

Changes in the Distribution of Income Volatility

Shane T. Jensen and Stephen H. Shore^{*†‡§}

March 7, 2008

Abstract

Research documenting an increase in income risk in recent decades has treated its proxy, income volatility – the variance of income changes – as if it were the same for everyone. In reality, some people face more risk than others. We develop a Markovian hierarchical Dirichlet process (MHDP) prior to model the evolution of the full distribution of income volatilities, not just the mean. This augments the recently-developed hierarchical Dirichlet process (HDP) prior to accommodate the serial dependence of panel data. We find that increases over time in mean volatility can be attributed solely to increases in volatility at the right tail; the risky are getting riskier while the rest of the distribution remains roughly unchanged. This has strong welfare implications, as the high-volatility individuals facing increasing volatility are likely those with the highest tolerance for risk or the best-sharing opportunities, as evidenced by their willingness to take on income risk in the first place.

*Jensen: Wharton School, Department of Statistics, stjensen@wharton.upenn.edu. Shore: Johns Hopkins University, Department of Economics, shore@jhu.edu. Please contact Shore at: 458 Merghenthaler Hall, 3400 N. Charles Street, Baltimore, MD, 21218; 410-516-5564.

†JEL Classification: D31 - Personal Income, Wealth, and Their Distributions; C11 - Bayesian Analysis; C14 - Semiparametric and Nonparametric Methods.

‡keywords: Markovian hierarchical Dirichlet process, income risk, income volatility, heterogeneity

§We thank Dylan Small for helpful comments, as well as seminar participants at the Wharton School of the University of Pennsylvania.

Contents

1	Introduction	1
2	Data	2
3	Income process	6
4	Heterogeneity	8
5	Estimation	11
5.1	Step 1: Sampling income process parameters (θ, ϕ, γ^2)	13
5.2	Step 2: Sampling realized shocks (ω, ε)	14
5.3	Step 3: Sampling volatility parameters (σ^2, τ^2)	14
6	Results	17
6.1	Basic results	18
6.2	Evolution of the volatility distribution	19
6.3	Heterogeneity or fat tails?	21
6.4	Whose incomes are increasingly volatile?	26
7	Conclusion	28

List of Figures

1	Model Hierarchy	10
2	Distribution of Volatility Parameters	18
3	Impulse Response Function	19
4	Evolution of Percentiles of Volatility Distribution	23
5	Evolution of Persistent Volatility	24

List of Tables

1	Summary Statistics	3
2	Distribution of Income	4
3	Autocorrelation of Income Changes	5
4	Year-by-Year Income Volatility Sample Moments	7
5	Summary of Notation	11
6	Basic Model Results	17
7	Year-by-Year Income Volatility Parameters	20
8	Volatility Trends	22
9	Year-by-Year Volatility Sample Moments by Past Income Volatility	25
10	Volatility Trends by Group	27

1 Introduction

Idiosyncratic income risk is arguably the largest source of risk people face. To the degree that people are risk-averse, it may carry substantial welfare costs. Income volatility – the variance of income changes – can be used to measure income risk. There has been a great deal of recent interest by journalists and economists in changes over time in the amount of income risk.¹ Most but not all of the papers on this topic argue that income changes have gotten larger in recent decades or that large drops in income have become more likely. Some work (Hacker, 2006) argues that since people are risk-averse, this increase in risk is a bad thing.

To date, research on the evolution of income volatility has treated volatility as if it were a single number in any given year. In reality, income risk will not be the same for everyone, even among those with similar demographics. Volatility will vary not just over time but also across individuals. Population volatility in a given year is a *distribution* of parameters and not a single parameter. Effectively, previous research has been estimating the evolution of the population *mean*. Changes in mean may not capture all important changes in the distribution. Volatility will be increasing for some people and at some times and decreasing for others (Meghir and Pistaferri, 2004). Since these changes need not cancel out, the 5th, 50th, and 95th percentiles of the volatility distribution may evolve very differently.

In this paper, we provide non-parametric estimates of the evolving distribution of individuals' income volatility parameters. We find that increases in the mean of income volatility can be attributed solely to the right tail of the volatility distribution moving farther right. While people at this risky end of the distribution has gotten riskier, volatility has not changed very much for the median-volatility person (or for person in the 10th, 25th, or 75th percentile in terms of income volatility). This has important welfare implications. To the degree that those who are most risk-tolerant will take on the most income risk (Schulhofer-Wohl, 2006), an increase in volatility for those people may not be a bad thing.

Our finding is consistent with Dynan, Elmendorf, and Sichel (2007) who find that increasing income volatility has been driven by the increasing magnitude of extreme income changes. The median absolute income change is no larger than before. In its reduced form, this paper shows that the increasing magnitude of extreme income changes is borne largely by individuals who are *ex-ante* likely to have volatile incomes. These volatile individuals are identified by their past large absolute income changes. The variance of income changes has grown more for these individuals than for others. Consistent with this, self-employed individuals (who have more volatile incomes on average) have experienced much more rapid increases in income volatility than the population at large.

¹Dynan, Elmendorf, and Sichel (2007) provide an excellent survey of research on this subject in their Table 2, including Gottschalk and Moffitt (1994); Moffitt and Gottschalk (1995); Daly and Duncan (1997); Dynarski and Gruber (1997); Cameron and Tracy (1998); Haider (2001); Hyslop (2001); Gottschalk and Moffitt (2002); Batchelder (2003); Gosselin (2004); Hacker (2006); Comin and Rabin (2006); Gottschalk and Moffitt (2006); Hertz (2006); Winship (2007); Bollinger and Ziliak (2007); Bania and Leete (2007); Dahl, DeLeire, and Schwabish (2007); Shin and Solon (2008).

We use panel data on labor income for men from the PSID to estimate the parameters of a flexible income process. This income process incorporates both permanent and transitory shocks, and allows them to enter and damp out at a rate that we estimate. In the absence of heterogeneity, our iterative solution method to estimate income volatility parameters would be relatively straightforward. We alternate between a) using a Kalman filter to decompose income changes into permanent and transitory shocks (given model parameters) and then b) using a regression to estimate model parameters (given these shocks).

The chief methodological innovation in this paper is in our use of hierarchical Bayes, so that volatility parameters for a given individual at an instant in time are drawn from a distribution of values. We begin with a Dirichlet process (DP) prior model that allows the distribution of volatility to take a very general shape. In a DP model, the parameters – in this case parameters that characterize volatility, the variance of income changes – can take any one of N values, where these values and the number of them are determined by the data. To allow for the possibility that the observations for a given individual tend to be more similar to each other than to observations from other individuals, we use the recently-developed hierarchical Dirichlet process (HDP, see Teh, Jordan, Beal, and Blei, 2007) prior model. This is the first application of an HDP model in economics.

The DP and HDP models are not set up for panel data, where repeated observations for an individual have a natural order: time. We extend the existing DP and HDP models to allow for Markovian dependencies between consecutive years within each individual. Unlike these models, our Markovian hierarchical Dirichlet Process (MHDP) model allows the distribution of individuals' volatility parameters to evolve over time (or not) based on the data.

For each year, our method gives a posterior estimate of any individual's volatility parameters and also the distribution such of parameters in the population. It is then straightforward to look at evolution over time of any individual's volatility parameters and also the the distribution in the population. Our chief result is that mean volatility has increased since 1968, but that this can be attributed exclusively to increases in volatility at the volatile tail of the distribution.

2 Data

Data are drawn from the core sample of the Panel Study of Income Dynamics (PSID). The PSID was designed as a nationally representative panel of U.S. households. It tracked families annually from 1968 to 1997 and in odd-numbered years thereafter; this paper uses data through 2005. The PSID includes data on education, income, hours worked, employment status, age, and population weights to capture differential fertility and attrition. In this paper, we limit the analysis to men age 22 to 60; we use annual labor income as a measure of income.² Table 1 presents summary

²Labor income in 1968 is labeled v74 for husbands and has a constant definition through 1993. From 1994, we use the sum of labor income (HDEARN94 in 1994) and the labor part of business income (HDBUSY94), with a constant definition through 2005. Note that data is collected on

Table 1: Summary Statistics

	mean	st. dev.	min	max
year	1986.3	10.0	1968	2005
age (years)	40.0	10.5	22	60
education (years)	13.1	2.9	0	17
# of observations/person	17.2	9.0	1	34
married (1 if yes, 0 if no)	0.80	.	.	.
black (1 if yes, 0 if no)	0.05	.	.	.
annual income (2005 \$s)	\$50,553	\$57,506	0	\$3,714,946
annual income (\$s)	\$29,277	\$46,818	0	\$3,500,000
family size	3.1	1.5	1	14

This table summarizes data from 52,181 observations on 3,041 male household heads.

statistics from these data.

We want to ensure that changes in income are not driven by changes in the top-code (the maximum value for income entered that can be entered in the PSID). The lowest top code for income in real terms was \$99,999 in 1982 (\$202,281 in 2005 dollars), after which the top-code rises to \$9,999,999. So that top-codes will be standardized in real terms, this minimum top-code is imposed on all years in real terms, so the top-code is \$99,999 in 1982 and \$202,281 in 2005. We also want to ensure that results for the log of income are not dominated by small changes in the level of income near zero (which will imply huge or infinite changes in the log of income). To address this concern, we replace income values that are very small or zero with a non-trivial lower bound. We choose as this lower-bound the income that would be earned from a half-time job (1,000 hours per year) at the real equivalent of the 2005 federal minimum wage (\$5.15 per hour). This imposes a bottom-code of \$5,150 in 2005 and \$2,546 in 1982. Note that the difference in log income between the top- and bottom-code is constant over time. The vast majority of the values below this bound are exactly zero. This bound allows us to exploit transitions into and out of the labor force. At the same time, the bound prevents economically unimportant changes that are small in levels but huge in logs (because the level of income is close to zero) from dominating the results. Results are robust to other values for this lower bound, such as the income from full-time work (2,000 hours per year) at the 2005 minimum wage (in real terms).

In this paper, we model the evolution of “excess” log income. This is taken as the residual from a regression to predict the natural log of labor income (top- and bottom-coded as described). The regression is weighted by the PSID-provided sample weights, with the weights normalized so that the average weight in each year is the

household “heads” and “wives” (where the husband is always the “head” in any couple). We use data for male heads so that men who are not household heads (as would be the case if they lived with their parents) are excluded.

