Comments on "Dissecting Idiosyncratic Income Risk" by Elin Halvorsen, Hans Holter, Serdarn Ozkan, and Kjetil Storesletten

Manuel Arellano
CEMFI

2nd Annual Research Conference of Banco de España
Taxes and Transfers

Madrid, September 4, 2018

Summary

What they do

- The goal is to document new facts about individual income risk from the combined data of several registers of the entire Norwegian population over more than 20 years.
- Four registers combined: taxes, social security, census (to group couples into households), and employer/employee contracts (to observe hours of work).
- The information on hours is not perfect: measurement error, missing values, absence of overtime hours, and unavailable for the first 10 years.
- They study the distributions of income growth over the life-cycle, across income groups, and by gender.
- Their data allows them to go beyond individual earnings:
  - They decompose earnings growth into (i) changes in wages and (ii) changes in hours.
  - They also look at household labor income, total income, and disposable income.

What they find

- The distribution of Norwegian earnings changes is left-skewed and highly kurtic as in the US. The main difference is that the variance of the shocks is larger in the US.
- Large earnings shocks are associated with larger movements in wages than in hours.
- Hour changes matter more to the left skewness of earnings growth than wage changes.
- Both wage and hour changes are highly kurtic.
- Moving in or out of long-term sickness, moving in or out of part-time employment, and job changes are the most important events associated with large earnings shocks.
- The negative skewness of earnings shocks is driven by job-stayers.
- Labor income of spouses and household disposable income contribute towards removing the negative skewness of individual earnings changes.

Facts in progress

- This is high quality work in progress using a fascinating range of data.
- The goal is to exploit the strengths of the data relative to other data sets used in ongoing or recent work to break new empirical ground in the most transparent way.
- They uncover many interesting results that call for further examination.

Related work

- An active area of research and a race to access better and better sources of data.
    - Hoffmann and Malacrino (2018) use Italian data to decompose earnings into hours and wages focusing on business cycle properties.
    - Busch, Domeij, Guvenen, and Madera (2016) use German, Swedish, and US data to learn to what extent family and government mitigate individual income risk in recessions.

Comments on the empirical methodology

- Comments intended to draw connections with recent work on income processes.
- Blending nonparametrics with the more principled objects of analysis that were targeted in traditional parametric linear approaches.
- There are strong complementarities between the exploratory research carried out in this paper and the estimation of more elaborate models of idiosyncratic income risk.
- Models of idiosyncratic income risk are a key input to macroeconomic modelling.

Marginal distribution of growth rates vs transmission of shocks

- The empirical literature in the wake of Guvenen et al (2014) looks at the distributions of earnings growth or wage growth for individuals or households.
- After some contritions, earnings changes are described as measures of earnings risk.
- Leaving aside issues of superior or inferior information, we tend to associate risk with shocks and their transmission, which informs us about their persistence.
- Predictive distributions of variables whose future outcomes are uncertain are the relevant objects in models of choice.
- A natural next step is to estimate those distributions nonparametrically while taking care of individual heterogeneity.

Nonlinear persistence and (conditional) skewness: two sides of the same coin

- We can think about skewness in terms of nonlinear transmission of shocks.
- Arellano, Blundell, and Bonhomme 2017 (ABB) find positive skewness for the income poor and negative skewness for the income rich in the distribution of PSID earnings given past earnings.
- This is isomorphic to the nonlinear persistence that they uncover.
- When high-y households are hit by an unusually negative shock, dependence of y_{it} on $y_{i,t-1}$ is low, with the result that they have a relatively large probability of outcomes far to the left from the central part of the distribution.
- Likewise, low-y households have a relatively large probability of getting outcomes far to the right of their distribution associated with low persistence episodes.
- Advances in computer technology not only facilitate the analysis of big administrative data sets but also the application of computer-intensive nonparametric methods to small survey data sets.

Skewness of log-earnings residuals growth

- As measures of fit of their model, ABB report estimates of the skewness, kurtosis, and densities of log-earnings residuals growth at various horizons.
- The data shows that log-earnings growth is non-Gaussian, displaying negative skewness and high kurtosis, similar to the results in Guvenen et al. (2015) for US administrative data and the current paper for Norwegian data.
- This shows both the similarity between the PSID and register data in terms of higher moments of earnings growth, and the ability of the ABB model to fit these features.

