Tax incentives to retirement saving and intertemporal income smoothing: Evidence from a kink design*

Raquel Carrasco† Ernesto Villanueva‡

February 28, 2024

Abstract

Tax incentives for retirement saving typically imply that monetary amounts invested in retirement saving products are tax-exempt, and 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 before retirement. To obtain exogenous variation in the incentive to use fiscal-favored products, we exploit the fact that maximum and minimum Social Security contribution limits introduce 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 60% of the fall in public pension retirement rates. Namely, while the fraction of tax payers with public replacement rates above 70% fall by .5%by each 1% increase in income above the maximum contribution, the fraction of tax payers whose total replacement rate is above 70% only falls by 0.20%. When tax incentives were not available (before 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 against a predictable fall in public pension income.

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

Keywords: Tax incentives, Retirement Saving, Regression Kink Design.

*Financial support from MINECO/FEDER, grants RTI2018-095231-B-I00, PID2022-138960OB-I00, MDM2014-0431, and Comunidad de Madrid, Grant No. MadEco-CM (S2015/HUM-3444), is gratefully acknowledged. The usual disclaimer applies.

†Department of Economics, Universidad Carlos III de Madrid. E-mail: rcarras@eco.uc3m.es

‡Banco de Espaa. E-mail: ernesto.villanueva@bde.es

1 Introduction

Tax incentives for 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 their life-cycle savings. 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 succeed in tilting the life-cycle profile of income around 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). Namely, 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. 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 progressivity of the income tax schedule, 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 happens 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 old age(as income 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, a comparison between the first and third steps 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 0.7% by each 1% increase in income, total replacement rates only fall by 0.2% When we compare the public pension replacement rate(.7) with the total income replacement rate (.25) we find that twothirds of the fall in public pension replacement rates are 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 before 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 several 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, Gale and Scholz 1996), either because savers are very sophisticated and reshuffle their portfolios (Engen, Gale and Scholz 1996), 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, 2019). 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 examining whether 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 crowd-out 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 Conceptual framework

Assume that a consumer lives for two periods (pre- and post- retirement). She faces no uncertainty and cannot choose retirement, that happens at the end of the first period. The public pension amounts to a fraction θ of her first period income, where $\theta < 1$ and we assume this to be constant for the time being. The interest rate is 0. Her utility function is $\frac{c_1^{1-\rho}}{1-\rho} + \beta\frac{c_2^{1-\rho}}{1-\rho}$ and her budget constraint can be expressed as $c_1 + c_2 \leq y_1(1+\theta)$.

And solving the budget constraint, consumption in the second period as a fraction of first-period income can be expressed as follows:

$$\frac{c_2}{y_1} = (1+\theta)\frac{\beta^{\frac{1}{\rho}}}{1+\beta^{\frac{1}{\rho}}}$$

Now, note that in Spain, as in other counties, there is a maximum pension. Hence, the pension an individual gets upon retirement is $\max\{y^{\max}, \theta y_1\}$. This feature implies that whenever $\theta y_1 > y^{\max}$ the effective replacement rate is $\frac{y^{\max}}{y_1}$, a decreasing function of y_1 or $\frac{d\theta}{dy} < 0$ In that circumstance, the ratio of second period consumption to first period income $\frac{c_2}{y_1}$ as income increases is flat for observations whenever $y_1 < \frac{y^{\max}}{\theta}$ and

$$\frac{\partial\frac{c_2}{y_1}}{\partial y_1} = 0 \qquad if \qquad y_1 < \frac{y^{\max}}{\theta}$$

while

$$\frac{\partial\frac{c_2}{y_1}}{\partial y_1} = \frac{\beta^{\frac{1}{\rho}}}{1+\beta^{\frac{1}{\rho}}}\frac{d\theta}{dy} < 0 \qquad if \qquad y_1 > \frac{y^{\max}}{\theta}$$

That is, the ratio of second period consumption to first period income falls with income, but at a rate that it is lower than $\frac{d\theta}{dy}$.

Under the introduction of tax incentives (the post period) contribution to pension funds are tax exempt. Each euro contributed to a pension fund effectively diminishes the tax liability in the first period by the marginal tax in that period τ_1. The euro then accumulates at the market rate 0 and is taxed upon withdrawal at retirement at a tax τ_2. To the extent that, for individuals in our sample, $\tau_2 > \tau_1$, the post-tax interest rate on saving becomes positive $r = \frac{1-\tau_2}{1-\tau_1}$. The higher interest rate induces a substitution effect and an incentive to tilt the consumption profile toward the old age. In other words, consumption in the second period represents a higher share of first income period income in a world where tax incentives are present. That

$$\frac{c_2^{post}}{y_1^{post}} = \frac{(1+\theta)}{1+[\beta R]^{-\frac{1}{\rho}}} > \frac{c_2}{y_1} \tag{1}$$

Secondly, as we document below, due to tax progressivity, $r = \frac{1-\tau_2}{1-\tau_1}$ increases with income. The reason is that, once $y_1 > \frac{y^{\max}}{\theta}$ pension income is at its maximum $\theta y^{\max}$ so the numerator is capped. Furthermore, τ_1 increases with income in the first period, due to tax progressivity.

Then, the combination of a capped second-period income and an increasing marginal income tax rate in the first period results in $r = \frac{1-\tau_2}{1-\tau_1}$ being an increasing function of y_1 at the right of the kink.

