Cyclical Labor Income Risk^{*}

MAKOTO NAKAJIMA

Federal Reserve Bank of Philadelphia

VLADIMIR SMIRNYAGIN

University of Minnesota

February 15, 2019

Abstract

In this paper, we systematically analyze several labor income definitions by drawing on PSID data, and estimate how the volatility and skewness of income shocks moves with the aggregate activity. We find that volatility is countercyclical, with individual and joint (head+wife) labor incomes being the most volatile, and with hourly wages and post-government joint labor income being the least ones. This suggests that *(i)* intra-family insurance is limited, *(ii)* taxes and transfers remove a large portion of fluctuations in risk, and that *(iii)* hours, not wages, make individual earnings fluctuate over the cycle. By re-estimating the volatility of earnings shocks on the “young” (ages 23-39) and “old” (ages 40-60) subsamples separately, we find that nearly all the countercyclicality of shocks comes from the young workers, the old subsample exhibits quantitatively muted fluctuations in risk. We then allow for time-varying skewness, and find that the probability of large negative events increases in recessions by more than the probability of large positive events. Taxes and transfers reduce the probability of tail events by a factor of 2 to 3 as compared to other income definitions considered.

Keywords: Labor income risk, business cycles

JEL: D31, E24, E32, H31, J31

^{*}We thank the Opportunity and Inclusive Growth Institute of the Federal Reserve Bank of Minneapolis, where this project was initiated. We thank seminar participants at MMM (Nashville) and LACEA (Guayaquil). The views expressed here do not necessarily reflect those of the Federal Reserve Bank of Philadelphia, the Federal Reserve Bank of Minneapolis, or the Federal Reserve System.

1 Introduction

How does labor income risk change between economic expansions and contractions? What drives the cyclical nature of individual earnings’ risk — wages or hours? How strong are intra-family insurance channels, and is there any evidence of the stabilizing role of the government? We address these questions by estimating the second and third moments of earnings shocks for several income types by drawing on the Panel Study of Income Dynamics (PSID) data.

At the conceptual level, the main objective of this paper is to systematically analyze several income definitions within an unified estimation framework. We adapt the econometric technique proposed by [Storesletten et al. \(2004\)](#), which is designed to estimate the parameters of income shocks with heteroscedasticity, and subsequently augment it to handle a time-varying third moment (skewness). The list of income definitions we consider includes individual (head’s) wages and labor income, joint (head and wife’s) total labor income, and post-government joint labor income. We also consider head’s labor income for a subsample of people with a strong labor market attachment (*inter alia* [Abowd and Card \(1989\)](#), [Meghir and Pistaferri \(2004\)](#), and [Guvenen et al. \(2014\)](#)), identified by labor income exceeding a certain income threshold. We find that the threshold matters quantitatively, and allows to reconcile several previous findings of the income risk literature. For the sake of convenience, hereafter we refer to the head’s labor income defined this way as a *narrowly defined* head’s labor income.

We obtain several results. First, by studying the cyclical nature of income risk (second moment), we find that business-cycle fluctuations in individual income risk are more pronounced than those of hourly wages, which implies that hours worked, likely a risk of unemployment — rather than per hour wage — is the driving force of countercyclical income risk. The intra-family insurance channel is not large: head’s and joint labor incomes exhibit similar swings in volatility over the business cycle. However, taxes and transfers are found to be quantitatively important: the post-government joint labor income has a more stable volatility across economic expansions and contractions.

Narrowly defined head’s labor income risk is fairly stable, with fluctuations in volatility resembling those of the hourly wage. This helps us reconcile findings of the previous literature: while some studies find income risk to be countercyclical (e.g., [Storesletten et al. \(2004\)](#), [Busch and Ludwig \(2016\)](#)), there are papers which estimate it to be acyclical (e.g., [Guvenen et al. \(2014\)](#), [Busch et al. \(2018\)](#)). Our findings suggest that it matters quantitatively how labor income is measured: while individual and joint labor incomes exhibit a pronounced countercyclical risk, hourly wages, narrowly defined labor income and post-government joint labor income show smaller fluctuations in risk over the business cycle.

By re-estimating the volatility of earnings shocks on the “young” (ages 23-39) and “old” (ages 40-60) subsamples separately, we find that nearly all the countercyclicality of shocks comes from the young workers, the old subsample exhibits quantitatively muted fluctuations in risk. In particular, while the overall dispersion of labor income shocks is higher for older workers, the cyclical movements in volatility are muted for that age group as compared to younger workers.

In the second part of the paper, we augment the estimation framework to accommodate both time-varying volatility and skewness. We make an identifying assumption that shocks are drawn from a *skew* normal distribution, which allows us to express both the second

and third moments as functions of parameters that govern volatility and skewness of labor income shocks.

We find that, during economic contractions, the shocks to all 5 income definitions we consider are drawn from PDFs that are more spread-out (higher volatility) and more skewed to the left (reflecting a higher probability of large negative events). Next, we estimate the implied probabilities of “tail-events” — large drops and increases in income. Our findings suggest that the probability of large negative events increases in recessions by more than the odds of large positive events: for example, the probability of a 100% drop at least doubles in recessions (quadruples for the head’s labor income), while the probability of a 100% income increase goes up by only 20% on average. We also find that taxes and transfers reduce the odds of tail events by a factor of 2 to 3 as compared to other income types.

The rest of paper is organized as follows. Section 2 reviews the literature. Sections 3 and 4 describe the data and lay out the estimation methodology. In Section 5, we study the heteroscedasticity case. We extend the methodology to allow for time-varying skewness in Section 6. We provide economic interpretation in Section 7, and Section 8 concludes.

2 Literature Review

The idea that the variance of an income shock can be time-dependent is not new, and originated in the literature located at the nexus of economics and finance (*inter alia* Constantinidas and Duffie (1996), Storesletten et al. (2007)). When placed into a business cycle context, a natural question is how shock structure changes with respect to aggregate conditions. Storesletten et al. (2004) proposed a novel estimation methodology of the countercyclical income risk: they find that standard deviation of labor income shocks is 80% higher in recessions than in expansions.

Storesletten et al. (2004) findings serve as benchmark in the literature. However, in a more recent contribution, Guvenen et al. (2014) reach a seemingly opposite conclusion. The authors draw from a large confidential panel dataset, and find that the second moment is remarkably stable across economic expansions and contractions. They also find that in “bad” times, the left tail of the income growth distribution expands, while the right tail gets compressed. In other words, during recessions large negative shocks become more likely than large positive ones. They conclude that what previous research was interpreting as a countercyclical variance turns out to be a procyclical skewness. In a follow-up work, Guvenen et al. (2016) document a related empirical regularity: they show that in the data, income growth rate is very small for most individuals, while there is a considerable mass of people with very large growth rates. Therefore, high-order moments (kurtosis in this case) are important features of income growth distribution.

An important difference between this paper and earlier studies of high-order income risk is that we adopt the tradition from labor/consumption literature to distinguish between transitory and permanent shocks to income (Deaton, 1992). This distinction is paramount for understanding how households are insured against permanent and transitory shocks (e.g., Kaplan and Violante (2010)). For example, in this paper we consider insurance channels acting within the household and through the government.

Idiosyncratic shocks lead to a more dispersed distribution of income in the cross-section.

Some studies looked into the welfare implications of rising inequality. For example, [Storesletten et al. \(2001\)](#) argue that welfare benefits from eliminating the countercyclicality of idiosyncratic shocks exceed the benefits of shutting down aggregate technology shocks, which underscores the importance of distributional effects. [Heathcote et al. \(2010\)](#) explore the welfare implications of the increasing wage volatility in the U.S., and find that it benefits recent generations of workers as higher educational premium improves college attainment and redistributes labor within the household.

In a separate but related strand of literature, several studies explored the implications of countercyclical risk for the production side of the economy. In a seminal paper, [Bloom \(2009\)](#) argues that uncertainty (second moment) can have real effects on the macroeconomy through the “wait-and-see” effect. [Bloom et al. \(2018\)](#) estimate productivities of U.S. manufacturing plants, and find robust evidence for the countercyclical nature of productivity shocks. On the same line, [Arellano et al. \(2012\)](#) and [Fernandez-Villaverde et al. \(2011\)](#) explore other propagation mechanisms of the countercyclical risk. [Salgado et al. \(2014\)](#) analyze several publicly available datasets, and consistently find that firm-level data displays the procyclicality of skewness: during recessions, the left tail of the growth rate distribution becomes heavier.

3 Data

We draw on the Panel Study of Income Dynamics (PSID) data, which is the longest available panel data on the U.S. population. PSID started in 1968 with more than 2,000 U.S. families being interviewed on a broad set of topics. The “split-off” families (when family members move out and establish their own households) are also interviewed. PSID spans the time period 1968-2014.

The advantage of this dataset for our purposes is the possibility to simultaneously observe several types of labor income. In particular, we consider 5 different definitions (types) of labor income:

1. head’s hourly wage¹,
2. head’s labor income,
3. head’s labor income (narrow definition),
4. joint labor income (head + wife combined),
5. post-government (taxes and transfers) joint labor income.

