 DOI: 10.1086/675535

Appendix B from Guvenen et al., "The Nature of Countercyclical Income Risk"
(JPE, vol. 122, no. 3, p. 621)

Sensitivity Analysis

A. Mean-Reverting Shocks

We now report the key statistics on standard deviation and skewness by quasi-differencing the data, $y_t - \rho y_{t-1}$, and using three possible values for $\rho = 0.80$, 0.90, and 0.95. As seen in figures B1–B8, the lack of cyclicality in the standard deviation is robust to these variations. The level of skewness is lower the lower is ρ, but the countercyclicality remains intact for all values of ρ.

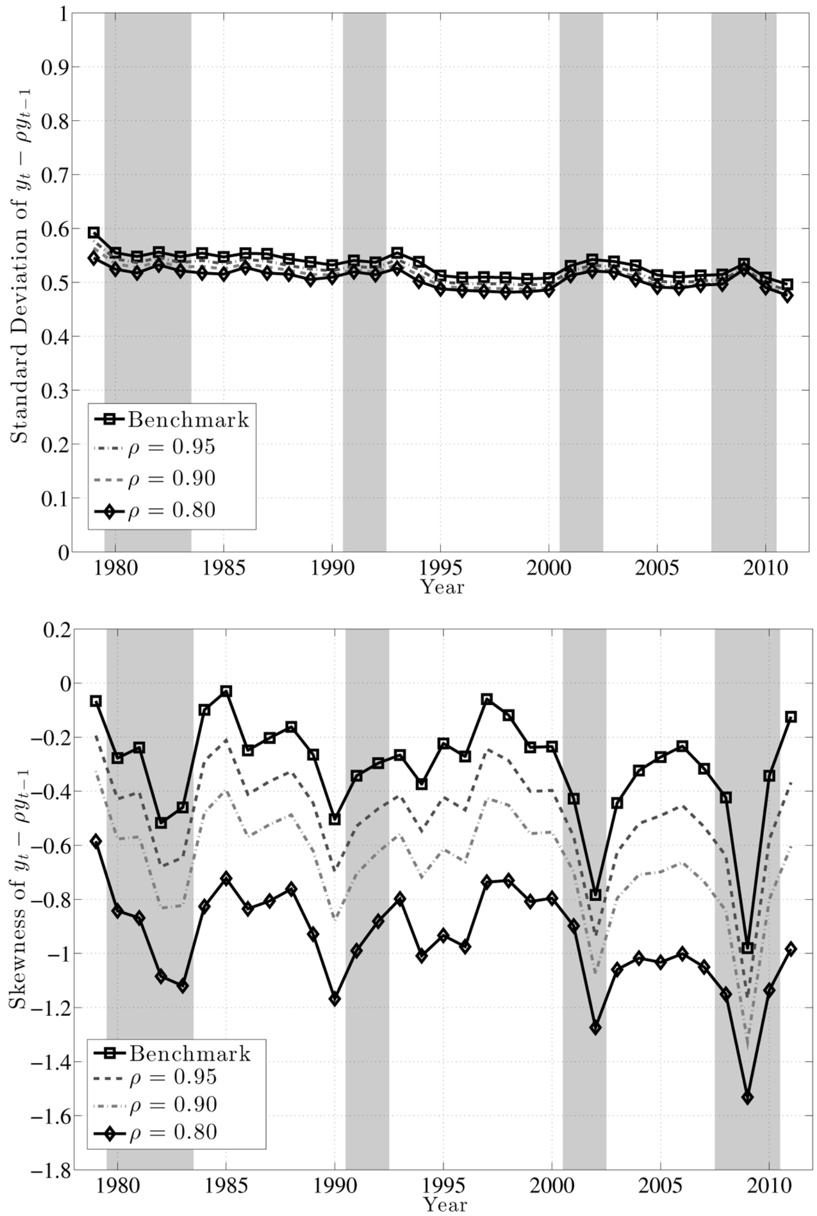

FIG. B1.—Sensitivity of business cycle variation to mean reversion. *Top*, standard deviation. *Bottom*, skewness.