Taking derivatives in (1) we find

$$\frac{\partial \frac{c_2^{post}}{y_1^{post}}}{\partial y_1} = 0 \qquad if \qquad y_1 < \frac{y^{\max}}{\theta}$$

$$\frac{\partial \frac{c_2^{post}}{y_1^{post}}}{\partial y_1} = \frac{1}{1+[\beta R]^{-\frac{1}{\rho}}}\frac{d\theta}{dy} + \frac{1}{\rho}(1+\theta)\frac{[\beta R]^{-\frac{1}{\rho}-1}}{[1+[\beta R]^{-\frac{1}{\rho}}]^2}\frac{dR}{dy_1} \qquad if \qquad y_1 > \frac{y^{\max}}{\theta} \tag{2}$$

The first term has been already discussed and reflects that if public pension replacement rates θ fall with income, consumption as a fraction of first-period income falls through an income effect.. The second term is new, and reflects that because of the presence of tax incentives, the net interest rate of contributors increases leading to more saving (and in turn mitigating the fall in $\frac{c_2^{post}}{y_1^{post}}$).

As a result, the discussion delivers two testable hypotheses regarding the saving behavior at the kink of the public pension replacement rate. As described below, we do not observe consumption c_2 but second-period income y_2 (in principle, they should be equivalent in our setting). We express the hypotheses in terms of y_2

- If individuals are forward-looking, the ratio between contributions to pension funds and first period income should change its relationship with income at the kink $y_1 = \frac{y^{\max}}{\theta}$. This is the consequence of two factors: when income is strictly below the kink, small differences in first-period income y_1 should maintain θ and R constant. Under our assumptions, this implies similar levels of the ratio of contributions to pension funds and first-period income. However, to the right of the kink, a slight increase in income diminishes *both* the public pension replacement rate θ (making saving more attractive to make up for the fall in public pension replacement rate) and the net interest rate increases $r = \frac{1-\tau_2}{1-\tau_1}$ (inducing saving through a substitution effect). Then the derivative (not the level) of contributions to pension funds with respect to first-period income should change at the kink.

- If individuals are forward looking, and for the same reasons, the ratio between second-period income and first period income should change its relationship with income at the kink $y_1 = \frac{y^{\max}}{\theta}$. Furthermore, a comparison between $\frac{\partial \frac{c_2^{post}}{y_1^{post}}}{\partial y_1}$ and $\frac{d\theta}{dy}$ (magnitudes both observable in the data) provide an estimate of the intertemporal smoothing role of pension funds. That role emanates in this model from the presence of a substitution effect - in turn, the second term in (2).

3 Empirical Strategy

To identify the causal effect of tax incentives on the intertemporal smoothing of income, we apply a Regression Kink Design (RKD) strategy. As noted by Card et al. (2016), a kink assignment rule enables us to identify, under certain conditions, the effect of interest that would be identified in a hypothetical randomized experiment. This is accomplished by comparing the relationship between the running variable and the outcome variable around the kink, with the kink in the "policy" rule.

3.1 Regression Kink Design

We exploit the kinked relationship between pre-retirement labor income, Y^L_{PRE}, and post-retirement labor income each year after retirement, Y^L_{POST}, and compare it with the kink in post-retirement total income (our "policy" variable, Y^T_{POST}). Specifically, we exploit two kinks that arise from the presence of a maximum (P_{MAX}) and a minimum (P_{MIN}) in the contribution to retirement public pensions. On either side of these thresholds, individuals face distinct incentives due to changes in the public pension replacement rates and the incentives to save.

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

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

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

In what follows, we present the analysis for the upper kink, at P_{MAX}, but the same applies to the lower kink, at P_{MIN}. 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 can calculate the coefficient of interest,

γ, as the ratio between the two slopes at both sides of the threshold:

$$\gamma = \frac{\lim_{Y_{PRE}^L \to P_{MAX}^+} \frac{\partial E(Y_{POST}^T | Y_{PRE}^L)}{\partial Y_{PRE}^L} - \lim_{Y_{PRE}^L \to P_{MAX}^-} \frac{\partial E(Y_{POST}^T | Y_{PRE}^L)}{\partial Y_{PRE}^L}}{\lim_{Y_{PRE}^L \to P_{MAX}^+} \frac{\partial E(Y_{POST}^L | Y_{PRE}^L)}{\partial Y_{PRE}^L} - \lim_{Y_{PRE}^L \to P_{MAX}^-} \frac{\partial E(Y_{POST}^L | Y_{PRE}^L)}{\partial Y_{PRE}^L}} \tag{4}$$

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

Our empirical strategy proceeds in three steps:

(i) Estimating the kink in the public pension replacement rate $\frac{d\theta}{dy}$. This step measures the incentive to save due to the presence of a maximum (and a minimum) in the public pension. It should be noted that this relationship is "deterministic" and does not involve a behavioral component. In addition to the change in the replacement rate, there is also a change in the incentives to save. This is because each euro contributed to pension funds is tax-exempted, accumulates at a tax-free rate, and is taxed upon withdrawal. Consequently, there is an increase in the returns to saving. Therefore, we also estimate the kink in the incentives to save at this stage.

(ii) Estimating the kink in contributions to pension funds (and in the probability of contributing). In this step, we analyze to what extent workers use the tax incentives in the presence of a decline in replacement rates.

(iii) Estimating the kink in total post-retirement income (including the public pension and other financial income) $\frac{\partial \frac{c_2^{post}}{y_1^{post}}}{\partial y_1}$. This step allows us to estimate how much of the income reduction resulting from the maximum (and the minimum) public pension is compensated by pension funds.

In the present version, we estimate the previous steps for cohorts of workers: those who retired after 1988, when pension funds were available.