As it was mentioned in Section 1, *narrowly defined* labor income refers to those observations for which the labor income exceeds some minimum threshold. This is intended

¹Throughout the paper, we stick to PSID terminology and call a male earner (husband) a household’s head, unless it is a family with a female being the only earner (in this case, wife is the head). A natural alternative is to mark the top-earner within the family as its “head”; this, however, will make our exercise not directly comparable to previous studies based on PSID data, and we, therefore, opt to use a conventional definition instead.

to capture individuals with a strong labor-market attachment. In particular, the income threshold is defined as half of an hourly minimum wage multiplied by 520 hours (13 weeks at 40 hours per week). The post-government joint labor income is equal to the joint labor income (head and wife’s labor income combined) plus government transfers (unemployment compensation, disability insurance and alike) minus taxes (federal and state). PSID provides imputed values of taxes for some years (1978-1990), but we opt to use as many years of data as possible and, therefore, use TAXSIM to obtain our own estimates of state and federal government tax liabilities for the sample of households. The TAXSIM, however, is capable of computing the taxes starting from 1978 — no state-level tax regulations are available prior to that year — which forces us to restrict the sample to years 1978-2014. Detailed explanations of how the variables have been constructed are delegated to Appendix A.1.

PSID is a representative survey data which suffers from well-known issues, such as top-coding of labor income, potential misreporting of income and small sample size. We, however, argue that PSID is an appropriate source of data for our exercise for several reasons. First, we are interested in several types of labor income, which necessitates the knowledge of labor income separately for each spouse, transfers received, wages and hours. This information is typically not simultaneously available in other datasets.

Second, while the top-coding problem is particularly acute for studying income inequality (especially at the right tail of distribution), in this paper we are primarily interested in general swings of income risk over the business cycle. We, therefore, do not expect this issue to contaminate our estimates considerably.

Finally, there is a technical issue with changing frequency: years 1969-1995 are covered with annual frequency, and 1996-2014 with biannual frequency. As it will become clear in Section 4, it is straightforward to handle the gaps in data using our methodology.

The rest of this section — Subsections 3.1 and 3.2 — discusses sample selection and identification of business cycles.

3.1 Sample Selection

We closely follow the sample selection strategy of Storesletten et al. (2004). In PSID, the object of analysis is a family unit (FU). We track heads of FUs as follows: if FU contains a married couple, then the husband is arbitrarily assigned to be the head. Only in incomplete FUs a woman can be the head. In our analysis, we treat split-off families as new independent families. That is, when the head of the household changes, we record it as a new family unit.

Next, we apply a series of selection criteria to construct our dataset. First, FU is in our sample as long as the age of the head is between 23 and 60 years old. By doing so, we pick only those households where the head has a sufficiently strong labor market attachment.

Second, we drop all families with zero or negative total labor income in any year. We also drop families with anomaly labor income growth rates². Observations with top-coded

²That is, we keep household i as long as

$$\ln \left(\frac{y_{it}}{y_{it-1}} \right) \in \left(\frac{1}{20}, 20 \right) \quad \forall t$$

is satisfied (Meghir and Pistaferri (2004) and Storesletten et al. (2004) impose analogous restrictions on the

TABLE 1: SUMMARY STATISTICS

	Wage	Head’s LI	Head’s LI (narrow)	Joint LI	Post-Govt LI
Min	0	0.21	248.05	0.23	372.06
Max	11.80	45,471.57	45,471.57	46,977.23	33,921.60
Median	3.13	6,912.93	6,953.65	9,148.92	8,056.87
Std	2.00	4,390.63	4,362.79	6,049.04	4,572.23
Bottom 5%	0.64	1,643.49	1,803.11	2,187.22	2,620.69
Top 5%	7.19	15,709.26	21,421.60	28,726.57	17,198.53

Notes: Table 1 report the summary statistics for the final dataset based on PSID raw data from years 1978-2014 (see Appendix A.2 for details on data construction). The description of income variables (wages, head’s LI, head’s LI (narrow), joint LI, post-government joint LI) is discussed in full detail in Appendix A.1. All variables are converted into real terms by dividing on CPI with 1968 being the base year.

values are also dropped.

Finally, we drop families which are part of the Survey of Economic Opportunity (1968 ID $\in [5000, 7000]$) or Latino subsample (1968 ID > 7000). This leaves us with approximately 55,000 observations. Table 1 provides the summary statistics for the final dataset.

Appendix A.2 provides richer details on the process of sample selection. Table A.9 reports the number of observations retained at each step of data preparation.

3.2 Identifying Business Cycles

There is no unique way to classify years into “expansions” and “contractions”. Even though PSID is the longest available panel data, its time span covers few recessionary periods as defined by NBER³. It has become a working standard in the literature to classify years into stages of business cycle based on whether the growth rate of some macro aggregate was above or below the long-run mean in that particular year: Storesletten et al. (2004) use GNP per capita, Lee and Mukoyama (2015) and Moreira (2016) use real GDP. We opt to use the real GNP per capita growth rate as a determinant of economic expansions and contractions in our estimation exercise, leaving the discussion of alternatives to Appendix B. One of the reasons we prefer GNP per capita growth rate is that it keeps our exercise close to Storesletten et al. (2004) which is important for comparison purposes. Besides, Appendix A.5 shows that classifications based on GDP and GNP per capita yield comparable results.

4 Estimation Methodology

We adopt the estimation methodology proposed by Storesletten et al. (2004) for three reasons. First, this is a parsimonious way to estimate conditional variance, and estimates reported by Storesletten et al. (2004) serve as a natural reference point. Second, the parametric assumptions this methodology relies on help mitigate small-sample size issues which are typical for easy-to-access datasets. Last but not least, this methodology can be extended

data).

³Recessionary years according to NBER are 1970, 1974-75, 1981-82, 1991, 2001, 2008-09.

to allow for a time-varying skewness (see Subsection 6.2). Third, the methodology can easily accommodate the change in the frequency of PSID from annual to biennial in the middle of the sample period (1996). We next give a brief summary of the estimation methodology in Subsection 4.1, and then provide an identification argument on the workings of this method in Subsection 4.2.

4.1 Overview

Let y_{it}^h be log labor income of household i of age h in year t . We first project log labor income on a set of observables (in line with Storesletten et al. (2004), Busch and Ludwig (2016) and others):

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

where x_{it}^h is the deterministic component of household-specific income attributable to age, education, and family size. Y_t is a measure of aggregate conditions at time t , which picks up the business cycle component of individual labor income.

The residual, u_{it}^h , is a random component, which under standard assumptions satisfies the orthogonality condition

$$\mathbb{E}(u_{it}^h | Y_t, x_{it}^h) = 0 \quad \forall t.$$

Intuitively, the residual captures variation in labor income which cannot be attributed to personal characteristics (like differences in education), and is not explained by the aggregate conditions (information contained in Y_t).

Next, several parametric assumptions are imposed. In particular, it is assumed that the idiosyncratic earnings component u_{it}^h follows the process:

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

Here α_i is a time-invariant fixed effect, which household i is assumed to draw at the beginning of (labor market) life. Next, ε_{it} is a purely transitory component, while z_{it} is a persistent earnings component which follows an AR(1) process. Innovations to both persistent and transitory components, as well as individual fixed effects are normally distributed:

$$\begin{aligned} \alpha_i &\sim \mathcal{N}(0, \sigma_\alpha^2), \\ \varepsilon_{it} &\sim \mathcal{N}(0, \sigma_\varepsilon^2), \\ \eta_{it} &\sim \mathcal{N}(0, \sigma_t^2). \end{aligned}$$

The model is capable of picking up the countercyclical risk of labor income, since it allows the variance of innovations to the persistent component η_{it} to be a function of the aggregate state:

$$\sigma_t^2 = \begin{cases} \sigma_E^2 & \text{if expansion at } t \\ \sigma_C^2 & \text{if recession at } t. \end{cases} \quad (3)$$

Therefore, there are 5 parameters to estimate:

$$\Theta = \{\rho, \sigma_\alpha^2, \sigma_\varepsilon^2, \sigma_E^2, \sigma_C^2\}.$$

