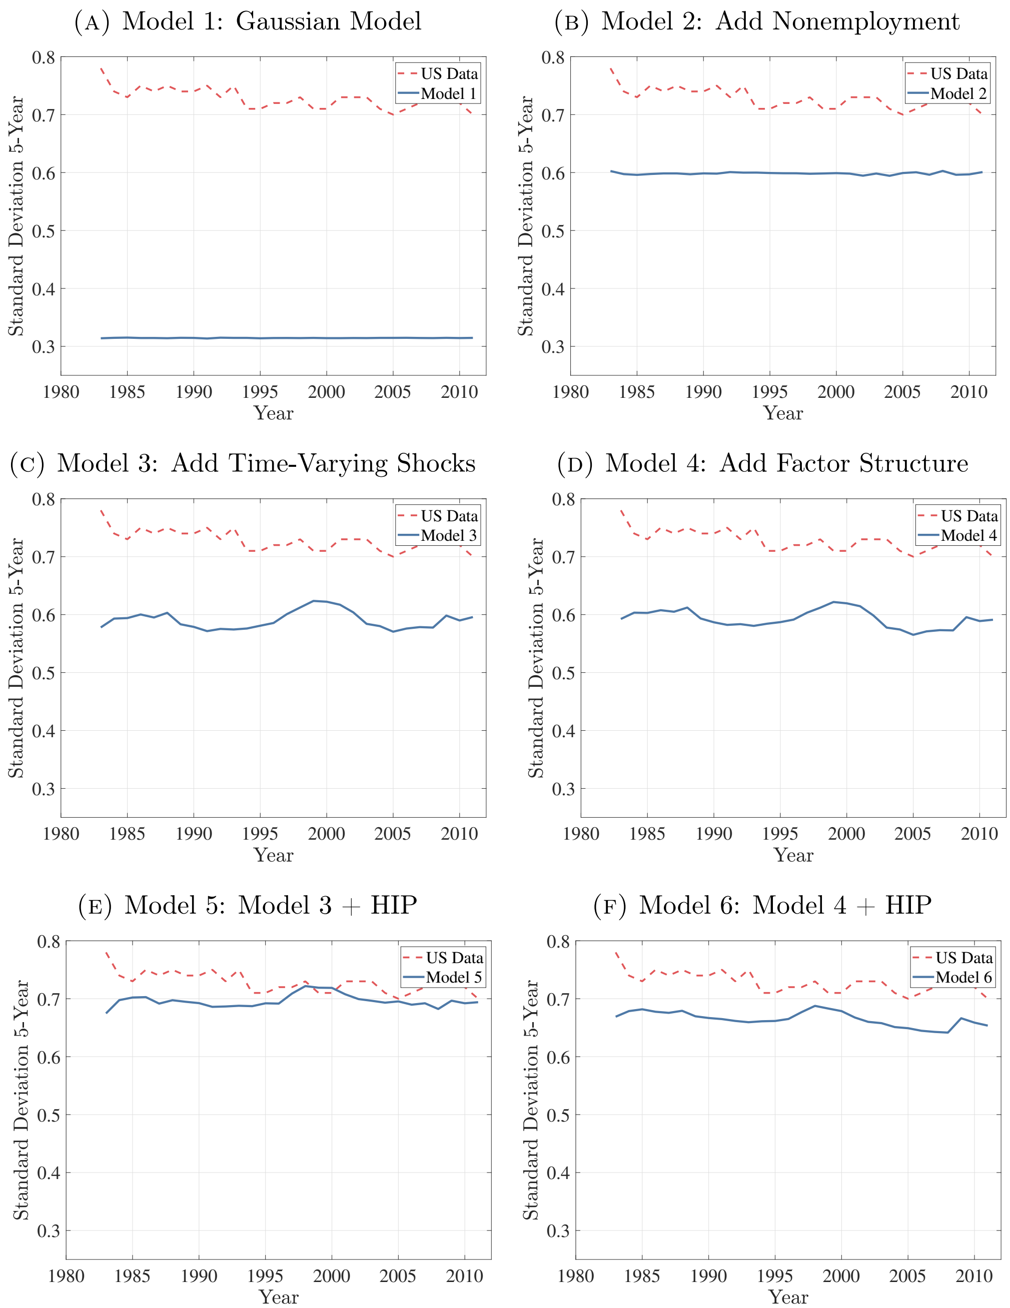

Note: This figure shows the fit to the time series of the standard deviation of five-year income growth for each model as specified in Table I. The dashed line represents the standard deviation of the empirical distribution of income growth with respect to 5 years prior (Guvenen et al. (2014)), and the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

The improvement in fit over Model 2 comes primarily from the time-varying moments. To demonstrate, we estimate a specification identical to Model 3 in all respects, but with a static normal mixture distribution. The results are displayed 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 2.33 to 2.07. The main improvements come from a better fit of the tails of the distribution and the kurtosis.

An alternative method to incorporate time-varying income risk would be to introduce business cycle variation through the transitory non-employment shocks. While the model is purely statistical, time-varying non-employment risk has the benefit of a straightforward economic interpretation. We estimate a specification in which the persistent shock is drawn from a static normal mixture, and the non-employment shock probability is a logistic function that depends on the aggregate state.[14] As seen in Appendix Tables IV and V, this process (Model 3") improves over an identical process without time-varying risk (Model 3'). However, allowing the non-employment shock to depend on the business cycle is less flexible than the time-varying mixture of normal distributions included in the baseline specification. In particular, the time-varying normal mixture models are better able to match the time series of procyclical skewness.