Our general empirical specification is the following:

$$E(Y^j_{POST}|Y^L_{PRE} = P_{MAX}) = \alpha^j_0 + \alpha^j_1 D + \sum_{p=1}^{P} \gamma^j_p (Y^L_{PRE} - P_{MAX})^p + \sum_{p=1}^{P} \delta^j_p D \times (Y^L_{PRE} - P_{MAX})^p \quad (5)$$

where Y^j_{POST} is the outcome of interest (income observed POST retirement) , $j = L, T$, Y^L_{PRE} is the running variable (income observed pre-retirement), P is the order of the polynomial, and $D = \mathbf{1}(Y^L_{PRE} > P_{MAX})$. This specification is estimated for observations with $\left|Y^L_{PRE} - P_{MAX}\right| \leq h$, where h is the bandwidth size.

Namely, when the dependent variable Y^j_{POST} is public pensions the expression estimates $\frac{d\theta}{dy}$ in the previous section. In particular, that specification delivers what we term γ^L_1 (the non-differentiability in the public replacement rate, or the incentive to save due to an income effect).

When the dependent variable Y^j_{POST} is total post-retirement income the expression estimates $\frac{\partial \frac{c_2^{post}}{y_1^{post}}}{\partial y_1}$ in the previous section. In particular, that specification delivers what we term γ^T_1 (the non-differentiability in the total income, or the result of public pensions plus the proceeds of any saving in period 1). Under our assumptions, the ratio $\frac{\gamma^T_1}{\gamma^L_1}$, informs about the income smoothing role of pension funds.

4 Data

We use an administrative dataset, the Spanish Panel of Tax Returns, which comprises a stratified 5% random sample of tax returns for the period 1982-1998. This dataset enables us to conduct an analysis both before and after the implementation of pension plans in Spain. It is a representative sample that accurately reflects the population of taxpayers.[1] The income tax samples are drawn from 15 of the 17 autonomous communities of Spain, including 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, while joint tax returns represent two spouses filing together or

[1] Income tax samples do not include individuals with no taxable income. Consequently, our analysis excludes individuals with either no income or very low income.

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

single-parent families with at least one child.[3]

This dataset includes almost the entire range of fiscal and socio-demographic information provided by taxpayers in their returns. Specifically, it provides detailed information on income from various sources (such as labor, financial, and self-employment), as well as yearly contributions to pension plans. Additionally, it includes information about certain household characteristics (such as the number of dependent relatives and disabilities). However, demographic information is limited (for example, age information was only available for 60% of the sample, so it was not used).

To determine whether or not an individual is retired we use information regarding their social security contributions, which should be zero for retirees. Specifically, we classify individuals as retired if their social security contributions are less than 6 euros per year. As the public pension income is based on an individual's labor income eight years before retirement, we selected a sample of retired individuals who were observed during at least that period.

We study the effects of the kinks resulting from both the maximum and minimum pension contributions by establishing two distinct subsamples. The first subsample focuses on contributions near the maximum limit, while the second subsample concentrates on contributions near the minimum limit.

4.1 Sample around maximum contribution limit

Our sample for this experiment includes only individuals who were continuously employed for at least 8 years before retirement (i.e., individuals who retired in 1993 or later) and excludes self-employed workers. The resulting sample contains 7,227 individuals. Table 1 presents the distribution of individuals based on the number of years observed before and after retirement. On average, individuals are observed for 9.6 years before retirement and 3.6 years after retirement.

[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.

Table 1: Number of years before and after retirement

Years before retir.	Nindiv.	%	Years after retir.	Nindiv.	%
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 nindiv.	7,227	100		7,227	100
			Total nobs.	24,324	

We define our running variable (RV) as the (log of the) ratio between the mean value of the real labor income measured in constant 2007 euros, Y_t^L, eight years before retirement, and the maximum contribution limit for the public pension the year of retirement (P_{MAX_0}):

$$\log(RV) = \log(\frac{1}{8}\sum_{t=-8}^{-1} Y_t^L) - \log P_{MAX_0} \tag{6}$$

Figure 1 displays the density of this variable. Among our sample of 7,227 individuals, 26.5% have an average real labor income of eight periods before retirement that exceeds the maximum contribution limit in the year they retire ($RV > 0$).

We define the replacement rate of public pensions for each post-retirement year, denoted as "s", θ_s, as follows:

$$\log\theta_s = \log(Y_s^L) - \log(\frac{1}{8}\sum_{t=-8}^{-1} Y_t^L),\ s \geq 0. \tag{7}$$

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

Tables 2 and 3 present basic summary statistics for the sample. The average pre-retirement labor income is approximately 22,000 euros, whereas the average post-retirement labor income is lower (around 15,000 euros) and more dispersed. On average 18% of the individuals in our sample make contributions to pension funds one year before retirement, while this figure is around 24% if we consider an extended period of eight years before retirement.

