Tax incentives to retirement saving and intertemporal income smoothing*

Raquel Carrasco[†] Ernesto Villanueva[‡]

February 27, 2023

Abstract

Tax incentives to retirement saving typically imply that monetary amounts invested in retirement saving products are tax exempt, accumulate at a tax free interest rate but are subject to income taxation upon withdrawal. We use longitudinal tax records between 1982 and 1998 to estimate if tax incentives increase the flow of post-retirement income relative to earnings prior to retirement. To obtain exogenous variation in the incentive to use fiscal favored products, we exploit the fact that máximum and mínimum Social Security contribution limits kinks in the relationship between public pension replacement rates and pre-retirement income. We find that tax incentives generate a stream of financial income that absorbs 2/3 of the fall in public pension retirement rates. Namely, while public pension replacement rates fall .7 by each 1% increase in income, total replacement rates only fall by .25%. When tax incentives were not available (prior to 1988) a fall in the public pension resulted in a one-to-one fall in total post-retirement income, indicating a limited ability to self-insure a drop in public pension income.

JEL classification: D14 (personal finance), D15 (intertemporal choice), H31 (fiscal policy: households)

Keywords: Tax incentives, Retirement Saving, Regression Kink Design

*We thank ... Research funding from the Ministry of Science, Innovation and Universities, Grant Nos. RTI2018-095231-B-I00, MDM2014-0431, and Comunidad de Madrid, Grant No. MadEco-CM (S2015/HUM-3444), is gratefully acknowledged.

[†]Department of Economics, Universidad Carlos III de Madrid. E-mail: raquel.carrasco@uc3m.es

[‡]Banco de España. E-mail: ernesto.villanueva@bde.es

1 Introduction

Tax incentives to retirement saving typically take the form of deferred taxation, which means that monetary amounts invested in retirement saving products are tax exempt, accumulate at a tax free interest rate and are subject to income taxation upon withdrawal. The rationale for those incentives is that the retirement decision is complex and that individuals may suffer from myopia. Hence, an incentive to save helps in guaranteeing appropriate standards during old age. However, tax incentives to retirement saving have been criticized on several grounds. The first is that sophisticated investors may reshuffle their portfolios to benefit from the tax treatment without necessarily altering life-cycle saving. The second is that saving behavior exhibits a lot of inertia (Chetty et al. 2014, Haliassos), and that many individuals fail to internalize the incentive to save implicit in deferred taxation. A metric of how successful tax incentives is then the additional saving generated by each euro contributed to tax favored products. The estimates vary across studies, but many of these estimate numbers at about 20 cents or less.

This paper takes a different route and estimates to what extent tax incentives to retirement saving tilt the life-cycle profile of income around retirement. Namely, we pose the following question: to what extent do tax incentives increase the flow of post-retirement income relative to earnings prior to retirement? We focus on the stream of post-retirement income (rather than on private wealth accumulated prior to retirement) because the ultimate objective of tax incentives is to raise the living standards of individuals after retirement. Hence, examining longitudinal data on income over the life-cycle of individuals permits a direct assessment of the costs of tax incentives (foregone revenues for the public sector) to their benefit: the extent to which those contributions to tax favored products increase living standards in the old age.

To obtain exogenous variation in the incentive to use fiscal favored products, we use a long-standing feature of the old-age Social Security system in Spain (and other countries). Public pension replacement rates exhibit kinks both at the maximum and minimum contribution limits during the working life of the worker. During the period we examine (1985-1998), the Spanish public pension was a constant fraction of the average of pre-retirement earnings during the last 8 years prior to retirement. However, once pre-retirement earnings exceeded a maximum contribution limit, further increases in pre-retirement income did not increase the post-retirement public pension -resulting mechanically in a fall of the public pension replacement rate. To the extent that it is predictable, that kink in the relationship between public replacement rates and pre-retirement income generates two different incentives to save during working life. The first is that lower public pension replacement rates generate a fall in post-retirement income that individuals may want to smooth using private savings. The second is that, due to progresivity, deferred taxation increases the return to saving by an amount that is proportional to marginal income taxes prior to retirement income and inversely proportional to income after retirement. A similar mechanism happpens at the *bottom* of the income distribution. Replacement rates at the bottom of the distribution are usually above one, generating little incentive to save to preserve living standards in the old age(as income actually increases once a worker retires). However, once pre-retirement income exceeds the minimum contribution limit, replacement rates fall below one, generating an incentive to save and smooth income over the life cycle.