Table 2: Distribution of Income, Excess Log Income, and Income Changes for Men

	Real Income	Excess Log Income	One-Year Change	Five-Year Change
Mean	\$50,553 (\$48,867)	0	0.0017	0.0043
St. Dev.	\$57,506 (\$34,943)	0.7307	0.4870	0.6863
Observations	52,181	52,181	43,261	34,972
Minimum	\$0 (\$5,150)	-2.9325	-3.6877	-3.8361
5 th Percentile	\$668 (\$5,150)	-1.6283	-0.7323	-1.3046
25 th Percentile	\$26,174	-0.2964	-0.1089	-0.2126
50 th Percentile	\$42,887	0.1246	0.0134	0.0653
75 th Percentile	\$62,012	0.4601	0.1442	0.3072
95 th Percentile	\$113,500	0.9757	0.6673	0.9764
Maximum	\$3,714,946 (\$202,381)	2.6435	3.5862	4.0678

Table 2 describes the distribution of labor income for men in the PSID over the period from 1968 to 2005. See Section 2 for a detailed description of the income variable and the top- and bottom-coding procedure. Column 1 shows the distribution of real annual income for men. The numbers in parentheses are the values with top- and bottom-coding restrictions. Column 2 shows the distribution of “excess” log income, the residual from the regression of log labor income on the covariates enumerated in Section 2. Column 3 presents the distribution of one-year changes in excess log income. Column 4 repeats the results for column 3, but presents five-year changes instead of one-year changes.

same. We use as regressors: a cubic in age for each level of educational attainment (none, elementary, junior high, some high school, high school, some college, college, graduate school), the presence and number of infants, young children, and older children in the household, the total number of family members in the household, and dummy variables for each calendar year. Including calendar year dummy variables eliminates the need to convert nominal income to real income explicitly. While this step is standard in the income process literature, it is not necessary to obtain our results. The results to follow are qualitatively the same and quantitatively similar when we use log income in lieu of excess log income.

Table 2 presents data on the distribution of real annual income in column 1 (imposing top- and bottom-code restrictions in parentheses). While the mean real income is nearly identical with and without top- and bottom-code restrictions (\$50,553 versus \$48,867), these restrictions on extreme values reduce the standard deviation of real income from \$57,506 to \$34,943. Column 2 shows the distribution of “excess” log income. Since excess log income is the residual from a regression, its mean is zero. The inter-quartile range of excess log income is -0.30 to 0.46 .

Column 3 presents the distribution of one-year changes in excess log income. Naturally, the mean of one-year changes is close to zero. The inter-quartile range of one-year changes is -0.11 to 0.14 ; excess income does not change more than 11 to 14 percent from year to year for most individuals. However, there are extreme

Table 3: Autocorrelation of Changes in Excess Log Income

correlation	$y_{i,t} - y_{i,t-1}$
$y_{i,t-1} - y_{i,t-2}$	-0.290
$y_{H,t-2} - y_{H,t-3}$	-0.073
$y_{H,t-3} - y_{H,t-4}$	-0.018
$y_{H,t-4} - y_{H,t-5}$	-0.027

Autocorrelation of one-year changes in excess log incomes at one- through four-year lags. The distribution of these one-year changes are shown in Table 2.

changes in income, so the standard deviation of changes to log income (0.49) is far great than the inter-quartile range. This implies either that changes to income have fat tails (so that everyone faces a small probability of an extreme income change), or alternatively that there is heterogeneity in volatility (so that a few people face a non-trivial probability of an extreme income change). Unless a model is identified from parametric assumptions, these are observationally equivalent in a cross-section of income changes. However, heterogeneity and fat tails have different implications for the time-series of volatility, and we exploit these in the paper.

Column 4 repeats the results from column 3, but presents five-year excess log income changes instead of one-year changes. These long-term changes have only slightly higher standard deviations than the one-year change, 0.69 vs. 0.49, suggesting some mean-reversion in income. This mean-reversion can be seen explicitly in Table 3, which shows the autocorrelation of $y_{i,t} - y_{i,t-1}$, the autocorrelation in one-year changes in men's excess log incomes. One-year autocorrelation in income is strongly negative (-0.29); income tends to increase in the year following a decrease. This negative autocorrelation between consecutive one-year changes in income is a central feature of earnings dynamics data; any process for labor income must accommodate negative autocorrelation. Autocorrelation at longer lags is negative but small and decreasing. The non-zero values for higher-order lags are driven primarily by the inclusion of individuals without income. Such individuals leave the labor force (income falls) and may then subsequently re-enter the labor force several years later (income rises), causing negative serial autocorrelation at longer lags. Excluding these observations eliminates higher-order autocorrelation almost entirely. Section 3 presents a standard income process that accommodates these features of the data.

This paper looks at the evolution of income volatility, the variance of income changes. Before presenting and estimating a statistical model, we show how volatility-related sample moments have evolved over time. These are shown in Table 4. The first three columns of numbers examine a sample moment that captures the variance of permanent income changes. This moment is the individual-specific product of two-year changes in excess log income (for example, between years t and $t-2$) and the six-year changes that span them (for example, between years $t+2$ and $t-4$). Meghir and Pistaferri (2004) show that this moment identifies the variance of permanent

income changes (between years $t-2$ and t) under fairly general conditions, including the income process we use in Section 3. The final three columns present two-year squared changes in excess log income, a raw measure of income volatility. The first and fourth columns present sample means, the second and fifth columns present sample medians, and the third and sixth columns present 95th percentiles.³

The key thing to note is that mean income volatility (columns 1 and 4) has increased over time while the median (columns 2 and 5) has not. This divergence can be explained by an increase in the magnitude of large unlikely income changes (columns 3 and 6). While not framed in this way, these features of the data have been identified in previous research, including (Dynan, Elmendorf, and Sichel, 2007).

3 Income process

Here, we present a standard process for excess log income for individual i at time t (following Carroll and Samwick, 1997; Meghir and Pistaferri, 2004, and many others):

$$\begin{aligned}
 y_{i,t} &= p_{i,t} + \xi_{i,t} + e_{i,t} & (1) \\
 p_{i,t} &= p_{i,0} + \sum_{k=1}^{t-\Omega} \omega_{i,k} + \sum_{k=t-\Omega+1}^t \theta_{t-k} \omega_{i,k} \\
 \xi_{i,t} &= \sum_{k=t-\epsilon+1}^t \phi_{t-k} \varepsilon_{i,k}
 \end{aligned}$$

Excess log income ($y_{i,t}$) is the sum of permanent income ($p_{i,t}$), transitory income ($\xi_{i,t}$), and measurement error ($e_{i,t}$). Permanent income is initial income ($p_{i,0}$) plus the weighted sum of past shocks to permanent income through year t ($\omega_{i,k}, 0 < k \leq t$). Transitory income is the weighted sum of recent shocks to transitory income ($\varepsilon_{i,k}$), but with the additional constraint that the weights sum to one ($\sum_k \phi_k = 1$). The permanent shock, transitory shock, and measurement error are assumed to be mean zero as well as independent of one another, over time, and across individuals.

Subsequently, “permanent variance” and “permanent volatility” (we use these two phrases interchangeably) refer to the variance of permanent shocks, $\sigma_{i,t}^2 \equiv E[\omega_{i,t}^2]$; “transitory variance” and “transitory volatility” refer to the variance of transitory shocks, $\tau_{i,t}^2 \equiv E[\varepsilon_{i,t}^2]$; “variance parameters” and “volatility parameters” refer to $\tau_{i,t}^2$ and $\sigma_{i,t}^2$ jointly. “Noise variance” refers to the variance of measurement error, $\gamma^2 \equiv E[e_{i,t}^2]$. Measurement error could be subsumed into transitory income; it is kept separate only to accommodate our estimation strategy.

³ All use weights from the PSID. The first row shows whole-sample results. The second row shows the percent change in the mean, median, or 95th percentile over the sample. This is merely calculated as coefficient of a weighted OLS regression of the year-specific sample moment on a time trend, multiplied by the number of years (2005 – 1968) and divided by the whole-sample value in the previous row. The coefficient and t-statistic from this regression are shown just below. Year-by-year values are then shown.

Table 4: Year-by-Year Income Volatility Sample Moments

Sample Moment	Permanent Variance			Squared Change		
	Mean	Median	95 th %	Mean	Median	95 th %
Average	0.1091	0.0099	0.8264	0.3561	0.0314	2.0042
% Change 1970-2003	49%	15%	92%	110%	19%	143%
Slope (t-stat)	0.0015 (4.11)	0.0000 (0.52)	0.0205 (8.76)	0.0106 (11.96)	0.0002 (1.26)	0.0775 (11.18)
1970	.	.	.	0.1555	0.0210	0.7709
1971	.	.	.	0.1823	0.0229	0.8004
1972	0.0665	0.0059	0.4003	0.2142	0.0277	1.1276
1973	0.0786	0.0048	0.4423	0.2296	0.0269	1.1500
1974	0.0792	0.0054	0.5090	0.2324	0.0264	1.1059
1975	0.0986	0.0129	0.6243	0.2496	0.0380	1.2286
1976	0.0997	0.0179	0.6749	0.3124	0.0498	1.6006
1977	0.0933	0.0095	0.7058	0.2983	0.0316	1.8058
1978	0.0706	0.0062	0.5958	0.2751	0.0296	1.3344
1979	0.0838	0.0061	0.6415	0.2931	0.0269	1.6711
1980	0.1388	0.0115	0.9270	0.2811	0.0292	1.4495
1981	0.1159	0.0123	0.8844	0.2932	0.0296	1.5200
1982	0.1004	0.0150	0.7256	0.2514	0.0305	1.2840
1983	0.0859	0.0150	0.6630	0.2912	0.0330	1.5820
1984	0.1220	0.0126	0.8786	0.3185	0.0331	1.8609
1985	0.1109	0.0118	0.7869	0.3283	0.0370	1.7499
1986	0.1002	0.0110	0.6905	0.3089	0.0358	1.5483
1987	0.1089	0.0093	0.7739	0.3015	0.0295	1.6058
1988	0.1224	0.0087	0.7969	0.3121	0.0300	1.6476
1989	0.1161	0.0077	0.8171	0.3278	0.0276	1.8996
1990	0.1174	0.0091	0.7770	0.2998	0.0261	1.5937
1991	0.1312	0.0121	0.9905	0.3523	0.0309	1.8485
1992	0.1013	0.0111	0.9119	0.3168	0.0295	1.7572
1993	0.1272	0.0112	1.0935	0.4166	0.0333	2.3561
1994	0.1083	0.0104	0.9270	0.4479	0.0347	2.6530
1995	0.1346	0.0077	1.1290	0.4914	0.0333	3.3055
1996	.	.	.	0.4768	0.0264	3.1923
1997	0.0898	0.0074	0.8660	0.4671	0.0282	2.9644
1999	0.1142	0.0080	0.9632	0.4539	0.0317	2.7189
2001	0.1190	0.0073	1.1174	0.4463	0.0271	2.9567
2003	0.1487	0.0182	1.2951	0.6348	0.0574	3.9098

The year t permanent variance is the product of two-year changes in excess log income (from $t - 2$ to t) and the six-year changes that span them (from $t - 4$ to $t + 2$). The year t squared change is from $t - 2$ to t . The first row shows full sample moments. The second row shows the percent change over the sample, calculated as the coefficient of a weighted OLS regression of year-specific sample moments on a time trend, multiplied by the number of years (2005-1968) and divided by the full sample moment. The coefficient and t-statistic are shown below.

