



ISSN 2282-6483

Alma Mater Studiorum - Università di Bologna
DEPARTMENT OF ECONOMICS

Revisiting wage, earnings, and hours profiles

Peter Rupert
Giulio Zanella

Quaderni - Working Paper DSE N°936



Revisiting wage, earnings, and hours profiles*

Peter Rupert
University of California, Santa Barbara[†]

Giulio Zanella
University of Bologna[‡]

October 2012

Abstract

We document empirical life cycle profiles of wages, earnings, and hours of work for pay from the Panel Study of Income Dynamics, following the same workers for up to four decades. For six of the eight cohorts we analyze the wage profile does not decline with age (not before 65, at least), while the earnings profile always does. The discrepancy is explained by a sharp drop in the hours of work for pay profile beginning shortly after age 50, when many workers start a smooth transition into retirement by working progressively fewer hours. This pattern is not an artifact of staggered abrupt retirement, and is robust to attrition and selection-correction (i.e., taking into account that the composition of our sample, for a given cohort, changes over time). We explore the nontrivial restrictions on dynamic models of the aggregate economy that this evidence suggests, and we provide numerical profiles that can be readily used in quantitative macroeconomic analysis.

JEL Classification Codes: E24; J13; J22; J24; J26

Keywords: life cycle, wage profile, labor supply, intensive margin, human capital, pre-retirement

*This is the working paper version of Rupert and Zanella (2014), “Revisiting wage and hours profiles”. The working paper was based on more limited data and provided less insight than the final version, although the conclusion is the same. We are grateful to Eric French, Andrea Ichino, Gueorgui Kambourov, Peter Kuhn, and Richard Rogerson for useful suggestions at the very early stages of this project, as well as to seminar participants at the universities of Aarhus, Alicante, Bologna, Chicago Booth, Humboldt-Berlin, IMT Lucca, Leicester, Rome-La Sapienza, Siena, the Stockholm School of Economics, as well as to participants to the EALE/SOLE 2010 meeting in London, the SED 2010 meeting in Montreal, and the EEA 2010 congress in Glasgow. Bonnie Queen provided outstanding research assistance. All errors are ours.

[†]E-mail: rupert@econ.ucsb.edu

[‡]Corresponding author. E-mail: giulio.zanella@unibo.it

1 Introduction

This paper is an empirical investigation into life cycle profiles of wages (hourly compensation of labor), hours (work for pay supplied to the labor market), and earnings (the product of the two), from the vantage point of four-decades-long panel data. Theoretical and empirical investigation of such profiles has a long tradition in labor economics and macroeconomics, because of their importance in understanding a variety of phenomena such as labor supply, retirement, the allocation of time, incentive contracts, training and human capital, saving, and inequality. The Panel Study of Income Dynamics (PSID) now covers more than 40 years and spans nearly the entire working lives of the cohorts that entered the labor market at the end of the 1960s, as well as substantial parts of the working lives of younger cohorts. We use these data to estimate life cycle profiles that are straightforward to interpret and that can be readily used in quantitative macroeconomic analysis. We then use these profiles to explore the restrictions the data impose on a benchmark life cycle model.

We find that wages typically do not fall (if they fall at all) until individuals are well into their 60's, in line with some previous empirical research. However, hours fall substantially (and so do earnings) beginning shortly after age 50. The drop in hours is generated by a smooth transition into retirement, in the form of less overtime work and passage from full- to part-time. Although not documented in a systematic way, [Mincer \(1974\)](#) had inferred this pattern from census data decades ago. Analyzing these data, Mincer observed that the weekly wage rate was not declining at the end of working life while annual earnings were. He wrote: “There is no visible decline at these later ages in weekly earnings. Apparently, declines in weeks worked per year are the main factor in the decline of annual earnings during preretirement years.” (p. 70). This suggestion remained somewhat overlooked. The pre-retirement pattern we document conforms to the more general notion of retirement advocated by [Heckman \(1976\)](#): “Retirement can more generally be defined as a period with few hours of work supplied to the market.” (p. S15). The pattern of non-declining wages and falling hours shortly after age 50 imposes nontrivial restrictions on quantitative aggregate models. In particular, such a pattern is inconsistent with a benchmark life cycle model with stationary uncertainty, complete markets, wages independent of past labor supply, and constant disutility of work along the life cycle. We discuss at the end of the paper the extent to which alternative departures from the benchmark model reconcile theory and data. In brief, the two departures that seem best able to reproduce the empirical pattern are endogenous wages from on-the-job training (leading to skills that do not depreciate until very late in life) and increasing marginal disutility of work (or increasing marginal utility of alternatives to work for pay) shortly after age 50—not necessarily because of changing preferences. Additionally, we provide the numerical life cycle profiles that quantitative models should reproduce.

It is important to clarify what this paper is *not* about. First, the focus here is on the intensive margin of labor supply. That is, when we construct the empirical hours profile and when we try to reproduce it in the model we consider hours *per worker*, not hours *per person*. The reason why we do not deal with the extensive margin here is that retirement choices are already well understood withing the life cycle model. We know from related research that technological features such as nonconvexities in the mapping between hours and labor services (e.g., [Rogerson and Wallenius \(2009\)](#), [Erosa et al. \(2011\)](#)) and institutional features such as social security rules (e.g., [French \(2005\)](#), [Erosa et al. \(2012\)](#)) allow a match of extensive margin behavior late in the life cycle. It is the intensive margin pattern we document here that is less explored. A first-order concern with

our intensive margin focus is that the pattern of declining hours in spite of non-declining wages may be an artifact of extensive margin movements. We will show that this is not the case. Second, we do not attempt to separate age, time, and cohort effects on life cycle profiles. Our analysis produces evidence that researchers should take into account, namely that workers in many cohorts substantially reduce labor supply along the intensive margin despite facing a non-declining wage profile, whatever the reasons why wages do not fall for these cohorts. We do not need to take a stand on these reasons. One possibility is that there are pure time or cohort effects at work. Another is that human capital depreciates only slowly, or not at all. Yet another explanation is that incentive contracts are in place and the observed wage profile does not track the individual productivity profile closely.¹ This is certainly a very interesting question for future research.

The remainder of the paper is organized as follows. [Section 2](#) connects the paper to the literature. [Section 3](#) describes the data. [Section 4](#) illustrates how these data restrict the theory. [Section 5](#) concludes. In the Appendix we report tables containing the numerical wage, hours, and earnings profiles (raw and smoothed profiles) that can be readily used in quantitative analysis. These, of course, are also available in electronic format at our research pages.

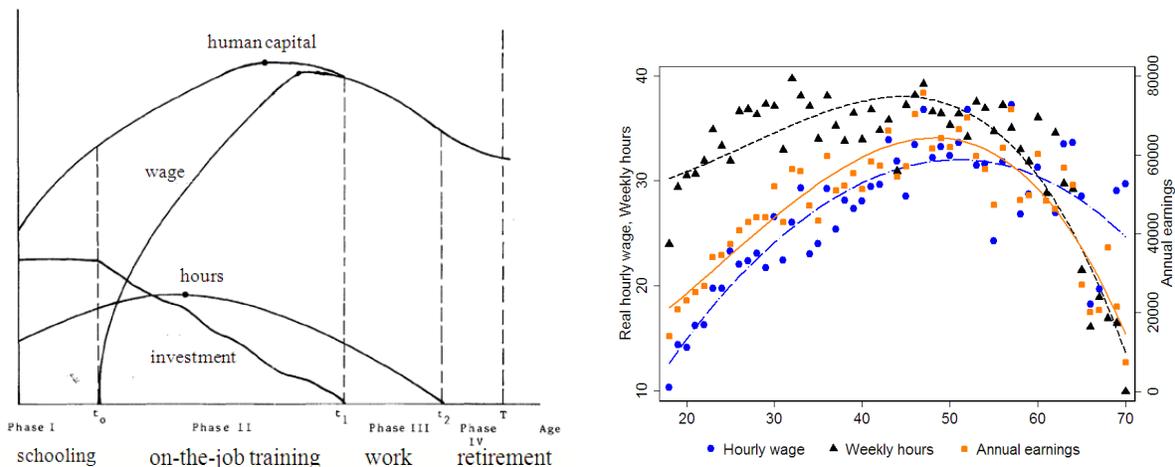
2 Connection to the literature

From a historical viewpoint, among the theories developed to characterize life cycle profiles, a prominent position is occupied by the human capital model, initiated by [Ben-Porath \(1967\)](#) and further developed by [Ghez and Becker \(1975\)](#), [Blinder and Weiss \(1976\)](#), [Ryder et al. \(1976\)](#), [Heckman \(1976\)](#), and [Rosen \(1976\)](#). Variants of this model have become a workhorse in dynamic macroeconomics and typically predict that wage, hours, and earnings profiles are “hump-shaped”. In this model the wage rate is the return on human capital, and grows as long as net investment increases the stock of skills. As the end of working life approaches, investment in human capital optimally falls below depreciation and the wage rate declines, tracing out a hump-shaped profile. The hours profile has the same shape because the substitution effect induces workers to work more when wages are higher. The earnings profile, then, is also hump shaped. The left portion of [Figure 1](#), reproduced from [Weiss \(1986\)](#), illustrates the typical dynamics generated by the human capital model, evidently to fit the data observed at the time it was developed, as summarized by [Weiss \(1986\)](#), p. 603: “The major stylized facts which the theory attempts to explain are: a life cycle earnings profile which is increasing at early ages and is declining towards the end of the working period. A wage profile which tends to increase over the life cycle with a weak tendency for wage reduction towards the end of the working period. An hours of work life cycle profile which is increasing at early ages and declining at older ages, with the peak occurring earlier than in the earnings or wage profiles.” When these empirical regularities were first isolated the main data sources were cross-sectional, so “the major stylized facts which the theory attempt[ed] to explain”

¹The works, among others, of [Becker and Stigler \(1974\)](#), [Lazear \(1979\)](#), [Lazear \(1981\)](#), [Freeman \(1977\)](#), [Medoff and Abraham \(1980\)](#), [Medoff and Abraham \(1981\)](#), [MacDonald \(1982\)](#), and [Harris and Holstrom \(1982\)](#) show that, at the individual level, wage growth is possible even in the absence of human capital (and productivity, more generally) growth. In this case the wage profile need not decline towards the end of the working life, even if individual productivity did. [Gibbons and Waldman \(1999\)](#) offer a thorough review of this literature.

were mostly derived from synthetic cohorts.² We produce and report in the right portion of Figure 1 the wage, hours, and earnings profiles that a researcher would have observed applying the synthetic cohort fiction to the 1970 cross-section of the PSID. They clearly resemble the theoretical profiles.

Figure 1: Theoretical and 1970 PSID synthetic cohort profiles



Notes: The figure on the left illustrates the theoretical life cycle profiles in the human capital model, and is reproduced from Weiss (1986). The figure on the right illustrates the empirical profiles constructed using the synthetic cohort method from the 1970 cross section of the PSID. The wage and earnings profiles result from adding 1% annual real wage growth to the raw cross sectional profiles, corresponding to the BLS estimate of productivity growth—this is roughly the average growth rate of labor productivity in the US since 1970. That is, the wage of an individual a years old in the cross section is multiplied by $(1.01)^{a-18}$, where 18 is the conventional age of the youngest worker in the 1970 cross section. The interpolating line is the best second-order fractional polynomial.

However, synthetic cohorts may produce a biased picture of the life cycle when productivity changes in time. Thornton et al. (1997) and Rubinstein and Weiss (2006) offer an illustration for annual and weekly earnings, respectively. When using repeated cross sections to form pseudo panels, many researchers have reached similar conclusions—i.e., wages decline in the second part of the working life and hours track them. For instance, Browning et al. (1985) apply this procedure to British household heads in the Family Expenditure Survey and find that wages peak around mid-working life and then decline, and similarly for hours. The increasing availability of longitudinal data has allowed a more direct look at actual portions of workers’ careers by following them over time. As far as the wage profile is concerned, these analyses have tended to confirm the suggestion of Mincer (1974) that the profile is not hump-shaped—see, for instance, Johnson and Neumark (1996) and references therein. In particular, Johnson and Neumark (1996) use panel data from the National Longitudinal Survey of Older Men to infer the dynamics of male wages during the late portion of the working life, and do not find clear evidence of negative wage growth until workers are in their 60s. They do not consider the associated labor supply behavior, however. Recent research in macroeconomics, based on micro panel data, has focused on earnings profiles, finding that they are hump-shaped. Examples include Huggett et al. (2007), and Heathcote et al.

²By synthetic cohort we mean the time series constructed from a cross-section at, say, time t using the $a + j$ years olds at time t as a counterfactual for the a years olds at time $t + j$.

(2008). The dynamics of earnings, however, results from both wages and hours of work, and it is of interest to look separately at these two distinct (although not necessarily independent) processes. We improve upon this research by taking a systematic approach to the analysis of life cycle profiles in longitudinal data. Ours is a joint analysis of wages, hours, and earnings profiles of both men and women in multiple cohorts. Furthermore, we consider the entire life cycle, not only the late working life, by exploiting the full longitudinal dimension of the PSID (1968-2009).

A number of additional recent papers in macroeconomics are directly or incidentally concerned with life cycle profiles. Among them, [French \(2005\)](#) studies the effect of health and social security rules on hours and retirement behavior of American male household heads. As a preliminary step in his structural investigation, French uses the annual portion of the PSID (1968-1997) to estimate wage and hours profiles. He finds that the wage profile is hump-shaped, although the most visible decline does not happen before age 62. The hours profile he estimates is also declining, with the trend changing a bit later than we estimate. French notes that, “Most of the variation in the wage and labor supply profiles is from individuals aged 55-65 [...]. [The] decline in hours coincides closely with the decline in wages” (p. 412). Apparently, this contrasts with our results despite using much of the same data. However, the samples used are different (Male head of households 1968–1997 in French, all individuals 1968–2009 here), as well as the estimation methods. For the wage profile, French uses the prediction from a fixed-effects regression of wages—corrected for tied wage-hours effect—on a quartic function of age, conditioning on health status and family size, and pooling all cohorts; we report simple raw profiles by cohort. Relative to these differences, the dissimilarities between ours and French’s estimates seem unimportant.

[Kuruscu \(2006\)](#) estimates the return to on-the-job training using NLSY data. This data set does not yet allow the observation of wages of older workers, so he *assumes* that the wage profile is not declining late in the working life, based on limited evidence from cross-sectional data as well as the PSID. Correspondingly, he assumes that human capital does not depreciate in his model. Our analysis confirms the soundness of such an assumption. In Kuruscu’s model, however, labor is inelastically supplied, and this puts aside the question of what happens to hours.

