Discussion of "Heterogeneous Life-Cycle Profiles, Income Risk and Consumption Inequality" by Giorgio Primiceri and Thijs van Rens

Jonathan Heathcote*

Federal Reserve Bank of Minneapolis
CEPR

October 17, 2008

This paper uses micro survey data on income and expenditure from the Consumer Expenditure Survey (CEX) to shed light on the nature of income fluctuations and the mapping from income inequality to consumption inequality. The main virtue of the analysis is that the CEX records income and consumption for the same households, so questions of cross-dataset comparability do not arise. Giorgio and Thijs explore both how inequality rises with age over the life-cycle and how cross-sectional inequality has changed over time. A central focus of their analysis is to quantify the fraction of permanent income fluctuations that is predicted by households, and thus does not translate to consumption.

Without question, some income fluctuations over the life-cycle are foreseen by individuals. Thus, as the authors recognize, the interesting questions are quantitative: (i) What fraction of fluctuations is foreseen? and (ii) How far in advance are fluctuations foreseen? A growing literature follows a variety of strategies to answer these questions, including Keane and Wolpin (1997), Jappelli and Pistaferri (2000), Cunha et al. (2005), Guvenen (2007), and Huggett et al. (2007).

Following the emphasis of the paper, I will focus my discussion on changes in predictability over time, rather than on how predictability impacts consumption inequality over the life-cycle.

Giorgio and Thijs argue that there has been a large increase over time in the variance of forecastable permanent innovations to household income. This conclusion is inescapable if one accepts the following three premises:

(i) The increase in the variance of log household income between 1980 and 2000 was very large, while the variance of log household consumption was relatively stable.
(ii) All of the increase in income inequality was permanent in nature.
(iii) Permanent shocks are largely uninsurable, so that an increase in their variance must generate an increase in the variance of consumption of similar magnitude.

I will critically examine each of these premises in turn. I will first make a case for an alternative characterization of the facts suggesting a much smaller increase in income inequality. Then I will argue that some of this smaller increase in income inequality was transitory in nature. Finally, I will question the extent to which permanent shocks to household income pass through to household consumption.

* heathcote@minneapolisfed.org The opinions expressed here are those of the author and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System

Trends in income inequality: Figure 3 in the paper indicates that over the sample period, the variance of cross-sectional household income in the authors' CEX sample rose by over 20 log points, while the variance of household consumption rose by only 2 log points. What accounts for this enormous increase in income inequality?

It is well known that between 1980 and 2000, there was a steady trend toward greater concentration at the top of the income distribution. However, trends in the variance of log income are driven primarily by what is happening at the bottom of the distribution. Thus, Giorgio and Thijs' analysis is mostly about what is happening at the bottom of the income and consumption distributions, rather than about the phenomenon of the rich getting richer.

The fact that the variance-of-log inequality metric is extremely sensitive to low income observations means that sampling and measurement error can have large effects on the estimated time-path for inequality, especially if the sample size or data collection and processing methods vary over time. When the authors trim the bottom 0.5% of income observations from their sample in each year, the increase in the sample variance of household income becomes smaller: 13 rather than 20 log points. Trimming 5% of observations reduces the increase further, to only 6 log points.[1] Trimming also largely eliminates the mid-1980s spike in inequality that is a central feature of the untrimmed data. Thus, the CEX picture for income inequality is quite sensitive to the inclusion or exclusion of low outliers—outliers that may or may not reflect accurate income reports.

In light of this, it is worth asking what other micro datasets say about income inequality over this period. Heathcote et al. (2008a) find that the variance of log disposable household income (per adult-equivalent, trimming the lowest 0.5% of observations) increased by 8 log points in both the March Current Population Survey (CPS) and the Michigan Panel Study of Income Dynamics (PSID) over the 1980–2000 period. The majority of this increase occurs in the two years between 1980 and 1982. At the bottom of the income distribution, inequality in hourly male wages (as measured by the 50^{th}–10^{th} percentile ratio in the March CPS) rose rapidly in the 1980s and then stabilized. However, after the 1980–82 recession, hours at the bottom of the earnings distribution rose even more rapidly than relative wages fell. Thus, the variance of log household earnings was essentially flat—in both the March CPS and the PSID—between 1982 and 2000. Why, then, did the variance of log disposable household income increase, albeit moderately, over this period? One part of the explanation is that the extent of social insurance provided by taxes and transfers was moderating as the unemployment rate steadily declined.

