Cyclical Dynamics in Idiosyncratic Labor-Market Risk*

Kjetil Storesletten,† Chris I. Telmer,‡ and Amir Yaron§

November 2002
(first draft: June 1996)

Abstract
We ask how idiosyncratic labor-market risk varies over the business cycle. A difficulty in addressing this question is the limited time-series dimension of existing panel data sets. We address this difficulty by developing a GMM estimator which conditions on the macroeconomic history experienced by each member of the panel. Variation in the cross-sectional variance between households with differing macroeconomic histories allows us to incorporate business-cycle information dating back to 1930, even though our data only begins in 1968. We implement this estimator using household-level labor-earnings data from the Panel Study on Income Dynamics. We estimate that idiosyncratic risk is (i) highly persistent, with an autocorrelation coefficient of 0.95, and (ii) strongly countercyclical, with a conditional standard deviation that increases by 75% (from 0.12 to 0.21) as the macroeconomy moves from peak to trough.

Forthcoming *Journal of Political Economy*

*This paper circulated previously as a component of the paper "Asset Pricing with Idiosyncratic Risk and Overlapping Generations." In addition to participants at seminars and conferences, we thank Dave Backus, Geert Bekaert, David Bowman, Martin Browning, John Cochrane, Wouter den Haan, Mick Devereux, Ron Gallant, René Garcia, Rick Green, John Heaton, Burton Hollifield, Mark Huggett, Narayana Kocherlakota, Per Krusell, Deborah Lucas, Pierre Mella-Barral, Bob Miller, Michael Parkin, Victor Ríos-Rull, Amlan Roy, Tony Smith, Paul Willen, Steve Zeldes, Stan Zin and three anonymous referees for helpful comments and suggestions. We have benefited from the support of NSF grant SES-9987602 and the Rodney White Center at Wharton. Valuable research assistance was provided by Xiaorong Dong. The current version of this paper is available at http://bertha.gsia.cmu.edu.

†Institute for International Economic Studies, Stockholm University and CEPR; storeslk@iies.su.se
‡Graduate School of Industrial Administration, Carnegie Mellon University; chris.telmer@cmu.edu
§Wharton School, University of Pennsylvania and NBER; yarona@savage.wharton.upenn.edu

1 Introduction

The interaction between cross-sectional risk and aggregate risk plays an increasingly important role in macroeconomics and finance. For example, Bernanke & Gertler (1989), Cooley, Marimon, & Quadrini (2001) and Carlstrom & Fuerst (1997) and others argue that financing constraints are an important propagation mechanism because they generate asymmetric effects of aggregate shocks on the the cross-sectional distribution of firms. Caballero & Hammour (1994), Davis & Haltiwanger (1992), Foote (1998), Tito (1996) and others examine cyclical dynamics in job creation and destruction. Rampini (2000) argues that incentive constraints associated with entrepreneurial activity give rise to countercyclical cross-sectional risk. Krebs (2002b), Krebs (2002a), Krusell & Smith (1999) and Storesletten, Telmer, & Yaron (2001) ask if countercyclical idiosyncratic risk is important for the welfare costs of business cycles. Finally, and most closely related to this paper, Alvarez & Jermann (2001), Balduzzi & Yao (2000), Brav, Constantinides, & Geczy (2002), Cogley (2002), Heaton & Lucas (1996), Krusell & Smith (1997), Lucas (1994), Lustig (2001), Marcet & Singleton (1999), Sarkissian (2001) and Storesletten, Telmer, & Yaron (1999) follow Constantinides & Duffie (1996) and Mankiw (1986) by asking if countercyclical idiosyncratic risk is important for asset pricing. A central question, therefore, is to what extent cyclical variation in the cross-sectional distribution is empirically plausible. A difficulty in asking this question involves data limitations. Even the Panel Study on Income Dynamics (PSID) — the panel data-set with the longest time dimension — covers only 4 or 5 business cycles.

This paper proposes an econometric technique to address this difficulty. Our method incorporates macroeconomic information dating back to 1930, in spite of the fact that our panel data begins in 1968. To understand what we do, consider two 'cohorts' of individuals, the first born in 1910 and the second in 1930. Everyone is subject to idiosyncratic labor-market shocks, some fraction of which are highly persistent. The conditional variance of these shocks is countercyclical, increasing during contractions and decreasing during expansions. Suppose that we have labor income data on each cohort when they are 60 years old. That is, we have 1970 data on the 1910 cohort and 1990 data on the 1930 cohort. What we will observe — given the high persistence and the countercyclical idiosyncratic risk — is more cross-sectional dispersion among the 1910 cohort than among the 1930 cohort. The reason is that the former worked through more contractionary years than the latter, including the Great Depression. More generally, we will see variation in the cross-sectional variance between any two cohorts of similar ages who have worked through different macroeconomic histories. This is the essence of our procedure. Even though the time-dimension of our panel data is limited to 1968-1993, we have a rich cross-section of ages in each year of the panel. We can therefore use macroeconomic data to characterize the working history of each household in the panel and then use cross-sectional variation between cohorts of similar ages to identify the cyclical idiosyncratic-risk effects we are after.

More specifically, we focus on the properties of household-level labor-market earnings from the PSID. We model idiosyncratic risk as a class of ARMA(1,1) processes with a regime-switching component in the conditional variance. The regime is identified by using macroeconomic data to classify each year between 1930 and 1993 as either a contraction or an expansion. A critical feature is finiteness: we make strong assumptions about initial conditions which allow us to base a GMM estimator on age-dependent moments.

We find robust evidence that idiosyncratic earnings risk is both highly persistent and countercyclical. Estimates of autocorrelation range from 0.94 to 0.96. Estimates of conditional standard deviations increase by roughly 75 percent (from 0.12 to 0.21) as the macroeconomy moves from expansion to contraction. We use graphical methods to make these GMM estimates transparent. We show that high autocorrelation is driven by a linearly-increasing pattern of cross-sectional variance with age.[1] Countercyclical volatility is driven by cohort effects in the cross-sectional variance, something which is clearly evident in a simple scatter plot.

The remainder of the paper is organized as follows. Section 2 presents our time-series model and Section 3 discusses our data. Section 4 presents a graphical analysis which shows that most of our results can be understood visually. Section 5 presents GMM estimates of the model in Section 2. Section 6 reconciles our estimates of variation in conditional variance with variation in the overall (age-pooled) cross-section. Section 7 concludes.

2 Time-Series Model for Idiosyncratic Risk

We begin with a time-series model for idiosyncratic labor earnings risk. Its most distinctive property — which will be emphasized in estimation — is that it depends explicitly on age. We denote Y_t as the logarithm of per-capita labor earnings at time t and y_{it}^h as the logarithm of labor earnings for household i of age h at time t. The latter is decomposed into two components:

$$y_{it}^h = g(x_{it}^h, Y_t) + u_{it}^h \quad . \tag{1}$$

The component $g(x_{it}^h, Y_t)$ incorporates aggregate earnings as well as x_{it}^h, deterministic components of household-specific earnings attributable to age, education and so on. The component u_{it}^h is the random component of a household's earnings which is idiosyncratic to them. It is identified by $E(u_{it}^h \,|\, Y_t, x_{it}^h) \equiv \tilde{E}_t(u_{it}^h) = 0$, for all dates t, where $\tilde{E}_t$ denotes the cross-sectional mean.

