Revisiting the Hump-Shaped Wage Profile: Implications for Structural Labor Supply Estimation

María Casanova*†

February 2013

Abstract

Structural labor supply models rely on accurate estimates of the age profile of offered wages. The typical specification for this profile is hump-shaped over the life-cycle. This paper shows that the hump-shaped profile results from aggregation over individuals who become partially retired at different ages. It is not a good proxy for the offered wage profile, which is flat at older ages. I show that misspecification of the offered wage profile for older workers leads to structural estimates of the intertemporal elasticity of substitution of labor supply that are biased upwards by 30 to 130 percent, depending on the specification.

JEL Codes: J31, J26

*University of California, Los Angeles (UCLA). Department of Economics. 8283 Bunche Hall. Los Angeles, CA, 90095. Email: casanova@econ.ucla.edu.

†An earlier version of this paper was circulated under the title "Wage and Earnings Profiles at Older Ages: Implications for the Estimation of Labor Supply Elasticities". I have benefited from insightful comments from Orazio Attanasio, Richard Blundell, Eric French, Michael Hurd, Hamish Low, Robin Lumsdaine, Maurizio Mazzocco, Lars Nesheim, Franco Peracchi, Andrew Samwick, and seminar participants at the University of British Columbia, Washington University- St Louis, the Michigan Retirement Research Center, the CASEE 2012 Quantitative Macroeconomics conference, the Aging Workshop at the 2012 NBER Summer Institute, the 2012 Meetings of the Society for Economic Dynamics and UCL PhD Alumni Conference. All errors are my own.

1 Introduction

Structural estimation and calibration of labor supply models have become more prevalent in recent years. These models rely on estimates for the laws of motion of exogenous state variables that are often obtained outside the model or taken from other studies in the literature. A key input to structural models of labor supply is the age profile of offered wages, also known as the *deterministic* or *predictable* component of wages. But what does this profile look like?

The most common specification found in the literature assumes that individuals face an inverted-U-shaped wage path that peaks in middle age and declines smoothly thereafter. The justification for this assumption goes back to the human capital model of Becker (1993), Ben-Porath (1967), and Mincer (1974)), which predicts a flattening off and eventual decline of a worker's productivity profile as he approaches retirement age. To the extent that wages reflect productivity, this translates into an inverted-U-shaped wage profile.

This paper first establishes that the wage path followed by the average individual does not decline smoothly towards the end of the working career, as implied by the popular inverted-U-shaped specification. Instead, the real hourly wage of the typical male over age 50 increases slightly with age for as long as he is employed full time. Two-thirds of workers transit directly from full-time work into full retirement. For these workers, no wage declines are observed at any age. The remaining one-third of workers transit into retirement more gradually, and go through at least one period of part-time work before withdrawing from the labor force -in what follows, I refer to this transitional period as *partial retirement*. The transition from full-time to part-time work is associated with a one-off 34% hourly wage drop. Thereafter, wages once more remain flat. Thus, the wage path followed by workers who partially retire is best represented by a step function. The smoothly declining wage profile often estimated in the literature is a reflection of the increasing proportion of part-time employees as workers age. These results imply that this profile is not a good proxy for the wage path followed by the average individual.

Having established that wage drops at older ages are driven by transitions into part-time work, the next step is to study the nature of these transitions. I propose and test 3 alternative models of retirement that differ in the forces driving retirement decisions and the underlying process for offered wages. In the *self-selection* model, offered wages follow a smooth, downward trend at older ages. In this case, transitions into partial and full retirement are a response to declining wages. Although wages decline on average, different workers are hit by different idiosyncratic shocks, and those who receive positive wage shocks are more likely to remain in full-time employment. Because of this positive self-selection, a declining offered wage profile can be reconciled with the flat wage profile for full-time workers observed in the data. The self-selection model is the only one that generates a wedge between offered and observed wages. In the next two models, the two overlap. Offered full-time and part-time wages are both flat with age. These two models differ on the determinants of retirement transitions. In the *involuntary retirement* model, factors outside the worker's control trigger separations from the full-time job with increasing probability every year. In this model, the offered wage profile is flat for full-time workers, but the expected wage profile for an individual approaching the end of his career declines smoothly. Finally, in the *voluntary retirement* model offered wages are flat and full and partial retirement transitions arise as optimal choices for workers who compare the relative value of work versus leisure. At every age, workers choose between a full-time job with a full-time wage w, a part-time job with a wage θw, where $\theta < 1$, and full retirement. Transitions into partial and full retirement are not driven by exogenous change in w, which is flat with age, but are instead triggered by changes in preferences (e.g., increasing taste for leisure) or the budget constraint (e.g., work disincentives after certain ages from the social security rules). In this model, wages and hours are chosen simultaneously.

The empirical evidence rejects the model of *self-selection*, indicating that there is no divergence between offered and observed wages. It follows that the correct specification for the offered wage profile is flat in age. Next, I provide evidence that between 20 and 30% of transitions out of full-time work are driven by events out of the worker's control, such as health events or plant closures. The substantial number of involuntary transitions underscores the key role of uncertainty in labor supply decisions of older workers. More importantly, the remaining 70 to 80% of partial and full retirement transitions are voluntary. This result indicates that the vast majority of retirements arise as the optimal choice for workers who have the option to remain in full-time employment with their full wage, but choose instead to trade more leisure for a lower hourly wage. As a consequence, their wages and hours are determined simultaneously, in sharp contrast to the canonical labor supply model, which posits that participation and hours choices are a function of an exogenous wage rate.

The last part of the paper illustrates the relevance of these results for structural estimation and calibration. I focus on a parameter that is key to labor and macro analyses of labor supply, namely the intertemporal elasticity of substitution of labor supply (*i.e.s.*). The magnitude of this parameter determines participation and hours responses to lifetime wages, and is crucial in studies of issues ranging from business cycles to tax or welfare reform. The *i.e.s.* has attracted considerable interest from both the labor and macro fields for the last 40 years (Lucas and Rapping (1969), Ghez and Becker (1975)). I show that estimates of the parameter that drives labor supply responses to lifetime wage changes, namely the intertemporal elasticity of substitution of labor supply (*i.e.s.*), are severely biased when a smoothly declining wage profile is used as a proxy for the flat offered wage path at older ages. I use artificial data generated from a structural model of consumption and labor supply choices of older workers to show that this type of misspecification leads to an upward bias in estimates of the *i.e.s.* of 30 to 130 percent, depending on the specification.

The implications of the results presented in this paper are wide ranging because key results in the labor supply literature depend on the correct specification of the age-wage profile. These include measures of income uncertainty; assessments of the relative importance of determinants of retirement decisions; computations of present discounted values of private and social security pensions; and results based on estimates of the *i.e.s.*, including estimates of labor supply responses to business cycles and life-cycle patterns of work and leisure.

Moreover, the results are also relevant to models of life-cycle consumption and savings. These models generally do not endogenize labor supply, and they use labor earnings, rather than hourly wages, as an input. I show that the earnings profile is also flat for full-time workers. As was the case for hourly wages, earning declines at older ages are driven by transitions into partial retirement. Specifically, earnings drop by 73% for a worker who switches into part-time work, and remain flat thereafter.

This paper builds on the literature that has studied the age-wage and age-earnings profile, dating back to Hurd (1971), Lillard and Willis (1978), Mincer (1974), and Welch (1979). Early cross-sectional studies find wage and earnings declines setting in as early as age 40 (see Willis (1985) for a survey). Longitudinal studies such as Honig and Hanoch (1985b), Johnson and Neumark (1996), and Rupert and Zanella (2012) find that wages peak later, and observe the sharpest declines around the social security ages. A few studies have pointed out that transitions into part-time work at older ages are often accompanied by significant wage cuts (Aaronson and French (2004), Hurd (1996)). Gustman an Steinmeier (1985) show that controlling for partial retirement status attenuates the rate of the decline of the cross-sectional wage profile, although it remains downward sloping in the last part of the life-cycle. Finally, Murphy and Welch (1990, 1992) estimate earnings profiles for the sub-sample of full-time workers. Consistent with the findings in this paper, they observe no or very small wage declines as workers age.

The paper proceeds as follows. Section 2 describes the data and the profiles of the variables of interest; section 3 presents estimates of wage, hours, and earnings paths for the average individual; section 4 characterizes the offered wage profile and tests alternative models of retirement; section 5 describes how structural estimates of the *i.e.s.* are affected by misspecification of the offered wage profile; and section 6 concludes.

2 Data and Variable Definition

A. Sample Description

The data used in the analysis are drawn from the University of Michigan Health and Retirement Study (HRS), a panel dataset of adults over 50 years of age and their spouses. Other longitudinal datasets, such as the PSID, collect information on older workers, but sample sizes are small for this age group. The HRS provides a large sample size and information on a uniquely extensive set of covariates, including employment status, job history, income, pensions, net worth, and demographic characteristics. HRS respondents are interviewed for the first time in 1992, and subsequently every two years. I use the 9 survey waves available up to 2008.[1]

I apply several restrictions to the data to produce a sample of workers that are relatively homogenous and have not yet started the retirement process. First, I limit the sample to men born between the years 1931 and 1941, which constitute the core HRS cohort; 4,895 HRS respondents fit these criteria. Second, only workers who are employed full time by the time they enter the panel are kept in the sample. This excludes 390 respondents who are partially retired in the first panel wave, and 1,152 who are fully retired. The reason for this restriction is that the wage change upon separation from the full-time job, which is the focus of the paper, is not observed for these individuals.[2] Third, because of concerns regarding the measurement of hourly wage rates for the self-employed, 1,015 individuals who report being self-employed for at least one period are also dropped. Appendix A.3 discusses the sensitivity of the results to this restriction. Finally, observations for a particular year are excluded if the hourly wage is lower that \$3 or higher than \$70 (in 2011\$), the number of

[1]Where possible, HRS variables are extracted from the RAND-HRS data file. For a description of this dataset, see "RAND HRS Data, Version J" (2010). A supplementary issue of the *Journal of Human Resources* (1995, vol. 30) provides more details about the HRS.

[2]See appendix A.1 for a description of the observable characteristics of individuals who are retired by the time they enter the HRS panel

hours worked cannot be established, or the individual is younger than 51 or older than 67 (due to small sample sizes outside those ages).[3] The final sample consists of 2,318 individuals and 12,036 person-year observations.

B. Cross-Sectional Wage and Hour Profiles

Below I describe the cross-sectional profiles for the main variables of interest in the HRS data. Due to the age structure of the HRS sample, these profiles span the last years of the working life only. For comparison, lifetime average profiles obtained from the census are also provided.[4]

The focus of the paper is on the real hourly wage (expressed in 2011$ using the CPI). The lifetime hourly wage profile, shown in figure 1, displays the well-known inverted U-shape.[5] Average wages increase with age in the early stages of the working career, flatten out for workers in their 40s, and decline from about age 50 until retirement. The second series of figure 1 shows the hourly wage profile for the pooled HRS sample. Due to differences in the definition of hourly wages -the census hourly wage rate includes commissions, cash bonuses, tips, and income from secondary jobs, while the HRS variable does not- the two series do not overlap. Despite the difference in levels, the HRS series tracks the rate of decline in the census remarkably well.

The other key variable in the empirical analysis is participation status, which is used to determine whether a worker is partially retired. Work status is defined as a function of the number of hours worked per week at the current job. Workers are classified as working full time (part time) if they work at least (less than) 35 hours per week.[6] The left panel of figure 2 shows life-cycle participation profiles from the census. Two main trends are noteworthy. First, participation rates start declining for workers in their late 40s, at approximately the same time as wages. Second, part-time participation increases threefold from age 50 to age 65. The proportion of workers in part-time employment rises from 4% to 28% between those ages. The right panel of figure 2 shows participation profiles for the pooled HRS sample.

[3]The bounds on hourly wages are intended to minimize the effect of outliers. Rows 1 to 3 of table 6 show that using an upper bound of $50, $100, or removing it altogether hardly affects the main results. This is not surprising, as only 2% of sample observations have wages above $70, and just below 6% have wages above $50.

[4]Data from the 1990 US census are used. For a description of the census sample, see Ruggles et al. (2010).

[5]The inverted-U-shaped wage profile indicates that the hump in lifetime earnings (described by Lillard and Willis (1978), Mincer (1974), Welch (1979) and Heckman, Lochner and Todd (2006), among many others) is not purely a reflection of declines in hours worked during the pre-retirement years.

[6]Using different definitions of full-time work, part-time work and inactivity, such as classifying as full-time workers those who work more than 30 hours per week, and considering those who work less than 2 or 5 hours per week as inactive, does not affect the results.

Figure 1: Average hourly wage by age, 1990 Census and HRS.

