

Life-Cycle Income and Consumption Variability

by Peter Rupert and Chris Telmer

Economic inequality is a major concern in policy circles, and it occupies a large field in the economics profession. Much of the literature has focused on income and wealth inequality; its main finding is that inequality has risen over time. Edward Wolff, for example, reports that in 1974 the richest 5 percent of American families earned 14.8 percent of total U.S. income, whereas by 1998, their share had risen to 20.7 percent.¹ When both income and wealth are taken into account, the growth in inequality becomes worse.

Only recently have researchers taken the logical next step, looking at the behavior of *consumption* inequality. Presumably, income and wealth inequality are so interesting because they affect consumption inequality and, as a result, inequality in economic welfare. It is important for policymakers to understand how and why income inequality affects consumption—there is little cause for action if income inequality does not manifest itself in consumption. If, on the other hand, it does, knowing how it operates may enable us to better compare the relative benefits of policies such as unemployment insurance and educational subsidies, which are designed to mitigate consumption inequality, but through completely different channels. Unemployment insurance reduces the effect of a loss in income, due to unemployment, on consumption (a process economists call risk-sharing). Educational policy might reduce the income inequality itself.

So where do income and consumption inequality come from? Are the determi-

nants fixed early in life through occupational choices or education? Or do they arise throughout life, in the form of unpredictable shocks to income, like layoffs, as individuals work their way through the labor market?

This *Economic Commentary* explores inequality in income and consumption over the life cycle, focusing on the importance of idiosyncratic shocks to labor market earnings to understand the determinants of consumption inequality. We begin with a striking empirical fact: Inequality in both income and consumption increases substantially with age. This evidence, in conjunction with a standard economic model of life-cycle savings, suggests that an important part of the uncertainty involved in knowing the level of income individuals will obtain throughout their lifetimes is realized during the working years and not before they enter the labor market.

The implications for economic theory and policymaking are important. For example, a popular theory of savings behavior—the “precautionary model”—suggests that individuals save not only for retirement, anticipated expenditures, and so on, but also to build up a “buffer stock” of wealth to help insure themselves against adverse shocks. The theory’s usefulness depends heavily on the extent to which individuals face such shocks in the first place. From a policy perspective, it is important to distinguish between the insurance role and the redistributive role of programs such as Social Security, unemployment insurance, and welfare. The magnitude and distribution of idiosyncratic labor-market risk are critical in making this distinction.

By all accounts, economic inequality is growing—the rich are getting richer, and the poor are getting poorer. This *Economic Commentary* explores inequality in income and consumption and asks whether inequality is determined early in life, before individuals enter the labor market, or whether it manifests itself during the working years.

Risk and Risk Sharing

Many models used by economists share a common feature derived from observing individual behavior: they assume individuals want to smooth consumption over time.² That desire may not be fulfilled, however, because of the uncertain path of future income, health, or other factors that can alter a person’s financial resources.

Some types of uncertainty involve *idiosyncratic risk*, meaning that the events hit only an individual or some subset of the population, rather than everyone at the same time.³ The term *risk* means that individuals are aware that there is some probability of the event occurring.

For example, a fire might destroy one’s home. Replacing it would require a huge spike in expenditure on durable goods such as house and furnishings, and, absent sufficient savings, a substantial reduction in the consumption of other goods.

To minimize the effects of idiosyncratic risk on consumption, financial markets make it possible to pool people together, through homeowners insurance, for example. Each individual pays a little every month to insure against the bad state of the world where the house burns down. Should this occur, funds are taken from the pooled resources and used to rebuild the house, and the consumption of other goods remains relatively constant.⁴

Another familiar example is unemployment insurance. Economic conditions might force a plant closure, leaving many workers out of a job. This can obviously have a severe impact on current earnings, and lifetime earnings would also be lower, the magnitude depending on the extent of the unemployment spell. Here again, individuals can join together and pay a little each month to insure that their consumption profile is not greatly affected by the unemployment shock.⁵ Because it is desirable to smooth consumption over time, individuals can be made better off, in the unlikely event that they experience a bad outcome, if they pool risks together.

The desire to smooth consumption over time means that *deterministic* sources of income inequality over the life cycle will result in a flat(er) age-consumption profile. Suppose, for example, that an individual knows he will take a year off from work ten years from now. It would be possible to save during the working years so that when income is zero during the nonworking year, he could consume the same amount as during the working years.

Life-Cycle Inequality

Figure 1 shows that both earnings inequality and consumption inequality increase over the working years of the life cycle, but only earnings inequality declines at retirement. Note that earnings and consumption inequality are roughly the same for young persons; however, earnings inequality increase at a substantially higher rate as people age. (See the sidebar for a description of the data used in the figures.)

