Journal of Monetary Economics 72 (2015) 114–130

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Revisiting wage, earnings, and hours profiles

Peter Rupert [a], Giulio Zanella [b,*]

[a] University of California, Santa Barbara, United States
[b] Department of Economics, University of Bologna, Piazza Scaravilli 2, 40126 Bologna (BO), Italy

ARTICLE INFO

Article history:
Received 28 November 2012
Received in revised form
3 February 2015
Accepted 3 February 2015
Available online 12 February 2015

Keywords:
Life cycle
Wage profile
Labor supply
Human capital
Pre-retirement

ABSTRACT

For the youngest cohorts whose entire working life can be observed, hours start falling much earlier than wages. Wages do not fall (if they fall at all) until one's late 60s. The data suggest that many workers start a smooth transition into retirement by working progressively fewer hours while still facing an upward-sloping wage profile. This pattern is not an artifact of staggered abrupt retirement or selection. This evidence imposes restrictions on dynamic models of the aggregate economy, and provide updated numerical profiles that can be readily used in quantitative macroeconomic analysis to incorporate this new pattern into aggregate models.

1. Introduction

An empirical investigation into life cycle profiles of wage rates and hours of market work is undertaken, from the vantage point of consistent, four-decades-long data, to examine their behavior over the life-cycle. 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, incentive contracts, human capital, and inequality.[1] The March Supplement of the Current Population Survey (March CPS) and the Panel Study of Income Dynamics (PSID) now cover in a consistent way more than 40 years, and so allow one to observe the entire working life or substantial portions of it for several cohorts.

Older cohorts (workers born before WW2 and who entered the labor market before the 1960s) are described by the traditional tracking between wages and hours that motivates the use of parallel, hump-shaped wage and hours profiles. However, for younger cohorts (individuals born during or after WW2 and who entered the labor market during the 1960s) such tracking disappears: hours per worker fall substantially (and so do earnings) beginning shortly after age 50, but wages do not fall—if they fall at all—until these workers are in their late 60s. Moreover, this pattern is not an artifact of selection out of employment, and we argue that this substantial drop in hours while wages are still growing is generated by a process that resembles a smooth transition into retirement, in the form of less overtime work and passage from full- to part-time.

The pattern observed imposes restrictions on a benchmark life cycle model. Alternative departures from the benchmark are explored that may reconcile theory and data. A departure that seems able to reproduce the fall in hours despite non-declining

* Corresponding author.
E-mail addresses: rupert@econ.ucsb.edu (P. Rupert), giulio.zanella@unibo.it (G. Zanella).

[1] The human capital model initiated by Ben-Porath (1967) and further developed by Ghez and Becker (1975), Blinder and Weiss (June 1976), Ryder et al. (1976), Heckman (1976) and Rosen (1976) occupies a prominent position among the theories developed to characterize life cycle profiles. Weiss et al. (1986) is still an excellent review.

http://dx.doi.org/10.1016/j.jmoneco.2015.02.001
0304-3932/

wages for low educated workers is an explicit role for human capital, such as endogenous wages from on-the-job training. For highly educated workers, increasing disutility of work shortly after age 55 (not because of health reasons) seems an important part of the story.

The life cycle profiles estimated are straightforward to interpret, and can be readily used in quantitative macroeconomic analysis. In particular, our wage profiles are an up-to-date, more direct description of the life cycle than the corresponding profile based on efficiency units weights provided by Hansen (1993). The latter are hump-shaped, with the decline beginning around age 55—i.e., 10–15 years earlier than they actually start declining for the younger cohorts in our samples.[2] There may be a host of reasons why wages fall much later than hours for these cohorts. We do not take a stand on these reasons, which go beyond the scope of this paper.[3]

The focus in this paper is on the intensive margin of labor supply—hours *per worker*. The reason the extensive margin is not dealt with is that retirement choices are already well understood within the life cycle model. There is related research that show technological features such as nonconvexities in the mapping between hours and labor services (Rogerson and Wallenius, 2009; Erosa et al., 2014) and institutional features such as social security rules (French, 2005; Erosa et al., 2012) allow one to match extensive margin behavior late in the life cycle. It is the intensive margin pattern documented here that is less explored.

Revisiting wage and hours profiles is important because data long enough to allow one to observe the entire life cycle in a consistent way were unavailable until recently. An early stratagem to overcome this limitation was synthetic cohorts from single cross-sections—i.e., using the $a+j$ year-olds at time t as a counterfactual for the a year-olds at time $t+j$. However, such synthetic cohorts may produce a biased picture of the life cycle when productivity changes over time. Thornton et al. (1997) and Rubinstein et al. (2006) offer an illustration for annual and weekly earnings, respectively. An alternative method was the construction of pseudo-panels from repeated cross-sections. For instance, Mincer (1974) analyzed census data and observed that the weekly wage rate was not declining at the end of working life, while annual earnings were. This method is also the starting point of our analysis, which confirms Mincer's observation. The increasing availability of longitudinal data has allowed a more direct look at actual portions of workers' careers by following them over time. Johnson and Neumark (1996) used 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 did not find clear evidence of negative wage growth until workers in their sample were in their 60s. They did 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. (2011) and Heathcote et al. (2010b). The dynamics of earnings, however, results from both wages and hours, and it is of interest to look separately at these two processes.

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, with the turning point between the ages of 55 and 60. The hours profile he estimates is also declining, with the trend changing at about these same ages. 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). The apparent contrast with our results has a simple explanation: the different pattern that we document here emerges with cohorts that are younger than those considered by French. 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, and produces out-of-sample (i.e., after age 40) predictions that result in markedly hump-shaped profiles. Nonetheless, as illustrated below, Imai and Keane (2004) offer a key insight: investment in human capital increases substantially the opportunity cost of time early in the life cycle, which offers an avenue for reconciling 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 leaving 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 one of the variants of the benchmark model explored below.

The remainder of the paper is organized as follows. Section 2 describes the data. Section 3 illustrates how these data restrict the theory. Section 4 concludes. In a Supplemental Web Appendix additional data is provided, as well as tables containing the numerical wage and hours profiles that can be readily used in quantitative analysis.

[2] Abbasoglu (2012) compares the profiles in Hansen (1993) to ones generated from the PSID, similar to ours.

[3] One possibility is that there are pure time and cohort effects at work. Another is that these effects reflect higher education and higher participation rates at older ages for these cohorts relative to older ones. It is also possible that human capital depreciates only slowly (or not at all) for these younger cohorts, or that the training content of an hour of work now decreases fast enough to compensate a declining productivity profile before retirement. Yet another explanation is that incentive contracts are in place. For instance, the work, among others, of Lazear (1979, 1981), Freeman (1977), Medoff and Abraham (1980), and Harris and Holstrom (1982) show that individual wage growth is possible even in the absence of individual productivity growth if dynamic incentives matter. Gibbons et al. (1999) provide a thorough review of this literature. A related question not addressed is what is the underlying productivity profile; as the focus is on *observed* wages rather than *unobserved* productivity because hours respond to the former, but of course these two are not necessarily the same.