[Imai and Keane \(2004\)](#) study a human capital model to reconcile small and large intertemporal elasticities of labor supply at the individual and aggregate level. The model is estimated using data from the NLSY, like [Kuruscu \(2006\)](#), and produces out-of-sample (i.e., after age 40) predictions that have the shape of the profiles in [Figure 1](#). As far as the shape of the wage profile is concerned, the difference between [Imai and Keane \(2004\)](#) and [Kuruscu \(2006\)](#) originates in different assumptions about whether human capital depreciates or not. Nonetheless, as we illustrate later, [Imai and Keane \(2004\)](#) offers a key insight, investment in human capital increases substantially the opportunity cost of time early in the life cycle, to reconcile a non-declining wage profile and a declining hours profile. [Rogerson and Wallenius \(2009\)](#), like [Imai and Keane \(2004\)](#), are interested in understanding the discrepancy between micro and macro elasticities of labor supply. The main feature of their model is a nonconvexity (a flat initial portion in the mapping between hours of work and labor services) that generates motives for entering and exiting the labor force at specific points in the life cycle. Their model takes the wage profile as exogenous and assumes it is hump-shaped. This implies that hours decline before retirement. In this model, the only way one can generate the decline in hours (and the other results) when wages do not decline is by assuming that the disutility of working increases at later stages in the life cycle. This is another possibility we explore at the end of this paper.

3 Data

Our main data source is the Panel Study of Income Dynamics (PSID). We also use the Current Population Survey (CPS) as a check against small sample size and attrition. As will become clear, we believe the advantages of using the PSID as base data outweigh the disadvantages. After the full release of the 2009 wave at the end of 2011, the PSID offers the unique opportunity to observe individual life cycle profiles spanning 41 years. This amounts to virtually the entire working life of the cohorts who entered the labor market at the end of the 1960s. We use individual-level data containing demographic, socioeconomic and labor market information from 1967 to 2008, for different cohorts.³ We compute averages of hourly wages, hours, and earnings for a given cohort and all available years to trace the life cycle profiles of individuals in that cohort. This raises a fundamental selection problem (we only observe wages for workers) which we postpone until [section 3.2](#). Attrition is also a potential problem, because we follow individuals for up to 41 years. This issue, too, is postponed until [section 3.2](#). We anticipate here that these processes do not affect our results. Wages and earnings are before-tax and are expressed in constant 2010 dollars using the CPI-U as deflator. In order to have a sufficient number of observations to compute meaningful averages, we define cohorts based on 5-year bins and use the central value of the bin to keep track of the age of the cohort. For instance, workers who were between 21 and 25 years old in 1967 we define as the “23 years old” cohort (or “cohort 23”, for short), those who were between 26 and 30 years old in 1967 are the “28 years old” cohort, etc. We can use eight such cohorts from the PSID. The youngest is the 23 years old cohort (i.e., those born between 1942 and 1946), and the oldest is the 58 years old cohort (i.e., those born between 1907 and 1911). The former is the cohort we analyze more closely, because it is the only one for which we can observe the entire working life. The older cohorts (as opposed to the younger ones, relative to cohort 23) are of interest because it is the shape of life cycle profiles during the late stages of the working life that we want to emphasize. Wage rates are directly observed for workers paid by the hour (about 1/3 of the total) and those that are estimated are done so by the ratio between annual hours and annual earnings. It is well known that this procedure, albeit standard, leads to measurement error in hourly wages. This problem here is vastly mitigated by averaging across individuals in the cohort. In constructing our final sample we apply the following selection criteria. First, we restrict our analysis to households belonging to the SRC sample.⁴ Second, self-employed individuals are excluded from the analysis both because their wage rate is more likely to be measured with error, and because their wage is a mix of labor and capital income. Third, we exclude individuals whose nominal wage rate in a given year is less than 50% the federal minimum wage in force that year. Fourth, in order to eliminate outliers we trim the distribution of wages, hours, and earnings at or above the top 1% and at or

³A year in our data set refers to the period a variable refers to, not the period in which the data were actually collected. For instance, the 1968 wave collected information referring to year 1967, the 1969 wave collected information referring to year 1968, and so on.

⁴The core sample of the PSID originates from two distinct samples: a cross-sectional national representative sample of about 3000 households (the 1968 SRC sample) and a non-representative sample of about 2000 low-income families located in metropolitan areas (the SEO sample). The SRC sample was representative of the U.S. population back in 1967, so there is no guarantee that it is representative during the following 41 years. However, this is the group that maximizes the representativeness of our sample across the four decades. Another way to make the PSID representative of the U.S. population is to apply individual sampling weights to the full core sample. Unfortunately, this cannot be done in a consistent way across all of the 41 waves. [Fiorito and Zanella \(2012\)](#) illustrate how representative the components of the PSID core sample are.

below the bottom 1% every year. This eliminates a handful of observations every year. Trimming at the top helps taking care of top-coded wage information, which we do not correct otherwise. [Table 1](#) reports summary statistics of relevant demographic and socio-economic characteristics for 8 different cohorts, over the entire 1967–2008 period.

Table 1: Summary statistics by cohort

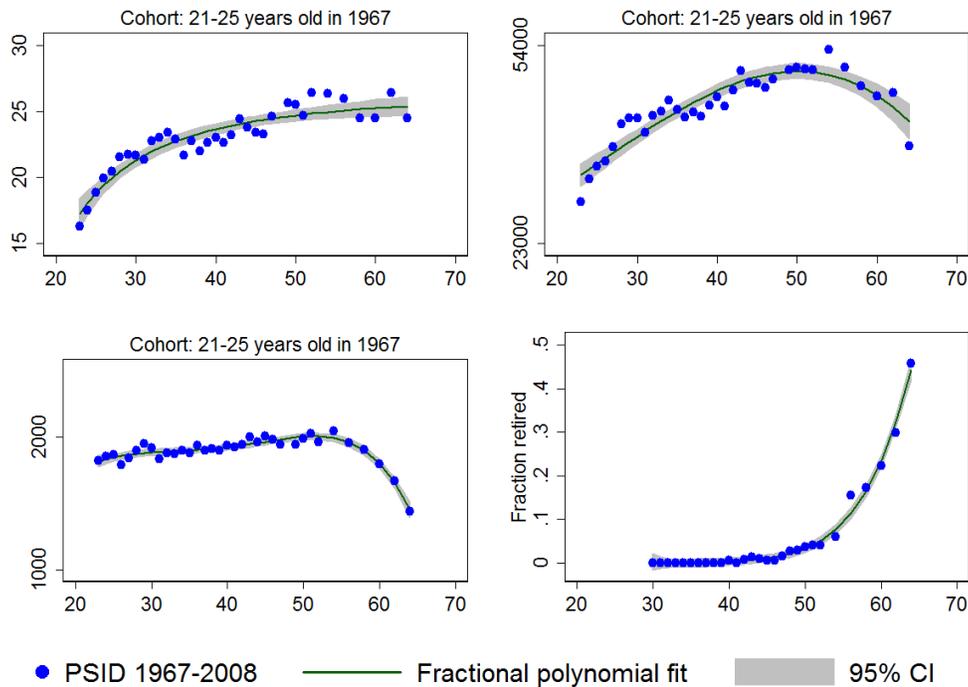
Variable	Cohort (age in 1967)							
	23	28	33	38	43	48	53	58
Sample size in 1967	497	470	484	580	500	504	415	386
Age	39.6	44.2	49.2	53.9	58.0	61.7	65.7	69.1
	(11.2)	(11.2)	(11.2)	(11.0)	(10.7)	(10.1)	(9.3)	(8.5)
Male	0.49	0.45	0.43	0.48	0.48	0.42	0.42	0.42
	(0.50)	(0.50)	(0.50)	(0.50)	(0.50)	(0.49)	(0.49)	(0.49)
White	0.92	0.87	0.85	0.90	0.88	0.89	0.88	0.89
	(0.26)	(0.34)	(0.36)	(0.30)	(0.32)	(0.31)	(0.32)	(0.31)
Black	0.05	0.09	0.09	0.08	0.08	0.08	0.09	0.09
	(0.21)	(0.28)	(0.29)	(0.26)	(0.27)	(0.26)	(0.29)	(0.28)
Disability	0.13	0.15	0.20	0.21	0.29	0.31	0.38	0.42
	(0.34)	(0.35)	(0.40)	(0.40)	(0.45)	(0.46)	(0.49)	(0.49)
Married	0.81	0.82	0.81	0.83	0.79	0.72	0.67	0.59
	(0.39)	(0.38)	(0.39)	(0.38)	(0.41)	(0.45)	(0.47)	(0.49)
Widowed	0.02	0.04	0.06	0.09	0.13	0.21	0.26	0.32
	(0.13)	(0.19)	(0.24)	(0.29)	(0.34)	(0.41)	(0.44)	(0.47)
Divorced	0.11	0.10	0.09	0.06	0.05	0.04	0.04	0.03
	(0.31)	(0.30)	(0.28)	(0.23)	(0.23)	(0.20)	(0.19)	(0.17)
Separated	0.02	0.02	0.02	0.01	0.01	0.01	0.01	0.01
	(0.15)	(0.14)	(0.14)	(0.10)	(0.09)	(0.09)	(0.10)	(0.08)
Children	1.15	1.15	0.99	0.70	0.42	0.28	0.14	0.10
	(1.25)	(1.43)	(1.48)	(1.23)	(0.97)	(0.82)	(0.56)	(0.49)
College or more	0.25	0.20	0.19	0.19	0.14	0.14	0.07	0.11
	(0.43)	(0.40)	(0.39)	(0.39)	(0.35)	(0.34)	(0.26)	(0.32)
High School	0.43	0.37	0.43	0.43	0.41	0.38	0.35	0.27
	(0.50)	(0.48)	(0.49)	(0.49)	(0.49)	(0.49)	(0.48)	(0.45)
Less than High School	0.14	0.23	0.26	0.26	0.36	0.35	0.50	0.51
	(0.34)	(0.42)	(0.44)	(0.44)	(0.48)	(0.48)	(0.50)	(0.50)
Total family income	84.01	87.04	85.16	84.00	69.58	62.74	50.61	47.20
	(71.32)	(90.35)	(79.75)	(84.29)	(54.07)	(53.06)	(50.60)	(73.08)

Notes: The table contains means and (in parentheses) standard deviations. A cohort is labelled by age in 1967, and results from pooling into a five-year bin individuals belonging to five adjacent one-year bins. For example, cohort 23 is composed of individuals who were between 21 and 25 years old in 1967. “Disability” is the share of individuals who report having a condition that limits the amount of type of work they can do. “Children” is the number of individuals below age 18 living in one’s household. “Total family income” is the sum of all individual gross incomes (from all sources, including losses) in one’s family, and is expressed in thousands of 2010 U.S. dollars.

3.1 Empirical profiles

We begin with an illustration of raw life cycle profiles in the PSID for the “23 years old” cohort, the only one for which we observe virtually the entire working life. Individuals in this cohort were between 21 and 25 years old at the first observation in 1967 and between 62 and 66 years old at the last available observation in 2008. Figure 2 illustrates the wage, earnings, and hours profiles of this cohort, as well as the share reporting “retired” as labor market status each year. The interpolating line is the best second-order fractional polynomial, except for hours where a fourth-order polynomial is used.⁵ The figure shows that the wage profile is increasing throughout the working life, and does not show any tendency to decline. Wage growth is sustained until about age 40 and modest between 40 and 50—two well-known facts—but virtually zero thereafter. The earnings profile, instead, has the familiar hump-shape: earnings reach a maximum around age 50 and then decline visibly. The hours profile shows that hours increase slowly until age 50 and then fall sharply, thus driving earnings down. Starting at this same age, more and more individuals in the cohort begin classifying themselves as retired.

Figure 2: Life cycle profiles: cohort 23

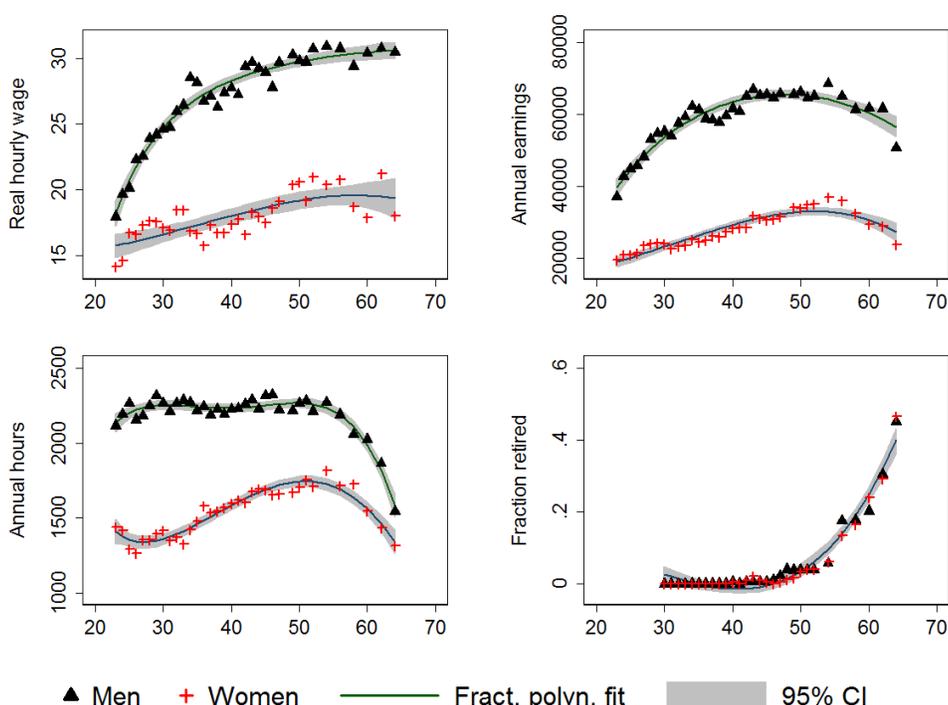


Notes: The figure reports wage, hours, earnings, and retirement profiles of individuals who were between 21 and 25 year old in 1967 (cohort 23). The interpolating line is the best second-order fractional polynomial, except for hours where a fourth-order polynomial is used.

⁵A fourth order polynomial is more appropriate for hours, given the sharp turning point. Since no such turning point is obvious for wages, a second order polynomial is preferable. We smooth earnings like wages to allow a more direct comparison of the two profiles.

We now decompose these profiles by gender and education. Figure 3 shows the gender decomposition. Wages, both in levels and growth rates, are lower for women, as are earnings. While earnings are hump-shaped, wages show only a modest tendency to decline for women. The hours profile of men is flat from shortly after the beginning of the working life until shortly after age 50, when it falls sharply. Women, instead, reduce hours until about age 30 and then increase them steadily until shortly after age 50—this is what drives the slow increase in the aggregate hours profile. At that point hours fall at a rate comparable to that of men. The hours profile of women clearly reflects the importance of part-time jobs in this group, as well as fertility choices. The figure also shows that there are no relevant gender differences in retirement rates in this cohort.