Figure 1: Density of the Running Variable around Pmax

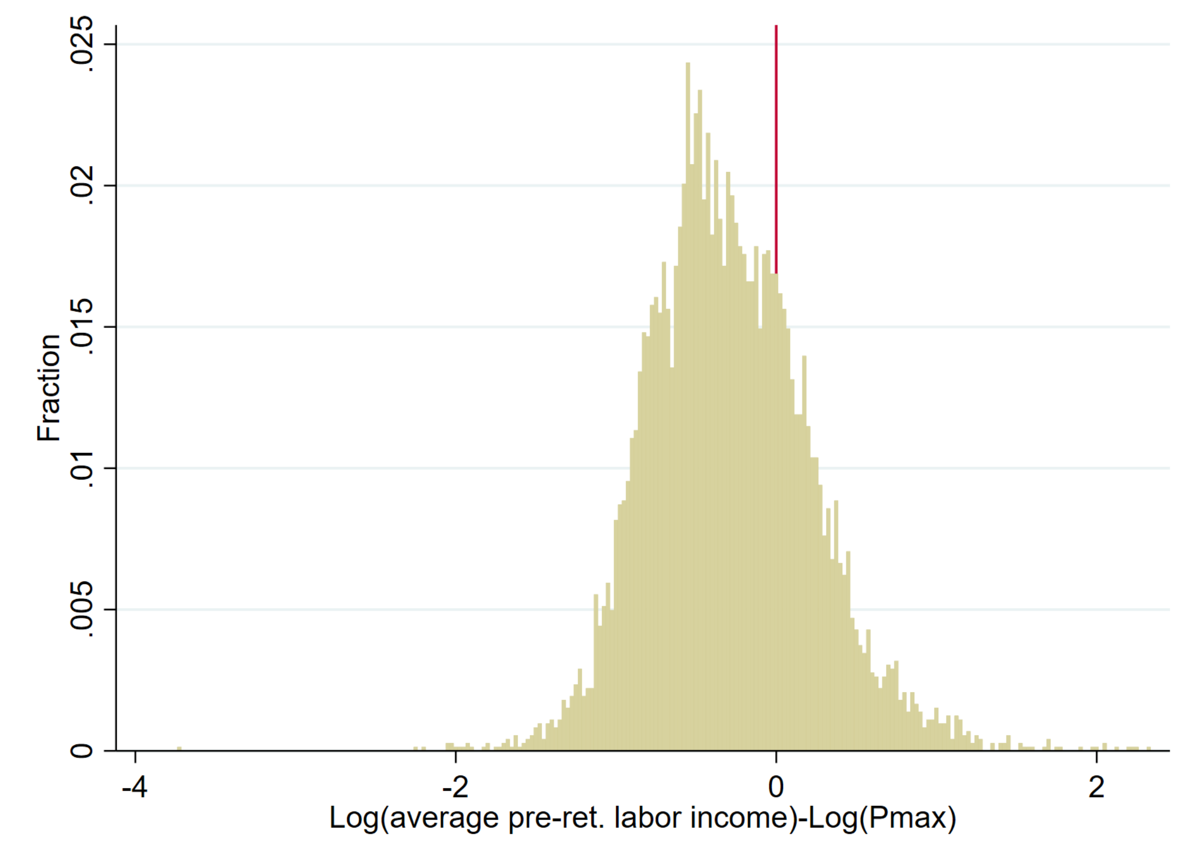

Table 2: Descriptive Statistics. Sample around Pmax

	Pre-ret.	Post-ret.	Post-ret.
	Labor income	Labor income	Total income
Mean	22,762	15,633	18,738
	(15,372)	(13,460)	(16,068)
By nof obs.			
post-ret.			
1	22,065	14,622	16,967
	(15,609)	(12,006)	(15,083)
2	22,852	15,165	18,160
	(16,736)	(12,799)	(15,554)
3	23,246	16,099	19,936
	(16,105)	(14,011)	(18,244)
4	23,422	16,168	19,866
	(17,431)	(18,581)	(22,014)
5	22,183	15,129	17,621
	(12,869)	(10,768)	(11,965)
6	23,828	16,485	20,088
	(12,594)	(10,850)	(13,148)

Notes: Standard deviations between brackets. Monetary magnitudes in 2001 euros. Sample of tax filers observed 8 years before retirement, for which their statutory public pension could be computed

Table 3: Descriptive Statistics. Pension Funds Contributions

	Prob. PF contrib.		Relative PF contrib.	
Nyears before ret.	1 year	8 years	1 year	8 years
Mean	0.18	0.24	0.035	0.062
	(0.39)	(0.42)	(0.424)	(0.181)
By nof obs. post-ret.				
1	0.20	0.29	0.034	0.086
	(0.40)	(0.45)	(0.245)	(0.212)
2	0.21	0.29	0.048	0.082
	(0.41)	(0.45)	(0.603)	(0.214)
3	0.19	0.25	0.026	0.067
	(0.39)	(0.44)	(0.226)	(0.179)
4	0.19	0.23	0.021	0.061
	(0.39)	(0.42)	(0.063)	(0.180)
5	0.14	0.17	0.028	0.036
	(0.35)	(0.38)	(0.498)	(0.128)
6	0.18	0.19	0.082	0.041
	(0.38)	(0.39)	(0.702)	(0.166)

Notes: Standard deviations between brackets. Sample of tax filers observed 8 years before retirement between 1988 and 1998, for which their statutory public pension could be computed. The first columns denote the fraction of filers who have contributed to pension funds The last two columns show the mean ratio of the amount contributed to the gross income of the tax unit.

4.2 Sample around minimum contribution limit

TO DO

5 Results

5.1 Assumptions of the RKD

One potential concern regarding the RKD design is the possibility of manipulating the running variable, which could lead to clustering around the kink and a non-smooth distribution of unobserved heterogeneity. However, the validity of the RKD can be assessed by examining the smoothness of the density of the running variable at the kink point. This test is similar to the examination of manipulation in assignment variables in Regression Discontinuity (RD) designs, as discussed by McCrary (2008).

In Figure 2, we present the probability density function of normalized pre-retirement income around both kinks. The plot displays a smooth density, suggesting that there is no manipulation of earnings at either kink point. The p-values obtained from the McCrary test for discontinuities are 0.20 and ? for the upper and lower kink, respectively. The null hypothesis of no discontinuity is not rejected at the conventional confidence levels.