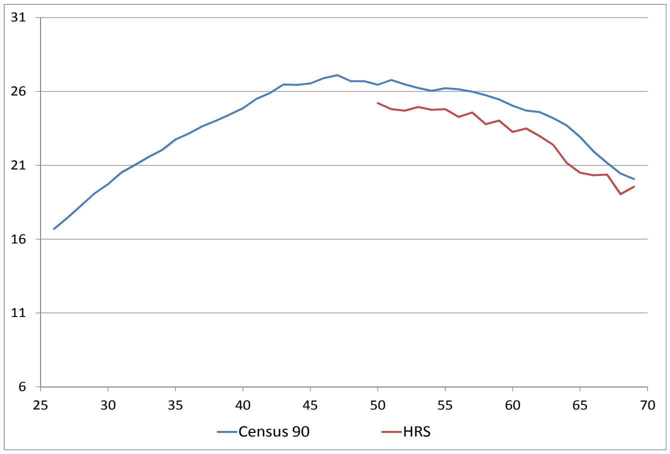

NOTE. - Hourly wages deflated to 2011 dollars using the CPI. Hourly wages lower than $3 and higher than $70 in real terms excluded. Sample includes all workers.

The same key trends -a rapid decline in total and full-time participation after age 50, and a simultaneous increase in part-time work- are apparent for these data.

Figure 2: Total, full-time, and part-time participation rates by age, 1990 Census and HRS.

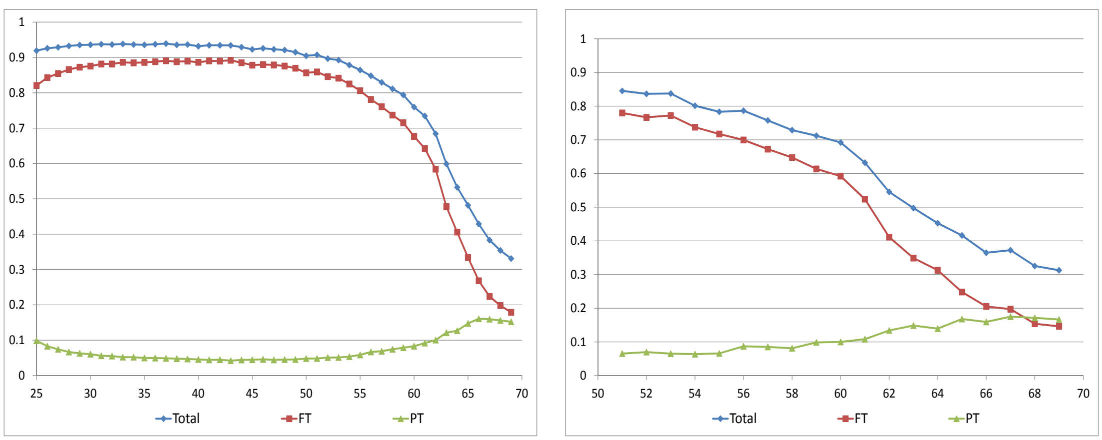

C. Definition of Partial Retirement

A worker is classified as partially retired the first time he transits from full-time into part-time work. Partial retirement is modeled as an absorbing state, that is, workers who go back to full-time work are still considered to be partially retired. This modeling assumption interprets the first exit from full-time work as a signal that the worker has entered the transitional period leading to full retirement. In practice, only a small proportion partially retired workers return to full-time work. As a result, the definitions of partial retirement and part-time work overlap closely, and in the rest of the paper I will refer to both interchangeably. According to this definition, 29% of workers in the sample are classified as partially retired in at least one period.

The measure of partial retirement used here is relatively parsimonious, requiring information on full- versus part-time status only. It will be easy to replicate in other contexts, given that most datasets that measure wages also provide information on the number of hours worked.[7]

3 Wage and Earnings Profiles

In this section, I first characterize the hourly age-wage profile for the typical male worker in the last stages of his working career. Next, I characterize the hours profile for the same worker. Combining the two sets of results, I construct the earnings profile for the average male above age 50.

A. Age-Wage Profile

I assume the following specification for the conditional mean of log wages:

$$w_{it} = W(Age_{it}) + X_{it}\beta + u_{it}, \tag{1}$$

where w_{it} is the logarithm of the real hourly wage of individual i in period t, W some function of age, X_{it} a vector of other observable characteristics, and u_{it} the error term.

Equation (1) can be interpreted as a standard Mincerian regression where wages, rather than earnings, are used as the dependent variable. The focus on wages allows to separate earning potential from hours choices, which are discussed separately below. The roles of education and experience are not made explicit here to maintain the focus on age, but both are implicit in the regression: education is captured by the fixed effect, while experience is embedded in the age effect.[8]

[7]Other studies have proposed somewhat different definitions of partial retirement. As a robustness check, I have re-run the analysis for 3 alternative definitions. The results are discussed in appendix A.3

[8]Using age, rather than potential experience, can confound the interpretation of the results if experience-wage profiles are not parallel across schooling levels over the age range covered in the sample. To investigate this possibility, I have run the main wage regressions separately by education category, that is, keeping potential experience constant. The results, shown in columns 7 to 9 of table 6, show that wage profiles are flat, and therefore parallel, for workers over age 50 in the 3 education groups considered. It is interesting to compare these results with those in Heckman, Lochner and Todd (2006), who reject parallel *average* earnings

I begin by estimating equation (1) by OLS, pooling all available wage observations. The results are shown in the second column of table 1. The implied wage profile, displayed as the grey line in figure 3, trends downwards for workers in their mid 50s and older. Specifically, the estimated coefficients imply a 3% decline in wages between ages 51 and 60 -which is not statistically significant according to the joint significance tests shown at the bottom of table 1-, followed by a strongly significant 22% decline between ages 60 and 67. The rapid wage decline at older ages is consistent with the inverted-U-shaped lifetime wage profile described in the previous section.

In the literature, the wage profile of interest is often the individual wage profile, that is, the one that measures returns to age for the average worker. The OLS profile estimated above measures something different, namely the mean wage for the sample of individuals that remain employed at every age. It is clear that, in general, these two profiles need not overlap. And yet it is common practice to use the OLS wage profile as a proxy for the individual one, particularly when working with cross-sectional datasets, such as the CEX, that preclude following workers over time.[9] This may be a reasonable approach for individuals in their prime working years, as the vast majority are employed full time. During the retirement years, however, the characteristics of the average worker are not stationary. To the extent that the retirement propensity is correlated with variables that also correlate with wages, such as education, pension type, or productivity, the OLS profile is a biased estimate of the returns to age. In particular, if variables associated to higher wages are predictors of early retirement, then the OLS wage profile understates the returns to age for the typical worker.

Unlike OLS, the fixed-effect estimator is robust in the presence of heterogeneous retirement propensities for workers with different observables.[10] Hence, I next estimate equation (1) by fixed-effects. Coefficients from this regression are presented in the second column

---

profiles. Several mechanisms can generate diverging education-specific average earnings profiles when the underlying *individual* wage profile is flat. The most salient ones are differences in retirement propensities across education categories, and a full/part-time wage gap that is education-specific. Both mechanisms are present in the data. Table 3 shows that (partial) retirement probabilities vary with educational attainment. Table 6 shows that the wage gap between full-time workers and partial retirees increases monotonically with education.

[9] See, for example, Gomes, Kotlifoff and Viceira (2008), Heathcote, Storesletten and Violante (2010), Guvenen (2009), Kopecky (2011), and Low (2005) for OLS profiles of hourly wages or hourly earnings. Attanasio and Weber (1995), Ghez and Becker (1975), Gourinchas and Parker (2002), Heckman (1976), Heckman, Lochner and Todd (2006), Krueger and Kubler (2006), Scholz, Seshadri and Khitatrakun (2006), and Storesletten, Telmer, and Yaron (2004) estimate OLS earnings profiles. As shown below, the conclusions regarding individual wage profiles extend to individual earnings profiles.

[10] It is worth noting that the fixed-effects estimator is not consistent if individuals retire as a result of wage shocks. The issue of self-selection is addressed in the next section.

of table 1. The results are also shown graphically as the dashed line in figure 3. The fixed-effects profile is mostly flat for workers in their 50s, and declines rapidly from age 60 onwards. Specifically, the estimated coefficients imply a borderline significant 2% increase in real hourly wages between ages 51 and 60, followed by a strongly significant 17% decline between ages 60 and 67. After controlling for composition bias, the rate of decline is attenuated with respect to the OLS profile, but hourly wages still fall noticeably over the retirement period.[11]

Figure 3: Log Wage Profiles for Different Specifications

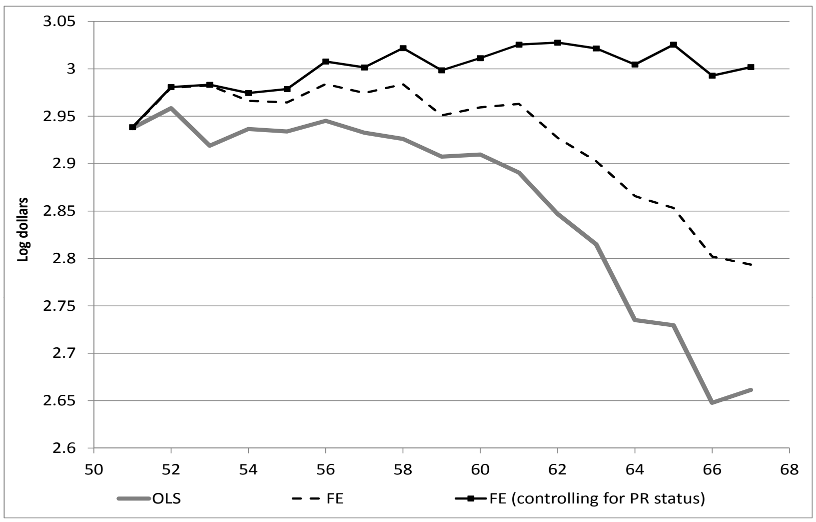

NOTE. - Series constructed using coefficients from columns 1 to 3 in table 1.

The fixed-effects profile has been used as a proxy for the individual one in studies where panel data are available.[12] Although an improvement with respect to OLS, it is still an *average* profile, measuring mean wage growth for employed individuals at every age. The increasing prevalence of part-time work during the retirement period naturally raises the question of whether average wage declines mask different wage growth rates between workers who remain in full-time employment and those who switch to a part-time job. Put differently, does the hourly wage of the typical worker decline smoothly over the retirement years, as suggested by the OLS and fixed-effect profiles? Or are declines in average wages driven by the increasing proportion of workers transiting into less well paid, part-time jobs during the retirement period?[13]

[11] Honig and Hanoch (1985b), Johnson and Neumark (1996) and Rupert and Zanella (2012) also find that the OLS profile declines more rapidly than the fixed-effects one.

[12] See, for example, Baker (1997).

[13] The well-known fact that part-time workers earn lower hourly wages is confirmed in figure 5 below. Lower

To answer this question, I re-estimate equation (1) adding as an additional regressor an indicator that is equal to 1 if the individual is partially retired. Since partial retirement is modeled as an absorbing state, its coefficient is identified from individuals who transit from full-time into part-time work for the first time, i.e., *entrants* into partial retirement. Estimates from this regression, reported in the third column of table 1, show that the transition into partial retirement is associated with a 34% decline in hourly wages. The estimated age-wage profile is shown as the solid line in figure 3. The contrast with the OLS and fixed-effects wage series is striking. After controlling for retirement status, the wage profile no longer declines with age. Wages now increase by 7% between ages 51 and 60, and by 0.7% from 60 to 67. The p-value for the test of joint significance of the age dummies spanning the retirement years is 0.70, indicating that a completely flat wage profile cannot be rejected for the late-career period.

Figure 4: Predicted Hourly Wage Profile (Example)

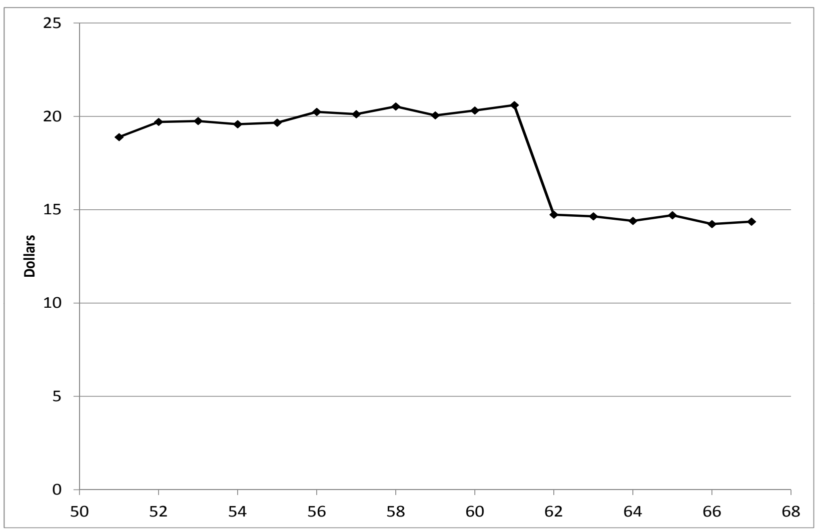

NOTE. - Series constructed using coefficients from column 3, table 1, assuming that the transition into partial retirement takes place at age 62.

The last set of results implies that average wage declines at older ages are fully driven by transitions into partial retirement. The wage path for the typical individual over age 50 is flat or increases slightly with age as long as he remains employed full time. For those who transit

wages for partially retired workers can arise as a result of the loss of tenure upon separation from the career job, the existence of a part-time wage penalty, or differences in job characteristics if partial retirees transit into occupations that require less effort or responsibility. Distinguishing among the different mechanisms is beyond the scope of this paper. See Johnson and Neumark (1996) and Ruhm (1990) for further discussions of the wage and job characteristics of partial retirees.

directly into full retirement, no wage declines are observed. Workers who partially retire, on the other hand, experience a discontinuous wage drop upon separation from the full-time job. Figure 4 illustrates, as an example, the predicted wage profile for an individual who transits into partial retirement at age 62. The resulting step-function profile is very different from the smoothly declining average wage profile discussed above. It is clear that the average profile obtained by aggregating over individuals who transit into partial retirement at different ages is not a good proxy for the wage path followed by the average individual.