B. Within-Group Variation by Age Group

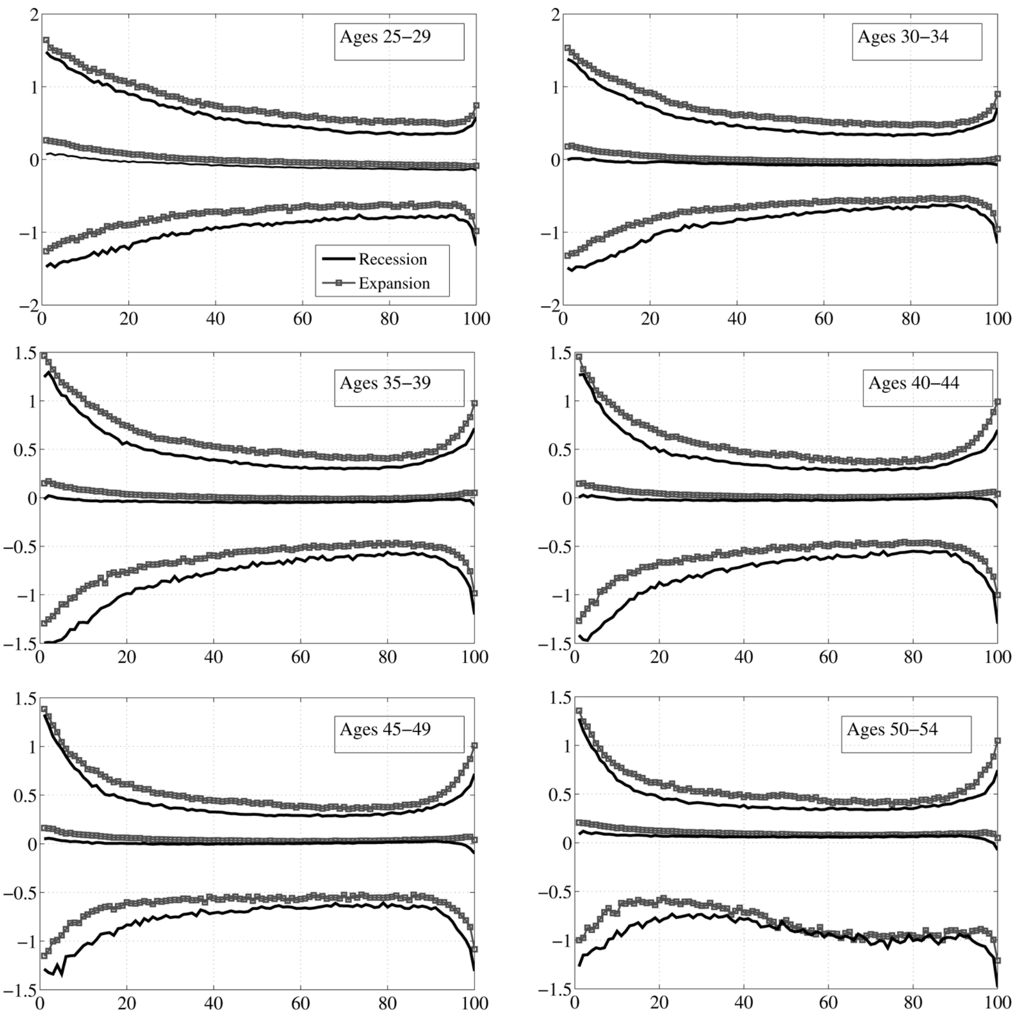

FIG. B2.—Standard deviation of 5-year earnings growth, by age groups

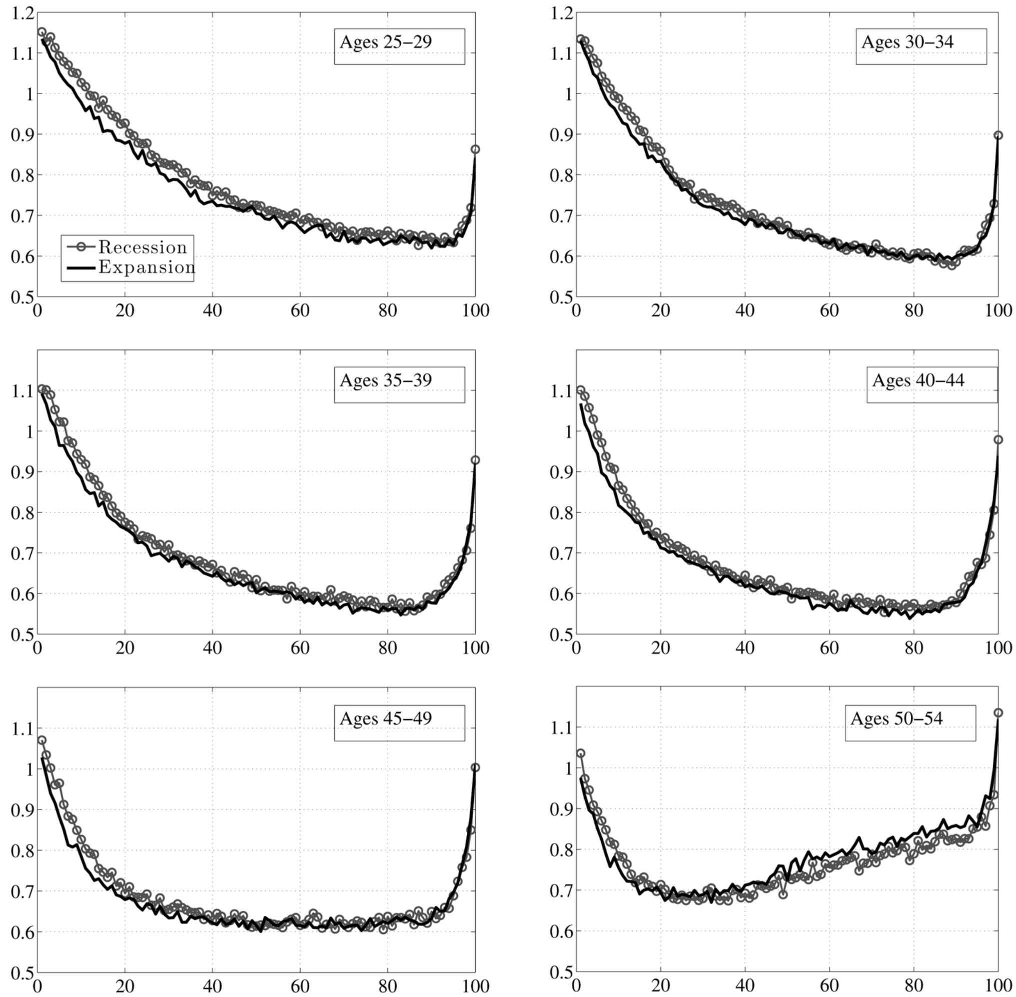

Fig. B3.—Standard deviation of 5-year earnings growth, by age groups

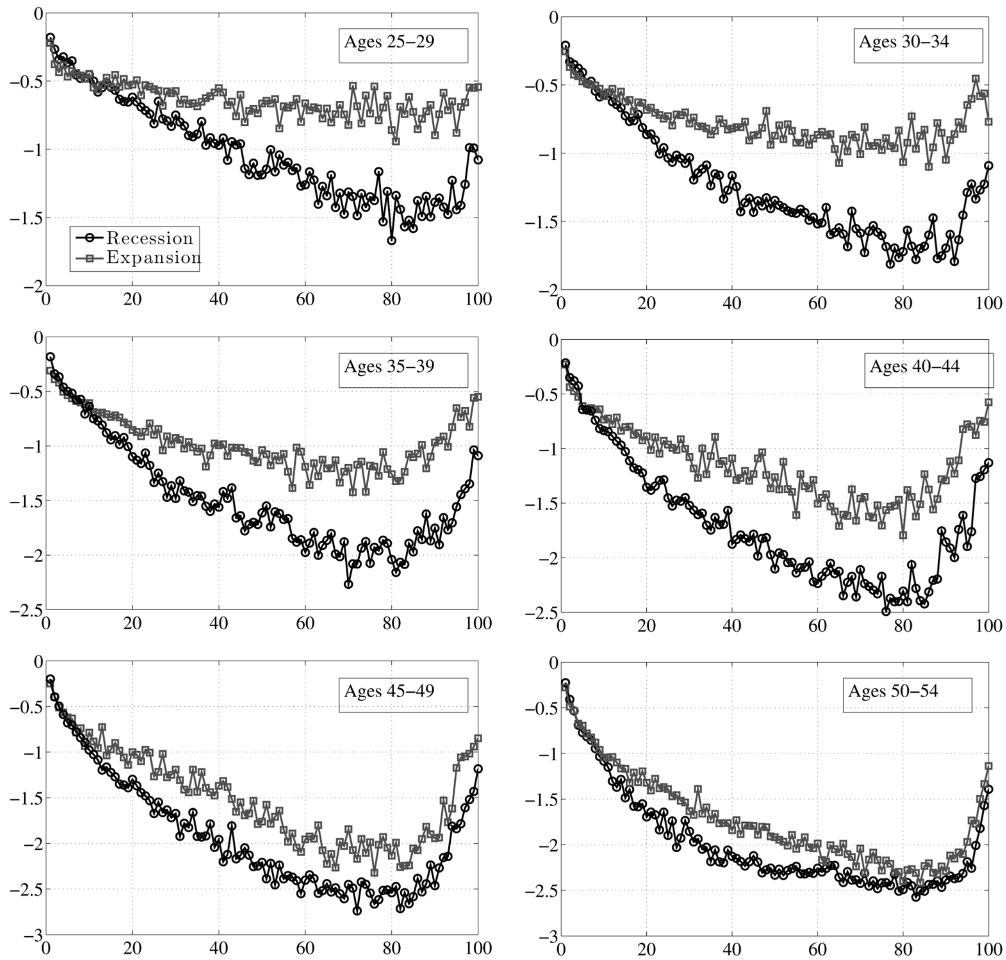

FIG. B4.—Skewness of 5-year earnings growth, by age groups

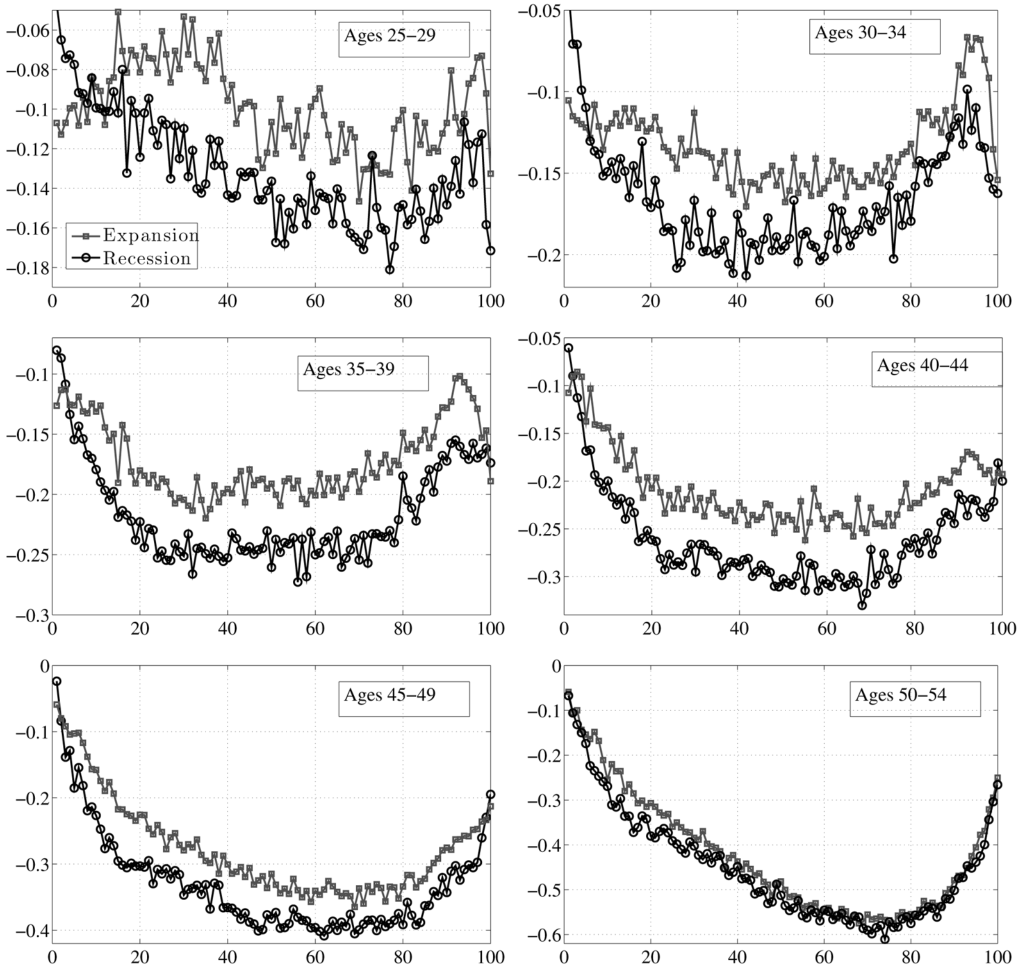

FIG. B5.—Kelley's skewness of 5-year earnings growth, by age groups

C. Between-Group Variation by Age Group

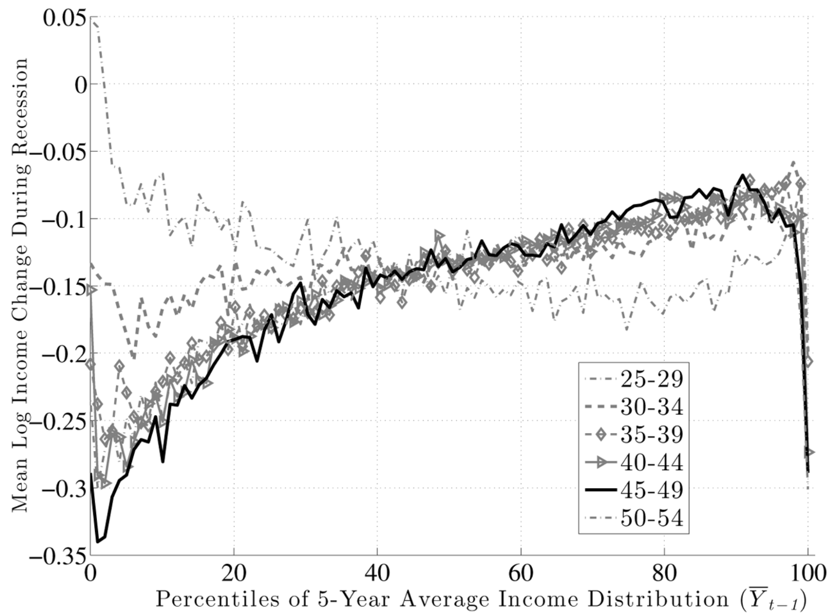

FIG. B6.—Growth in log average earnings during the Great Recession (2007–10)