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,

[14]Specifically, the probability of a non-employment shock is specified as $\frac{\exp(\beta_0^{\xi}+\beta_1^{\xi} w_t)}{1+\exp(\beta_0^{\xi}+\beta_1^{\xi} w_t)}$.

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. 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 little importance.

The estimated parameters of the factor structure can be interpreted as follows. Rearranging equation 1 for a value of $\gamma_i + z_{i,t} < \bar{q}$ we have

$$y_{i,t} = w_t + [1 - \alpha_1 w_t] (\gamma_i + z_{i,t}) + \tilde{\zeta}_{i,t},$$

where we have used $\kappa_i = 0$ for simplicity. With an estimate of $\alpha_1 = -0.4$ and a difference in w_t of 0.15, which is the change in w_t from 2007 to 2010, one finds that the effects of persistent idiosyncratic shocks are amplified by a factor of 1.06 in 2010 as compared to 2007. The parameter $\bar{q}$ is a threshold that is on average near the 80th percentile of the distribution of $\gamma_i + z_{i,t}$.[15]

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.[16] The specification of the factor structure that we estimate is intended to be parsimonious, but a better fit is possible if the model includes more aggregate shocks thereby allowing for more differences across business cycle episodes. For example, we estimated a version of

[15] The exact percentile varies over time as the distribution of $\gamma_i + z_{i,t}$ changes with the time-varying distribution of income risks.

[16] Note that Model 3 predicts that all households are equally exposed to the aggregate component and therefore has no factor structure. If we were to replicate Figure 10 using Model 3, we would find a flat line.

the model with an index of equity prices as an additional aggregate shock with a separate factor structure.[17] This specification can easily match the large negative effect of the 2000-2002 recession.

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.[18]

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 C (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. In light of this tension, it is worth revisiting the argument for $\rho = 1$. The interpretation of the cross-sectional variance profile as evidence of near unit root behavior of log income hinges on the assumption the assumption that there is no deterministic heterogeneity across households. We therefore consider alternative models in which $\rho < 1$ and there is deterministic heterogeneity in income profiles (HIP). As we will see, these alternatives allow for a better fit to the data in some dimensions but not in others. Therefore we do not necessarily view these models as superior to the restricted income profile models already presented.

[17] This would appear in the income process specification as an additional term, e.g. $g(\gamma_i + z_{i,t})s_t$.

[18] While this paper aims to fit a broad set of income dynamics, the factor structure can be better fit if it receives more priority in the objective function. Using a greater weight on factor structure moments can improve the factor structure fit by about 50% at the cost of increasing the distance to the other moments in the model by about 35%.

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 0.84. 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 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 $\rho = 0.82$ is similar to what we found in Model 5. The improvements to the model fit due to HIP also apply in this case (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)