The effect of aggregation can also be illustrated using the results of separate regressions for the subsample of non-retired workers (the *full-time* profile, shown on the left series in figure 5) and of partially retired workers (the *part-time* profile, shown in the right series in figure 5). Coefficients for these regressions are presented on the last two columns of table 1. Joint significance tests provide no evidence of wage declines for full-time workers. For partially retired workers there are both positive and negative wage changes at older ages, but no discernible trend. The average wage profile is a weighted average of these two series. The increasing weight on part-time wages as workers age and more of them become partially retired explains its downward slope.

Figure 5: Full-Time and Part-Time Log Wage Profiles

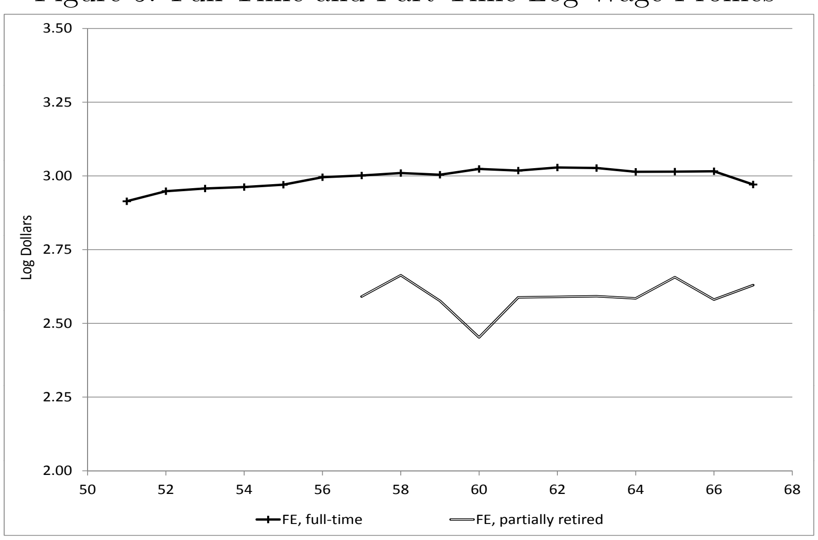

NOTE. - Series constructed using coefficients from columns 4 and 5 in table 1.

B. Age-Earnings Profile

Because structural models of consumption and saving often use the earnings profile (rather than the hourly wage profile) as an input, this section characterizes the earnings profile that results after controlling for retirement transitions. To do so, I proceed in two steps: I first analyze the (weekly) hours profile, and then combine the results with those already presented for hourly wages to characterize the (weekly) earnings profile, constructed as the product of the two.[14]

I use the same specification for the hours profile as I did for hourly wages. I estimate equations analogous to (1) using log weekly hours as the dependent variable. As before, I compare results from OLS and fixed-effects regressions estimated by pooling observations for full-time and partially retired workers to those obtained from a fixed-effects regression that controls for partial retirement status. The coefficients are presented in columns 1 to 3 of table 2. The implied hour profiles (in levels) are plotted in figure 6.

Figure 6: Hours Profiles for Different Specifications

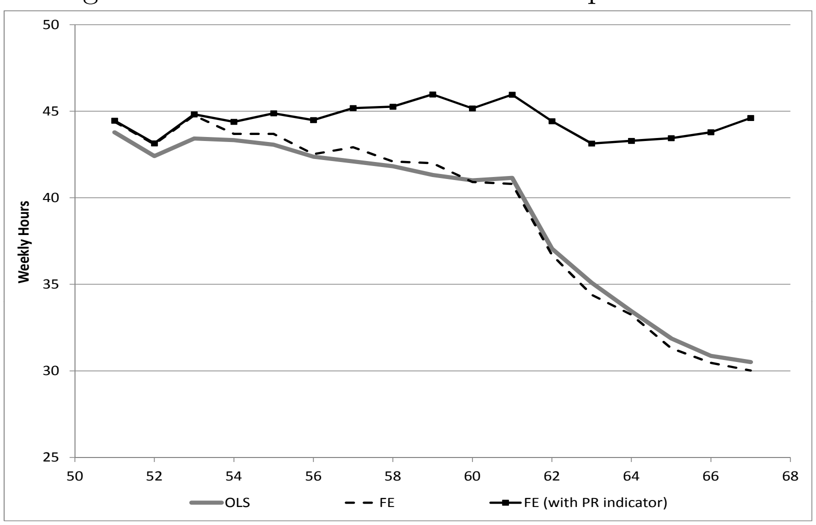

NOTE. - Series constructed using coefficients from columns 1 to 3 in table 2.

The OLS and fixed-effect profiles are nearly identical, suggesting that the composition effects that are evident for wages are not a concern for hours regressions. That is, even though individuals with different observables are heterogeneous in their propensity to retire, on average they all work the same number of hours. Thus, changes in the characteristics

[14] I do not use the yearly earnings measure from the HRS in this analysis because it does not allow to accurately measure the earnings drop upon partial retirement. This is because the HRS collects earnings information for the whole year before the time of the interview. For workers who partially retired during that year, the earnings measure includes some earnings obtained while employed full time. Because information on hourly wages and hours during the pre-retirement period is generally not available, it is not possible to accurately divide yearly earnings between pre- and post-retirement jobs. As a result, estimates of the earnings profile using yearly earnings as the dependent variable will understate the drop in earnings associated to the partial retirement transition.

of the average worker over the retirement years leave the mean number of hours worked unchanged. Regarding the regression that adds a partial retirement indicator, several findings are noteworthy. First, the transition into partial retirement is associated with a 64% drop in hours worked. Second, most of the decline in average hours at older ages is driven by transitions into partial retirement. This is illustrated in figure 6. Average hours, represented by the dashed line, decline steeply with age, dropping by close to 8% from ages 51 to 61, and by a further 31% from ages 61 to 67. The acceleration in the downward trend takes place at the social security early retirement age (62). By contrast, the hour profile obtained after controlling for transitions into partial retirement, represented by the solid line, shows no decline in hours before age 60, and a small 1% decline from 61 to 67.[15]

The smoothly declining average profile masks a sharply discontinuous labor supply profile at the individual level. The two findings discussed above are inconsistent with a gradual reduction of hours worked in the years leading to full retirement. Instead, labor supply adjustments are concentrated in two points in time. Hours drop from full time to 0 for workers who transit directly into partial retirement. For workers who partially retired, there is a further discontinuous drop in hours, illustrated on the left panel of figure 7 for a hypothetical worker who becomes partially retired at age 62. As was the case for wages, the average hour profile is not a good proxy for the labor supply path of the average individual.

One final point worth noting is the acceleration in the downward trend of the average hour profile around the social security early and full retirement ages (62 and 65, respectively). The coefficients in column 2 of table 2 show an 11% decline in average hours at age 62, and a further 6% decline at age 65. Once the transition into partial retirement is accounted for, the coefficient on the full retirement age dummy falls to zero and becomes insignificant. The coefficient on the early retirement dummy remains significant, but its magnitude is reduced by more than two-thirds. These results confirm previous findings indicating that a large number of transitions into partial retirement take place at the social security ages. Previous studies have argued that transitions into part-time work at these ages are driven by incentives from the social security rules (See Aaronson and French (2004), Blau (2008), French (2005) and Rust and Phelan (1997)). Since these rules are exogenous from the point of view of the individual, in the next section I will use them as exclusion restrictions for partial and full retirement probabilities.

[15]The age dummies spanning the social security ages are jointly statistically significant, suggesting that some hour adjustments precede the retirement decision. This is confirmed by the regression run using only the sample of full-time workers (column 4 of table 2), which shows a 2% hour drop at age 62.

Figure 7: Predicted Weekly Hour (left) and Weekly Earnings (right) Profiles.

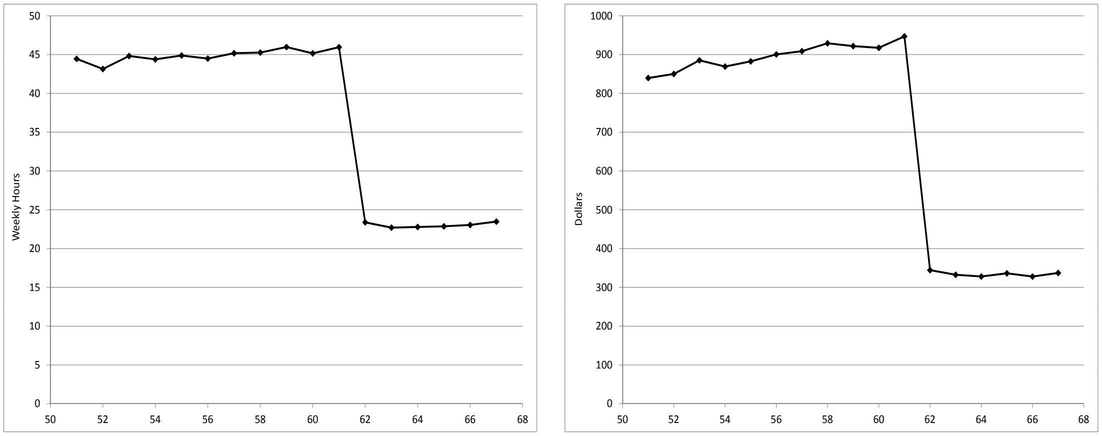

NOTE. - Left series constructed using coefficients from column 3, table 2. Right series constructed combining coefficients from column 3 in tables 1 and 2. Both series assume that the transition into partial retirement takes place at age 62.

Finally, the results for wages and hours can be combined to offer some insight on the shape of the earnings profile. The right panel of figure 7 shows predicted weekly earnings for an individual who transits into partial retirement at age 62. Not surprisingly, as weekly earnings are obtained as the product of the wage and hours profiles shown above, the earnings profile is also a step function.

The definition of labor earnings used in this analysis does not include bonuses, tips, or income from secondary jobs. It is conceivable that these may be adjusted before partial retirement. However, in view of the evidence in figure 7, it seems safe to conclude that the drop in earnings upon entrance into partial retirement accounts for the vast majority of the decline in earnings at older ages.

4 Offered Wages

The previous section characterizes the *ex post* wage path followed by a typical individual, which is conditional on labor supply decisions. That profile captures the large wage drop associated with transitions into part-time work for workers who choose to partially retire, but it is not informative about the wage an individual would receive *were he to remain employed full time*. Notice that a wage profile that is not confounded by hour choices is the primitive of interest in a structural model that endogenizes labor supply decisions. This section characterizes such a profile.

The step function wage profile described in the previous section is consistent with at least three models of retirement. These differ in the forces driving the retirement decision and the underlying process for offered wages. In order to characterize the offered-wage profile, I test the empirical implications of the three models to determine which of them is compatible with the data.

In the first model, which I will refer to as the *self-selection model*, offered wages trend downwards at older ages; that is, the lifetime offered wage profile is hump-shaped. Becker (1993), Ben-Porath (1967) and Mincer (1974) provide the microfoundations for this model by assuming that offered wages are a function of an unobserved variable, such as productivity, that declines smoothly with age. To the extent that individuals allocate more work hours to periods of higher wages, they will switch from full-time work into partial or full retirement once their wage falls below some threshold. Thus, in this framework, transitions into partial and full retirement at older ages are a response to declining offered wages. The key testable implication from this model is a discrepancy between offered and observed wages. Recall that the evidence presented in the previous section shows that observed full-time wages are flat or increase slightly with age. For this to be compatible with declining offered wages, individuals who receive positive wage shocks must be more likely to remain in full-time employment. Hence, this scenario implies a positive self-selection bias.

Alternatively, if there is no selection on unobservables, offered wages are the same as observed wages, and both the full-time and part-time age-wage profiles are nondecreasing in age (as shown in figure 5). In this case, we can broadly define two retirement models depending on whether retirement transitions are voluntary.

In the *involuntary retirement model*, transitions out of full-time work are due to factors outside the worker's control. For simplicity, suppose that they are driven by a discrete Markov process that the individual takes as exogenous. In each period workers are hit by a shock leading to job separation with positive probability (e.g., they become disabled or are fired).[16] In this framework, while the *ex-post* wage profile for full-time workers is flat, the *expected* wage profile for an individual approaching the end of his career declines smoothly as long as the probability of involuntary job separation increases continuously with age.

Finally, in the *voluntary retirement model* offered wages are flat and transitions into partial and full retirement are voluntary, i.e., driven by the relative preference for leisure rather than an exogenous shock that prevents work. In a model of voluntary transitions, workers evaluate different bundles of wages and hours (e.g., wage $w_t > 0$ for full-time work, wage θw_t, with $0 < \theta < 1$, for part-time work, and no wage for full retirement), and effectively

[16]Models of deferred compensation with mandatory retirement are included in this category.

choose the two variables simultaneously. In the absence of changes in offered wages with age, other variables must be responsible for shifting the relative value of leisure and generating transitions into part-time work and out of the labor force. These variables can operate through preferences (e.g., increasing taste for leisure) or through the budget constraint (e.g., work disincentives after certain ages from the social security rules). In either case, retirement transitions are not a response to declining wages in this model. Rather, wage and hour changes are determined simultaneously when workers choose to transit from full-time into part-time work.