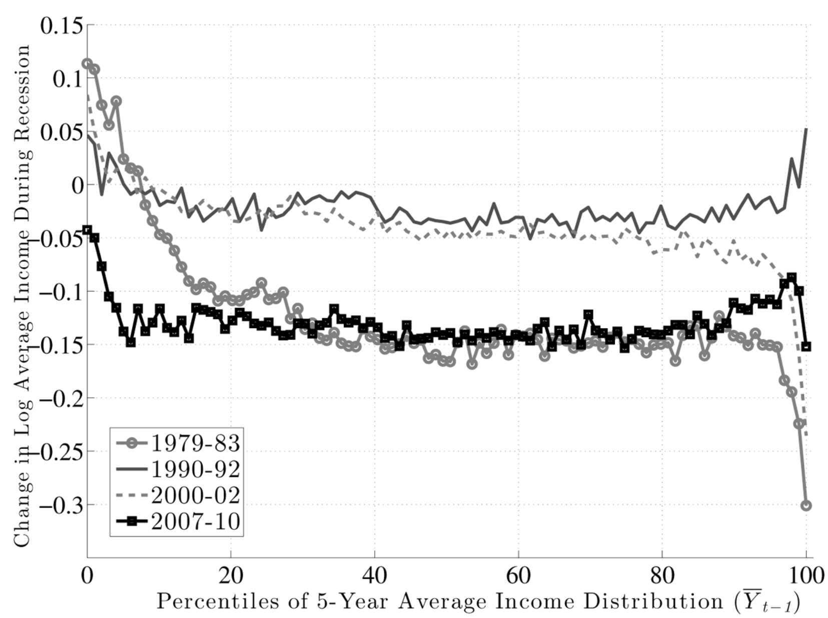

FIG. B7.—Growth in log average income during recessions, young (25–34) males

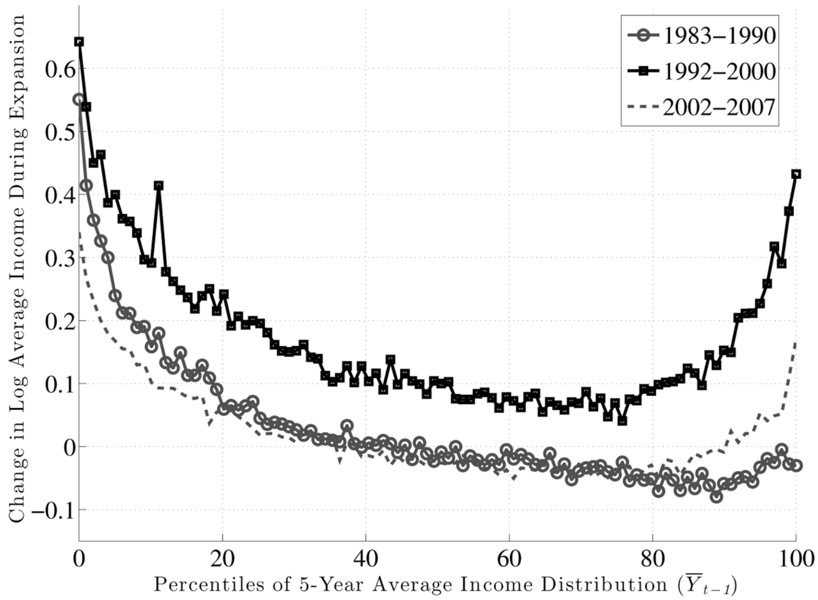

FIG. B8.—Growth in log average income during expansions, young males

D. Alternative Measure of Factor Structure: f_1

We now construct the alternative measure of average earnings growth, f_1, described in the main text (Sec. VI). Recall that f_1 differs from f_2 in two important ways. First, f_1 excludes individuals with zero earnings in either year t or year $t+k$. Because the probability of full-year nonemployment rises in recessions most strongly for low-income individuals, dropping them will tend to increase f_1 below the median relative to f_2. Second, because f_1 is based on the *average of log* earnings whereas f_2 is based on the *log of average* earnings, the latter will tend to be higher within quantiles that have a wider dispersion of earnings growth rates (due to Jensen's inequality). So, we would expect this force to raise f_2 relative to f_1 below the median level of $\overline{Y}_{t-1}$, where the variance of shocks is higher, as well as at the very top end for the same reason.