Figure 3: Life cycle profiles: cohort 23, by gender

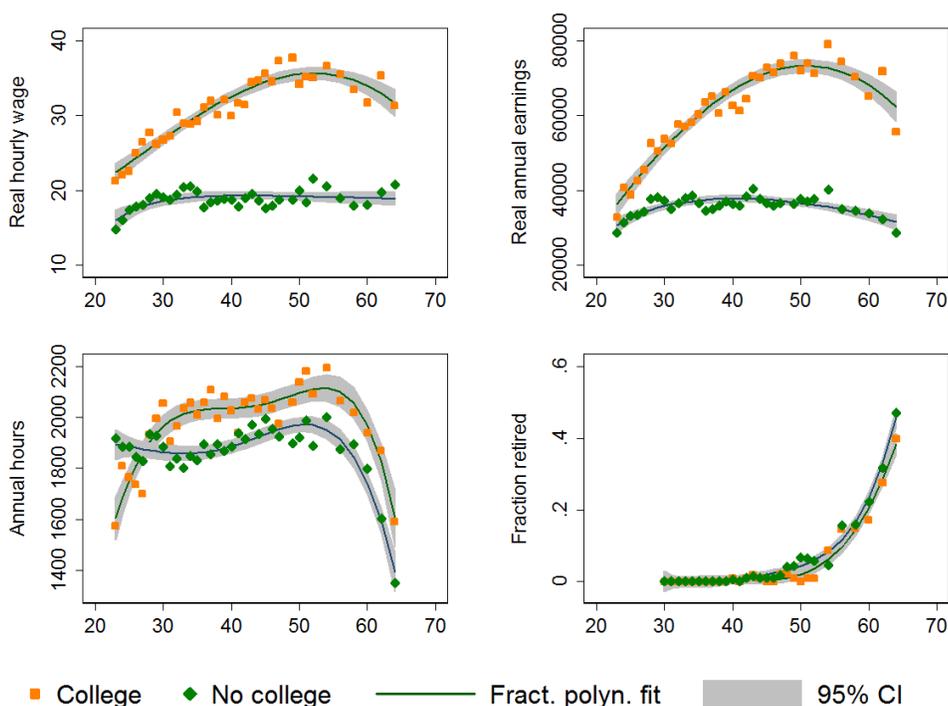


Notes: The figure reports wage, hours, earnings, and retirement profiles by gender of individuals who were between 21 and 25 years old in 1967 (cohort 23). The interpolating line is the best second-order fractional polynomial, except for hours where a fourth-order polynomial is used.

Figure 4 shows the decomposition by education. We consider two groups: college degree and high school degree or less. The wage profiles of the two groups grow at a comparable rate until about age 30. Thereafter, the profile of the less educated is virtually flat and that of college graduates keeps growing at a sustained rate until age 50. At that point there is some negative wage growth. Both earnings profiles decline after age 50, when hours fall sharply at a similar rate. The hours profile of college graduates, however, is above the profile of non graduates (the difference is about 200 hours), except during a worker's 20s, when college graduates are allocating most of their time to education and training and so work less. Finally, the picture suggests that the less educated tend to retire earlier.

Table 2 summarizes the percentage changes of raw wage, earnings, and hours between ages 23 and 51, and between ages 51 and 64 in this cohort, for the whole sample and by gender and education.

Figure 4: Life cycle profiles: cohort 23, by education



Notes: The figure reports wage, hours, earnings, and retirement profiles by level of education (college degree and high school degree or less, respectively) of individuals who were between 21 and 25 years old in 1967 (cohort 23). The interpolating line is the best second-order fractional polynomial, except for hours where a fourth-order polynomial is used.

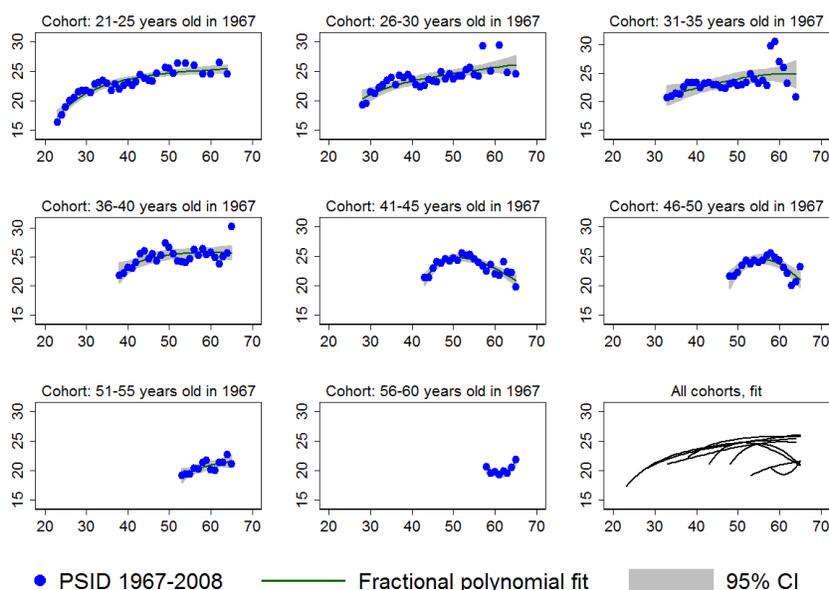
Table 2: Cumulative variations in the raw data, cohort 23

	% wage change		% earnings change		% hours change	
	23-51	51-64	23-51	51-64	23-51	51-64
All	51.3	-0.6	70.6	-24.0	11.2	-29.1
Males	65.5	2.5	75.0	-21.7	4.0	-32.4
Females	35.5	-6.2	78.8	-31.7	21.9	-25.2
College	59.3	-10.8	81.4	-24.8	20.6	-27.2
No college	24.1	13.1	28.8	-22.4	3.7	-32.1

Notes: The table reports the percentage change in hourly wages, earnings, and hours of market work for individuals who were between 21 and 25 years old in 1967 (cohort 23), for the whole sample as well as by gender and education. The variations are between age 23 and 51, and between age 51 and 64, and refer to the raw PSID data plotted in Figure 2, Figure 3, and Figure 4.

So far we have considered the youngest cohort for which we can observe the entire life cycle. While we can conjecture that the *shape* of the wage profile is invariant across cohorts in the early stages of the life cycle, the shape at the end is a more sensitive issue. It could be that cohort 23 benefited from positive aggregate shocks in the final years of the career, between 2000 and 2008. Absent such shocks the wage profile might have exhibited a hump-shape. We can look at older cohorts to check that the pattern we have isolated is not driven by shocks that affected cohort 23. Figure 5, Figure 6, and Figure 7 illustrate the wage, earnings and hours profiles of the eight cohorts summarized in Table 1.⁶ These pictures reproduce the pattern observed in Figure 2, with only two exceptions as far as the wage profile is concerned: cohort 43 and cohort 48. The wage decline for *both* these cohorts begins around 1980, when the two groups are between 54 and 58, and 59 and 63 years old, respectively, and continues afterward until they reach age 63-67 (i.e., age 65 in the picture). This may be due to cohort-specific shocks, in the form of large recessions near retirement age: it is possible that these older workers may have been more adversely affected by the 1980-1982 recessions. The last panel of each picture overlaps the profiles of the 9 cohorts. While the hours profiles are almost perfectly overlapped (with some increase in hours for the two of the youngest cohorts⁷) the wage and earnings profiles shift upward for younger cohorts. This, of course, reflects productivity growth.

Figure 5: Wage profile, all cohorts

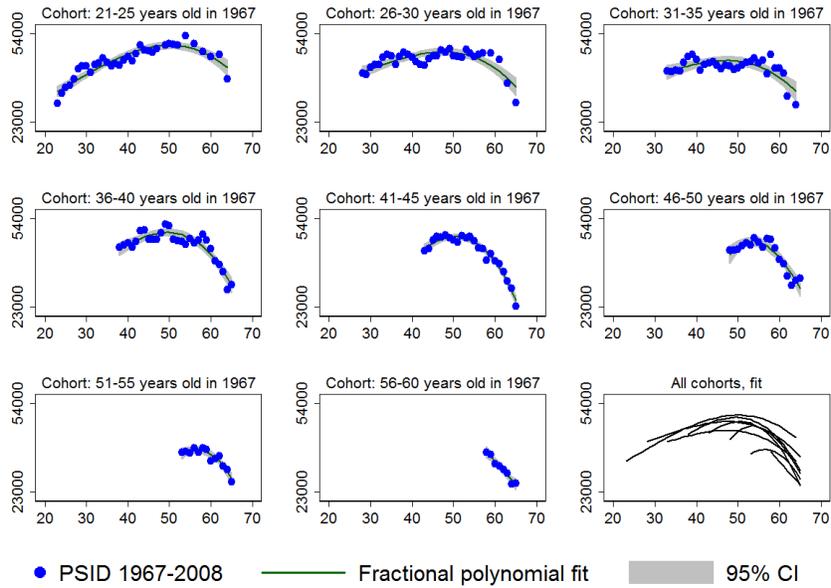


Notes: The figure reports the wage profiles of all the cohorts who were in the labor market in 1967. The interpolating line is the best second-order fractional polynomial.

⁶These profiles are truncated at age 65 (i.e., when individuals in the cohort are between 63 and 67 years old): in our data set the group of workers in a specific cohort after that age is too small to provide reliable estimates.

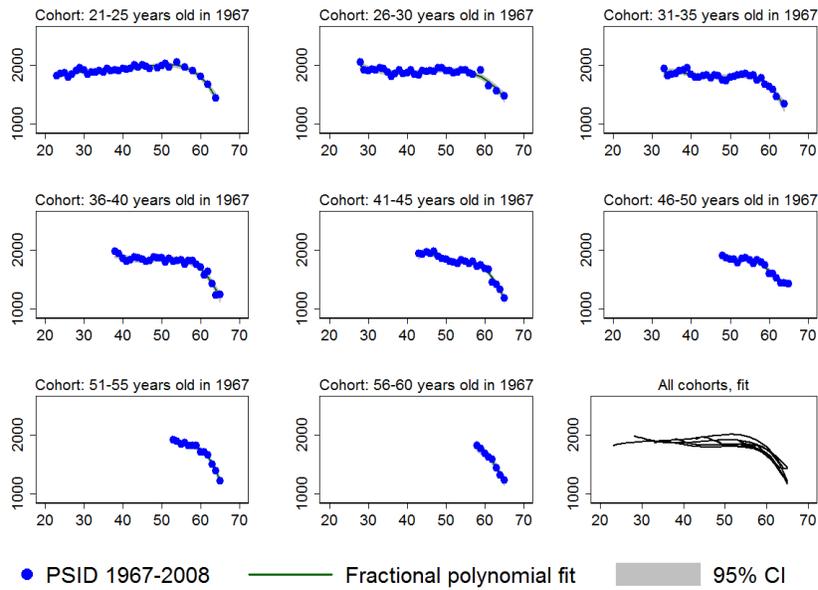
⁷This is consistent with the evidence reported in papers that analyze time use trends across different cohorts. See, for instance, figures 21 and 22 in Attanasio et al. (2012)

Figure 6: Earnings profile, all cohorts



Notes: The figure reports the earnings profiles of all the cohorts who were in the labor market in 1967. The interpolating line is the best second-order fractional polynomial.

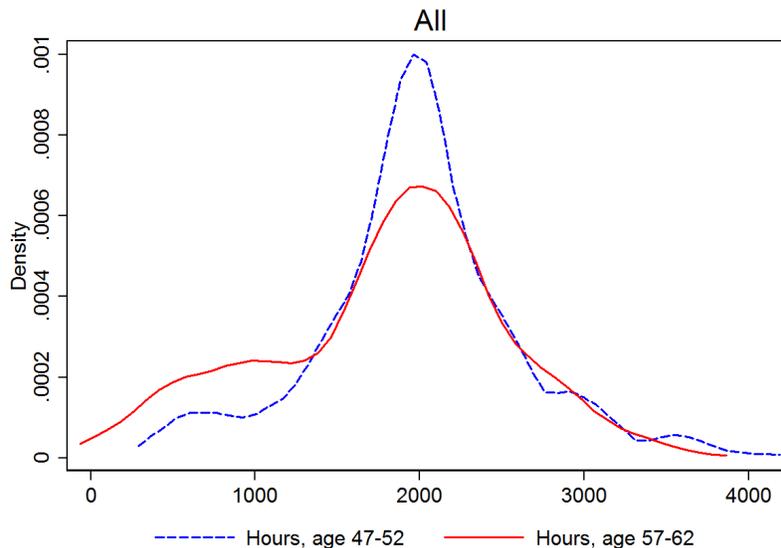
Figure 7: Hours profile, all cohorts



Notes: The figure reports the hours profiles of all the cohorts who were in the labor market in 1967. The interpolating line is the best fourth-order fractional polynomial.

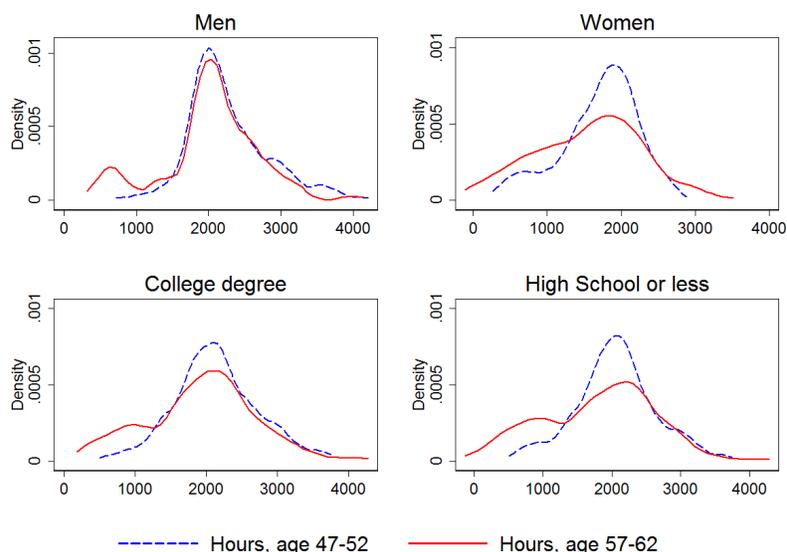
We next investigate how workers reduce hours along the life cycle, a process we refer to as “pre-retirement”. The main margins of adjustment are likely two: switching from full-time to part-time and reducing overtime hours. Figure 8 shows that this is exactly what happens for the 23-years old cohort. The figure shows kernel density estimates of the distribution of hours when the cohort is between 47 and 52 year old (that is, at the peak of hours) and ten years later, when the cohort is between 57 and 62 year old. The figure shows that the distribution at the early stage is concentrated around the full-time level of hours, and has fat tails, indicating that a non-negligible portion of workers are employed part-time or do overtime work. However, ten years later there is a large decrease in the density at full-time, the right tail is thinner, and the distribution looks bimodal—there is a second peak, likely around part-time hours. Therefore, ignoring for a moment selection into employment (an issue we tackle below), this picture suggests that many senior workers switch to part-time mode and reduce overtime work. This pattern is confirmed when we decompose the distribution by gender and education, as illustrated in Figure 9. Given the possibly relevant role of moves from full- to part-time during the pre-retirement years, the existence of a negative part-time premium on hourly wages documented by Aaronson and French (2004) suggests that pre-retirement adjustments may lead to underestimating productivity growth late in the working life. This reinforces the pattern we are documenting.

Figure 8: Distribution of hours



Notes: The figure shows the kernel density estimate (Epanechnikov kernel) of hours worked on the market by individuals who were between 21 and 25 years old in 1967 at two different points in time: when they are between 47 and 52 years old and ten years later, when they are between 57 and 62 years old.

Figure 9: Distribution of hours, by gender and education



Notes: The figure shows, by gender and education, the kernel density estimate (Epanechnikov kernel) of hours worked on the market by individuals who were between 21 and 25 years old in 1967 at two different points in time: when they are between 47 and 52 years old and ten years later, when they are between 57 and 62 years old.