Figure 2: McCrary Test

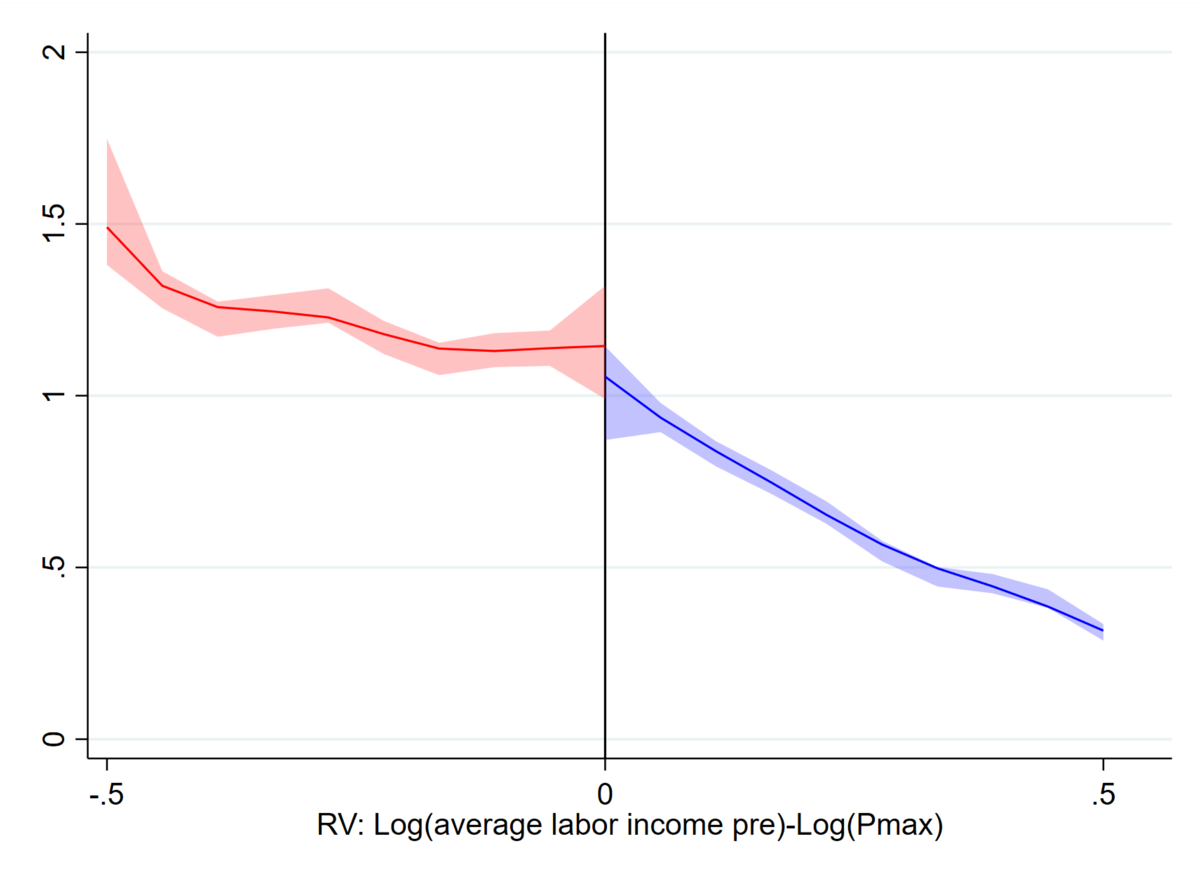

The second concern associated with the RKD relates to the lack of data on characteristics that are determined prior to Y^L_{PRE} or Y^T_{PRE}, and should exhibit a smooth conditional distribution without a kink at $Y^L_{PRE} = P_{MAX}$ or $Y^L_{PRE} = P_{MIN}$. In particular, for the RKD to be valid, it is necessary for the derivatives of the conditional expectation functions of those covariates with respect to Y^L_{PRE} to be continuous at the kink points. This implication can also be tested and is similar to the test emphasized by Lee (2008) in the context of regression discontinuity design. To perform this test, following the approach outlined by Card et al. (2014), we construct a "covariate index" by predicting the value using a simple linear regression of the log of the logarithm of the total replacement rate on a set of predetermined covariates, such as retirement age, filling status, standard deviation of pre-retirement labor income, and first observed labor income. Figure 3 displays the average values of the covariate index for different bins of the running variable, and they exhibit a relatively smooth evolution across both the upper kink.