I conclude that the evidence presented in the paper in favor of a rise in the cross-sectional variance of log disposable household income an order of magnitude larger than the corresponding increase for consumption is not conclusive. If, in fact, the true increase in income inequality was smaller, this would imply a smaller increase in the variance of either insurable or predictable income fluctuations.

Transitory versus persistent shocks: Next I shall discuss very briefly the authors' finding that all of the increase in income inequality over the period appears to be permanent in nature. My only comment on this point is that the Consumer Expenditure Survey only asks the same

[1] I thank the authors for computing these numbers.

household about income twice, and these questions refer to partially overlapping time periods. This limited panel dimension means that the authors can only estimate a simple model for income with purely permanent and purely transitory components. Moreover, the auto-covariance of income is the only moment in the data that can be used to estimate the relative importance of these two pieces. Most work on estimating income dynamics has used the PSID. Because the PSID repeatedly surveys families over many consecutive years, it contains much more information about the persistence of income changes. This work, beginning with Gottschalk and Moffitt (1994) has typically attributed around one-third of the increase in income inequality in recent decades to an increase in the variance of transitory shocks. To what extent does the authors' "entirely permanent" characterization represent a real difference between the CEX and PSID income data, and to what extent does it reflect the fact that the authors fit a different set of moments? This question could be answered by applying the estimation procedure used for the CEX directly to PSID data.

If, as the PSID evidence suggests, around one-third of the increase in income inequality over the sample period was truly transitory and thus insurable in nature, then the implied increase in the variance of persistent / permanent shocks that did not translate to consumption inequality would be further reduced.

Permanent income shocks and consumption: The problem with the authors' third premise is that even highly persistent or permanent shocks to earnings need not translate one-for-one into consumption. I will outline three reasons why: (i) the existence of insurance mechanisms beyond risk-free borrowing and lending, (ii) the possibility that some of what the authors identify as permanent shocks are really persistent but mean-reverting, and (iii) the fact that in a life-cycle context, permanent income responds less than one-for-one to "permanent" shocks to earnings.

First, an important contribution of the paper is to show that, within the environment considered, predictability and insurability are observationally equivalent. Thus, it could be that permanent changes to income do not completely pass through to consumption because a range of insurance mechanisms act as a buffer, rather than because the income changes were foreseen. Giorgio and Thijs note that potentially important forms of insurance against shocks to earnings include offsetting changes in asset income, taxes and public transfers. They argue that the extent of this insurance must be relatively modest in practice, because there is no significant difference between the average response of consumption to a change in earnings versus the response of consumption to a change in disposable income.

In contrast to the authors' findings, however, Blundell et al. (forthcoming) find that taxes and transfers do play an important role in insuring permanent shocks to earnings. Why the difference? One reason why it is hard to directly measure insurance provided through financial markets or the government is that survey respondents drastically underreport most forms of income other than labor earnings (see Heathcote et al. 2008a). Moreover, Michael (1995, p. 403) suggests that income underreporting at the bottom end of the distribution is particularly problematic in the CEX, relative to the March CPS or the PSID. In addition, the CEX contains self-reported taxes, while the CPS and PSID use an imputation procedure.