The positive relationship between inequality and age indicates that households receive persistent idiosyncratic earnings shocks throughout their working lives. In addition, the fact that con-

FIGURE 1 CONSUMPTION AND EARNINGS INEQUALITY OVER THE LIFE CYCLE^a

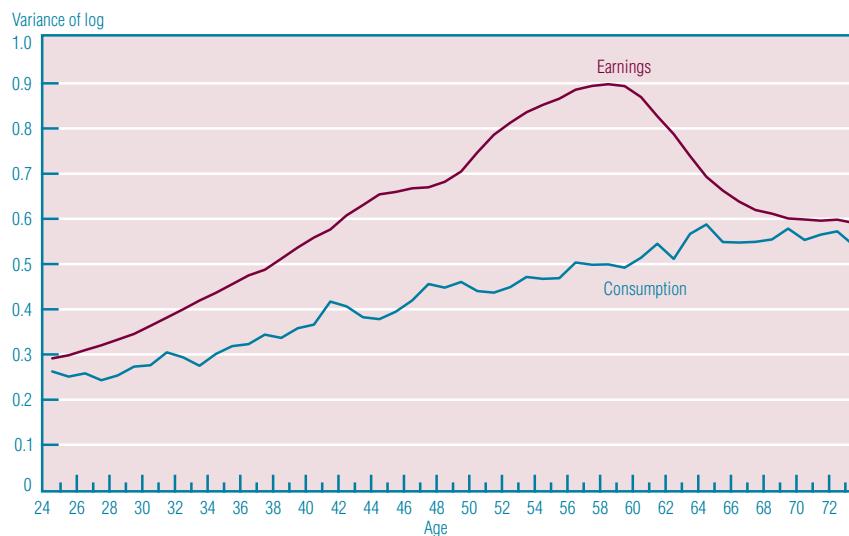



FIGURE 2 EARNINGS INEQUALITY OVER THE LIFE CYCLE, BY EDUCATIONAL ATTAINMENT^b



a. As measured by the cross-sectional variance of consumption and earnings.

b. As measured by the cross-sectional variance of educational attainment.

SOURCES: U.S. Department of Labor, Bureau of Labor Statistics, Consumer Expenditure Survey; and University of Michigan, Institute for Social Research, *Panel Study of Income Dynamics*.

sumption inequality increases alongside earnings inequality suggests incomplete risk sharing, since complete risk sharing would result in a flat consumption profile.

This is not the only interpretation of the data. In terms of earnings inequality, for example, heterogeneity of skills might produce such a pattern. Suppose that wages for highly skilled workers grew faster than those for less skilled workers. Also suppose (although it is contrary to

the interpretation put forward in this *Commentary*) that most of the idiosyncratic risk relates to the acquisition of skills or to the return to acquiring those skills, so that it derives from decisions (or circumstances) made before individuals begin their working careers. In other words, if idiosyncratic risk *within* a skill cohort is relatively unimportant—one would expect to see inequality increasing across the

THE DATA BEHIND FIGURES 1 AND 2

Figures 1 and 2 present data drawn from two sources, the Panel Study of Income Dynamics (PSID) and the Consumer Expenditure Survey (CEX). The PSID, begun in 1968, is arguably the highest quality source of income-related data. It is a panel study, meaning that it tracks individuals over time. However, it is deficient in data on consumption, in that only food consumption is recorded. The construction of the PSID data used in this analysis is well documented in several papers by Storesletten, Telmer, and Yaron.^a

The CEX, on the other hand, contains detailed information on consumption categories, but is inferior to the PSID in terms of income measures. The necessary information on expenditures from the CEX is taken from Deaton and Paxson and covers 1980–1990.^b

Figure 1 presents the data from these two sources and represents the key motivation for the studies listed above and for this *Economic Commentary*. It shows the cross-sectional variance of the logarithm of earnings and consumption. That is, it depicts the variation within an age category, and since logarithms are used, the variability is expressed in percent.

Earnings consist of labor-market earnings plus transfers; transfers consist of such things as unemployment insurance, workers compensation, transfers from nonhousehold family members, and so on. In addition, the measures of dispersion are net of “cohort effects”—that is, if dispersion among young households in the 1982 cohort were uncharacteristically high, it would be identified and removed.^c

a. See Kjetil Storesletten, Chris Telmer, and Amir Yaron, “Asset Pricing with Idiosyncratic Risk and Overlapping Generations,” Carnegie-Mellon University, GSIA Working paper no. 1998-E226, and “The Risk-sharing Implications of Alternative Social Security Arrangements,” *Carnegie-Mellon Rochester Conference Series on Public Policy*, vol. 50, no. 99 (June 1999), pp. 213–99.

b. Angus Deaton and Christina Paxson, “Intertemporal Choice and Inequality,” *Journal of Political Economy*, vol. 102, no. 3 (June 1994), pp. 437–67.

c. The details can be found in Storesletten, Telmer, and Yaron (2000); see footnote 6.

entire population but not across households with similar skill levels.

