A Parsimonious Income Process for Business Cycle Analysis

Fatih Guvenen* Alisdair McKay†

February 15, 2016

Preliminary and Incomplete

Abstract

In this paper, we estimate a parsimonious income process that is consistent with several key features of how income risk varies over the business cycle. In particular, the estimated process generates year-to-year income changes that (i) have flat and acyclical variance, (ii) have volatile and procylical skewness, (iii) have very high excess kurtosis, and (iv) imply a moderate rise in cross-sectional inequality over the life cycle consistent with the US data. Furthermore, and importantly, the process also captures the predictable nature of business cycle income risk: income changes during a business cycle episode is partly predicted by income levels before that episode. The estimated process features a mixture of normals as well as a factor structure, both of which are driven by a latent process capturing the business cycle.

JEL Codes: D31, E24, E32, H31
Keywords: Idiosyncratic income risk, business cycles, skewness, kurtosis.

*University of Minnesota, FRB of Minneapolis, and NBER; guvenen@umn.edu; www.fguvenen.com
†Boston University; amckay@bu.edu

1 Introduction

While business cycles show modest fluctuations in average incomes, there are considerable changes in the distribution of income changes. Cyclical changes in the income risks that households face have been shown to have important implications for a variety of questions in macroeconomics and finance. For example, Krebs (2007) and De Santis (2007) show the importance of cyclical income risk for the welfare cost of business cycles; Constanides and Duffie (1996), Storesletten et al (2007), Constanides and Ghosh (2014), and Schmidt (2015) show that cyclical income risk can explain several features of asset prices; McKay (2015) shows that time-varying idiosyncratic risk can generate large movements in aggregate consumption; Braun and Nakajima (2012) and Dew-Becker et al (2015) show that cyclical risk has important implications for the conduct of monetary policy.

Recently, Guvenen et al (2014) have documented a variety of facts on the nature of cyclical income risks using data from the Social Security Administration. Those authors emphasize that it is the skewness of the distribution that appears to be procyclical as opposed to the more traditional view of a coutner-cyclical skewness. The first contribution of this paper is to present an estimated income process that is consistent with a variety of evidence from the Social Security Administration data including the cyclicality of earnings risk. We hope that these estimates will facilitate future work on the quantitative analysis of the implications of cyclical income risk.

The Social Security data show that low-income and very high-income households are more exposed to the business cycle than those in the center of the income distribution. A second contribution of the paper is to introduce an income process with this type of factor structure in earnings. With such an income process, the business cycle has considerably different distributional implications, which will introduce a new channel in the analysis of the welfare cost of business cycles and the impact on aggregate consumption. These tdistributional issues are central to the design of optimal fiscal policy over the cycle as shown by Bhandari et al (2015).

Our work is related to Storesletten et al (2004), but our focus is on pro-cyclical skewness as opposed to counter-cyclical variance. Our work is also related to McKay (2015), but relative to that paper we introduce a different specification and new data such as the rise in inequality over the life-cycle and the kurtosis of earnings growth rates, which lead to rather different estimates. Furthermore we add the factor structure to the earnings risk.

2 First Income Process

2.1 Assumptions

We start with the income process from McKay (2015) in which earnings are given by

$$Y_{i,t} = \begin{cases} \exp[w_t + \theta_{i,t} + \xi_{i,t}] & \text{if } n_{i,t} = 1 \\ 0 & \text{if } n_{i,t} = 0 \end{cases} \tag{1}$$

where w_t is an aggregate component, $\theta_{i,t}$ is a persistent idiosyncratic component, $\xi_{i,t}$ is a transitory idiosyncratic component, and $n_{i,t}$ is an indicator for whether the individual is employed in this period. $\theta_{i,t}$ evolves as a random walk

$$\theta_{i,t} = \theta_{i,t-1} + \eta_{i,t} \tag{2}$$

where $\eta_{i,t}$ is drawn from a mixture of normals

$$\eta_{i,t} \sim \begin{cases} N(\mu_{1,t}, \sigma_1) & \text{with prob. } p_1, \\ N(\mu_{2,t}, \sigma_2) & \text{with prob. } p_2, \\ N(\mu_{3,t}, \sigma_3) & \text{with prob. } p_3, \\ N(\mu_{4,t}, \sigma_4) & \text{with prob. } p_4. \end{cases}$$

The first component of this mixture will control the center of the distribution, the second and third components will control the mass in the right and left tails, respectively, and the fourth component will control the kurtosis of the distribution. The tails move over time as driven by the latent variable x_t such that

$$\begin{aligned} \mu_{1,t} &= \bar{\mu}_t, \\ \mu_{2,t} &= \bar{\mu}_t + \mu_2 - x_t, \\ \mu_{3,t} &= \bar{\mu}_t + \mu_3 - x_t, \\ \mu_{4,t} &= \bar{\mu}_t. \end{aligned}$$

$\bar{\mu}_t$ 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 the mean income.

The transitory component is distributed $\eta_{i,t} \sim N(-\sigma_\xi^2/2, \sigma_\xi)$ and employment status is first-order Markov with transition probabilities λ_t for the job-finding rate between period t-1 and t and ζ_t the corresponding job-separation rate.

2.2 Moments and first estimates

The model period is one quarter. The parameters are selected to match year-by-year the median earnings growth, the dispersion in the right tail (90 - 50), and the dispersion in the left-tail (50-10) for one, three, and five year earnings growth rates. These are 279 moments (three moments times 33 observations for one-year differences, 31 observations for three-year differences, and 29 observations for five-year differences).

The procedure simulates quarterly data using the observed job-finding and -separation rates and then aggregates to annual income and computes these moments. To simulate the income process, we require time series for x_t and w_t. We assume that these series are linearly related to observable labor market indicators. The time-varying risk factor x_t is assumed to be linearly related to four quarterly labor market series: the fraction of the labor force unemployed with duration less than 15 weeks (this is closely related to unemployment inflows), the fraction unemployed with durations over 15 weeks (related to the job-finding rate), an index of average weekly hours, and the labor force participation rate. We express these four series in their principle components and denote the factor loadings with $\{\beta_i^x\}_{i=1}^4$.

The time series for w_t is assumed to be of the form

$$w_t = \beta_0^w + \beta_1^w \lambda_t + \beta_2^w \zeta_t + \beta_3^w m_t,$$

where m_t is the mean one-year income growth for the year containing quarter t scaled by 0.25.

The restrictions $p_2 = p_3$ and $\sigma_2 = \sigma_3$ are imposed in all our estimates. We begin by also including the restriction $p_4 = 0$. We simulate 20 years for burn in with $w_t = x_t = 0$ and λ_t and ζ_t equal to their sample averages. Our focus is on the 35 years 1977 to 2011 and most of our moments concern differences from t to $t + k$ where t ranges from 1978 to 2010, which is 33 years. The first row of Table I has the fitted parameter values.

Table I: Fitted parameters. * μ_4 is constrained to 0.

	$p_2 = p_3$	p_4	σ_ξ	σ_1	$\sigma_2 = \sigma_3$	σ_4	μ_2	μ_3	μ_4
Version 1	0.0424		0.261	0.00698	0.192		0.339	–0.309	
Version 2	0.00227		0.329	0.0479	0.159		0.222	-0.599	
Version 3	0.00727	6.27e-5	0.354	0.0403	0.0966	3.574	0.266	-0.184	0*
Version 4	0.00969	1.68e-4	0.293	0.0469	0.135	3.267	0.249	-0.240	0*
Version 5	0.00751	9.69e-5	0.327	0.0399	0.0757	3.376	0.219	-0.150	0*

	st. dev.		skew.		kurt.	
	model	data	model	data	model	data
One-year	0.410		-0.083		7.13	
Three-year	0.528		-0.037		5.05	
Five-year	0.617		-0.001		3.57	

Table II: Moments of stationary earnings growth rate distributions.

2.3 Implied moments

Table II shows moments of the one-, three-, and five-year growth rates. These are computed by simulating quarterly data, aggregating to annual observations and then taking appropriate differences.

We now look at the evolution of the cross-sectional variance over the life cycle. To be clear, there is no life cycle component of this income process. We consider the cross-sectional variance of log income year by year for a cohort who start in the first quarter of the first year with $\theta_{i,t} = 0$ for all i. The income process implies a very large increase in variance over 35 years equal to 1.7 log points.