If the menu of insurance instruments is restricted to risk-free savings, agents can still partially insure shocks that are persistent but not fully permanent. Even with the longer panel dimension of the PSID it is difficult to distinguish statistically between shocks that are permanent and

autoregressive shocks that are quite persistent. From an economic standpoint, however, the precise persistence matters. To illustrate this point, consider an individual who starts work at 20, retires at 60, and collects a pension equal to half average lifetime earnings until death at 80. A positive $1.00 innovation to earnings at age 20 increases average yearly lifetime income by $0.83 if this shock is permanent (i.e. if it raises expected earnings by $1.00 for each year of working life). Thus the pass-through coefficient from shocks to income to consumption (i.e. the marginal propensity to consume) is 0.83 for 20 year-olds according to the simplest version of the permanent-income life-cycle hypothesis.[2] However, if the initial shock decays with an annual autoregressive coefficient of 0.95, average yearly income will increase by only $0.36. Thus, permanent income consumers should respond with a $0.36 increase in consumption.

Finally, in a realistic life-cycle environment, there is really no such thing as a fully permanent shock to earnings. In particular, as one approaches retirement age, fluctuations in earnings have an ever smaller effect on permanent income. Continuing the example outlined above, a "permanent" $1.00 increase in earnings at age 50 increases average income over the remainder of the life-cycle by only $0.42. Kaplan and Violante (2008) compute aggregate pass-through coefficients for a range of calibrated life-cycle models when shocks are alternatively permanent or highly persistent.

To recapitulate, the authors assume that permanent shocks translate one-for-one into consumption. However, within a fully articulated life-cycle model, pass-through will be smaller, and smaller still if shocks are autoregressive, or if households can access alternative sources of insurance. Thus, one does not need to appeal to foresight to account for imperfect pass-through.

Putting the three pieces of this discussion together, I conclude that the authors' CEX sample likely exaggerates the true increase in income inequality over the period, and the fraction of this increase that reflected permanent income changes at the household level. At the same time, the authors' model for consumption likely exaggerates the predicted response of consumption to permanent shocks. So what really happened to the variance of predictable income fluctuations? To conclusively answer this question requires future refinements in the use of micro-data, and in the estimation of structural models that can differentiate between foresight and insurance. In the meantime, my own prior, based on recent work with Kjetil Storesletten and Gianluca Violante (Heathcote et al. 2008b), is that forecastability has not improved. We find that a life-cycle model with a variety of explicit insurance mechanisms (savings, labor supply, reallocation with the household, a public pension system) broadly replicates the observed time paths for both income and consumption inequality, even though all idiosyncratic wage changes arrive as unforeseen shocks.

Blundell, R., Pistaferri, L., Preston, I.. Forthcoming. Consumption inequality and partial insurance. American Economic Review.

Cunha, F., Heckman, J., Navarro, S., 2005. Separating uncertainty from heterogeneity in life cycle earnings. Oxford Economic Papers 57(2), 191–261.

[2] In particular, assuming quadratic utility, with $\beta=1$ and $r=0$.

Gottschalk, P., Moffitt, R.A., 1994. The growth of earnings instability in the U.S. labor market. Brookings Papers on Economic Activity 25(1994-2), 217–72.

Guvenen, F., 2007. Learning your earning: Are labor income shocks really very persistent? American Economic Review 97(3), 687–712.

Heathcote, J., Perri, F., Violante, G.L., 2008a. Cross-sectional facts for macroeconomists: United States (1967–2006). Mimeo, New York University.

Heathcote, J., Storesletten, K., Violante, G.L., 2008b. The macroeconomic implications of rising wage inequality in the United States. Mimeo, New York University.

Huggett, M., Ventura, G., Yaron, A., 2007. Sources of lifetime inequality. Mimeo, Georgetown University.

Jappelli, T., Pistaferri, L., 2000. Using subjective income expectations to test for excess sensitivity of consumption to predicted income growth. European Economic Review 44(2), 337–58.

Kaplan, G., Violante, G.L., 2008. How much insurance in Bewley models? Mimeo, New York University.

Keane, M.P., Wolpin, K.I., 1997. The career decisions of young men. Journal of Political Economy 105(3), 473–522.

Michael, R.T., 1995. Measuring poverty: A new approach. Panel on Poverty and Family Assistance, National Academy Press, Washington DC.