Figure B9 plots f_1 for each of the four recessions. A quick comparison to figure 13 shows that the two measures reveal the same qualitative patterns. The clear upward-sloping factor structure is there for all recessions. Quantitatively, the slope is somewhat smaller: a difference of 10 log points between the 90th and 10th percentiles during the Great Recession versus 17 log points under f_2. Inspecting the two graphs shows that the difference mainly comes from the steeper drop in f_2 between the 20th and 1st percentiles, probably because of the increased chance of unemployment in this range mentioned above. Between the 20th and 90th percentiles, the two graphs differ by little. The other recessions show slopes that are also slightly lower than before. Another difference to note is that under f_1, the 1980–83 recession looks less favorable to individuals in the top 10 percent: their earnings growth pattern resembles the recent recessions more closely. This suggests that the strong performance of this group revealed by f_2 was affected by some large gains at the right tail, which dominated the mean earnings measure for these groups in 1983.

Overall, the two measures are quite comparable. In the main text, we focus on f_2 so as to capture the total earnings risk, which includes the risk of long-term unemployment rising during recessions.

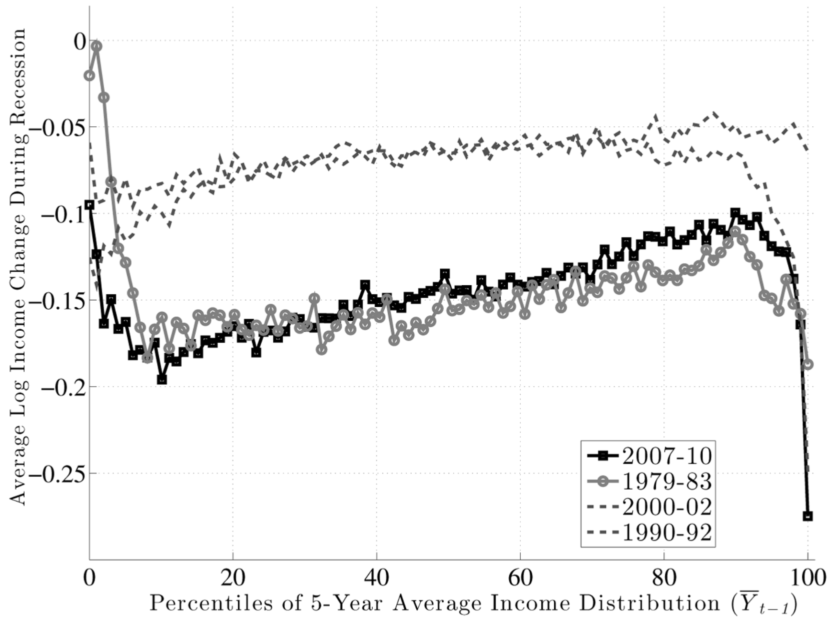

FIG. B9.—Average growth in log earnings during recessions (f_1), prime-age males

E. Rising Stars versus Stagnant Careers

We now control for three characteristics simultaneously: age, $\overline{Y}^i_{t-1}$, and $\Delta_5(y^i_{t-1})$. Because the 1979–83 period does not allow us to construct the pre-episode growth rate, we drop it from the analysis of this section.

We first sort individuals within an age group according to their $\overline{Y}^i_{t-1}$ and $\Delta_5(y^i_{t-1})$ (independently in each dimension) and compute 50- and 40-quantile thresholds, respectively. We use these thresholds to assign each individual into groups formed by the intersection of age, pre-episode average earnings (indexed by j), and earnings growth (indexed by p) categories. To give an idea about the bounds of a typical group, for the analysis of the Great Recession, one such group will consist of individuals who (i) were between the ages of 35 and 39 in year 2006, (ii) earned average annual earnings ($\overline{Y}^i_{t-1}$) between $32,033 and $33,455 from 2002 to 2006, and (iii) experienced an annual earnings growth rate between 1.30 percent and 1.49 percent per year from 2002 to 2006. Clearly, this is a finely defined group of individuals. For each of these 2,000 cells, we compute the average labor earnings: $y^{j,p}_t$ and $y^{j,p}_{t+k}$.[29] We then regress

$$y^{j,p}_{t+k} - y^{j,p}_t = \sum_{j=1}^{50} \alpha_j d^j_{\overline{Y}} + \sum_{p=1}^{40} \gamma_p d^p_{\Delta y^i} + u^{j,p}_t, \tag{B1}$$

where $d^j_{\overline{Y}}$ is a dummy variable that equals one if the group on the left-hand side belongs in the jth quantile of the $\overline{Y}_{t-1}$ distribution and zero otherwise. The dummy $d^p_{\Delta y^i}$ is defined analogously for the quantiles of $\Delta_5(y^i_{t-1})$. The 90 dummies are estimated via ordinary least squares.

The main findings are as follows. First, the additional control for $\Delta_5(y^i_{t-1})$ has virtually no effect on the results presented in the main text, where we conditioned only on $\overline{Y}^i_{t-1}$. This can be seen clearly in figure B11, which plots the original graph ($f_2(\overline{Y}_{t-1})$) superimposed on the new one ($f_2(\overline{Y}_{t-1}|\Delta Y_{t-1})$). Second, the main finding is that pre-episode earnings growth has a significant effect on future growth. This is shown in figure B10, which plots average earnings growth during expansions (blue line with circle markers) and recessions (red line with square markers). While mean reversion is apparent in both cases, the gap between the two graphs is smallest in the middle and expands at both ends. This is clearly seen in the right panel,

[29] Because the two variables can be correlated, there is no presumption that every cell will contain the same number of observations (unlike the previous experiment with a single characteristic). Therefore, we drop cells that have less than 30 percent of the maximum number of observations.

