

TWO VIEWS OF INEQUALITY OVER THE LIFE CYCLE

Jonathan Heathcote
Georgetown University

Giovanni L. Violante
New York University

Kjetil Storesletten
University of Oslo

Abstract

Data on the life-cycle profiles of inequality in wages, earnings, hours worked, and consumption contain precious information for answering questions about the ability of households to insure labor market risk and about the sources of this risk. This paper demonstrates that the choice of whether to control for cohort effects or for time effects has a drastic impact on the estimated age profiles for inequality and, thus on the answers to those questions. It also shows that time effects are required to account of the observed trends in inequality in 30 years of U.S. data, whereas there is no evidence that cohort effects have been important. (JEL: C13, D31, D91, J22, J31)

1. Introduction

This paper measures how inequality in wages, hours worked, and consumption varies over the life cycle in the United States. The steepness of the age profile for wage inequality is informative about whether wage dispersion is primarily driven by fixed-effects or life-cycle shocks. Contrasting the age profile for wage inequality with age profiles for inequality in consumption and labor supply can shed light on the degree to which agents can insure wage shocks, and can also be useful for differentiating between alternative specifications for preferences (Heathcote, Storesletten, and Violante 2004a).

The task of documenting how inequality changes with age is complicated by the sharp rise in cross-sectional inequality in U.S. wages and earnings since 1980. In order to isolate the dynamics of inequality with age, we need to take a

Acknowledgments: We thank Dirk Krueger and Fabrizio Perri for help with the CEX consumption data. Heathcote and Violante also thank the Economics Program of the National Science Foundation for financial support. Violante also thanks the CV Starr Center for support. Storesletten thanks the Institute for International Economic Studies and the Frisch Centre for support. The authors also thank the Federal Reserve Bank of Minneapolis for its kind hospitality.

E-mail addresses: Heathcote: jhh9@georgetown.edu; Storesletten: kjstore@econ.uio.no; Violante: glv2@nyu.edu

Journal of the European Economic Association April–May 2005 3(2–3):765–775
© 2005 by the European Economic Association

stand on what accounts for these changes over time. Wage inequality could potentially have increased in recent decades due to either (1) time effects—changes in the economic environment that have increased wage inequality within every age group, or (2) cohort effects—younger cohorts were more unequally endowed with labor market skills than older cohorts.¹ To compute life-cycle facts we must disentangle the relative contributions of time and cohort effects. This decomposition is of independent interest, since it sheds light on what factors are behind the observed increases in cross-sectional inequality.²

To identify age effects and disentangle time and cohort effects we must impose restrictions on how age, time and cohort effects can interact, since these are not independent.³ We follow Juhn, Murphy, and Pierce (1993) and Ameriks and Zeldes (2001) and assume that the effects of age, time, and cohort are additively separable, and that only two of the three effects are operative.⁴ We then consider two linear regression models: one in which unrestricted age and cohort effects are present, but no time effects, and one with age and time effects, but no cohort effects. We find that age profiles of inequality are quite different depending on whether one controls for cohort or time effects. These differences lead to very different interpretations of the driving wage process, including the fraction of wage inequality that is predetermined at the time of labor market entrance. We will argue that the picture of life-cycle inequality that arises from the time effects model is consistent with a standard life-cycle model with endogenous labor supply. Moreover, the life-cycle profiles suggest that while some fraction of wage inequality is insurable, insurance markets are incomplete.

We then examine the data more closely to gauge whether the observed trends in cross-sectional inequality are best captured by time effects or by cohort effects. According to the cohort effect model, growth in within-cohort inequality should be the same at each date. We find that the cohort model is contradicted by the data; growth in within-cohort inequality varies strongly with time. This holds for wage inequality, earnings inequality, and the covariance between wages and earnings. By contrast, looking at the age profile of inequality in repeated cross-sections,

1. A third possibility is that wage inequality increases with age, and the average age of the population has risen over time. However, this cannot be the source of increases in cross-sectional inequality in our sample, since the age composition of our sample is stable over time.