To sum up the finding documented in this section, the empirical regularity that emerges from our analysis shows that the life cycle earnings profile is increasing at early ages and is declining after age 50. The wage profile, however, increases over the life cycle with no clear tendency to decline before age 65 (with the exception of 2 cohorts out of 8). Women and college graduates are a possible exception, with some negative wage growth after age 60. Therefore, earnings fall because of a reduction in hours of market work along the intensive margin. Such a reduction is sizeable (about 1/3, on average) and constitutes a smooth transition into retirement via reduction of overtime work and switch to a part-time work mode.

A concern is that this pattern may be an artifact of staggered retirement, selection out of employment, and attrition. We next address these three threats to the picture we are rendering.

3.2 Is this pattern an artifact?

So far we have ignored three important processes that may give the false appearance of a smoothly declining hours profile and a non-declining wage profile. They are all manifestations of the selection problem. That is, they all have to do with the crucial fact that we observe wages only for individuals who are interviewed and work in a given year and, for the sub-sample of these that leave employment during that year, we measure hours only up to the day they stop working.

First, the pattern of a smooth decline of the aggregate hours profile in spite of a non-declining wage profile could result from staggered, abrupt retirement at points different from the end of the year. The following example illustrates what this means. Suppose that workers either work full-time at a given hourly wage or do not work at all, and that individuals can retire only on July

1st of any year. In this case the true individual hours profile is flat at 2000 hours until the year before retirement, has an intermediate point at 1000 hours in the retirement year, and is flat at zero hours afterwards. The wage profile, instead, is flat. Now suppose that 10% of workers retire at age 57, 20% of those remaining retire at 58, 30% at 59, and so on, until all remaining workers retire at 65. This would induce a smoothly declining aggregate hours profile (in correspondence to non-declining wages) while the true underlying individual profiles are discontinuous.

Second, our sample changes over time because of selection into and out of employment. The selection problem in this context is easy to see: as [Figure 2](#) illustrates, around age 50 workers begin self-selecting into retirement. This process is very likely non-random with respect to wages, earnings, and hours. Consider the wage profile and suppose that low-wage workers retire first, or that it is precisely a negative wage shock that induces workers to retire. Then our sample would change in a way that leads to overestimating wage growth late in the working life: we would render a non-declining selected profile while the true underlying profile is actually declining.

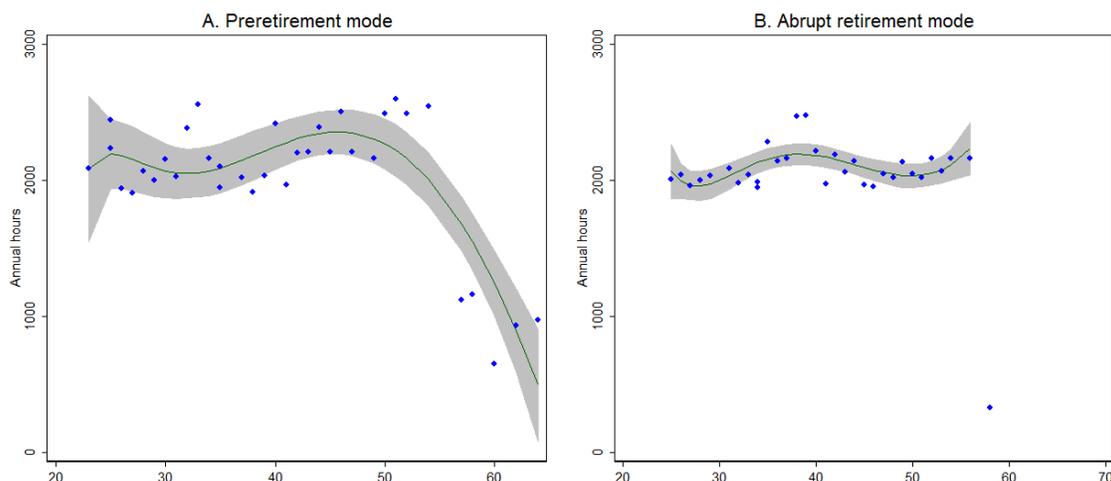
Third, our sample changes over time because of attrition. Focusing on specific cohorts in a panel that included about 3000 households (the SRC sample) at its outset in 1967 means working with relatively few individuals, as attested by the first line of [Table 1](#). The fact that we follow them for up to 40 years makes things worse because the attrition rate is substantial. Consider the 23 years old cohort. In 1967 this comprised about 500 individuals. At the last data point, in 2008, little more than 300.

Is the empirical regularity we have isolated an artifact of these three processes? The answer is no. In what follows we carefully motivate this answer.

3.2.1 Pre-retirement vs. abrupt retirement during the year

[Rogerson and Wallenius \(2011\)](#) claim on the basis of PSID data that “abrupt retirement” (i.e., moving from, say, full-time hours to zero hours with an intermediate level of hours in the year a worker retires) is by far the most important retirement mode in the U.S.—more than 70% of workers. This, apparently, contrasts with the preretirement pattern we have documented above. To understand the difference between these two retirement modes, consider [Figure 10](#). These pictures reproduce the *individual* profiles of two workers in the 23 years old cohort. Panel A refers to a worker (a man) who is retiring along a preretirement pattern. He works at a full-time level (i.e., about 2000 hours per year) until age 55, and then switches to a part-time level (i.e., about 1000 hours), with a declining trend thereafter. Panel B, instead, refers to a worker (another man) who moves abruptly from full-time work to zero hours at age 58, through an intermediate step of 330 hours—these are presumably worked during the first two months of the retirement year, which is year 2000. There are 78 individuals in cohort 23 for whom we can inspect individual hours profile like the two we have just illustrated (i.e., 78 individuals in the cohort who work every year from 1967 until 1997 and who remain in the survey and may move to zero hours afterwards.). Of these 78 profiles, 26 look like case A in [Figure 10](#), 28 look like case B, 15 are non-declining, and 9 are hard to interpret. This indicates that pre-retirement is more important a mode than the analysis in [Rogerson and Wallenius \(2011\)](#) suggests, although abrupt retirement is important too. The contrast with [Rogerson and Wallenius \(2011\)](#) is in fact only apparent, because they restrict the sample to individuals who worked at least 1750 hours after age 56. This selection procedure removes workers who are already on a preretirement path like the worker in panel A of [Figure 10](#).

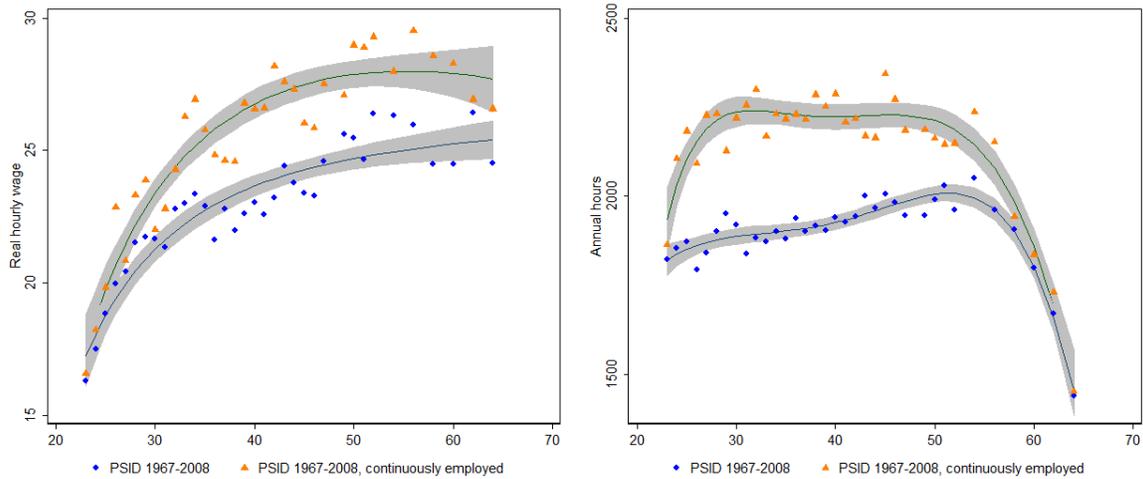
Figure 10: Retirement modes: abrupt retirement vs. preretirement path



Notes: The figure shows two individual profiles of workers who were between 21 and 25 years old in 1967. Worker in panel A moves into retirement through a preretirement stage at part-time hours. Worker in panel B moves into retirement by switching abruptly to zero hours. The interpolating line is the best fourth-order fractional polynomial.

However, given the relevance of abrupt retirement one may still worry that the pattern of a smoothly declining aggregate hours profile and non-declining wage profile may be an artifact of extensive margin movements, as the example given at the beginning of this section illustrates. There is a simple way to demonstrate that this is not the case: shut off the extensive margin entirely. This can be done by considering the group of individuals in the cohort who are employed every year from 1967 to 2008. These are referred to as the “continuously employed.” For these individuals, there are no extensive margin movements at the annual frequency in the period of observation. Their number is of course tiny in a single cohort: in the 23 years old cohort they are just 38. Despite such small sample size, if the shape of the wage and hours profiles of this group is not different from that of the larger sample then we have evidence that the pattern we have isolated is not an artifact of self-selection out of employment towards the end of the life cycle. Figure 15 illustrates that this is indeed the case. The levels are of course different, and reflect the different labor market attachment in the two groups. Those more attached—i.e., those who do not experience unemployment nor leave the labor force for more than 11 months—have higher wages and work more hours. But the shape of the profiles is the same in the two groups. In particular, those in the cohort who are not yet retired by 2008 are clearly moving along a pre-retirement path despite no clear downward trend in wages, similar to the pattern in the full sample.

Figure 11: Wage and hours profiles: cohort 23, full sample vs. continuously employed



Notes: The figure reports wage and hours profiles of individuals who were between 21 and 25 years old in 1967 (cohort 23), for the full sample and for the sample of those who are employed every year from 1967 to 2008 (i.e., the “continuously employed”). The interpolating line is the best second-order fractional polynomial for wages and the best fourth-order fractional polynomial for hours.

3.2.2 Attrition

To show that attrition does not affect the pattern, we first provide two informal checks based on a comparison with the Current Population Survey (CPS) and with the no-attrition group in the PSID, respectively, and then perform formal tests. The first informal check consists of constructing equivalent life cycle profiles for the same cohorts from the Outgoing Rotation Groups of the CPS. Attrition is not an issue in the CPS: finding a similar pattern would suggest it is not an issue in the PSID either, as far as the shape of life cycle profiles is concerned.⁸ To perform such a check we use the CPS data set constructed by Heathcote et al. (2010). These data result from applying selection criteria similar to those we applied to the PSID, and cover the 1967–2005 period. Because the sample size is now much larger we can expand (relative to Figure 5, Figure 6, and Figure 7) the upper limit of the age interval to 70 years without having to worry about unreliable estimates. Figure 12 illustrates the resulting wage profile. This picture confirms the pattern found in the PSID, at least for the four youngest cohorts. The wage profiles of the four oldest cohorts look hump-shaped. Wage growth becomes negative shortly before 60 for cohorts 43 and 48, consistently with the PSID, but not before 65 for the two oldest cohorts. Figure 13, and Figure 14 illustrate the earnings and hours profiles. Except for the 23 years old cohort (which we observe only until age 61 in this CPS data set) the CPS also replicates the shape of earnings and hours profiles in the PSID.⁹

⁸Using CPS data also provides a useful check with respect to small sample size in the PSID.

⁹There is an issue about the *level* of hours, which turn out to be slightly higher in the CPS than in the PSID—between about 30 and 120 hours. This difference may reflect the different wording of the survey questions (hours *usually* worked per week in the CPS versus hours *actually* worked per week in the PSID).

Figure 12: Wage profile in the CPS, all cohorts

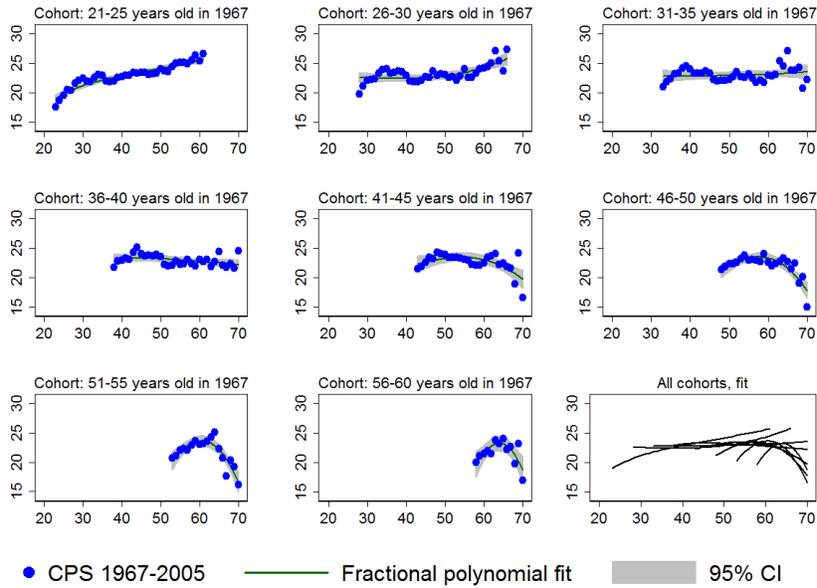


Figure 13: Earnings profile in the CPS, all cohorts

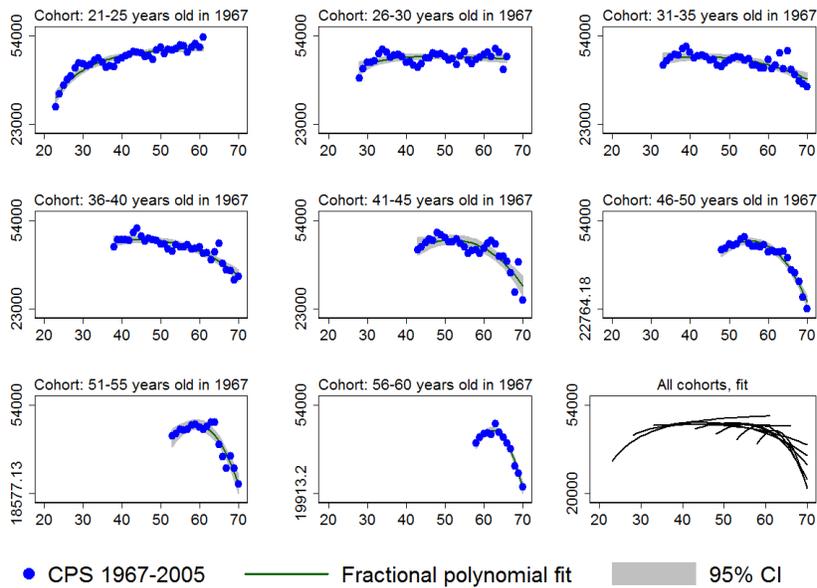
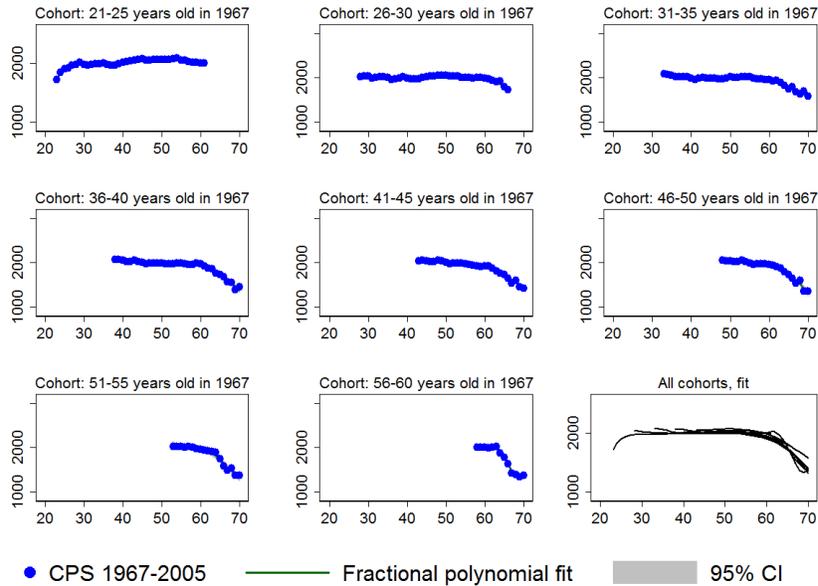