Each of these 3 models of retirement has empirical implications that can be tested. The remainder of the section examines which of them is compatible with the data.

A. Self-selection into Retirement

To test whether transitions out of full-time work are driven by selection on unobservables, I follow the procedure described in Wooldridge (1995), which extends Heckman's (1979) two-step estimator to a panel-data context.

The equation of interest is the age-profile of offered wages for non-retired (i.e., full-time) workers, namely:

$$w_{it}^{FT^*} = f_i + A(Age_{it}) + X_{it}\delta + \varepsilon_{it}, \tag{2}$$

where w^{FT^*} is the logarithm of the real offered full time hourly wage, A is some function of age, X a vector of observables, f an individual-specific effect, and ε an idiosyncratic error term.

Two alternative specifications for the selection process are considered. The first one lumps partial and full retirement together: The selection indicator d^R is equal to 1 if the individual is not working full time. It is determined by the following probit equation:

$$\begin{aligned} d_{it}^R &= I[Z_{it}\beta + \bar{Z}_i\Pi > -v_{it}], \\ v_{it}|Z_i &\sim \text{Normal}(0,1) \end{aligned} \tag{3}$$

where $Z \supset X$ is a vector of observables that contains all the variables in the vector X plus the variables used as exclusion restrictions, and $\bar{Z}_i$ a vector of within-individual averages.[17]

[17]The vector is defined as follows: $\bar{Z}_i \equiv \frac{1}{T_i}\sum_{j=1}^{T_i} Z_{it}$, with T_i equal to the number of periods individual i is observed in the panel.

v_{it} is a normally distributed error term.[18] Full-time wages are observed for individuals who are neither partially nor fully retired, that is, $w_{it}^{FT} = w_{it}^{FT^*} * (1 - d_{it}^{R})$.[19]

The selection equation is estimated by probit regression, and the results are used to generate the inverse Mill's ratio, $\hat{\lambda}_{it}^{R}$. In the second step, $\hat{\lambda}_{it}^{R}$ is added as a regressor to the wage equation, which is estimated by fixed-effects, using observations for full-time workers only. Wooldridge (1995) shows that a valid test of the null hypothesis of no selection is a t-statistic on the coefficient of $\hat{\lambda}_{it}^{R}$.

In line with the retirement literature that characterizes partial and full retirement as two separate statuses (Honig and Hanoch (1985a)), I also consider an alternative specification for the selection process that allows for observables to affect the two decisions differentially. In this case, selection is based on two indices, d^{PR} and d^{FR}, which equal 1 if the individual is partially and fully retired, respectively. Full-time wages are observed when both indicators are equal to 0. The specifications for the selection indices are analogous to (3), with the residuals in the two probit equations assumed jointly normal conditional on Z_i. The first stage is estimated by multinomial probit. Under the assumption of independent idiosyncratic error terms in the partial and full retirement equations, the two inverse Mill's ratios, $\hat{\lambda}_{it}^{PR}$ and $\hat{\lambda}_{it}^{FR}$, enter additively the wage equation and their coefficients are used to perform the selection test.

First-stage results.

Results from the first-stage regressions are shown on table 3. The first 2 columns show estimates for the binomial probit specification, and the next 4 for the multinomial probit one. For ease of interpretation, marginal effects are reported next to the coefficients for each variable. As has been done before in the literature, the social security early and full retirement ages (ERA and FRA) are used as exclusion restrictions.[20] At the ERA and FRA, individuals face strong incentives to reduce their labor supply (see Aaronson and French (2004) for a detailed discussion). At the same time, these exogenous policy variables can be plausibly excluded from the second stage as long as employers are not systematically lowering

[18]No further functional form assumption is required on the joint distribution of ε_{it} and v_{it} in this context where the aim is to test (and not to correct) for sample selection.

[19]As discussed in Wooldridge (1995), equation (3) can be viewed as the reduced form of a selection equation with an unobserved, time-constant individual effect which is specified as a correlated random effect (as in Chamberlain (1980)), where $\bar{Z}_i$ is intended to capture its correlation with observables.

[20]See, for example, Aaronson and French (2004) and Bernheim, Skinner and Weinberg (2001). The ERA is 62 for all individuals in my sample, while the FRA varies from 65 years to 65 years and 8 months. Since I use yearly wage observations, I make no distinction between those who reach the FRA at the beginning or closer to the end of their 65^{th} year.

wages of employees who become eligible for social security benefits. To my knowledge, there is no evidence that this is the case. In order to use the 62 and 65 age dummies as exclusion restrictions, I make the additional assumption that the participation and wage profiles are well approximated by a smooth function of age.

The estimates in table 3 show that the social security ERA and FRA are statistically significant in all specifications and have the expected sign. Specification (1) shows that full-time participation declines by 14 percentage points at the ERA. The results from specification (2) indicate that up to four-fifths of this decline are attributable to transitions into full retirement, and the rest to transitions into part-time work. Instead, the decline in full-time participation at age 65 is mostly explained by transitions into partial retirement.

The rest of the regressors are divided in two groups, namely time-variant and time-invariant ones. The latter, which include time averages of time varying regressors together with variables that are constant in time, such as education, are intended in this framework to capture time-constant, individual effects. Focusing on time-invariant variables first, we can see that average age has a significant, negative impact on both partial and full retirement. Given the design of the the HRS panel, later-born workers enter the sample at younger ages and hence have a lower average age. This result therefore confirms Cahill, Giandrea and Quinn's (2006) finding that later-born workers are more likely to take on bridge jobs, and suggests that the increase in older men's labor force participation over the last two decades (Schirle (2008)) may be partly attributable to increases in part-time participation.

Not surprisingly, the lower an individual's health type -proxied by the average number of periods in which he reports suffering from a health condition that limits his ability to work- the lower the likelihood of full-time work. A one-standard deviation increase in the average health limitation variable is associated with a 2.2% increase in the probability of part-time work and a 5.5% increase in the probability of retirement. Lifetime wealth, measured here by the average of wealth holdings over the period, is not significant in the partial retirement equation, but it is positively associated with full retirement: A one standard deviation increase in lifetime wealth increases the full retirement probability by 4.1%.

College-educated individuals are more likely than those without a high school degree (the omitted category) to work part time, and less likely to fully retire. The corresponding marginal effects for high school graduates go in the same direction, but their magnitude is smaller and they are only significant at 10%. The type of pension plan an individual holds is also a strong predictor of retirement. Workers who have a defined benefit plan in their full-time job are 6 percentage points less likely to partially retire, and slightly more likely to fully retire than those who have no pension (the omitted category). Those with a defined contribution plan in their full-time job are the most likely to continue to work full time. Mother's education, which aims to proxy for individual ability, is not associated with retirement propensities.

Turning now to the time-variant regressors, we can see that the probability of both partial and full retirement increases by around 4 percentage points every year. The onset of a health limitation decreases the probability of partial retirement by 4 percentage points while increasing the probability of full retirement by more than 20 percentage points. Finally, transitory changes in wealth have essentially no effect on retirement transitions.

The conclusion from the first stage results is that retirement decisions are strongly associated with observables. The fact that these same observables correlate with wage levels generates composition effects that lead to the differences between the OLS and fixed-effects wage profiles discussed in section 3.

Second-stage results.

Results for the second stage are presented in table 4. The two columns correspond to the specifications with one and two margins of selection, respectively. In both cases the coefficients on the estimated inverse Mill's ratios are statistically insignificant. In the second column, the p-value for the joint significance test of the two inverse Mill's ratios is 0.61. Thus the null hypothesis of no selection bias cannot be rejected for any specification.

Despite the statistical rejection, it is worth considering whether the magnitude of the coefficients on the estimated Mill's ratios would be negligible in practice. This in important in order to ascertain that the rejection of self-selection is not simply due to a lack of precision in the estimates for those coefficients. In the presence of selection, the inverse Mill's ratio is an omitted variable in the regressions estimated in section 3. To the extent that the omitted variable is correlated with the age coefficients, we would expect to observe changes in the estimates of the age profile once the inverse Mill's ratio is added to regression. Table 4 shows that the age coefficients are significant in the two specifications, indicating a concave profile. However, when we analyze the data graphically (see figure 8) we can see that the concavity results from fitting the initial increase and subsequent flattening of hourly wages with a quadratic polynomial. Indeed, the age-wage profiles from the two second-stage specifications remain very close to the one estimated in section 3 controlling for transitions into partial retirement. Hence there is no evidence that the flat full-time wage profile suffers from omitted variable bias arising from self-selection into retirement as a result of wage

Figure 8: Log Wage Profiles from Second Stage Regression Estimates

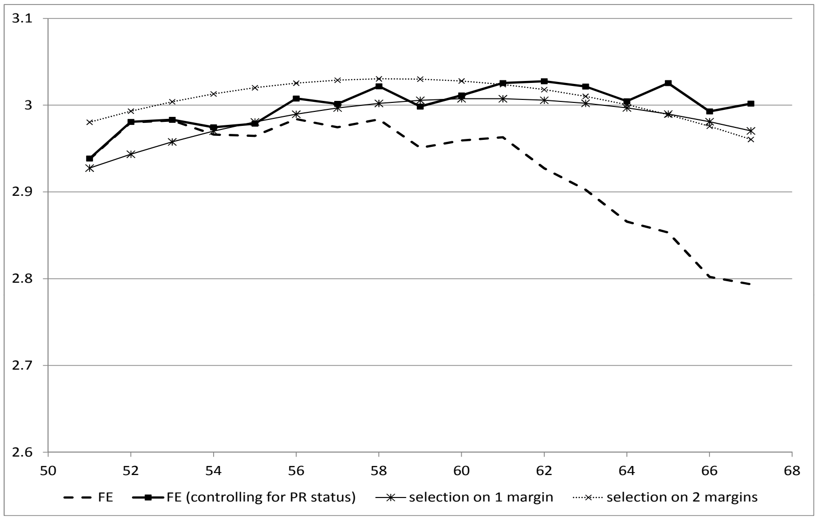

NOTE. - "F.E". and "FE(controlling for PR status)" series constructed using age coefficients from columns 2 and 3 in table 1. "Selection on 1 margin" and "Selection on 2 margins" constructed using age coefficients from columns 1 and 2 in table 4.

shocks.[21]

The evidence presented above indicates that selection on unobserved worker characteristics such as productivity is not an important driver of retirement transitions. A straightforward implication of the absence of selection on unobservables is that the age profile of offered wages is nondecreasing in age.[22]

B. Voluntary vs. Involuntary Retirement

Having established that retirement transitions are not driven by declines in offered wages, I now address the question of whether they are voluntary. In other words, are transitions out of full-time work driven by a relative preference for leisure or by exogenous employment

[21]It is important to keep in mind that the objective of the regressions in this section is to test, but not to correct, for selection. Given that selection on unobservables is rejected, the preferred estimate for the age-wage profile is still the one presented in column 3 of table 1, which uses a less restrictive age specification.

[22]At first sight, this results appears at odds with evidence showing that physical and mental abilities decline with age, potentially leading to a loss of general human capital and thus productivity (see, for example, Agarwall et al. (2009)). The two can be reconciled in two ways. First, older workers may accumulate enough job-specific human capital to compensate for general human capital declines. Under this scenario, job-specific productivity does not decline with age. Second, nominal wage rigidities or long-term contracts may prevent employers from adjusting wages to reflect productivity drops. In this case, older workers' wages would be above their marginal product. Hellerstein, Neumark and Troske (1996) provide evidence consistent with this possibility. Distinguishing between the two explanations for flat offered wages is beyond the scope of this paper.

constraints? Distinguishing between the two is not straightforward, given that retirement transitions are influenced by a myriad of motives and that the definition of "voluntary" retirement is somehow subjective. Consider, for example, two workers who are diagnosed with arthritis. One of them may view his retirement as involuntary, the result of an unexpected health condition that makes work too painful. The second one, who has higher savings and access to a pension plan, may claim that he voluntarily retired to avoid the rising cost of work. The aim of this section is not to provide a formal definition of what constitutes a "voluntary" retirement. Rather, I review several pieces of evidence from different sources that, on the whole, point to a majority of retirements being a choice of workers who had the option but chose not to remain in their full-time jobs.

The first sources of evidence are papers showing that most workers plan to either partially or fully retire at some point in the future (Abraham and Houseman (2004)) and are able to predict their retirement date with high accuracy, even when asked years in advance (Bernheim (1989), Haider and Stephens (2004)). These studies indicate that most retirement transitions are, at least, expected. While this is reassuring, it is not enough to conclude that they are also voluntary.

More direct evidence comes from self reports obtained from surveys of individuals who have already retired. Szinovacz and Davey (2005) find that, when asked directly, 70% of male HRS retirees perceive their retirement as "wanted" rather than "forced." In interpreting these results, we may be concerned that responses to this type of question could be contaminated by justification bias if workers are reluctant to admit that they were coerced out of their job. I complement Szinovacz and Davey's evidence by analyzing a series of questions in the HRS survey regarding the reasons that lead workers to leave a job. These questions are asked of all HRS respondents who change employer between two waves, regardless of retirement status. Since workers are not asked to give a subjective evaluation of the reason why they left their previous job, their answers to these questions are arguably less susceptible to justification bias.