2. Examples of time effects include changing returns to ability and trade liberalization. Examples of cohort effects are variations in the size of birth cohorts, and in the dispersion in education quality. Finally, factors such as a changing returns to education, deunionization, and changes in the minimum wage (which binds mostly for the young) are examples of changes in the environment that combine both cohort and time-effects.

3. To see this, let a denote age, t denote time, and c denote cohort birth year. Let the statistic of interest be $x(a, t, c)$ for agents of age a at date t . By appropriate substitutions of the identity $c = t - a$ it is easy to see that any model of the form $x(a, t, c) = f(a, t, c)$ is equivalent to three other models that are functions of only two of the three arguments.

4. In addition to simplicity, this approach has the advantage that collinearity would likely remain a problem in any approach incorporating all three effects, even if identification were formally achieved.

the slope of the age profile is relatively constant, suggesting that cohort effects can safely be abstracted.

Section 2 describes the data, Section 3 discusses cross-sectional inequality over the period 1979–1996, Section 4 presents our two sets of estimates for inequality over the life cycle, and Section 5 discusses the relative strengths and weaknesses of the time effect and cohort effect views.

2. Data Description

Our empirical analysis is based on two sources of data: the 1968–1997 waves of the *Panel Study of Income Dynamics* (PSID) for wages, hours worked and earnings, and the 1980–1997 waves of the *Consumer Expenditures Surveys* (CEX) for consumption.

We use the PSID data sample of Heathcote, Storesletten, and Violante (2004b, HSV henceforth), and the CEX data sample of Krueger and Perri (2002, KP henceforth) and refer to those papers for details on sample selection. Most importantly, the sample is restricted to households whose “head” (“head of household” in the PSID, the “reference person” in the CEX) is aged 20–65 and works $h \in [520, 5096]$ annual hours. Hourly wages are computed as annual labor earnings divided by annual hours worked. KP’s “total consumption” is nondurables plus services from durables.

There are three important data issues that we only briefly mention here:

Measurement Error in the PSID: Several papers based on the PSID Validation Studies argue that in PSID data, hours are measured with error. Since wages are defined as the ratio of earnings to hours, one implication is that the covariance between hours and wages can be underestimated (“division-bias”). If the measurement error is “classical”, then it will affect the levels of variances and covariances, but not their slopes over time or over the life cycle, which are the focus here.

Measurement Error in the CEX: While KP and Slesnick and Ulker (2004) focus exclusively on the survey data sets in the CEX, Attanasio, Battistin, and Ichimura (2004) find that CEX diary data suggest a somewhat larger increase in consumption inequality over time.

Unit of Analysis: Wages, hours, and earnings are recorded at the level of the individual worker, while consumption is measured only at the level of the household. We therefore convert household consumption data into adult equivalents using the scaling factor, $\hat{c} = \sqrt{\#A + 0.5\#C}$, where $\#A$ and $\#C$ are the number of adults and children living in the household (see Slesnick and Ulker 2004).⁵

5. See HSV for a discussion of the pros and cons of focussing on the individual versus the household as the unit of analysis.

3. Cross-Sectional Inequality over Time (1967–1996)

Despite the sharp rise in cross-sectional inequality in U.S. wages and earnings since 1980, inequality in consumption and hours worked has been relatively stable. However, the underlying consumption-hours distribution has changed significantly: the covariance between consumption and hours has decreased, while the covariance between wages and hours has increased (see HSV and KP).

In HSV we use the panel dimension of the PSID to estimate a time varying stochastic process for individual wages. We find that a calibrated Bewley–Huggett–Aiyagari incomplete-markets economy with endogenous labor supply and the estimated wage dynamics can account—quantitatively—for all the trends described above. The model predicts a minor increase in the variability of consumption because our wage process attributes some of the increase in wage-inequality to transitory shocks that are easy to self-insure. Moreover, it replicates the observed rise in the wage-hours covariance: As the variance of transitory shocks increases, labor supply tracks wages more closely. Consequently, the model can account for the fact that earnings inequality has increased by more than wage inequality.