Skewness of log-earnings residuals growth (continued)

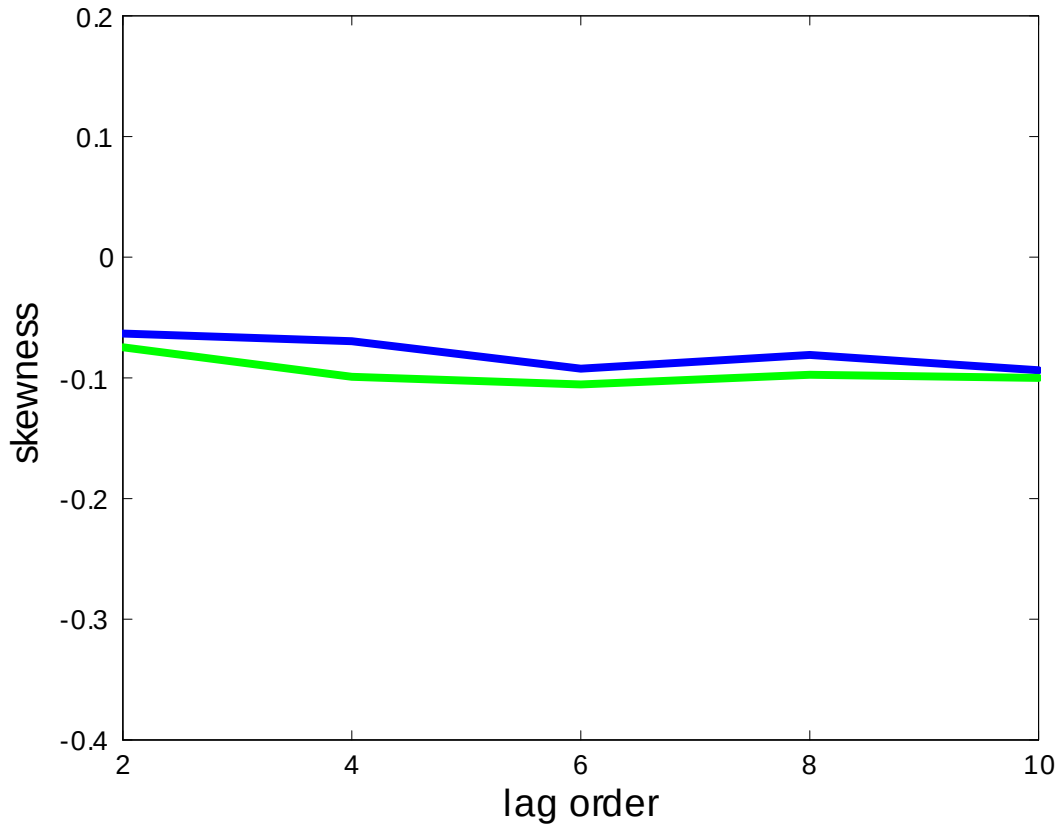

Note: Bowley/Kelley measure of unconditional skewness of $y_{it} - y_{i,t-s}$ for $\tau = 11/12$, at various horizons s, from 2 to 10 years (reported on the x-axis). Blue is PSID data, red is nonlinear model. Source: ABB 2017, Econometrica Supplement, Figure S6, p. 12.

ABB results on Norwegian data

- To corroborate their PSID findings on a larger data set, ABB estimated their income process using a sample of 2873 households from the 2000–2005 Norwegian administrative data.
- The results are for the disposable income of married males aged 30–60 and for household disposable income.
- The Norwegian population register data present a similar pattern of conditional skewness to the PSID and similar nonlinear persistence in the persistent component.
- The Norwegian results in ABB were provided as part of the Blundell, Graber, and Mogstad project on "Labour Income Dynamics and the Insurance from Taxes, Transfers and the Family' (Blundell, Graber, and Mogstad 2015).