Results are presented in table 5. The first two columns show the reasons for leaving their last employer cited by workers who fully and partially retired between the last two waves, respectively. For comparison, the third and fourth columns show the answers of workers who moved between full-time jobs (i.e., they changed employer but did not retire) and those who were already retired in the first wave (most of these workers are transiting from part-time jobs into full retirement). I group the reasons for leaving an employer according to how likely they are to indicate an involuntary job separation. Focusing first on involuntary transitions, the comparison of columns 1 and 2 shows that full and partial retirees are almost equally likely to have left their last job due to employment constraints (i.e., being laid off or displaced because of business closure) or to negative changes in their working conditions. However, full retirees are four time as likely as partial retirees to have been forced to retire because of poor health or a disability, consistent with the results from the first stage regression estimation in the previous section. Overall, one-third of full retirees and one-fifth of partial retirees report having left their previous full time job for involuntary reasons. Voluntary job separations are explained in a majority of cases (87% of full retirees and 70% of partial retirees) by a desire to enjoy more leisure.

By comparison, full-time workers who switch employer but do not retire (column 3) are much more likely than partial and full retirees to have suffered exogenous changes in their working conditions (mainly layoffs and displacements) and less likely to have had health problems. Finally, transitions from partial into full retirement (column 4) seem to respond to very similar motives as transitions from full-time work into partial retirement.

Even though not shown separately in the table, one of the possible justifications for job transitions is that wages would have been cut had the worker stayed with the previous employer. It is worth pointing out that only 1 out 672 workers who report having entered full retirement voluntarily, and none of those who transited voluntarily into partial retirement, mention this reason for their work transition.

The evidence presented in this section confirms Szinovacz and Davey's (2005) results and indicates that the largest share of full-time work exits is voluntary. The number of involuntary retirements, between 20 and 30% of the total, is far from negligible, and points to the key role of uncertainty in the labor supply decisions of older workers. However, at least 70% of full retirements, and 80% of partial retirements, arise as the optimal choice for workers who could have remained employed full time at their previous wage. For all those workers, the utility associated to increasing leisure outweighs that associated to full-time employment.

The conclusions from this section are twofold. First, the offered full-time wage profile is nondecreasing in age. Second, hours choices, and in particular retirement transitions, are not a response to declining wages. Rather, hours and *ex-post* wages are determined simultaneously.

5 Implications for Structural Estimation

Structural estimation and calibration have a long tradition in economics, and have become more prevalent with the advent of more powerful computers. When estimating or calibrating a structural model, it is common practice to rely on estimates for some of the model primitives that are obtained outside the model. Needless to say, if those externals estimates are biased, this will affect all subsequent estimates and predictions.

The age profile of offered wages is an example of a model primitive that is often estimated outside structural models of labor supply, or taken from other studies in the literature. The age-wage profile of interest for a model that endogenizes labor supply is the one measuring the wage increase or decrease anticipated by a worker, should he remain employed for an extra period. As has been discussed in previous sections, it is common to proxy this *offered* wage profile with the *average* wage profile, even though the two series will generally not overlap, particularly for older workers. In this section, I consider the consequences of such misspecification of the age-wage profile for the estimation of labor supply models. In particular, I study how it affects estimates of the *i.e.s.*.

Because of its crucial role determining intertemporal labor supply allocations, the estimation of the *i.e.s.* has long been a topic of interest in both the labor and macro literatures. Early studies identified the parameter from the covariation in observed hours and wages over the life-cycle (e.g., MaCurdy (1981), Browning, Deaton and Irish (1989) and Altonji (1986)) and concluded that its value was close to zero. More recently, several papers have argued that this covariation does not allow to determine whether labor supply responses (or lack thereof) are driven by the *i.e.s*, or by alternative factors such as liquidity constraints (Domeij and Flodén (2006)), human capital accumulation (Imai and Keane (2004)), or precautionary saving motives (Low(2005)). These papers have structurally estimated the *i.e.s.*, and argued that its value is considerably higher than suggested by the earlier literature.

In what follows, I discuss how sensitive structural estimates of the *i.e.s.* are to misspecification of the lifetime profile for offered wages. I do this in the context of a life-cycle model of consumption and labor supply choices that is consistent with the *voluntary retirement* model described in section 4.[23] I generate artificial data from the model, and use them to estimate the *i.e.s.* for different specifications of the offered wage profile.

[23]Involuntary retirements are ignored to focus on those workers whose retirement decision responds most directly to wages, as opposed to labor supply constraints.

Model

Individuals consumption and leisure to maximize expected discounted utility:

$$\max_{\{c_t\}_{t=t_0}^{T}, \{h_t\}_{t=t_0}^{R<T}} E_{t_0} \sum_{t=t_0}^{T} \beta^{(t-t_0)} \left\{ \frac{c_t^{(1-\rho)}}{1-\rho} + B_t \frac{l_t^{(1-\frac{1}{\gamma})}}{1-\frac{1}{\gamma}} \right\}, \tag{4}$$

where c_t is consumption and l_t is leisure. Leisure is a linear function of hours of market work (h_t) that will be specified below. There are three possible choices for work hours: full-time work (h^F), part-time work (h^P), and zero. This coarse classification is sufficient to capture the choice between full-time work, partial and full retirement faced by older workers. The utility of consumption is a standard CRRA. Consumption and leisure are assumed separable for simplicity, but the marginal rate of substitution between leisure and consumption is allowed to vary with age through the parameter B_t.

Since I will want to compare my results to previous analysis of the *i.e.s.*, I have chosen a standard parameterization for the utility of leisure that follows Chang et al. (2011) and Rogerson and Wallenius (2009 and forthcoming), among others. The parameter γ measures the intertemporal elasticity of substitution of *leisure*. Notice that the intertemporal elasticity of substitution of labor supply is equal to γ multiplied by a non-negative factor. Hence all results regarding biases in estimates of γ translate directly into biases in the *i.e.s.*

Individuals start their life at t_0 and die with certainty in period T. They make consumption and hours decisions until age $R < T$. After age R they become permanently retired, and only choose consumption.

Lifetime utility is maximized subject to the following borrowing constraint:

$$A_{t+1} + c_t = \exp(w_t) h_t + SS_t + (1+r) A_t, \tag{5}$$

where w_t is the log offered hourly wage, A_t are asset holdings, r is the interest rate -assumed constant-, and SS_t represents retirement benefits. The latter are modeled as a fixed proportion of the deterministic component of wages, including the fixed effect. For simplicity, I abstract from issues related to pension eligibility and assume that retirement benefits are available at any age.[24] Individuals cannot borrow against future labor income or retirement income; that is, the borrowing constraint $A_t \geq 0$ is imposed in all periods.

Full-time log wages evolve according to the following process:

[24] Because I only model the behavior of older workers, this assumption implies that retirement coverage is available from age 50 onwards.

$$w_{it} = f_i + W(t) + u_{it},$$
$$u_{it} \sim \text{Normal}(0, \sigma_u), \tag{6}$$

where f_i is an individual fixed effect and W some function of age.[25] This specification is consistent with the equations estimated in section 3, and is also common in the literature (see, for example, French (2005)).[26] There is a wage penalty associated to part-time work, which is specified as in Chang et al. (2011): part-time workers' hourly wage is equal to $(1 - \alpha)w_{it}$.

So far, the model closely resembles others that have been used to study the sensitivity of estimates of the *i.e.s.* to different modeling assumptions. Chang et al. (2011), for example, use a similar model -with infinitely-lived households- to match the average part-time participation rate of prime-age men and women. When focusing on prime-age workers, a steady-state equilibrium concept is appropriate. In this case, instead, the focus is on the labor supply changes that take place over the retirement period. In order to capture these changes, I add two additional features to the model.

First, the parameter B_t, which in steady-state models is assumed constant, is allowed to change with age. As in French and Jones (2011), I assume that preferences for leisure increase linearly with age, that is:

$$B_t = b_0 + b_1 t \tag{7}$$

The increasing taste for leisure is enough to fit the declining trend in full-time participation with age, but not to fit both the proportion of workers that are retired and those working part time. Hence I also introduce a cost of work that varies with age and with the number of hours worked, i.e., it is different for part- and full-time workers.[27] This is con-

[25] In the model, age and time are indistinguishable.

[26] An alternative representation of the error assumes a random walk process instead of a fixed effect (Meghir and Pistaferri (2004)), allowing for persistent shocks to wages. This model focuses on individuals approaching retirement, who will be employed for a short number of years. With a short time horizon, the effects of persistent and transitory shocks to wages are similar, and for simplicity I choose to focus on the latter only.

[27] Chang et al. (2011) exploit a different nonconvexity to fit part-time participation rates. Specifically, they calibrate a different value for the part-time wage penalty for each of the baseline values of the parameter γ. Here the value of the part-time wage penalty is fixed at the level estimated in section 3. I use the differential cost of work for part- and full-time workers as an alternative way of varying the relative value of part- and full-time work.

sistent with plenty of evidence from the retirement literature suggesting that workers value the flexibility associated to part-time jobs increasingly as they age (Hurd, 1996, Abraham and Houseman, 2004). As in French (2005) and French and Jones (2011), the cost of work is modeled as a loss of leisure. The amount of leisure enjoyed in every period is given by:

$$l_t = L - h_t - \phi_t, \tag{8}$$

where L is the leisure endowment, and the cost of work ϕ_t is parameterized as:

$$\phi_t = q_0 + q_1 t + q_2 h_t + q_3 h_t t, \quad \text{with } \phi_t \in [0, L - h_t] \tag{9}$$

B. Calibration

I consider 4 baseline economies characterized by the value of γ, assumed to be equal to 0.25, 0.5, 0.75, and 0.95 in each one of them, respectively.

Some parameter values are common to all baseline scenarios, and their values are taken from existing estimates. The set of external parameters is shown in table 8. The coefficient of relative risk aversion, ρ, is set equal to 2. This implies a value for the intertemporal elasticity of substitution that is consistent with the evidence provided by Attanasio and Weber (1995). The yearly interest rate is fixed at 3%. The discount factor is set at at 0.95, somewhat lower than estimates such as French (2005) because these papers account separately for mortality risk.

Since the model is intended to simulate the behavior of individuals approaching retirement, the initial age t_0 is taken to be 50. The maximum retirement age R is set at 75, and the terminal age T is 90. Retirement benefits are assumed to replace 40% of the deterministic component of wages, and the part-time penalty α is set at 0.3, consistent with the estimate obtained in section 3.

Consistent with the findings in the previous sections, the age wage profile of offered wages is taken to be constant in the baseline specifications. In order to match the average full-time hourly wage in the data, which is approximately \$20, I set $W(t) = W = \ln(20) = 3$.

This leaves 6 remaining parameters, the 2 taste for leisure parameters b_0 and b_1 and the 4 cost of work parameters q_0, q_1, q_2 and q_3. They are calibrated to fit 2 sets of employment targets. The first one is the proportion of individuals that are fully retired at every age between 51 and 69. The second one is the proportion of workers who are employed part time at every age between 51 and 69. The last set of moments is key in this context, because the increasing proportion of part-time workers with age is the driver of the declining wage profiles often estimated in the literature. Since there are more moment conditions (38) than parameters (6), I use a standard minimum distance criterion to obtain the calibrated values.

The 6 parameters of interest are calibrated separately for each baseline value of γ. Results are presented in table 9, and the implied values for the taste of work and the difference between the cost of full-time and part-time work are plotted in figure 9. The graph shows that the taste for leisure and the relative cost of full-time versus part-time work increase with age in all baseline economies.

Figure 10 shows the predicted proportions of workers employed full time, part time and retired (solid lines) versus the corresponding ones in the data (dashed lines) for the 4 baseline economies. The figure shows that this coarse model is able to provide a good fit for the main participation trends.

C. Results

I use the model to generate hours and (if employed) wage data for each individual at every age. Using the same definition of partial retirement, and the same age criterion for sample selection as in section 3, I re-run the wage regressions in all baseline scenarios. Figure 11 shows the age-wage coefficients for the baseline economy where $\gamma = 0.5$ -the results are very similar for the others. As was the case when using HRS data, I find that the age-wage profile from a fixed-effects regression that does not control for participation status declines with age. Once the partial retirement indicator is included as an additional regressor, the estimated age-wage profile is completely flat. The first row of table 10 shows the estimated coefficients on the partial retirement indicators for each of the baseline economies. The wage drop associated to the transition into partial retirement is approximately 34%. The next row shows the equivalent coefficients from the hours regressions. The transition into partial retirement is associated with a 56% decline in hours. These results follow from the values assumed for the part-time penalty and the number of hours worked by full- and part-time workers.

The next step is to investigate how estimates of the parameter γ are affected by mis-specification of the age-wage profile. In particular, I will use the smoothly declining wage profile obtained from regressions that pool observations of full- and part-time workers as a proxy for the offered wage profile. This is the procedure commonly used in the literature. To maintain the same average wage at all ages, in the counterfactual the part-time wage penalty is set to zero.[28] I fix all other model parameters except for γ to their baseline values.