We now show how the income process captures the change in risk over the business cycle. Figure 1 shows the fit to the mean, median, and the 10th and 90th percentiles year by year. Figure 2 shows the evolution of Kelly's skewness over the cycle. and Figure 3 shows the evolution of the unemployment rate. Overall the fit to is pretty good although the standard deviations are a bit on the low side.

2.4 Summing up

The income process does pretty well in matching the cyclical properties of risk, but it over predicts the rise in the cross-sectional variance over the life-cycle and it does not generate enough Kurtosis.

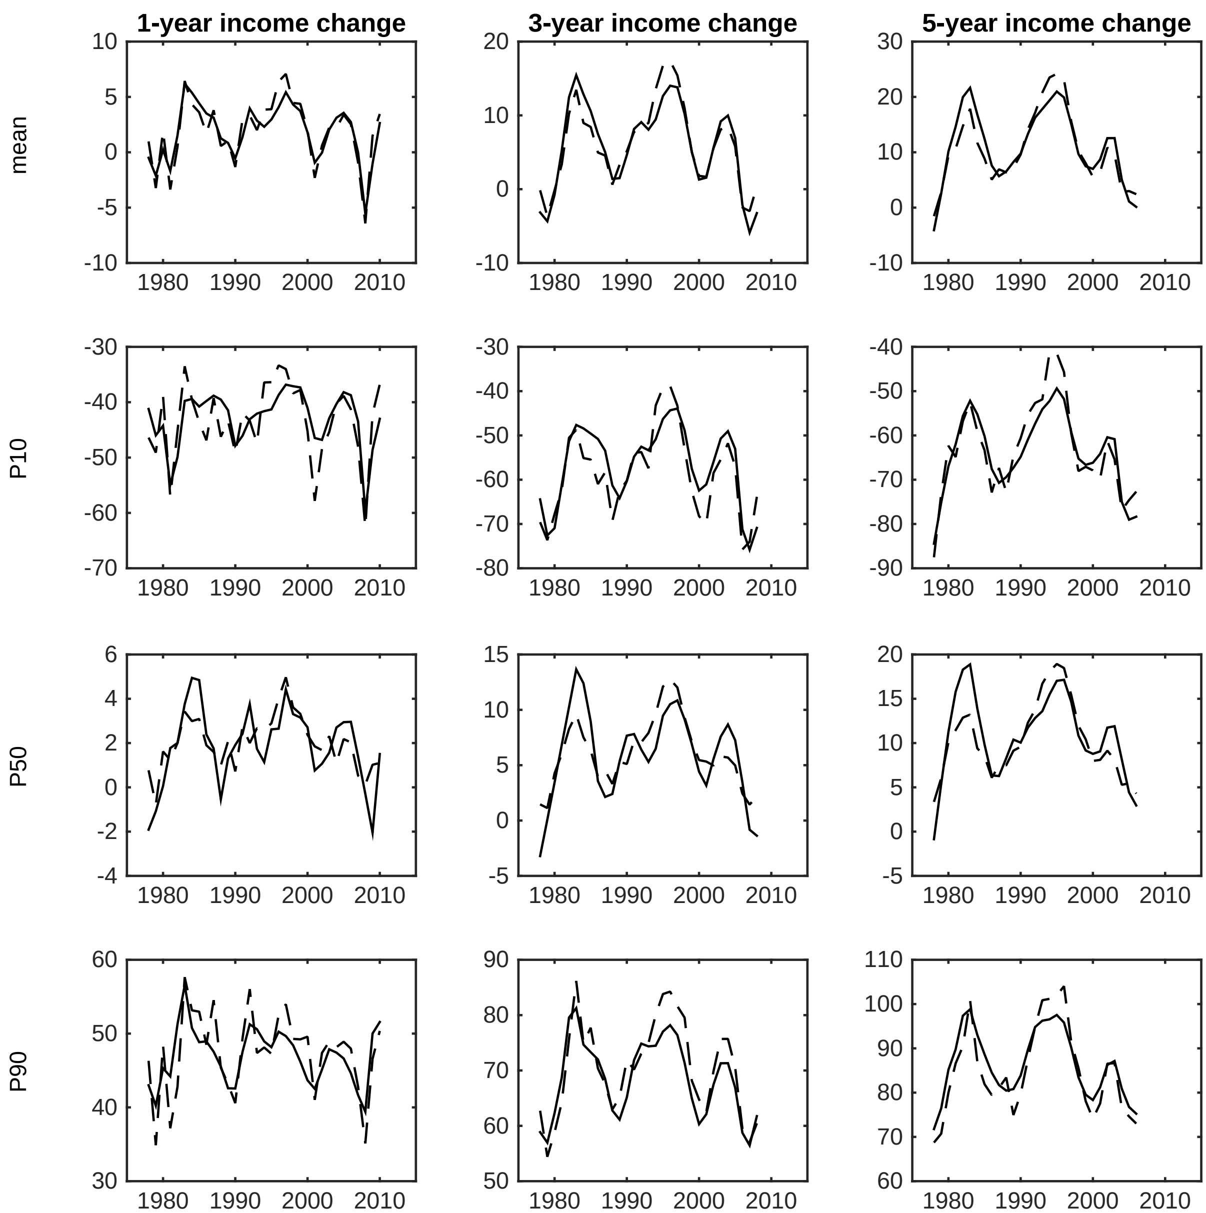

Figure 1: Moments of the earnings growth distributions. Solid = simulated, dashed = data.

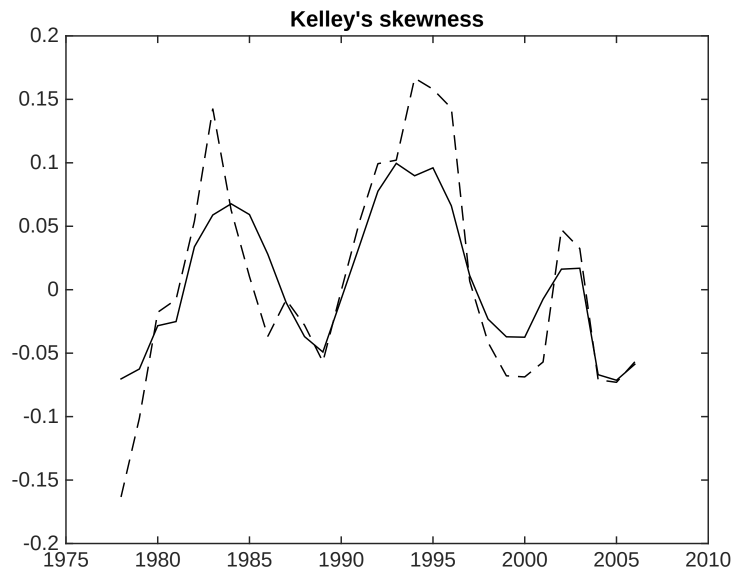

Figure 2: Kelly's skewness. Solid = simulated, dashed = data.

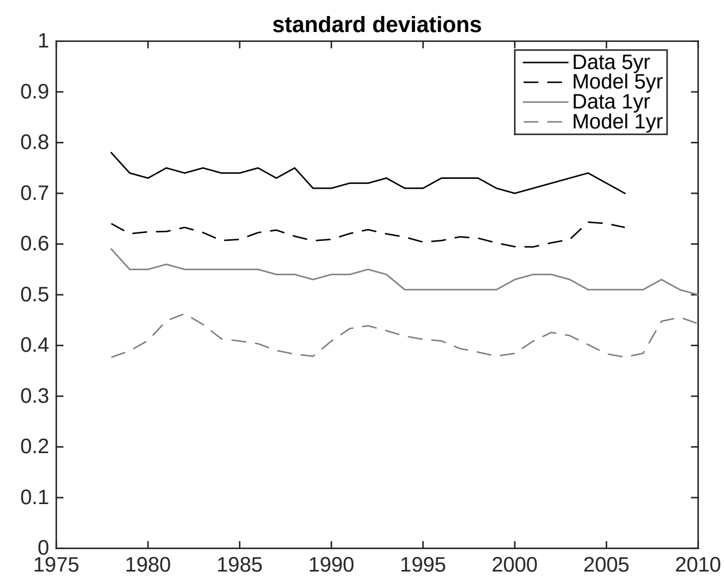

Figure 3: Standard deviation of earnings growth distribution.