In HSV our specification for the individual wage process allows for time effects but not for cohort effects. If cohort effects were important, abstracting from them could lead to misleading predictions for the dynamics of inequality over time in other variables. Fortunately, our analysis of life-cycle data in the next section will lead us to conclude that it is reasonable to focus on time effects when modeling changes in the wage-generating process.

4. Cross-Sectional Inequality over the Life Cycle

We start by constructing variances of wages, hours, total and nondurable consumption, and the covariance between hours and wages for observations grouped by year and age. An individual is defined to be of age a if her actual age lies between $a - 2$ and $a + 2$. The typical “cell” contains several hundred observations, although the size range is large. We then estimate the life-cycle profile for each moment, removing year and cohort effects successively. To remove year effects from the typical moment $x(a, t, c)$ (e.g., the variance of wages of individuals with age index a at time t) we assume away cohort effects and run a linear regression with a full set of year and age dummies. To remove cohort effects from the moment we assume away time effects and run a linear regression with a full set of cohort and age dummies.

Figure 1 reports the results. The “time view” and the “cohort view” yield very different results for the variances of wages and earnings and the covariance between wages and hours. The rise in wage and earnings inequality over the life

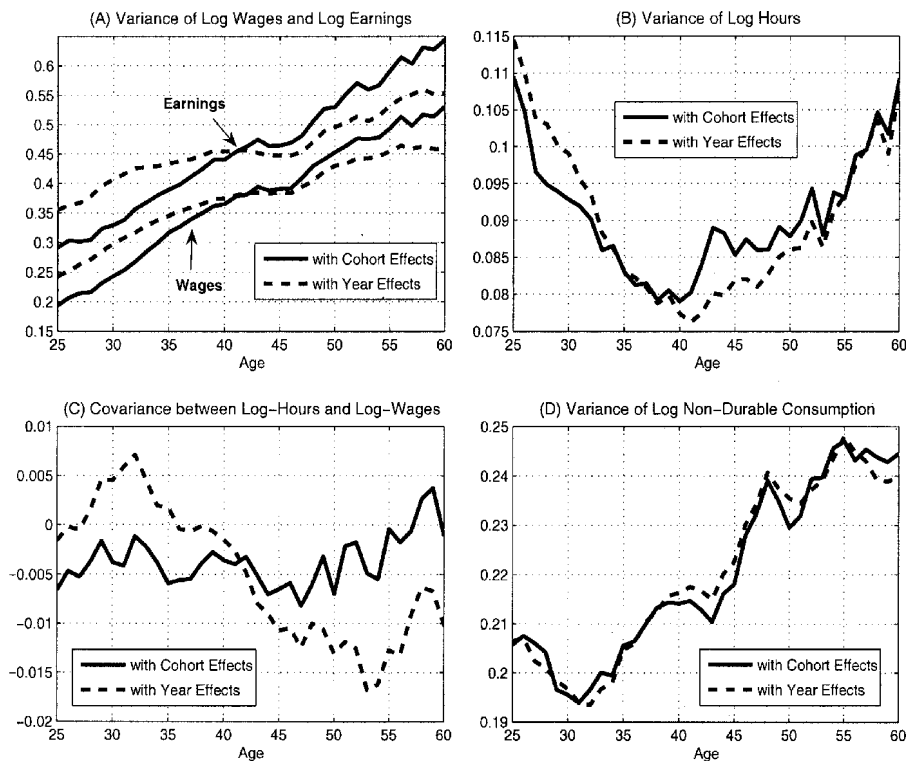


FIGURE 1. Age profiles of inequality: Cohort vs. time effects. Note: The figure portrays the 36 age dummies from running linear regressions abstracting from cohort effects and time effects, respectively. Each graph is normalized to match the unconditional average for the corresponding cross-sectional moment over the period 1980–1997.