which plots the annualized gap between expansions and recessions. The implication is that workers with the highest and lowest earnings growth rates prior to an episode do better during expansions than recessions. This is related to the fact documented earlier that the top of the earnings shock distribution collapses during recessions. Consequently, the earnings growth rate of those individuals whose earnings would have grown faster during expansions actually slows down during a recession.[30]

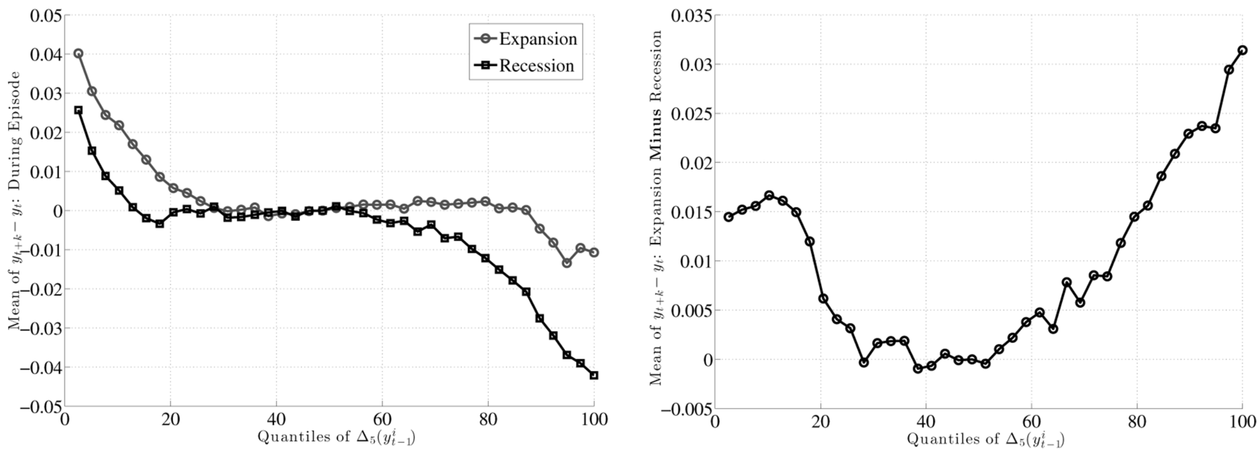

FIG. B10.—Growth in log average earnings by quantiles of recent growth rate. *Left*, expansion versus recession. *Right*, expansion minus recession.

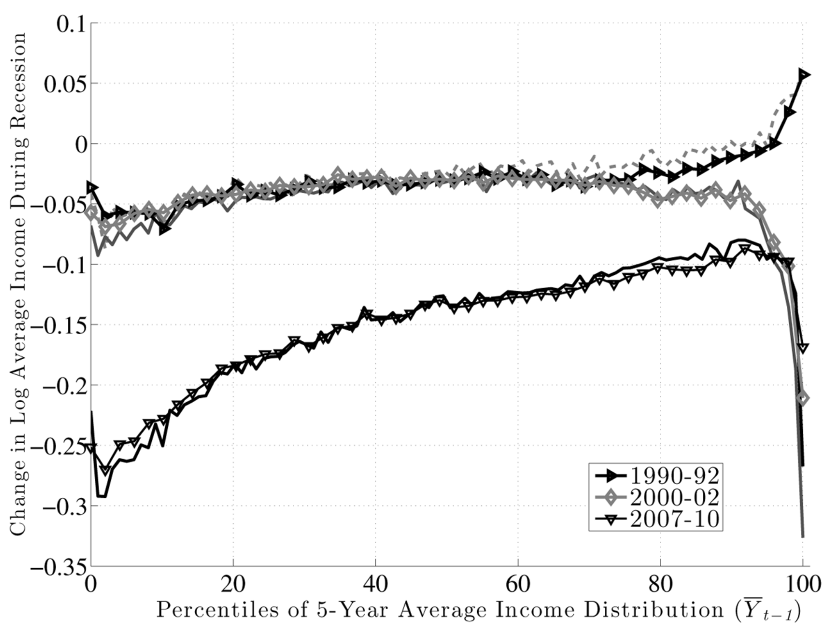

FIG. B11.—Comparing $f_2(\overline{Y}_{t-1})$ (from fig. 13) to $f_2(\overline{Y}_{t-1}|\Delta Y_{t-1})$

[30] Incidentally, controlling for past earnings growth has virtually no effect on the relationship between the quantiles of average earnings and future earnings growth documented above. Thus, further conditioning does not alter the relationship documented so far. These figures are available on request.

F. What Role Does Unemployment Play?

How much of the countercyclicality of left-skewness is due to the fact that unemployment rises in recessions, so more individuals experience large negative earnings changes, because they are part-year unemployed? Here, we address this question.

Recall that the MEF data set does not contain information on labor hours or unemployment. However, providing an upper bound on the potential effects of unemployment is still possible. To begin with, notice that unemployment (or nonemployment) can affect our results through two separate channels. First, workers who are full-year nonemployed are excluded from the sample in that year. This creates a truncation at the bottom end of the earnings growth distribution, whose severity varies over the business cycle. Second, many part-year unemployed individuals are still included in our sample as long as their annual earnings remains above $Y_{\min}$. (Incidentally, both of these assumptions are precisely the same ones made in the bulk of existing literature on income risk.) It is useful to discuss whether and, if so, how they might be affecting our findings on skewness.

1. The Effect of Part-Year Unemployment

First, recall that the countercyclicality of left-skewness is due to both (i) the compression of positive earnings growth changes toward the median and (ii) the expansion of negative earnings growth rates toward the bottom end (figs. 4, 7, 10, and 12). The compression at the top is unlikely to be related to unemployment. So even if the bottom half remained unchanged, skewness would be more negative during recessions because of the compression at the top alone.[31]

Second, countercyclical left-skewness is also evident in 5-year earnings changes. Because recessions last less than 5 years, (the incidence of) unemployment is only slightly higher in $t+5$ than in t. This can be seen in the left panel of table B1, which reports the fraction of 35–54-year-old males with an unemployment spell longer than $x = 0$, 13, and 26 weeks in a given year, computed from the CPS (Integrated Public Use Microdata Series).