Conditional skewness of log-earnings residuals and persistent component

Log earnings residuals y_{it}

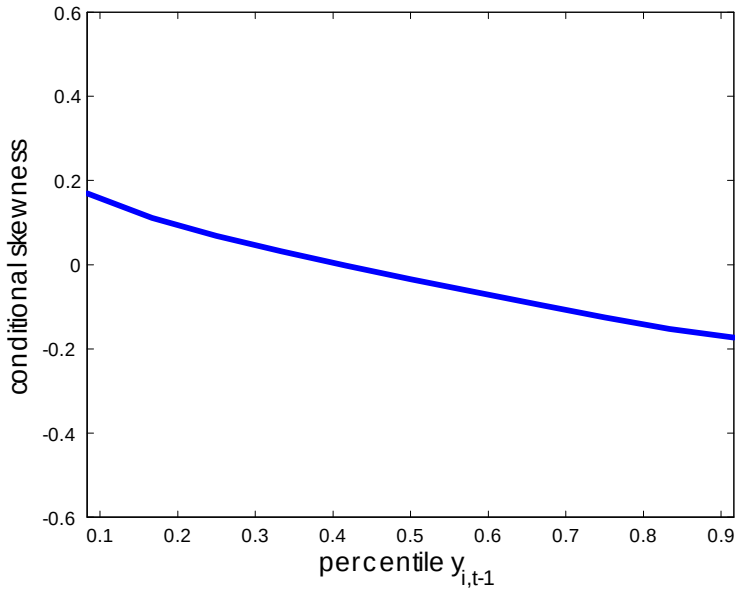

Persistent component η_{it}

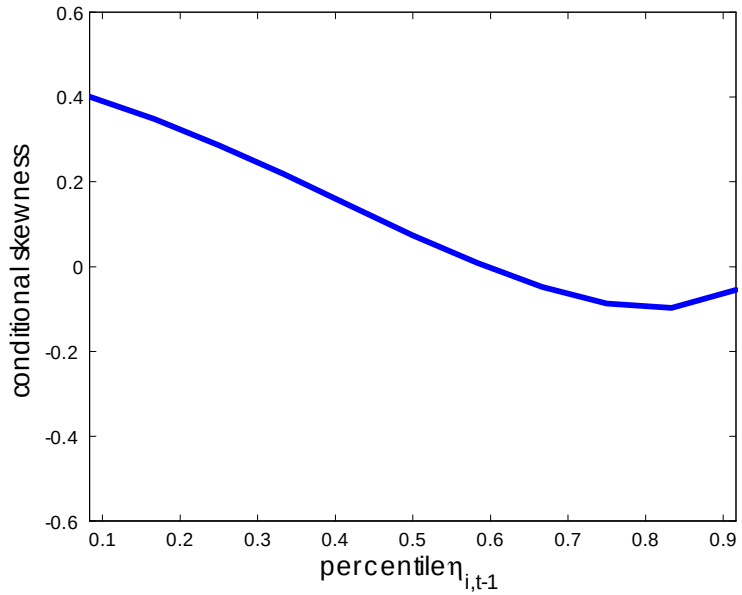

Note: Bowley/Kelley conditional skewness of log-earnings residuals (data, left) and η component (right) for $\tau = 11/12$. Random subsample of 2,873 households, from 2000 – 2005 Norwegian administrative data, non-immigrant residents, age 30 to 60.
Source: ABB 2017, Econometrica Supplement, Figure S18, p. 21.