3 Adding kurtosis and life-cycle variance as a moments

In addition to the 279 moments described above, we now also target an increase in the cross-sectional variance of earnings over 35 years equal to 0.45 (see Figure A2 in Guvenen et al, 2015), a kurtosis of 20 for one-year earnings changes, and a kurtosis of 12 for five-year earnings changes (see Figure 10 in Guvenen et al, 2015). We compute kurtosis of one-year (five-year) changes for the 33 (29) years and average the resulting values across years to arrive at a single kurtosis value for one-year changes and another for five-year changes.[1] The fitted parameters appear as "Version 2" in Table I.

3.1 Implied moments

Table III shows moments of the one-, three-, and five-year growth rates.

The income process implies a increase in variance over 35 years equal to 0.46 log points. Figure 4 shows the fit to the mean, median, and the 10th and 90th percentiles year by year. Figure 5 shows the evolution of Kelly's skewness over the cycle. and Figure 6 shows the evolution of the unemployment rate. The fit has deteriorated relative to version 1. Figure 7 shows a histogram of one-year earnings growth rates [computed with $w_t = x_t = 0$]. The histogram shows that the model generates too much dispersion near the center and too little dispersion in the tails.

	st. dev.		skew.		kurt.	
	model	data	model	data	model	data
One-year	0.471		-0.452		16.0	
Three-year	0.620		-0.846		18.0	
Five-year	0.731		-0.875		12.7	

Table III: Moments of stationary earnings growth rate distributions.

[1] All deviations of moments from their targets are expressed in percentage terms with the medians expressed as percent of the 90th percentile. We give the increase in the cross-sectional variance a relative weight of 93 to have comparable weight as the other moments for which we use a time series. Similarly we give the two kurtosis moments relative weights of 93/2.

3.2 Summing up

Adding the additional moments greatly reduced the increase cross-sectional variance over the life cycle and it generated a more reasonable level of kurtosis (although still somewhat low for one-year changes). Also notice that the standard deviations have increased closer to the values for the data. On the other hand, the aspects of earnings risk seems to fit less well.

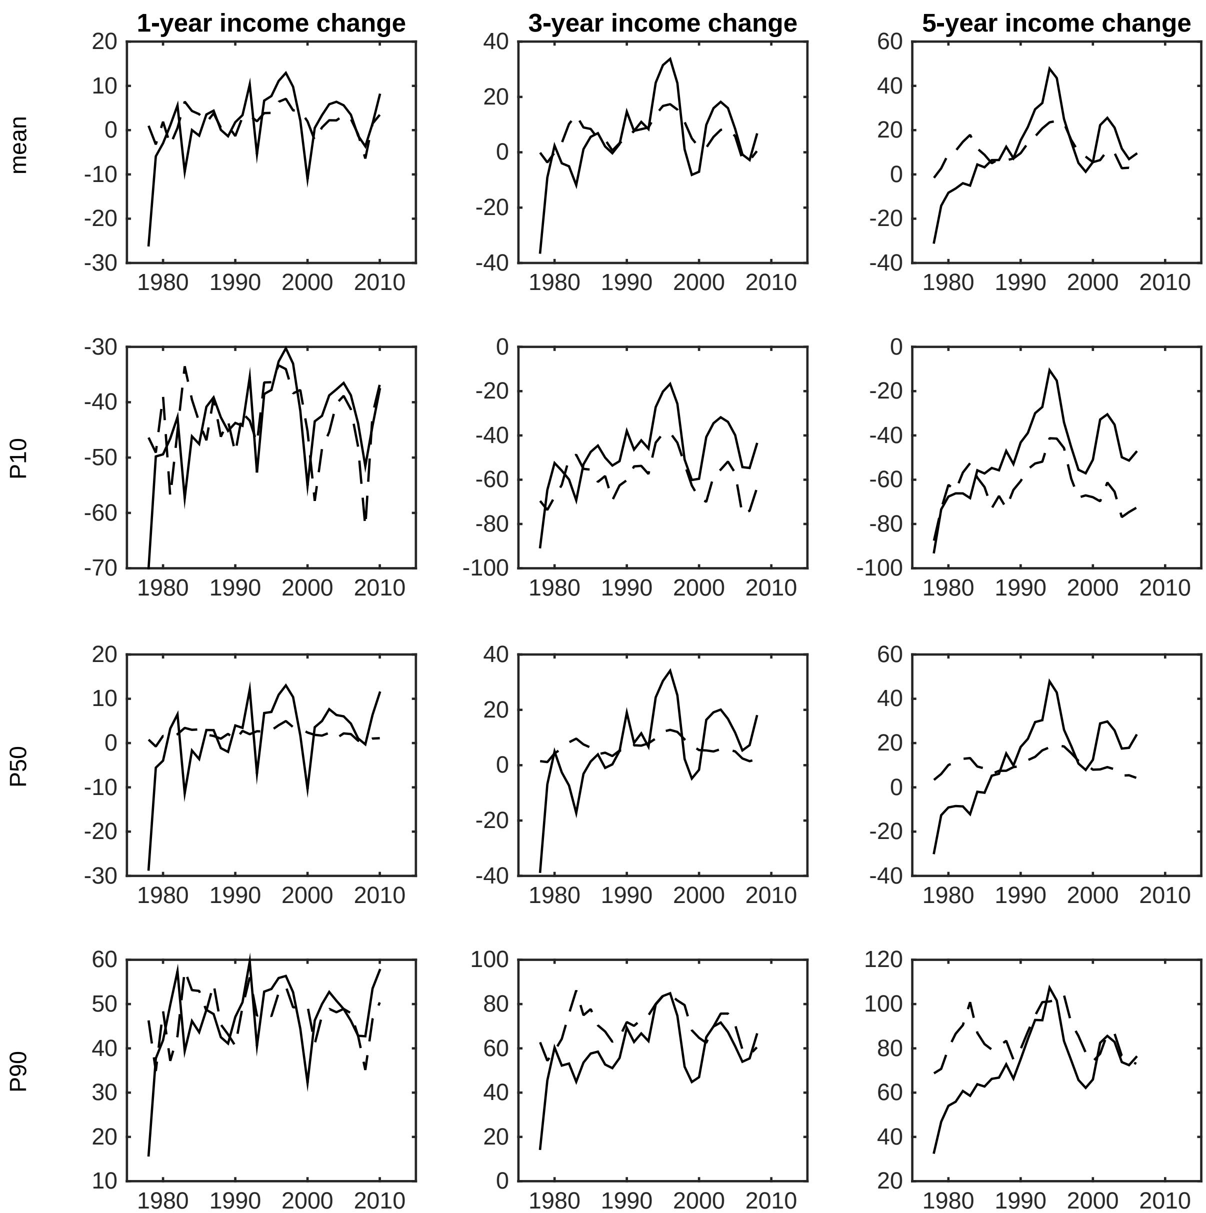

Figure 4: Moments of the earnings growth distributions. Solid = simulated, dashed = data.

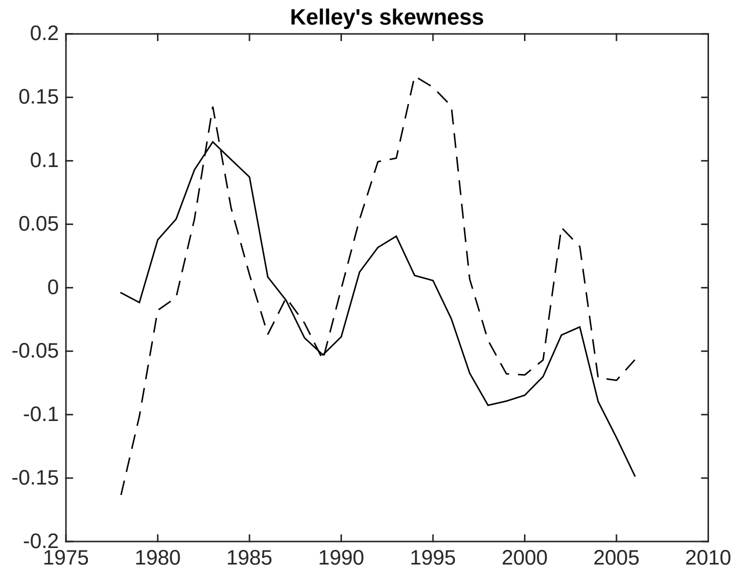

Figure 5: Kelly's skewness. Solid = simulated, dashed = data.