Our empirical strategy is built around a series of regression kink designs (RKD). First, we leverage on longitudinal data from a 5% sample of all tax returns during the period 1985-1998 to estimate whether the relationship between public pension replacement rates and pre-retirement labor income presents kinks around the maximum and minimum contribution limits. Secondly, we assess if saving in tax-favored products responds to falls in the replacement rate by estimating the relationship of contributions to pension funds and income at both sides of the kink. The third step estimate the relationship between *total* post-retirement income and pre-retirement income around the kink. As total retirement income includes concepts like dividends, rental income or other returns to life-cycle saving, this step allows us to identify how much of the fall in public pension replacement rates detected in step 1 is compensated by contributions to tax-favored retirement products. We conduct a similar analysis at the bottom of the income distribution.

Our findings can be summarized as follows. We document that the relationship between public pension replacement rates and pre-retirement income is flat (around 70%) below the maximum contribution limit and decreases strongly with pre-retirement income above that maximum contribution (1% increase in pre-retirement income reduces replacement rates by .7%). That strong change in the incentive to save for life-cycle purposes is reflected in increases in saving for retirement. An increase of 1% in pre-retirement income increases the propensity to save (amount saved) in tax favored products by 17pp (4.7% of income) below the maximum contribution limit and by 23pp (6.1%) above the limit. Finally, we find that tax incentives generate an additional financial income stream that absorbs 2/3 of the fall in public pension retirement rates. Namely, while public pension replacement rates fall by .7% by each 1% increase in income, total replacement rates only fall by .2. When we compare the public pension replacement rate(.7) with the total income replacement rate (.25) we find that two thirds of the fall in public pension replacement rates is compensated by self-insurance through saving in tax favored saving products (.65=.25/.7-1). Importantly, the relationship between public replacement rates and income (on one hand) and total replacement rates and income (on the other) was very similar for cohorts who retired prior to the introduction of favored saving products. In other words, prior to 1988, a fall in the public pension results in a one-to-one fall in total post-retirement income. In the absence of tax favored products, individuals had a very limited success in generating life-cycle saving that could compensate for the fall in public pension replacement rates.

We contribute to serveral strands of the literature. First, a substantial literature estimates to what extent contributions to tax favored products increases or not pre-retirement wealth. Several studies document very limited contribution of contribution to tax-favored pension products (Chetty et al, 2014, Engen and Gruber), either because savers are very sophisticated and reshuffle their portfolios (Engen and Gruber), or because they are inertial and fail to internalize tax incentives (Chetty et al, 2014). We provide a direct assessment of whether tax incentives achieve the objective of increasing post-retirement income, which is the basis of welfare calculations (see Moser and Olea, 2020). In particular, the fact that we identify a large response of tax incentives may be related to the saliency of the fall in replacement rates around the contribution limits. Studies like Chetty et al rely on changes in marginal taxes, that individuals may or may not understand. On the contrary, the fact that replacement rates fall abruptly around contribution limits, coupled with the availability of tax favored products may trigger saving responses.

Secondly, we contribute to the literature that examines if private saving responds to changes in public pensions (Attanasio and Brugiavini, 2003, Jappelli, Padula, Garcia- Miralles and Leganza 2022). Under some conditions, those studies can identify the crowdout effect of private saving and public pensions. We quantify the degree of actual substitution between public and private post-retirement income. In addition, working with income streams lends naturally to cost-benefit analysis of public vs private provision of retirement income. Indirect utility functions are defined over income (not over saving), so we can evaluate alternative income streams. Our current work is building a set-up to evaluate formally the cost and benefits of those interventions.

2 Data

We use an administrative dataset, the Spanish Panel of Tax Returns, containing a (stratified) 5% random sample of tax returns for the period 1982-1998, which allows us to perform the analysis before and after the introduction of pension plans in Spain. It is a representative sample of the population of taxpayers.[1] The income tax samples are drawn from 15 of the 17 autonomous communities of Spain, in addition to the two autonomous cities, Ceuta and Melilla.[2] Our unit of observation is the individual or married couple, depending on filing status, which can be either single or joint. Single tax returns are filed at the individual level, whereas joint tax returns represent two spouses filing together, or single-parent families with at least one child.[3]

This dataset includes almost the complete set of fiscal and socio-demographic information taxpayers provide in their returns. In particular, it provides detailed information of income from different sources (labor, financial, self-employment), as well as the yearly contributions to pension plans. There is also information about some household characteristics (number of dependent relatives, disability). Nonetheless, there is only limited demographic information (for instance, the Tax Agency provided age only for 60% of sample, so this information has not been used).

To determine whether or not an individual is retired we use information about their social security contributions, which should be zero for retirees. In particular, we classify as retired those individuals for which the social security contributions are smaller than 6 euros per year. Given that the public pension income depends on the individual labor income eight years prior to retirement,

[1] Income tax samples do not include those individuals with no income subject to the tax. Thus, our analysis excludes the population who have either no income or a very low income.

[2] Two autonomous regions, Basque Country and Navarra, are excluded, as they do not belong to the Common Fiscal Regime (Régimen Fiscal Común), because they manage their income taxes directly.