Here, permanent shocks come into effect over Ω periods, and transitory shocks fade completely after ϵ periods. As an example of our notation, θ_2 denotes the weight placed on a permanent shock from two periods ago in current log income; ϕ_2 denotes the weight placed on a permanent shock from two periods ago in current excess log income. While we use the word “shock” for parsimony, these innovations to income may be predictable to the individual, even if they look like shocks in the data.

While this model provides a useful stylized process for income, income does not evolve exactly according to this process. Most obviously, labor income will be zero – and log income undefined – when an individual is not in the labor force; transitions into and out of employment the labor force are not modeled. While bottom-codes for income allow zeros, the income process does not account for them explicitly.

A challenge when taking any income process to data is that income data are missing for some individuals in some years. PSID surveyors are occasionally unable to contact study participants. More pervasively, the PSID did not collect data in even-numbered years following 1997. Since our estimation procedure is not set up to explicitly handle missing data, we will fill in missing observations with bootstrapped guesses of income. We examine the two-year change in excess log income that spans any single-year of missing data. We identify the set of two-year excess log income changes with a similar magnitude elsewhere in the data and select one at random.⁴ We drop individuals with longer spans of missing data.

4 Heterogeneity

Note that the permanent and transitory variance ($\sigma_{i,t}^2, \tau_{i,t}^2$) describe the magnitude of income changes for an individual (i) at a point in time (t). These volatility parameters can and will differ across individuals and over time.⁵ To date, the literature on the evolution of income volatility has assumed homogeneity at any point in time, $\sigma_{i,t}^2 = \sigma_t^2$ and $\tau_{i,t}^2 = \tau_t^2$. While a few papers have considered heterogeneity in income processes across individuals or over time, these were not designed to estimate an evolving distribution of parameters. (Alvarez, Browning, and Ejrnaes, 2001; Meghir

⁴This bootstrapped draw has an intermediate value which is used to fill in the missing data. For example, if excess log income for a given individual is 0.1 in 1999 and 0.5 in 2001 and (since the PSID did not gather data in the intervening year) missing in 2000. From the set of all sample observations with two-year excess log income changes of roughly 0.4, we select one at random. In general, this observation will be drawn from a different individual than the one with the missing data. Imagine that the individual-years drawn at random have excess log incomes of 0.6, 0.7, and 1.0 in 1972, 1973, and 1974, respectively. We then fill in the original individual’s missing data in 2000 with 0.2 (0.1+0.7-0.6).

⁵In this paper, we assume that the noise variance (γ^2) that describes measurement error (e) is homogeneous; γ^2 is the same for all individuals and at all times. Measurement error could be subsumed into transitory income (the transitory variance can vary across individuals and over time); we separate them only accommodate our estimation strategy. As a result, restricting γ^2 to be homogenous does not preclude the possibility of differences across groups or over time in what another researcher might refer to as measurement error; this will merely be identified as variation in what we call the transitory variance.

and Pistaferri, 2004)

We could relax the assumption of homogeneity while imposing strong parametric assumptions, for example that $(\sigma_{i,t}^2, \tau_{i,t}^2)$ are jointly drawn from an inverse Wishart distribution. The problem with such an approach is that it imposes a particular parametric shape on the distribution of volatilities that cannot change over time. This would preclude us from looking for, much less finding, changes in the shape of the volatility distribution.

Instead, we develop a new methodology based on recent advances in non-parametric Bayesian modeling. Specifically, we model $(\sigma_{i,t}^2, \tau_{i,t}^2)$ pairs as being drawn from a completely unknown (non-parametric) distribution. Sharing of information across individuals and across years is accomplished by using a Dirichlet process (DP) prior distribution. In a DP prior model, observations are placed in clusters, with all observations in the same cluster having the same parameters. The number of clusters and the parameters that describe each cluster are not known *a priori* but rather identified from the data. This allows volatility parameters for a given individual in a given year to be identical to values for other person-years, or to take on unique values.

While the DP prior model can be used to approximate any other distribution, it does have an implicit structure. For a given observation, the probability of taking on particular parameter values is proportional to the number of other observations with those particular parameter values. A common metaphor-driven name for this model is the “Chinese restaurant process”. When the first patron arrives at an empty Chinese restaurant, she chooses a seat at her preferred table. When the next patron arrives, he must choose whether to sit with the first patron, or to sit at an unoccupied table. The third patron arrives and decides to sit with one of the first two, or sit at an unoccupied table, and so on. People entering this restaurant prefer popular tables, and so are more likely to join a table that already contains many patrons. In this case, each patron refers to an individual in a year and tables correspond to possible values of the volatility parameters for that person-year.⁶

The standard DP model is appealing because it allows the distribution of volatility to take a general shape. However, this one-level model is inadequate for our purposes because it does not model systematic differences between people in income volatility. The standard DP model assumes that all pairs of observations are *a priori* equally likely to have the same parameters, regardless of whether or not they come from the same person or from different people. This shortcoming is overcome in the two-level hierarchical Dirichlet process (HDP) model, recently introduced in Teh, Jordan, Beal, and Blei (2007). This generalization of the standard DP model allows groups of observations (in our case, multiple observations from one individual) to be *a priori* more likely to share parameters. Unfortunately, the HDP model cannot account for panel data, where an individual’s observations are naturally ordered by time. In an HDP model, all pairs of observations from the same person are *a priori* equally likely to share parameters, regardless of whether or not they are close together in time. In other words, that model cannot accommodate autocorrelation.

⁶See Muller and Quintana (2004) for a general introduction to Bayesian non-parametric modeling with the Dirichlet process.

Figure 1: Model Hierarchy




Diagram describes evolution of volatility parameters. The numbers 1, 2, and 3 in circles at each decision node correspond to the levels of the hierarchy described on page 10. The numbers (7) through (12) identify the equation number giving the probability of reaching that branch.

We allow for autocorrelation by extending the DP model to a three-level Markovian hierarchical DP (MHDP) model, illustrated in Figure 1. In our model, each individual begins with a given volatility (or more precisely a pair of volatility parameters, $\sigma_{i,t}^2$ and $\tau_{i,t}^2$), the distribution of which may come from the standard DP process. In subsequent years, the volatility parameter will evolve according to this three-level hierarchy (also shown in Figure 1, with the number of the paper's corresponding equation in parentheses):

- Level 1 A volatility parameter can remain unchanged from the *previous period* (7) or can change (8). If the volatility parameter changes, it can change to;
- Level 2 A parameter from the set of *other values for that individual* (9) or can take on new value (a value the individual has not had before) (10). If the volatility parameter takes on a new value, this value can be;
- Level 3 A volatility parameter held by *other individuals* (11) or can be a new value not shared with other individuals (12).

Level 3 of our MHDP model would also be found in a DP or HDP model. Here, an observation may or may not take on the same parameter value as other observations in the sample. Level 2 of our MHDP model would also be found in an HDP model (but not a standard DP model). Here, an observation may take on the same parameter value as other observations for that individual. Level 1 is unique to our model. Here, an observation may remain unchanged, namely take on the same parameter as the immediately previous observation. Allowing a Markovian dependence on the

Table 5: Summary of Notation

0. Data used to estimate parameters and shocks:

\mathbf{y} (N by $T + 1$ matrix) : Excess log incomes

1. Parameters, constant over time and across individuals:

$\boldsymbol{\theta}$ (Ω by 1 matrix): Rate at which permanent shocks ($\omega_{i,t}$) enter

$\boldsymbol{\phi}$ (ϵ by 1 matrix): Rate at which transitory shocks ($\varepsilon_{i,t}$) damp out

γ^2 : Noise variance, homogeneous variance of measurement error ($e_{i,t}$)

2. Shocks:

$\boldsymbol{\omega}$ (N by T matrix) : Realized permanent shocks ($\omega_{i,t}$)

$\boldsymbol{\varepsilon}$ (N by T matrix) : Realized transitory shocks ($\varepsilon_{i,t}$)

3. Parameters, vary over time and across individuals:

$\boldsymbol{\sigma}^2$ (N by T matrix) : Variances of permanent shocks ($\omega_{i,t}$)

$\boldsymbol{\tau}^2$ (N by T matrix) : Variances of transitory shocks ($\varepsilon_{i,t}$)

See Section 5 for details on these variables. Vectors and matrices are written in bold.

immediately previous state is a novel methodological development that allows us to consider time-series or panel data.

5 Estimation

The model uses annual data on excess log income for N individuals over T years. Our results include estimates of the parameters that describe the income process, the permanent and transitory shocks that drive changes in income, and the permanent and transitory volatility parameters for each individual in each year. These are summarized in Table 5 and described in detail below. Vectors and matrices are written in bold:

0. Data (\mathbf{y}): We use \mathbf{y} to denote the N by T matrix of excess log incomes where the element in the i -th row of the t -th column is the excess log income for individual i in year t , $y_{i,t}$. This matrix will be ragged because not all individuals have data in all years.

1. Income Process Parameters ($\boldsymbol{\theta}, \boldsymbol{\phi}, \gamma^2$): We use $\boldsymbol{\theta}$ to denote the Ω by 1 vector whose k -th element is θ_k , the rate at which permanent shocks enter into income; we use $\boldsymbol{\phi}$ to denote the ϵ by 1 vector whose k -th element is ϕ_k , the rate at which transitory shocks damp out of income. These are constant over time and across individuals. The scalar noise variance, γ^2 , is the variance of measurement error.

2. Shocks ($\boldsymbol{\omega}, \boldsymbol{\varepsilon}$): We use $\boldsymbol{\omega}$ to denote the N by T matrix of realized (but not directly observed) permanent shocks where the element in the i -th row of the t -th

column is the permanent shock for individual i between years $t - 1$ and t , $\omega_{i,t}$. We use ε to denote the N by T matrix of realized (but not directly observed) transitory shocks where the element in the i -th row of the t -th column is the transitory shock for individual i between years $t - 1$ and t , $\varepsilon_{i,t}$. Like \mathbf{y} these matrices will be ragged; without data there is no shock to decompose into permanent and transitory components.