As in [Storesletten et al. \(2004\)](#), we estimate Θ by GMM, using the moment conditions that relate the cross-sectional variance of estimated residuals, $\hat{u}_{it}^h$, with the history of shocks households experienced throughout their labor market life⁴. Using normality and independence assumptions, we can express the variance of a labor income shock of family i with the head aged h in year t as:

$$\begin{aligned} \text{Var}(u_{it}^h) &= \text{Var}(\alpha_i + z_{it} + \varepsilon_{it}) \\ &= \sigma_\alpha^2 + \sigma_\varepsilon^2 + \text{Var}(\rho z_{it-1} + \eta_{it}) \\ &= \sigma_\alpha^2 + \sigma_\varepsilon^2 + \rho^2 \text{Var}(\rho z_{it-2} + \eta_{it-1}) + [I_t \sigma_E^2 + (1 - I_t) \sigma_C^2] \\ &= \sigma_\alpha^2 + \sigma_\varepsilon^2 + \rho^4 \text{Var}(\rho z_{it-3} + \eta_{it-2}) + \rho^2 [I_{t-1} \sigma_E^2 + (1 - I_{t-1}) \sigma_C^2] + [I_t \sigma_E^2 + (1 - I_t) \sigma_C^2] \\ &= \sigma_\alpha^2 + \sigma_\varepsilon^2 + \sum_{j=0}^{h-1} \rho^{2j} [I_{t-j} \sigma_E^2 + (1 - I_{t-j}) \sigma_C^2]. \end{aligned} \quad (4)$$

The sample analogue of the population moment (4) takes the form:

$$\frac{1}{N_{ht}} \sum_{i=1}^{N_{ht}} \left\{ (u_{it}^h)^2 - (\sigma_\alpha^2 + \sigma_\varepsilon^2) - \sum_{j=0}^{h-1} \rho^{2j} [I_{t-j} \sigma_E^2 + (1 - I_{t-j}) \sigma_C^2] \right\} = 0. \quad (5)$$

Here, N_{ht} is the number of families at time t with a head aged h . Note that σ_α and σ_ε are not identified separately: [Storesletten et al. \(2004\)](#) use extra moment conditions (with autocovariances of u_{it}^h) in order to disentangle these two parameters. In this paper, we are not interested in either of those parameters separately, and, hence, estimate the sum of σ_α^2 and σ_ε^2 .⁵

There are in total $H \times T$ moments of type (5), with H denoting the number of different ages in the data, and T — the number of available years.⁶ Furthermore, we “aggregate” the moment conditions so that the number of observations in any $\{H, T\}$ -cell does not fall below 100. To accomplish this, we

- break down all feasible ages 23-60 into 4 age groups indexed by $h \in \{25, 35, 45, 55\}$,
- make each group contain ages ± 5 years with group $h = 25$ being an exception.⁷

These adjustments help us balance the two opposing forces: on the one hand, the more moments conditions we use, the more information we extract from the data; on the other hand, more moment conditions lead to some (age, year)-cells being too small.⁸

⁴We assume that individuals enter labor market at the age of 23.

⁵As we show below, our estimates of $\sigma_\alpha^2 + \sigma_\varepsilon^2$ are close to what [Storesletten et al. \(2004\)](#) report.

⁶This amounts to $37 \times 38 = 1,406$ moment conditions.

⁷Precise distribution of ages across 4 groups is as follows: group $h = 25$ contains ages 23-29, group $h = 35$ contains ages 30-39, group $h = 45$ encompasses ages 40 – 49, and group $h = 55$ aggregates the remaining ages 50-60.

⁸[Busch and Ludwig \(2016\)](#) are able to use more age groups since the cross-section is larger in their

4.2 Identification

The way estimation is set up in Subsection 4.1 highlights its benefits (see Equation (4)): even though there are very few families in the dataset whose working life we observe entirely — from the year when its head enters the labor market till the year when he/she retires — we can still incorporate the entire history of business cycle fluctuations that every household experienced over its lifetime into the estimation. In other words, the use of cross-sectional moments for identification allows us to exploit macroeconomic information that predates the micro panel, thereby incorporating more business cycles in the analysis than covered by the sample.

The basic idea behind the entire approach is to exploit how the distribution of persistent idiosyncratic shocks accumulates over time: if the income process is persistent (values of ρ are close to 1 in Equation (2)), then as a cohort ages, the cross-sectional income distribution at any age becomes a function of the sequence of shocks experienced by the cohort’s members. If the variance of income shocks is higher in recessionary years than in expansionary ones, then a cohort that lived through more contractions will have a higher income variance at a given age than a cohort of the same age that lived through fewer contractions. Figure A.11 in Appendix A.3 illustrates this intuition: it shows that the cross-sectional variance of $\hat{u}$ tends to increase as the share of labor market life spent in recessions rises. Each green dot corresponds to the variance of $\hat{u}$ computed across households from the same cohort; the location of dots along the horizontal axis is determined by the share of working life a particular cohort spent in recessionary years. For example, a cohort of 23 year old workers in 2008 will be assigned a value of 1 (it is their first year on the labor market, and 2008 was a recessionary year), and a cohort of 24 year old workers in 2005 will get a value of 0 (both 2004 and 2005 were expansionary years). There are no cohorts among older workers with corresponding shares being strictly either 0 or 1 (see bottom right panel): mature workers have a long enough labor market history which necessarily covers both expansions and contractions. Figure A.11 confirms that the aforementioned upward sloping relationship is present in all 4 age groups, with a somewhat more pronounced pattern for older groups.

Our extension of that approach in Subsection 6.2 is based on the insight that a similar “accumulation” argument holds for skewness. If the probability of a large positive income shock is lower during an aggregate contraction, then the skewness of the shock in a recessionary period will be smaller (more negative) than in an expansion. Therefore, by way of comparing two cohorts of the same age, the distribution of residual income for the cohort that lived through more recessions will exhibit a smaller (more negative) cross-sectional skewness. Figure A.12 in Appendix A.4 confirms this conjecture: the skewness of income shocks decreases as the share of labor market life spent in recessions rises. Again, this pattern is visible for all 4 income definitions, and is more pronounced for mature workers (ages 50-60).

dataset.

5 Volatility of Idiosyncratic Earnings Risk

In this section, we study how the volatility of labor income risk fluctuates across economic expansions and contractions. Conceptually, our exercise is reminiscent of [Storesletten et al. \(2004\)](#) in that we estimate the same parameters using the same moment conditions, but we diverge from them in that we explore the nature of fluctuations in riskiness of multiple income definitions ([Storesletten et al. \(2004\)](#) use joint labor income *after* transfers but *before* tax). Following a bulk of literature on income risk ([Abowd and Card \(1989\)](#), [Meghir and Pistaferri \(2004\)](#) and [Guvenen et al. \(2014\)](#) among others), we also analyze the narrowly defined individual labor income.

By studying different types of labor income, we are able to shed more light on the origins of income risk fluctuations. For example, by moving from hourly wage to head’s labor income, we can speak to the quantitative importance of hours (employment and unemployment) in shaping labor income risk. The intra-family insurance channel can be evaluated through the juxtaposition of risk between head’s and joint (head and wife) labor incomes. Finally, in order to quantify the role of government policy — including both taxes and transfer — in possibly alleviating the cyclical nature of labor income risk, we assess to what extent (pre-government) joint labor income are more volatile than post-government income.

We first conduct a graphical analysis in Subsection 5.1, before providing the estimates in Subsection 5.2. In Subsection 5.3, we look at subgroups, in particular, the young and the old.

5.1 Graphical Analysis

In order to shed light on the (counter)cyclical nature of idiosyncratic income shock volatility, first we need to obtain the residuals u_{it}^h . We estimate Equation (1) by running a pooled regression (we also experimented with estimating a panel regression, but the results did not change significantly).

We consider the following specification of function $g(\cdot)$:

$$g(x_{it}^h, Y_t) = \theta_0 + \theta_1' D(Y_t) + \theta_2' x_{it}^h + \theta_3' f(x_{it}^h, Y_t),$$

where x_{it}^h includes the following list of observables: cubic polynomial in age, education of head, and the size of the family. Aggregate effects are absorbed in two ways: first, we include a full set of year dummy-variables $D(Y_t)$. Second, we allow for educational premium to vary over the business cycle, and therefore include a quadratic (in education) polynomial $f(\cdot)$:

$$f(x_{it}^h, Y_t) = D(Y_t) \times (\text{Education}_{it}^h + [\text{Education}_{it}^h]^2).$$

Results in Table 2 are consistent with a wide body of literature: the earnings age profile is concave and increasing in education, large family sizes are associated with high labor income. All estimates are statistically significant and have the expected sign.

We subsequently retrieve $\hat{u}_{it}^h$ as residuals from the estimated Equation (1). Figure 1, which is close to Figure 1(d) in [Storesletten et al. \(2004\)](#), plots the evolution of the cross-sectional

TABLE 2: ESTIMATION OF EQUATION 1

	DEPENDENT VARIABLE $\log(y_{it})$				
	Wage	Head's LI	Head's LI (narrow)	Joint LI	Post-govt LI
Age	0.136 (0.013)	0.168 (0.016)	0.186 (0.014)	0.186 (0.015)	0.191 (0.012)
Age ²	-0.003 (0.000)	-0.003 (0.000)	-0.004 (0.000)	-0.004 (0.000)	-0.004 (0.000)
Age ³	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
Education	0.052 (0.005)	0.057 (0.007)	0.057 (0.006)	0.063 (0.006)	0.049 (0.005)
Family size	0.041 (0.002)	0.067 (0.002)	0.065 (0.002)	0.122 (0.002)	0.113 (0.002)
N	53,210	55,209	54,840	55,209	55,209
Adj. R^2	0.14	0.12	0.14	0.16	0.21

Notes: Table 2 reports the results of OLS estimation, and is based on Family Files from PSID over the period 1978-2014. **Age** is the age of a household's head, **Education** is a number of completed (by the head) years of college. Standard errors are in parentheses. Regressions also include yearly dummies and time-dependent educational premium (not reported). All reported coefficients are significant at 1% level.

mean of $\log$ (post-government) joint labor income and the standard deviation of $\hat{u}_{it}^h$.⁹ The figure displays a robust feature of the data: the standard deviation of joint labor income shocks is countercyclical, decreasing in expansions and increasing in recessions. The labor income risk increases during all NBER recession years within the sample period: 1980-1981, 1991, 2001, and 2007-2009.