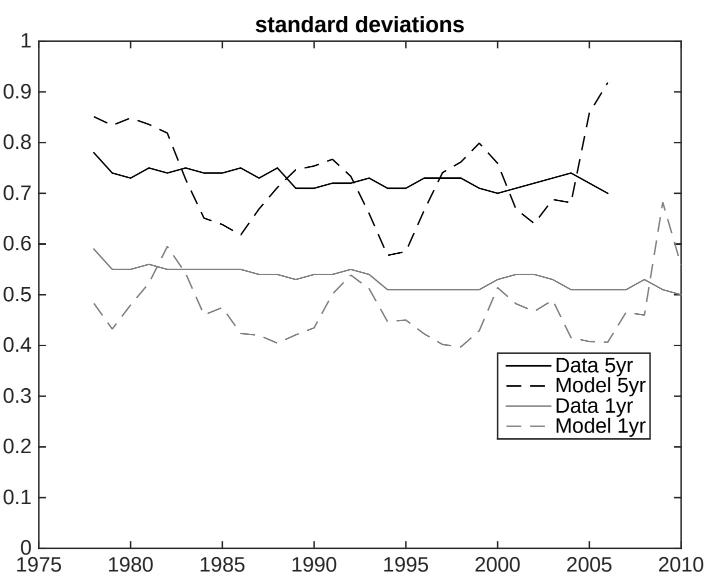

Figure 6: Standard deviation of earnings growth distribution.

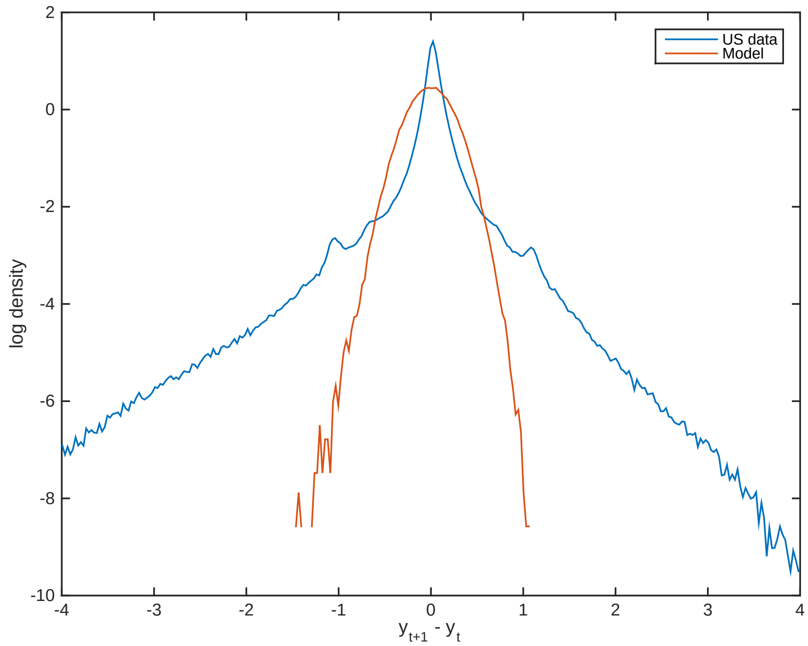

Figure 7: Histogram of one-year earnings growth rates.

4 Adding a fourth component to the mixture distribution

We target the same moments as in the previous version, but now we modify the model so that there is a fourth component of the mixture distribution for the distribution of $\eta_{i,t}$. That rationale behind this change is to allow the model to fit the kurtosis with some extreme innovations while still having the second and third components of the mixture distribution determine the 10th and 90th percentiles. Another change from the previous version is that we spent more effort searching for a global minimum of the objective function.

The estimated parameter values appear as "Version 3" in Table I. The fourth mixture component occurs with low probability and has a very large standard deviation. As the model can create kurtosis in this way, the second and third components do not need to create as much kurtosis. The probability of these components rises and their standard deviation falls.

4.1 Implied moments

Table IV shows moments of the one-, three-, and five-year growth rates. There is now more kurtosis as one might expect and the model almost hits the targets of 20 and 12 for the one-year and five-year kurtosis.

The income process implies an increase in variance over 35 years equal to 0.45855 log points, which is very close to the 0.45 target. Figure 8 shows the fit to the mean, median, and the 10th and 90th percentiles year by year. Figure 9 shows the evolution of Kelly's skewness over the cycle, and Figure 10 shows the evolution of the standard deviations. Overall, the fit is much better than version 2. In Figure 8, there is a bit too much volatility in the median, but perhaps not too much to worry about. Figure 9

	st. dev.		skew.		kurt.	
	model	data	model	data	model	data
One-year	0.46		0.93		20.40	
Three-year	0.56		0.54		16.59	
Five-year	0.63		0.43		11.92	

Table IV: Moments of stationary earnings growth rate distributions.

shows a good fit to Kelly's skewness. Figure 10 shows that there is still somewhat too much volatility in the standard deviations, but not as much as in version 2. Figure 11 shows a histogram of one-year earnings growth rates [computed with w_t and x_t from 1995-1996 to match the data]. The histogram is does better in the shoulders than version 2. One noticeable defect is that it is not as sharply peaked in the center as the empirical distribution.

4.2 Summing up

This is our prefered version among those without a factor structure. The two things that one might hope to improve are the fit to the time series of the standard deviations and the center of the histogram.