cycle is substantially larger when one controls for cohort effects (by a factor of 1.6 for wages and 1.7 for earnings). For the covariance the differences are even more pronounced. When we remove time effects, the covariance falls with age, when we remove cohort effects, it rises slightly. By contrast, the age profiles for the variance of hours and the variance of consumption are largely invariant to the assumptions made about time versus cohort effects.⁶

These age profiles contain information about the driving income process (see also Storesletten, Telmer, and Yaren 2004). For example, assume that the log of an individual's wages can be decomposed into a fixed effect, a persistent autoregressive shock, and a transitory shock. The fact that both age profiles for wages are linear suggests that the persistent component is close to a unit root.

6. To save space the picture for total consumption is omitted from Figure 1. Its age profile is similar to that for nondurable consumption until age 48, but then it flattens out. Moreover, age profiles controlling for time and cohort effects are essentially identical, as for nondurable consumption.

The steepness of the profile identifies the conditional variance of the permanent shock. Thus, the time view suggests substantially less life-cycle risk than the cohort view, implying that a larger fraction of wage and earnings inequality is “pre determined” prior to labor market entrance.

In Heathcote, Storesletten, and Violante (2004a), we estimate an incomplete-markets life-cycle model with endogenous labor supply. That model is quantitatively consistent both with the life-cycle profiles according to the time view, and with the evolution of cross-sectional inequality as described in Section 3. As agents age, accumulated permanent shocks account for a larger fraction of the (within-cohort) cross-sectional variance of wages. When preferences are such that agents increase leisure in response to a permanent wage increase, many of the patterns in Figure 1 emerge: the variance of consumption increases less than the variance of earnings, and the covariance between hours and wages declines with age.⁷

Figure 1 seems inconsistent with the existence of complete markets against wage risk. Complete insurance implies that workers with good (bad) wage shocks should supply more (fewer) hours (see Heathcote, Storesletten, and Violante 2004c).⁸ Consequently, the complete markets model has the following sharp implication: as a cohort grows older and shocks account for a larger fraction of within-cohort wage inequality, the correlation between wages and hours should increase. However, Figure 1C reveals that the wage-hours correlation is close to zero at all ages.

5. Time or Cohorts Effects?

Given the sharp differences between the cohort-view and the time-view documented here, we now try to measure the relative importance of time and cohort effects in the determination of the dynamics of inequality in the US over the last 30 years.

Consider the following general model for the moment $x(a, t, c)$ in which we assume separability between the effects age a , time t and cohort $c = t - a$:

$$x(a, t, t - a) = g_1(a) + g_2(t) + g_3(t - a).$$

7. We have not formally estimated the life-cycle model under the assumption that the empirical moments are those arising from the cohort view. However, since (1) the accumulated permanent shocks must account for an even larger fraction of wage-variance than under the time view, and (2) the cohort view implies an increasing empirical age profile for the wage-hours correlation, we conjecture that the life-cycle model of labor supply will be rejected.

8. This holds even if preferences are non separable between consumption and leisure. Non-separability is needed in order to account for the rising age profile of consumption inequality (Storesletten, Telmer, and Yaron 2001).

In the regressions described in the previous section, the function g_1 , g_2 , and g_3 are implemented as full sets of regression dummies.