While Figure 1 indicates that the countercyclical income risk is a robust feature of the data, it does not say much about where this cyclicity originates from. This motivates us to repeat the exercise for other income definitions.

A visual inspection of Figure 2 suggests that the countercyclical nature of income risk is not an artifact of the post-government joint labor income: all five series appear to be negatively correlated with mean income. One can also notice that despite exhibiting similar dynamics, there is a visibly pronounced heterogeneity in fluctuation of risk across income types. For example, head's labor income and (pre-government) joint labor income (red and orange lines) track each other fairly closely, while the line corresponding to the post-government joint labor income is visually twice less volatile.

While Figure 2 is suggestive, it is barely useful to compare how the countercyclical nature of different income definitions relate to each other. We, therefore, categorize every sample year into one of 4 bins, depending on the growth rate of real GNP per capita in that year: if GNP per capita grew by a lot (in top quantile of growth rate distribution), we place that year in bin 4. Conversely, if the growth was in a bottom quantile of the growth rate distribution, that year falls in bin 1. Subsequently, we take the average (across years which are sorted in

⁹The main difference is that, while Storesletten et al. (2004) use joint labor income plus government transfers, and do not subtract taxes, we take into account both government transfers and taxes.

FIGURE 1: MEAN LOG EARNINGS AND STANDARD DEVIATION OF $\hat{u}_{it}^h$

Notes: Figure 1 is based on Family Files from PSID over the period 1978-2014. The mean labor income in year t is a cross-sectional mean of log (post-government) joint labor income in a corresponding year. The standard deviation of joint labor income shocks in year t is a cross-sectional standard deviation of $\hat{u}_{it}^h$ - residuals from the estimated Equation (1). We also subtract linear trends from the series, which chiefly eliminates the long-run mean (the slope coefficient is nearly zero). Grey bars represent NBER recessions.

a particular bin) deviation of a corresponding statistic from its long-run mean, and do it for all 5 different income definitions (Figure 3).

Figure 3 anticipates several findings. At the very least, it reiterates the message of Figure 2: the volatility of income shocks exhibits countercyclicality, and this pattern is robust across different labor income definitions. However, we can say more than that. First, consistently with findings of Guvenen et al. (2014), the narrowly-defined head's labor income exhibits relatively modest fluctuations in risk over the cycle. Second, head's wages are less volatile and less cyclical than head's labor income, pointing at the importance of hours in driving the labor income risk fluctuations. Third, joint labor income, if anything, exhibits fluctuations in income risk which are comparable to those of the head's labor income. And, finally, post-government income risk fluctuations are moderate, quantitatively similar to those of the narrowly-defined individual labor income.

These cyclical properties we observe are robust to alternative definitions of the business cycle. In Appendix B, we provide analogous figures where we categorize years based on mean income growth rate (Subsection B.2) and by the NBER definition of recessions (B.1).

FIGURE 2: STANDARD DEVIATION OF IDIOSYNCRATIC INCOME COMPONENT

Notes: Figure 2 is based on Family Files from PSID over the period 1978-2014. The mean labor income in year t is a cross-sectional mean of log post-government joint labor income in a corresponding year. The standard deviation of labor income shocks in year t is a cross-sectional standard deviation of $\hat{u}_{it}^h$ - residuals from the estimated Equation (1), where the left-hand side variable is one from the set $y_{it} \in \{\text{head's wage, head's income, head's income (narrow), joint labor income, post-government joint labor income}\}$. We also demean the resulting series. Grey bars represent NBER recessions.

5.2 Estimation Results

In Subsection 5.1, we provided suggestive evidence on the countercyclical nature of income shocks volatility. In this section, we take a step forward and estimate the vector of structural parameters which govern the income process (2). As it has been discussed above, there are in total $H \times T$ moment conditions of type (5). In Section 4 we argued that we cannot use all of them, as the sample size of certain age-year cells becomes too small to obtain precise estimates. Instead, we focus on a subset of moment conditions which correspond to ages 25, 35, 45 and 55. We check that there are at least 100 observations in each cell.

Our baseline business cycle indicator is based on real GNP per capita. That is, we set $I_t = 1$ if in year t real GNP per capita growth rate exceeded the sample mean, and assign $I_t = 0$ otherwise (Appendix A.5 compares the allocation of sample years into “expansions” and “contractions” based on alternative aggregate measures).

Table 3 provides GMM estimates for all 5 income definitions. Our results reconcile the findings of previous studies with seemingly contradicting results. On one hand, post-government joint labor income exhibits a sizable countercyclical risk. The ratio of our estimates is somewhat lower than what Storesletten et al. (2004) report, but still within the range of estimates they provide. The discrepancy might arise because of taxes: Storesletten

FIGURE 3: VOLATILITY OF IDIOSYNCRATIC INCOME RISK BY GNP PER CAPITA GROWTH QUANTILE

Notes: Figure 3 is based on Family Files from PSID over the period 1978-2014. Each year from the period 1978-2014 is classified in one out of 4 bins, depending on which quantile the growth rate of GNP per capita in that year falls into. Quantile 1 contains years with the lowest growth rate of GNP per capita, while quantile 4 contains years with the highest growth rates. The standard deviations shown are averages over years in the bin. Each quantile contains standard deviations for 5 measures of labor income: head's wage, head's labor income, joint labor income, post-government joint labor income, and head's labor income (narrow definition).

TABLE 3: GMM ESTIMATION RESULTS: TIME-VARYING VOLATILITY ONLY

Type of Income	ρ	σ_E	σ_C	$\sqrt{\sigma_\alpha^2 + \sigma_\varepsilon^2}$	$\sigma_C - \sigma_E$
Head Hourly Wage	0.91 (0.00)	0.10 (0.01)	0.15 (0.01)	0.69 (0.01)	0.05
Head Labor Income	0.84 (0.01)	0.11 (0.01)	0.20 (0.01)	0.79 (0.01)	0.09
Head LI (narrow definition)	0.84 (0.00)	0.11 (0.01)	0.15 (0.01)	0.76 (0.01)	0.04
Joint (Head+Wife) LI	0.77 (0.01)	0.16 (0.03)	0.23 (0.02)	0.78 (0.03)	0.07
Post-Govt Joint LI	0.81 (0.01)	0.09 (0.03)	0.14 (0.01)	0.69 (0.01)	0.05

Notes: Table 3 reports the estimation results for Θ by GMM based on the moment conditions of type (5). In particular, we used moment conditions corresponding to ages 25, 35, 45 and 55. Equation (4) do not include autocorrelation moments, and, therefore, only a sum of σ_α and σ_ε is estimated. Standard errors are in parentheses.

et al. (2004) definition of labor income corresponds to joint labor income *after* government transfers but *before* taxes, while we incorporate both types of government redistribution policies. Heathcote et al. (2010) study the distributional effects of taxes and transfers and find that they compress the earnings inequality, especially at the bottom of the distribution.

On the other hand, the narrowly-defined head’s labor income exhibits countercyclical variance, but its countercyclicality is noticeably weaker than that of head’s labor income. Guvenen et al. (2014) find that the second-moment of income risk is flat with respect to the business cycle, while our findings suggest moderate countercyclical fluctuations. This difference might arise from several sources, including the way we identified the income shock (residual from OLS regression, rather than income growth), the estimation approach (parametric, rather than non-parametric), and different data used (PSID vs. Social Security Administration records). This highlights the heterogeneity of income risk fluctuations across different definitions.

Indeed, the head’s labor income risk shows the greatest countercyclicality, with the recessionary standard deviation being 2 times larger than the expansionary one. Wage rate does not show such a pronounced countercyclicality, hinting towards an important quantitative role of hours (most likely, employment and unemployment). This finding mirrors the observation from Figure 3. Moving from head’s labor income to joint labor income, we see that intra-family insurance channel through added worker effect reduces fluctuation in risk from 0.09 down to 0.07. The limited quantitative role of this channel could also have been anticipated from Figure 3. Finally, the taxes and transfer by the government further mitigate fluctuations in risk, from 0.07 to 0.05.