The specification of $g(x_{it}^h, Y_t)$ is critical for identifying u_{it}^h but, beyond that, it is not central. We defer its discussion until the estimation section. For u_{it}^h we use an ARMA(1,1) with a regime-switching conditional variance:

$$\begin{aligned} u_{it}^h &= \alpha_i + z_{it}^h + \varepsilon_{it} \\ z_{it}^h &= \rho z_{i,t-1}^{h-1} + \eta_{it} \quad , \end{aligned} \tag{2}$$

where $\alpha_i \sim \mathrm{Niid}(0, \sigma_\alpha^2)$, $\varepsilon_{it} \sim \mathrm{Niid}(0, \sigma_\varepsilon^2)$, $\eta_{it} \sim \mathrm{Niid}(0, \sigma_t^2)$,

$$\sigma_t^2 \;=\; \sigma_E^2 \;\text{ if }\; \text{aggregate expansion at date } t$$

[1] Deaton & Paxson (1994) have also documented this evidence using a different data source, the Consumer Expenditure Survey.

$$= \quad \sigma_C^2 \text{ if aggregate contraction at date } t \ ,$$

and $z_{it}^0 = 0$. The variable α_i is a 'fixed effect:' a shock received once at birth and then retained throughout life (hence the absence of a time subscript). The variables z_{it}^h and ε_{it} are persistent and transitory shocks, respectively (we make age h explicit only when the conditional distribution of a variable depends upon it). What we mean by 'countercyclical volatility' is that $\sigma_C > \sigma_E$. We do not impose this *a priori*.

We include fixed effects and transitory shocks so that we can better measure the primary objects of interest, σ_t^2 and ρ. For example, the way we will identify these parameters involves how the cross-sectional variance of u_{it}^h increases with age. However, much of the overall variation in u_{it}^h is common to households of all ages. Were we to exclude fixed effects α_i, we would overstate the magnitude of σ_t^2 and, therefore, overstate the amount of idiosyncratic risk that is relevant for economic decision making. Along similar lines, the transitory shocks ε_{it} are included to mitigate the extent to which measurement error (and 'true' transitory shocks) overstates our estimate of σ_t^2.

The essence of our approach lies in the initial conditions: $z_{it}^0 = 0$ and $\alpha_i \sim \text{Niid}(0, \sigma_\alpha^2)$. These are strong assumptions. They rule out, for example, time-variation in the distribution of α_i and dependence between α_i and, for a given i, subsequent realizations of z_{it}^h and ε_{it}. The benefits, however, are substantial. The initial conditions allow us to interpret equations (2) as a collection of finite processes and, therefore, allow us to condition on age. The implications are most apparent in the cross-sectional variances, for each age h, which will underly our GMM estimator:

$$\tilde{Var}(u_{it}^h) = \sigma_\alpha^2 + \sigma_\varepsilon^2 + \sum_{j=0}^{h-1} \rho^{2j}(I_{t-j}\sigma_E^2 + [1 - I_{t-j}]\sigma_C^2) \ , \tag{3}$$

where $\tilde{Var}$ denotes the cross-sectional variance and $I_t = 1$ if the economy is in an expansion at date t and $I_t = 0$ otherwise. If $|\rho| < 1$, the summation term converges to the familiar $\sigma^2/(1-\rho^2)$, where σ^2 is a probability weighted average of σ_E^2 and σ_C^2. For finite h, however, $\tilde{Var}(u_{it}^h)$ is increasing in h, at a rate determined by the magnitude of ρ. This, in conjunction with a (roughly) linearly increasing age-profile of sample moments for $\tilde{Var}(u_{it}^h)$, is what drives the relatively large estimate of ρ we will arrive at in Section 5.

Age dependence in equations (3) is also critical for how we assess countercyclical volatility. Consider, for instance, the cohort of 60 year old workers in the first year of our panel, 1968. According to equation (3), the cross-sectional variance for this cohort involves indicator variables I_{t-j} dating as far back as 1930 (we assume that $h = 1$ corresponds to a 23 year old worker). By classifying each year between 1930 and 1968 as either an expansion or a contraction, we can form history-dependent cross-sectional moments based on (3) which incorporate aggregate shocks dating back to 1930, far beyond the 1968-1993 confine of our panel. Doing so across all 26 years of our panel is what identifies σ_E and σ_C. To see this, compare the 60 year old workers in 1968 to the 60 year old workers in 1993. Each have worked for 38 years. However the 1968 cohort have worked during more contractions, including the Great Depression. The process (2) predicts that, if ρ is sufficiently large and

$\sigma_C > \sigma_E$, then the cross-sectional variance of a particular cohort will be increasing in the number of contractionary years that they worked through.[2] Therefore, inequality among the 1968 cohort should exceed that of the 1993 cohort. This is exactly what we document below and (anecdotally speaking) it is what drives our results on countercyclical volatility.

To summarize, we are primarily interested in autocorrelation, ρ, and countercyclical volatility, $\sigma_C > \sigma_E$. What will be pivotal for the former is how sample analogs of $\tilde{Var}(u_{it}^h)$ increase in h. This will also determine the *average* of σ_C and σ_E. What will be pivotal for the latter is variation in the cross-sectional variation between households of similar ages who have worked through different macroeconomic histories.

3 PSID Data

We use annual PSID data, 1968-1993, defined at the household level. We define 'earnings' as wage receipts of all adult household members, plus any transfers received such as unemployment insurance, workers compensation, transfers from non-household family members, and so on. Transfers are included because we wish to measure idiosyncratic risk *net* of the implicit insurance mechanisms which these payments often represent. Along similar lines, we study the household as a single unit so as to abstract from shocks which are insured via intra-household variation in labor force participation.

We depart from the common practice of using a longitudinal panel: an equal number of time series observations on a fixed cross-section of households. A longitudinal panel necessarily features an average age which increases with time. This is problematic for our approach which emphasizes restrictions between age and aggregate variation. For instance, were we to use a longitudinal panel for the years 1968-1993, a large fraction of the household heads will have been retired, or at least been in their late earning years, during the last 2 of only 5 business cycles witnessed during 1968-1993.[3] A longitudinal panel will also feature a relatively small sample size and be subject to survivorship bias in that only the relatively stable households are likely to remain in the panel for all 26 years.

We overcome these issues by constructing a sequence of overlapping three-year sub-panels. For each of the years 1968-1993, we construct a three year sub-panel consisting of households which reported strictly positive total household earnings (inclusive of transfers) for the given year and the next 2 consecutive years in the survey. For example, our 1971 sub-panel is essentially a longitudinal panel on 2,036 households over the years 1971, 1972 and 1973. Doing this for all years results in sequence of 24 overlapping sub-panels. The overall data structure contains more than enough time-series information to identify the parameters in equation (2), while at the same time survivorship bias is mitigated and the cross-sectional distribution of age is quite stable. The mean and standard deviation of the average age in

[2] It is not enough for ρ to simply be positive. Consider, for example, two cohorts of the same age, where one worked through many contractionary years early in life and the other worked through only one contraction, but late in life. If ρ is sufficiently small then the cross-sectional variance among the latter can be greater than among the former.

[3] Specifically, average age would increase from 39 to 64 over the years 1968-1993.

each sub-panel are 40.8 and 1.0, respectively. The mean and the standard deviation (across panels) of the number of households is 2449 and 322, respectively. Additional details, including a number of filters related to measurement error and demographic stability, are outlined in Appendix A.