Figure 3: Running Variable and Covariate Index

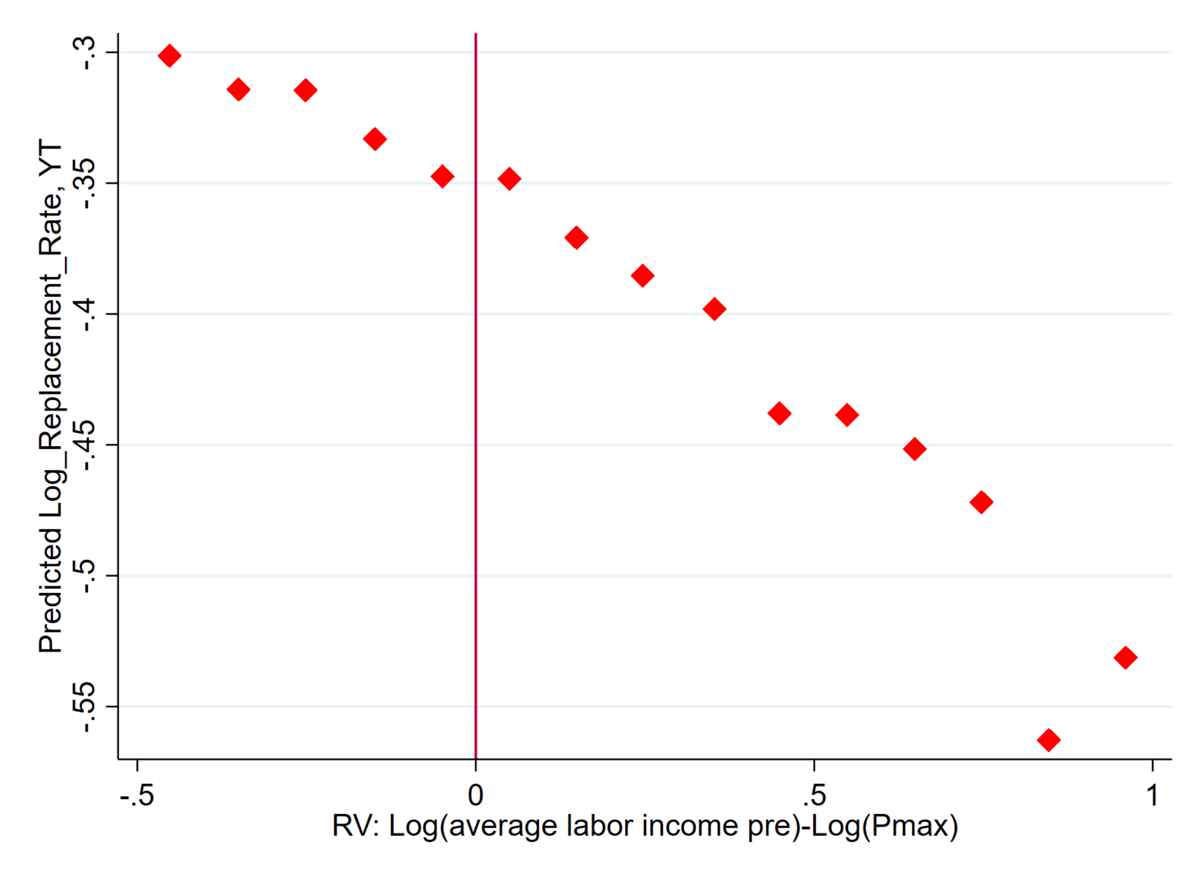

5.2 Graphical evidence and estimation results

In this section, we provide the estimation results from the various steps of our empirical strategy. Specifically, we present both visual evidence and estimation results for equation (5) concerning the different outcomes of interest. The implementation of the RKD relies on two considerations: the selection of the polynomial degree and the choice of bandwidth. We present results using both a linear and a quadratic polynomial, along with two bandwidth choices: $h = 1$ and $h = 0.5$.

5.2.1 Incentives to save

The existence of a maximum (and a minimum) pension generates a kink in the relationship between pre-retirement labor income and the post-retirement replacement rate. Consequently, it becomes possible to predict not only the decline in income during retirement but also the change in the rate at which public pensions replace labor income.

As an initial step in our RKD analysis, Figure 4 illustrates the relationship between pre-retirement labor income and the public pension replacement rate around the upper kink. On the horizontal axis, we represent the deviation of the (log of) average income in the eight years preceding retirement from the logarithm of the contribution cap in the retirement year. The vertical axis represents the (log of) the public pension replacement rate, which is defined as the ratio between the yearly post-retirement public pension income and the average pre-retirement income eight years before retirement (both logged). We plot the average pre-retirement labor income, normalized within bins centered around the kink. It is important to note that in this figure, as well as in subsequent figures, our focus is narrowed to a range of (-0.5, 0.5) and (-1, 1) around the kink, with bin widths of 0.071 and 0.5 log-points to the left and right of the kink, respectively. Additionally, we include a linear regression fit between the two variables.[4]

Although Figure 4 shows a kink in the empirical relationship between post-retirement and pre-retirement labor income, displaying a decreasing slope as pre-retirement income passes through the threshold (maximum pension), we would have anticipated a flatter relationship before the kink. Specifically, individuals with an average income eight years before retirement below the cap typically experience a public replacement rate of approximately 70%, which should remain constant regardless of the income level. However, our data does not align with this. A possible factor accounting for this disparity is that the distribution of income within the bins has fat tails, that result in a distorted estimate of the mean replacement rate.[5]

For that reason, we have defined an alternative outcome variable as a dummy taking the value 1 if the public replacement rate is greater than 0.7. Figure 5 presents the relationship between the average of this variable (that is, the probability that the public replacement rate is greater than 70%) and the running variable for the upper kink. As expected, for individuals whose average income eight years before retirement is below the maximum cap the relationship is flat

[4]The figure is constructed using the "rdplot" Stata command by restricting the support to the neighborhood around the cutoff defined by the choice of bandwidth (in this case, equal on both sides). We then adjust the (global) fit in the RD plot to match the local polynomial point estimation performed by "rdrobust" in that neighborhood. We choose a first-order polynomial.

[5]Namely, in the bins further away from the maximum, a few observations hit the minimum contribution limit. To preserve differentiability, our running variable uses actual rather than the legal base income. When one individual has actual income below the legal minimum contribution, the observed replacement rate is 70% of the average pre-retirement income where incomes below the minimum income are replaced by that minimum. Conversely, our running variable does not apply those limits, which may result in an overestimation of the public pension replacement rate.