5.3 Analysis of Subgroups

In this subsection, we break down the sample into 2 subgroups — “young” (ages 22-39) and “old” (ages 40-60) — and re-estimating the parameters on those subsamples, in order to uncover a striking heterogeneity between the age groups (see Table 4). The column “overall” corresponds to the case when we estimate the parameters without differentiating between expansions and contractions. While the “overall” estimates of risk for old subsample are higher than the ones for an young subsample (by a factor of 2-3 for head’s hourly wage and head’s labor income), it appears that all business cycle fluctuations in risk come from young people. Households aged 40-60 at best exhibit only minor countercyclicality in income risk across all 5 income definitions.

6 Skewness of Idiosyncratic Income Risk

In this section, we extend our analysis to allow for a time-varying skewness of income shocks. While countercyclical variance can tell us that tail events (large positive and negative shock realizations) become more likely during economic downturns and had been the primary focus of the literature until recently, there is a growing body of literature highlighting the importance of a third moment (*inter alia* Salgado et al. (2014), Guvenen et al. (2014), Busch et al. (2018)). Non-zero skewness implies that some extreme shock realizations are likely to be either positive or negative — depending on the sign of the coefficient of skewness. This

TABLE 4: GMM ESTIMATES OF EARNINGS SHOCK VOLATILITY BY AGE GROUP

All Households, Age 22-39				
	Overall	Exp	Rec	Diff
Head Hourly Wage	0.09 (0.02)	0.06 (0.02)	0.15 (0.02)	0.09
Head Labor Income	0.12 (0.03)	0.09 (0.04)	0.21 (0.02)	0.12
Head LI (narrow definition)	0.10 (0.02)	0.06 (0.02)	0.11 (0.03)	0.06
Joint (Head+Wife) LI	0.12 (0.03)	0.08 (0.03)	0.18 (0.04)	0.10
Post-Govt Joint LI	0.14 (0.02)	0.02 (0.01)	0.16 (0.06)	0.14
All Households, Age 40-60				
	Overall	Exp	Rec	Diff
Head Hourly Wage	0.26 (0.06)	0.26 (0.06)	0.28 (0.07)	0.02
Head Labor Income	0.22 (0.05)	0.22 (0.06)	0.26 (0.06)	0.04
Head LI (narrow definition)	0.21 (0.08)	0.22 (0.09)	0.25 (0.09)	0.05
Joint (Head+Wife) LI	0.15 (0.05)	0.13 (0.05)	0.21 (0.04)	0.08
Post-Govt Joint LI	0.13 (0.02)	0.11 (0.02)	0.17 (0.01)	0.06

Notes: Table 4 is based on PSID Family Files over the period 1978-2014. The top part of the table corresponds to the “young” subsample (22-39 years old), the bottom part — to the “old” subsample (40-60 years old). “Overall” columns contain estimates when σ_t is assumed to be time-invariant. “Diff” column contains the distance between estimates of σ_E and σ_R . Standard errors are in parentheses.

also implies that constant skewness — something our analysis has implicitly assumed so far — can mask a rich heterogeneity between left- and right-tail events. We proceed in the same way as in the previous section on the cyclicity of variance. We first graphically confirm that these phenomena are present in our dataset (Subsection 6.1). Then we explore the quantitative magnitude of skewness fluctuations over the business cycle for all 5 different income definitions (Subsection 6.2). By analyzing 5 income definitions separately, we can answer whether some income definitions exhibit stronger procyclicality of skewness than others.

6.1 Graphical Analysis

Throughout our analysis of skewness, we consider the following 2 conventional measures:

FIGURE 4: SKEWNESS OF IDIOSYNCRATIC LABOR INCOME RISK, BY GNP PER CAPITA GROWTH QUANTILE

Notes: Figure 4 is based on PSID Family Files over the period 1978-2014. Panel A plots the third central moment, Panel B plots the Kelly measure.

1. third central moment:

$$\text{Third moment}_t = \frac{1}{n_t} \sum_i (\hat{u}_{it}^h - \bar{\hat{u}}_{it}^h)^3,$$

2. Kelly's measure:

$$\text{Kelly}_t = \frac{(P90_t - P50_t) - (P50_t - P10_t)}{P90_t - P10_t}.$$

The interpretation of these statistics is straightforward. The first one is a sample analogue of the third central moment, which is a measure of skewness by definition. The second one (Kelly measure) is a function of several percentiles of $\hat{u}_{it}^h$ -distribution, which makes it robust to “extreme” observations (note that it is independent from the first and last deciles of the underlying distribution). Intuitively, Kelly's measure computes the difference in inequality between the right ($P90 - P50$) and left ($P50 - P10$) tails, and relates it to the overall variation in the sample ($P90 - P10$). If the right tail is heavier than the left one (underlying distribution of $\hat{u}_{it}^h$ has a positive skew), then the Kelly measure is positive. And the other way around, a heavier left tail makes the Kelly's measure negative.

Procyclical skewness implies that during economic upturns (downturns), the right (left) tail of income shocks thickens, leading to a disproportionate bigger fraction of large positive (negative) shocks. At the same time, the odds of receiving a large negative (positive) shock go down.

Figure 4 illustrates the evolution of two measures of skewness over the sample period: Panel A plots the third central moment, while Panel B plots the Kelly's measure. Several observations emerge. First, both measures of skewness are procyclical, increasing in expansions and decreasing in contractions. Both measures of skewness tend to go down during NBER recession dates (with an exception of the head's wage during 1992 and 2001 recessions).

FIGURE 5: SKEWNESS OF IDIOSYNCRATIC LABOR INCOME RISK, BY GNP PER CAPITA GROWTH QUANTILE

Notes: Figure 5 is based on PSID Family Files over the period 1978-2014. Panel A plots the third central moment, Panel B plots the Kelly measure. Each year from the period 1978-2014 is classified in one out of 4 bins, depending on which quantile the growth rate of GNP per capita in that year falls into. Quantile 1 contains years with the lowest growth rate of GNP per capita, while quantile 4 contains years with the highest growth rates. The measures of skewness shown are averages over years in the bin. Each quantile contains skewness measures for 5 different types of labor income: head's wage, head's labor income, joint labor income, post-government joint labor income and head's labor income (narrow definition).

Second, the third moment is a more volatile measure of skewness as compared to the Kelly's measure — something one could have anticipated from our discussion about their differences above. Panel A of Figure 4 shows that the skewness of head's and joint labor incomes fluctuates the most across income definitions we consider. This implies, for example, that probabilities of tail events co-move very strongly with aggregate conditions for these income types. The remaining definitions also exhibit procyclical skewness, but cyclical movements of their tail events are by an order of magnitude weaker. Kelly's measure in Panel B is more stable across income types, which suggests that tail observations (above the 90th and below the 10th percentiles) are quantitatively important drivers of the cyclical skewness.

In order to better understand which income definitions exhibit stronger/weaker cyclical skewness, we adapt a similar to Section 5 graphical approach. In Figure 5, we plot the average (deviation from the trend of) skewness of income shocks by GNP per capita growth quantile. Panel A plots the third central moment, while Panel B plots the Kelly measure. The figure confirms that the skewness is procyclical, which is robust across types of income and measures of skewness. That means that during economic downturns, a large negative income shock is more likely than a large positive one (the left tail becomes thicker).

A closer inspection of Figure 5 reveals that the skewness of post-government joint labor income turns out to be among the most stable across types of income over the cycle. It implies that the odds of getting a very negative shock for that income definition co-move with the odds of getting a very positive shock. We can also visually confirm that the Kelly's measure ranks post-government joint labor income as the most “stable” definition.

When we look at household’s income without transfers and taxes (joint labor income), we find that the cyclical nature of skewness increases with the difference being the starkest in the right panel of Figure 5: when we discard the taxes and transfers, the probability of getting a large negative income shock increases by more than the odds of getting a large positive shock during economic contractions. This observation can be interpreted as suggestive evidence of the insurance or stabilizing role of the government.

It is hard to rank the remaining income definitions, but it is possible to see that the head’s labor income exhibits the strongest fluctuations in skewness. This observation might reflect the fact that spouse provides a significant amount of intra-family insurance through an added worker effect: during economic downturns the probability of getting laid off increases, and the wife can step in and compensate for the head’s job loss (by working more hours, getting an extra job, etc.). The head’s wage exhibits relatively moderate fluctuations in skewness (Kelly’s measure is almost acyclical for head’s hourly wage).

Figures B.14 and B.16 in Appendix B display similar patterns and confirm that the above observations are robust to alternative ways of business cycle identification: with respect to the growth rate of mean income, and according to the NBER definition of recessions.

Appendix C explores the cyclical nature of shocks when those are identified as the growth rate of income (Guvenen et al. (2014), Busch et al. (2018) among others). This approach is non-parametric, and allows to study fluctuations in risk with very few identifying assumptions.¹⁰ While in this paper, we opt to use the parametric approach given the size of the dataset, it is important to establish the connection between the two approaches given the data at hand. Taking into account the fact that PSID became biannual starting from 1996, we decided to compute the 2-year growth rate as our alternative measure of income shocks. Figure C.17 shows that the countercyclical nature of labor income shocks carries over to that alternative way of shock identification. Surprisingly, the ranking of income types with respect to their cyclical nature is retained.