[3] The filing status is chosen by the taxpayer. Joint tax returns typically benefit couples in which one partner earns little or no income, as well as single-parent families with dependent children.

we select a sample of retired individuals who are observed at least during that period.

2.1 Sample around maximum contribution limit.

Our final sample for this experiment includes only individuals continuously employed (retired in 1993 or afterwards) and excludes self-employed workers. The resulting sample contains 7,227 individuals. Table 1 presents the distribution of individuals according to the number of years observed before and after retirement. On average the individuals are observed during 9.6 years before and 3.6 years after retirement.

Table 1: Number of years before and after retirement.

Years before retir.	N^o indiv.	%	Years after retir.	N^o indiv.	%
8	2,172	30.05	1	1,424	19.7
9	1,908	26.40	2	1,180	16.33
10	1,129	15.62	3	929	12.85
11	797	11.03	4	1,303	18.03
12	815	11.28	5	1,805	24.98
13	406	5.62	6	586	8.11
Total n^o indiv.	7,227	100		7,227	100
			Total n^o obs.	24,324	

The running variable is defined as the log of the ratio between the mean value of the (real) labor income eight years before retirement ($\sum_{t=-8}^{0} \frac{Y_t^L}{\pi_t}$) and the maximum contribution limit (public pension) the year of retirement ($P\max_0$):

$$RV = \log(\frac{1}{8}\sum_{t=-8}^{0} \frac{Y_t^L}{\pi_t}) - \log \frac{P\max_0}{\pi_0} \qquad (1)$$

Figure 1 shows the density of the running variable. Out of our sample of 7,227 individuals, 1,919 have average income eight periods before retirement above the prevailing maximum contribution limit the year they retire ($RV > 0$).

Figure 1

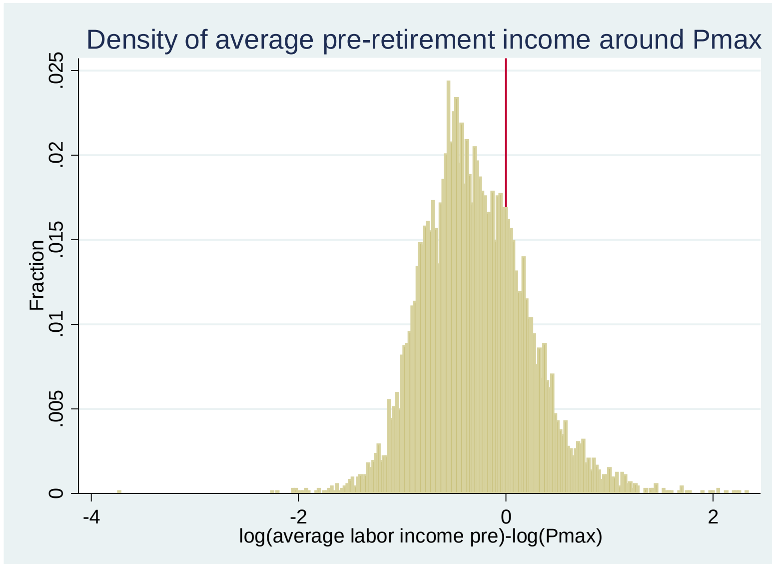

Figure 1: Density of the RV

We calculate for each post-retirement year, s, the rate at which public pensions replace labor income. The (log of the) replacement rate is defined as

$$\log \theta_s = \log(\frac{Y_s^L}{\pi_s}) - \log(\frac{1}{8}\sum_{t=-8}^{0}\frac{Y_t^L}{\pi_t}),\ s > 0. \tag{2}$$

We pool the observations for all post-retirement years. Therefore, our analysis is performed with a sample of 24,324 post-retirement observations corresponding to the 7,227 individuals observed at least 8 years before retirement.

Table 2 reports basic summary statistics for the sample. Mean pre-retirement labor income is about 15,000 euros, while mean post-retirement labor income is smaller (around 13,000 euros) and more disperse.

3 Institutional framework

jjjjjjjjjj

Table 2: Descriptive Statistics

	Mean	Std. Dev
Pre-retirement Labor income (euros)	15,449.26	10881.86
Post-retirement Labor income (euros)	13,750.09	11960.69
Post-retirement Total income (euros)	16,472.10	14278.12
N^{o} individuals	7,227	

4 Empirical Strategy

To identify the causal effect of tax incentives on the intertemporal smoothness of income, we apply a Regression Kink Design (RKD) strategy. As has been noted by Card et al. (2012), a kink assignment rule allows to identify, under certain conditions, the effect of interest that would be identified in a hypothetical randomized experiment. This is achieved by the comparison between the kink in the relationship between the running variable and the outcome variable with the kink in the "policy" rule.