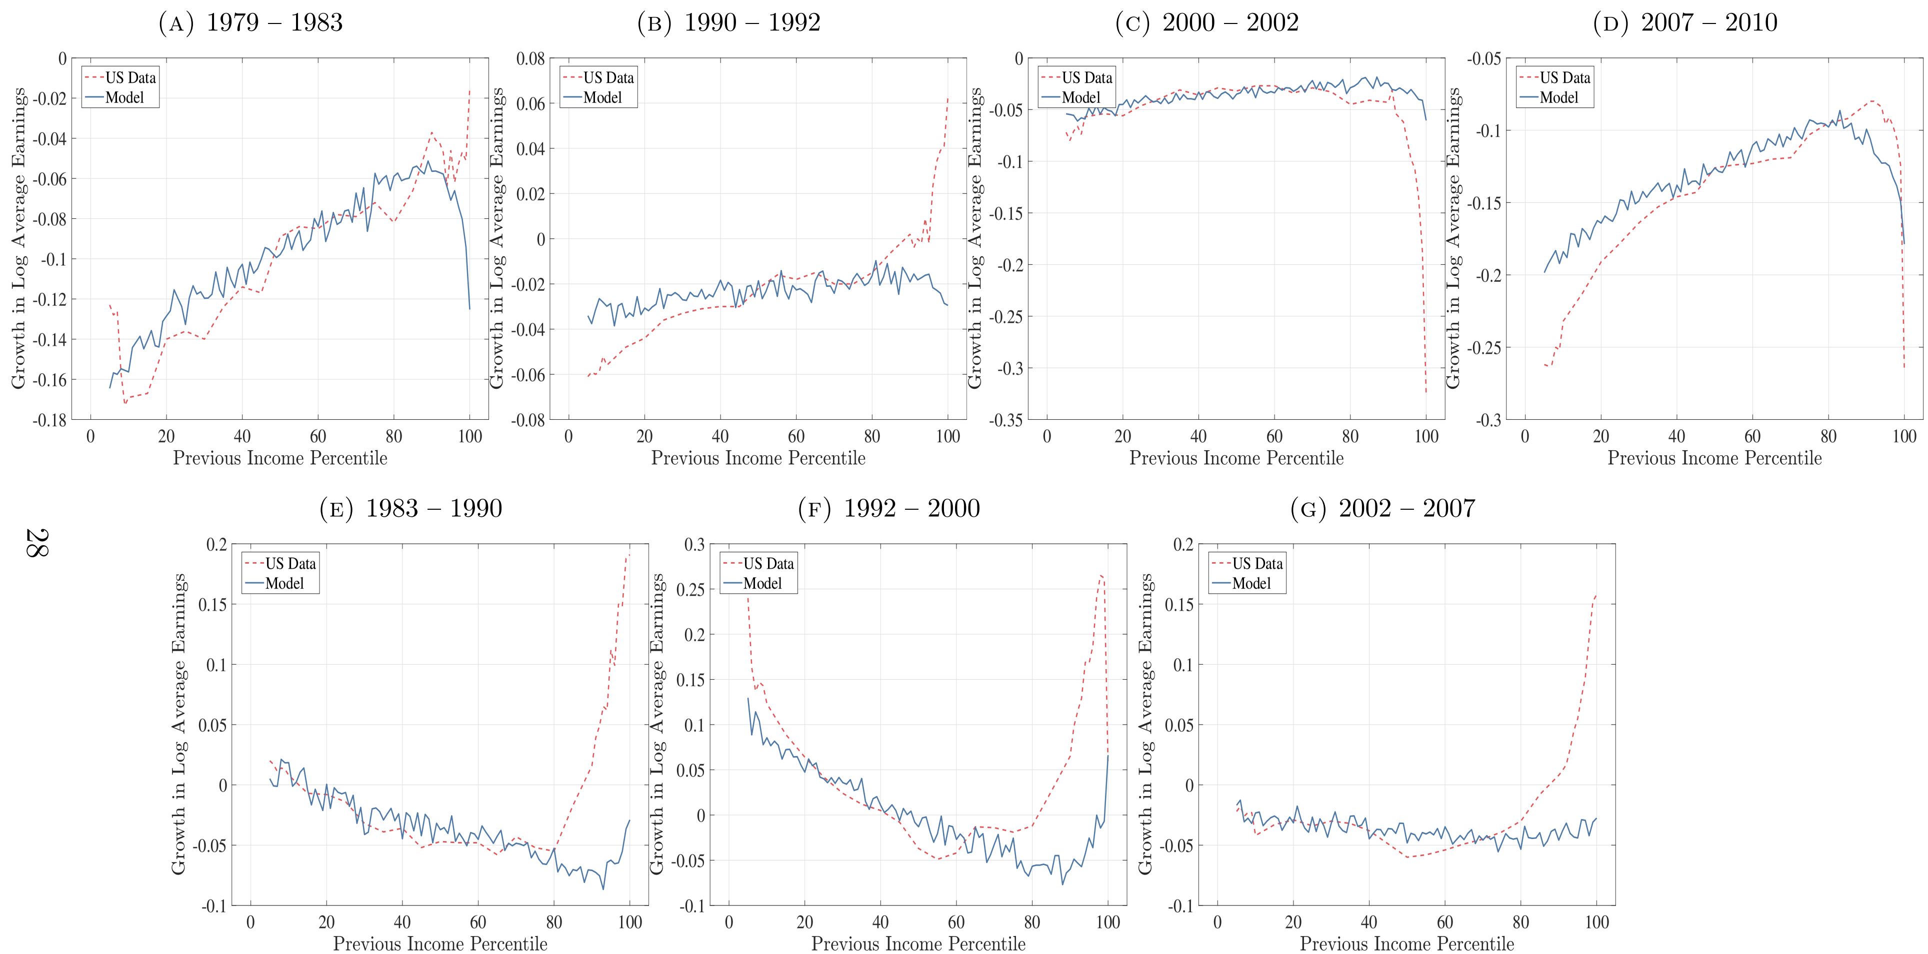

Note: This figure shows the model fit to the factor structure for Model 4 as specified in Table I. The dashed line is the growth in log income over each expansion or recession period for individuals ranked by the income percentile for the 5 years prior to the beginning of the period (Guvenen et al. (2014)). 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)