Figure 4: Pre-retirement labor income and public pension replacement rate

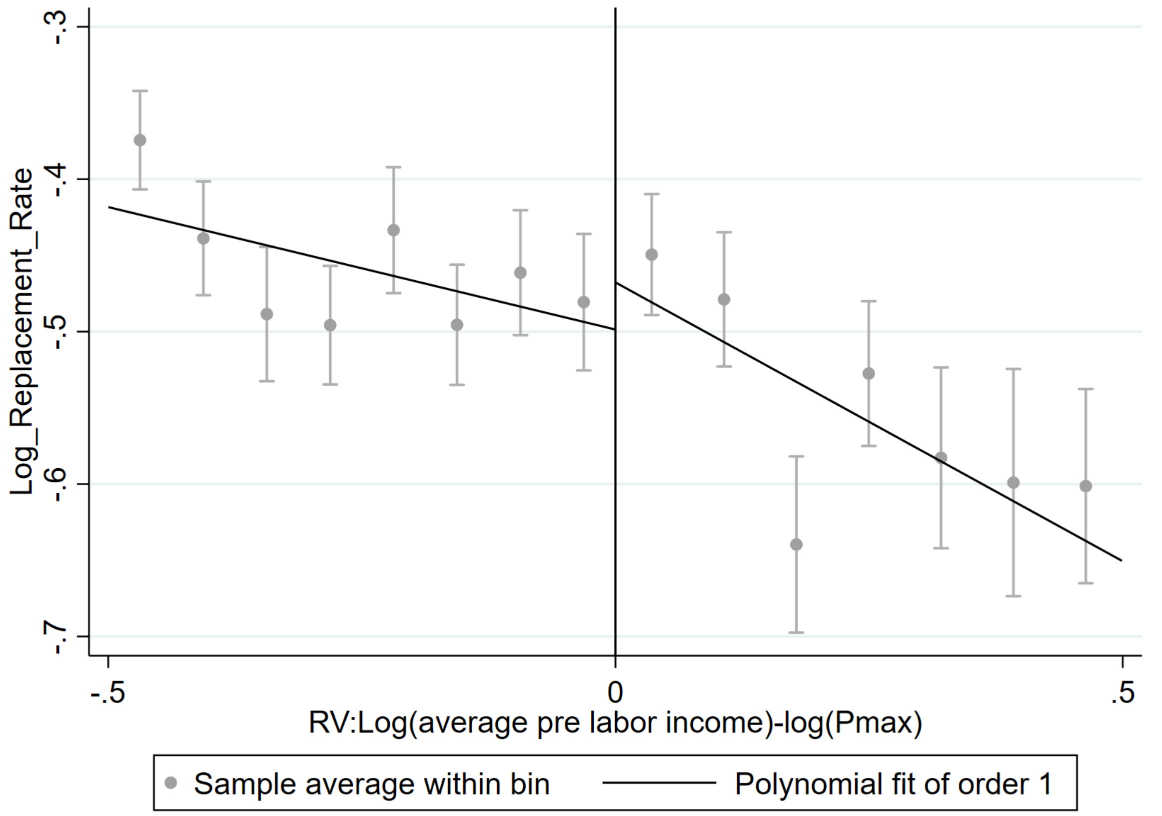

and shows a kink with a change in slope as preretirement income passes through the threshold. The probability that the replacement rate is above 70% is around 50% for individuals below the cap, and it does not vary with income. For individuals with pre-retirement earnings above the cap, the probability falls with pre-retirement earnings.

The change in the slope at 0 represents a change in the incentive to save. Consider two individuals with pre-retirement income below the maximum pension (0 in the horizontal axis). Both will experience a similar income drop at retirement. However, the public pension replacement rate changes with income at a different rate when pre-retirement incomes exceed the cap.

Table 4 presents the estimates of the kink in our policy variable (probability that the replacement rate is greater than 70%) around P_{MAX}. We estimate the specification in equation (5) using a linear polynomial and a bandwidth of $h = 1$ and $h = 0.5$, and a triangular kernel. Regardless of the bandwidth used, and in line with the results shown in Figure 5 suggest a statistically significant fall in the proportion of individuals whose public pension exceeds the legal limit of 70%.