4 Graphical Analysis

Most of our econometric results can be anticipated by simply looking at the data. The important dimensions are variation in the cross-sectional variance according to age and according to time. Loosely speaking, variation across age is what drives our estimate of autocorrelation, ρ, while variation across time, and how it relates to age, is what drives our estimates of σ_E and σ_C.

To show this, we use a more reduced-form representation of the data — a dummy variable regression similar to that in Deaton & Paxson (1994) — and ask what it suggests about our model's parameters. Denoting the cross-sectional variances in equation (3) as $\sigma^2_{h,t} \equiv \tilde{Var}(u^h_{it})$, we decompose their variance into cohort and age effects:

$$\sigma^2_{h,t} = a_c + b_h + e_{h,t} \quad , \tag{4}$$

where $c = t-h$ denotes a 'cohort' (*i.e.*, birth year).[4] The parameters a_c and b_h are cohort and age effects and $e_{h,t}$ are residuals. Our model represents restrictions on the age/time/cohort effects captured in the a_c and b_h parameters. We impose and test them in Section 5. Here, we show what the unrestricted estimates imply about the parameters in (3).

We begin with age. Figure 1-A plots the age effects, b_h, from the regression (4).[5] The graph shows that variance among the young is substantial and that it increases by a factor of 2.3 between ages 23 and 60. Moreover, the increase is approximately linear. Deaton & Paxson (1994) report similar results using data from the Consumer Expenditure Survey. Inspection of equation (3) suggests that the initial dispersion can identify $\sigma^2_\alpha + \sigma^2_\varepsilon$, the rate of increase can (loosely speaking) identify the average of σ_E and σ_C, and the linear shape can identify ρ. The latter suggests a value of ρ very close to unity. Section 5 formalizes this and shows that the near-unit-root implications are robust to the incorporation of information not represented in the graph, most notably autocovariances.

Figure 1-B plots the cohort coefficients, but in a different fashion. An implication of our model is that, controlling for age (as equation (4) does), a cohort that has worked through more contractions than another should on average exhibit greater cross-sectional dispersion (as long as ρ is large and $\sigma_C > \sigma_E$). Figure 1 shows that this is a characteristic of our panel. The cohort coefficients a_c are plotted against the fraction of years during the

[4] As is well known, an alternative decomposition is age and time effects. We choose cohorts effects because, as Figure 1-B will show, they are closely associated with the essence of our procedure. Moreover, Section 6 deals with time effects, albeit in a somewhat different fashion.

[5] We plot the variance instead of the standard deviation because it will clearly delineate the case of $\rho = 1$, which implies linearity for the former but not for the latter. The coefficients b_h are scaled so the graph passes through the unconditional (across cohorts) variance of agents of age 40.

working life of each cohort that were contractions (defined as the NIPA measure of GDP growth being below average for the elders of each cohort). Although there are a number of cohorts who worked through relatively few contractions and nevertheless exhibit substantial cross-sectional dispersion (*i.e.*, the cluster of points in the south-west corner of the graph), the positive relationship is apparent. The OLS slope coefficient is 0.94 with a standard error of 0.10. This relationship between cross-sectional variance and macroeconomic history is the main reason that our conclusions regarding variation in σ_t^2 differ from those of previous studies, in particular Heaton & Lucas (1996) who found relatively small effects.

In Figure 1-C we move beyond the variance decomposition (4) and impose some structure from our model. Suppose that $\rho = 1$. Then the cross-sectional variances (3) imply

$$\begin{aligned} \sigma_{h,t}^2 &= \sigma_\alpha^2 + \sigma_\varepsilon^2 + (h - n_{h,t})\sigma_E^2 + n_{h,t}\sigma_C^2 \\ \Longrightarrow \frac{\sigma_{h,t}^2}{h} &= \frac{\sigma_\alpha^2 + \sigma_\varepsilon^2}{h} + \sigma_E^2 + f_{h,t}(\sigma_C^2 - \sigma_E^2) \ , \end{aligned} \tag{5}$$

where $n_{h,t}$ is the number of contractions that households of age h in panel-year t have worked through, and $f_{h,t}$ expresses $n_{h,t}$ as a fraction of age h. We set $\sigma_\alpha^2 + \sigma_\varepsilon^2 = 0.3$, using the intercept from Figure 1-A. Figure 1-C, is a scatter-plot of sample moments of $h^{-1}(\sigma_{h,t}^2 - 0.3)$ versus $f_{h,t}$. Although the data are obviously much noisier than Figure 1-B — we are now essentially plotting the raw data from the regression (4) — a positive relationship is apparent. The OLS slope coefficient is 0.023 which, according to equation (5), is an estimate of countercyclical volatility, $(\sigma_C^2 - \sigma_E^2)$. This is close to the value that we will estimate in Section 5.

Figure 1-D provides one last piece of evidence on countercyclical volatility. It pools the age effects highlighted in Figure 1-A and plots the year-by-year 'pooled' cross-sectional standard deviation in the PSID. It also plots the cross-sectional mean. Both the mean and the standard deviation are linearly detrended. Even at this informal level we see striking evidence of countercyclical volatility. The correlation between the mean and the standard deviation is -0.74. The magnitude of the changes, however, must be interpreted with caution. In Section 6 we examine the pooled cross-sectional variance in more detail. We show that, because it is a close cousin of the unconditional variance, it will always understate (time) variation in what we are ultimately interested in: the conditional variance σ_t^2.

5 Estimation

Recall that equation (1) decomposed log earnings into two components, $y_{it}^h = g(x_{it}^h, Y_t) + u_{it}^h$, where $g(x_{it}^h, Y_t)$ captures deterministic cross-sectional variation x_{it}^h and aggregate variation Y_t, and u_{it}^h captures idiosyncratic cross-sectional variation. We specify the first component as

$$g(x_{it}^h, Y_t) = \theta_0 + \theta_1^\top D(Y_t) + \theta_2^\top x_{it}^h \tag{6}$$

where $D(Y_t)$ is a vector of year-dummy variables, $t = 1968$ to 1993, and x_{it}^h is a vector comprised of age, age squared, age cubed, and educational attainment for household i

of age h at date t. Educational attainment is measured as the number of school years completed by the household head. Table 1 reports estimates of θ. The estimates and fit are quite similar to existing studies, implying a concave earnings function in age and education (*e.g.*, Hubbard, Skinner, & Zeldes (1994)).

Equation (6) allows us to identify the idiosyncratic component of income, u_{it}^h, which follows the process (2). We estimate its parameters using GMM. For the sake of transparency, we begin with an exactly-identified system which corresponds closely with the graphical analysis of the previous section. Afterwards, we add overidentifying restrictions. The cross-sectional variances, equations (3), represent one variance for each age/time pair, (h, t):

$$\tilde{E}_t \left[(u_{it}^h)^2 - (\sigma_\alpha^2 + \sigma_\varepsilon^2) - \sum_{j=0}^{h-1} \rho^{2j} (I_{t-j}\sigma_E^2 + [1 - I_{t-j}]\sigma_C^2) \right] \quad = \quad 0 \quad , \quad \text{for all } h, t \tag{7}$$

where h runs from 23 to 60 and t runs from 1968 to 1993. We do not have sufficient data to formulate sample moments for each h, t-pair. We therefore form age-cells of width 3 years and interpret the age of a household in a particular cell as the midpoint. Furthermore, in order to mimic the graph in Figure 1-A (*i.e.*, one variance per age), we aggregate the moments (7) over time and form sample analogs:

$$\frac{1}{T} \sum_{t=1}^{T} \frac{1}{N_{ht}} \sum_{i=1}^{N_{ht}} \left[(u_{it}^h)^2 - (\sigma_\alpha^2 + \sigma_\varepsilon^2) - \sum_{j=0}^{h-1} \rho^{2j} (I_{t-j}\sigma_E^2 + [1 - I_{t-j}]\sigma_C^2) \right] \quad = \quad 0 \quad , \tag{8}$$

where N_{ht} is the sample size of households aged h at time t, and h is now to be interpreted as an age-cell.