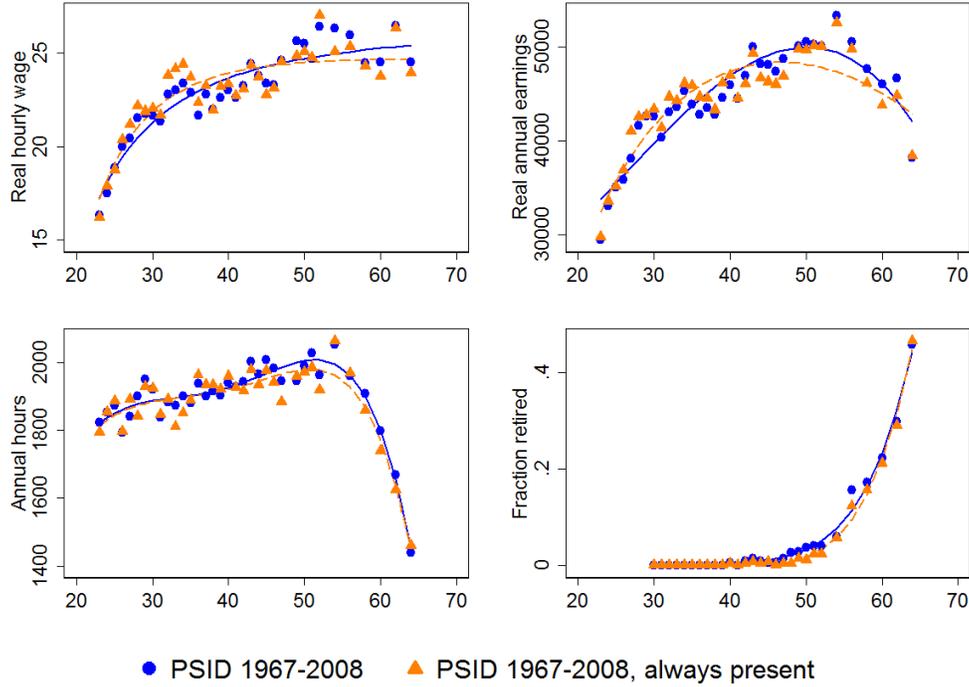
Figure 14: Hours profile in the CPS, all cohorts



Notes: Figure 12, Figure 13, and Figure 14 are the CPS analog of Figure 5, Figure 6, and Figure 7. They illustrate the hours, earnings, and hours profiles of all the cohorts who were in the labor market in 1967, from the March CPS. The interpolating line is the best second-order fractional polynomial, except for hours where a fourth-order polynomial is used.

The second informal check consists of comparing the profiles of the full, but selected sample (because of attrition), with the profiles of the group of individuals in the cohort who are present in the survey every year from 1967 to 2008. We refer to these individuals as the “always present”, and there are about 270 such individuals in the cohort. If the profiles of these two groups are different, then we have evidence that attrition is selective with respect to the three variables of interest. If instead they do not differ, then we have evidence that, with respect to these variables, individuals leave the survey at random. This comparison is illustrated in Figure 15. We do not observe systematic differences in this picture.

Figure 15: Life cycle profiles: cohort 23, full sample vs. no-attrition group



Notes: The figure reports wage, hours, earnings, and retirement profiles of individuals who were between 21 and 25 years old in 1967 (cohort 23), for the full sample and for the sample of individuals who are present in the survey every year from 1967 to 2008 (i.e., the no-attrition group). The interpolating line is the best second-order fractional polynomial, except for hours and the fraction retired where a fourth-order polynomial is used.

Finally, we test for the presence of attrition bias formally using a parametric, Heckman-type two-step procedure. We follow [Wooldridge \(2010\)](#), chapter 19.9.3, who suggests first-differencing the data and working with the usual pair of equations, one for outcome and one for selection:

$$\Delta y_{it} = \Delta \mathbf{x}_{it} \boldsymbol{\beta} + \Delta \varepsilon_{it}, \quad (1)$$

$$s_{it} = \mathbb{I}[\mathbf{q}_{it} \boldsymbol{\gamma} + e_{it} > 0], \quad (2)$$

where y_{it} is the variable of interest (wage rate, earnings, or hours), \mathbf{x}_{it} is a vector of time-varying covariates, $\mathbb{I}[\cdot]$ is the indicator function, s_{it} is the selection indicator, keeping track of whether individual i is present ($s_{it} = 1$) in the survey or not ($s_{it} = 0$), \mathbf{q}_{it} is a vector containing variables that may correlate with the attrition process and that are observable for at least one period after attrition, and e_{it} is an unobservable term with a standardized normal distribution.¹⁰

¹⁰We include in \mathbf{x}_{it} the following variables: age, whether married, number of children living in the household, whether college degree or more, non-labor income, whether lives with an employed partner, whether unemployed, whether disabled, and whether relocated during the year. As for \mathbf{q}_{it} , we include in this vector the first lag of \mathbf{x} (i.e., \mathbf{x}_{it-1}) as well as six permanent characteristics: whether ever obtained a college degree, maximum number of children ever lived in the household, whether white, whether ever been unemployed, whether ever reported a disability, whether

The test is based on three assumptions. First, attrition is an absorbing state, i.e. $s_{it+j} = 0$ for all $j > 0$ when $s_{it} = 0$. This is straightforward to impose on the data. Second, the conditional mean of $\Delta \varepsilon_{it}$ is linear: $\mathbb{E}[\Delta \varepsilon_{it} | e_{it}, s_{it}] = \rho_t e_{it}$. Third, both $\Delta \varepsilon_{it}$ and e_{it} are independent of \mathbf{q}_{it} . Then it follows from these assumptions that

$$\mathbb{E}[\Delta y_{it} | \Delta \mathbf{x}_{it}, \mathbf{q}_{it}, s_{it-1} = 1] = \Delta \mathbf{x}_{it} \beta + \rho_t \lambda(\mathbf{q}_{it} \gamma), \quad (3)$$

where $\lambda(\cdot)$ is the inverse Mills ratio. The latter can be estimated for each period via pooled Probit and interacted with time dummies to allow for a time-varying ρ in Eq. (3). This equation is then estimated via pooled two-stage least squares (clustering standard errors at the individual level and allowing for heteroskedasticity), using a subset of \mathbf{q}_{it} as an instrument for $\Delta \mathbf{x}_{it}$.¹¹ Then, the null hypothesis of no attrition bias is equivalent to the joint hypothesis that $\rho_t = 0$ in each period, which is straightforward to test. Since Eq. (3) is in first-difference, we base the test on the annual portion of the PSID (1967-1996). We cannot reject the null for either wages (P-value: 0.15), earnings (P-value: 0.14), or hours (P-value: 0.24).¹²

Based on this informal and formal evidence we conclude that attrition, although non-negligible, is unlikely to have serious consequences for inferring the shape of life cycle profiles in our data. This conclusion is consistent with what is known about attrition in the PSID at large, at least until the end of the 1990s: Fitzgerald et al. (1998) and Lillard and Panis (1998) conclude that despite being selective, attrition has not undermined the representativeness of the PSID and that ignoring it when estimating the dynamics of variables of interest has probably only mild consequences.

3.2.3 Self-selection

We finally address selection. In Section 3.2.1 we have already provided an informal check based on a comparison with the group of continuously employed individuals in the PSID. This is conceptually equivalent to the second informal check on attrition bias illustrated in Figure 15. That is, we compared the profiles of the full, but selected sample (because of self-selection into employment), with the profiles of continuously employed workers. Figure 11 suggests that selection affects the levels but not the shape of the profiles. Here we provide a more formal check by producing a selection-corrected wage profile. We use the procedure for selection-correction in panel data developed by Wooldridge (1995),¹³ and we keep focusing on the 23 years old cohort for illustrative purposes. The procedure is an extension of standard two-stage methods valid for cross sectional data such as “Tobit” and “Heckit”. Specifically, we predict wages of individuals who are not employed in a given year—but for whom we observe covariates in that year as well as in other periods—using the coefficients of a wage regression consistently estimated. Consistency is

ever relocated.

¹¹We use \mathbf{x}_{it-1} as instruments. The rationale for using lags of levels to instrument first differences is illustrated in Arellano (1989).

¹²The following, much simpler test confirms this result: estimate in levels using the fixed-effects estimator and including s_{it+1} as an additional right-hand variable; then test the null hypothesis that the coefficient on s_{it+1} is zero. The intuition is that if the attrition process is random with respect to wages, earnings, and hours, then the attrition indicator in the next period should not help predicting any of these variables.

¹³Correcting the hours profile (and so the earnings profile as well) is harder with the procedure we use below: it’s very hard to maintain credibly that the unobservable determinants of hours are uncorrelated with the unobservable determinants of the decision to supply a positive amount of hours.

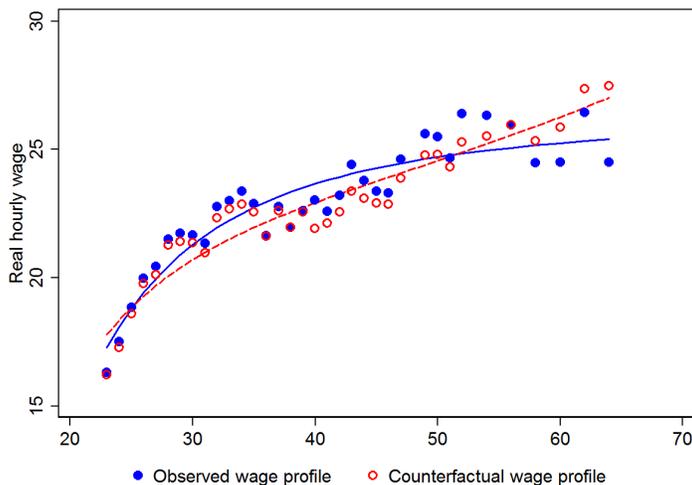
achieved as follows. As in the case of the attrition test illustrated above, the model consists of a pair of equations, one for outcome (the wage rate) and one for selection into employment—the latter now takes the Tobit form because selection is determined by a corner solution in the labor supply problem.

$$w_{it} = \mathbf{x}_{it}\boldsymbol{\beta} + \theta_i + \varepsilon_{it}, \quad (4)$$

$$h_{it} = \max(0, \mathbf{X}_i\boldsymbol{\gamma} + v_{it}), \quad (5)$$

where w_{it} is the wage rate, h_{it} hours, \mathbf{x}_{it} time-varying covariates, \mathbf{X}_i contains vectors \mathbf{x}_{it} at all leads and lags (in practice, feasibility requires us to replace these with the corresponding longitudinal averages), and v_{it} is normally distributed and independent of \mathbf{X}_i . It is understood that w_{it} is observed if and only if $h_{it} > 0$. Under the mean conditional independence assumption $\mathbb{E}[\varepsilon_{it} | \mathbf{X}_i, \theta_i, h_{it} > 0] = 0$ $\boldsymbol{\beta}$ can be consistently estimated via pooled OLS after including in Eq. (4) the Tobit residuals from Eq. (5) estimated for each period as additional regressors. Using this procedure, we obtain a counterfactual wage profile—i.e., the profile we would have observed under the maintained identifying assumptions had those workers who retired remained employed. The observed and counterfactual wage profiles for the 23-years old cohort are reported in Figure 16.

Figure 16: Selection-corrected wage profile



Notes: The figure reports the observed average wage profile of individuals who were between 21 and 25 years old in 1967 (cohort 23), and the corresponding profile after prediction of missing wages through a selection-correction procedure. The interpolating line is the best second-order fractional polynomial.

This figure shows that the non-declining wage profile we have documented is not an artifact of selection. During the early part of the working life there seems to be “negative selection” out of employment. That is, the counterfactual wage profile is below the observed one, suggesting that individuals who do not work would earn a lower hourly wage relative to those who work. However, towards the end of the working life there is “positive selection” into retirement: workers with higher wages leave employment first. This evidence is consistent with the findings of Casanova

(2010), who studies the wage process of older workers using panel data from the Health and Retirement Study. She, too, finds evidence of positive selection into retirement. This evidence is, however, inconsistent with French (2005), who finds that selection-corrected wages are 7% and 11% lower at age 62 and 65, respectively, relative to the uncorrected ones. However, in addition to the differences we have already pointed out in Section 2, French uses a very different correction procedure, based on assuming that selection bias in the PSID is the same as in his model, and comparing the wages of workers and non-workers in the model. In any case, as we show next, even negative wage growth of the order of 10% at the end of the working life can hardly be replicated in a benchmark life cycle model unless the intensive margin elasticity of labor supply is calibrated to an implausibly large value.

4 A benchmark life cycle model

We now perform a very simple quantitative exercise aimed at understanding how accurately a basic life cycle model explains the empirical regularities we have documented so far. The model is stylized and intentionally simple. Not surprisingly, in its basic version it is unable to replicate the pattern of falling hours when calibrated with reasonable parameters and the empirical wage profile we have estimated. We explore which variants of the benchmark model, among those more frequently considered in the macroeconomic literature, are able to reconcile facts and theory.