To help identify time and cohort effects, we compute three different measures of changes in inequality:

(i) The change within age group a between t and $t + 1$:

$$\Delta x_{t,t+1}^a = g_2(t+1) - g_2(t) + g_3(t+1-a) - g_3(t-a) = \Delta g_2(t) + \Delta g_3(t-a).$$

The average (across age groups) within-age change is

$$\bar{\Delta} x_{t,t+1}^a = \Delta g_2(t) + \bar{\Delta} g_3(t).$$

(ii) The change within cohort $c = t - a$ between t and $t + 1$:

$$\Delta x_{t,t+1}^c = g_1(a+1) - g_1(a) + g_2(t+1) - g_2(t) = \Delta g_1(a) + \Delta g_2(t).$$

The average (across cohorts) within-cohort change is

$$\bar{\Delta} x_{t,t+1}^c = \bar{\Delta} g_1 + \Delta g_2(t).$$

(iii) The within period $t + 1$ change between age a and $a + 1$:

$$\begin{aligned} \Delta x_{a,a+1}^t &= g_1(a+1) - g_1(a) + g_3((t+1) - (a+1)) - g_3((t+1) - a) \\ &= \Delta g_1(a) - \Delta g_3(t-a). \end{aligned}$$

The average (across age groups) within-period change is

$$\bar{\Delta} x_{a,a+1}^t = \bar{\Delta} g_1 - \bar{\Delta} g_3(t).$$

Juhn, Murphy, and Pierce (1993) note that $\bar{\Delta} x_{t,t+1}^a$ and $\bar{\Delta} x_{t,t+1}^c$ have a common component, the time effect. Thus, in the presence of important time effects and negligible cohort effects, these two measures should be strongly correlated. This is the first hypothesis we examine. For each moment Table 1 reports the average within-cohort and within-age-group changes for six periods (four periods for the variance of consumption). First, consider the PSID moments. The correlation between average within-cohort and within-age-group changes (reported in the last column of the table) is always high and positive. In particular, note that in the three time periods from 1973–1987, when most of the rise in wage inequality took place, the magnitude of the within-age and within-cohort changes are remarkably similar for all moments.⁹

A positive correlation between within-cohort and within-age changes might occur even in the absence of time effects if the age effect $\bar{\Delta} g_1$ and the cohort

9. The correlation coefficient restricted to the periods 1973–1977, 1978–1982, and 1983–1987 are above 0.92 for every PSID moment.

TABLE 1. Average annual change within cohort, within age group, and between successive age groups.

	1968–1972	1973–1977	1978–1982	1983–1987	1988–1992	1993–1996	Correlation
	Variance of log wages						
Within cohort $\Delta \times^c_{i,t+1}$	-0.2727	-0.4166	0.9909	1.2466	0.8135	0.5877	
S.E.	0.113	0.127	0.111	0.180	0.217	0.246	
Within age group $\Delta \times^a_{i,t+1}$	0.5426	-0.4642	0.9262	0.9266	0.1813	-0.0226	0.647
S.E.	0.190	0.225	0.210	0.319	0.274	0.395	
Between age group $\Delta \times^f_{a,t+1}$	0.5485	0.5861	0.3925	0.6137	0.8134	0.8029	
S.E.	0.113	0.127	0.111	0.180	0.217	0.246	
	Variance of log earnings						
Within cohort $\Delta \times^c_{i,t+1}$	-0.3994	-0.2912	1.7410	1.6375	0.6589	0.6158	
S.E.	0.157	0.169	0.174	0.237	0.274	0.331	
Within age group $\Delta \times^a_{i,t+1}$	0.5948	-0.2434	1.4004	1.0915	-0.1149	-0.0696	0.691
S.E.	0.326	0.294	0.289	0.414	0.421	0.518	
Between age group $\Delta \times^f_{a,t+1}$	0.4787	0.4946	0.3338	0.4990	0.8192	0.8261	
S.E.	0.157	0.169	0.174	0.237	0.274	0.331	
	Variance of log hours						
Within cohort $\Delta \times^c_{i,t+1}$	-0.1354	-0.2044	0.4152	-0.0181	0.3669	-0.2475	
S.E.	0.059	0.057	0.059	0.070	0.072	0.081	
Within age group $\Delta \times^a_{i,t+1}$	0.1298	-0.1594	0.2361	-0.0846	0.2358	-0.4560	0.804
S.E.	0.136	0.113	0.102	0.125	0.124	0.159	
Between age group $\Delta \times^f_{a,t+1}$	-0.0954	-0.0588	-0.0391	0.0050	-0.0049	-0.0226	
S.E.	0.059	0.057	0.059	0.070	0.072	0.081	