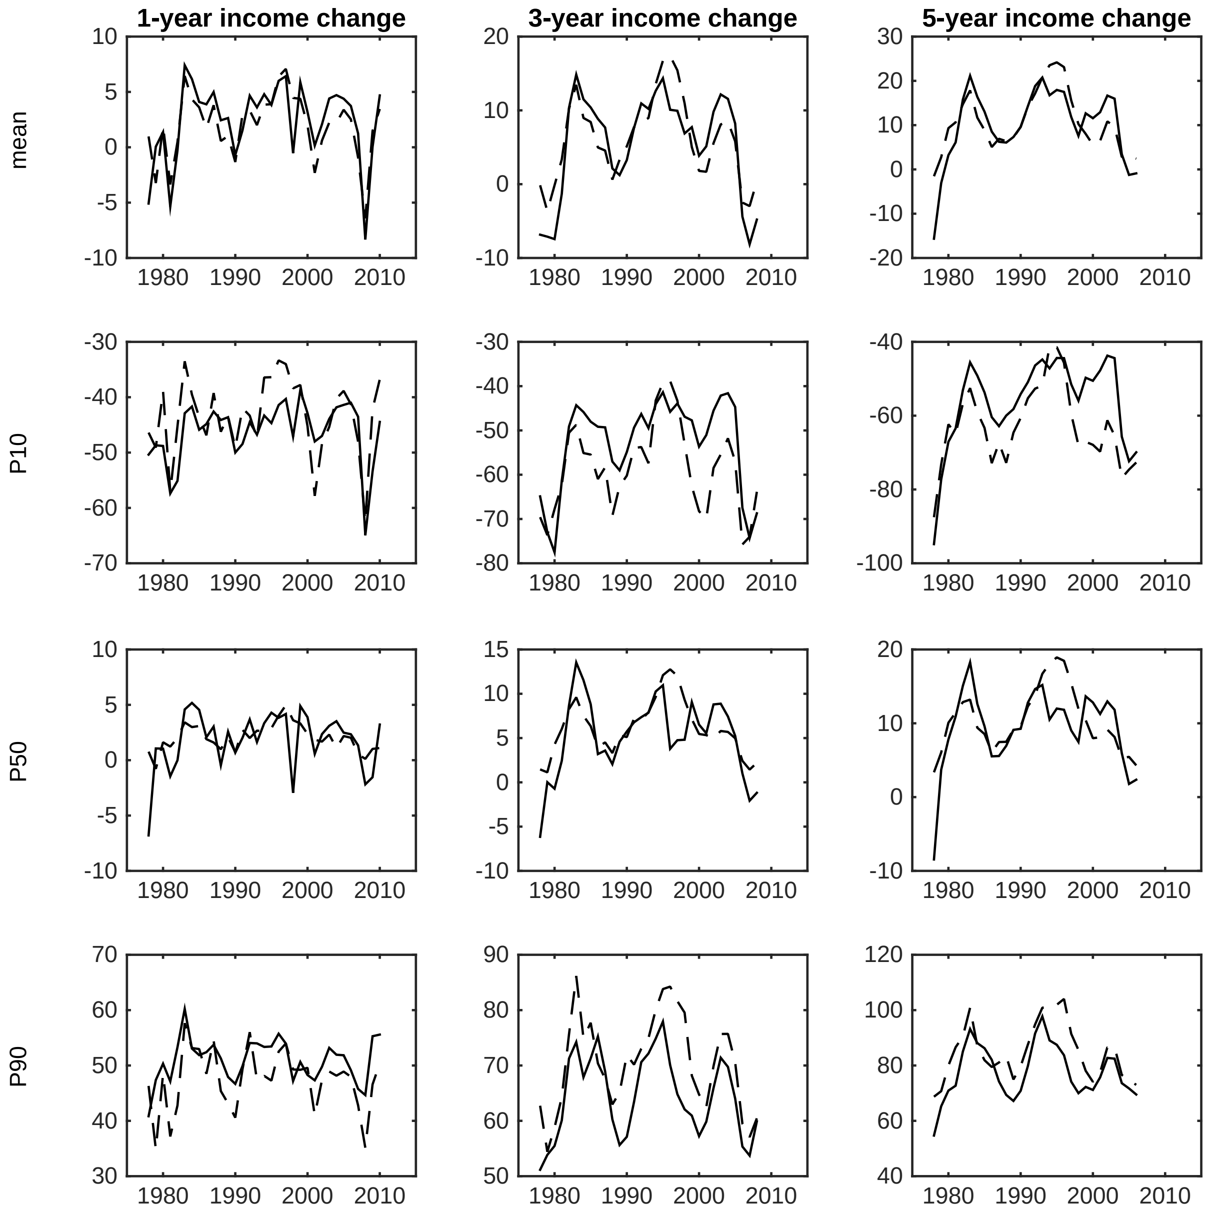

Figure 8: Moments of the earnings growth distributions. Solid = simulated, dashed = data.

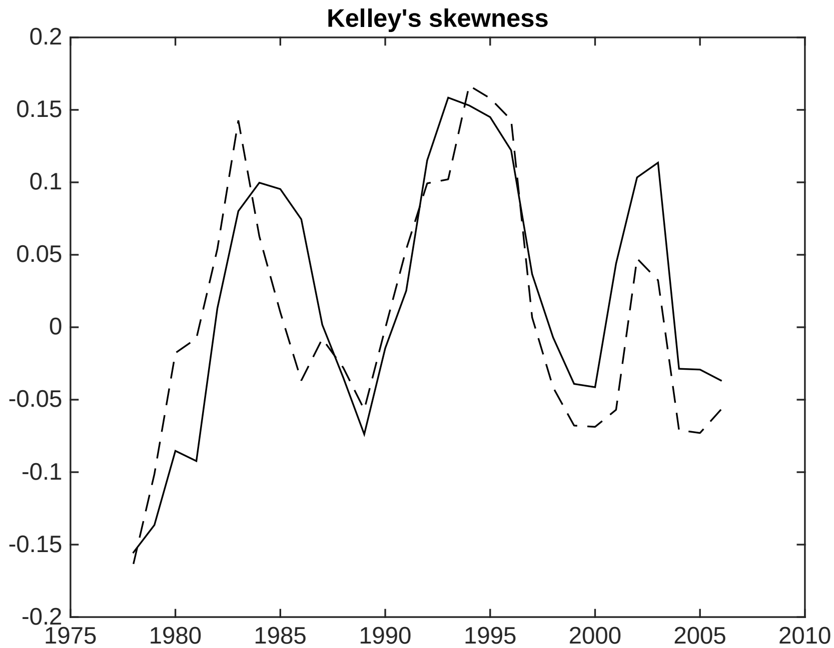

Figure 9: Kelly's skewness. Solid = simulated, dashed = data.

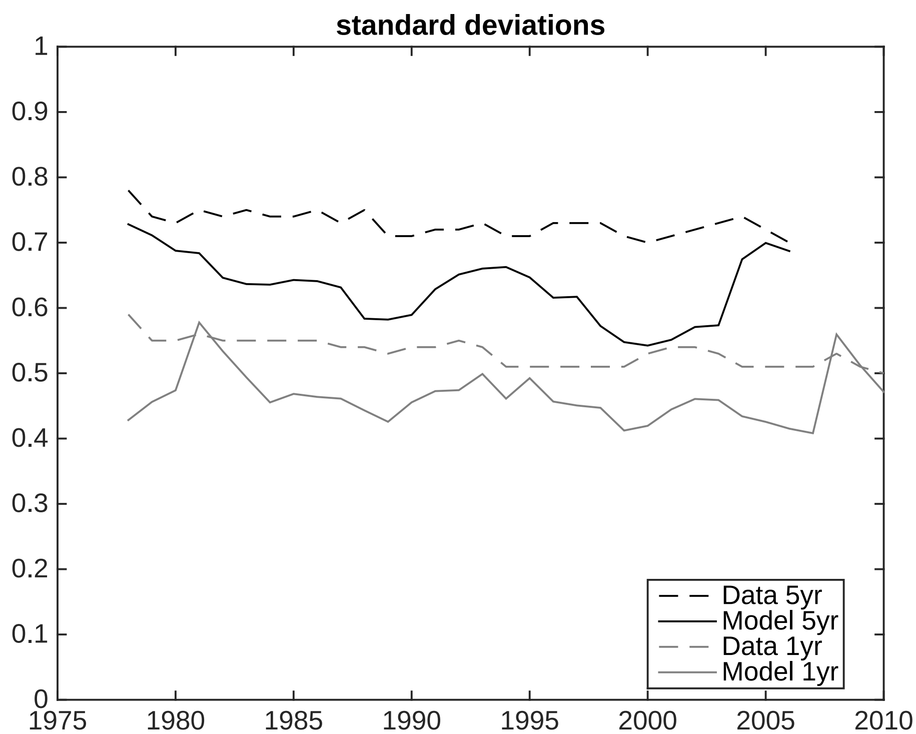

Figure 10: Standard deviation of earnings growth distribution.

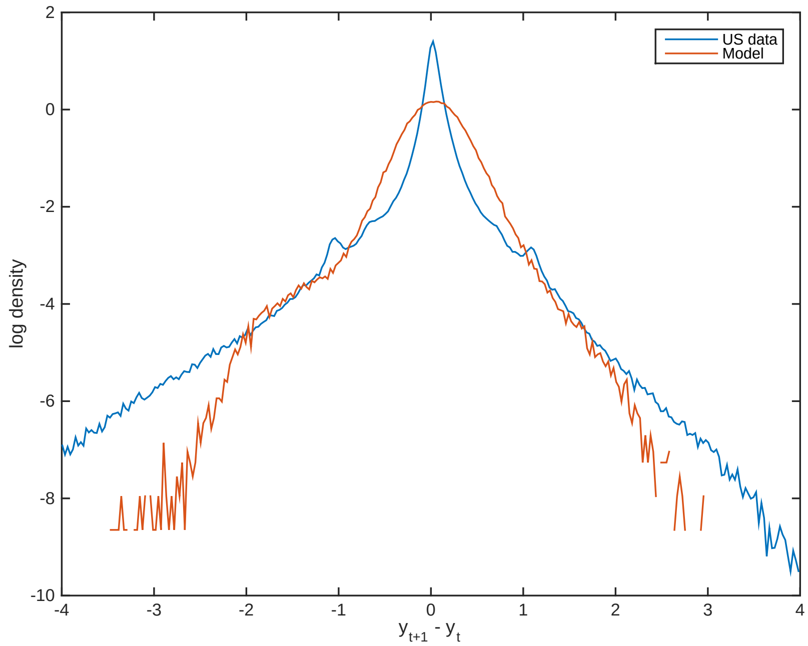

Figure 11: Histogram of one-year earnings growth rates.

5 Factor Structure (The path not taken)

We discussed a version in which income is given by

$$Y_{i,t} = \begin{cases} \exp[w_t + \theta_{i,t} + f(\theta_{i,t})z_t + \xi_{i,t}] & \text{if } n_{i,t} = 1 \\ 0 & \text{if } n_{i,t} = 0 \end{cases} \tag{3}$$

where $f(\cdot)$ is a cubic polynomial. We then intended to fit this using an indirect inference procedure where we fit a cubic polynomial to the factor strucutre. My attempts with this ran into the problem that results did not match the strong upward sloping tendency observed in the data. Fitting the behavior of the upper and lower tails seemed to be somewhat at odds with the fit to the middle of the distribution. I also experienced some problems related to explosive cubic behavior. For these reasons I decided to simplify to the version above.