4.1 Regression Kink Design

We exploit the kinked relationship between pre-retirement labor income (our running variable, $Y^L_{PRE} = \log(\frac{1}{8}\sum_{t=-8}^{0}\frac{Y^L_t}{\pi_t})$ and post-retirement labor income each year ($Y^L_{POST,s} = \log(\frac{Y^L_s}{\pi_s})$) and compare it with the kink in the post-retirement total income (our "policy" variable, $Y^T_{POST,s} = \log(\frac{Y^T_s}{\pi_s})$). In particular, we exploit two kinks that arise due to the existence of a minimum (P_{MIN}) and a maximum (P_{MAX}) in the contribution to retirement public pensions. At both sides of the threshold, individuals face different incentives due to changes in the public pension replacement rates and in the incentives to save.

The relationship between the variables in the RKD can be described by the following model:

$$Y^T_{POST} = \gamma Y^L_{POST} + G(Y^L_{PRE}) + \varepsilon, \qquad (3)$$

where Y^L_{PRE} is the observed "running variable" that affects Y^T_{POST} through a smooth function $G(Y^L_{PRE})$, and $Y^L_{POST} = F(Y^L_{PRE})$ is assumed to be a continous (and deterministic) function of Y^L_{PRE} with a policy-induced kink at $Y^L_{PRE} = P_{MAX}$. The identifying assumption is that, given the smootheness of $G(Y^L_{PRE})$ and the kink in Y^L_{POST}, there should be also a kink in the relationship between Y^L_{PRE} and Y^T_{POST} at the point $Y^L_{PRE} = P_{MAX}$.

Under the assumption that $G(Y^L_{PRE})$ and $E(\varepsilon \mid Y^L_{PRE} = P_{MAX})$ have derivatives that are continuous at $Y^L_{PRE} = P_{MAX}$, we have that the coefficient of interest, γ, can be calculated as the ratio between the two slopes at both sides of the threshold:

$$\gamma = \frac{\lim_{Y^L_{PRE} \to P^+_{MAX}} \frac{\partial E(Y^T_{POST} | Y^L_{PRE})}{\partial Y^L_{PRE}} - \lim_{Y^L_{PRE} \to P^-_{MAX}} \frac{\partial E(Y^T_{POST} | Y^L_{PRE})}{\partial Y^L_{PRE}}}{\lim_{Y^L_{PRE} \to P^+_{MAX}} \frac{\partial E(Y^L_{POST} | Y^L_{PRE})}{\partial Y^L_{PRE}} - \lim_{Y^L_{PRE} \to P^-_{MAX}} \frac{\partial E(Y^L_{POST} | Y^L_{PRE})}{\partial Y^L_{PRE}}}. \quad (4)$$

The numerator in (4) is the change in the replacement rate of the total post-retirement income at the kink, while the denominator is the change in the replacement rate of the post-retirement labor income at the kink. If the relationship between Y^L_{POST} and Y^L_{PRE} although deterministic, depends on other (unobserved or unknown) variables in addition to the primary running variable, or measurement errors in Y^L_{PRE} or Y^L_{POST}, a fuzzy RKD design can be used (see Card et al., 2015). In that case, the denominator needs also to be estimated.

Our empirical strategy proceeds in three steps:

(i) Estimation of the kink in the public pension replacement rate. This would measure the incentive to save because of the existence of a maximum (and a minimum) in the public pension. Notice that this is a "deterministic" relationship, in the sense that there is no a behavioral component involved. In addition to the change in the replacement rate, there is also a change in the incentives to save. The reason is that each euro contributed to pension funds is tax-exempted, and it accumulates at a tax-free rate and is taxed at withdrawal. This generates an increase in the returns to saving. Thus, we also estimate the kink in the incentives to save at this stage.

(ii) Estimation of the kink in the contributions to pension funds (and in the probability to contribute) to analyze to what extent workers use the tax incentives in the presence of a fall in replacement rates.

(iii) Estimation of the kink in the total post-retirement income (which includes the public pension and other financial income). This allows us to estimate how much of the fall in income due to the maximum (and the minimum) public pension is compensated by pension funds.

We estimate previous steps both for cohorts of workers retired before 1988, that is, workers for which pension funds were not available, and for cohorts retired afterwards.

A first concern with the RKD is the manipulation of the running variable. This could lead to bunch at the kink and to a non-smooth distribution of unobserved heterogeneity around it. Nonetheless, as in a regression discontinuity design (RD) (see for instance Lee and Lemieux, 2010), the validity of the RKD can be tested by showing that the density of the running variable is smooth at the kink point (this is analogous to the test of manipulation of the assignment variable for RD designs, discussed in McCrary, 2008).

The second concern with the RKD is the lack of data on characteristics that are determined prior to Y^L_{PRE} or Y^T_{PRE}, which conditional distribution is smooth and do not exhibit a kink at $Y^L_{PRE} = P_{MAX}$. In particular, for a valid RKD we need that the derivatives of the conditional expectation functions of those covariates with respect to Y^L_{PRE} are continuous at the kink point. This implication can also be tested and is analogous to the test emphasized by Lee (2008) for a regression discontinuity design. We will deal with both concerns in next section.