In this quantitative exercise we stick to the fiction of a representative worker who makes intensive margin choices only. We do not attempt to replicate the important heterogeneity along gender and education we documented above, nor extensive margin decisions. The reason is that the model we want to use is too simple to capture such heterogeneity or retirement behavior in any meaningful way. Furthermore, the analysis is built on a single cohort, so time and age are indistinguishable in what follows. The representative worker enters the labor market at time 0 and dies at time T . For the moment, assume there is no uncertainty. The problem is to choose sequences of consumption, $\{c_t\}$, labor supply, $\{h_t\}$, and assets, $\{a_{t+1}\}$, given an exogenous wage profile, a sequence of real interest rates, $\{r_t\}$, and initial and terminal assets, to maximize utility over the life cycle:

$$\begin{aligned} \max_{\{c_t, h_t, a_{t+1}\}} \quad & \sum_{t=0}^T \beta^t \left[v(c_t) - \gamma \frac{h_t^{1+\frac{1}{\varepsilon}}}{1+\frac{1}{\varepsilon}} \right] \\ \text{s.t.} \quad & c_t + a_{t+1} \leq a_t(1+r_t) + w_t h_t \\ & h_t \leq \bar{h} \\ & a_0 = a \\ & a_{T+1} = 0 \end{aligned}$$

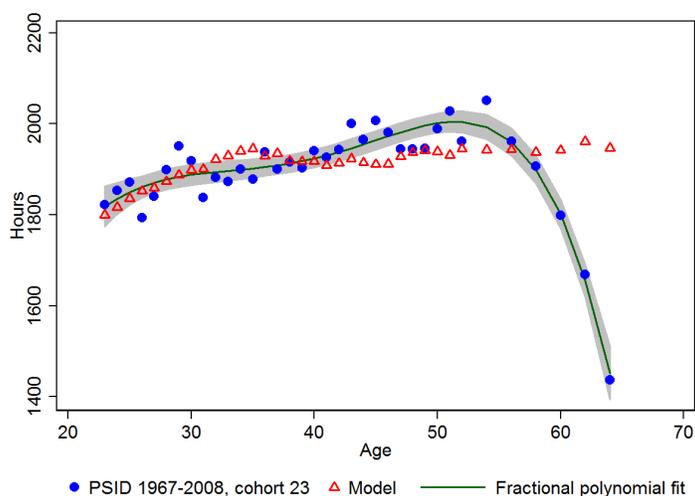
where β is the discount factor, $v(\cdot)$ a quasi-concave increasing function, ε the intertemporal elasticity of substitution of labor, and \bar{h} the time endowment. From the first-order conditions, the dynamics of hours in this model, for $t < T$, is given by:

$$\frac{h_{t+1}}{h_t} = \left[\frac{1}{\beta(1+r_{t+1})} \frac{w_{t+1}}{w_t} \right]^\varepsilon. \quad (6)$$

The term $1/\beta(1+r_{t+1})$ is the ratio between the multipliers on the budget constraint in two adjacent periods and captures the income effect: the worker is moving along a known wage profile and is accumulating assets optimally. The term w_{t+1}/w_t , on the other hand, captures the substitution effect. If the income effect is not too strong (i.e. $\beta(1+r_{t+1}) \leq 1$) then hours rise if wages do, or fall if wages fall.

Eq. (6) can be used to predict the hours profile of the 23 years old cohort in the PSID. We set $\beta = 0.98$, and $\varepsilon = 0.12$. This is the point estimate of *intensive* margin responses in the SRC subsample of the PSID produced by Fiorito and Zanella (2012). Finally, we estimate the series of r_t as the difference between the annualized 6-Month Treasury Bill interest rate and the CPI-U inflation rate from 1967 to 2008. We set h_0 to 1800 hours per year and simulate hours forward based on the actual PSID wage profile. The result of this simulation is shown in Figure 17. The model reproduces reasonably well the empirical wage profile until age 50 and very badly afterwards, because of the inability to generate the sharp fall in hours observed after this age.¹⁴ Given a non-declining wage profile, the tracking between hours and wages prevents the model from generating such a fall.

Figure 17: Actual and simulated hours profile



Notes: The figure reports the observed and simulated (from Eq. (6)) wage profiles of individuals in cohort 23.

Some departure from this benchmark model is needed in order to match the data. We consider four possible departures that can potentially reconcile the benchmark model and the empirical profiles: (1) uncertainty about future wages; (2) borrowing constraints; (3) human capital investment via on-the-job training; (4) time-varying disutility of labor. In all these cases hours can potentially fall even if wages do not. We do not consider the incentives embedded in the Social Security system. While it is well known that these affect labor supply, they primarily affect the extensive margin. And when they do affect the intensive margin—e.g., by giving incentives to work less hours—they do not do so until age 62.

¹⁴Gomme et al. (2004) also find that a standard life cycle model does not perform well when trying to explain the hours of older workers.

4.1 Uncertainty

As illustrated by [Low \(2005\)](#), if one introduces wage uncertainty in the benchmark model, then it is possible to have declining hours in spite of an increasing wage profile. That is, it is possible to observe high labor supply early in life when wages are low, and low labor supply later in life when wages are high. Uncertainty breaks the close tracking between hours and wages. The intuition is straightforward: when the wage rate is uncertain, workers are subject to a precautionary effect and delay leisure until the late stages of the life cycle. The point can be illustrated analytically as follows. Without uncertainty, the growth rate of hours, from [Eq. \(6\)](#), is approximately equal to:

$$\Delta \ln h_{t+1} = -\varepsilon \ln \beta (1 + r_{t+1}) + \varepsilon \Delta \ln w_{t+1}, \quad (7)$$

where $\Delta \ln h_{t+1} = \ln h_{t+1} - \ln h_t$, and similarly for $\Delta \ln w_{t+1}$. If the wage rate is uncertain, then the first-order conditions do not boil down to anything like [Eq. \(6\)](#). Instead, they boil down to:

$$\frac{h_t^{\frac{1}{\varepsilon}}}{w_t} = \beta (1 + r_t) \mathbb{E}_t \left[\frac{h_{t+1}^{\frac{1}{\varepsilon}}}{w_{t+1}} \right]. \quad (8)$$

Taking a second-order expansion of $\frac{h_{t+1}^{\frac{1}{\varepsilon}}}{w_{t+1}}$ around h_t, w_t and rearranging, we obtain the analog of [Eq. \(7\)](#) under uncertainty:

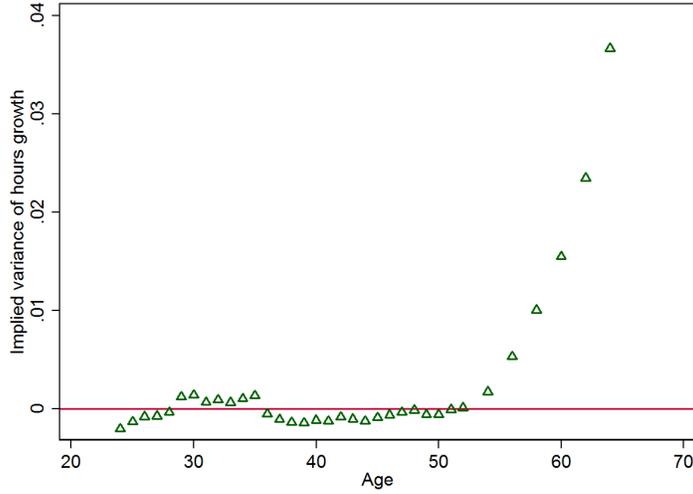
$$\mathbb{E}_t[\Delta \ln h_{t+1}] = -\varepsilon \ln \beta (1 + r_{t+1}) + \varepsilon \mathbb{E}_t[\Delta \ln w_{t+1}] - \frac{1 - \varepsilon}{2\varepsilon} \mathbb{E}_t[\Delta \ln h_{t+1}^2] - \varepsilon \mathbb{E}_t[\Delta \ln w_{t+1}^2]. \quad (9)$$

The variance of the wage rate and the consequent variance of hours (i.e., the last two terms) counterbalance the effect of positive wage growth on the hours profile. If the variance increases late in the life cycle, then hours may fall even if the wage rate does not. However, it is obvious that such variance should increase substantially. To illustrate, we solve [Eq. \(9\)](#) for $\mathbb{E}_t[\Delta \ln h_{t+1}^2]$ after replacing the left-hand side with $\Delta \ln h_{t+1}$ and recover the implied series of the variance of the growth rate of hours under the same calibrated parameters.¹⁵ Given the variance of wages, such variance in the series for hours growth allows the model to perfectly replicate the PSID trend. This series is illustrated in [Figure 18](#). Even ignoring the fact that the model actually requires *negative* variance of the growth rate of hours over most of the early part of the life cycle,¹⁶ it needs the variance to grow sharply after age 50 to induce the observed fall in hours. The benchmark model with uncertainty, therefore, raises the question of why the variance of hours (or the variance of wages, or both) increase sharply after this age. [Figure 18](#) implies an implausible degree of non-stationarity late in the life cycle. Notice that this conclusion does not depend on the assumption of intratemporal separability between consumption and leisure. Such an assumption, as [Low \(2005\)](#) makes clear, has consequences for the shape of the consumption profile, which is not the focus of this paper.

¹⁵Via the formula $\mathbb{V}_t[\Delta \ln h_{t+1}] = \mathbb{E}_t[\Delta \ln h_{t+1}^2] - \mathbb{E}_t[\Delta \ln h_{t+1}]^2$.

¹⁶The reason is that the benchmark model without uncertainty reproduces the trend in hours reasonably well until age 50. With uncertainty, the variance of the wage profile reduces hours growth. Therefore, one would need negative variance of hours to keep hours growing at the rate they would grow without uncertainty.

Figure 18: Implied variance of the hours growth rate



4.2 Borrowing constraints

Erosa et al. (2011) are able to replicate the fall in hours worked in the PSID by calibrating a model with both uncertainty and a borrowing constraint. The constraint takes the form of a zero lower bound on assets. Workers engage in precautionary savings and work more when young (more than in Low (2005), where there is no borrowing constraint) even if their future wage rate will be higher. Later in the life cycle, the stock of assets is so large that even small negative wage shocks induce workers to retire. Since the time period in their model is a third of a year, retirement generates a decline in hours on a one-year basis. Notice that this explanation relies on an active extensive margin. In particular, the argument implies that a declining hours profile of workers is an artifact of staggered abrupt retirement, in the sense discussed at the end of section 3.1. However, we have documented there that this is not the case in the PSID: pre-retirement transitions to fewer hours are an important phenomenon, empirically.

It is nonetheless of interest to understand whether a borrowing constraint improves the empirical performance of the benchmark model along the intensive margin of labor supply. To isolate the role of such a constraint, we assume that there is no uncertainty. Adding inequality $a_{t+1} \geq 0$ to the constraint set and denoting by μ_t the associated multiplier, the Euler equation becomes:

$$\lambda_t = \beta(1 + r_{t+1})\lambda_{t+1} + \mu_t, \quad (10)$$

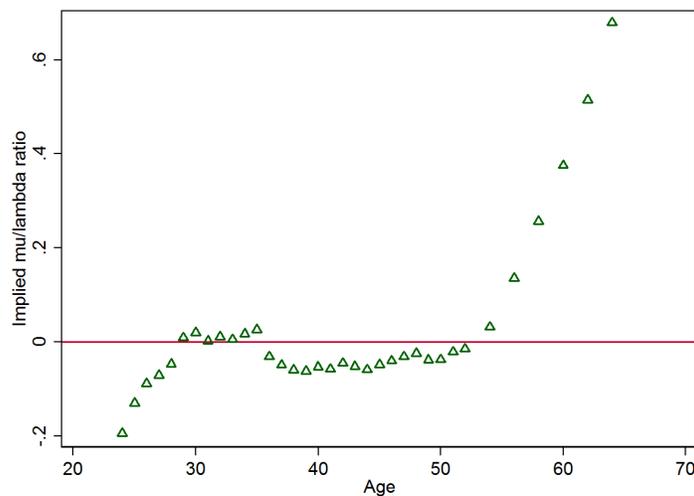
meaning that when the borrowing constraint is binding ($\mu_t > 0$) a worker cannot smooth consumption, so $\lambda_t > \beta(1 + r_{t+1})\lambda_{t+1}$ and consumption at time t is lower than it would have been otherwise. The dynamics of hours is now given by:

$$\frac{h_{t+1}}{h_t} = \left[\frac{1 - \frac{\mu_t}{\lambda_t}}{\beta(1 + r_{t+1})} \frac{w_{t+1}}{w_t} \right]^\varepsilon. \quad (11)$$

That is, hours grow at a slower (possibly negative) rate when the borrowing constraint is binding, and at the same rate as in the basic model when it is not. This may sound counterintuitive: we expect an individual for whom the borrowing constraint is active—i.e., for whom the marginal

utility of consumption is higher than in the unconstrained case and so an extra hour of work is more valuable—to work more. However, this is a statement about growth rates, not levels. A borrowing-constrained worker does, in fact, work more hours in the model when the constraint is binding compared to when it’s not. As a consequence, the hours growth rate is defined relative to a higher level and so is lower. From Eq. (11) we can back-out the series of $\frac{\mu_t}{\lambda_t}$ that allows the model to reproduce the data. This is illustrated in Figure 19. Not surprisingly, we incur in the same difficulty as in the previous case of uncertainty. The μ_t/λ_t ratio is required to be negative during most of the early part of the life cycle—which of course cannot be. Furthermore, to explain the hours decline along the intensive margin through this channel we would need a binding borrowing constraint after age 55. This is unlikely given the empirical fact that after that age workers dissave. The model of Erosa et al. (2011), in fact, reproduces the fall in hours per worker worse than the fall in hours per person.

Figure 19: Implied $\frac{\mu_t}{\lambda_t}$



4.3 On-the-job training

A falling hours profile can be generated out of a non-declining wage profile in the benchmark model if wages depend on past labor supply. This is what happens in the presence of accumulation of human capital via on-the-job training: in this case it is by working that an individual obtains wage growth over the life cycle. Such a mechanism breaks the tracking between hours and wages because it is no longer true that young workers want to work less than older workers in response to their initially low wages. However, the extent of the fall in hours late in the working life depends on how fast the effect of current hours of work on future earnings (that is, the return to on-the-job training) declines. To illustrate, assume that the current wage rate depends on past labor supply:

$$w_t = w_t(\mathbf{h}^{t-1}), \quad \text{where} \quad \mathbf{h}^{t-1} = \{h\}_{j=0}^{t-1}.$$

As illustrated by Keane (2011), the consequence is that the opportunity cost of time increases. The increase is equal to the present value of the cumulative effect of current labor supply on all

future wage rates. To see this, notice that now the first-order condition for hours is no longer $h_t^{\frac{1}{\varepsilon}} = v'(c_t)w_t$, but:

$$h_t^{\frac{1}{\varepsilon}} = v'(c_t)w_t + \sum_{j=t+1}^T \beta^{j-t} v'(c_j) h_j \frac{\partial w_j}{\partial h_t}, \quad (12)$$

where the second term on the right-hand side is the present value of the lifetime increase in earnings caused by an extra hour of work at time t . Eq. (6) becomes:

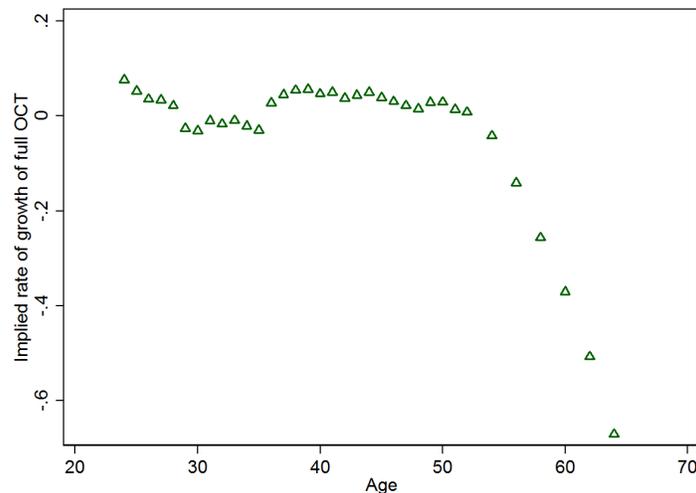
$$\frac{h_{t+1}}{h_t} = \left[\frac{1}{\beta(1+r_{t+1})} \frac{\tilde{w}_{t+1}}{\tilde{w}_t} \right]^{\varepsilon}, \quad (13)$$

where

$$\tilde{w}_t = w_t + \sum_{j=t+1}^T \frac{1}{\prod_{k=t+1}^j (1+r_k)} h_j \frac{\partial w_j}{\partial h_t}. \quad (14)$$

It is clear from these expressions that hours decline, even if the current wage rate does not, if the full opportunity cost of time \tilde{w}_t declines fast enough. We can infer from the data the growth rate of \tilde{w}_t that makes this variant of the model consistent with the empirical hours trend. The series is illustrated in Figure 20. Given that the wage profile is essentially flat after age 50, this variant of the benchmark model implies that in order to reproduce the data, the marginal return from on-the-job training (i.e., $\tilde{w}_t - w_t$) should decline by about 7% per year from age 52 to age 64. This number matches remarkably well the estimate of the decline in the rate of return provided by Kuruscu (2006). In that paper, the cumulative decline in the marginal return from on-the-job training between 30 and 45 years of experience (age 53 and 68 in our paper) is about 6% per year (see Figure 5 in Kuruscu (2006)).

