

Referee report

“Life-Cycle Profiles of Income and Consumption Inequality”, by Guozhong Zhu

Review of Economics and Statistics

This paper revisits an age-old identification problem between age, time and cohort effects. The author notes that this identification problem seems to be important in practice for the literature on income and consumption inequality. The age profile in income inequality controlling for time but not cohort effects is much flatter than controlling for cohort but not time effects. The age profile in consumption inequality on the other hand, is very similar for both estimates. This leads to substantially different conclusions depending on whether the researcher believes cohort or time effects are more important.

The main contribution of the paper is to argue that:

1. There is evidence for age-specificity in the time effects in income inequality. In other words: the evolution of income inequality over time is different for different age groups.
2. Controlling for age-specific time effects, the estimates of the age profiles in income and consumption inequality are “less biased than the traditional ones” (p.4).

This argument implies that it is important to allow for age-specific time effects in applied work on income and consumption inequality for two reasons. First, these patterns in the data may be interesting in themselves. Second, in this more general model, the always-present identification problem between age, time and cohort effects is mitigated, which allows to draw more robust conclusions.

First of all, I want to compliment the author on the way the paper is written. The paper is well motivated, very clear and generally a pleasure to read. I also am in complete agreement with the author that this is an interesting problem, which is likely to be important for many applications, including the literature on consumption and income inequality.

The main result of the paper, however, seems too good to be true. The suggestion is that by estimating a more flexible specification, we can -if not solve- mitigate an identification problem, without adding data or imposing restrictions. Having thought about this, I convinced myself that the result is in fact not true. My argument goes in two steps. First, I claim that age-specific time effects are the same as time-specific age profiles. Second, the slope of ‘the’ age profile, controlling for time-specific age profiles, depends on what is the reference year. It seems that here the reference year is (implicitly) chosen such that the age profile in income inequality in that year is steeper than the average, bringing it closer to the profile we would get controlling for cohort rather than time effects.

To make this argument precise, assume that all effects are linear. Then, equation (1) for income inequality x_{ht} simplifies to

$$\begin{aligned}x_{ht} &= \alpha + \psi h + \beta t + \gamma t * h [+ \delta (t - h)] + \varepsilon_{ht} \\ &= \alpha + \psi h + (\beta + \gamma h) t + [+ \delta (t - h)] + \varepsilon_{ht}\end{aligned}$$

where h is age, t is time (both in years), so that $t - h$ is birth year, ψ is the slope of the age profile, β of the time effects, and δ of the cohort effects, which are dropped from the regression. The variation in time effects across age groups is captured by the coefficient γ

on the interaction between time and age, which is made explicit on the second line. But of course we could also have rewritten the expression as,

$$x_{ht} = \alpha + (\psi + \gamma t) h + \beta t + [\delta(t - h)] + \varepsilon_{ht}$$

making it immediately clear that age variation in the time effects and time variation in the age profiles are not separately identified. With marginal modifications, this argument goes through when we allow for non-linear time and age effects as in the paper.

If we do not allow for time variation in the age effects, then ψ captures the average age profile. Controlling for time variation in these profiles, ψ captures the age profile in the reference year. Since there must be values both above and below the mean, the age profile controlling for time-specific age effects can be steeper or flatter than the unconditional age profile, depending on which reference year we choose.

I do not dispute that there is time variation in age profiles of income inequality (or variation across age groups in the evolution of income inequality over time). In fact, I think this fact is potentially interesting, and I encourage the authors to rewrite the paper to focus on this finding and to think about potential explanations. However, I am not convinced by the argument in the current version of the paper, which focuses a lot on the fact that this more general model mitigates the identification problem between cohort time and age.