Figure 2 plots earnings inequality by educational attainment, which is used here as a proxy for skills. Somewhat surprisingly, earnings inequality is roughly the same within broad educational groups, casting some doubt on the assumption that the increase in inequality derives from differences in skill acquisition.

Using the techniques of Kjetil Storesletten, Chris Telmer, and Amir Yaron,⁶ it is possible to compute a simple decomposition of how much earnings uncertainty is resolved early in life (because of skill acquisition, for example) and how much is distributed throughout the life cycle. Their results suggest that roughly 60 percent of lifetime uncertainty is resolved before individuals enter the labor market, while the remaining 40 percent is resolved during the working years.

These results stand in contrast with those of Michael Keane and Kenneth Wolpin, who argue, based on a model of occupational choice, that 90 percent of lifetime uncertainty is resolved before individuals enter the labor market.⁷

Keane and Wolpin’s model addresses the central issue—the extent to which inequality is driven by *decisions* versus *shocks*—in a richer, more explicit manner than the model of Storesletten, Telmer, and Yaron. The latter paper, however, points out that although inequality in labor earnings may be explained by occupational choice, educational attainment, or both, explaining inequality in consumption is likely to present a greater challenge, precisely because of the consumption-smoothing motive outlined above. If an individual knows that—because of her inherent skills or previous decisions—she will enjoy relatively high

earnings in the future, then she will borrow against those earnings in the present, causing future earnings inequality to manifest itself in current consumption inequality. Financial-market frictions, such as borrowing constraints, will mitigate this effect.

■ Inequality, Risk Sharing, and Policy

The questions asked by Keane and Wolpin and Storesletten, Telmer, and Yaron have important, wide-reaching policy implications; therefore, distinguishing between them is crucial. (This is where the short lesson on risk-sharing pays off.) Recall that if the uncertainty is resolved *before* a person enters the labor market, then consumption inequality tends to be flat across the life cycle due to life-cycle smoothing motives.⁸ Such an environment suggests that financial markets have a minimal insurance role because there would be no reason (or only a small one) for precautionary savings to protect consumption against shocks to income over the life cycle. Further, it suggests that policies to combat inequality should focus almost exclusively on schoolchildren.

But consumption inequality does increase over the life cycle, suggesting that risk sharing is incomplete. To generate such consumption inequality, it is necessary to have persistent, idiosyncratic, life-cycle shocks.

Policies aimed at combating inequality must recognize the importance of both sources of it. Some inequality can be addressed by focusing on schooling and, thereby, on occupational choice. Other policies, however, must focus on understanding why the insurance provided through financial markets and precautionary savings by individuals seem insufficient to eliminate the increase in consumption inequality over the life cycle.

Footnotes

1. Edward N. Wolff, "The Rich Get Richer...and Why the Poor Don't," *American Prospect*, vol. 12, no. 3 (Spring 2001), p. 15–17.
2. The conditions usually call for smoothing marginal utility over time, but smoothing consumption is a good first approximation.
3. Aggregate risk, where the entire population is affected, is not considered in the *Commentary*.
4. With the existence of insurance, individuals may take fewer precautions to prevent a bad outcome; this is known as moral hazard, and its ramifications are ignored here.

5. It does not matter who actually pays the unemployment insurance, that is, the firm or worker, the outcome will be the same.

6. See Kjetil Storesletten, Chris Telmer, and Amir Yaron, "Accounting for Idiosyncratic Risks over the Life Cycle: Theory and Evidence," Carnegie-Mellon University, unpublished manuscript, 2000.

7. See Michael P. Keane and Kenneth I. Wolpin, "The Career Decisions of Young Men," *Journal of Political Economy*, vol. 105, no. 3 (June 1997), pp. 473–522.

8. Liquidity constraints have been ignored here.

Peter Rupert is a senior economic advisor at the Federal Reserve Bank of Cleveland; Chris Telmer is an assistant professor of financial economics at Carnegie-Mellon University. This Commentary was printed on July 2, 2001.

The views expressed here are those of the authors and not necessarily those of the Federal Reserve Bank of Cleveland, the Board of Governors of the Federal Reserve System, or its staff.

Economic Commentary is published by the Research Department of the Federal Reserve Bank of Cleveland. To receive copies or to be placed on the mailing list, e-mail your request to 4d.subscriptions@clev.frb.org or fax it to 216-579-3050. Economic Commentary is also available at the Cleveland Fed's site on the World Wide Web: www.clev.frb.org/research, where glossaries of terms are provided.

We invite comments, questions, and suggestions. E-mail us at editor@clev.frb.org.

**Federal Reserve Bank of Cleveland
Research Department
P.O. Box 6387
Cleveland, OH 44101**

Return Service Requested:
Please send corrected mailing label to the above address.

Material may be reprinted if the source is credited. Please send copies of reprinted material to the editor.

**PRSRT STD
U.S. Postage Paid
Cleveland, OH
Permit No. 385**