Figure 20: Implied growth rate of full opportunity cost of time (OCT)



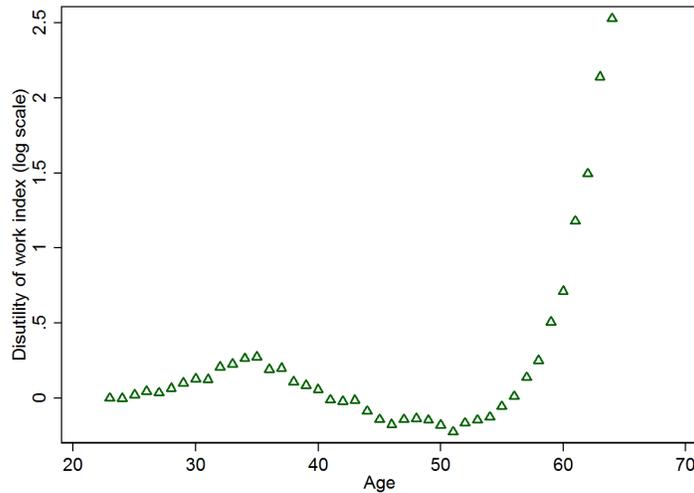
4.4 Increasing disutility of work/utility of leisure

Finally, predicted hours may fall in the benchmark model if we allow the disutility of labor/value of leisure to increase after age 50. If γ is not constant, then Eq. (6) becomes:

$$\frac{h_{t+1}}{h_t} = \left[\frac{1}{\beta(1+r_{t+1})} \frac{w_{t+1}}{w_t} \frac{\gamma_t}{\gamma_{t+1}} \right]^\varepsilon. \quad (15)$$

Clearly, if γ_t increases sufficiently fast after age 50 then hours fall even if wages are growing. For instance, when analyzing wage and hours profiles in the March CPS between 1977 and 1989 and finding shapes in line with those we report in Figure 12 and Figure 14, Card (1994) comments: “The life-cycle profiles of wages and hours are far from parallel. Of course this does not refute the life-cycle model, because tastes may vary systematically with age” (p. 56). Like before, we can ask how the taste parameter γ evolve in this model to reproduce the trend of the observed hours profile (i.e., back out the implied γ_t profile). The result is illustrated in Figure 21, using the log scale. This picture reveals an interesting pattern. γ_t increases until the mid 30s (by about 30% at the local maximum at age 35), then declines throughout one’s 40s (by more than 60%, relative to the peak at 35, at the minimum, at age 51), and then picks-up very quickly. That is, the disutility of working (or the value of alternatives to market work, or both) increases during the early years in the labor market, declines during the second half of prime age, and then increases fast. For junior workers, this may reflect adaptation and careers; for prime-age workers, work habits; for seniors, instead, factors such as health deterioration, stress, or a desire to work less and devote time to other activities may play an important role.

Figure 21: γ_t index



Health, however, explains little of the discrepancy between the actual and simulated profiles. Table 3 reports the results from a fixed-effects regression of hours on a quartic polynomial in age and the “disabled” dummy described in Section 3. The regression reveals that deteriorating health (as captured by this variable) induces a decline of only about 100 hours—less than 20% of the discrepancy between the data and the model at age 64 in cohort 23. Although this particular longitudinal measure of disability in the PSID may be a poor proxy for health conditions, the data

suggest that an increase in the marginal disutility of labor via health deterioration alone is unlikely to explain the fall in hours after age 50. This is consistent with French (2005), who concludes that “health status alone must have a small causal role in the decline in the number of hours worked by workers near retirement.” (p. 410)

Table 3: Effect of disability on hours of workers

	1	2
	hours	hours
<i>disabled</i>	-98.2* (23.8)	-102.1* (6.4)
<i>age</i>	384.1* (104.4)	-13.8 (10.7)
<i>age</i> ²	-1509.2* (391.4)	141.2* (38.1)
<i>age</i> ³	265.3* (63.1)	-21.0* (5.7)
<i>age</i> ⁴	-17.3* (3.7)	0.36 (0.31)
<i>Const.</i>	-1773.5 (1007.0)	1570.9* (106.5)
Cohort	23 years old	all
Observations	9427	143,156
Individuals	510	15,016

Notes: The table contains the results from a regression of hours on a fourth-degree polynomial in age and variable *disabled*. This is a dummy variable that takes value 1 if a person responded “yes” to the question “Do you have any physical or nervous condition that limits the type of work or the amount of work that you can do?”, and 0 otherwise. Only individuals reporting strictly positive hours and included in the regression. Standard errors in parenthesis. An asterisk denotes significance at the 1% confidence level or better

Another possibility is that the value of alternatives to market work increases during the pre-retirement years. For instance, in a companion paper to this one (Rupert and Zanella (2010)) we show that, with unchanging preferences over the life cycle, grandparenting goes a long way towards reconciling falling hours and non-declining wages. The reason is that the appearance of grandchildren increases the value of “leisure” during the preretirement years. It is interesting to note that in Figure 21 the value of alternatives to market work increases in those stages of the life cycle when there are young children in the household or the family. This endogenous process is observationally equivalent, in the benchmark model, to an exogenous increase in γ late in the life cycle. This example indicates how the evidence documented in this paper restricts quantitative macroeconomic models. Except for the dependence of γ on age and the wage schedule, this variant of the model is the discrete-time version of Rogerson and Wallenius (2009). They assume that the wage profile is hump-shaped. Their results can be obtained with a concave, non-declining wage profile and a J-shaped profile for the value of leisure similar to the one in Figure 21. Our results indicate that the latter is likely a better assumption.

5 Conclusions

We have documented that life-cycle wage profiles derived from the Panel Study of Income Dynamics show no tendency to decline until late into one's 60's. In contrast, hours begin to fall after age 50, as do earnings. These results challenge a common assumption in macroeconomic models, that the wage profile is hump-shaped and that it is its decline late in the working life that induces a reduction of labor supply along the intensive margin. A benchmark life cycle model is unable to replicate this pattern. The model and the data can be reconciled by departing from such benchmark. The mere presence of uncertainty and borrowing constraints seem unable, by themselves, to do so. However, we have shown that the life cycle model is able to replicate the empirical pattern under two circumstances, in particular. First, the benefits of accumulating human capital via on-the-job training are an important component of the opportunity cost of time and decline relatively fast after age 50—in line with previous empirical research. Second, the value of alternatives to work for pay increase fast after that same age. Research about this channel is more scant. While deteriorating health plays a role, it seems unlikely that it plays a decisive role. The possible desire by part of older workers to allocate more time to activities different from work for pay—perhaps because new “goods” become available—seems a promising direction for future research.

6 Appendix

In this appendix we report tables containing the numerical life cycle profiles illustrated in the paper, for the eight cohorts we have analyzed and for the whole sample. That is, the profiles illustrated in [Figure 5](#), [Figure 6](#), and [Figure 7](#). Numerical profiles by gender and education and available from the authors upon request.

We report both the raw profiles and the profiles smoothed using a fractional polynomial fit (degree 2 for wages and earnings, degree 4 for hours). The initial value normalized to 1 for cohort 23. For older cohorts, the initial value is normalized to the corresponding value, at that same age, of cohort 23. This way growth rates can be compared across cohorts. Furthermore, the initial smoothed value is set equal to the initial raw value. This way growth rates can be better compared across the raw and smoothed series as well.

Notice that average earnings are not equal, of course, to the product of average wages and average hours. Researchers interested in using the latter measure in particular applications should compute it out of the wage and hours tables.

Table 4: Wage profiles, PSID, cohorts 23–38

Age	Cohort 23		Cohort 28		Cohort 33		Cohort 38	
	Raw	Smooth	Raw	Smooth	Raw	Smooth	Raw	Smooth
23	1.000	1.000						
24	1.074	1.109						
25	1.157	1.153						
26	1.225	1.191						
27	1.254	1.225						
28	1.319	1.255	1.319	1.255				
29	1.333	1.282	1.329	1.279				
30	1.329	1.306	1.364	1.300				
31	1.310	1.327	1.381	1.320				
32	1.397	1.347	1.392	1.337				
33	1.411	1.365	1.423	1.354	1.411	1.365		
34	1.433	1.381	1.442	1.369	1.422	1.376		
35	1.404	1.395	1.438	1.383	1.458	1.388		
36	1.328	1.408	1.402	1.396	1.477	1.400		
37	1.397	1.421	1.466	1.408	1.488	1.411		
38	1.348	1.432	1.472	1.420	1.522	1.423	1.348	1.432
39	1.387	1.442	1.486	1.430	1.542	1.435	1.358	1.464
40	1.412	1.452	1.460	1.440	1.537	1.447	1.393	1.493
41	1.386	1.460	1.388	1.450	1.500	1.459	1.411	1.518
42	1.424	1.468	1.375	1.459	1.568	1.470	1.422	1.539
43	1.497	1.476	1.363	1.468	1.575	1.481	1.453	1.558
44	1.459	1.483	1.384	1.476	1.589	1.492	1.473	1.574
45	1.434	1.489	1.382	1.484	1.561	1.503	1.469	1.588
46	1.429	1.495	1.396	1.491	1.485	1.514	1.432	1.600
47	1.509	1.501	1.439	1.499	1.471	1.524	1.497	1.611
48	1.540	1.506	1.454	1.506	1.458	1.534	1.504	1.620
49	1.571	1.511	1.445	1.513	1.480	1.543	1.518	1.627
50	1.563	1.516	1.432	1.519	1.478	1.552	1.491	1.634
51	1.513	1.520	1.435	1.526	1.493	1.561	1.418	1.639
52	1.619	1.524	1.442	1.532	1.540	1.569	1.405	1.644
53			1.788	1.538	1.555	1.576	1.392	1.647
54	1.614	1.531	1.560	1.545	1.546	1.583	1.414	1.650
55			1.504	1.551	1.532	1.589	1.411	1.653
56	1.592	1.538	1.465	1.557	1.535	1.594	1.426	1.654
57			1.504	1.563	1.542	1.599	1.471	1.655
58	1.501	1.544			1.912	1.603	1.486	1.656
59			1.535	1.574	1.669	1.605	1.477	1.657
60	1.503	1.549			1.609	1.607	1.463	1.657
61			1.579	1.586	1.567	1.608	1.466	1.656
62	1.621	1.554			1.608	1.608	1.473	1.656
63			1.564	1.598			1.826	1.655
64	1.503	1.558			1.642	1.604	1.594	1.654
65			1.563	1.610			1.537	1.652

Table 5: Wage profiles, PSID, cohorts 43–58

Age	Cohort 43		Cohort 48		Cohort 53		Cohort 58	
	Raw	Smooth	Raw	Smooth	Raw	Smooth	Raw	Smooth
43	1.497	1.476						
44	1.505	1.546						
45	1.610	1.602						
46	1.694	1.647						
47	1.670	1.681						
48	1.726	1.706	1.540	1.506				
49	1.701	1.723	1.541	1.571				
50	1.737	1.733	1.583	1.626				
51	1.711	1.738	1.666	1.670				
52	1.790	1.737	1.727	1.704				
53	1.771	1.731	1.682	1.729	1.599	1.502		
54	1.772	1.721	1.727	1.746	1.614	1.531		
55	1.723	1.708	1.705	1.755	1.615	1.557		
56	1.679	1.692	1.732	1.757	1.701	1.580		
57	1.635	1.673	1.794	1.751	1.684	1.600		
58	1.580	1.652	1.819	1.739	1.778	1.618	1.501	1.544
59	1.655	1.629	1.765	1.720	1.809	1.635	1.420	1.495
60	1.543	1.605	1.728	1.696	1.681	1.649	1.438	1.465
61	1.523	1.579	1.646	1.667	1.662	1.662	1.402	1.455
62	1.691	1.551	1.577	1.632	1.776	1.673	1.447	1.466
63	1.566	1.523	1.424	1.592	1.777	1.683	1.423	1.498
64	1.558	1.494	1.465	1.548	1.893	1.691	1.492	1.554
65	1.390	1.464	1.654	1.499	1.761	1.699	1.592	1.634

Table 6: Earnings profiles, PSID, cohorts 23–38

Age	Cohort 23		Cohort 28		Cohort 33		Cohort 38	
	Raw	Smooth	Raw	Smooth	Raw	Smooth	Raw	Smooth
23	1.000	1.000						
24	1.122	1.175						
25	1.188	1.204						
26	1.217	1.233						
27	1.294	1.262						
28	1.412	1.291	1.412	1.291				
29	1.445	1.320	1.393	1.308				
30	1.446	1.348	1.486	1.324				
31	1.370	1.376	1.514	1.340				
32	1.461	1.404	1.521	1.356				
33	1.480	1.430	1.610	1.372	1.480	1.430		
34	1.538	1.456	1.645	1.387	1.468	1.446		
35	1.489	1.481	1.623	1.402	1.483	1.462		
36	1.451	1.506	1.523	1.417	1.475	1.476		
37	1.476	1.529	1.619	1.430	1.587	1.491		
38	1.453	1.551	1.670	1.443	1.663	1.504	1.453	1.551
39	1.513	1.572	1.641	1.456	1.689	1.516	1.479	1.582
40	1.558	1.591	1.604	1.467	1.625	1.528	1.503	1.611
41	1.510	1.609	1.549	1.477	1.488	1.538	1.454	1.639
42	1.592	1.626	1.520	1.486	1.558	1.547	1.521	1.664
43	1.696	1.641	1.508	1.494	1.574	1.555	1.645	1.687
44	1.635	1.654	1.586	1.501	1.600	1.561	1.659	1.707
45	1.631	1.666	1.628	1.506	1.553	1.566	1.548	1.725
46	1.608	1.676	1.626	1.510	1.506	1.569	1.544	1.739
47	1.653	1.684	1.697	1.512	1.547	1.571	1.546	1.750
48	1.676	1.688	1.669	1.512	1.544	1.570	1.625	1.758
49	1.700	1.693	1.718	1.510	1.495	1.567	1.723	1.762
50	1.714	1.694	1.626	1.506	1.511	1.562	1.707	1.761
51	1.706	1.693	1.618	1.500	1.550	1.555	1.551	1.757
52	1.700	1.690	1.608	1.492	1.587	1.545	1.529	1.748
53			1.709	1.481	1.592	1.533	1.522	1.734
54	1.810	1.676	1.657	1.468	1.642	1.518	1.486	1.715
55			1.614	1.452	1.576	1.501	1.556	1.691
56	1.713	1.651	1.644	1.433	1.611	1.480	1.508	1.661
57			1.662	1.412	1.443	1.456	1.542	1.625
58	1.617	1.614			1.688	1.429	1.604	1.583
59			1.664	1.359	1.514	1.398	1.538	1.535
60	1.562	1.565			1.511	1.364	1.442	1.480
61			1.579	1.294	1.451	1.326	1.300	1.418
62	1.582	1.503			1.158	1.285	1.250	1.348
63			1.285	1.214			1.170	1.271
64	1.297	1.427			1.045	1.190	0.962	1.187
65			1.051	1.119			1.022	1.094