5 Graphical evidence

5.1 Changes at the contribution limit

The existence of a maximum pension generates a kink in the relationship between labor income prior to retirement and the post-retirement replacement rate. To the extent that retirement is predictable, this implies not only a predictable income fall, but a predictable change in the rate at which public pensions replace labor income.

As a starting point for our RKD analysis, Figure 2, shows the relationship between preretirement labor income and the public pension replacement rate around the top kink. The horizontal axis represents the deviation of the (log of) average income in the eight years prior to retirement from contribution cap in the year of retirement. The vertical axis shows the (log of) public pension replacement rate: the median difference between yearly log post-retirement public pension income and the mean log pre-retirement income 8 years prior to retirement. We plot average preretirement labor income, in years, for bins normalized at the kink. Note that in these (and subsequent) figures we have narrowed our focus to a range of $(-0.5, 1)$ around the kink (using bins of 0.1 log-points).

As expected, Figure 2 shows a kink in the empirical relationship between postretirement labor income and pre-retirement labor income, with a decrease in slope as preretirement income passes through the threshold (the maximum pension). We include the fit of a linear regression between both variables. The public replacement rate for individuals whose average income eight years prior to retirement is below the cap is 76%. The ratio between yearly (log) retirement income and (log) income prior to retirement is $-.3$, and it does not vary with income. For individuals with pre-retirement earnings above the cap, the public replacement rate falls monotonically with pre-retirement earnings (the differential slope is $-.71$, while the actual slope is $-.57$).

The change in the slope at 0 represents a change in the incentive to save.

Figure 2: Change in the replacement rate

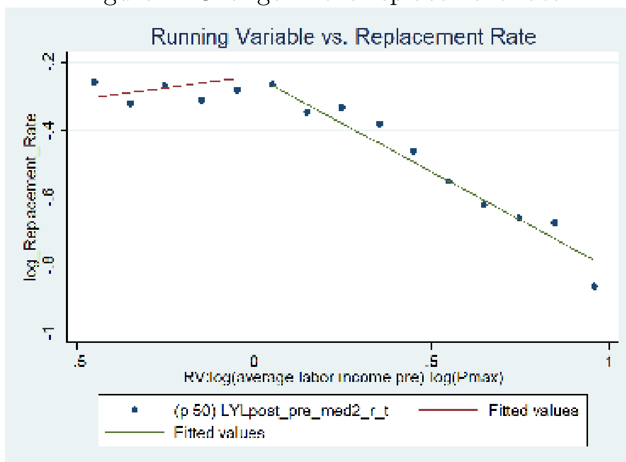

Consider two individuals with pre-retirement income below the maximum pension (0 in the horizontal axis). Both will experience a similar 24 pp. income drop at retirement. The public pension replacement rate changes with income at a different rate when pre-retirement incomes exceed the cap: an individual right at the maximum public pension expects a replacement rate of about 76%, while an individual with 50% higher income has a replacement rate of 64%.

Figure 2

Figure 3 presents a parallel figure for another change at the contribution limit: the post-tax return to saving increases (price effect). An euro contributed to pension funds is tax-exempted. It accumulates at a tax-free rate and is taxed at withdrawal. This generates an increase in the returns to saving. If contributions are tax exempted at a marginal tax $MTAXpre$ and tax liable at $MTAXpost$. Figure 3 presents the returns to saving by plotting the effective post-tax return on saving $R \times (1 - MTAXpost)/(1 - MTAXpre)$ against the running variable (the deviation of the (log of the) average income in the eight years prior to retirement from contribution cap in the year of retirement). It shows that there is a well-defined kink in the relationship between preretirement labor income and the returns to saving.