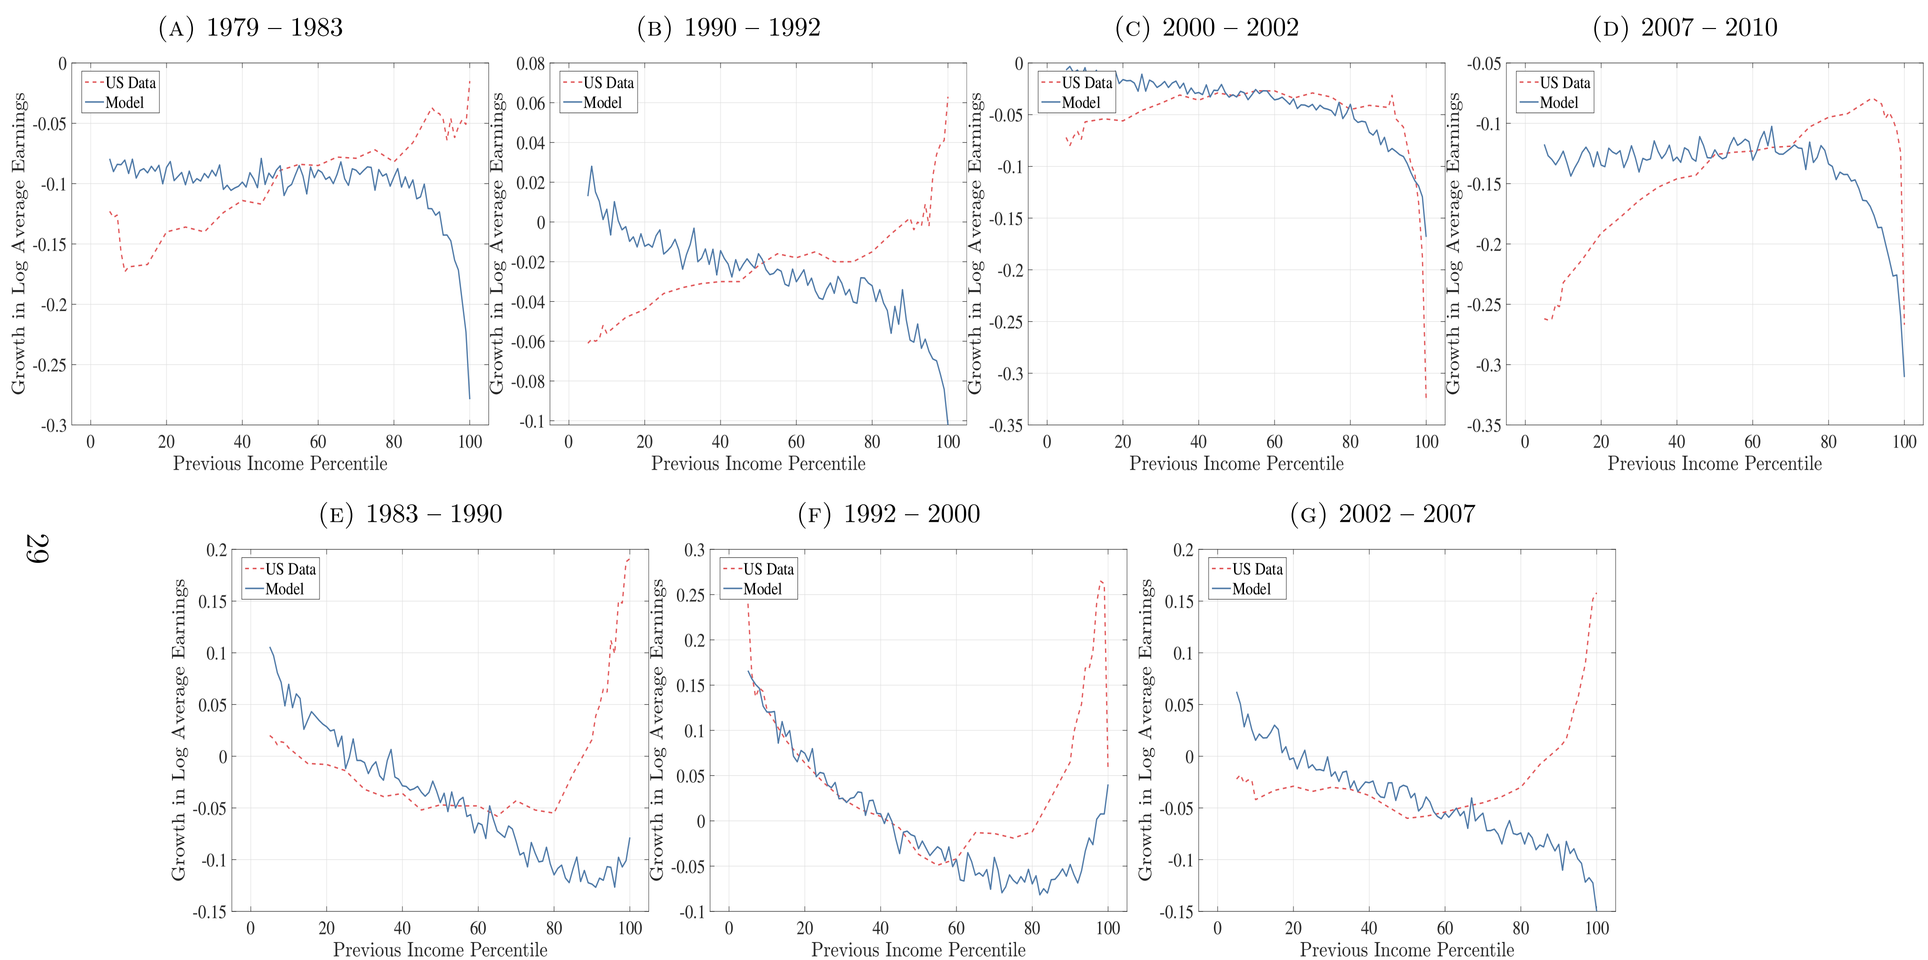

(a) 1979 – 1983 (b) 1990 – 1992 (c) 2000 – 2002 (d) 2007 – 2010

(e) 1983 – 1990 (f) 1992 – 2000 (g) 2002 – 2007

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 Moment Sensitivity

In the previous section, we discuss each component of the income process in the context of the set of moments that motivate the particular component. However, the parameters of the model are estimated jointly to fit all the moments discussed in Section 4. We now formally investigate the sensitivity of the parameter estimates to variation in the empirical moments using the method proposed by Andrews et al. (2017). The intuition for this analysis is to compute the gradient of the estimated parameters with respect to small deviations in the moments.