Panel A of Table 2 reports exactly-identified estimates of ρ, σ_C, σ_E and $\sigma_\alpha + \sigma_\varepsilon$ based on equations (8), using ages 25, 35, 45 and 55. Since cross-sectional variances alone cannot distinguish σ_ε and σ_α, we report their sum. The indicator variables I_t are identified by classifying each year, 1930-1993, as expansion or contraction according to whether the growth rate in U.S. GNP per-capita was above or below its sample mean (alternative definitions are reported later). Appendix A provide specific details of all GMM estimators, including how we construct the weighting matrix to incorporate differing sample sizes N_{ht}, the overlapping structure of our panel, the first-stage regression, and several other issues.

The results in Table 2 accord well with the graphs in Figure 1. Much of the overall level of dispersion gets attributed to some combination of fixed effects, σ_α and transitory shocks/measurement error, σ_ε. We estimate that $\sigma_\alpha^2 + \sigma_\varepsilon^2 = 0.56^2$, which corresponds closely to the intercept of $0.55^2 = 0.30$ from Figure 1-A. For the persistent shock z_{it}^h, we estimate an autocorrelation of $\rho = 0.963$, slightly lower that what we inferred graphically. For the conditional volatilities we estimate $\sigma_C = 0.162$ (contraction) and $\sigma_E = 0.088$ (expansion) which, on average, are consistent with the size of the increase in variance in Figure 1-A.[6] Viewed on their own, the estimates indicate a substantial amount of countercyclical volatility: an increases of 84% from peak to trough.

[6]Specifically, the conditional variance in Figure 1-A increases by 0.4 between ages 25 and 55. With a

Next, we identify σ_α from σ_ε using cross-sectional autocovariances analogous to the variances (7):

$$\tilde{E}_t \left[u_{it}^h u_{it-1}^{h-1} - \sigma_\alpha^2 - \rho \sum_{j=1}^{h-1} \rho^{2(j-1)} (I_{t-j}\sigma_E^2 + [1 - I_{t-j}]\sigma_C^2) \right] = 0 \ \text{ for all } h, t \ . \tag{9}$$

Again, we aggregate these moments over t for each age-cell h and then form sample counterparts analogous to equations (8). Panel B of Table 2 reports exactly-identified estimates obtained by adding just one moment from equations (9) to those underlying Panel A. We use $h = 40$ but the results are not sensitive to the particular age. The estimates of primary interest, ρ, σ_C and σ_E, change only marginally relative to Panel A. For the other parameters, the majority of the previously-combined variance gets attributed to fixed effects: $\sigma_\alpha = 0.44$ and $\sigma_\varepsilon = 0.35$ (these add-up to roughly the same variance as Panel A).

The estimates in Panel A and B of Table 2 are transparent in the sense that, given the graphical evidence of Section 4, it's clear which features of the data are driving them. The remainder of the table examines how robust these findings are once we add overidentifying restrictions and climb inside the usual black box of GMM. We use sample analogs of the (h, t)-pair variances and first-order autocovariances in equations (7) and (9), along with a set of second-order autocovariances analogous to (9). That is, we eliminate the time-aggregation in equations (8) (and Panels A and B) and just use the primitive moments.[7] Panel C of Table 2 uses the same definitions of 'expansion' and 'contraction' as Panels A and B (based on GNP growth), whereas Panels D, E and F use definitions based on NBER business cycle dates, the unemployment rate, and the stock market. We see slightly lower autocorrelation and, with the exception of Panel F, higher conditional volatility for the persistent shock. The latter comes at the expense of the transitory and fixed-effect shocks, where the combined volatility drops from roughly 0.32 to the neighborhood of 0.20. Apparently, the additional autocovariances have re-classified an important part of the variance from transitory/measurement error shocks into persistent shocks. Finally, the extent to which the persistent-shock conditional variance is countercyclical is relatively stable, ranging from a 70% increase (expansion to contraction) in Panels C and D to a 90% increase in Panel F.

To summarize, much of what we can infer graphically appears robust to the inclusion of information which is not represented in the graph. In particular, the age pattern in the cross-sectional variances suggests near unit-root behavior and this is robust to the inclusion of what are typically used to identify persistence: autocovariances.

homoskedastic unit-root process for the persistent shock, this implies a conditional variance of 0.4/(55-25) = 0.013, or a standard deviation of 0.114. The average of our point estimates, σ_C^2 and σ_E^2, is $(0.162^2 + 0.088^2)/2 = 0.017 = 0.13^2$.

[7] Specifically, Panels C through F use variances and first and second-order autocovariances for ages 25, 35, 45 and 55, over years 1968 through 1993. The ages were chosen to ensure a sample size of at least 100 households per (h, t)-pair. This makes for 312 moments and 307 degrees-of-freedom.

5.1 Consumption Data

To this point we've focused on labor earnings. For many questions, most notably asset pricing, we are equally interested in consumption and how its cross-sectional distribution is related to aggregate variation. Table 3 replicates the estimation in Panels C, D and F of Table 2 using data on food expenditure from the PSID (the only consumption data available). Our estimates of autocorrelation are lower, ranging from 0.86 to 0.91. Estimates of the fixed-effect variance are also lower, something we might expect given that we are limited to food data (*i.e.*, overall inequality in earnings is likely to exceed that of food consumption). Finally, estimates of the persistent-shock conditional variance still indicate countercyclical volatility, however the estimates are more sensitive to the definition of expansion/contraction, are somewhat lower on average, and the difference between expansion and contraction is muted, with the peak-to-trough increase in the neighborhood of 30% instead of 75%. The extent to which these estimates are indicative of the behavior of a broader definition of consumption, obviously, remains to be seen.

6 Temporal Variation in the Cross-Sectional Variance

Figure 1-B indicates that, over the period 1968-1993, the largest change in the age-pooled cross-sectional standard deviation of log earnings was about 10 percent. Similarly, the largest change from one year to the next was 4.4%. At first blush, this may seem inconsistent with our estimates, which indicate a change of roughly 75% over the course of the business cycle. The missing link is that the standard deviation in Figure 1-B is closely associated with the unconditional cross-sectional distribution of earnings, while the estimates in Table 2 are associated with the conditional distribution. In this section we ask whether the two can be quantitatively reconciled, thereby providing another corroborative check on our estimates.

The pooled cross-sectional variance in Figure 1-B is a weighted average of the conditional variances σ_E^2 and σ_C^2. This is true in two senses. First, equation (3) indicates that, in the usual AR(1) sense, the cross-sectional variance among agents of a given age h is a moving average of the variances of past innovations. The only wrinkle is that the terms in the moving average change over time, depending on the macroeconomic history which the cohort has lived through. Second, the pooled variance is a weighted average of the cohort-specific variances, where the weights are the population shares. Algebraically,

$$v_t \equiv \tilde{Var}(u_{it}) = \sum_{h=1}^{H} \varphi_h \left(\sigma_\epsilon^2 + \sigma_\varepsilon^2 + \sum_{j=0}^{h-1} \rho^{2j} (I_{t-j}\sigma_E^2 + [1 - I_{t-j}]\sigma_C^2) \right) \quad , \tag{10}$$

where φ_h are population shares and v_t denotes the pooled cross-sectional variance.