Figure 3: Change in the returns to saving

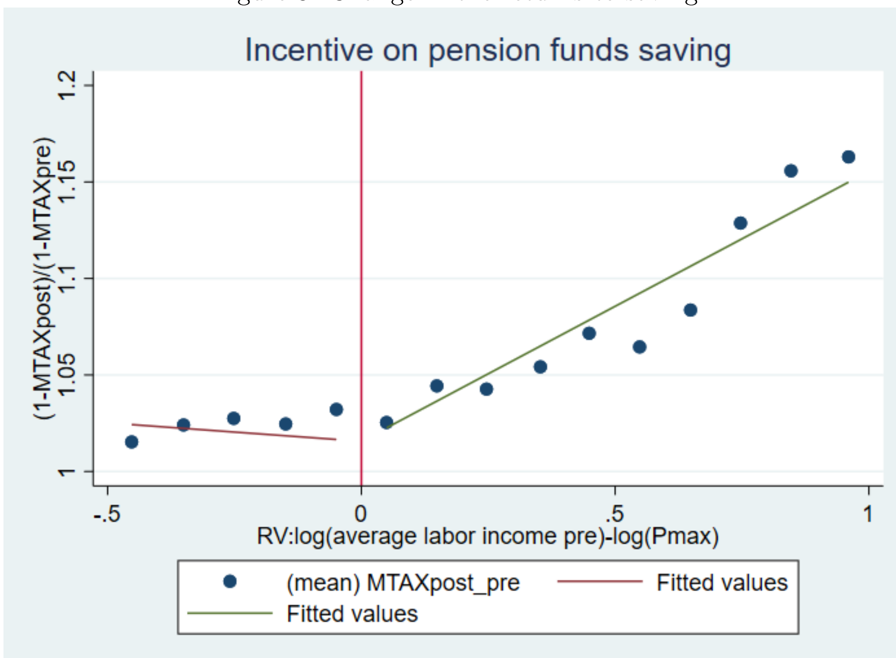

Figure 3

5.2 Changes in saving behavior

Previous graphs show that there are two incentives to increase saving behavior around the social security cap: the decrease in the replacement rate and the increase in the returns to saving. Next issue we analyze graphically is whether or not taxpayers change their saving behavior around the cap.

We analyze two different outcomes: the probability of contributing (at least one of the two years before retirement) and the amount contributed (two years before retirement). For the latter we compute the yearly contributions relative to income and take averages.

Figure 4 shows that the semi-elasticity of contributing to pension funds out of income increases by .23 pp. at the cap. That is, above the cap, marginal propensity to contribute doubles relative to below the cap (.40 vs .17). As to the fraction of income contributed (Figure 5) we obtain that above (below) the cap,

Figure 4: Change in the porpensity to contribute to PF

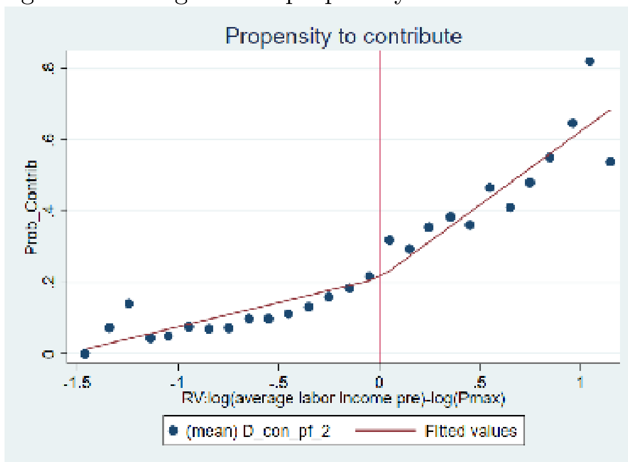

Figure 5: Change in the mean contribution to PF

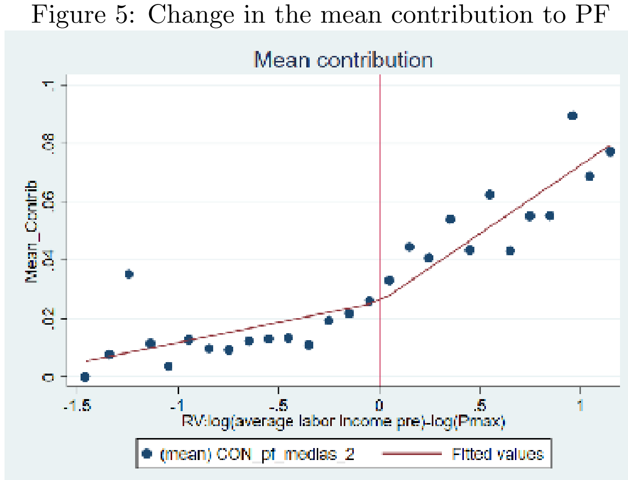

an increase in 1% in income increases the fraction of income saved in pension funds by 1.1% (.47%).

Figure 4

5.3 Intention to treat: The response of post-retirement income to tax incentives

Given that our purpose is to analyze the success of tax incentives to retirement saving to change the slope of income along the life-cycle (do individuals tilt their income profiles towards more income at retirement?) we compute all income

Figure 6: Change in post-retirement total income

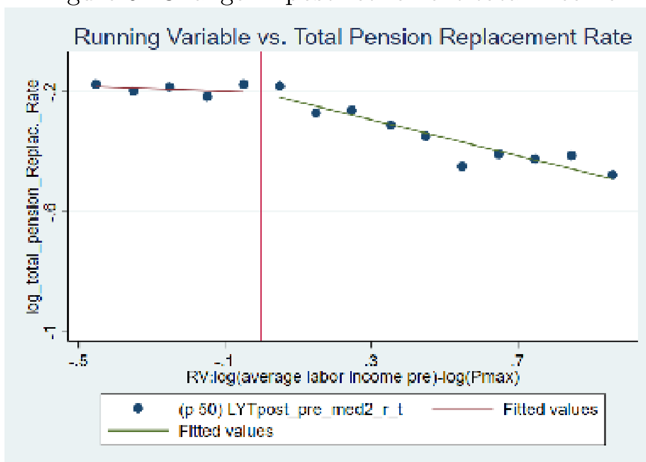

sources at retirement as a measure of whether or not they succeed. Thus, we calculate the total income replacement rate by including public pensions, but also all other income from private investments (financial income, rents, self-employment income).