3. Volatility Parameters σ^2 to denote the N by T matrix of permanent variances where the element in the i -th row of the t -th column is $\sigma_{i,t}^2$, the permanent variance of individual i between years $t - 1$ and t ; we use τ^2 to denote the N by T matrix of transitory variances where the element in the i -th row of the t -th column is $\tau_{i,t}^2$, the transitory variance of individual i between years $t - 1$ and t . The rows of σ^2 and τ^2 , σ_i^2 and τ_i^2 respectively, represent the distribution of volatility parameters $\sigma_{i,t}^2$ and $\tau_{i,t}^2$ over time t for an individual i ; the columns of σ^2 and τ^2 , σ_t^2 and τ_t^2 respectively, represent the distribution of volatility parameters ($\sigma_{i,t}^2, \tau_{i,t}^2$) at time t across all individuals i . This paper aims to identify the evolution of the permanent and transitory variance (σ^2, τ^2) over time. Like \mathbf{y} , these matrices will be ragged.

We take a fully Bayesian approach to the estimation of these parameters. We estimate the joint posterior distribution of all unknown parameters conditional on our observed data as:

$$p(\boldsymbol{\theta}, \boldsymbol{\phi}, \gamma^2, \boldsymbol{\omega}, \boldsymbol{\varepsilon}, \boldsymbol{\sigma}^2, \boldsymbol{\tau}^2 | \mathbf{y}) \propto p(\mathbf{y} | \boldsymbol{\theta}, \boldsymbol{\phi}, \gamma^2, \boldsymbol{\omega}, \boldsymbol{\varepsilon}) \cdot p(\boldsymbol{\omega}, \boldsymbol{\varepsilon} | \boldsymbol{\sigma}^2, \boldsymbol{\tau}^2) \cdot p(\boldsymbol{\sigma}^2, \boldsymbol{\tau}^2) \quad (2)$$

Following Bayes rule, the distribution of parameters given the data – $p(\boldsymbol{\theta}, \boldsymbol{\phi}, \gamma^2, \boldsymbol{\omega}, \boldsymbol{\varepsilon}, \boldsymbol{\sigma}^2, \boldsymbol{\tau}^2 | \mathbf{y})$ – is proportional to the product of the distribution of the data given those parameters – $p(\mathbf{y} | \boldsymbol{\theta}, \boldsymbol{\phi}, \gamma^2, \boldsymbol{\omega}, \boldsymbol{\varepsilon})$ – and the probability of those parameters – $p(\boldsymbol{\omega}, \boldsymbol{\varepsilon} | \boldsymbol{\sigma}^2, \boldsymbol{\tau}^2) \cdot p(\boldsymbol{\sigma}^2, \boldsymbol{\tau}^2)$. We will estimate the posterior distribution of our unknown parameters by Markov Chain Monte Carlo (MCMC) simulation, specifically the Gibbs sampler. (Geman and Geman, 1984) The Gibbs sampler estimates the full posterior distribution in equation (2) by iteratively sampling a value for each unknown parameter conditional on the current values of the other unknown parameters. In other words, we iterate over the following steps. The number of each step corresponds to subsection number in this section, so that Subsection 5.1 describes Step 1 in detail and so on. The notation for variables estimated in a given step are summarized in Table 5, so that entry 1 in Table 5 summarizes the variables estimated in Step 1 and so on.

Step 1: Sample new values of $(\boldsymbol{\theta}, \boldsymbol{\phi}, \gamma^2)$ from $p(\boldsymbol{\theta}, \boldsymbol{\phi}, \gamma^2 | \mathbf{y}, \boldsymbol{\omega}, \boldsymbol{\varepsilon}, \boldsymbol{\sigma}^2, \boldsymbol{\tau}^2)$

Step 2: For each person i and year t , sample new values of $(\omega_{i,t}, \varepsilon_{i,t})$ from $p(\omega_{i,t}, \varepsilon_{i,t} | \mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\phi}, \gamma^2, \boldsymbol{\sigma}^2, \boldsymbol{\tau}^2)$

Step 3: For each person i and year t , sample new values of $(\sigma_{i,t}^2, \tau_{i,t}^2)$ from $p(\sigma_{i,t}^2, \tau_{i,t}^2 | \mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\phi}, \gamma^2, \boldsymbol{\omega}, \boldsymbol{\varepsilon})$ using three-level MHDP model.

These sampling steps form a Markov chain that is iterated until the set of all parameters has converged to their joint posterior distribution. In the following subsections, we outline in detail the sampling process for the three steps given above.

5.1 Step 1: Sampling income process parameters $(\boldsymbol{\theta}, \boldsymbol{\phi}, \gamma^2)$

In this step, we take realized shocks $(\boldsymbol{\omega}, \boldsymbol{\varepsilon})$ as well as excess log income data (\mathbf{y}) as given, to estimate the rate at which shocks enter into and damp out of income $(\boldsymbol{\theta}, \boldsymbol{\phi})$.

Our model for the income process is based on a decomposition of the excess log income into permanent and transitory components, as outlined in Section 3. Reorganizing equation (1) and setting limits of $\Omega = 3$ periods for permanent shocks to come into effect and $\epsilon = 3$ periods for transitory shocks to fade completely (a conservative choice according to Abowd and Card, 1989), we get the following dynamic linear model,

$$y_{i,t} = \sum_{k=0}^{t-3} \omega_{i,k} + \sum_{k=t-2}^t \theta_{t-k} \omega_{i,k} + \sum_{k=t-2}^t \phi_{t-k} \varepsilon_{i,k} \quad (3)$$

For each individual i , the dynamic linear model for their excess log income (\mathbf{y}_i) is a combination of the homogeneous parameters $(\boldsymbol{\theta}, \boldsymbol{\phi})$ and realized shocks $(\boldsymbol{\omega}, \boldsymbol{\varepsilon})$. In our Gibbs sampling model implementation, we take advantage of the fact that sampling new values of the homogeneous parameters conditional on fixed values of the realized shocks is relatively simple, and vice versa.

If we are given values of the realized shocks $(\boldsymbol{\omega}_i, \boldsymbol{\varepsilon}_i)$, we can calculate the scalar $y_{i,t}^*$ and the 1×6 (since $\Omega + \epsilon = 6$) vector $X_{i,t}$,

$$y_{i,t}^* \equiv y_{i,t} - \sum_{k=0}^{t-3} \omega_{i,k} \quad X_{i,t} \equiv (\omega_{i,t-2}, \omega_{i,t-1}, \omega_{i,t}, \varepsilon_{i,t-2}, \varepsilon_{i,t-1}, \varepsilon_{i,t})$$

Let \mathbf{y}^* be the $N(T-3) \times 1$ vector of all $y_{i,t}^*$ across individuals i and time t , and let \mathbf{X} be the $N(T-3) \times 6$ matrix whose rows are all $X_{i,t}$ across individuals i and time t . We can then write equation (3) as a simple linear regression model,

$$\mathbf{y}^* = \mathbf{X} \cdot \boldsymbol{\beta} + \mathbf{e} \quad \text{where } \mathbf{e} \sim \text{Normal}(\mathbf{0}, \gamma^2 \cdot \mathbf{I})$$

where $\boldsymbol{\beta} = (\theta_2, \theta_1, \theta_0, \phi_2, \phi_1, \phi_0)$ are the homogeneous parameters of interest. Note that this is the stage at which we use measurement error (\mathbf{e}) as distinct from transitory shocks. Measurement allows us to capture the fact that any attempt to explain income with estimated shocks will necessarily be imperfect. Economically, \mathbf{e} is merely an additional transitory shock, an i.i.d. homogenous shock that has been separated from the heterogeneous transitory shock $(\boldsymbol{\varepsilon})$ purely to estimate the model.

Under our Bayesian approach, we use non-informative prior distributions for both γ^2 and $\boldsymbol{\beta}$, which leads to the following simple posterior distributions (the Bayesian analog of a least-squares estimate):

$$\begin{aligned} \gamma^2 &\sim \text{Inv - Gamma} \left(\frac{TN}{2}, \frac{(\mathbf{y}^* - \mathbf{X}\hat{\boldsymbol{\beta}})'(\mathbf{y}^* - \mathbf{X}\hat{\boldsymbol{\beta}})}{2} \right) \\ \boldsymbol{\beta} &\sim \text{Normal} \left(\hat{\boldsymbol{\beta}}, \gamma^2 \cdot (\mathbf{X}'\mathbf{X})^{-1} \right) \end{aligned} \quad (4)$$

where $\hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y}^*$ as in a least-squares regression. We sample new values of

γ^2 and $(\boldsymbol{\theta}, \boldsymbol{\phi})$ from the distributions in (4), but with the additional constraint that $\sum_k \phi_k = 1$. For a detailed review of the Bayesian approach to regression models, see Gelman, Carlin, Stern, and Rubin (1995).

5.2 Step 2: Sampling realized shocks $(\boldsymbol{\omega}, \boldsymbol{\varepsilon})$

In this step, we take excess log income data (\mathbf{y}) and the homogeneous parameters $(\boldsymbol{\theta}, \boldsymbol{\phi})$ – that determine the rate at which shocks enter into and damp out of income – as given. We use these to sample realized shocks $(\boldsymbol{\omega}, \boldsymbol{\varepsilon})$.

If we are now given values of the homogeneous parameters $(\boldsymbol{\theta}, \boldsymbol{\phi})$, then the only unmeasured variables in our dynamic linear model (3) are the realized shocks $(\boldsymbol{\omega}_i, \boldsymbol{\varepsilon}_i)$. The Kalman filter (Kalman, 1960) is a popular approach for calculating maximum likelihood estimates of the realized shocks in our dynamic linear model. The maximum likelihood estimates from the Kalman filter can then be used to sample new values of the realized shocks $(\boldsymbol{\omega}_i, \boldsymbol{\varepsilon}_i)$, as outlined in (Carter and Kohn, 1994). Given the homogeneous parameters $(\boldsymbol{\theta}, \boldsymbol{\phi}, \gamma^2)$ and the collection of volatility parameters $(\boldsymbol{\sigma}^2, \boldsymbol{\tau}^2)$, each individual’s income process is independent, which means that we can run the Kalman filter and sampling procedure for the realized shocks $(\boldsymbol{\omega}_i, \boldsymbol{\varepsilon}_i)$ for each individual i separately.