I estimate γ by matching full-time and part-time participation rates generated in the counterfactual simulations to those obtained for the baseline simulations. Row (5) in table 10 shows the results. All the estimated values are higher than the baseline value of the parameter, and the bias increases from close to 30% in the first column to over 130% in the last one. Row (3) shows that in this scenario wages do not change upon entrance into partial retirement -all declines are attributed to age- and, given the maintained assumption of discrete labor supply, row (4) indicates that hours still change discretely upon transition out of full-time work. The coefficients in rows (3) and (4) provide the intuition for the upward bias in the counterfactual estimates of γ. In the counterfactual economies, the same drops in hours have to be generated in the absence of the large wage changes observed at baseline. These hour changes are attributed to a high elasticity of leisure that interacts with the increasing taste for leisure to produce exits from full-time work. Thus the estimates of γ obtained using the misspecified age-wage profile are biased upwards.

6 Conclusions

This paper shows that the offered hourly wage profile does not decline at older ages. The hump-shaped profile often estimated in the literature results from pooling observations of full- and part-time workers, and is not a good proxy for the wage path followed by the typical individual.

This finding has important implications for structural models of labor supply, which often use the hump-shaped estimates found in the literature as a proxy for the age path of offered wages. I illustrate the consequences of misspecifying the offered wage profile by showing that it leads to biases of up to 130% in estimates of the intertemporal elasticity of substitution of labor supply.

My results have implications for all results in the labor supply literature that rely on the correct specification of the age-wage profile, including measures of income risk; studies of retirement decisions; computations of present discounted values of private and social security pensions; and results based on estimates of the i.e.s., from labor supply responses to business cycles to life-cycle patterns of work and leisure.

[28]Keeping the part-time penalty makes part-time wages lower, on average, in the counterfactual economies, which are then not directly comparable to their baseline counterparts. However, to show that my results are not driven by the absence of the part-time penalty, I have estimated γ in the counterfactual scenario maintaining the part-time penalty, and the direction and size of the bias is similar.

Moreover, the results are also relevant to models of life-cycle consumption and savings. These models generally do not endogenize labor supply, and they use labor earnings, rather than hourly wages, as a model primitive. They often take estimates of earnings profiles from external sources. As was the case for wages, the paper shows that the smoothly declining earnings profile often estimated in the literature is not a good proxy for the earnings path of the typical older worker. Misspecification of the earnings profile is bound to affect the estimates and predictions obtained from these models.

References

[1] Aaronson, D. and E. French. 2004. "The Effect of Part-Time Work of Wages: Evidence from the Social Security Rules". *Journal of Labor Economics*. Vol. 22, pp. 329-352.

[2] Abraham, K. and S. Houseman. 2005. "Work and retirement plans among older Americans", in R. Clark and O. Mitchell (eds.), *Reinventing the Retirement Paradigm*, Oxford: Oxford University Press.

[3] Agarwal, S., J. Driscoll, X. Gabaix, X. and D. Laibson. 2009. "The Age of Reason: Financial Decisions over the Life Cycle and Implications for Regulation", *Brookings Papers on Economic Activity*, pp. 51-117.

[4] Altonji, J. G. 1986. "Intertemporal Substitution in Labor Supply: Evidence from Micro Data". *Journal of Political Economy*. Vol. 94(3), pp. S176-S215.

[5] Attanasio, O. and G. Weber. 1995. "Is Consumption Growth Consistnent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey". *Journal of Political Economy*. Vol. 103, pp. 1121-1157.

[6] Baker, M. 1997. "Growth-Rate Heterogeneity and the Covariance Structure of Life-Cycle Earnings". *Journal of Labor Economics*. Vol. 15(2), pp. 338-375

[7] Becker, G. S. 1993. *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education. 3rd Edition.* The University of Chicago Press.

[8] Ben-Porath, Y. "The Production of Human Capital and the Life-Cycle of Earnings". *Journal of Political Economy*. 75 (1967) 352-365.

Bernheim, B. D. 1989. "The Timing of Retirement: A Comparison of Expectations and Realizations", in D. A. Wise (ed.), *The Economics of Aging*, Chicago and London: Unversity of Chicago Press.

[9] Bernheim, B. D., J. Skinner and S. Weinberg. 2001. "What Accounts for the Variation in Retirement Wealth among U.S. Households?" *The American Economic Review*. Vol. 91(4), pp. 832-857

[10] Blau, D. 2008. "Retirement and Consumption in a Life Cycle Model". *Journal of Labor Economics*. Vol. 26(1), pp. 35-71.

[11] Browning, M., A. Deaton and M. Irish. 1985. "A Profitable Approach of Labor Supply and Commodity Demands over the Life-Cycle". *Econometrica*. Vol. 53, pp. 503-544.

[12] Cahill, K. E., M. D. Giandrea and J. F. Quinn. 2006. "Retirement Patterns From Career Employment". *The Gerontologist*. Vol. 46(4), pp. 514-523

[13] Chamberlain, G. 1980. "Analysis of Covariance with Qualitative Data". *Review of Economic Studies*. Vol. 47, pp. 225-238.

[14] Chang, Y., S.-B. Kim, K. Kwon, and R. Rogerson. 2011. "Interpreting Labor Supply Regressions in a Model of Full- and Part-Time Work". *The American Economic Review*. Vol. 101, pp. 476-481.

[15] Domeij, D. and M. Flodén. 2006. "The Labor-Supply Elasticity and Borrowing Constraints: Why Estimates Are Biased". *Review of Economic Dynamics*. Vol. 9(2), pp. 242-262.

[16] French, E. 2005. "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour". *Review of Economic Studies*. Vol. 72, pp. 395-427

[17] French, E. and J. B. Jones. 2011. "The Effects of Health Insurance and Self-Insurance on Retirement Behavior". *Econometrica*. Vol. 79(3), pp. 693-732

[18] Ghez, G. and G. S. Becker. 1975. "The Allocation of Goods Over the Life Cycle". In *The Allocation of Time and Goods over the Life Cycle*. New York: National Bureau of Economic Research, Columbia University Press.

[19] Gomes, F. J., L. J. Kotlikoff, L. M. Viceira. 2008. "Optimal Life-Cycle Investing with Flexible Labor Supply: A Welfare Analysis of Life-cycle Funds". *The American Economic Review*. Vol. 98 (2), pp. 291-303. Papers and Proceedings of the One Hundred Twentieth Annual Meeting of the American Economic Association.

[20] Gourinchas, P. O. and J. A. Parker. 2002. "Consumption Over the Life Cycle". *Econometrica*. Vol. 70(1), pp. 47-89.

[21] Gustman, A. and T. Steinmeier. 1985. "The Effect of Partial Retirement on Wage Profiles of Older Workers". *Industrial Relations*. Vol. 24, pp. 257-265.

[22] Guvenen, F. 2009. "An empirical investigation of labor income processes". *Review of Economic Dynamics*. Vol. 12(1), pp. 5879

[23] Haider, S. J. and M. Stephens Jr. 2007. "Is there a retirementconsumption puzzle? Evidence using subjective retirement expectations". *Review of Economics and Statistics*. Vol. 84, pp. 24764.

[24] Heathcote, J., K. Storesletten and G. L. Violante. 2010. "The Macroeconomic Implications of Rising Wage Inequality in the United States". *Journal of Political Economy*. Vol. 118(4), pp. 681-722.

[25] Heckman, J. J. 1976. "A Life-Cycle Model of Earnings, Learning, and Consumption". *Journal of Political Economy*. Vol. 84(4), pp.S11-S44.

[26] Heckman, J. J. 1979. "Sample selection bias as a specification error". *Econometrica*. Vol. 47, pp. 153-61.

[27] Heckman, J. J., L. J. Lochner and P. Todd. 2006. "Fifty Years of Mincer Earnings Regrressions". In Hanushek, E. and F. Welch, *Handbook of Education Economics. Vol. 1*. (The Netherlands: Elsevier)

[28] Hellerstein, J. K., D. Neumark and K. R. Troske. 1996. "Wages, Productivity and Worker Characteristics: Evidence From Plant Level Production Function and Wage Equations". *Journal of Labor Economics*. Vol. 17, pp. 409446

[29] Honig, M. and G. Hanoch. 1985a. "Partial Retirement as a Separate Mode of Retirement Behavior". *The Journal of Human Resources*. Vol. 20, pp. 21-46.

[30] Honig, M. and G. Hanoch. 1985b. "'True' Age Profiles of Earnings: Adjusting for Censoring and for period and Cohort Effects". *The Review of Economics and Statistics*. Vol. 67, pp. 383-394.

[31] Hurd, M. 1971. "Changes in Wage Rates Between 1959 and 1967, *The Review of Economics and Statistics*. Vol. 53, pp. 189-199.

[32] Hurd, M. 1996. "The Effect of Labor Market Rigidities on the Labor Force Behavior of Older Workers" in D. A. Wise (ed.), *Advances in the Economics of Aging*, The University of Chicago Press

[33] Imai, S. and M. P. Keane. 2004. "Intertemporal Labor Supply and Human Capital Accumulation". *International Economic Review*. Vol. 45(2), pp. 601-641.

[34] Johnson, R. W. and D. Neumark. 1996. "Wage Declines among Older Men. *The Review of Economics and Statistics*. Vol. 78, pp. 740-748.

[35] Kopecky, K. A. 2011. "The Trend in Retirement". *International Economic Review*. Vol. 52(2), pp. 287316

[36] Krueger, D. and F. Kubler. 2006. "Pareto-Improving Social Security Reform when Financial Markets Are Incomplete?". *The American Economic Review*. Vol. 96(3). pp. 737-755.

[37] Lillard, L. A. and R. J. Willis. 1978. "Dynamic Aspects of Earning Mobility". *Econometrica*. Vol. 46, pp. 985-1012.

[38] Low, H. W. 2005. "Self-insurance in a life-cycle model of labour supply and savings". *Review of Economic Dynamics*. Vol. 8(4), pp. 945-975.

[39] Lucas, R. E., Jr., and L. A. Rapping. 1969. "Real Wages, Employment, and Inflation." *Journal of Political Economy*. Vol. 77(5), pp. 721-754.

[40] MaCurdy, T. E. 1989. "An Empirical Model of Labor Supply in a Life-Cycle Setting". *Journal of Political Economy*. Vol. 89, pp. 1059-85.

[41] Meghir, C. and L. Pistaferri. 2004. "Income Variance Dynamics and Heterogeneity". *Econometrica*. Vol. 72(1), pp. 1-32.

[42] Mincer, J. 1974. *Schooling, Experience, and Earnings*. New York: Columbia University Press.

[43] Murphy, K. and F. Welch. 1990. "Empirical Age-Earnings Profiles". *Journal of Labor Economics*. Vol. 8(2), pp. 202-229

[44] Murphy, K. and F. Welch. 1992. "The Structure of Wages". *Quarterly Journal of Economics*. Vol. 107(1), pp. 285-326.

[45] "RAND HRS Data, Version J". Produced by the RAND Center for the Study of Aging, with funding from the National Institute on Aging and the Social Security Administration. Santa Monica, CA, 2010.

[46] Rogerson, R. and J. Wallenius. 2009. "Micro and Macro Elasticities in a Life Cycle Model with Taxes", *Journal of Economic Theory*. Vol. 144, pp. 2277-2292.

[47] Rogerson, R. and J. Wallenius. Forthcoming. "Nonconvexities, Retirement and the Elasticity of Labor Supply". *American Economic Review*.

[48] Ruhm, C. J. 1990. "Bridge Jobs and Partial Retirement". *Journal of Labor Economics*. Vol. 8, pp. 482-501.

[49] Rupert, P. and G. Zanella. 2012. "Revisiting Wage, Earnings and Hours Profiles". UCSB Working Paper.

[50] Rust J. and C. Phelan. 1997. "How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets". *Econometrica*. Vol. 65, pp. 781831.

[51] Scholz, J. K., A, Seshadri and S. Khitatrakun. 2006. "Are Americans Saving Optimally for Retirement?". *Journal of Political Economy*. Vol. 114(4), pp. 607-643.

[52] Schirle, T. 2008. "Why Have the Labor Force participation Rates of Older Men Increased since the Mid-1990s?". *Journal of Labor Economics*. Vol. 26(4), pp. 549-594.

[53] Storesletten, K., C. I. Telmer and A. Yaron. 2004. "Cyclical Dynamics in Idiosyncratic Labor Market Risk". *Journal of Political Economy*. Vol. 112(3), pp. 695-717

[54] Szinovacz, M. E. and A. Davey. 2005. "Predictors of Perceptions of Involuntary Retirement". *The Gerontologist*. Vol. 45, pp. 36-47.

[55] Welch, F. 1979. "Effects of Cohort Size on Earnings: The Baby Boom's Babies Financial Bust". *Journal of Political Economy*. Vol. 87, pp. 65-98.

[56] Willis, R. J. 1985. "Wage Determinants: A Survey and Reinterpretation of Human Capital Earnings Functions. In O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics*. North Holland, Amsterdam.

[57] Wooldridge, J. M. 1995. "Selection Correction for Panel Data Models Under Conditional Mean Independence Assumptions". *Journal of Econometrics*. Vol. 68, pp. 115-132.