6 Factor Structure

Households at different points in the earnings distribution show differential sensitivity to the business cycle (see Figures 13 and 14 in Guvenen et al 2014). We now replace 1 with a specification that can capture this factor structure in earnings

$$Y_{i,t} = \begin{cases} \exp[w_t + \theta_{i,t} + f(\theta_{i,t})w_t + \xi_{i,t}] & \text{if } n_{i,t} = 1 \\ 0 & \text{if } n_{i,t} = 0. \end{cases} \tag{4}$$

The function $f(\cdot)$ controls the factor structure. Suppose high income households gain more when average income increases–this would be represented with an increasing $f(\cdot)$. We use equation 3 with $f(\theta_{i,t})$ being piecewise linear. Specifically, we assume

$$f(\theta_{i,t}) = \begin{cases} \alpha_1\theta_{i,t} & \text{if } \theta_{i,t} < \bar{\theta} \\ \alpha_2(\theta_{i,t} - \bar{\theta}) + \alpha_1\bar{\theta} & \text{if } \theta_{i,t} \geq \bar{\theta}. \end{cases}$$

This piecewise-linear specficiation is intended to capture the fact that both low- and high-income earners show more cyclicality in earnings than those in the middle of distribution.

This factor structure has three parameters, α_1, α_2, and $\bar{\theta}$. We form the objective based on the slope coefficients of linear regressions of the factor structure between the 10th and 90th percentiles and between the 91st and 100th percentiles. This gives us eight moments (two slopes for each recession). In choosing $\bar{\theta}$ our intention is that $\bar{\theta}$ should coincide with the 90th percentile of the five-year earnings distribution. While the distribution of $\theta_{i,t}$ is highly correlated with the five-year earnings distribution this correlation is nonetheless imperfect. For a business cycle episode (expansion or recession), we form a metric that captures how closely $\bar{\theta}$ coincides with the 90th percentile of the five-year earnings distribution, which we call the classification error. To construct the classification error, we use the values of $\theta_{i,t}$ from the end of the episode and for each i we say there is an error if $\theta_{i,t} < \bar{\theta}$ and the individual is above the 90th percentile of the five-year earnings distribution or if $\theta_{i,t} \geq \bar{\theta}$ and the individual is below the 90th percentile. The classification error for this episode is the fraction of individuals that are wrongly classified. In addition our objective function above we add the sum of squared classification errors for the four recessions.[2]

We have data for the factor structure during four recessions. For a given episode starting in t_1 and ending in t_2 we rank households by their average incomes over the years $t_1 - 5$ to $t_1 - 1$. We then create 100 bins by splitting the population at each percentile. Within each bin we compute the mean of the log difference of incomes between t_1 and t_2. This procedure yields $g_{j,e}$, which is the average growth rate for bin $j \in \{1, \ldots, 100\}$ and episode $e \in \{1, \ldots, 4\}$. For each episode we regress $g_{j,e} = a_e^0 + a_e^1 \times j$ for $j \in \{11, \cdots, 90\}$ and we regress $g_{j,e} = a_e^2 + a_e^3 \times j$ for $j \in 91, \cdots, 100$. We perform the same regressions using the values for $g_{j,e}$ from the data and form the objective by assessing the difference between a_e^1 and a_e^3 and their empirical counterparts, denoted $a_e^{1,data}$ and $a_e^{3,data}$. We then add the following to the objective function

$$\frac{93}{8} \sum_{e=1}^{4} \left[\left(\frac{a_e^1 - a_e^{1,data}}{a_e^{1,data}} \right)^2 + \left(\frac{a_e^3 - a_e^{3,data}}{a_e^{3,data}} \right)^2 \right].$$

The results show that the model is able to match the increasing slope of the factor structure in the four recessions. Figure 12 shows the factor structure during recessions for both model and data. The empirical estimates reported by GOS remove life-cycle effects from the intercept. As we are interested mainly in the slope across percentiles of the income distribution, we normalize the empirical estimates so the model and empirical

[2]We apply weight 93/7 to this component of the objective function.

Figure 12: Factor structure for recessions. (Black = model, gray = data)

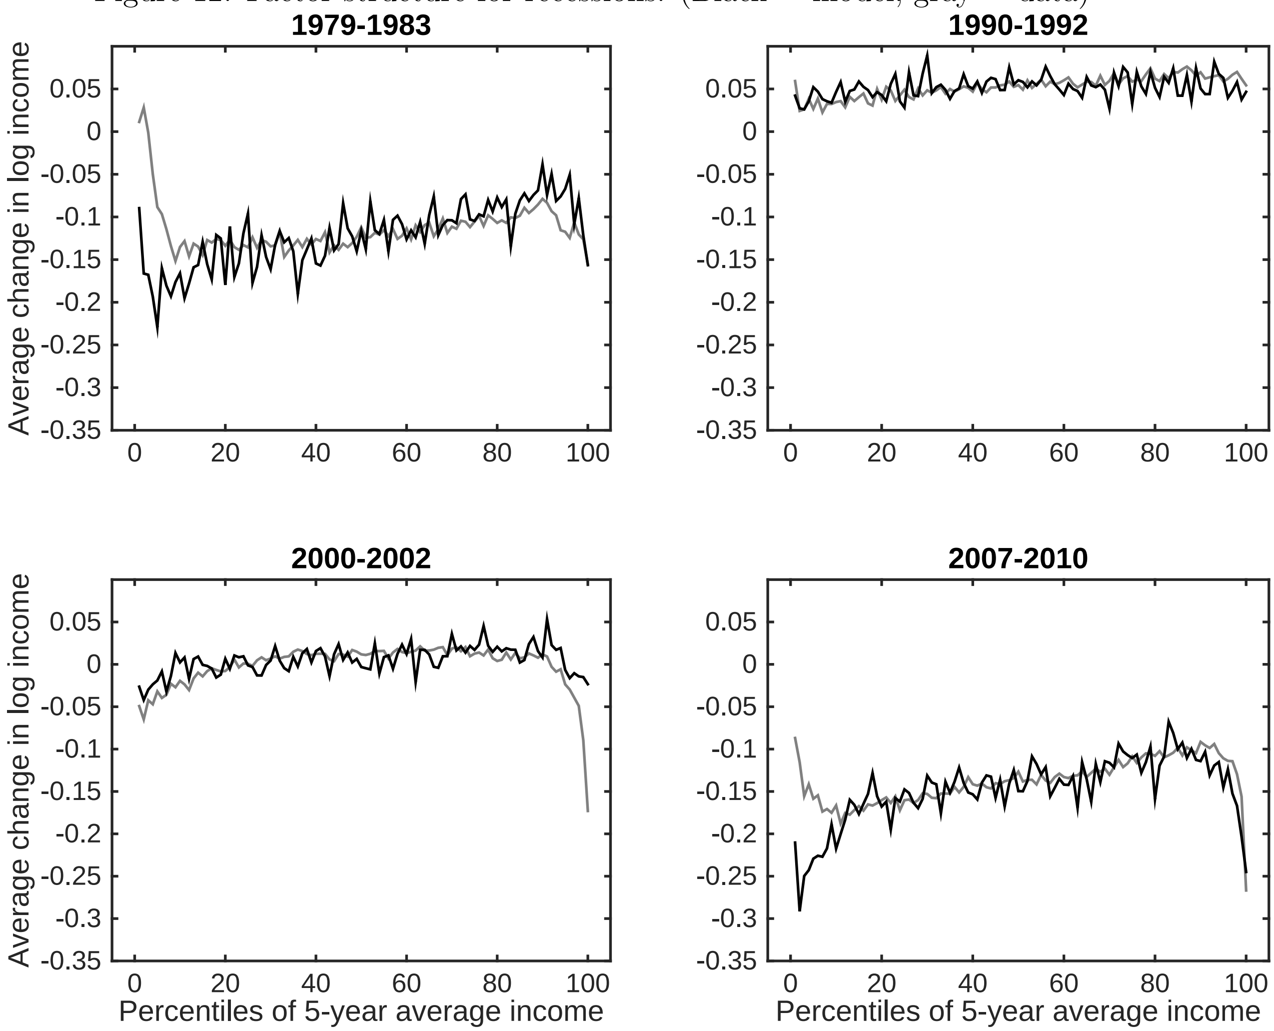

series have the same mean over the 11th to 90th percentiles. The model fits the slope of the central part of the factor structure very well. Except for 2000-2002 it fits the behavior of the right tail well, too. 2000-2002 did not involve a large drop in average income, but perhaps the right-tail was particularly affected by the dot-com bust (or in other words, the right-tail was really fat in the dot-com boom and then things returned to normal). The other noticeable errors are the left tails for 1979-83 and 2007-2010. 13 shows similar plots for the expansions. For these we do not have data to compare to.