5.3 Step 3: Sampling volatility parameters $(\boldsymbol{\sigma}^2, \boldsymbol{\tau}^2)$

In this step, we take sampled realized shocks $(\boldsymbol{\omega}, \boldsymbol{\varepsilon})$ as given and use these to sample estimates of volatility parameters $(\boldsymbol{\sigma}^2, \boldsymbol{\tau}^2)$.

Our primary parameters of interest are the permanent and transitory variance $(\sigma_{i,t}^2, \tau_{i,t}^2)$ for each person i in each year t . The information about $(\sigma_{i,t}^2, \tau_{i,t}^2)$ comes from our sampled permanent and transitory shocks $(\omega_{i,t}, \varepsilon_{i,t})$ as well as the permanent and transitory variance parameters from other years and other people, which we denote $(\boldsymbol{\sigma}_{-(i,t)}^2, \boldsymbol{\tau}_{-(i,t)}^2)$. They are linked through the posterior distribution of the volatility parameters,

$$p(\sigma_{i,t}^2, \tau_{i,t}^2 | \omega_{i,t}, \varepsilon_{i,t}, \boldsymbol{\sigma}_{-(i,t)}^2, \boldsymbol{\tau}_{-(i,t)}^2) \propto p(\omega_{i,t}, \varepsilon_{i,t} | \sigma_{i,t}^2, \tau_{i,t}^2) \cdot p(\sigma_{i,t}^2, \tau_{i,t}^2 | \boldsymbol{\sigma}_{-(i,t)}^2, \boldsymbol{\tau}_{-(i,t)}^2) \quad (5)$$

The first term of equation (5) comes from the likelihood of our realized shocks $(\omega_{i,t}, \varepsilon_{i,t})$ from our dynamic linear model,

$$p(\omega_{i,t}, \varepsilon_{i,t} | \sigma_{i,t}^2, \tau_{i,t}^2) \propto (\sigma_{i,t}^2 \tau_{i,t}^2)^{-\frac{1}{2}} \exp\left(-\frac{1}{2} \frac{\omega_{i,t}^2}{\sigma_{i,t}^2} - \frac{1}{2} \frac{\varepsilon_{i,t}^2}{\tau_{i,t}^2}\right) \quad (6)$$

The second term of equation (5) is our Markovian hierarchical Dirichlet process (MHDP) prior $p(\sigma_{i,t}^2, \tau_{i,t}^2 | \boldsymbol{\sigma}_{-(i,t)}^2, \boldsymbol{\tau}_{-(i,t)}^2)$ which consists of a weighted mixture of the other volatility parameters in the population. Sampling new values $(\sigma_{i,t}^2, \tau_{i,t}^2)$ from the posterior distribution (5) is a multi-step process that acknowledges the structure of our population. First, we sample volatility parameter proposal values $(\sigma_{\star}^2, \tau_{\star}^2)$ from a continuous distribution $f(\cdot)$. This is a new pair of possible volatility parameter

values not found in the population. We will set $(\sigma_{i,t}^2, \tau_{i,t}^2)$ equal to these proposal values only if we cannot find a suitable set of values from among our population of other $(\sigma_{-(i,t)}^2, \tau_{-(i,t)}^2)$ values across time within the same person and across other people.

5.3.1 Level 1: Is volatility unchanged from last year?

We first consider sampling volatility parameters $(\sigma_{i,t}^2, \tau_{i,t}^2)$ to be the same as the previous year $(\sigma_{i,t-1}^2, \tau_{i,t-1}^2)$,

$$p((\sigma_{i,t}^2, \tau_{i,t}^2) = (\sigma_{i,t-1}^2, \tau_{i,t-1}^2)) \propto m_{t-1} \cdot p(\omega_{i,t}, \varepsilon_{i,t} | \sigma_{i,t-1}^2, \tau_{i,t-1}^2) \quad (7)$$

$$p((\sigma_{i,t}^2, \tau_{i,t}^2) \neq (\sigma_{i,t-1}^2, \tau_{i,t-1}^2)) \propto p(\omega_{i,t}, \varepsilon_{i,t} | \sigma_{\star}^2, \tau_{\star}^2) \quad (8)$$

The choice here is between leaving volatility parameters unchanged from last year ($(\sigma_{i,t-1}^2, \tau_{i,t-1}^2)$ in equation (7)) or rejecting last year's value in favor of the new proposal value ($(\sigma_{\star}^2, \tau_{\star}^2)$ in equation (8)). This choice is made stochastically by flipping a weighted coin with weights equal to the probabilities in equations (7) and (8). Note that these probabilities are based on the likelihood of our realized shocks – equation (6) – given the set of candidate volatility parameters. However the probability that parameters are unchanged from the previous period (equation 7) gets an extra weight equal to the current “streak” of those parameter values; m_{t-1} is the number of consecutive years through $t-1$ with parameter values $(\sigma_{i,t-1}^2, \tau_{i,t-1}^2)$. This weighting allows for within-individual autocorrelation in the choice of parameter values $(\sigma_{i,t}^2, \tau_{i,t}^2)$.

If the choice in equation (7) is selected, volatility parameters are left unchanged at $(\sigma_{i,t}^2, \tau_{i,t}^2) = (\sigma_{i,t-1}^2, \tau_{i,t-1}^2)$. Otherwise the choice in equation (8) is selected, in which case we go to Level 2 to identify volatility parameters $(\sigma_{i,t}^2, \tau_{i,t}^2)$.

5.3.2 Level 2: Is volatility the same as in another year?

If the choice from equation (8) is selected in Level 1, volatility parameter values for individual i in year t will change from those held year $t-1$. In Level 2, we consider whether or not parameter values change to values individual i held in another year. Assume there are K unique values $(\sigma_{i,k}^2, \tau_{i,k}^2)$ held by individual i over the sample. We next consider sampling $(\sigma_{i,t}^2, \tau_{i,t}^2)$ equal to one of these values,

$$p((\sigma_{i,t}^2, \tau_{i,t}^2) = (\sigma_{i,k}^2, \tau_{i,k}^2)) \propto n_k \cdot p(\omega_{i,t}, \varepsilon_{i,t} | \sigma_{i,k}^2, \tau_{i,k}^2) \quad k = 1, \dots, K \quad (9)$$

$$p(\sigma_{i,t}^2, \tau_{i,t}^2 \neq \text{any } (\sigma_{i,k}^2, \tau_{i,k}^2)) \propto p(\sigma_{i,t}^2, \tau_{i,t}^2 | \sigma_{\star}^2, \tau_{\star}^2) \quad (10)$$

n_k is the number of time-points within person i that currently have the values $(\sigma_{i,k}^2, \tau_{i,k}^2)$ as their volatility parameters. These probabilities are still based on the likelihood of our realized shocks in equation (6) given the set of candidate volatility parameters. We add an additional weight n_k to allow popular or common volatility parameter values for person i to have a greater chance of being selected.

The choice here is between one of volatility values the individual held in another year ($(\sigma_{i,k}^2, \tau_{i,k}^2)$ in equation (9)) and the new proposal value ($(\sigma_{\star}^2, \tau_{\star}^2)$ in equation(10)). Again, this choice is made stochastically by flipping a weighted coin with weights equal

to the probabilities given above. If one of the choices in equation (9) is selected, we change the volatility parameter values to those the individual held at another time $(\sigma_{i,t}^2, \tau_{i,t}^2) = (\sigma_{i,k}^2, \tau_{i,k}^2)$. Otherwise the choice in equation (10) is selected, in which case we go to Level 3 to identify volatility parameters $(\sigma_{i,t}^2, \tau_{i,t}^2)$.

5.3.3 Level 3: Is volatility the same as another person's?

If the choice from equation (10) is selected in Level 2, volatility parameter values for individual i in year t will be different from any other values for that individual in other years. In Level 3, we consider whether or not parameter values change to values held by an individual other than person i . Assume there are L unique values (σ_l^2, τ_l^2) across all time points within all people in the population of individuals other than person i . We next consider sampling $(\sigma_{i,t}^2, \tau_{i,t}^2)$ equal to one of these values,

$$p((\sigma_{i,t}^2, \tau_{i,t}^2) = (\sigma_l^2, \tau_l^2)) \propto n_l \cdot p(\omega_{i,t}, \varepsilon_{i,t} | \sigma_l^2, \tau_l^2) \quad l = 1, \dots, L \quad (11)$$

$$p((\sigma_{i,t}^2, \tau_{i,t}^2) \neq \text{any } (\sigma_l^2, \tau_l^2)) \propto p(\omega_{i,t}, \varepsilon_{i,t} | \sigma_\star^2, \tau_\star^2) \quad (12)$$

n_l is the number of time-points in the population of all people outside of person i that currently have the values (σ_l^2, τ_l^2) for their volatility parameters. These probabilities are still based on the likelihood of our realized shocks (6) given the set of candidate volatility parameters. We add the weights n_l to allow popular or common volatility values to have a greater chance of being selected.

The choice here is between one of the volatility values held by another individual (σ_l^2, τ_l^2) in equation (11) and the new proposal value $(\sigma_\star^2, \tau_\star^2)$ in equation (12). Again, this choice is made stochastically by flipping a weighted coin with weights equal to the probabilities given above. If one of the choices in equation (11) is selected, we change the volatility parameter values to those held by another individual $(\sigma_{i,t}^2, \tau_{i,t}^2) = (\sigma_l^2, \tau_l^2)$. Otherwise the choice in equation (12) is selected, in which case we use the new proposal value $(\sigma_{i,t}^2, \tau_{i,t}^2) = (\sigma_\star^2, \tau_\star^2)$. These are values that have not yet been seen in the population.

The three steps outlined above produce new sampled values for the permanent and transitory variance $(\sigma_{i,t}^2, \tau_{i,t}^2)$ for year t of person i . These steps must be repeated for all other years and all other people in the population. The entire sampling procedure outlined in this section is designed to allow popular values of $(\sigma_{i,t}^2, \tau_{i,t}^2)$ to be shared over time within an individual as well as across the population of individuals, while still allowing for new values to be introduced into the population if needed to effectively model the realized shocks sampled in Section 5.2.

This algorithm is programmed in Python and run on a grid cluster of computers. One run of this model (with 2,000 iterations) takes approximately one week, though multiple runs can be done simultaneously. These multiple runs were used to evaluate convergence of the algorithm to a reasonable set of samples from the posterior distribution of all parameters.