Inspection of equation (10) indicates that v_t is a moving average which varies between an upper and a lower bound. Given that $\sigma_C > \sigma_E$, the upper bound coincides with a long sequence of aggregate contractions (*i.e.*, $I_{t-h} = 0$ for all $0 < h \leq H$), and is equal to $\sigma_\epsilon^2 + \sigma_\varepsilon^2 + \sigma_C^2/(1-\rho^2)$ for large H and $\rho < 1$. Similarly, the lower bound coincides with a long sequence of expansions and is equal to $\sigma_\epsilon^2 + \sigma_\varepsilon^2 + \sigma_E^2/(1-\rho^2)$. The way that v_t varies between these extremes depends on persistence in both aggregate and idiosyncratic shocks, and the economy's demographic structure (the φ_h's). Figure 2 graphs this for one particular realization of the process (2), using U.S. data for the φ_h's. It plainly illustrates the main point of this section, that large changes in the variance of the conditional distribution are necessarily associated with smaller changes in the pooled cross-section.

We use this fact as another check on the plausibility of our estimates. We conduct a Monte Carlo experiment where we obtain a large number of replications of the process (2). Each replication is a panel with a time dimension matching our our PSID data (26 years) and a large number of individuals i in the cross-section. We define two test statistics. The first is based on the difference between the maximum and the minimum values of v_t which are observed in any given replication:

$$\max(v_t) - \min(v_t) \ . \tag{11}$$

The second is based on the maximal absolute change between any two adjacent years:

$$\max(|v_t - v_{t-1}|) \ . \tag{12}$$

We use the parameter estimates in Panel C of Table 2, with aggregate shocks following the transition matrix [2/3 1/3;1/3 2/3]. We obtain the following results for the two test-statistics. For the first, based on equation (11), the fraction of cases in which the maximal change across the 26 years exceeds the 10% value from Figure 1-B is 28%. For the second, based on equation 12, the fraction of cases in which the change from one year to the next exceeds 3% (3.5%) is 5% (< 1%).

To summarize, the estimated large changes in the conditional variance reported in Table 2 are quantitatively consistent with the relatively small changes we see in the pooled cross-section, Figure 1-B. This serves to put the validity of our estimates on firmer ground, especially to the extent that variation in the pooled cross-section is more easily and accurately measured.

7 Last Thoughts

Our main finding is that PSID data on idiosyncratic labor-earnings risk exhibit robust evidence of high persistence and countercyclical volatility. We now conclude by comparing our estimates to those of related work and offering some economic interpretation.

We estimate an autocorrelation of roughly 0.95. Many other papers also advocate high persistence. Abowd & Card (1989), Carroll (1997), Gourinchas & Parker (2002) and MaCurdy (1982) — prominent papers on labor-income dynamics — *restrict* labor-income to be a unit-root process. Hubbard *et al.* (1994) don't impose a unit root, but instead estimate an autocorrelation of 0.95. An exception is Heaton & Lucas (1996) who obtain estimates closer to 0.5.[8] What distinguishes our work is an emphasis on age and its relation to cross-sectional variance. The linearly increasing pattern suggests near unit-root

[8]The main differences between our approach and Heaton & Lucas's (1996) is that they do not condition

behavior. This is robust to the incorporation of what most other studies have emphasized: autocovariances. Our findings, therefore, represent new evidence in support of Hubbard *et al.*'s (1994) estimates as well as the restrictions imposed by the other papers.

At face value, our estimates of transitory shocks seem large. We get an estimated standard deviation of roughly 0.25 which, in 2002 dollars, translates to an annual standard deviation of $11,500 for the average worker in our panel, who earns just over $45,000. While we can't say how much of this represents measurement error and how much represents actual shocks, we can say that it is comparable to related work. Gourinchas & Parker (2002) and Hubbard *et al.* (1994), for example, obtain 0.21 and 0.17 respectively. Our estimates are somewhat higher, the most likely reason being additional measurement error in our relatively-broad selection criterion.

We estimate the standard deviation of fixed effects to be roughly 0.40. The average young worker earns just over $21,000 (in 2002 dollars), so this translates into inequality among the young with a standard deviation of roughly $8,800. That this is plausible is apparent in Figure 1-A where the standard deviation among the young is 0.55 (or $12,480). We get a lower number here because (a) the transitory and persistent shocks also contribute to the initial variance (see equation (3) with $h = 1$), and (b) the overidentifying moments in Panels C-D of Table 2 generate a reduction in overall variability, as the GMM estimator gives weight to moments not represented in Figure 1-A.

Our estimates of the conditional standard deviations of the persistent shock are roughly 0.21 in a contraction and 0.12 in an expansion. The frequency-weighted average is 0.17. Because the autocorrelation is almost unity, these are approximate estimates of the standard deviations of labor-income *growth rates*. Equivalently, they translate into contraction/expansion conditional standard deviations of $9,500 and $5,300, based on the average worker's 2002-dollar earnings of $45,000. By any measure, idiosyncratic risk is large. Related work reaches similar conclusions (relative to our average). Gourinchas & Parker (2002) estimate 0.17 for high-school graduates and 0.15 for their overall sample. Carroll & Samwick (1997) obtain almost the same numbers. Hubbard *et al.* (1994) obtain 0.16 for high-school graduates and 0.14 for college graduates. Our estimates, therefore, agree with related work on average, and also suggest that cyclical variation is an important part of the story.

Finally, how large are the shocks relative to each other? We present a variance decomposition based on total lifetime uncertainty. Consider an unborn individual who is to receive lifetime labor income according to equation (1). Using a constant discount factor (for simplicity), we compute the variance of the present value of their future income (in levels, not logs):

$$Var\left(\sum_{h=23}^{62} 1.04^{-h}\exp(y_{it}^{h})\right) \quad . \tag{13}$$

on age as we do, and they estimate one intercept term per-household whereas we make the stronger assumption that fixed effects are drawn from a single-parameter distribution. They also use the more traditional longitudinal panel whereas we use the sequence of overlapping panels described in Section 3.

We compute this variance by simulation.[9] We do so three times, first with only the fixed-effect shocks set to zero, second with only the transitory shocks set to zero and third with only the persistent shocks set to zero. We then compute the ratio of each of these three variances to the total variance in equation (13). We find that the fixed effects contribute 54.8% to the total, the persistent shocks contribute 44.5% and the transitory shocks contribute 0.7%. The 'size' of the fixed effects, then, is only slightly larger than that of the (relatively low conditional variance) persistent shocks, once the total working life is accounted for.

[9]To simulate, we use the parametric specifications in equations (1) (2) and (6), the parameter estimates in Tables 1 and 2-C, and a two-state Markov chain with transition matrix [2/3 1/3;1/3 2/3] for the process 'expansion/contraction' needed to compute the indicator variables in equations (2).

References

Abowd, J. & D. Card, (1989), On the covariance structure of earnings and hours changes, *Econometrica* 57, 411–455.

Alvarez, F. & U. Jermann, (2001), Quantitative asset pricing implications of endogeneous solvency constraints, *Review of Financial Studies* 14, 1117–1152.

Balduzzi, P. & T. Yao, (2000), Does heterogeneity matter for asset pricing, Unpublished manuscript, Boston College.