Introducing the factor structure does not diminish the model's ability to match the other targets by much. Table V shows moments of the one-, three-, and five-year growth

Figure 13: Factor structure for expansions. (Model)

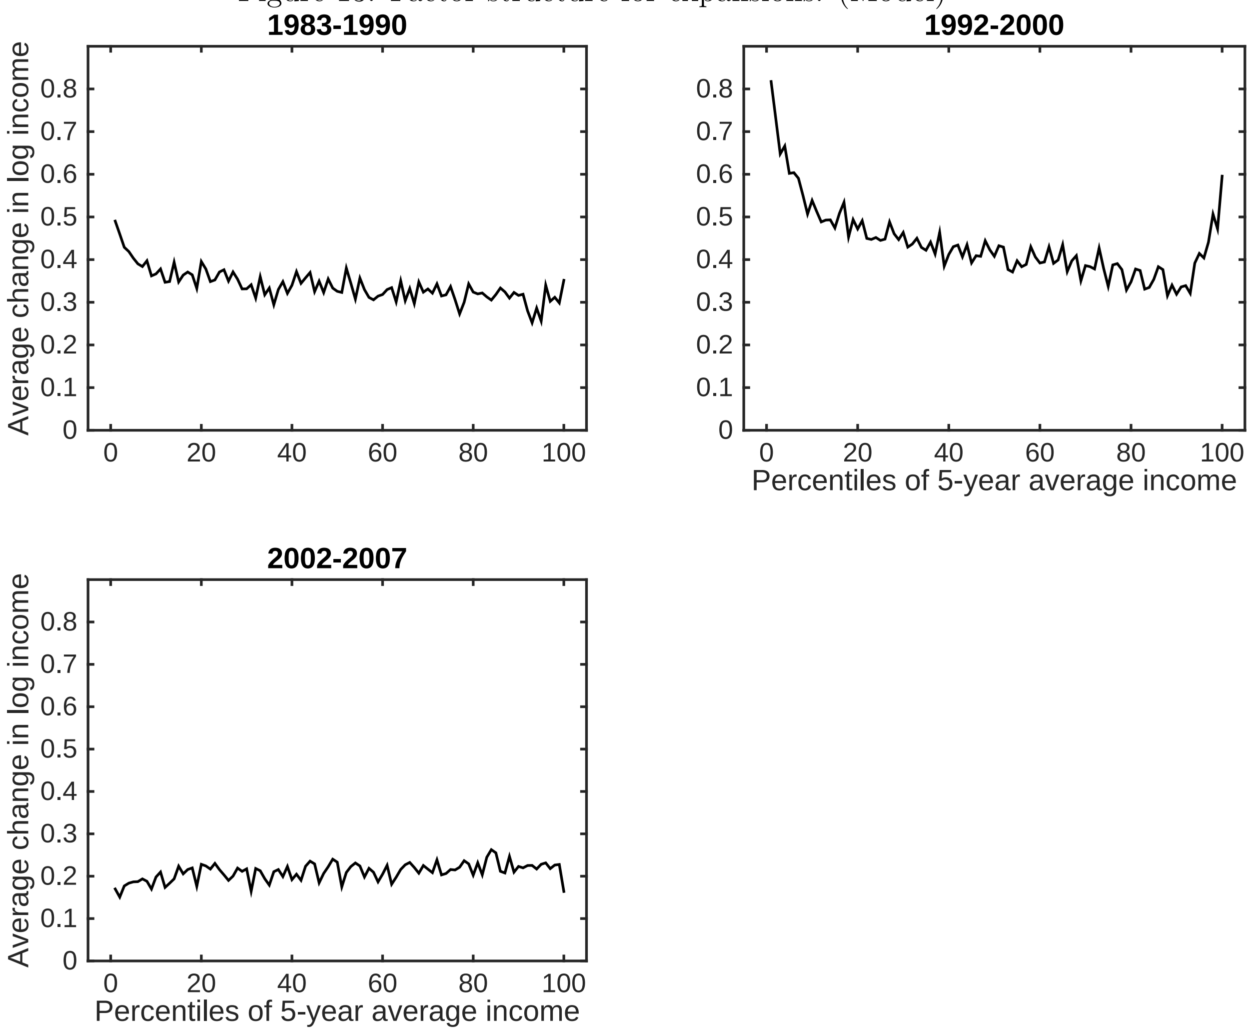

rates. The model almost hits the targets of 20 and 12 for the one-year and five-year kurtosis.

The income process implies an increase in variance over 35 years equal to 0.46197 log points, which is again close to the of 0.45. Figure 14 shows the fit to the mean, median, and the 10th and 90th percentiles year by year. Figure 15 shows the evolution of Kelly's skewness over the cycle, and Figure 16 shows the evolution of the standard deviations. In Figure 14, there is a bit too much volatility in the mean and median and the 10th and 90th percentiles inherit this to some degree. I believe this reflects the fact that w_t controls the center of the distribution and the extent to which their is a factor structure

	st. dev.		skew.		kurt.	
	model	data	model	data	model	data
One-year	0.47		-0.84		20.47	
Three-year	0.61		-0.47		14.73	
Five-year	0.70		-0.27		11.68	

Table V: Moments of stationary earnings growth rate distributions.

(recall x_t in the factor structure is set equal to w_t). By asking the model to generate the factor structure it is pulled in the direction of generating larger fluctuations in average income growth. Figure 15 shows a good fit to Kelly's skewness. Figure 16 shows that there is still somewhat too much volatility in the standard deviations, but this is similar to version 4. Figure 17 shows a histogram of one-year earnings growth rates [computed with w_t and x_t from 1995-1996 to match the data]. The histogram isdiffers from those in previous version in that there is a noticeable shoulder around -2.