Fluctuations in skewness are also preserved when shocks are measured by growth rates (Figure C.18). Both the third central moment and Kelly measure are strongly procyclical, rising sharply during aggregate expansions. Overall, our key observations are robust to the non-parametric way of measurement of income risk.

6.2 Joint Estimation of Cyclical Volatility and Skewness

The objective of this subsection is to jointly estimate both the volatility (second moment) and skewness (third moment) of income shocks. Technically, we still assume that the labor income shock follows process (2), but innovations to a persistent component η_{it} are now drawn from a *skew normal distribution*, which is a generalization of a normal distribution to the case with a non-zero skewness. The skew normal distribution is a family of probability distributions governed by 3 parameters: location ($\zeta \in \mathbb{R}$), scale ($\omega \in \mathbb{R}_{++}$), and shape

¹⁰Guvenen et al. (2014) differentiate between transitory and persistent components of income. *Transitory* component is measured as 1-year growth rate ($\log(y_{it}) - \log(y_{it-1})$). *Persistent* component is a 5-year growth rate ($\log(y_{it}) - \log(y_{it-5})$).

FIGURE 6: SKEW NORMAL DENSITY FOR DIFFERENT VALUES OF ν

Notes: Figure 6 plots the skew normal density for 5 values of $\nu \in \{-1, -0.5, 0, 0.5, 1\}$. The shape parameter ω is 0.1, and location parameter ζ is set to 0.

($\nu \in \mathbb{R}$).¹¹ Innovations to persistent component are drawn from the distribution:

$$\eta_{it} \sim \mathcal{SN}(\zeta, \omega_t, \nu_t). \quad (6)$$

We assume that the location parameter is business cycle invariant, and we normalize it to 0. Figure 6 shows how the shape parameter ν governs the third moment of the skew normal distribution: the corresponding density tends to be skewed towards more positive values (positive skew) for positive ν , and towards negative values (negative skew) for negative values of ν . The shape parameter ω is set equal to 0.1 across all 5 shown densities.

Crucially, we make variance and skewness state-dependent, allowing the shock structure to change between expansions and contractions. The support of both ω_t and ν_t consists of two points:

$$\omega_t = \begin{cases} \omega_E & \text{if expansion at } t \\ \omega_C & \text{if recession at } t \end{cases}$$

and

$$\nu_t = \begin{cases} \nu_E & \text{if expansion at } t \\ \nu_C & \text{if recession at } t. \end{cases}$$

Next, we estimate a set of 6 parameters jointly. First, analogously to the case with a state-dependent volatility, it is straightforward to show that the following $H \times T$ moment

¹¹The p.d.f. of the skew normal distribution is $f(x) = 2\phi(x)\Phi(\nu x)$, where $\phi(x)$ and $\Phi(x)$ are p.d.f. and c.d.f. of the standard normal distribution, respectively. In case of symmetric distribution ($\nu = 0$), the formula collapses to a standard normal p.d.f.

TABLE 5: GMM ESTIMATION RESULTS: TIME-VARYING VOLATILITY AND SKEWNESS

Type of Income	ω_E	ω_C	δ_E	δ_C	ρ	$\sqrt{\sigma_\alpha^2 + \sigma_\varepsilon^2}$
Head Hourly Wage	0.10 (0.01)	0.17 (0.02)	-0.54 (0.01)	-0.57 (0.01)	0.91 (0.01)	0.69 (0.02)
Head Labor Income	0.12 (0.06)	0.25 (0.11)	-0.67 (0.03)	-0.70 (0.02)	0.76 (0.06)	0.78 (0.12)
Head LI (narrow definition)	0.13 (0.05)	0.21 (0.11)	-0.67 (0.03)	-0.70 (0.03)	0.75 (0.05)	0.74 (0.05)
Joint (Head+Wife) LI	0.16 (0.02)	0.25 (0.04)	-0.56 (0.02)	-0.62 (0.02)	0.86 (0.02)	0.78 (0.02)
Post-Govt Joint LI	0.10 (0.05)	0.17 (0.08)	-0.58 (0.02)	-0.64 (0.02)	0.78 (0.04)	0.69 (0.03)

Notes: Table 5 reports the estimation results for Θ by GMM based on the moment conditions of type (7) and (8). In particular, we used moment conditions corresponding to ages 25, 35, 45 and 55. Parameters σ_α and σ_ε are not identified separately; only their sum is estimated. Standard errors are in parentheses.

conditions must be satisfied:¹²

$$\mathbb{E}_t \left[(u_{it}^h)^2 - \sigma_\alpha^2 - \sigma_\varepsilon^2 - \sum_{j=0}^{h-1} \rho^{2j} \left\{ I_{t-j} \left(\omega_E^2 \left[1 - \frac{2\delta_E^2}{\pi} \right] \right) + (1 - I_{t-j}) \left(\omega_C^2 \left[1 - \frac{2\delta_C^2}{\pi} \right] \right) \right\} \right] = 0, \quad (7)$$

where $\delta_i = \frac{\nu_i}{\sqrt{1+\nu_i^2}}$ and $i \in \{E, C\}$.

We use a set of extra $H \times T$ moment conditions (in addition to moments of type (7)), in order to identify the parameters ν_E and ν_C by way of relating empirical skewness to theoretical one:

$$\mathbb{E}_t \left[\text{skew}(u_{it}^h) - \underbrace{\text{skew}(\sigma_\alpha^2 + \sigma_\varepsilon^2)}_{=0} - \sum_{j=0}^{h-1} \rho^{3j} \left\{ I_{t-j} \gamma_E + (1 - I_{t-j}) \gamma_C \right\} \right] = 0, \quad (8)$$

where $\gamma_i = \frac{4-\pi}{2} \frac{(\delta_i \sqrt{2/\pi})^3}{(1-2\delta_i^2/\pi)^{\frac{3}{2}}}$ and $i \in \{E, C\}$.

In total, we have $2 \times H \times T$ moment conditions to estimate the following 7 parameters: $\rho, \omega_E, \omega_C, \delta_E, \delta_C, \sigma_\alpha, \sigma_\varepsilon$. Similarly to the case with time-varying heteroscedasticity only (Section 5), we do not estimate σ_α and σ_ε separately but only their sum.¹³ That makes the number of parameters to be estimated equal 6.

Table 5 reports the parameter estimates for 5 definitions of labor income. Note, however, that parameters governing the skew normal distribution do not map one-to-one to parameters

¹²If a random variable $Y \sim \mathcal{SN}(\zeta, \omega, \nu)$ has a skewed normal distribution, then $\mathbb{E}[Y] = \zeta + \omega \delta \sqrt{\frac{2}{\pi}}$ and $\text{Var}[Y] = \omega^2 \left(1 - \frac{2\delta^2}{\pi} \right)$, where $\delta = \frac{\nu}{\sqrt{1+\nu^2}}$. Skewness can then be expressed as $\gamma = \frac{4-\pi}{2} \frac{(\delta \sqrt{2/\pi})^3}{(1-2\delta^2/\pi)^{\frac{3}{2}}}$.

¹³Even though we do not estimate these parameters separately, we still confirm in Table 5 that their sum is close to what Storesletten et al. (2004) report.

of the normal distribution: for example, the variance is now a function of both ω and δ (see Footnote 12). To deal with this issue, in Table 6 we report the variance and skewness for all 5 income definitions that are implied by the parameter estimates shown in Table 5.

Let us start with volatility of labor income risk implied by estimated parameters. We are interested in seeing how consistent they are with respect to our results from Subsection 5.2. Overall, our previous findings remain intact when we allow for a time-varying skewness. First, we find that the implied estimates of volatility are very close to what we reported when allowed for time-varying second moment only (Section 5). This result is not mechanical as the variance is now a function of 2 parameters, of which one affects both the volatility and skewness. In this sense, the objective function for GMM is not “block-recursive”. We also note that the standard errors (Table 5) became larger, pushing the significance of some of estimated parameters to borderline values. On possible interpretation of the countercyclical variance result of Storesletten et al. (2004) is that the countercyclical variance is obtained because skewness is not allowed to be time-varying, but this hypothesis is not supported by our findings; we obtain countercyclical volatility of labor income risk even if we allow skewness to be state-dependent as well.

Second, not only that, our estimated parameters imply that the head’s labor income has the strongest cyclicity of income volatility, as we found in Subsection 5.2. The fifth column of Table 6 (labeled $\Delta\sigma$) shows the differences between the implied standard deviations in recessions and expansions. The column shows that the difference in the standard deviation is the largest with head’s labor income. It is 0.10 when skewness is also allowed to be state-dependent, and 0.09 when it is not.

Third, cyclicity of income variance becomes weaker if we use narrower definition of head’s labor income, as we found in Subsection 5.2. The difference in the standard deviation of head’s labor income declines from 0.10 to 0.06 when we switch to the narrower definition of head’s labor income, while the decline is from 0.09 to 0.04. Indeed, narrowly-defined head’s labor income risk is still among the most stable definitions over the economic cycle, which is in line both with our previous findings.