The output of the analysis is a 13×133 matrix that shows the sensitivity of the 13 parameters to the 133 moments. To summarize the analysis, we report summary measures that perturb the moment vector in particular directions of interest. We briefly describe these here with additional explanation in Appendix Section B. We consider two deviations in the skewness time series: an increase in the amplitude of skewness fluctuations and the directional derivative most influential to the relationship between the aggregate shock and time-varying shape of permanent income risk (β). We consider deviations in the 10^{th}, 50^{th}, and 90^{th} percentiles of the 1-year earnings growth rate, and deviations in the 1-year and 5-year kurtosis. We consider a deviation in the level of the cross-sectional variance at every point in the life-cycle (25, 35, 45, and 55 years of age) and a deviation in the slope of cross-sectional variance over the life-cycle holding fixed the variance at age 25. We consider two deviations in the slopes of the factor structure: above and below the 80th percentile of income. Finally, we consider a deviation that increases the mass in both tails of the 1-year earnings growth histogram and a deviation that increases the slope of both tails.

We display the results for Model 4 in Figure 12. Each panel contains the magnitude of the directional derivatives described above for 12 parameters.[19] We indicate with a star the two directional derivatives with the greatest magnitude for each parameter. The results of the analysis are consistent with the intuition regarding which moments are important for which parameters. The standard deviation of the individual fixed effect σ^{γ} is sensitive to the level of cross-sectional variance, and the probability and variance of the non-employment shocks are sensitive to the kurtosis and tails of the histogram. Many of the parameters have a straightforward connection to the moments. For example, the

[19] To conserve space, we exclude the estimate of $\bar{q}$which simply targets the 80th percentile of the income distribution.

parameters governing the slopes of the factor structure below (α_1) and above (α_2) the 80th percentile are most sensitive to the respective factor structure slope moments.

Figure 12 – Sensitivity Analysis

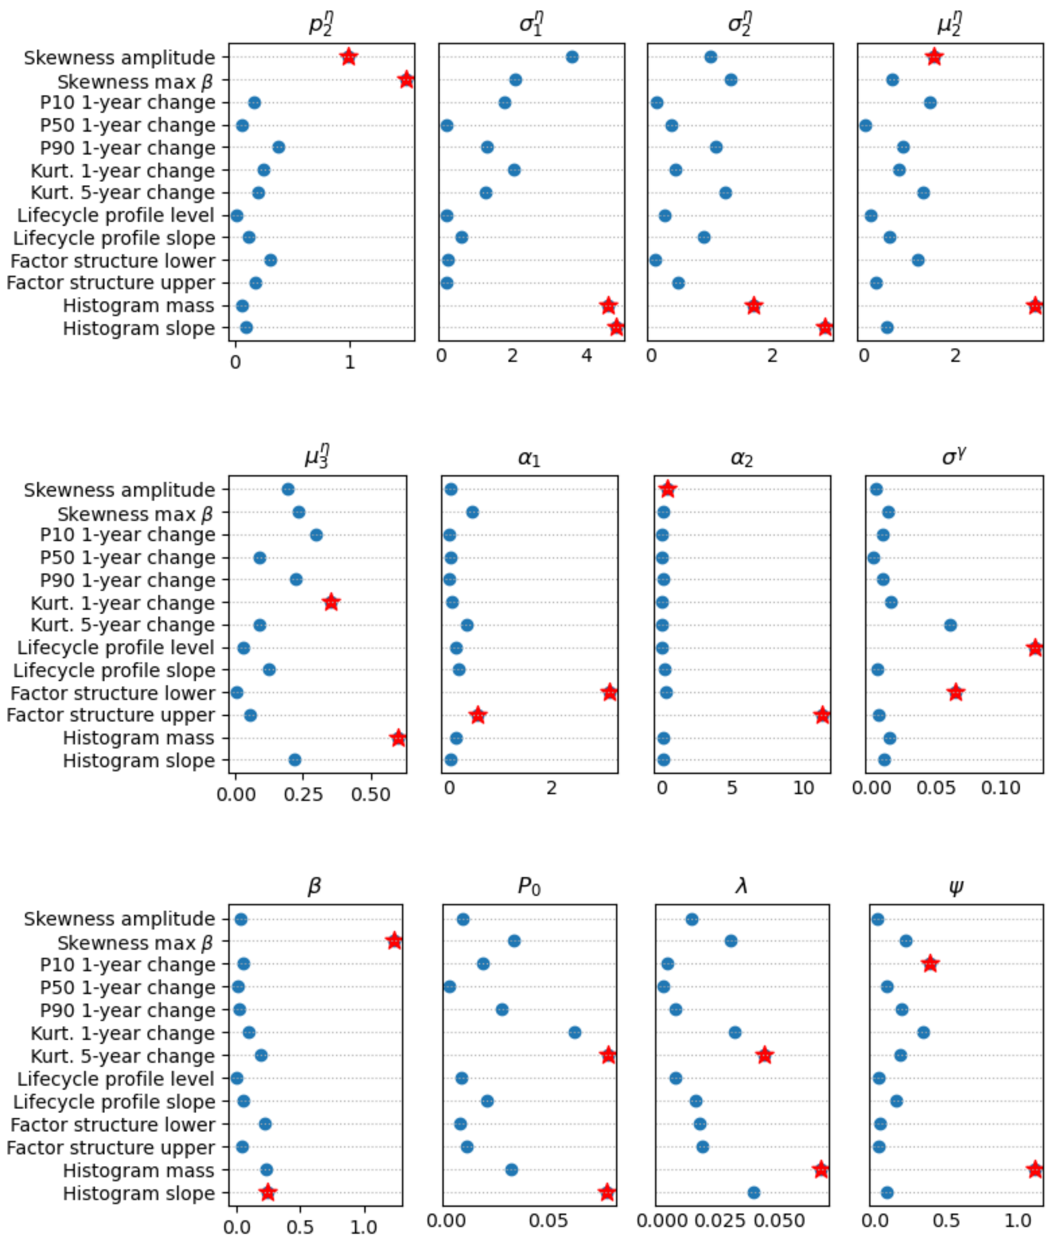

Note: This figure shows the results of a sensitivity analysis following Andrews et al. (2017). Each figure displays the absolute value of 13 directional derivatives of the parameters with respect to percent deviations in groups of moments described on the y-axis. See Appendix Section B for more detail. The stars indicate the two groups of moments to which each parameter is the most sensitive.