Consider spells longer than 13 weeks (third column). Only 5.4 percent of prime-age males are in this group in year $t+5$ (averaging over 1984, 1994, 2004, and 2010). Now let us assume that (i) none of these individuals spent any time in unemployment in year t and (ii) their actual wages and hours remained the same in t and $t+5$ while they were employed. Then, for these individuals, unemployment reduces their annual earnings by at least 25 log points between t and $t+5$. So this would appear as a negative earnings shock of 25+ log points. Similarly, the average incidence in year t is 3.9 percent, so by the same computation, these individuals will appear as having received a positive shock of 25+ log points between t and $t+k$. So the net effect on skewness depends on the gap: $5.4 - 3.9 = 1.5$ percent of individuals who get more negative shocks than positive in year $t+k$. If we assume for the moment that these individuals are evenly spread across the $\overline{Y}_{t-1}$ distribution, it would amount to a 1.5 percent net change of the sample within each quantile, which is a small number. Further, the same computation can be repeated for $x = 0$ or $x = 26$ weeks, with nearly identical results.[32]

Overall, this analysis suggests that the direct effect of unemployment is likely to be small for the results on skewness. The cyclical change in unemployment for prime-age males is simply too small to account for the countercyclicality of skewness, which is observed across the entire range of past earnings levels and earnings growth rates.

Table B1. Incidence of Unemployment over the Business Cycle, Prime-Age Males

	CPS Data				SSA Data	
Year	$x > 0$ (Weeks)	$x > 13$	$x > 26$	Period	$E \to N$	$N \to E$
1979	10.5	4.4	1.5	1979–83	5.1	4.4
1984	11.4	6.3	2.8	1983–90	4.1	3.8
1989	10.0	4.7	1.8	1990–92	3.8	2.7
1994	9.5	5.2	2.3	1992–2000	3.1	2.8

[31] Notice also that in this scenario, the variance of shocks would go down in recessions, and thus the variance would be *procyclical*.

[32] In addition, the case described here relies on some unlikely assumptions. For example, the probability of unemployment is a strongly decreasing function of past income, so the change in incidence among individuals with $\overline{Y}_{t-1} \in$ P90 will be much smaller than the 1.5 percent average figure. Yet, the shift to negative skewness among that group is as large as among workers who have $\overline{Y}_{t-1} \in$ P50 as well as $\overline{Y}_{t-1} \in$ P30 (see fig. 10). As an alternative

Table B1 (*Continued*)

	CPS Data				SSA Data	
Year	$x > 0$ (Weeks)	$x > 13$	$x > 26$	Period	$E \rightarrow N$	$N \rightarrow E$
1999	6.0	3.0	1.1	2000–2002	3.7	2.3
2004	6.5	3.6	1.4	2002–7	3.2	2.7
2005	6.7	3.6	1.4	2007–10	4.5	2.3
2010	10.4	6.6	3.2			
Average t	8.3	3.9	1.5	Expansion	3.5	3.1
Average $t + 5$	9.4	5.4	2.4	Recession	4.3	2.9

Note.— The left panel reports the incidence of unemployment with duration exceeding x weeks. The first column in the right panel reports the fraction of individuals who are full-year nonemployed in $t + 1$ (denoted N) conditional on being employed in t (denoted E). The last column shows the opposite transition.

2. Excluding Zeros (Full-Year Nonemployed)

A second and separate issue relates to our exclusion of full-year nonemployed individuals. If anything, this assumption is truncating the actual downside risk in recessions and is understating the countercyclicality of skewness. This can be seen as follows. Using our sample, we compute the fraction of individuals who are in the sample in year t but not in $t + 1$ for every year of the sample. Then for each business cycle episode, we report the average figure in the right panel of table B1. Not surprisingly, we are dropping more individuals from the sample in each recession (given that the likelihood of full-year nonemployment rises). On average, we are dropping 4.3 percent of individuals from our sample in year $t + 1$ during recessions and 3.5 percent during expansions. If these excluded individuals were included (e.g., by assigning them a nominal earnings level, say \$100 in that year), this would register as a large earnings drop in recessions and increase the left-skewness in recessions. However, because the change over the business cycle is small, the effect would also be small.

G. Broadening the Definition of Business Cycles

So far in the analysis, we have viewed business cycles as consisting of recessionary and expansionary episodes. But some important macroeconomic variables do not perfectly synchronize with these episodes. For example, as also mentioned earlier, unemployment peaked in 1993 and 2003—two years that are part of expansions. Similarly, the stock market experienced a significant drop in 1987, again during an expansion. With these considerations in mind, this section explores the robustness of our results to alternative indicators of business cycles.

For a given quantile j of $\overline{Y}_{t-1}$, we regress the change between t and $t+1$ in log average earnings (f_2^j) on alternative measures of business cycles, denoted by x:

$$f_2^j(t, t+1) = \alpha^j + \beta^j x + \epsilon_t.$$

We consider three choices for x: (log) growth rate in GDP per capita, the annual return on the US stock market (as measured by the S&P 500 index), and the annual change in the male unemployment rate (denoted ΔU). Table B2 displays the estimated β^j's for several key quantiles and for two time periods: the full sample (1978–2009) and one that excludes the double-dip recession (1985–2009).