Figure 6 shows in the horizontal axis the deviation of the (log of) average labor income in the eight years prior to retirement from contribution cap. The vertical axis shows the (log of) total pension replacement rate: the difference between yearly log post-retirement total pension income and the mean log pre-retirement income. According to the results, we have evidence to reject self-insurance: even after the introduction of tax incentives, private income does not compensate for the fall in public pension (the slope above the cap is negative)

To study to what extent the gap is closed, we compare the elasticity of public income (only public pension) to pre-retirement income (Figure 2), which is $-.71$ to the elasticity of total income to pre-retirement income (Figure 6) which is $-.25$. Therefore, pension funds absorb 65% ($1 - \frac{0.25}{0.71}$) of the fall in the replacement rate.

Figure 6

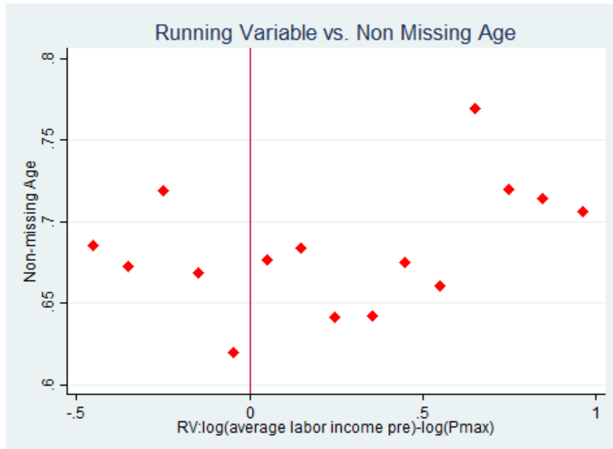

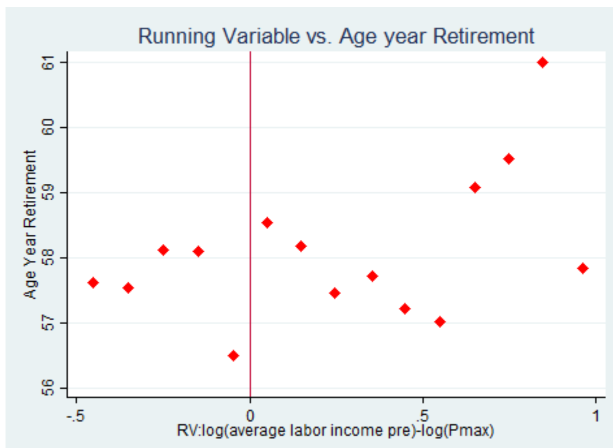

5.4 Predetermined covariates around the kink

We analyze the pattern of some predetermined covariates around the kink. We plot the relationship between our running variable and a set of predetermined variables: non-missing age, age at retirement, joint versus individual tax return the year of retirement, standard deviation of pre-retirement labor income, and first labor income observed. The graphs in Figure ? show average values of each characteristic for each bin of the running variable. Most of the graphs show lack of or little indication of a kink in the conditional mean function.

To increase the power of this analysis, we construct a "covariate index". In particular, we obtain the predicted outcome (the log of total pension replace-

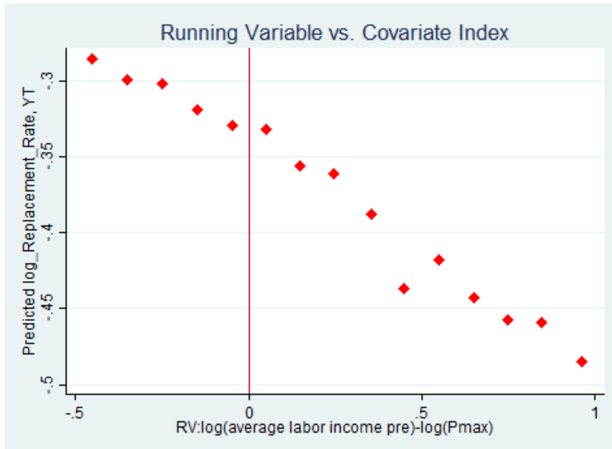

ment rate) using previous covariates as regressors from a simple linear regression model. As pointed out by Card et. al (2012) this index function can be interpreted as "best linear prediction" of mean log total pension replacement rate, given the vector of predetermined variables. Figure ? plots the conditional mean values of the estimated covariate index around the kink. The graph shows that the mean predicted outcome evolve smoothly through the kink.