7 Application in heterogeneous-agent models

We now provide 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.[20] 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'. [21]

We now present results from a simple application along the lines above. We define S as the scalar aggregate business cycle component w_t, which we assume follows an an

[20]The main alternative to quadrature is to discretize the income process using a Markov chain. Kirkby (2025) discusses methods for doing so in the context of income processes with innovations drawn from a normal mixture distribution. The challenge in applying that approach to our application is to incorporate the time-varying aspect of the innovation distribution.

[21]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).

AR(1) process. Fitting an AR(1) to our empirical measure of w_t (after removing a cubic trend) results in persistence 0.913 and innovation standard deviation 0.0254. We assume a perpetual-youth structure in which households have a survival probability of 35/36. We assume constant relative risk aversion preferences with a coefficient of relative risk aversion of 1. Households save in an annuity with an expected return of 2 percent per year. We include a simple social insurance system that provides a minimum income at 6% of the average income. We calibrate the discount factor β so that the aggregate wealth in the stationary equilibrium of the economy is 3 times aggregate income, which determines β. We solve the model with the endogenous grid method (Carroll, 2006).

We solve this consumption-savings problem four times with the income process specified as in Models 1-4 above. We then simulate a time series of aggregate consumption by aggregating the consumption decisions of a simulated population of households. As an input to the simulation, we feed in a time path of w_t such that the simulation recreates the observed time series of aggregate labor income as measured by detrended real wage and salary compensation from the National Income and Product Accounts (NIPA).[22] Figure 13 shows detrended real aggregate consumption of non-durable goods and services from NIPA. Each panel of the figure shows the simulation of our consumption-savings problem with a different income process. The results show that Models 3 and 4 give a remarkably good fit to the observed aggregate consumption data while Models 1 and 2 predict aggregate consumption that is quite a bit smoother than the observed data. The empirical standard deviation of aggregate consumption is 0.020, while the predicted standard deviations are 0.008, 0.009, 0.016, 0.015, for Models 1-4 respectively. The predictions for all four models have a correlation with the data between 0.89 and 0.91.

The results of this simulation study suggest that time-varying risk and precautionary savings motives can contribute meaningfully to aggregate consumption dynamics. One should be a bit cautious about interpreting the results of this simulation as saying Models 3 and 4 are superior to Models 1 and 2 because there are other forces affecting the economy besides just fluctuations in aggregate income; for example changes in interest rates. As we omit those forces, one would not expect our simulation to match aggregate consumption even if the income process were specified perfectly.

[22]We choose a path for $w_t$that matches aggregate income in the simulation of Model 4, we then keep this path of $w_t$fixed as we consider Models 1-3. In all cases where we detrend aggregate time series we remove a cubic time trend.