Fourth, we find that both intra-household insurance and government taxes and transfers ameliorate cyclicity of labor income risk, no matter if we include state-dependency of skewness in our estimations, although the estimated size of the effects is slightly different. In the case with time-varying skewness of income shocks, the standard deviation of labor income risk is ameliorated (from 0.10 to 0.07) when we consider incomes of the wife, while the role of government taxes and transfers in lowering cyclicity of income risk seem weaker (from 0.07 to 0.06). If we do not allow state-dependent skewness, both intra-household insurance and government insurance lower the difference in the standard deviation of labor income risk by 2 percentage points.

Next, we look at the skewness implied by the estimated parameters. The third and the fourth columns of Table 6 show the implied skewness in expansions and recessions, respectively, and the last column shows their difference. Let us make three comments. First, skewness for both expansions and recessions are the highest with head’s labor income and skewness of procyclical (probability of a large decline increases in recessions). Even with the narrower definition of head’s labor income, the levels of skewness in both expansions and recessions are not significantly affected, and skewness is equally procyclical in both normal and narrow definitions of head’s labor income. It is not surprising, therefore, that Guvenen

TABLE 6: GMM-IMPLIED VOLATILITY AND SKEWNESS

Type of Income	σ_E	σ_C	skew $_E$	skew $_C$	$\Delta\sigma$	Δskew
Head Hourly Wage	0.09	0.15	-0.04	-0.06	0.06	-0.02
Head Labor Income	0.10	0.20	-0.10	-0.13	0.10	-0.03
Head LI (narrow definition)	0.11	0.17	-0.11	-0.13	0.06	-0.02
Joint (Head+Wife) LI	0.14	0.21	-0.05	-0.08	0.07	-0.03
Post-Govt Joint LI	0.09	0.15	-0.06	-0.09	0.06	-0.03

Notes: Table 6 reports the standard deviation and skewness of income shocks implied by GMM estimates from Table 5 using the mapping from parameters into moments from Footnote 12. The last two columns report the differences in the estimated standard deviations and coefficients of skewness across expansions and contractions.

et al. (2014) find skewness is procyclical using narrow definition of head’s labor income, while they find acyclical variance.

Second, having wife’s income lowers the level of skewness in both expansions and recessions but its cyclicalities is not significantly affected. Third, on the other hand, incorporating government taxes and transfers does not seem to affect both the level and the cyclicalities of skewness of income shocks. Finally, head’s hourly wage exhibits the lowest level of skewness, which implies that houses, most likely due to (un)employment, play an important role in shaping the downside risk of income shocks, in both expansions and recessions, but the role of hours does not seem to change depending on the aggregate state.

While we do not see significant heterogeneity in the cyclicalities of skewness across different income definitions, it does not imply, however, that tail events evolve similarly across different income definitions in expansions and recessions. The key reason of why it is the case is based on the fact that it is both the volatility and skewness of income shocks which simultaneously determine the probability of tail events. In the next section, we explore the implications of the obtained estimates, and put them into economic context.

7 Economic Interpretation

In this section, we provide economic interpretations for our estimates. First, in Subsection 7.1, we show with the help of graphical apparatus, how shock distributions change depending on the aggregate economic activity. In particular, we illustrate that the resulting distributions are visually very different across the stages of the business cycle.

Second, in Subsection 7.2 we quantify the probabilities of tail events for different income definitions: in particular, we ask by how much the probability of a large increase/decrease in income goes up in recessions as compared to expansions. Our results will recover a substantial heterogeneity in probabilities of tail events across income types, despite the fact that estimates are visually quite similar.

7.1 Estimated Distributions of Income Shocks

In this subsection, we plot the resulting distribution of shocks η_{it} for parameter estimates from Table 5. Figure 7 shows distributions of shocks η_{it} in expansions (solid red lines) and in recessions (dashed blue lines), for 5 labor income definitions. The most striking difference between expansion and contraction distributions of income shocks is observed in case of the head’s labor income (panel B): while the solid red curve (expansion) is fairly symmetric around 0 with a small variance, the blue dashed line is substantially more dispersed with the left tail being visually heavier than the right one. This observation implies that individuals are more likely to be hit by a large negative (rather than positive) shock during economic contractions. However, when the economy is expanding, tail events, both positive and negative, become almost equally likely, although the overall probability of large shocks becomes smaller.

Panel A of Figure 7 confirms again that it is the number of hours worked rather than hourly wages that drives a significant portion of countercyclical head’s labor income risk. In particular, we see that the recessionary distribution of wage shocks gets substantially closer to the expansionary one, compared with head’s labor income (Panel A), even though we still observe a strongly pronounced countercyclicality of wage shocks. We will make this point clear below when we estimate the probabilities of tail events.

Intra-family insurance channels are visible in Panel D of Figure 7 — variance is more stable, and the probability of tail events drops as compared to head’s labor income — however, it appears that this channel has a limited quantitative role. Government transfers and taxes (Panel E) smooth out countercyclical risk and procyclical skewness significantly, even after spousal channel is taken into account — this result has been anticipated since Section 5. Finally, labor income risk for people with a strong market attachment (narrow definition) exhibits moderate swings in risk and skewness over the cycle, if anything, the distributions in Panel C are similar to those in panel A (wages).

7.2 Asymmetries in Tail Events

What are the implications of our findings? We ask: by how much does the probability of a tail event (50% and 100% increase/decrease in residual income) change in recessions as compared to expansions for all 5 different income definitions? In order to accomplish this, we first simulate a large number of individual (residual) income histories ($N=100,000$), assuming that their incomes follow Equation (2) with shocks distributed according to Equation (6). The number of simulated periods is 1000, where each period is either a recession or an expansion. At the simulation step, we use the corresponding “expansionary” or “recessionary” parameters of income shocks depending on the contemporaneous aggregate state.

We end up having a large panel dataset ($100,000 \times 1000$). At this stage, we have sufficient information to compute income changes, $u_{it} - u_{it-1}$, and assess the probabilities of tail events.¹⁴ We sort all observations into one of 2 bins. The first bin contains all income changes which happened during “good” aggregate times, while the second one — during bad

¹⁴Technically, we obtain a distribution of (residual) income changes, from which we get the probabilities of tail events.

FIGURE 7: ESTIMATED DISTRIBUTIONS OF SHOCKS, EXPANSIONS AND RECESSIONS

Notes: Figure 7 plots the estimated distributions of residual income shocks (Equation (6)) for economic expansions (solid red) and contractions (dashed blue). The parameters of those distributions are taken from Table 5.

times. We then plot the histograms of income changes in Figure 8. The blue area represents the histogram in recessions while the red area represents the histogram in expansions.

Several observations are in order. First, and as expected, the distributions of income changes at recessionary times (in blue) have heavy tails. Second, the head's labor income (Panel B) exhibits the most pronounced difference between the red and blue distributions: the probability of large income drops (more than 100%) is the largest among the income types considered. Third, head's wages, post-government joint labor income as well as narrowly-

FIGURE 8: SIMULATED DISTRIBUTIONS OF (RESIDUAL) INCOME CHANGES

Notes: Figure 8 plots the simulated distributions of (residual) income changes for all 5 income definitions considered in this paper. Histograms in red correspond to the case of economic expansions, while in blue — to contractions. See Subsection 7.2 for the simulation details.

TABLE 7: ESTIMATED PROBABILITIES OF TAIL EVENTS, %

	Expansions		Recessions	
	$\leq -50\%$	$\leq -100\%$	$\leq -50\%$	$\leq -100\%$
Head Hourly Wage	7.9	0.3	10.6	0.6
Head LI	7.4	0.2	13.3	1.0
Head LI (narrow)	7.9	0.3	11.4	0.7
Joint (Head+Wife) LI	8.6	0.4	12.9	1.0
Post-Govt Joint LI	7.8	0.3	10.6	0.5
	$\geq +50\%$	$\geq +100\%$	$\geq +50\%$	$\geq +100\%$
Head Hourly Wage	9.4	0.4	9.0	0.4
Head LI	10.4	0.5	9.4	0.6
Head LI (narrow)	10.0	0.4	9.3	0.5
Joint (Head+Wife) LI	10.8	0.6	10.3	0.7
Post-Govt Joint LI	9.0	0.4	8.7	0.4

Notes: Table 7 is based on PSID Family Files over the period 1978-2014. The results are based on the simulation exercise described in Subsection 7.2. The top part of the table reports the estimated probabilities of income drops exceeding 50 and 100%, and the bottom part reports the estimated probabilities of income increases exceeding 50 and 100%. The left part of the table reports the corresponding probabilities for aggregate economic expansions, and the right part — for contractions.

defined head’s labor income seem to be changing the least among other income definitions, but their recessionary histograms still have heavy tails.