Table 1: Age-Log Wage Profiles for Different Specifications

	Dependent Variable				
	w_{it}	w_{it}	w_{it}	w_{it}^{FT}	w_{it}^{PT}
	OLS	FE	FE	FE	FE
	(1)	(2)	(3)	(4)	(5)
PR=1			-0.337***		
			(0.025)		
age≥52	0.021	0.042	0.042	0.035	
	(0.039)	(0.030)	(0.029)	(0.023)	
age≥53	-0.040	0.002	0.002	0.003	
	(0.034)	(0.027)	(0.025)	(0.019)	
age≥54	0.018	-0.016	-0.009	0.011	
	(0.030)	(0.022)	(0.020)	(0.016)	
age≥55	-0.003	-0.002	0.004	0.007	
	(0.029)	(0.020)	(0.019)	(0.015)	
age≥56	0.011	0.019	0.029*	0.021	
	(0.025)	(0.017)	(0.016)	(0.012)	
age≥57	-0.013	-0.009	-0.006	0.010	
	(0.024)	(0.016)	(0.015)	(0.012)	
age≥58	-0.007	0.009	0.020	-0.003	0.072
	(0.024)	(0.015)	(0.015)	(0.012)	(0.074)
age≥59	-0.019	-0.033**	-0.023	0.003	-0.087
	(0.024)	(0.016)	(0.015)	(0.012)	(0.063)
age≥60	0.002	0.008	0.013	0.012	-0.124**
	(0.024)	(0.017)	(0.015)	(0.012)	(0.060)
age≥61	-0.019	0.004	0.014	0.001	0.136**
	(0.024)	(0.016)	(0.016)	(0.013)	(0.062)
age≥62	-0.044	-0.036**	0.002	-0.002	0.002
	(0.028)	(0.018)	(0.017)	(0.015)	(0.049)
age≥63	-0.032	-0.025	-0.006	0.014	0.002
	(0.033)	(0.021)	(0.020)	(0.017)	(0.041)
age≥64	-0.080**	-0.037	-0.017	-0.024	-0.007
	(0.036)	(0.022)	(0.021)	(0.018)	(0.040)
age≥65	-0.006	-0.013	0.021	0.014	0.072*
	(0.041)	(0.026)	(0.024)	(0.023)	(0.037)
age≥66	-0.082*	-0.051	-0.033	-0.007	-0.076*
	(0.047)	(0.033)	(0.031)	(0.029)	(0.041)
age≥67	0.014	-0.008	0.009	-0.044	0.049
	(0.048)	(0.035)	(0.033)	(0.035)	(0.041)
individual-year obs.	7,915	7,500	7,500	6,277	830
# of individuals		1,834	1,834	1,666	287
R^2	0.22	0.05	0.15	0.01	0.02
Tests of Joint Significance (p-value):					
Age≥52-Age≥60	0.695	0.059	0.08	0.012	
Age≥61-Age≥67	0.000	0.000	0.618	0.733	0.01

NOTE. - Robust standard errors in parentheses. *, ** and *** indicate significance at 10, 5, and 1%, respectively. All regressions include an intercept and a measure of the unemployment rate to capture yearly changes in labor market conditions. OLS regression includes controls for marital status, education (defined as less than high school/high school/college), and race (defined as black/hispanic/white). OLS standard errors clustered at the individual level. For fixed-effects estimates, reported R^2 is for the within regression. Dummies for ages 52 to 57 are eliminated from PT wage regression because of the low number of observations available.

Table 2: Age-Log Hours Profiles for Different Specifications

	Dependent Variable				
	$hours_{it}$	$hours_{it}$	$hours_{it}$	$hours_{it}^{FT}$	$hours_{it}^{PT}$
	OLS	FE	FE	FE	FE
	(1)	(2)	(3)	(4)	(5)
PR=1			-0.642***		
			(0.025)		
age≥52	-0.032**	-0.030	-0.030	-0.013	
	(0.016)	(0.022)	(0.021)	(0.013)	
age≥53	0.024*	0.038*	0.038**	0.002	
	(0.013)	(0.020)	(0.019)	(0.010)	
age≥54	-0.002	-0.024	-0.010	0.002	
	(0.011)	(0.018)	(0.015)	(0.009)	
age≥55	-0.006	-0.000	0.011	-0.007	
	(0.012)	(0.017)	(0.014)	(0.009)	
age≥56	-0.016	-0.027*	-0.009	-0.013*	
	(0.011)	(0.016)	(0.013)	(0.008)	
age≥57	-0.007	0.009	0.015	-0.001	
	(0.011)	(0.016)	(0.013)	(0.007)	
age≥58	-0.007	-0.019	0.002	-0.008	0.068
	(0.013)	(0.017)	(0.014)	(0.008)	(0.075)
age≥59	-0.012	-0.002	0.016	0.011	-0.042
	(0.014)	(0.018)	(0.015)	(0.008)	(0.081)
age≥60	-0.008	-0.026	-0.018	-0.015**	0.037
	(0.015)	(0.017)	(0.014)	(0.007)	(0.064)
age≥61	0.004	-0.003	0.017	0.004	-0.017
	(0.014)	(0.017)	(0.014)	(0.007)	(0.066)
age≥62	-0.105***	-0.106***	-0.034**	-0.021***	-0.049
	(0.021)	(0.021)	(0.016)	(0.008)	(0.061)
age≥63	-0.055*	-0.065**	-0.030	0.002	-0.109
	(0.030)	(0.029)	(0.022)	(0.009)	(0.063)
age≥64	-0.048	-0.034	0.004	-0.002	0.045
	(0.037)	(0.033)	(0.026)	(0.011)	(0.061)
age≥65	-0.048	-0.060*	0.004	-0.013	0.004
	(0.041)	(0.035)	(0.029)	(0.013)	(0.057)
age≥66	-0.032	-0.027	0.008	-0.008	0.005
	(0.047)	(0.040)	(0.034)	(0.016)	(0.054)
age≥67	-0.012	-0.015	0.019	0.011	0.077
	(0.050)	(0.043)	(0.037)	(0.018)	(0.056)
individual-year obs.	7,915	7,500	7,500	6,277	830
# of individuals		1,834	1,834	1,666	287
R^2	0.10	0.15	0.42	0.02	0.02
Tests of Joint Significance (p-value):					
Age≥52-Age≥60	0.000	0.021	0.341	0.043	
Age≥61-Age≥67	0.000	0.000	0.000	0.001	0.01

NOTE. - Robust standard errors in parentheses. *, ** and *** indicate significance at 10, 5, and 1%, respectively. All regressions include an intercept and a measure of the unemployment rate to capture yearly changes in labor market conditions. OLS regression includes controls for marital status, education (defined as less than high school/high school/college), and race (defined as black/hispanic/white). OLS standard errors clustered at the individual level. For fixed-effects estimates, reported R^2 is for the within regression. Dummies for ages 52 to 57 are eliminated from PT wage regression because of the low number of observations available.

Table 3: First-stage binomial (1) and multinomial (2) probit estimates.

	(1)		(2)			
	PT+Fully Ret.		PT		Fully Ret.	
	Coeff	Mg Effects	Coeff	Mg Effects	Coeff	Mg Effects
Time-Varying Regressors						
age	-0.118	0.040***	0.167	0.035***	-0.271	0.042***
	(0.191)	(0.002)	(0.271)	(0.004)	(0.233)	(0.003)
age^2	0.002		0.000		0.004**	
	(0.002)		(0.002)		(0.002)	
health limit	0.672***	0.165***	0.252***	-0.043***	1.107***	0.206***
	(0.052)	(0.011)	(0.084)	(0.008)	(0.068)	(0.013)
current wealth (in 100k)	0.022	0.005	-0.001	0.002	0.044	0.007*
	(0.021)	(0.003)	(0.032)	(0.004)	(0.030)	(0.004)
current wealth^2	-0.000		0.002		-0.001	
	(0.001)		(0.002)		(0.002)	
Exclusion Restrictions						
over ERA	0.531***	0.143***	0.564***	0.025**	0.768***	0.119***
	(0.060)	(0.015)	(0.082)	(0.010)	(0.074)	(0.015)
over FRA	0.150*	0.036**	0.263***	0.022**	0.168*	0.014
	(0.079)	(0.016)	(0.097)	(0.011)	(0.087)	(0.015)
Time-Constant Regressors						
avg. age	1.185***	-0.024***	1.578**	-0.027***	1.544***	-0.023***
	(0.270)	(0.003)	(0.617)	(0.004)	(0.538)	(0.003)
avg. age^2	-0.011***		-0.014***		-0.014***	
	(0.002)		(0.005)		(0.004)	
avg. health limit	0.675***	0.158***	0.427**	0.077**	1.090***	0.189***
	(0.077)	(0.024)	(0.182)	(0.033)	(0.145)	(0.025)
avg. wealth (in 100k)	0.073***	0.011**	0.044	0.004	0.122***	0.015***
	(0.024)	(0.005)	(0.053)	(0.006)	(0.043)	(0.005)
avg. wealth^2	-0.004**		-0.004		-0.006*	
	(0.002)		(0.004)		(0.003)	
education = college	-0.104**	-0.024	0.113	0.030**	-0.254**	-0.053***
	(0.043)	(0.016)	(0.120)	(0.013)	(0.102)	(0.016)
education = high school	-0.028	-0.007	0.113	0.021*	-0.109	-0.027*
	(0.040)	(0.015)	(0.111)	(0.012)	(0.091)	(0.015)
mother's education	-0.014	-0.003	0.046	0.008	-0.050	-0.009
	(0.015)	(0.006)	(0.042)	(0.008)	(0.035)	(0.006)
pension type = DB	-0.143***	-0.034*	-0.532***	-0.064***	-0.023	0.030*
	(0.045)	(0.019)	(0.123)	(0.013)	(0.106)	(0.016)
pension type = DC	-0.381***	-0.089***	-0.532***	-0.032***	-0.493***	-0.057***
	(0.048)	(0.018)	(0.132)	(0.012)	(0.114)	(0.017)
Obs.	11,877		11,877		11,877	

NOTE. - *, ** and ** indicate significance at 10, 5, and 1%, respectively. Robust standard errors reported for regression coefficients. Standard errors for marginal effects are computed using 1,000 bootstrap replications clustered at the individual level in order to account for the panel structure of the data. All regressions include an intercept and a measure of the unemployment rate to capture yearly changes in labor market conditions. Observations with wealth above $2M (1.31% of the sample) are dropped from specifications where wealth appears as regressor.

Table 4: Second-stage estimates.

	(1)	(2)
age	0.107***	0.109***
	(0.032)	(0.036)
age^2	-0.001***	-0.001***
	(0.000)	(0.000)
health limit	-0.031*	-0.016
	(0.017)	(0.022)
wealth (in 100k)	0.005	0.006
	(0.005)	(0.005)
wealth^2	-0.000	-0.000
	(0.000)	(0.000)
$\hat{\lambda}^R$	-0.012	
	(0.029)	
$\hat{\lambda}^{PR}$		-0.057
		(0.057)
$\hat{\lambda}^{FR}$		0.028
		(0.035)
individual-year observations	6,244	6,244
# of individuals	1,661	1,661

NOTE. - *, ** and ** indicate significance at 10, 5, and 1%, respectively. Standard errors (in parentheses) computed from 1,000 bootstrap replications clustered at the individual level to account for panel structure of the data. The first and second stages are bootstrapped jointly to obtain consistent estimates of the standard errors on inverse Mill's ratios in the second stage. All regressions include an intercept and a measure of the unemployment rate to capture yearly changes in labor market conditions. Observations with wealth above $2M (1.31% of the sample) are dropped from specifications where wealth appears as regressor.

Table 5: Reasons left last employer (in percentage points)

	i not retired in $t-1$			i retired in $t-1$
	i fully retired in t	i partially retired in t	i not retired in t	
Business closed	4.20	3.97	18.30	6.31
Laid off/let go	9.13	10.60	24.55	13.46
Retired, negative change in working conditions	2.92	2.65	0.67	2.33
Poor health/disabled	17.26	4.64	2.23	15.45
Total involuntary retirements	**33.51**	**21.86**		
Leisure	53.15	50.33	11.16	36.05
Voluntary quits	3.20	15.89	34.82	14.62
Financial incentive	4.11	4.64	1.56	1.83
Change in external circumstances	0.91	0.66	2.01	2.16
Total voluntary retirements	**61.37**	**71.52**		
Other	5.11	6.62	4.69	7.81
Obs.	1,095	151	448	602

Source: author's calculations combining answers to the questions "why did you leave [your previous] employer?" and "did your employment situation change in some way that encouraged you to leave?" The category "Retired, negative change in working conditions" includes workers who cite retirement as the reason they left their last employer and, when asked about the change in their working conditions, report that a) their supervisor or coworkers encouraged their departure, b) their wages or hours had or would have been cut if they had stayed, c) they would have been laid off, d) they had new job duties, e) they had to move to a new job location or f) their employer had changed their health insurance. The category "Leisure" includes workers who say they left their employer because they wanted to travel, or those who claim to have retired and report no changes in their working conditions. The category "Voluntary quits" includes workers who report having left their employer because they a) quit, b) found a better job, c) sold their business, or d) handed over responsibilities to other family members. The category "Financial incentive" includes workers who took an early retirement incentive, became eligible for a pension, found it financially advantageous not to work, or found a job with better pay or benefits. The category "Change in external circumstances" includes those who had to stop working to care for a family member, because their family moved, or because they separated or divorced. Finally, the category "Other" includes workers who do not give the reason why they left their previous employer, together with those who provide reasons that are difficult to classify as voluntary or involuntary, namely a) transportation, distance to work, b) respondent/spouse was transferred (but we don't observe which one) and c) didn't get along with coworkers.