Figure 13 – Observed and Simulated Aggregate Consumption

(a) Model 1: Gaussian Model

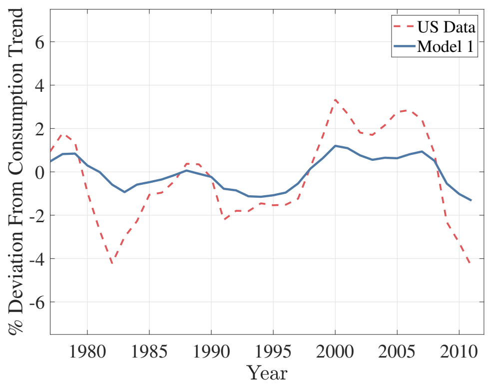

(b) Model 2: Add Nonemployment

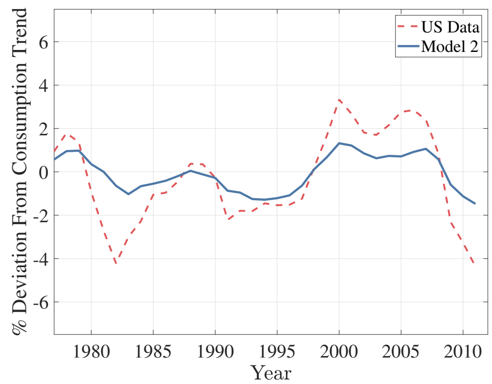

(c) Model 3: Add Time-Varying Shocks

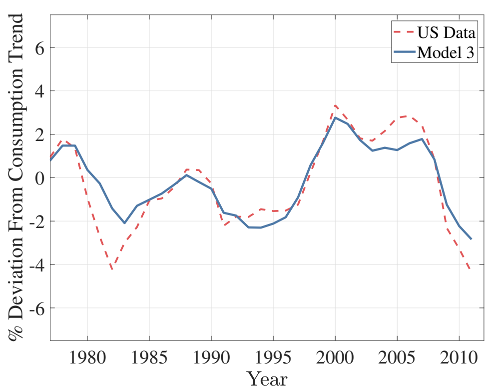

(d) Model 4: Add Factor Structure

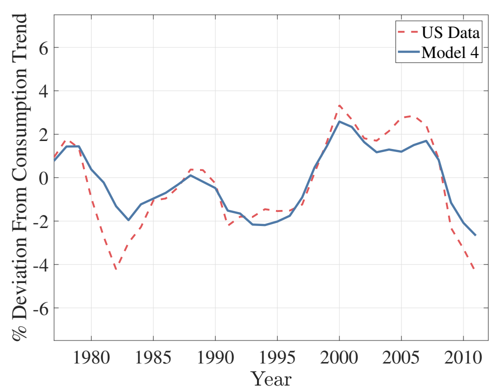

Note: This figure shows a simulation of aggregate consumption by combining each model of the income process as specified in Table I with a partial-equilibrium consumption-savings problem. The dashed line is the detrended data constructed by aggregating non-durable goods and service consumption from NIPA and the solid line is the simulated aggregate consumption in each model.

8 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

Almuzara, M., M. Arellano, R. W. Blundell, and S. Bonhomme (2025): "Nonlinear micro income processes with macro shocks," *FRB of New York Staff Report*.

Andrews, I., M. Gentzkow, and J. M. Shapiro (2017): "Measuring the Sensitivity of Parameter Estimates to Estimation Moments*," *The Quarterly Journal of Economics*, 132, 1553–1592.

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.

Carroll, C. D. (2006): "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economics Letters*, 91, 312–320.

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.

Geweke, J. and M. Keane (2000): "An Empirical Analysis of Earnings Dynamics Among Men in the PSID: 1968-1989," *Journal of Econometrics*, 96, 293–356.

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.

Hubmer, J. (2018): "The job ladder and its implications for earnings risk," *Review of Economic Dynamics*, 29, 172–194.

Huckfeldt, C. (2022): "Understanding the scarring effect of recessions," *American Economic Review*, 112, 1273–1310.

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

Kirkby, R. (2025): "Discretizing earnings dynamics: implications of Gaussian-mixture shocks for life-cycle models," *The Japanese Economic Review*, 76, 493–519.

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 $P^{Mod}_{90,5}$ be the model-implied value. We then form

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

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

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

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

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

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

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

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

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

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 Sensitivity Test

Let θ represent the parameter vector and let $g(\theta)$ represent the vector of 133 moments. The estimator can be written as

$$\hat{\theta} = \arg\min_{\theta} g(\theta)' W g(\theta)$$

where W is a weighting matrix. The first order condition is given by

$$G(\theta)'Wg(\theta) = 0$$

where

$$G(\theta) \equiv \frac{\partial g(\theta)}{\partial \theta}.$$

The objective of the sensitivity test is to determine how perturbing the moment conditions would affect the estimate of θ. Andrews et al. (2017) show that, locally, a perturbation of the moments from $g(\theta)$ to $g(\theta) + dg$ will cause a shift in the estimated parameters $\hat{\theta}$ given by

$$\Lambda = -\left(G'WG\right)^{-1} G'W.$$

In the most flexible specifications (i.e. Models 4 and 6), there are 13 parameters and 133 moments. Thus, Λ is a 13×133 matrix in which each element corresponds to the derivative of one parameter with respect to one moment. The units of Λ reflect the scale of the parameters and the moments. For comparability, we scale the parameters based on the absolute value of the point estimate. Where appropriate, we similarly scale the empirical moments. As a result, each deviation we consider represents a marginal percentage change in the moments.[23] The elements of Λ reflect percentage changes in the parameter estimates with respect to percent changes in the moments.

Rather than report all 133 moments, we report derivatives in a few directions of interest in Figure 12. Specifically we report Λx for various x that satisfy $||x|| = 1$. The values of x are described below.

- Skewness: The raw moments are the time-varying Kelley skewness of 1-year, 3-year, and 5-year earnings growth. We consider two summary measures of these time series. The first ("Skewness amplitude") considers increasing the amplitude of the fluctuations in skewness. Specifically, we multiply the (demeaned) skewness time series by a constant and take the derivative with respect to that constant. The second ("Skewness max-β") is to compute the unit-length deviation in the skewness time series that would have the greatest impact (according to Λ) on the estimate of the aggregate wage loading in the time-varying mean of permanent income risk, β.

[23] We do not scale the empirical moments that are near zero. In practice, we scale Λ as $\text{Diag}(\hat{\theta})^{-1}\Lambda W^{-1}$, where W is constructed as explained in our estimation procedure.

- Percentiles: We consider the 10th, 50th, and 90th percentiles of the 1-year earnings growth rate as three separate directions.
- Kurtosis: Similarly, we consider the kurtosis of 1-year and 5-year earnings growth as two separate directions.
- Cross-sectional Variance: The life-cycle variance profile is captured by four raw moments, the cross-sectional variance at ages 25, 35, 45, and 55. We consider two deviations: increasing all four uniformly ("Lifecycle profile level") and increasing the slope in a linear fashion ("Lifecycle profile slope").[24]
- Factor Structure: The raw moments include the slopes of the factor structure below and above the 80th percentile for each business cycle episode. "Factor structure lower" is the unit-length deviation in the slopes below the 80th percentile that would have the largest impact on α_1. "Factor structure upper" is the unit-length deviation in the slopes above the 80th percentile that would have the largest impact on α_2.
- Histogram: We consider two deviations corresponding to the two types of moments concerning the tails of the histogram. The histogram moments include the mass in each tail, and we consider a deviation that increases both of these masses simultaneously ("Histogram mass"). The histogram moments also include the slopes in each tail. We consider a deviation that makes both tails steeper ("Histogram slope").