Bernanke, B. & M. Gertler, (1989), Agency costs, net worth, and business fluctuations, *American Economic Review* 79 (1), 14–31.

Brav, A., G. M. Constantinides, & C. Geczy, (2002), Asset pricing with heterogeneous consumers and limited participation: Empirical evidence, *Journal of Political Economy* 110, 793–824.

Caballero, R. & M. Hammour, (1994), The cleansing effect of recessions, *American Economic Review* 144, 1350–1368.

Carlstrom, C. & T. Fuerst, (1997), Agency costs, net worth, and business fluctuations: A computable general equilibrium analysis, *American Economic Review* 87 (5), 893–910.

Carroll, C., (1997), Buffer stock savings and the life cycle and permanent income hypothesis, *Quarterly Journal of Economics* 112, 1–55.

Carroll, C. & A. Samwick, (1997), The nature of precautionary wealth, *Journal of Monetary Economics* 40, 201–220.

Cogley, T., (2002), Idiosyncratic risk and the equity premium: Evidence from the consumer expenditure survey, *Journal of Monetary Economics* 29, 309–334.

Constantinides, G. M. & D. Duffie, (1996), Asset pricing with heterogeneous consumers, *Journal of Political Economy* 104, 219–240.

Cooley, T., R. Marimon, & V. Quadrini, (2001), Aggregate consequences of limited contract enforceability, Working Paper: New York University.

Davis, S. & J. Haltiwanger, (1992), Gross job creation and gross job destruction and employment reallocations, *Quarterly Journal of Economics* 107, 819–863.

Deaton, A. & C. Paxson, (1994), Intertepmoral choice and inequality, *Journal of Political Economy* 102, 437–467.

Foote, C., (1998), Trend employment growth and the bunching of job creation and destruction, *Quarterly Journal of Economics* 113(3), 809–834.

Gordon, R. J., ed., (1986), *The American Business Cycle: Continuity and Change*, University of Chicago Press.

Gourinchas, P. O. & J. Parker, (2002), Consumption over the life cycle, *Econometrica* 70 (1), 47–89.

Heaton, J. & D. J. Lucas, (1996), Evaluating the effects of incomplete markets on risk sharing and asset pricing, *Journal of Political Economy* 104, 443–487.

Hubbard, R. G., J. Skinner, & S. P. Zeldes, (1994), The importance of precautionary motives in explaining individual and aggregate savings, *Carnegie Rochester Conferance Series on Public Policy* 40, 59–126.

Krebs, T., (2002a), Growth and welfare effects of business cycles in economies with idiosyncratic human capital risk, Working Paper, Brown University.

Krebs, T., (2002b), Human capital risk and economic growth, Forthcoming, *Quarterly Journal of Economics*.

Krusell, P. & A. A. Smith, (1997), Income and wealth heterogeneity, portfolio choice, and equilibrium asset returns, *Macroeconomic Dynamics* 1, 387–422.

Krusell, P. & A. A. Smith, (1999), The welfare effects of eliminating business cycles, *Review of Economic Dynamics* 2, 245–272.

Lucas, D. J., (1994), Asset pricing with undiversifiable risk and short sales constraints: Deepening the equity premium puzzle, *Journal of Monetary Economics* 34, 325–341.

Lustig, H., (2001), The market price of aggregate risk and the wealth distribution, Working Paper, Stanford University.

MaCurdy, T. E., (1982), The use of time series processes to model the error structure of earnings in a longitudinal data analysis, *Journal of Econometrics* 18, 83–114.

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

Marcet, A. & K. J. Singleton, (1999), Equilibrium assets prices and savings of heterogeneous agents in the presence of portfolio constraints, *Macroeconomic Dynamics* 3, 243–277.

Ogaki, M., (1993), *Generalized Method of Moments: Econometric Applications*, chapter 17, 455–488, Elsevier Science Publisher, North-Holland, In *Handbook of Statistics*, G.S Maddala, C. R. Rao, and H. D. Vinod editors.

Rampini, A., (2000), Entrepreneurial activity, risk and the business cycle, Working Paper, Northwestern University.

Sarkissian, S., (2001), The cross-country consumption dispersion and the world business cycle, forthcoming, *Review of Financial Studies*.

Storesletten, K., C. I. Telmer, & A. Yaron, (1999), Asset pricing with idiosyncratic risk and overlapping generations, Working paper, Carnegie Mellon University.

Storesletten, K., C. I. Telmer, & A. Yaron, (2000), Consumption and risk sharing over the life cycle, NBER working paper 7995.

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

Tito, B., (1996), Is job turnover counter-cyclical?, *Journal of Labor Economics* 34, 603–625.

White, H., (1980), A heteroskedasticity-consistent covariance matrix estimator and a direct test of heteroskedasticity, *Econometrica* 48, 817–38.

A Data Appendix

Additional details regarding the data selection procedure, specific characteristics of the data and the estimation procedure based on equations (1) and (7) are as follows.

Selection Criteria

A household is selected into a panel if the following conditions are met: (i) head of households is no younger than 22 and not older than 60 (ii) total earnings are positive in each of the sample years, (iii) total earnings growth rates are no larger than 20 and no lower than 1/20 in any consecutive years, (iii) the households was not part of the Survey of Economic Opportunity.

Individual Data

The individual earnings data are based on the 1969-1994 Family Files of the PSID. We use the PSID's Individual Files to track individuals across the different years of the Family Files. The definition of earnings includes wage earnings by head of household plus female wage earnings plus total transfers to the household. Total transfers include unemployment insurance, workers compensation, transfers by non-household family members, and several additional (minor) categories. Total earnings are then deflated to common 1968 dollars using the CPI index. Earnings are then converted to per-household-member rates by dividing total earnings by family size.

Consumption data is based on total food consumption (the sum of expenditures of food for meals at home, meals eating outside, and value of food stamps used). These questions were not available for panel years 1973, 1988, and 1989.

3-year repeated panels

Agents are selected into a panel if they meet the selection criteria above for the two years following the base year of the panel. We start with the 1969 PSID file (and therefore follow agents through PSID files from 1970 and 1971) which we denote panel 1968, since each PSID file provides information on the previous year's income. Our last 3-year panel, which we denote as the 1991 panel, starts with the 1990 PSID file and ends with 1994 PSID file. Hence, we end up with 24 3-year repeated panels. The table below reports summary statistics of the 3-year repeated panels we use in our estimations.

Table A1
Summary Statistics of Panels

Year	# of Households	Age of Head Household	
		Mean	Std. Deviation
1968	1821	43.22	12.18
1969	1923	42.83	12.39
1970	1971	42.27	12.63
1971	2036	41.57	12.69
1972	2088	41.22	12.61
1973	2124	40.75	12.64
1974	2216	40.87	12.73
1975	2282	40.57	12.79
1976	2337	40.22	12.76
1977	2406	40.12	12.70
1978	2452	40.18	12.59
1979	2535	39.75	12.52
1980	2565	39.76	12.49
1981	2584	39.82	12.36
1982	2619	39.75	12.19
1983	2646	39.99	12.02
1984	2728	39.93	11.99
1985	2768	40.12	12.07
1986	2778	40.38	11.83
1987	2819	40.55	11.65
1988	2819	40.90	11.44
1989	2856	40.73	11.29
1990	2806	40.78	11.18
1991	2610	41.80	11.23

The entries correspond to 3 year panels. The entries under the year column correspond to the base year of each panel. The age statistics correspond to the head of the household.