Conditional skewness in Norwegian administrative data 2005-2006

Family income

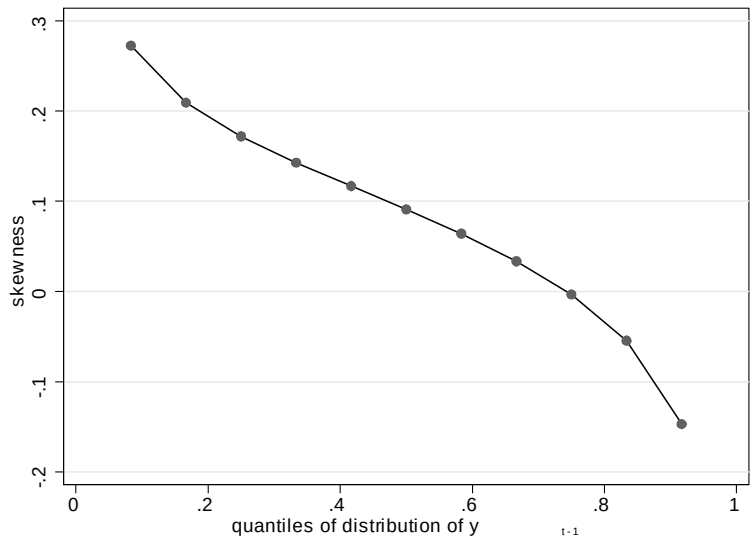

Individual income

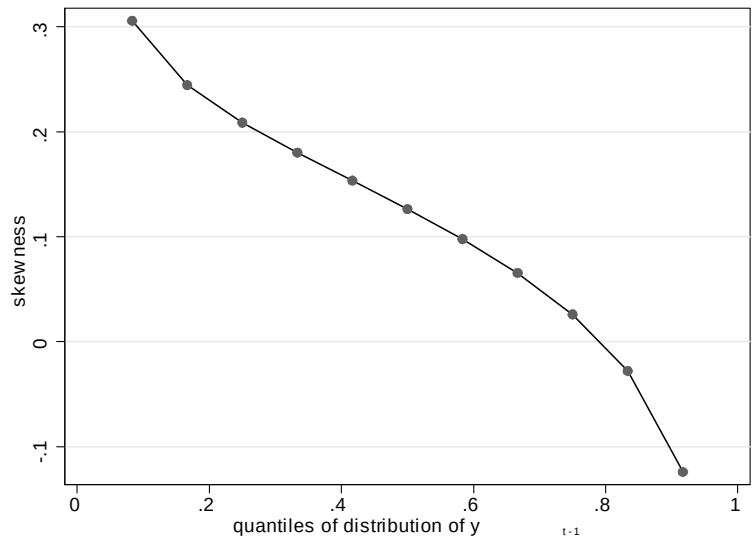

Note: Bowley/Kelley conditional skewness of log-earnings for $\tau = 1/10$. Age 25-60, years 2005-2006. Source: ABB 2017, Econometrica Supplement, Figure S16, p. 19.

Average recent earnings as a control for heterogeneity

- The authors compute an age-specific measure of average recent earnings based on the previous 5 years.
- Moments calculated at each t are conditioned on recent average earnings.
- The motivation for conditioning on previous average earnings and age is to deal with "ex ante identical individuals".
- This is a simple, easily describable, way of controlling for heterogeneity when documenting earnings growth, which has been used in several papers.
- The standard downside is that conditioning on average recent earnings may confound persistence with fixed effects.

Persistent/transitory decompositions in the era of big administrative data

- One- and five-year log-earnings changes are calculated as measures of "transitory earnings shocks" and "persistent earnings shocks", similarly to Guvenen et al 2014.
- The motivation is that in a permanent/transitory income process, the variance of a short-term change puts a bigger weight on transitory shock variances, whereas in a longer-term difference the variance tends to reflect more of the permanent shocks.
- The literature on Gaussian linear income processes has traditionally relied on persistent/transitory formulations.
- The standard argument for favoring a two-error process has been that it has different implications for the response of consumption to income than an observationally equivalent single-error model, since persistence is key.
- From modern deconvolution theory, we now know that the distributions of persistent and transitory shocks are nonparametrically identified, so that it is possible to investigate departures from normality without resorting to parametric assumptions.
- Moreover, deconvolution estimators should work well in very large data sets such as those employed here.

Persistent/transitory decompositions in the era of big administrative data (cont.)

- Two-error decompositions help establishing the persistence of large non-normal shocks.
- In the ABB model, skewness and excess kurtosis of log-earnings growth at long horizons are mostly due to the non-Gaussianity of the transitory component.
- Incidentally, an easily interpretable quantile-based measure of kurtosis is, for some $\alpha < 1 - \tau$,

$$kur_t(\tau, \alpha) = \frac{Q_t(1-\alpha) - Q_t(\alpha)}{Q_t(\tau) - Q_t(1-\tau)}.$$

Transitory earnings in Norway

- ABB find that the dispersion of the transitory component is much smaller in the Norwegian data, suggesting either the presence of large measurement error in the PSID or smaller true transitory innovations in Norway.

What are large earnings shocks?

- Why not ask what life events predict large earnings changes? For example, calculate $\Pr\left(\Delta y < -0.5 \mid \text{long term sickness}\right)$ instead of $\Pr\left(\text{long term sickness} \mid \Delta y < -0.5\right)$.
- They could be very different:

$$\frac{\Pr\left(\Delta y < -0.5 \mid \text{long term sickness}\right)}{\Pr\left(\text{long term sickness} \mid \Delta y < -0.5\right)} = \frac{\Pr\left(\Delta y < -0.5\right)}{\Pr\left(\text{long term sickness}\right)}$$

- This formulation makes it easier to connect with the thinking in the quasi-experimental literature.
- It may be interesting to exploit the long panel dimension to quantify the persistence of particular observed shocks, such as parental leave and long term sickness.

Change of employer

- Among the job switchers, we want to distinguish between quits and layoffs. That is, between job moves and job losses.
- This may help understand the seeming paradox that negative skewness of earnings shocks is driven by job-stayers.
- Collective dismissals may be identifiable from the universe of social security records.
- One would expect the probability of having more large-negative earnings changes than large-positive earnings changes to be larger for layoffs than for stayers, and larger for stayers than for quitters.

Uncoupling wage rates and hours of work

Hours of work data

- To what extent variation in effective work hours as opposed to contractual hours can be expected to be less well captured for high earners than for low earners, and for different occupations?

Co-skewness

- Proposition 3: The mean of $x + y$ enters the formula. Are not variables demeaned? To what extent the role of co-skewness is affected by a positive covariance between hours and wages?

- Hoffmann and Malacrino (2018) use a similar decomposition on Italian data. They do not find a large role for the cross term. In their case most of the skewness comes from hours of work.