Table 7: Earnings profiles, PSID, cohorts 43–58

Age	Cohort 43		Cohort 48		Cohort 53		Cohort 58	
	Raw	Smooth	Raw	Smooth	Raw	Smooth	Raw	Smooth
43	1.696	1.641						
44	1.720	1.675						
45	1.839	1.706						
46	1.890	1.732						
47	1.884	1.754						
48	1.914	1.771	1.676	1.688				
49	1.878	1.783	1.675	1.755				
50	1.845	1.789	1.687	1.807				
51	1.817	1.789	1.735	1.844				
52	1.899	1.782	1.773	1.869				
53	1.871	1.769	1.751	1.880	1.795	1.646		
54	1.889	1.749	1.847	1.879	1.810	1.676		
55	1.834	1.722	1.787	1.867	1.784	1.696		
56	1.729	1.686	1.715	1.844	1.869	1.706		
57	1.721	1.643	1.832	1.810	1.799	1.706		
58	1.560	1.590	1.828	1.766	1.867	1.694	1.617	1.614
59	1.654	1.529	1.706	1.712	1.839	1.670	1.587	1.540
60	1.554	1.459	1.546	1.649	1.647	1.634	1.442	1.467
61	1.511	1.378	1.488	1.577	1.689	1.584	1.411	1.395
62	1.403	1.287	1.326	1.496	1.730	1.521	1.348	1.323
63	1.272	1.186	1.191	1.408	1.558	1.444	1.291	1.251
64	1.172	1.073	1.264	1.311	1.497	1.351	1.134	1.181
65	0.918	0.949	1.294	1.207	1.288	1.243	1.145	1.112

Table 8: Hours profiles, PSID, cohorts 23–38

Age	Cohort 23		Cohort 28		Cohort 33		Cohort 38	
	Raw	Smooth	Raw	Smooth	Raw	Smooth	Raw	Smooth
23	1.000	1.000						
24	1.017	1.009						
25	1.027	1.016						
26	0.984	1.022						
27	1.010	1.027						
28	1.042	1.031	1.042	1.031				
29	1.070	1.034	0.972	1.019				
30	1.053	1.037	0.967	1.008				
31	1.008	1.038	0.978	0.998				
32	1.033	1.040	0.974	0.990				
33	1.027	1.042	0.991	0.983	1.027	1.042		
34	1.043	1.043	0.985	0.977	0.968	1.052		
35	1.031	1.045	0.955	0.973	0.980	1.056		
36	1.063	1.047	0.920	0.970	0.982	1.055		
37	1.042	1.049	0.945	0.968	1.008	1.051		
38	1.051	1.052	0.972	0.966	1.016	1.044	1.051	1.052
39	1.044	1.055	0.942	0.966	1.038	1.037	1.031	1.036
40	1.065	1.058	0.947	0.967	0.979	1.029	0.989	1.022
41	1.057	1.062	0.975	0.969	0.951	1.022	0.964	1.011
42	1.066	1.066	0.938	0.971	0.952	1.015	0.972	1.002
43	1.098	1.071	0.933	0.974	0.962	1.010	0.997	0.996
44	1.078	1.076	0.965	0.977	0.970	1.006	0.994	0.991
45	1.101	1.081	0.962	0.980	0.948	1.003	0.983	0.988
46	1.087	1.086	0.970	0.983	0.970	1.002	0.960	0.987
47	1.067	1.091	0.962	0.986	0.963	1.002	0.964	0.987
48	1.067	1.095	0.992	0.989	0.928	1.003	1.000	0.989
49	1.067	1.099	0.995	0.991	0.918	1.005	0.996	0.991
50	1.091	1.102	0.969	0.993	0.954	1.007	0.992	0.993
51	1.112	1.103	0.969	0.993	0.956	1.010	0.957	0.995
52	1.076	1.103	0.948	0.992	0.970	1.011	0.986	0.996
53			0.958	0.990	0.975	1.012	0.963	0.996
54	1.125	1.096	0.977	0.986	0.983	1.010	0.967	0.994
55			0.973	0.981	0.968	1.006	0.974	0.990
56	1.076	1.078	0.950	0.973	0.972	0.999	0.931	0.983
57			0.937	0.962	0.927	0.988	0.967	0.972
58	1.046	1.044			0.947	0.972	0.965	0.956
59			0.977	0.932	0.887	0.950	0.932	0.935
60	0.987	0.990			0.867	0.922	0.906	0.908
61			0.840	0.889	0.839	0.886	0.837	0.874
62	0.915	0.909			0.774	0.841	0.868	0.831
63			0.795	0.829			0.759	0.780
64	0.789	0.797			0.710	0.723	0.655	0.718
65			0.749	0.749			0.661	0.645

Table 9: Hours profiles, PSID, cohorts 43–58

Age	Cohort 43		Cohort 48		Cohort 53		Cohort 58	
	Raw	Smooth	Raw	Smooth	Raw	Smooth	Raw	Smooth
43	1.098	1.071						
44	1.090	1.094						
45	1.106	1.096						
46	1.093	1.085						
47	1.113	1.068						
48	1.069	1.051	1.067	0.000				
49	1.048	1.035	1.047	0.000				
50	1.039	1.022	1.031	1.095				
51	1.015	1.013	1.029	1.060				
52	1.012	1.007	1.000	1.044				
53	0.998	1.003	1.036	1.050	1.143	1.117		
54	1.032	1.001	1.047	1.053	1.125	1.096		
55	1.021	0.998	1.028	1.052	1.100	1.083		
56	0.995	0.994	0.989	1.044	1.110	1.074		
57	1.017	0.987	1.023	1.029	1.084	1.066		
58	0.970	0.975	1.006	1.006	1.082	1.057	1.046	1.044
59	0.985	0.958	0.975	0.976	1.086	1.043	1.020	1.002
60	0.953	0.932	0.896	0.941	1.021	1.022	0.967	0.973
61	0.945	0.898	0.894	0.903	1.015	0.991	0.935	0.938
62	0.821	0.854	0.852	0.867	0.986	0.946	0.911	0.889
63	0.799	0.798	0.807	0.835	0.897	0.887	0.827	0.827
64	0.746	0.731	0.807	0.815	0.829	0.809	0.759	0.759
65	0.667	0.650	0.798	0.813	0.727	0.710	0.712	0.706

References

- Aaronson, Daniel, and French, Eric. “The Effect of Part-Time Work on Wages: Evidence from the Social Security Rules.” *Journal of Labor Economics* 22 (April 2004): 329–352.
- Arellano, Manuel. “A note on the Anderson-Hsiao estimator for panel data.” *Economics Letters* 31 (December 1989): 337–341.
- Attanasio, Orazio, Hurst, Erik, and Pistaferri, Luigi. “The evolution of income, consumption, and leisure inequality in the US, 1980-2010.” (2012). Working Paper.
- Becker, Gary S., and Stigler, George. “Law Enforcement, Malfeasance, and Compensation of Enforcers.” *Journal of Legal Studies* 3 (1974): 1–18.
- Ben-Porath, Yoram. “The Production of Human Capital and the Life Cycle of Earnings.” *Journal of Political Economy* 75 (1967): 352–365.
- Blinder, Alan S, and Weiss, Yoram. “Human Capital and Labor Supply: A Synthesis.” *Journal of Political Economy* 84 (June 1976): 449–72.
- Browning, Martin, Deaton, Angus, and Irish, Margaret. “A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle.” *Econometrica* 53 (1985): 503–43.
- Card, David. “Intertemporal labor supply: An assessment.” In Christopher Sims, editor, “Advances in Econometrics, Sixth World Congress,” Cambridge University Press (1994): 49–78.
- Casanova, Maria. “The Wage Process of Older Workers.” (2010). Manuscript.
- Erosa, Andres, Fuster, Luisa, and Kambourov, Gueorgui. “Towards a Micro-Founded Theory of Aggregate Labor Supply.” (2011). Working Paper, IMDEA Social Sciences Institute and University of Toronto.
- Erosa, Andres, Fuster, Luisa, and Kambourov, Gueorgui. “Labor Supply and Government Programs: A Cross-Country Analysis.” *Review of Economic Studies* 59 (2012): 84–107.
- Fiorito, Riccardo, and Zanella, Giulio. “The Anatomy of the Aggregate Labor Supply Elasticity.” *Review of Economic Dynamics* 15 (2012): 171–187.
- Fitzgerald, John, Gottschalk, Peter, and Moffitt, Robert. “An Analysis of Sample Attrition in Panel Data: The Michigan Panel Study of Income Dynamics.” *Journal of Human Resources* 33 (1998): 251–299.
- Freeman, Smith. “Wage trends as performance displays productive potential: a model and application to academic early retirement.” *Bell Journal of Economics* 8 (1977): 419–443.
- French, Eric. “The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour.” *Review of Economic Studies* 72 (2005): 395–427.
- Ghez, Gilbert, and Becker, Gary. *The Allocation of Time and Goods over the Life Cycle*. NBER, Cambridge (1975).

- Gibbons, Robert, and Waldman, Michael. "Careers in organizations: Theory and evidence." In O. Ashenfelter, and D. Card, editors, "Handbook of Labor Economics," volume 3 of *Handbook of Labor Economics*. Elsevier (April 1999): 2373–2437.
- Gomme, Paul, Rogerson, Richard, Rupert, Peter, and Wright, Randall. "Home Production in a Life-cycle Model." In "NBER Macroeconomics Annual," Cambridge, Mass. and London: MIT Press (2004): 415–592.
- Harris, Milton, and Holstrom, Bengt. "A Theory of Wage Dynamics." *Review of Economic Studies* 49 (1982): 315–333.
- Heathcote, Jonathan, Perri, Fabrizio, and Violante, Giovanni L. "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 1967-2006." *Review of Economic Dynamics* 13 (2010): 15–51.
- Heathcote, Jonathan, Storesletten, Kjetil, and Violante, Giovanni. "The Macroeconomic Implications of Rising Wage Inequality in the United States." Working Papers Series 14052, NBER (2008).
- Heckman, James J. "A Life-Cycle Model of Earnings, Learning, and Consumption." *Journal of Political Economy* 84 (August, Part 2 1976): S11–44.
- Huggett, Mark, Ventura, Gustavo, and Yaron, Amir. "Sources of Lifetime Inequality." NBER Working Papers 13224, National Bureau of Economic Research, Inc (July 2007).
- Imai, Susumu, and Keane, Michael P. "Intertemporal Labor Supply and Human Capital Accumulation." *International Economic Review* 45 (05 2004): 601–641.
- Johnson, Richard W, and Neumark, David. "Wage Declines among Older Men." *The Review of Economics and Statistics* 78 (1996): 740–48.
- Keane, Michael. "Income Taxation in a Life Cycle Model with Human Capital." (2011). Working Paper, ARC Centre of Excellence in Population Ageing Research (CEPAR), Australian School of Business, University of New South Wales.
- Kuruscu, Burhan. "Training and Lifetime Income." *American Economic Review* 96 (2006): 832–846.
- Lazear, Edward P. "Why Is There Mandatory Retirement?" *Journal of Political Economy* 87 (December 1979): 1261–84.
- Lazear, Edward P. "Agency, Earnings Profiles, Productivity, and Hours Restrictions." *American Economic Review* 71 (September 1981): 606–20.
- Lillard, Lee A., and Panis, Constantijn W. A. "Panel Attrition from the Panel Study of Income Dynamics: Household Income, Marital Status, and Mortality." *Journal of Human Resources* 33 (1998): 437–457.
- Low, Hamish. "Self-Insurance in a Life-Cycle Model of Labor Supply and Savings." *Review of Economic Dynamics* 8 (2005): 945–975.

- MacDonald, Glenn. “Information in Production.” *Econometrica* 50 (1982): 1143–1162.
- Medoff, James, and Abraham, Katharine. “Experience, Performance, and Earnings.” *Quarterly Journal of Economics* 95 (1980): 703–736.
- Medoff, James, and Abraham, Katharine. “Are Those Paid More Really More Productive?” *Journal of Human Resources* 16 (1981): 186–216.
- Mincer, Jacob. *Schooling, Experience and Earnings*. Columbia University Press (1974).
- Rogerson, Richard, and Wallenius, Johanna. “Micro and macro elasticities in a life cycle model with taxes.” *Journal of Economic Theory* 144 (November 2009): 2277–2292.
- Rogerson, Richard, and Wallenius, Johanna. “Nonconvexities, Retirement and the Elasticity of Labor Supply.” (2011). Working Paper, Princeton University and Stockholm School of Economics.
- Rosen, Sherwin. “A Theory of Life Earnings.” *Journal of Political Economy* 84 (August 1976): S45–67.
- Rubinstein, Yona, and Weiss, Yoram. “Post Schooling Wage Growth: Investment, Search and Learning.” In Erik Hanushek, and F. Welch, editors, “Handbook of the Economics of Education,” volume 1 of *Handbook of the Economics of Education*. Elsevier (April 2006): 1–67.
- Rupert, Peter, and Zanella, Giulio. “Grandparenting and Labor Supply.” (2010). Working Paper, University of California, Santa Barbara.
- Ryder, Harl E., Stafford, Frank P., and Stephan, Paula E. “Labor, Leisure and Training over the Life Cycle.” *International Economic Review* 17 (October 1976): 651–674.
- Thornton, R., Rodgers, J., and Brookshire, M. “On the Interpretation of Age-Earnings Profiles.” *Journal of Labor Research* 18 (1997): 351–365.
- Weiss, Yoram. “The Determination of Life Cycle Earnings: A Survey.” In O. Ashenfelter, and R. Layard, editors, “Handbook of Labor Economics,” Amsterdam: North Holland (1986): .
- Wooldridge, Jeffrey M. “Selection corrections for panel data models under conditional mean independence assumptions.” *Journal of Econometrics* 68 (July 1995): 115–132.
- Wooldridge, Jeffrey M. *Econometric analysis of cross section and panel data*. Cambridge, MA: The MIT Press (2010).



Alma Mater Studiorum - Università di Bologna
DEPARTMENT OF ECONOMICS

Strada Maggiore 45
40125 Bologna - Italy
Tel. +39 051 2092604
Fax +39 051 2092664
<http://www.dse.unibo.it>