Aggregate Data

The aggregate GNP levels used to define the indicator functions described in (7) were derived by merging GNP data from Gordon (1986) for the years 1910-1958 with CITIBASE data for the years 1959-1993. These were converted to 1968 dollars using the CPI and to per-capita terms by dividing by total population figures from CITIBASE.

The NBER based business cycles are derived from the monthly NBER definitions of contractions and expansions. We converted these monthly definitions into yearly ones. Our criteria, corresponding to the NBER definition, is to define years of recession, as those years for which the majority of the months are contractionary. For cases where contractions spanned less than 12 months but more than 6 months over two calendar years, we designated the first year as a recession. Our results are robust to close alternatives.

The 'Unemployment' definition of cycles is based on annual unemployment rates. This is constructed from two sources. For the years 1910-1970 we use Historical Statistics of the United States (page 135). For 1971-1993, we used the Economic report of the president (February 1999, page 376). To convert these rates into contractionary years we used information about the levels of unemployment as well as direction of changes. Recessions are defined as years for which unemployment rate was greater than 7.5%. If the rate was greater than 7.5% but fell more than 3% relative to the previous year, it was not counted as a recession year. In addition years for which the unemployment rate had risen by more than 3% were defined as recession years.

The 'Returns' definition is based on the real (CPI-deflated) Value Weighted Return from CRSP. We define contractionary years to be ones in which the real return is negative.

Estimation Procedure

Our estimation procedure has two distinct steps. In the first stage we estimate equation (1). In the second stage we estimate the system given in (7).

To recover u_{it}^h in the first stage, we need to specify $g(Y_t, x_{it}^h)$. We let $g(\cdot)$ be represented by a simple linear regression using as regressors year dummies (corresponding to Y_t) and age, age squared, age cubed, and education (corresponding to individual specific attributes x_{it}^h). Let the parameters of this first stage regression be summarized by θ_1. That is $\psi_1(y_{i,t}^h, Y_t, x_{it}^h, \theta_1) \equiv y_{it}^h - \theta_1'[1, Y_t, x_{it}^h]$, All of the coefficients are significant and the R^2 is 0.23. These are displayed in Table 1

Next, let the parameters of the system in (7) be denoted by θ_2 (that is ρ, σ_E, σ_C, σ_ϵ, σ_α). The joint system we estimate can be written compactly as

$$E \begin{bmatrix} \psi_1(y_{it}^h, Y_t, x_{it}^h, \theta_1) \\ \psi_2(y_{it}^h, Y_t, \theta_1, \theta_2) \end{bmatrix} = 0 \tag{14}$$

where ψ_1 and ψ_2 are the moment conditions corresponding to (1) and (7) respectively. The triangular structure of the moment condition allows us to get consistent estimates of θ_1 using only ψ_1. The weighting matrix is trivially the identity as this is always an exactly identified system. We then estimate θ_2 using moment conditions ψ_2. This second step incorporates the standard errors in estimating θ_1 using standard two-step GMM procedure – as in Ogaki (1993).

There are two additional complication that arise in our set-up one being the overlapping structure of our repeated panels when we use autocovariances, and the other the non-balanced panel.

For each moment condition based on say panel year t an MA(2) correction is added to the estimate of the covariance matrix associated with moment conditions ψ_2 .

Specifically define the empirical residuals of the moment conditions ψ_2 to be,

$$\begin{aligned}
\psi_{2,i,t}^{h,0} &\equiv [(u_{i,t}^{h,t})^2 - (\sigma_\varepsilon^2 + \sum_{j=0}^{h-1} \rho^{2j}(I_{t-j}\sigma_C^2 + [1 - I_{t-j}]\sigma_E^2))] \\
\psi_{2,i,t}^{h,1} &\equiv [u_{i,t}^{h,t} u_{i,t+1}^{h+1,t} - (\rho \sum_{j=1}^{h-1} \rho^{2(j-1)}(I_{t-j}\sigma_C^2 + [1 - I_{t-j}]\sigma_E^2))] \\
\psi_{2,i,t}^{h,2} &\equiv [u_{i,t}^{h,t} u_{i,t+2}^{h+2,t} - (\rho^2 \sum_{j=1}^{h-1} \rho^{2(j-2)}(I_{t-j}\sigma_C^2 + [1 - I_{t-j}]\sigma_E^2))]
\end{aligned}$$

where the superscript t in u denotes the base year of the panel from which this agent is selected. The superscript j denotes whether the moment is based on variances ($j = 0$) or first and second autocovariances ($j = 1, 2$). By assumption $\psi_{2,i,t}^{h,j}$ is not correlated with $\psi_{2,i,t+k}^{h,l}$ $\forall k \neq 0$, and $j, l = 0, 1, 2$. It can be easily shown that due to the overlap of the sample, for each t, $\psi_{2,i,t}^{h,0}$, $\psi_{2,i,t}^{h,1}$ and $\psi_{2,i,t}^{h,2}$ are correlated. We stack the repeated 3 moment conditions and use sample counterparts to estimate these covariance terms – where the covariance matrix is block-diagonal and each 3×3 block has non-empty off-diagonal elements. In practice our results do not seem to be sensitive to this correction.

Because of the block-diagonal structure described above, the non-balanced panel allows for the construction of standard GMM statistics. Let N^* be minimum sample size across all moments, and let other sample moments, say moment j be of size κ_j times N^*. The standard GMM asymptotics follow from assuming that sample sizes grow in these ratios. The covariance matrix used is constructed using each sample moment then scaled by its κ_j.

Table 1
First-Stage Regression

	constant	Age	Age^2	Age^3	Education	Family Size
Estimate	7.512	-0.013	0.147	-0.019	0.161	0.070
Std. Error	(0.087)	0.006	0.015	0.001	0.011	0.001

Entries are estimates of the parameters of the first-stage regression, Section 5, equation (6). The regression also includes year-dummy variables, which are not reported here. The regression R^2 is 0.23.

Table 2
Idiosyncratic Earnings Process: Parameter Estimates

ρ	σ_C	σ_E	σ_ϵ	σ_α	P-value
Panel A: Exactly Identified (no autocovariance)					
0.963	0.162	0.088	0.562	–	–
(0.05)	(0.06)	(0.05)	(0.09)	–	–
Panel B: Exactly Identified (with autocovariances)					
0.957	0.163	0.094	0.351	0.448	–
(0.06)	(0.03)	(0.04)	(0.06)	(0.094)	
Panel C: Overidentified ($H \times T$)					
0.952	0.211	0.125	0.255	0.378	0.899
(0.020)	(0.034)	(0.044)	(0.021)	(0.057)	
Panel D: Overidentified, NBER Indicators					
0.943	0.201	0.119	0.255	0.386	0.939
(–)	(0.041)	(0.022)	(0.025)	(0.064)	
Panel E: Overidentified, Unemployment Indicators					
0.938	0.246	0.138	0.257	0.308	0.459
(0.011)	(0.026)	(0.034)	(0.020)	(0.062)	
Panel F: Overidentified, CRSP Value-Weighted Return Indicators					
0.939	0.159	0.084	0.235	0.423	0.967
(0.037)	(0.031)	(0.019)	(0.011)	(0.019)	

Entries are GMM estimates of the process (2), based on PSID earnings data 1968-1993. The parameters σ_E^2 and σ_C^2 denote the conditional standard deviation of the persistent component, conditional on an aggregate expansion, E, or contraction, C. Estimates in Panels A,B, and C define expansion/contraction in terms of the growth rate of U.S. GNP being above/below its mean, those in Panels D are based on the NBER definitions, those in Panel E are based on U.S. unemployment data, and those in Panel F are based on equity-index returns. Panel A uses moments (7) with ages 25,35,45,55. Panel B adds one more parameter and one more moment, the first autocovariance, equation (10) for age 40. Panels C-F use an overidentified system by using equations (7) - (9) in addition to second-order autocovariances. Standard errors are computed using the White (1980) estimator and incorporate sampling uncertainty from the first-stage regression based, equation (1) (reported in Table 1). The p-value is for the overidentifying test of the moment conditions. Further details, including all data sources, are available in Appendix A.