Hours and skewness: mixed results in different datasets

- Hoffmann and Malacrino (2018) find a dominant role for employment time in generating the procyclical skewness of earnings growth. They use third moments.
- In contrast, Busch et al (2016) find that skewness in wage growth is responsible for skewness of annual earnings growth in German data. They use Kelley's skewness.
- These discrepancies suggest considering quantile-based Bowley/Kelley measures of skewness at different percentiles to find out what kind of skewness matters.

- Complementary good survey data on hours of work may help.
- ABB 2018 provide some PSID evidence of skewness in hourly wages (below).

Nonlinear persistence in PSID hourly wages

- ABB 2018 provide preliminary evidence that the nonlinear persistence ABB 2017 uncovered in family earnings data is also evident in hourly wage data.
- This suggests that nonlinear persistence may be a key feature of life-cycle labor supply.

Wages, PSID

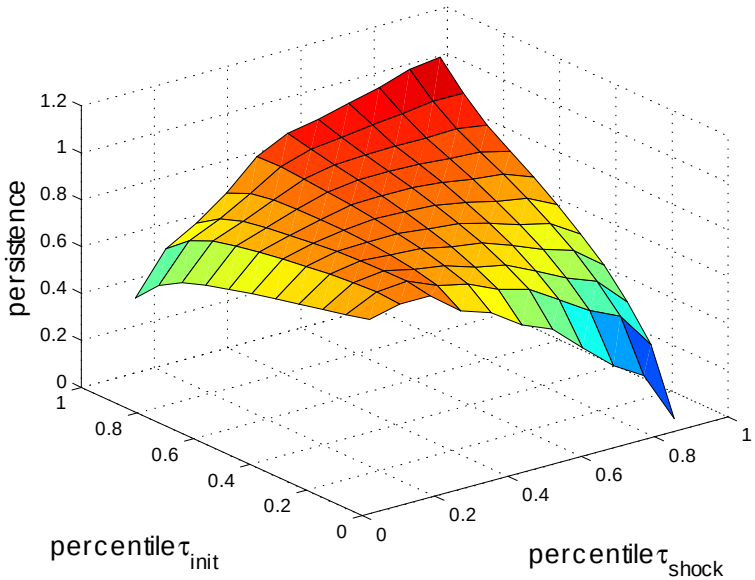

Persistent component

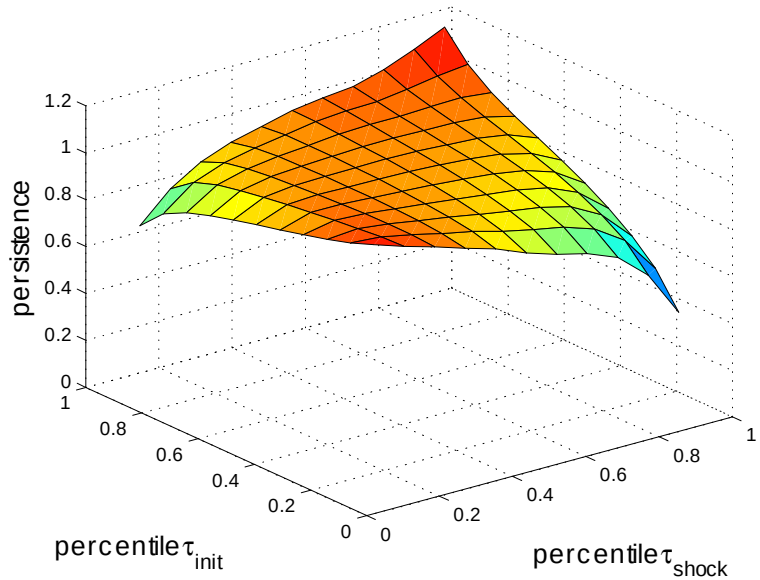

Note: Log male hourly wages, Age 30–60 PSID 1999–2009. Estimates of the average derivative of the conditional quantile function.

Source: Arellano, Blundell, and Bonhomme (2018): "Nonlinear Persistence and Partial Insurance: Income and Consumption Dynamics in the PSID," AEA Papers and Proceedings, 108: 281–286.

Couples, business and capital income, taxes and transfers

- The paper contains a wealth of distributional results on alternative income concepts, which speak of the roles of couples and government as providers of insurance. I quite like these results.

- To what extent the large reduction in income variance from being a couple is the result of conditioning on two-earner couples, if that is the case?

- More generally, we are still short of productive ways of incorporating movements between zero earnings and positive earnings in documenting earnings risk.

Concluding remarks

- This paper has the ingredients to make an important contribution to an exciting research agenda.

- Congratulations to the authors on their excellent work!