C Results from Additional Specifications

[24] The slope deviation increases the variance at age τ by $d(\tau - 25)$ where d is a constant such that $||x|| = 1$.

Table IV – Estimated Parameters

Parameters		(2')	(3')	(3")	(4')
σ_γ	St. dev. of fixed effects	0.639	0.604	0.600	0.607
p^ζ	Probability of full-year empl.	0.509	0.506	–	0.631
λ	Transitory exponential parameter	3.681	3.483	3.103	2.977
ψ	Scarring effect of transitory shock	–	0.160	0.163	0.166
p_2^η	Mix. probab. for persist. innov 2	–	0.122	0.094	0.042
p_3^η	Mix. probab. for persist. innov 3	–	0.122	0.094	0.042
σ_1^η	Std. dev. for persistent innov. 1	0.119	0.025	0.035	0.004
σ_2^η	Std. dev. for persistent innov. 2	–	0.117	0.126	0.205
σ_3^η	Std. dev. for persistent innov. 3	–	0.117	0.126	0.205
μ_1^η	Center for persistent component 1	–	0.024	0.021	0.011
μ_2^η	Center for persistent component 2	–	-0.063	-0.054	-0.107
μ_3^η	Center for persistent component 3	–	0.211	0.232	0.344
α_1	Factor struct. slope, low income	–	–	–	–0.905
α_2	Factor struct. slope, high income	–	–	–	.565
$\overline{q}$	Factor structure threshold	–	–	–	0.921
β_0^ξ	Employment prob, constant	–	–	-0.249	–
β_1^ξ	Employment prob, agg wage loading	–	–	-8.788	–

Notes: We estimate four supplemental specifications which represent small deviations from the baseline estimates discussed in the main body of the paper. Specification 2' is identical to specification 2, but without a scarring effect. Specification 3' and 4' are identical to specifications 3 and 4 but without allowing the normal mixture distribution for the persistent income shock to vary by time. Specification 3" is similar to specification 3', but allows the non-employment shock process to be time-varying instead of the normal mixture.

Table V – Targeted and Fitted Moments

Moments	US Data	(2)	(2')	(3)	(3')	(3")	(4)	(4')
		Model Specifications						
P10, one-year change	–0.434	–0.438	-0.453	-0.418	-0.450	-0.453	-0.420	-0.420
P10, three-year change	–0.585	–0.485	-0.501	-0.476	-0.496	-0.498	-0.475	-0.474
P10, five-year change	–0.631	–0.524	-0.540	-0.501	-0.530	-0.530	-0.498	-0.508
P50, one-year change	0.020	0.011	0.002	0.019	0.012	0.015	0.019	0.018
P50, three-year change	0.061	0.015	0.007	0.026	0.013	0.017	0.027	0.030
P50, five-year change	0.103	0.030	0.023	0.046	0.028	0.031	0.044	0.041
P90, one-year change	0.474	0.431	0.457	0.427	0.440	0.447	0.447	0.425
P90, three-year change	0.705	0.493	0.515	0.514	0.511	0.513	0.525	0.509
P90, five-year change	0.848	0.566	0.586	0.605	0.584	0.585	0.620	0.583
Kurtosis, one-year chg	20.00	23.40	22.759	24.247	22.453	22.452	24.233	24.045
Kurtosis, five-year chg	12.00	17.61	17.164	17.135	16.616	16.942	17.306	17.739
Cross-sect. var., age 25	0.595	0.555	0.570	0.545	0.541	0.554	0.562	0.570
Cross-sect. var., age 35	0.719	0.692	0.709	0.695	0.683	0.695	0.712	0.708
Cross-sect. var., age 45	0.814	0.834	0.851	0.838	0.830	0.844	0.850	0.848
Cross-sect. var., age 55	0.905	0.969	0.987	0.958	0.974	0.986	0.981	0.982
Left-tail mass	0.024	0.023	0.022	0.023	0.024	0.024	0.023	0.025
Left-tail slope	1.260	1.292	1.318	1.299	1.282	1.250	1.308	1.301
Right-tail mass	0.015	0.020	0.022	0.018	0.018	0.019	0.018	0.019
Right-tail slope	–2.035	–1.538	-1.420	-1.599	-1.606	-1.598	-1.592	-1.538
Objective value		2.333	2.451	1.544	2.074	1.884	1.176	1.853
Quantiles		0.279	0.229	0.252	0.240	0.234	0.235	0.268
Kurtosis		0.247	0.204	0.228	0.163	0.185	0.240	0.270
Cross-sect. var. profile		0.116	0.122	0.125	0.170	0.154	0.122	0.110
Histogram		0.444	0.703	0.277	0.279	0.299	0.305	0.375
Skewness time series		0.749	0.691	0.160	0.722	0.517	0.154	0.640
Factor Structure		0.498	0.502	0.502	0.500	0.495	0.119	0.190

Notes: This table shows the model fit for each estimated supplemental models described in Table IV, along with baseline specifications 2, 3, and 4 for comparison. 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.