Table 3
Idiosyncratic Consumption Process: Parameter Estimates

ρ	σ_C	σ_E	σ_ϵ	σ_α	P-value
		Panel A: Overidentified, $H \times T$			
0.893	0.222	0.172	0.283	0.267	0.999
(0.015)	(0.059)	(0.043)	(0.004)	(0.013)	(0.041)
		Panel B: Overidentified NBER Indicators			
0.862	0.172	0.133	0.288	0.246	0.999
(0.021)	(0.073)	(0.071)	(0.043)	(0.013)	
Panel C: Overidentified, CRSP Value-Weighted Return Indicators					
0.913	0.127	0.095	0.295	0.303	> 0.999
(0.050)	(0.016)	(0.023)	(0.005)	(0.057)	

Entries are GMM estimates, based on food consumption data from the PSID. We define a first stage regression as in (1) except that the dependent variable is now log food consumption. We then employ a methodology identical to that used Table 2, where we treat the residual as the idiosyncratic consumption process. The PSID does not report consumption data for the years 1973, 1988 and 1989.

Figure 1
Cross-Sectional Variance of Earnings

A: Cross-Sectional Variance by Age

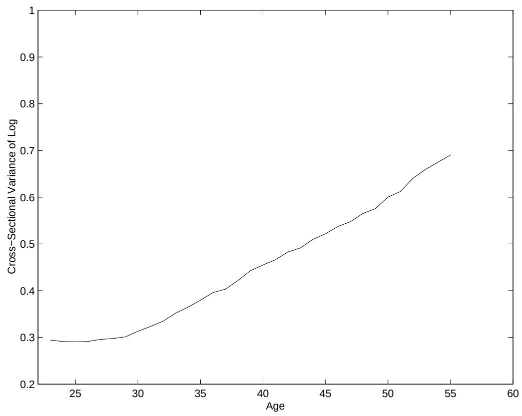

B: Macroeconomic History versus Cross-Sectional Variance (cohorts)

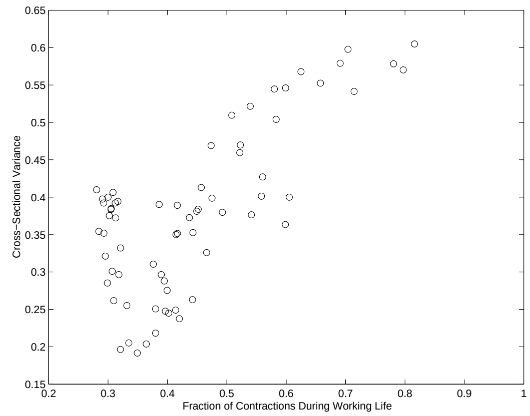

C: Macroeconomic History versus Cross-Sectional Variance

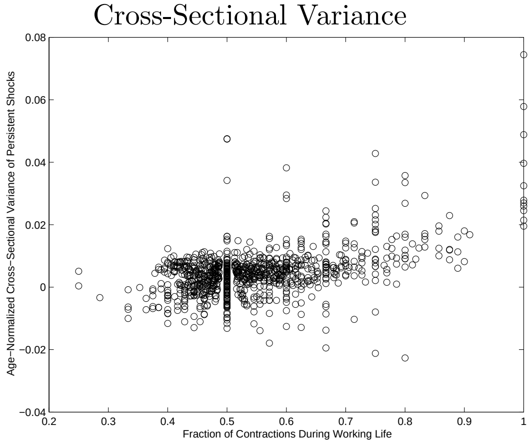

D: Cross-Sectional Moments By Time

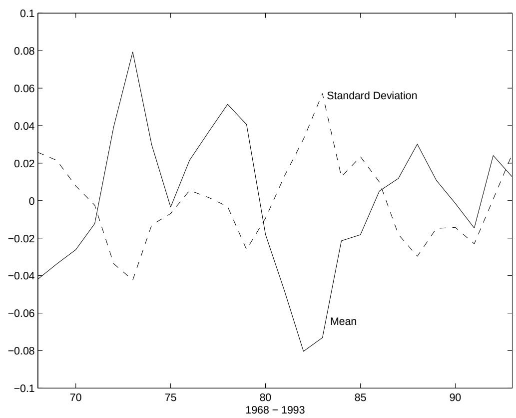

Cross-sectional variances are based on the idiosyncratic component (u_{it}^h from equation (2)) of log labor earnings plus transfers from the PSID, 1968-1993. The variances in Figure A control for 'cohort effects' by regressing cohort-age-specific cross-sectional variances on cohort-age dummy variables as in Deaton & Paxson (1994) and Storesletten, Telmer, & Yaron (2000). The points in the graph are the age coefficients, rescaled to match the level of variance at age 40. Figure B plots the cohort coefficients from the dummy-variable regression underlying Figure A against the fraction of contractionary years during which the oldest members of that cohort were of working age. Figure C plots the age-normalized cross-sectional variance of the persistent shock underlying equation (5) against the fraction of contractionary years during which members of that age panel-year worked through. Figure D reports the linearly-detrended cross-sectional mean of log income (y_{it}^h from equation (1)) and the linearly-detrended standard deviation of u_{it}^h, pooled across all ages, for each year 1968-1993. The standard deviation is additively scaled for graphical reasons. The correlation coefficient between the two series is -0.74. Figure B is robust to (i) alternative methods of detrending the mean and (ii) using the coefficient of variation instead of the standard deviation. Further details are given in the section 3.

Figure 2
Simulated Cross-Sectional Standard Deviation

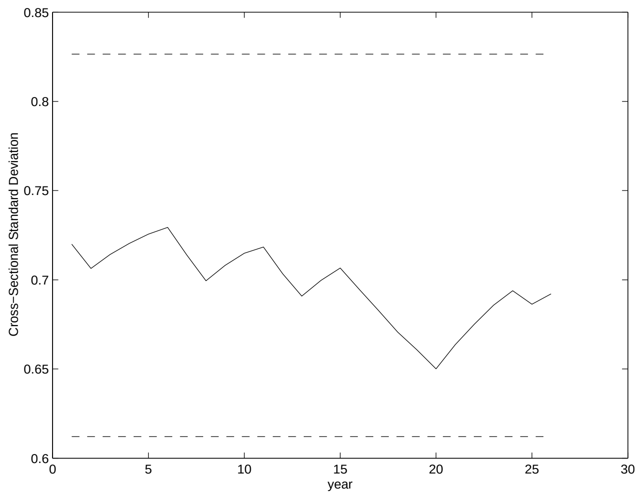

The solid line is one particular realization of the (population) pooled cross-sectional standard deviation of income shown in Section 6, equation (10), based on the process (2) from Section 2. Parameter estimates are from Panel C of Table 2. The upper line is the limit point, should all idiosyncratic shocks be drawn from the high variance (contractionary) distribution. The lower line is the analogous lower limit point. Further details are provided in section 6 of the text.