Important Note: Because of the substantial time needed to estimate the model, the volatility parameter values shown here are not the most up-to-date. In particular,

Table 6: Basic Model Results

Distribution of Variance Parameters			Shocks' Rate of Entry/Exit		
	Permanent Variance	Transitory Variance	lag, k	θ_k	ϕ_k
Mean	0.0542	0.2811	$k = 0$	0.517	0.755
St. Dev.	0.3821	1.0449	(s.e.)	(0.002)	(0.024)
N	62,881	62,881	$k = 1$	0.947	0.190
1 st %	0.0198	0.0508	(s.e.)	(0.002)	(0.016)
5 th %	0.0208	0.0535	$k = 2$	1.025	0.054
10 th %	0.0213	0.0546	(s.e.)	(0.001)	(0.014)
25 th %	0.0220	0.0552			
50 th %	0.0229	0.0564			
75 th %	0.0238	0.0679			
90 th %	0.0250	0.3736			
95 th %	0.0311	1.2338			
99 th %	0.4801	5.4330			

Distribution of posterior means (σ^2, τ^2)

θ_k : impact of permanent shock from k periods ago
 ϕ_k : impact of transitory shock from k periods ago

The left panel presents the estimates of the permanent and transitory variance distribution (σ^2, τ^2). These are the distribution of posterior means estimated from the data. These posteriors of the permanent variance and transitory variance are calculated for each individual in each year, as described in Section 5. The distributions presented here consider all years and all individuals together. The right panel of this table presents θ and ϕ . θ refers to the rate at which permanent shocks enter into income data, so that θ_k shows the impact of a permanent shock from k periods ago in current excess log income. These are assumed to be 1 for $k > \Omega$. ϕ refers to the rate at which transitory shocks damp out of income data, so that ϕ_k shows the impact of a transitory shock from k periods ago in current excess log income. These are assumed to be 0 for $k > \epsilon$. The impact of these coefficients is detailed in equation 1.


they use a more expansive sample (including the low-income oversample of the PSID, explaining the larger sample size) and a different set of age restrictions and controls. Initial checks suggest results are not affected substantially by these differences.

6 Results

Here, we present the model parameters estimated using the methods from Section 5. The chief object of interest is the evolution of the cross-sectional distribution of permanent and transitory volatility parameters (σ_t^2, τ_t^2) over time. These are shown in Section . Our main finding is that the increase in mean volatility can be attributed solely to an increase in volatility at the volatile tail; those who in other years would be expected to have large income changes are now likely to have even larger ones.

We begin with more basic results. In subsection 6.1, we present estimates of the

Figure 2: Distribution of Permanent and Transitory Variance



This figure presents the distribution of τ^2 and σ^2 . These are the distribution of posterior means estimated from the data, as presented numerically in Table 6. These posteriors of the permanent variance and transitory variance are calculated for each individual in each year, as described in Section 5. The distributions presented here show all years and individuals together. Values are truncated at the 95th percentile for the permanent variance and at the 90th percentile for the transitory variance. Mean and median of the truncated part of each distribution is given.

homogeneous parameters (θ, ϕ) that determine the rate at which shocks enter into and exit from income and also the overall distribution of volatility (σ^2, τ^2) .


6.1 Basic results

Table 6 presents the basic parameter estimates obtained from fitting our model to the PSID income data described in Section 5. The left panel shows the distribution of risk in the population, τ^2 and σ^2 . Formally, we present the distribution of posterior means of permanent and transitory variance parameters. The right panel show estimates of the rate at which permanent shocks enter into income, θ , and the rate at which transitory ones exit, ϕ . While we estimate a distribution of volatility parameters, allowing them to vary over time and across individuals, we impose the constraint that θ and ϕ are constant.

The main thing to note about the distribution of volatility parameters (σ^2, τ^2) (shown in the left panel of Table 6) is the extreme skew and fat tails (kurtosis). While medians are modest, means far exceed medians. At the median, transitory shocks have a standard deviation of approximately 12% annually; permanent shocks have a standard deviation of just over 20% annually. However, the highest volatility observations imply shocks with standard deviations well above 100% annually. Figure 2 plots these skewed and fat-tailed distributions by truncating the right tail.

The main thing to note about estimates of (θ, ϕ) (shown in the right panel of Table 6) is that permanent shocks enter in quickly (θ_k are close to one) while transitory

Figure 3: Impulse Response Function for Permanent and Transitory Shocks



This figure presents an estimated impulse response function for a permanent (left panel) and transitory (right panel) shock.

shocks damp out quickly (ϕ_k fall to zero). The impact of a shock on the evolution of income is presented in Figure 3. These present impulse response functions for a permanent (left panel) and transitory (right panel) shock. Shocks were calibrated as a one standard-deviation shock for an individual with volatility parameters at the estimated means (pulled from Table 6).

6.2 Evolution of the volatility distribution

Here, we show how the distribution of posterior means of variance parameters has evolved over time. This evolution is shown in Tables 7 and 8 and also in Figure 4. Table 7 shows the year-by-year distribution of volatility parameters (σ_t^2, τ_t^2) posterior means. This table mirrors Table 4, with volatility parameter ($\sigma_{i,t}^2, \tau_{i,t}^2$) posterior means replacing reduced form moments. The first three columns show results for the permanent variance parameter, σ^2 ; the final three columns show results for the transitory variance parameter, τ^2 . The first and fourth columns present means of the permanent and transitory variance parameter posterior means, the second and fifth columns present medians of parameter posterior means, and the third and sixth columns present 95th percentiles. All are weighted with weights from the PSID. The first row shows whole-sample results. The second row shows the percent change in the mean, median, or 95th percentile over the sample.⁷ The coefficient and t-statistic from this regression are shown just below. Year-by-year values are then shown.

Table 7 shows that the mean of permanent and transitory parameters have increased substantially over the sample (by 71 and 100 percent, respectively) while the

⁷This is calculated as coefficient of a weighted OLS regression of the year-specific moments from below on a time trend, multiplied by the number of years (2005-1968) and divided by the whole-sample value in the previous row.

Table 7: Year-by-Year Income Volatility Parameters

Sample Moment	Permanent Variance, σ^2			Transitory Variance, τ^2		
	Mean	Median	95 th %	Mean	Median	95 th %
Average	0.0542	0.0229	0.0311	0.2811	0.0579	1.2338
% Change	76%	0%	44%	96%	1%	143%
Slope (t-stat)	0.0011 (6.92)	0.0000 (2.13)	0.0004 (8.59)	0.0073 (6.82)	0.0000 (5.02)	0.0475 (6.06)
1970	0.0374	0.0228	0.0253	0.1478	0.0575	0.3824
1971	0.0335	0.0228	0.0254	0.1949	0.0576	0.6204
1972	0.0352	0.0229	0.0265	0.2102	0.0576	0.9897
1973	0.0412	0.0228	0.0258	0.1831	0.0577	0.6658
1974	0.0374	0.0228	0.0274	0.1922	0.0579	0.6183
1975	0.0415	0.0228	0.0265	0.1911	0.0579	0.7106
1976	0.0432	0.0229	0.0287	0.2616	0.0581	1.1555
1977	0.0380	0.0229	0.0296	0.2488	0.0578	1.1268
1978	0.0434	0.0229	0.0297	0.1995	0.0576	0.8308
1979	0.0611	0.0229	0.0298	0.2440	0.0578	1.1638
1980	0.0407	0.0229	0.0305	0.2103	0.0578	0.8399
1981	0.0465	0.0229	0.0288	0.1938	0.0577	0.7386
1982	0.0383	0.0228	0.0288	0.2098	0.0579	0.9679
1983	0.0581	0.0229	0.0316	0.2738	0.0578	1.3297
1984	0.0570	0.0229	0.0298	0.2153	0.0578	0.9589
1985	0.0474	0.0228	0.0290	0.2321	0.0579	1.0527
1986	0.0439	0.0229	0.0287	0.2521	0.0578	1.1314
1987	0.0600	0.0229	0.0299	0.2639	0.0579	1.2514
1988	0.0580	0.0229	0.0308	0.2279	0.0579	0.9420
1989	0.0444	0.0229	0.0304	0.2613	0.0578	1.3572
1990	0.0488	0.0228	0.0305	0.2454	0.0577	1.0738
1991	0.0826	0.0229	0.0369	0.2671	0.0578	1.3342
1992	0.0610	0.0228	0.0325	0.2580	0.0580	1.1469
1993	0.0626	0.0229	0.0381	0.4030	0.0581	2.3549
1994	0.0711	0.0229	0.0386	0.4464	0.0581	2.4692
1995	0.0663	0.0229	0.0392	0.4305	0.0580	2.4381
1996	0.0438	0.0229	0.0346	0.4244	0.0579	2.2832
1997	0.0582	0.0229	0.0346	0.3897	0.0580	2.0100
1999	0.0662	0.0229	0.0336	0.3062	0.0580	1.3047
2001	0.0699	0.0229	0.0322	0.2689	0.0579	1.0350
2003	0.0731	0.0229	0.0389	0.4832	0.0580	2.5074
2005	0.0802	0.0228	0.0399	0.3918	0.0585	2.2470

The construction of posterior means for σ^2 and τ^2 for each individual in each year is detailed in the text. The first row shows the full sample distribution, so that the second column shows the median value of the posterior mean of σ^2 over all individual-years. The second row shows the percent change over the sample, calculated as the coefficient of a weighted OLS regression of year-specific sample moments on a time trend, multiplied by the number of years (2005-1968) and divided by the full sample value. The coefficient and t-statistic are shown below.

medians have not (0 and 1 percent increases, respectively). This divergence can be explained by an increase in the magnitude of permanent and transitory variance parameters at the right tail, among individuals with the highest parameters (the 95th percentile values increasing 43 percent and 101 percent, respectively). Colloquially, the kind of people whose incomes had always moved around a lot are moving around even more than they used to; the median person’s income does not move more than it used to.

The same pattern is also apparent in Table 8. This shows percent changes over time in many quantiles (not just the mean, median, and 95th percentile as in Table 7) but omits year-by-year values. A large increase in permanent and transitory variance parameters is apparent above the 90th percentile, but not below it.

This pattern can be seen graphically in Figure 4, which shows the year-by-year evolution of many quantiles of the distribution of permanent and transitory variance posterior means. In the bottom panels of Figure 4, we plot the 1st, 5th, 10th, 25th, 50th, and 75th percentile values of the posterior mean of the permanent (σ^2 , left) and transitory (τ^2 , right) variance parameters by year. These are very stable and show no clear upward trend. The size of this increase is extremely small economically. Looking at all but the “risky” tail of the distributions, the distributions look very stable.