(continued)

effect $\bar{\Delta}g_3(t)$ were similar. The hypothesis that age effects and cohort effects happen to coincide, and that time effects are small can be examined in two ways. First, in the absence of time effects, the within-cohort change $\bar{\Delta}x_{t,t+1}^c$ should be constant across time periods. This is rejected in our data: the majority of the within cohort changes for any given variable are statistically different between periods. Second, the within-period between-age change $\bar{\Delta}x_{a,a+1}^t$ should be zero, since the age effect and the cohort effect enter with opposite offsetting signs in the expression. However, in our sample the within-period between-age change is statistically different from zero in every period. Thus both tests suggest that we should reject the hypothesis that time effects are small.

To go one step further, note that if cohort effects are small, then $\bar{\Delta}x_{a,a+1}^t$ should be constant across periods. This hypothesis cannot be rejected by the PSID moments. We conclude that the evidence speaks in favor of time effects, rather than cohort effects.

We do not analyze changes in inequality in consumption between cohorts and age-groups because changes in the variance of log consumption are essentially not significant and are therefore not useful in distinguishing between time and cohort effects.

6. Concluding Remarks

In this paper we have shown that the dynamics of cross-sectional inequality across time and age groups are consistent with the presence of time effects and absence of cohort effects, but inconsistent with the presence of cohort effects and absence of time effects.

Furthermore attributing rising inequality over time solely to cohort effects paints a misleading picture of inequality over the life-cycle. Inequality rose dramatically within cohorts during the 1980s, so a model without time effects must attribute this increase in inequality to a very steep age profile for inequality. By contrast, a model with time effects attributes much of the rise in inequality to time effects, implying a flatter “steady-state” wage profile for inequality. This might explain why Deaton and Paxson (1994), who abstracted from time effects, found a dramatic increase in consumption inequality by age using data from the 1980s—a period of sharply rising wage inequality—while we and Slesnick and Ulker (2004) find a smaller increase when extending the sample to include the 1990s.

References

- Ameriks, John and Stephen P. Zeldes (2001). “How Do Household Portfolio Shares Vary with Age?” Manuscript, Columbia University.
- Attanasio, Orazio, Erich Battistin, and Hide Ichimura (2004). “What Really Happened to Consumption Inequality in the U.S.?” NBER Working Paper 10338.

- Deaton, Angus and Christina Paxson (1994). "Intertemporal Choice and Inequality." *Journal of Political Economy*, 102, 437–467.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2004a). "A Tractable Framework for understanding Dispersion in Labor Supply and Consumption." Working paper, Georgetown University.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2004b). "The Macroeconomic Implications of Rising Wage Inequality in the United States." CEPR Working Paper 4296.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2004c). "Insurance and Opportunities: The Welfare Implications of Rising Wage Dispersion. Working paper, Georgetown University.
- Juhn, Chinhui, Kevin M. Murphy, and Brooks Pierce (1993). "Wage Inequality and the Rise in Returns Skill." *Journal of Political Economy*, 101, 410–442.
- Krueger, Dirk and Fabrizio Perri (2002). "Does Income Inequality Lead to Consumption Inequality? Evidence and theory." Working paper, Stanford University.
- Slesnick, Daniel T. and Aydogan Ulker (2004). "Inequality and the Life Cycle: Age, Cohort Effects, and Consumption." Working paper, University of Texas, Austin.
- Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron (2001). "How Important are Idiosyncratic Shocks? Evidence from Labor Supply." *American Economic Review*, 91(2), 413–417.
- Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron (2004). "Consumption and Risk Sharing over the Life Cycle." *Journal of Monetary Economics*, 51, 609–633.