We next quantify the probabilities of tail events based on our simulation exercise, and report the results in Table 7. Table 7 implies that the probability of tail events (100% increase/decline) increases in recessions (last column) as compared to expansions (second column) for all 5 income types, which is consistent with the observation that income shocks are more dispersed in recessions as compared to expansions.

Moreover, when we compare probabilities of tail events of different signs *within* income definitions, we find that, in expansions, large positive events (+50% and +100%) are more likely than large negative events (-50% and -100%). In recessionary times, however, the pattern is reversed: income “catastrophes” turn out to be more probable than income “miracles”. These facts combined reflect the procyclical nature of skewness of income shocks, which holds across all income types.

To better understand the implications of time-varying volatility and skewness, in Table 8 we report the changes in probabilities of tail events within income definitions and across aggregate economic expansions and contractions. The reported numbers are the ratios of the corresponding probabilities from Table 7. For example, number 35 in the first row reflects a 35% increase in probability of having a 50% wage cut in recessions as compared to expansions. Head’s labor income shows the most pronounced movements in tail events probabilities across income definitions: while 100% wage increase is only 20% more likely in recessions, the probability of wage loss increases fourfold. All income definitions (apart from post-government joint labor income) exhibit strong asymmetries in tail events (a large increase in probability of negative events in recessions as opposed to large positive events). We note that taxes and transfers reduce the odds of tail events by a factor of 2 to 3 as

TABLE 8: CHANGE IN TAIL EVENTS PROBABILITIES (RECESSIONS VS. EXPANSIONS), %

Type of Income	-50%	-100%	+50%	+100%
Head Hourly Wage	35	100	-4	0
Head LI	80	400	-10	20
Head LI (narrow)	44	130	-7	25
Joint (Head+Wife) LI	50	150	-5	17
Post-Govt Joint LI	36	67	-3	0

Notes: Table 8 reports the change in probability (in %) of tail events in recessions as compared to expansions. The table is based on the estimated probabilities in Table 7.

compared to other income types.

8 Conclusion

In this paper, we systematically analyzed the volatility and skewness of income shocks over the business cycle. We first allowed for only a time-varying risk (volatility) of income shocks, and found that the quantitative results crucially depend on the income definition considered. Head’s labor income turns out to be the most volatile type of income, with the standard deviation of shocks being almost twice larger in recessions than in expansions. Intra-family insurance channel has a limited role, as the joint (head + wife) labor income still exhibits pronounced fluctuations in risk. We also investigated the reason why individual labor income risk is so volatile: we found fluctuations in hourly wage risk to be small, which highlights a large quantitative role of hours employed. Government taxes and transfers are quantitatively important, and mitigate a good portion of the countercyclical risk in household’s income. Finally, we estimate income risk for a subsample of people with a strong labor market attachment; our results indicate that the volatility of risk of that income type is small, which mirrors findings of previous research.

We re-estimated the model on “young” (ages 23-39) and “old” (ages 40-60) subsamples separately, and found a remarkable heterogeneity: while the overall dispersion of income shocks is larger for the old subsample, fluctuations in risk for young workers are stronger.

After quantifying a countercyclical income risk, we augment the analysis to allow for both time-varying volatility and skewness. Technically, we make an identifying assumption that income shocks are drawn from the skew normal distribution, and subsequently extend the approach of [Storesletten et al. \(2004\)](#) by including a set of moment conditions for skewness which allow us to estimate an extended set of parameters. After showing that the implications of countercyclical volatility carry over to the case where we allow for both time-varying moments, we quantify the probabilities of “tail event” — large positive and negative changes in income. Our results suggest that income shocks exhibit a pronounced asymmetry, making large income drops more likely than large increases in recessions. Individual labor income shows the strongest asymmetry: in recessions, a large increase is only 20% more likely (as compared to expansions), while the odds of a large drop go up by 400%. Moreover, we find that taxes and transfers reduce the probability of tail events by a factor of 2 to 3 as compared to other income definitions.

More work is needed to fully understand the origins and consequences of income risk. We see at least 3 directions for future research. First, more empirical work is required to estimate income risk in a systematic manner: economists typically face a trade-off where, on the one hand, data is comprehensive but with few variables (e.g., Social Security data), and, on the other hand, contains lots of information but for a small sample of population (survey data). Obtaining a large dataset with more observables will make the estimation more precise, and, therefore, make it feasible to estimate the income process for narrow “slices” of population - by age groups, race, gender, etc.

Another fruitful direction for future research can be focused on quantifying the impact of countercyclical risk on macro aggregates. In this paper, we did not pursue this line, since our intention was to systematically analyze several income definitions within a unified estimation framework. A quantitative macro model can be the focus of an independent study.

Finally, an effort should be put forward to explore the opposite direction of causality: how can a macro shock generate realistic fluctuations in income risk? A promising exercise would be to “unpack” empirically realistic time-varying risk into several economic forces, instead of assuming an *a priori* counter-cyclical risk and procyclical skewness. We leave these issues to future research.

References

- ABOWD, J., CARD, D. (1989) “On the Covariance Structure of Earnings and Hours Changes” *Econometrica*, 57(2), 411-445.
- ARELLANO, C., BAI, Y., KEHOE, P. (2012) “Financial Markets and Fluctuations in Volatility”, manuscript.
- BLOOM, N. (2009) “The Impact of Uncertainty Shocks” *Econometrica*, 77(3), 623-685.
- BLOOM, N., FLOETOTTO, M., JAIMOVICH, N., SAPORTA-EKSTEN, I., TERRY, S. (2018) “Really Uncertain Business Cycles” *Econometrica*, 86(3), 1031-1065.
- BLUNDELL, R., PISTAFERRI, L., PRESTON, I. (2008) “Consumption Inequality and Partial Insurance” *American Economic Review*, 98(5), 1887-1921.
- BUSCH, C., DOMEIJ, D., GUVENEN, F., MADERA, R. (2018) “Asymmetric Business Cycle Risk and Social Insurance”, manuscript.
- BUSCH, C., LUDWIG, A. (2016) “Labor Income Risk in Germany Over the Business Cycle”, manuscript.
- CONSTANTINIDES, G., DUFFIE, D. (1996) “Asset Pricing with Heterogeneous Consumers” *Journal of Political Economy*, 104(2), 219-240.
- DEATON, A. (1992) “Understanding Consumption”, Clarendon Press.
- ELSBY, M., HOBIJN, B., ŞAHIN, A. (2010) “The Labor Market in the Great Recession”, *Brookings Panel on Economic Activity*.
- FERNANDEZ-VILLAYERDE, J., GUERRON, P., RUBIO-RAMIREZ, J., URIBE, M. (2011) “Risk Matters: The Real Effects of Volatility Shocks” *American Economic Review*, 101, 2530-2561.
- GUVENEN, F., KARAHAN, F., OZKAN, S., SONG, J. (2016) “What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Dynamics?”, manuscript.
- GUVENEN, F., OZKAN, S., SONG, J. (2014) “The Nature of Countercyclical Income Risk” *Journal of Political Economy*, 122(3), 621-660.
- HEATHCOTE, J., STORESLETTEN, K., VIOLANTE, G. (2010) “The Macroeconomic Implications of Rising Wage Inequality in the United States” *Journal of Political Economy*, 118(4), 681-722.
- HEATHCOTE, J., PERRI, F., VIOLANTE, G. (2010) “Unequal We Stand: An Empirical Analysis of Economic Inequality In the United States 1967-2006” *Review of Economic Dynamics*, 13, 15-51.
- KAPLAN, G., VIOLANTE, G. (2010) “How Much Consumption Insurance beyond Self-Insurance?” *American Economic Journal: Macroeconomics*, 2(4), 53-87.

- LEE, Y., MUKOYAMA, T. (2015) “Entry and Exit of Manufacturing Plants over the Business Cycle” *European Economic Review*, 77, 20-27.
- LOW, H., MEGHIR, C., PISTAFERRI, L. (2010) “Wage Risk and Employment Risk over the Life Cycle” *American Economic Review*, 100(4), 1432-1467.
- MEGHIR, C., PISTAFERRI, L. (2004) “Income Variance Dynamics and Heterogeneity” *Econometrica*, 72(1), 1-32.
- MOREIRA, S. (2016) “Firm Dynamics, Persistent Effects of Entry Conditions, and Business Cycles”, manuscript.
- SALGADO, S., GUVENEN, F., BLOOM, N. (2014) “Skewed Business Cycles”, manuscript.
- STORESLETTEN, K., TELMER, C., YARON, A. (2001) “The Welfare Cost of Business Cycles Revisited: Finite Lives and Cyclical Variation in Idiosyncratic Risk” *European Economic Review*, 45(7), 1311-1339.
- STORESLETTEN, K., TELMER, C., YARON, A. (2004) “Cyclical Dynamics in Idiosyncratic Labor Market Risk” *Journal of Political Economy*, 112(3), 695-717.
- STORESLETTEN, K., TELMER, C., YARON, A. (2007) “Asset pricing with idiosyncratic risk and overlapping generations” *Review of Economic Dynamics*, 10(4), 519-548.