In the middle and upper panels of Figure 4, we show the evolution of the “risky” tail of the distribution of posterior means. In this case, variance parameters increase strongly and significantly. This increase in the right tail of the distribution explains the increase in the mean completely.

6.3 Heterogeneity or fat tails?

So far, we have shown that the increases in income volatility can be attributed solely to increases in the right tail of the volatility distribution. In other words, individuals whose incomes have always moved around a lot move around even more than they used to. We find almost no change in income volatility at the median.

To obtain this result, the model makes several assumptions. Chief among these is that the distribution of shocks is normal conditional on the volatility parameters. When the unconditional distribution of shocks is fat-tailed (has high kurtosis), this is automatically attributed to heterogeneity in volatility parameters. An alternative hypothesis is that there is little or no heterogeneity in volatility parameters, but that shocks are conditionally fat-tailed.

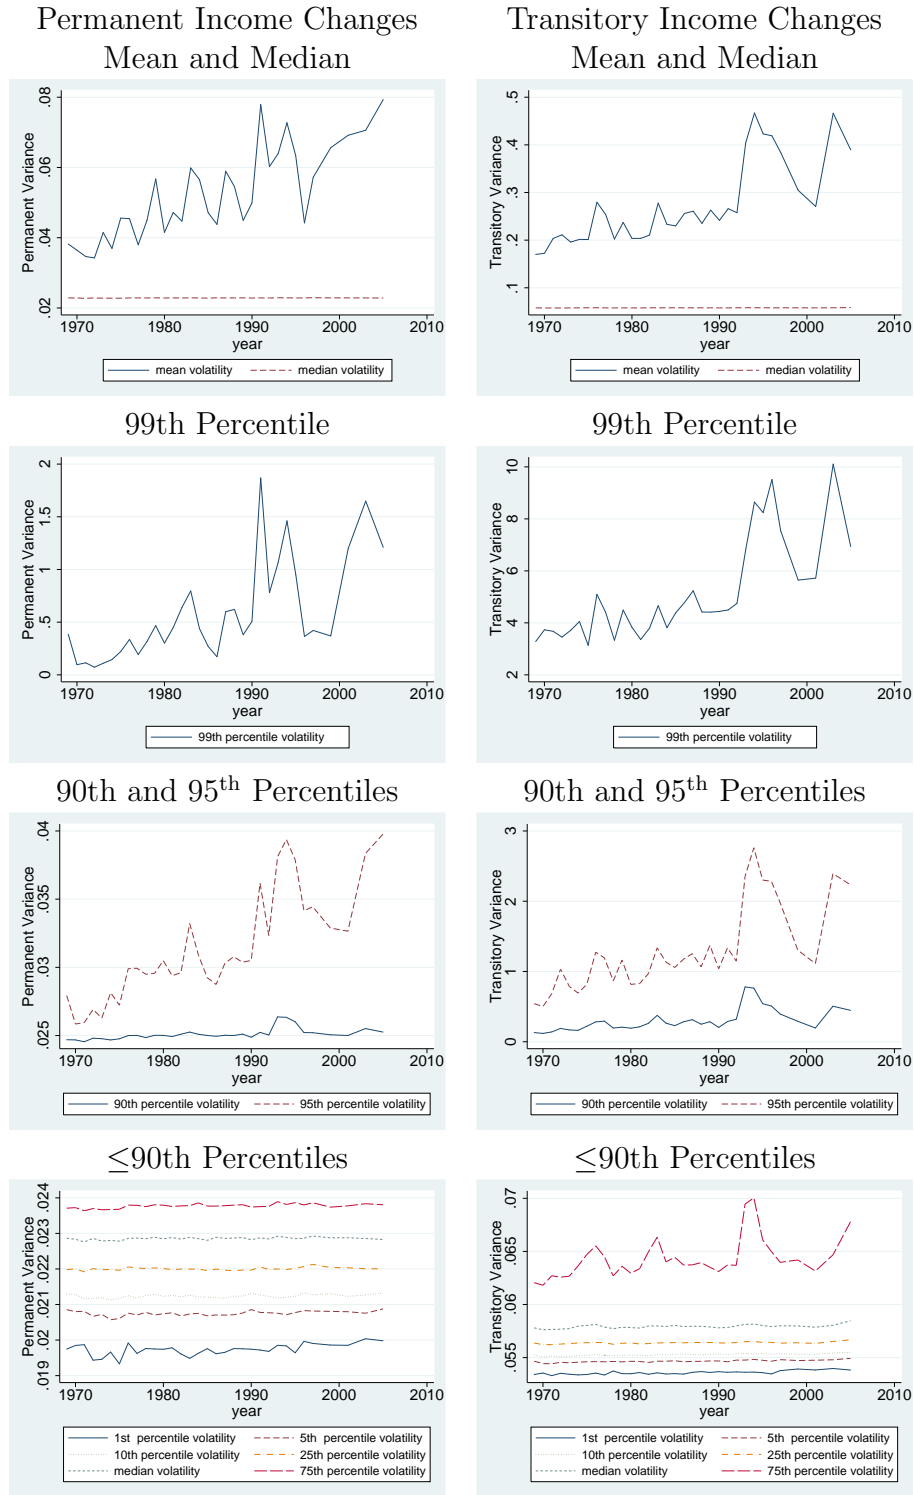
When looking at the cross-section of income changes, heterogeneity in volatility parameters with conditionally normal shocks and no heterogeneity in volatility parameters and conditionally fat-tailed shocks are observationally equivalent; they both imply a fat-tailed unconditional distribution of income changes. By examining serial dependence, it is possible to reject the hypothesis that everyone has the same volatility parameter. If shocks are conditionally fat-tailed but everyone has the same volatility parameters, then those with large past income changes should be no more likely than others to experience large subsequent income changes. If individuals differ in their volatility parameters and those volatilities are persistent, then individuals with large

Table 8: Percent Change in Volatility Over Sample

Percent Change in	Permanent Volatility	Transitory Volatility
mean (t-stat)	72.9% (7.84)	64.7% (6.84)
1st percentile (t-stat)	1.8% (4.67)	0.4% (3.63)
5 th percentile (t-stat)	0.6% (3.71)	0.2% (2.50)
10th percentile (t-stat)	0.3% (2.89)	0.2% (1.98)
25 th percentile (t-stat)	0.2% (2.72)	0.1% (0.73)
median (t-stat)	0.2% (2.25)	-0.3% (0.70)
75 th percentile (t-stat)	-0.5% (3.42)	-0.4% (0.13)
90th percentile (t-stat)	3.5% (2.73)	125.6% (4.18)
95 th percentile (t-stat)	25.8% (5.49)	98.3% (6.66)
99th percentile (t-stat)	279.0% (4.99)	69.3% (7.21)


Values shown are the coefficients from a quantile regression used to predict posterior mean volatility parameter estimates with a linear time trend. These coefficients are then multiplied by the number of years in the sample (2005-1968) and divided by the average volatility parameter for this quantile over time. Values are then multiplied by 100 to give the percent change in volatility for a given quantile over the sample. The first row shows the same, where an OLS regression is used instead of a quantile regression. t-statistics are in parentheses. These are bootstrapped by randomly permuting individuals' posterior means and running quantile regressions on these permuted data. Regressions are unweighted.

Figure 4: Evolution of Percentiles of Volatility Distribution



These figures show the evolution of various percentiles of the posterior mean of the permanent (left) and transitory (right) variance for various percentiles of the distribution of variance parameters.

Figure 5: Comparing Sample Variances for Those With and Without Large Past Income Changes



Following Meghir and Pistaferri (2004), the sample permanent variance is calculated as the product of two-year changes in excess log incomes (between years t and $t-2$) and the six-year changes that span them (between years $t+2$ and $t-4$). The sample transitory variance is calculated as the square of two-year changes in excess log income. Individuals are defined as low past variances when their sample variance (permanent or transitory, respectively) four years ago is below median; individuals are defined as high past variance when their sample variance four years ago is above the 95th percentile. Weighted averages for these groups are presented in each year for which data is available for permanent variance (left panel) and transitory variance (right panel).

past income changes will be more likely than others to have large subsequent income changes.

This possibility is investigated in Table 9 and shown graphically in Figure 5. These compare the sample variance of income changes for individuals with and without large past income changes. The left two columns in Table 9 and the left panel in Figure 5 examine the permanent variance, calculated as the product of short-term (two-year) changes in excess log income times the long-term (six-year) changes that span them. (Meghir and Pistaferri, 2004) The right two columns in Table 9 and the right panel in Figure 5 examine the variance of raw income changes, calculated as the square of two-year income changes in excess log income.

In each year, a cohort without large income changes is formed as the set of individuals whose measure of variance, either permanent variance or squared income change, was below median four years ago; a cohort with large income changes is formed as the set of individuals whose measure of variance was above the 95th percentile four years ago. This four-year period is chosen so that income shocks are far enough apart to be uncorrelated. (Abowd and Card, 1989) Results for the low past-volatility cohorts are shown in the first and third columns of Table 9 or in the solid lines in Figure 5. Results for the high past volatility cohorts are shown in the second and fourth columns of Table 9 or in the dashed lines in Figure 5.

Table 9: Year-by-Year Volatility Sample Moments by Past Income Volatility

Sample Moment	Permanent Variance		Squared Change	
	Mean		Mean	
Past Variance	Low	High	Low	High
Average	0.0883	0.3923	0.2859	0.6902
% Difference	99%		44%	
Slope (t-stat)	0.000096 0.13	0.019 4.14	0.0084 8.10	0.022 4.68
1974	0.1645	0.4291	.	.
1975	0.2152	0.3374	.	.
1976	0.1180	0.3575	0.2403	0.4384
1977	0.1034	0.4189	0.2333	0.8105
1978	0.0428	0.1308	0.2174	0.4622
1979	0.0595	0.2730	0.2482	0.6020
1980	0.0763	0.2993	0.1841	0.4536
1981	0.0572	0.2294	0.1742	0.9885
1982	0.0810	0.1523	0.1937	0.3602
1983	0.0876	0.1372	0.2431	0.5713
1984	0.1523	0.2951	0.2667	0.6789
1985	0.0803	0.2906	0.2447	0.5956
1986	0.1282	0.1534	0.2610	0.5591
1987	0.0963	0.3567	0.2575	0.4961
1988	0.0826	0.3033	0.2507	0.5903
1989	0.0668	0.4335	0.2546	0.6841
1990	0.1125	0.4831	0.2491	0.5228
1991	0.1025	0.5864	0.2533	0.9688
1992	0.0751	0.3595	0.2642	0.4703
1993	0.0952	0.6386	0.3689	0.7722
1994	0.0589	0.2812	0.4079	0.8099
1995	0.1255	0.4024	0.3500	0.7534
1996	0.3885	1.1607	.	.
1997	0.0398	0.9939	0.3495	0.8679
1999	0.0726	0.6292	0.3469	1.0065
2001	0.1015	0.3618	0.3653	0.8549
2003	0.1039	0.8478	0.5005	1.3810

The year t permanent variance is the product of two-year changes in excess log income (from $t - 2$ to t) and the six-year changes that span them (from $t - 4$ to $t + 2$). The first and third columns show sample means for the cohort of individuals whose permanent variance and squared change, respectively, was below median in the year four years prior. The second and fourth columns show the same, but for the cohorts with past values above the 95th percentile four years prior. The first row shows full sample moments. The third and fourth rows present the coefficient and t-statistic from a weighted OLS regression of year-specific sample means on a time trend. The difference in these two coefficients, divided by their average, is the % difference in the second row. Year-by-year means are shown below.