NOTE: The proportions in the table are computed using only nonmissing responses to the relevant questions.

Table 6: Robustness checks

	Sample definition	coeff. on PR=1	Joint sig. tests (p-value) Ages 52-60	Ages 61-67	N (total)	N (# ind.)	# partial retirements
0	Baseline sample	-0.337***	0.08	0.62	7,500	1,834	524
	Wage limits						
1	$w_{it} \leq \$50$	-0.347***	0.05	0.88	7,211	1,789	508
2	$w_{it} \leq \$100$	-0.333***	0.043	0.752	7,592	1,849	527
3	$w_{it} \leq \infty$	-0.303***	0.017	0.400	7,653	1,857	532
	Definition of partial retirement						
4	Keep all old-age self-employed (DEF1)	-0.342***	0.025	0.677	8,915	2,155	769
5	Keep PT, old-age self-employed (DEF2)	-0.350***	0.005	0.455	8,436	2,051	665
6	Use tenure changes (DEF3)	-0.278***	0.002	0.146	7,500	1,834	752
	Results by education category						
7	H.S. dropouts	-0.206***	0.819	0.942	1,661	431	104
8	H.S. graduates	-0.322***	0.387	0.700	2,788	683	205
9	(Some) College	-0.412***	0.527	0.196	3,038	716	213

NOTE. - *, ** and ** indicate significance at 10, 5, and 1%, respectively.

Table 7: Descriptive Statistics by Work Status in First Sample Wave. Non self-employed only.

	Fully Retired	PT Workers	FT Workers
Age	57.37	56.26	55.37
	(4.26)	(4.10)	(3.27)
Education (%)			
Less than High School	35.60	22.17	23.18
High School Graduate	34.90	36.79	36.57
(Some) College	29.50	41.04	40.26
Health Status (%)			
Disabled	31.94	2.79	0.15
Bad Health	23.70	18.14	12.05
Good Health	44.36	79.07	87.80
Median Wealth by Health Status (in 2011$)			
Disabled	24,500		
Bad Health	41,000	36,500	67,300
Good Health	123,000	99,000	100,000
Observations	1,152	215	2,663

NOTE. - Wealth is measured at the household level.

Table 8: External Parameters

Parameters		
Risk aversion	ρ	2.0
Discount factor	β	0.95
Yearly interest rate	r	0.03
Initial age	t_0	50
Max. retirement age	R	75
Terminal age	T	90
Retirement benefits replacement rate		0.40
Part time penalty	α	0.30
Age-Wage profile	$W_t\ (= W)$	3.00

Table 9: Calibrated parameters

	$\gamma = 0.25$	$\gamma = 0.50$	$\gamma = 0.75$	$\gamma = 0.95$
b_0	-5.36E-4	4.88E-3	5.45E-3	5.12E-3
b_1	3.18E-4	4.24E-4	4.82E-4	5.99E-4
q_0	8.88E-2	-5.25E-2	-6.25E-2	-1.02E-1
q_1	3.02E-2	-3.29E-3	-8.17E-3	-1.13E-2
q_2	9.73E-2	5.08E-2	2.25E-1	3.75E-1
q_3	-9.35E-2	1.50E-2	2.22E-2	2.56E-2

Table 10: Simulation results

		$\gamma = 0.25$	$\gamma = 0.50$	$\gamma = 0.75$	$\gamma = 0.95$
I. Baseline					
(1)	Δw_{it} upon partial retirement	**-0.343**	**-0.341**	**-0.340**	**-0.340**
(2)	ΔH_{it} upon partial retirement	**-0.563**	**-0.559**	**-0.556**	**-0.555**
II. Declining age-wage profile, discrete labor supply					
(3)	Δw_{it} upon partial retirement	-0.005	-0.003	-0.004	-0.004
(4)	ΔH_{it} upon partial retirement	**-0.579**	**-0.576**	**-0.579**	**-0.578**
(5)	$\hat{\gamma}$	0.321	1.00	1.551	2.203
(6)	GMM criterion	1.14	0.418	0.415	0.481

Appendix - For Online Publication

A.1 Analysis of the Sample of "Previously Retired" Workers

The wage change at the time of separation from the full-time job is not observed for workers who are already fully or partially retired by the time they enter the HRS panel (in what follows, I will refer to these workers as the "previously retired"). Notice that the wage change information may not be recovered even if we were to make use of retrospective SSA earnings records (which are available for a subsample of HRS respondents). Consider a worker who switches from full-time to part-time work in year t. His SSA records would show a decrease in earnings for that year. However, because retrospective hour measures are not available, we would not be able to map the earnings drop into an hourly wage drop.

Table 7 analyzes the characteristics of workers entering the HRS panel for the first time, classified according to their participation status. Workers in the first column are fully retired, those in the second column are partially retired, and those in the third one are working full time. Observations from the first two columns make up the "previously retired" sample, and are excluded from the main analysis.

Not surprisingly, "previously retired" workers are on average older. However, there are other noticeable differences across groups. In general, partially retired workers and full-time workers have similar characteristics, whereas full retirees are very different. The latter are less educated and considerably more likely to be disabled or in bad health. The correlation between health status and net worth is striking. The wealth holdings of disabled individuals are about a fifth of those of individuals in good health. Those in bad health hold between one- and two-thirds as much wealth as those in good health.

The table suggests that the "previously retired" can be divided in two broad categories: one with low wealth and in bad health or disabled, and another with high wealth and good health. We can speculate that workers in the first category are more likely to have left their full-time job due to their health conditions, while those in the second one are more likely to have done so because they had the financial means to afford more leisure in the form of early retirement. It is important to bear in mind that these two motives for retirement are present in the sample used in the main analysis. In particular, the number of disability insurance recipients more than doubles before age 60. So while the empirical sample undersamples workers who retired very early in their career through disability insurance, disability benefit recipients are by no means excluded from the analysis.

A.2 Robustness Checks

Definition of partial retirement

Different studies have used different definitions of partial retirement. Maestas and Zissimopoulos (2010), for example, have pointed out that in some cases self-employment may be a form of partial retirement. In particular, workers who want to reduce hours and cannot do so at their current job, may opt for becoming self-employed. With this consideration in mind, I contemplate two alternative definitions of partial retirement. DEF1 classifies a worker as partially retired the first time he switches from full-time work into part-time work *or* self employment. DEF2 defines partial retirement as a transition from full-time work into part-time work or *part-time* self-employment. DEF0 is the definition used in the main analysis.

Notice that because of differences in the treatment of the self-employed, the 3 definitions impose different restrictions on the sample. DEF0 is the most restrictive. Any worker who reports being self-employed at any point during his career is dropped from the sample. Under DEF2, workers who enter the panel as full-time employees and transit into part-time self-employment before fully retiring are kept in the sample. DEF 1, the least restrictive, keeps observations of workers who enter the panel as full-time employees and transit into either full-time or part-time self-employment.

The 3 definitions also lead to differences in the prevalence of partial retirement. 29% of workers in the DEF0 sample are classified as partially retired in at least one period, versus 36% and 42% for the DEF2 and DEF1 samples, respectively.

Following Cahill, Giandrea and Quinn (2006), I also consider an alternative definition of partial retirement (DEF3) that makes use of job tenure information. In their paper, they consider that a worker becomes partially retired whenever he moves to a new job that he will keep for less than 10 years before fully retiring. This definition imposes stringent data requirements, as it is necessary to follow workers for at least 10 years after their last work transition. In the HRS, this is not possible in most cases. Under DEF3, I assume that all workers who switch jobs in their 50s become partially retired. Given that few of those who switch jobs in their 50s will end up accumulating more than 10 years of tenure in the new job before they retire, DEF3 should be a good approximation to the definition of partial used by Cahill et al. (2006). According to DEF3, 41% of workers become partially retired at some point in their career

As a robustness check, I have re-run the main regressions in the paper using definitions 1 to 3. Results are shown in rows 4 to 6 of table 6. For DEF1 and DEF2, the key results remain virtually unchanged, indicating that the conclusions regarding the wage profile are not sensitive to the treatment of the self-employed. For DEF3, the wage drop associated to entry into partial retirement is several percentage points smaller. Recall that this definition interprets all job-to-job transitions after age 50 as partial retirements. It is likely that some of these late-career job moves are undertaken by workers who are not entering retirement, which explains the dampening of the partial retirement coefficient. All other results regarding the shape of the wage profile stand unchanged; in particular, the absence of wage drops during the retirement years cannot be rejected.

A.3 Graphical Analysis of Calibration Results

Figure 9: Calibrated taste for work (left) and cost-of-work difference between part time and full time (right)

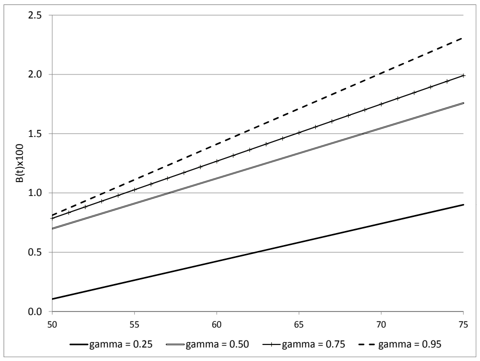

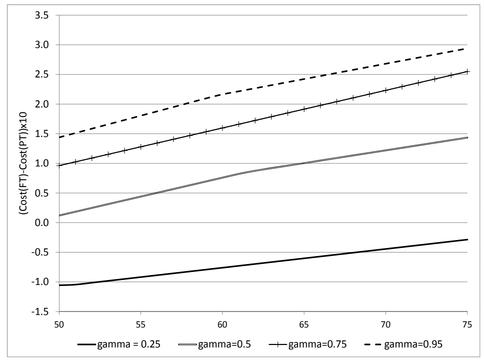

Figure 10: Percentage Employed FT, PT and Retired at Every Age. Data vs. Simulations.

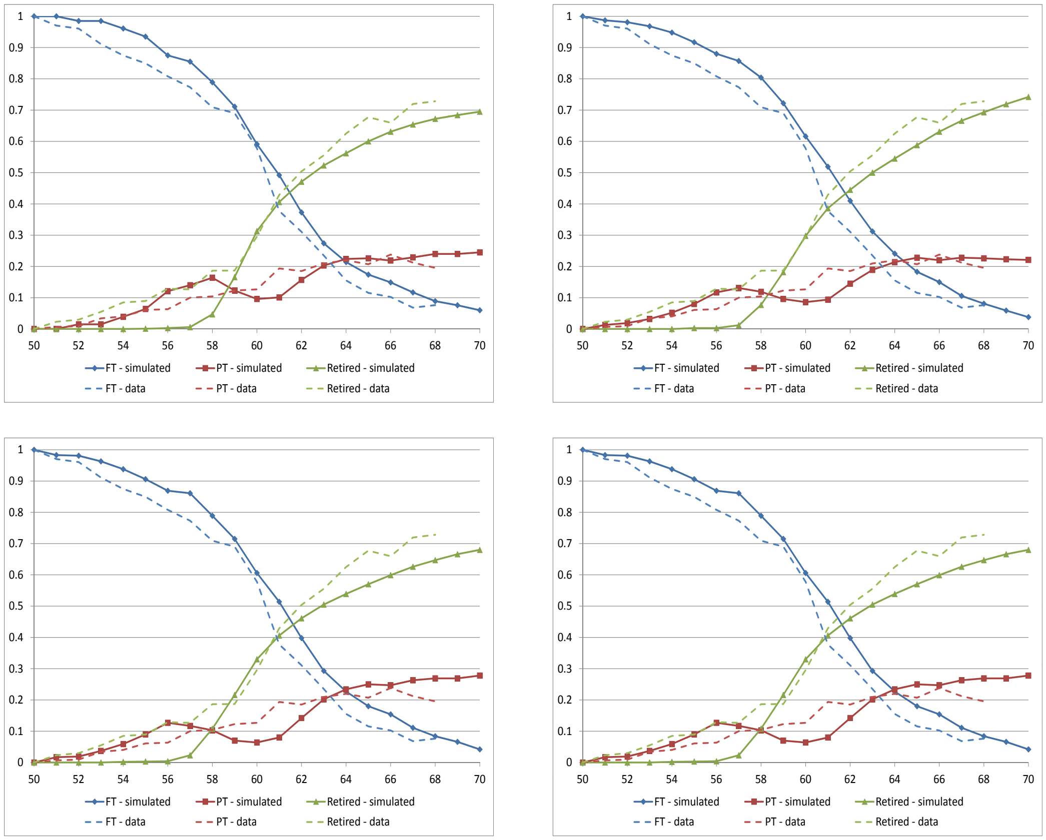

Figure 11: Log Wage Profiles for Different Specifications Using Simulated Data. $\gamma = 0.5$