Another important identifying assumption is the smoothness of the density of the the running variable. This assumption is less likely to hold if a discontinuity can be observed in its density around the kink. In Figure **??** we plot the probability density function of this variable to show the smoothness of its distribution at the kink point. This smoothness is evidence against the possibility of manipulation of earnings at the kink point.

Figure

5.5 A benchmark: cohorts retiring prior to 1988

As a benchmark, we perform a similar analysis but with the individuals retired before 1988. Prior to 1998 pension funds were not available, so taxpayers had to rely on other saving products. To the date we only have data from 1985, so we can only use three years of income before retirement to compute our

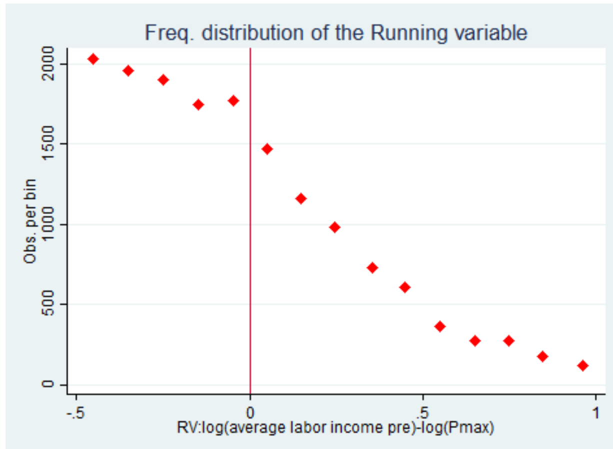

running variable, and not eight. But otherwise we present similar figures as in previous section, that is: (i) public pension replacement rates out of three pre-retirement years (i.e., public pension after retirement as a function of three last years of income pre-retirement); (ii) total income replacement rates out of 3 pre-retirement years (i.e., total income after retirement as a function of three last years of income pre-retirement). If taxpayers smooth income intertemporally, the slope of (ii) should be smaller in absolute value than that of (i).

Figures 7 and 8 present our graphical results. Although not using the full eight years to compute the running variable results in some noise, Figure 7 shows an elasticity of public income (only public pension) to pre-retirement income of $-.376$, while according to Figure 8 the elasticity of total income (public pension plus all sources of income) to pre-retirement income is $-.343$. Therefore, it seems that saving vehicles (other than pension funds) absorb only 8% $(1 - \frac{0.343}{0.376})$ of the fall in the replacement rate.

Figure 7

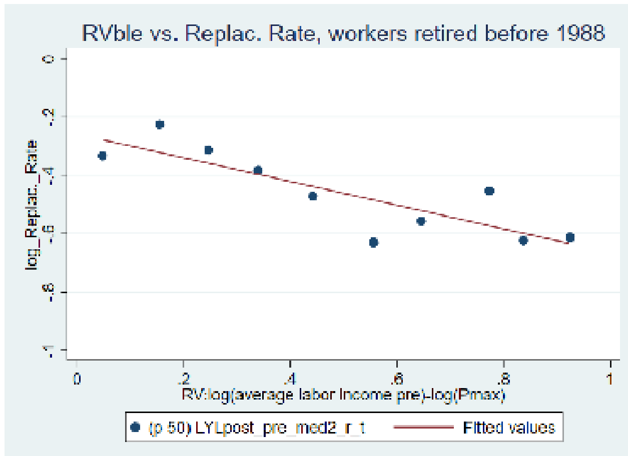

Figure 7: Change in the replacemente rate prior to 1988

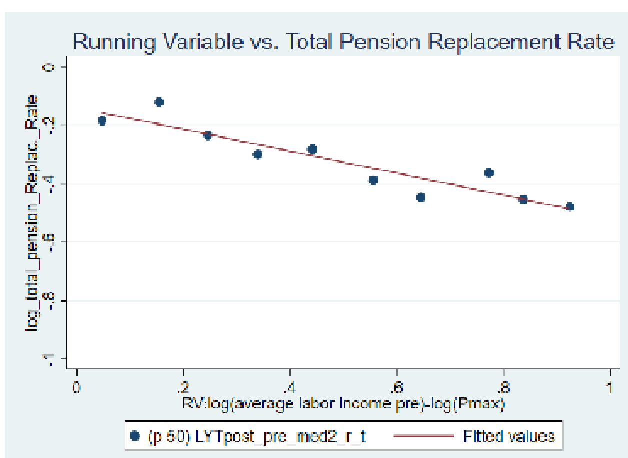

Figure 8: Change in post-retirement total income prior to 1988