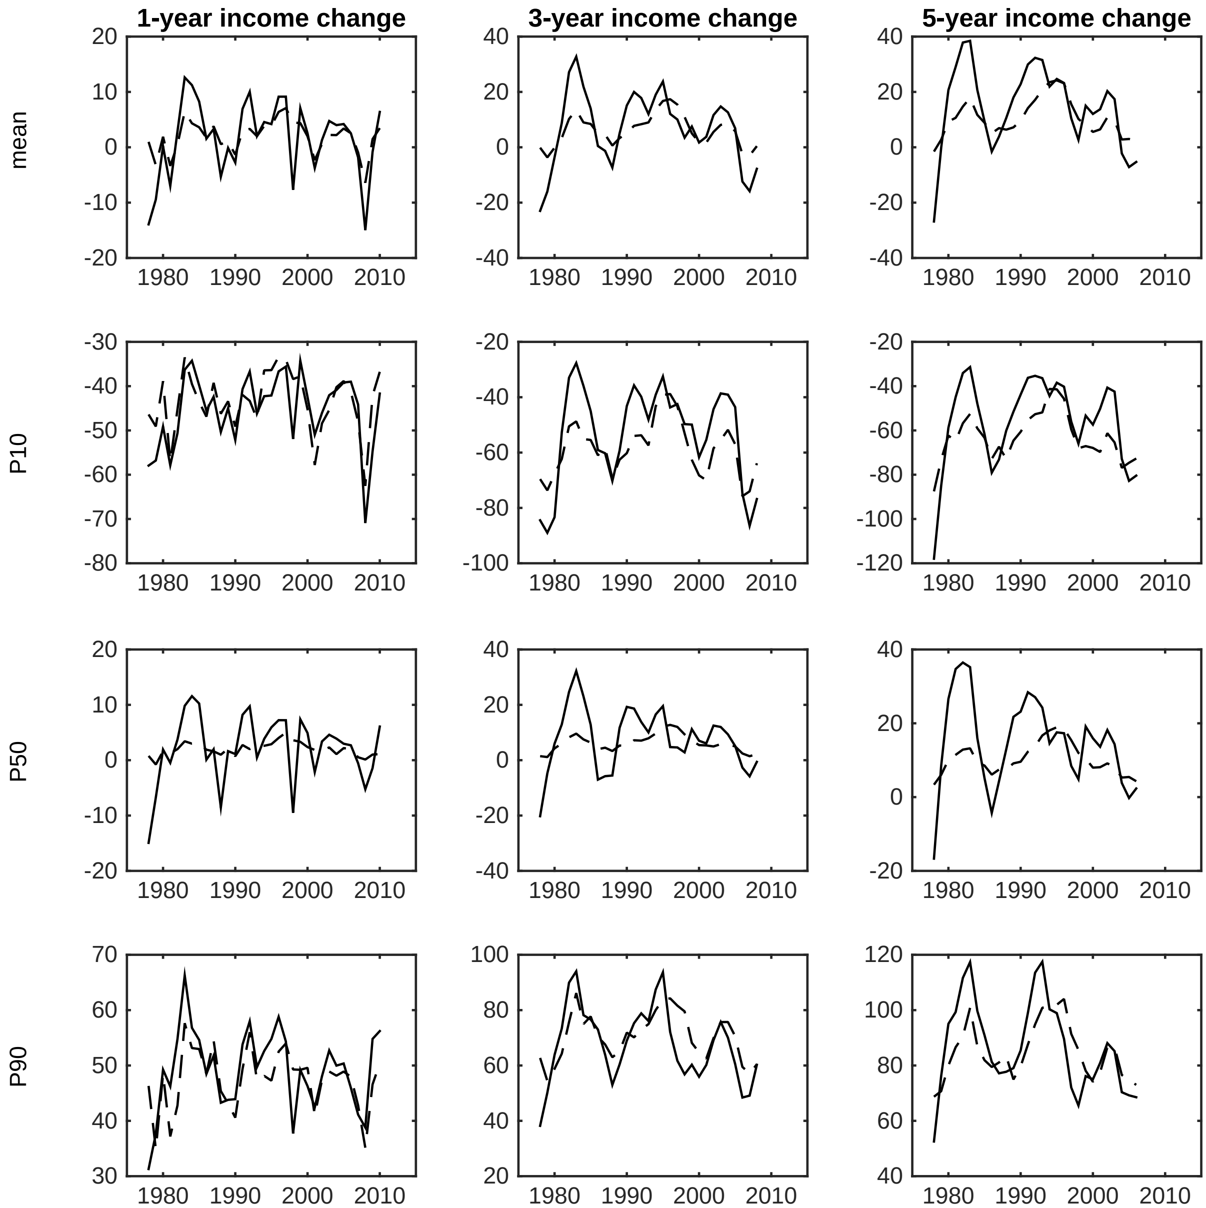

Figure 14: Moments of the earnings growth distributions. Solid = simulated, dashed = data.

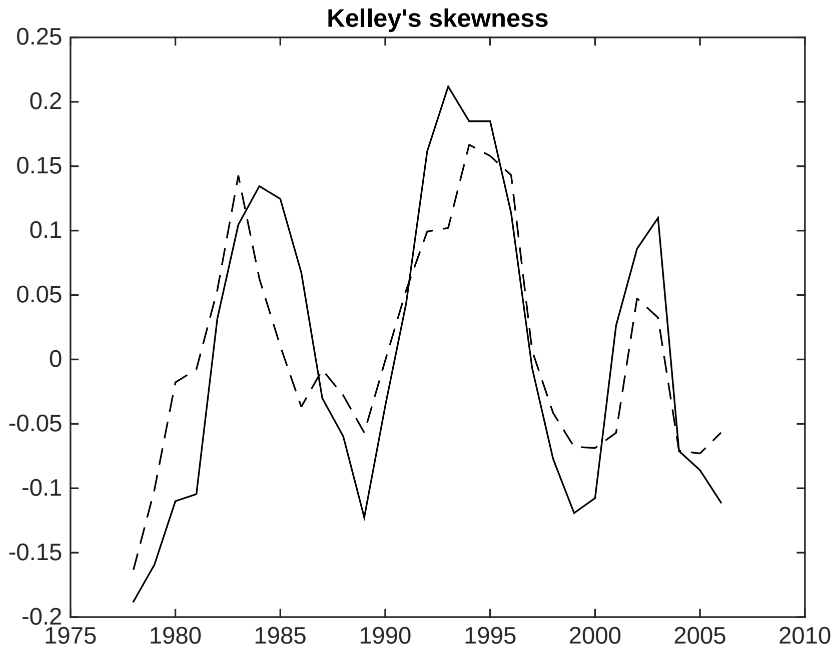

Figure 15: Kelly's skewness. Solid = simulated, dashed = data.

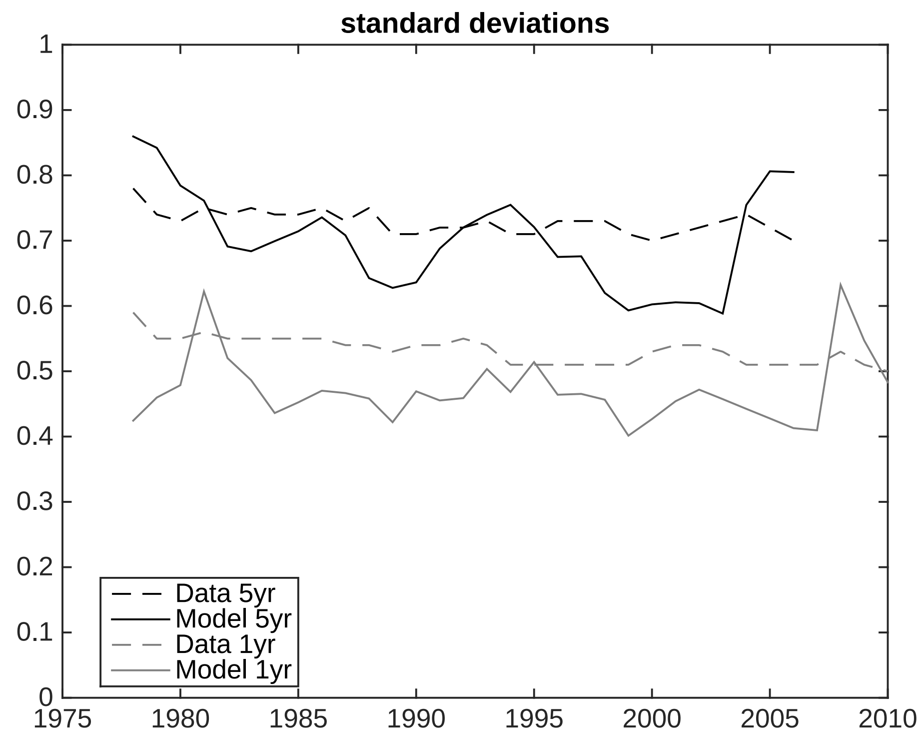

Figure 16: Standard deviation of earnings growth distribution.

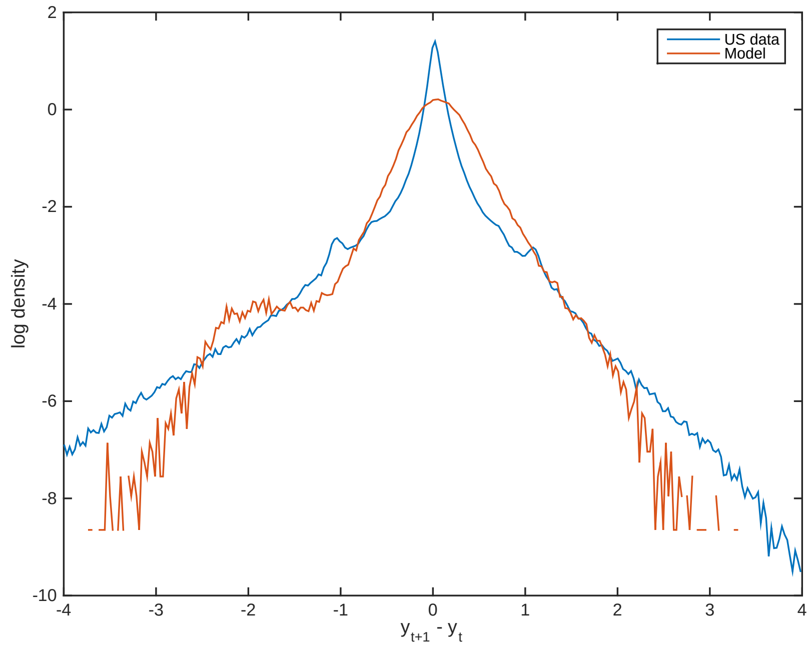

Figure 17: Histogram of one-year earnings growth rates.