Several observations are worth noting. First, cyclicality is U-shaped across earnings quantiles, regardless of the business cycle variable chosen. This is consistent with the conclusion of Section VI.B above, summarized in figure 15. It is also consistent with Parker and Vissing-Jørgensen's (2010) analysis using repeated cross sections and synthetic earnings

sensitivity analysis, we repeat the computation of skewness, but this time using the 1980–85, 1990–95, and 2000–2005 periods and excluding the Great Recession. With this timing, the ending year is well into the expansion, so the incidence of unemployment of 13 weeks or longer is only 0.4 percent higher in $t + 5$ compared with t. The grey dashed-dotted line in the left panel of fig. 11 plots Kelley's skewness under these assumptions, which is still significantly more negative during these three recessions.

groups. Second, cyclicality increases after 1985, especially at the very top of the earnings distribution and especially when business cycles are measured by GDP growth or the unemployment rate. Cyclicality is pretty flat in the middle of the earnings distribution (e.g., between P25 and P75) and increases slightly at the bottom end. Third, the comovement of the earnings growth of top earners with GDP growth and stock returns is quite striking. For example, after 1985, a 1 percentage point rise in the male unemployment rate has been accompanied with an average earnings decline of 6.87 percent for individuals who were in P99.9 before the shock. Similarly, a 1 percentage point slowdown in GDP/capita growth implies a 4.55 percent decline in the earnings of the same individuals.[33] For comparison, the corresponding numbers for individuals with median earnings are 1.08 and −1.77.

Table B2. Cyclicality of Earnings Growth, Prime-Age Males

	Dependent Variable x: f_2^j					
	1978–2009			1985–2009		
j	ΔGDP	$R^S_{t,t+1}$	ΔU	ΔGDP	$R^S_{t,t+1}$	ΔU
P99.9	3.07	.43	−4.76	4.55	.46	−6.87
P99	1.45	.20	−2.42	2.09	.22	−3.34
P90	1.48	.06	−1.17	1.70	.06	−1.21
P75	.75	.06	−1.22	.75	.05	−1.13
P50	1.04	.09	−1.77	1.09	.08	−1.74
P25	1.63	.14	−2.80	1.78	.14	−2.86
P10	1.85	.17	−3.22	2.06	.16	−3.34
Standard deviation (x)	2.10	16.80	1.23	1.81	17.78	1.10

Note.—Each cell reports the β^j estimated for individuals in earnings group j and for business cycle variable x; $R^S_{t,t+1}$ is the annual realized return on the S&P 500 index (data obtained from Robert Shiller's website at Yale University). All regression coefficients are significant at the 0.1 percent level when the regressor is the GDP growth or change in unemployment rate and are significant at 1 percent for stock returns.

[33] The corresponding figures for the whole sample period are 4.76 percent and 3.07 percent, respectively.

H. Cyclicality of Top 1 Percent Using f_1

Figure B12 plots the counterpart of figure 16 using a different measure of earnings growth (f_1). The same pattern discussed in Section VI.C is visible here with an even larger 5-year loss for all individuals in the top 1 percent.

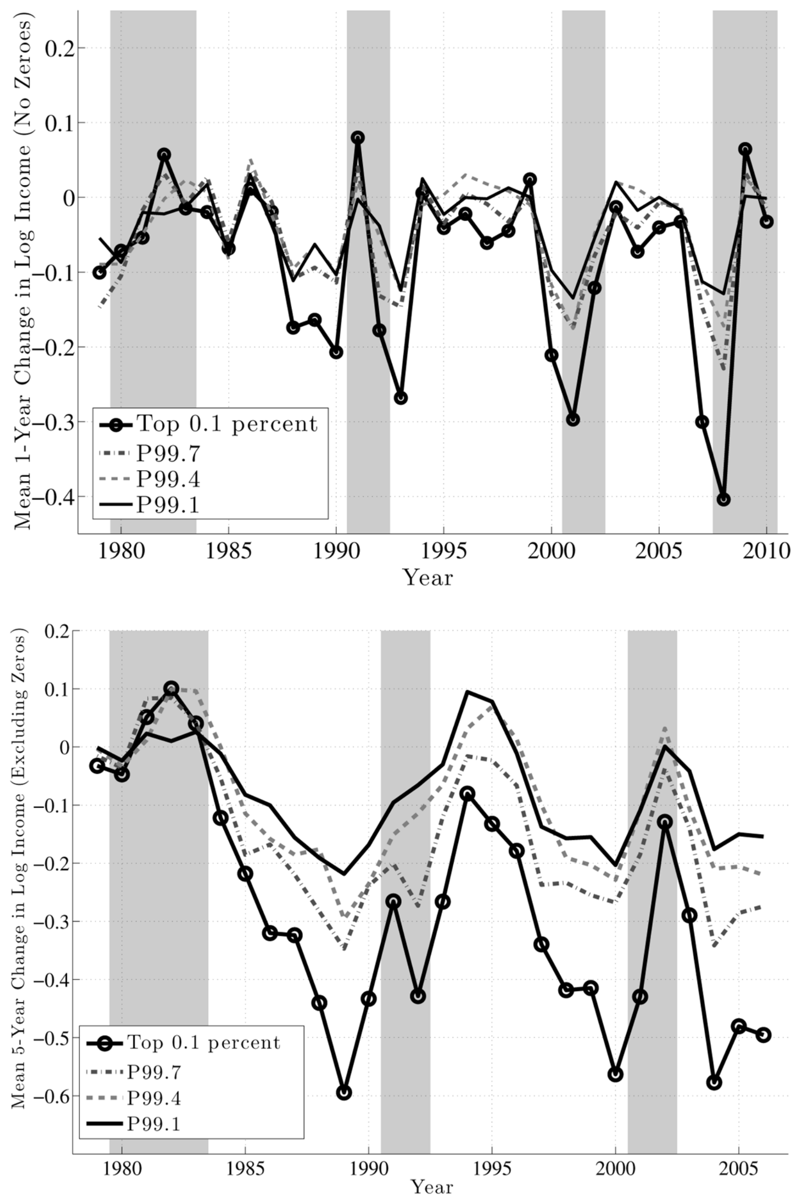

FIG. B12.—5-year earnings growth, top 1 percent of individuals. *Top*, average 1-year change in log earnings (f_1). *Bottom*, average 5-year change in log earnings (f_1).