Note that individuals with large past income changes tend to have larger subsequent income changes. The tendency to have large income changes is persistent, which indicates that some individuals have *ex-ante* more volatile incomes than others.

This paper attributes the rise in unconditional tail-fatness (the increasing magnitude of extreme income changes) to an increase in volatility among high-volatility individuals. An alternative hypothesis is that extreme income changes are getting more extreme for everyone. These are observationally equivalent when looking only at repeated cross-sections of income changes. Since the model assumes normality, it cannot rule out this alternative. Again, serial dependence can be used to reject the hypothesis that the results are driven by tail-fatness increasing for everyone. If (as we argue) volatility is increasing for high-volatility individuals but not for low-volatility individuals, then the gap in the sample variance between those with and without large past income changes should be increasing over time. Alternatively, if volatility is increasing for everyone because tail-fatness is increasing, then that gap should not be increasing.

The divergence over time in volatility between past low- and high-volatility cohorts is clear in both the table and figure. The magnitude of income changes has been increasing more for those with large past income changes (who are more likely to be inherently high-volatility) than for those without such large past income changes (who are not). This rules out the possibility that the increasing average magnitude of income changes affects everyone equally. This increase in volatility falls primarily on those who could be expected to have volatile incomes to begin with.

It is worth noting that the permanent variance has not increased at all for those whose past permanent variance suggests their incomes are unlikely to move very much. However, squared income changes have increased for those with and without large past changes; our finding is just that this increase is larger for those with income changes in the past. This is to be expected. Past squared income changes cannot perfectly identify future expected squared income changes. Not only will expected risk change for some individuals, but also realized squared income changes to not perfectly identify expected squared income changes. Some of those in the cohort with high realized volatility actually have low expected volatility, and *vice versa*.

6.4 Whose incomes are increasingly volatile?

This paper has shown that increases in income volatility over time in the U.S. can be attributed to an increase in volatility for those with the most volatile incomes. While the variance of income changes has changed little for the median individual (the individual with the median variance of income changes), the incomes of individuals with volatile incomes have become even more volatile. So far, we have only identified a latent distribution of volatility parameters; we have said nothing about the other attributes of the high-volatility individuals whose volatility is increasing. Who are these individuals?

Table 10 predicts the posterior mean variance (volatility) estimates described earlier with a linear time trend. The “change” row shows the coefficient on calendar time; the “percent change” row shows the expected percent change over the sample

Table 10: Volatility Time Trend by Group

Permanent Variance					
sample	self-employed	not self-employed	more than high school	high school	less than high school
change per year	0.0048	0.0005	0.0014	0.0004	0.0010
% change '68-'05	225%	38%	92%	28%	64%
(t-stat)	(6.96)**	(3.70)**	(5.18)**	(1.77)	(4.04)**
N	7,797	63,149	28,759	22,498	19,689

Transitory Variance					
sample	self-employed	not self-employed	more than high school	high school	less than high school
change per year	0.0256	0.0054	0.0084	0.0065	0.0051
% change '68-'05	165%	87%	102%	100%	74%
(t-stat)	(11.98)**	(15.43)**	(11.91)**	(10.48)**	(8.23)**
N	7,797	63,149	28,759	22,498	19,689

Results from a weighted OLS regression to predict the posterior mean variance (volatility) estimate with a linear time trend. The “change” row shows the coefficient on calendar time; the “percent change” row shows the expected percent change over the sample implied by this coefficient. This is (100 percent) times (2005 minus 1968) times (the coefficient on calendar time) divided by (the average posterior mean in the sample). The top panel presents results for the permanent variance; the bottom panel presents results for the transitory variance. Each column presents results for a different sub-sample.

implied by this coefficient. The top panel presents results for the permanent variance; the bottom panel presents results for the transitory variance. Each column presents results for a different sub-sample. By comparing the first two columns, note that risk has increased dramatically more for self-employed people than for others. These individuals have much higher average levels of volatility, but their percentage change in volatility is still higher than for other individuals. Self-employed individuals account for a substantial proportion of the overall increase in income volatility.

The final three columns of Table 10 show results for different education groups. Income volatility increases in similar amounts for various levels of educational attainment.

7 Conclusion

We have shown that increases in the size of income changes in the PSID can be attributed solely to the “right tail” of the volatility distribution. In other words, those who would have been high risk in past years are now even higher risk. Everyone else has had no substantial change in risk. This has important welfare implications. Those with the most income risk are presumably also the most risk tolerant; more risk averse individuals would select into occupations or life-paths with less risk as much as possible. Since the increase in risk has been limited to those who are risk-tolerant or even risk-seeking (e.g., the self-employed as shown in Section 6.4), there will be much lower welfare costs of increased risk. There may even be welfare gains associated with an increase in income volatility.

References

- ABOWD, J., AND D. CARD (1989): “On the Covariance Structure of Earnings and Hours Changes,” *Econometrica*, 57(2), 411–445.
- ALVAREZ, J., M. BROWNING, AND M. EJRNAES (2001): “Modelling income processes with lots of heterogeneity,” University of Copenhagen Working Paper.
- BANIA, N., AND L. LEETE (2007): “Income Volatility and Food Insufficiency in U.S. Low-Income Households, 1992-2003,” Institute for Research on Poverty Discussion Paper No. 1325-07.
- BATCHELDER, L. L. (2003): “Can Families Smooth Variable Earnings?,” *Harvard Journal on Legislation*, 40(2), 395–452.
- BOLLINGER, C., AND J. P. ZILIAK (2007): “Welfare Reform and the Level, Composition, and Volatility of Income,” Working Paper.
- CAMERON, S., AND J. TRACY (1998): “Earnings Variability in the United States: An Examination Using Matched-CPS Data,” Working Paper.
- CARROLL, C. D., AND A. A. SAMWICK (1997): “The Nature of Precautionary Wealth,” *Journal of Monetary Economics*, 40(1), 41–71.
- CARTER, C., AND R. KOHN (1994): “On Gibbs Sampling for State Space Models,” *Biometrika*, 81, 541–553.
- COMIN, DIEGO, E. L. G., AND B. RABIN (2006): “Turbulent Firms, Turbulent Wages?,” NBER Working Paper No. 12032.
- DAHL, M., T. DELEIRE, AND J. SCHWABISH (2007): “Trends in Earnings Variability Over the Past 20 Years,” Working Paper: Congressional Budget Office.
- DALY, M. C., AND G. J. DUNCAN (1997): “Earnings Mobility and Instability, 1969-1995,” Working Paper.
- DYNAN, K. E., D. W. ELMENDORF, AND D. E. SICHEL (2007): “The Evolution of Earnings and Income Volatility,” Working Paper, Federal Reserve Board.
- DYNARSKI, S., AND J. GRUBER (1997): “Can Families Smooth Variable Earnings?,” *Brookings Papers on Economic Activity*, 0(1), 229–303.
- GELMAN, A., J. CARLIN, H. STERN, AND D. RUBIN (1995): *Bayesian Data Analysis, Chapter 8*. Chapman and Hall/CRC, Boca Raton, FL.
- GEMAN, S., AND D. GEMAN (1984): “Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images,” *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 6, 721–741.

- GOSSELIN, P. G. (2004): “The Poor Have More Things Today C Including Wild IncomeSwings,” *Los Angeles Times*.
- GOTTSCHALK, P., AND R. MOFFITT (1994): “The Growth of Earnings Instability in the U.S. Labor Market,” *Brookings Papers on Economic Activity*, 0(2), 217–254.
- (2002): “Trends in the Transitory Variance of Earnings in the United States,” *Economic Journal*, 112, C68–73.
- (2006): “Trends in Earnings Volatility in the U.S.: 1970-2002,” Working Paper.
- HACKER, J. (2006): *The Great Risk Shift: The Assault on American Jobs, Families, Health Care, and Retirement; And How You Can Fight Back*. Oxford University Press, Oxford.
- HAIDER, S. J. (2001): “Earnings Instability and Earnings Inequality in the United States, 1967-1991,” *Journal of Labor Economics*, 19(4), 799–836.
- HERTZ, T. (2006): “Understanding Mobility in America,” Center for American Progress Discussion Paper.
- HYSLOP, D. R. (2001): “Rising U.S. Earnings Inequality and Family Labor Supply: The Covariance Structure of Intrafamily Earnings,” *American Economic Review*, 91(4), 755–777.
- KALMAN, R. E. (1960): “A New Approach to Linear Filtering and Prediction Problems,” *Transaction of the ASME—Journal of Basic Engineering*, 82, 35–45.
- MEGHIR, C., AND L. PISTAFERRI (2004): “Income Variance Dynamics and Heterogeneity,” *Econometrica*, 72(1), 1–32.
- MOFFITT, R. A., AND P. GOTTSCHALK (1995): “Trends in the Covariance Structure of Earnings in the U.S.,” Working Paper.
- MULLER, P., AND F. A. QUINTANA (2004): “Nonparametric Bayesian Data Analysis,” *Statistical Science*, 19, 95–110.
- SCHULHOFER-WOHL, S. (2006): “Heterogeneity, Risk Sharing and the Welfare Costs of Risk,” Working Paper.
- SHIN, D., AND G. SOLON (2008): “Trends in Mens Earnings Volatility: What Does the Panel Study of Income Dynamics Show?,” Working Paper.
- TEH, Y., M. JORDAN, M. BEAL, AND D. BLEI (2007): “Hierarchical Dirichlet Processes,” *Journal of the American Statistical Association*, 101(476), 1566–1581.
- WINSHIP, S. (2007): “Income Volatility, The Great Risk Shift, and the Democratic Agenda,” Working Paper.