Figure 5: Probability Replacement rate¿70%

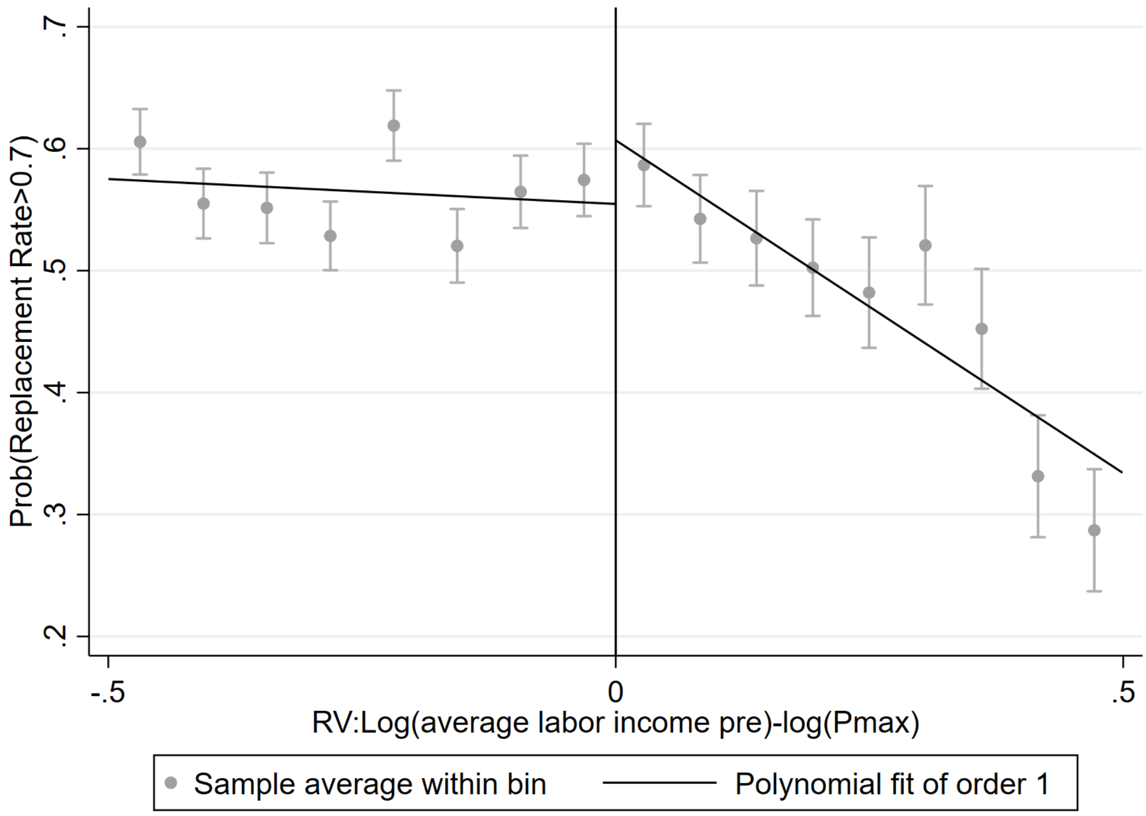

Table 4: RKD Estimates, Public Pension Repl. Rate

	$Kink = P_{MAX}$	
	$p = 1$	
	$h = 0.5$	$h = 1$
$\widehat{\gamma}_1^L$	-0.486	-0.284
	(0.144)	(0.067)
Total nobs.	13,845	22,007

Notes: Standard errors between brackets. Models estimated using local linear regressions centered at the maximum pension with two bandwidths: 0.5 and 1.

Figure 6 visually illustrates another change at the contribution limit: the increase in post-tax return to saving (price or substitution effect). In this context, when a euro is contributed to pension funds, it is exempt from taxes. It accumulates at a tax-free rate and is subject to taxation upon withdrawal. Consequently, this generates an increase in the returns to savings. The figure 6 displays the returns to saving by plotting the effective post-tax return on saving, represented as $R \times (1 - MTAXpost)/(1 - MTAXpre)$, against the running variable, which corresponds to the deviation of the (log of the) average income in the eight years before retirement from the contribution cap in the retirement year. The figure depicts a kink in the relationship, highlighting the impact of the contribution limit on the post-tax returns.

Figure 6: Incentives on Pension Funds savings

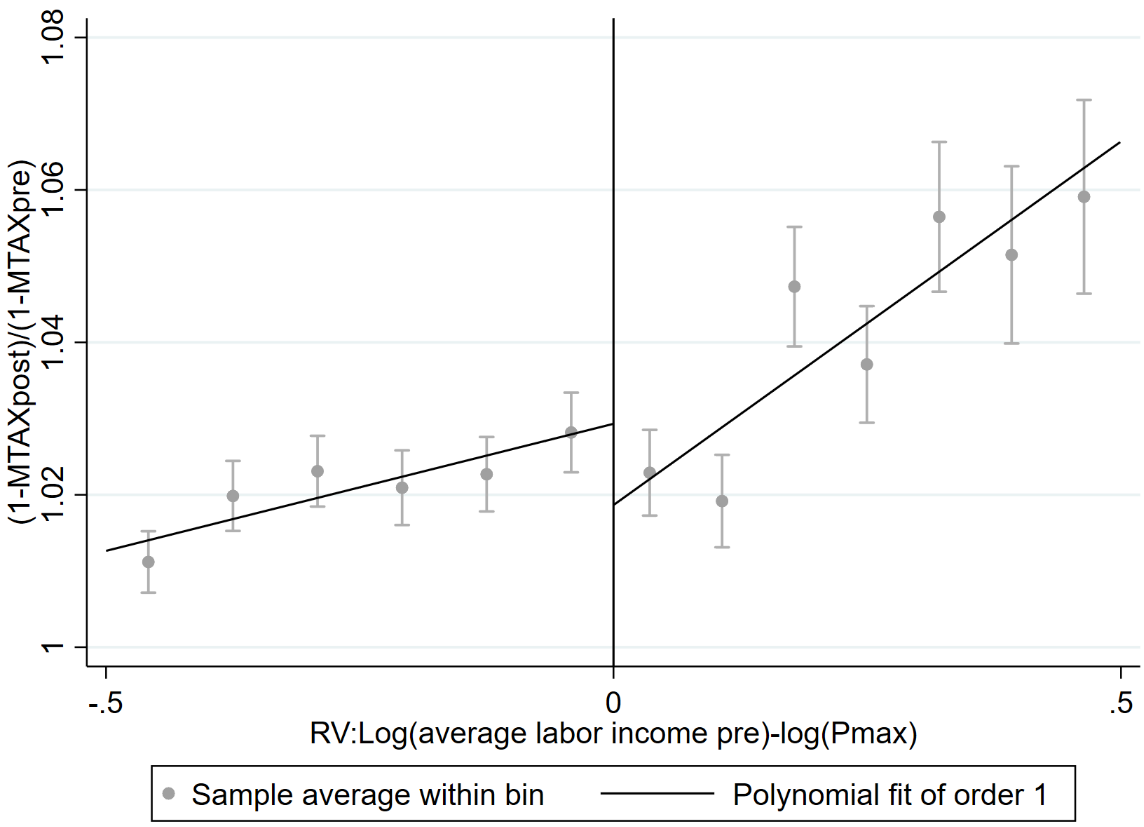

Table 5 shows the corresponding estimates of the previous relationship.

Table 5: Kink MTAX

	$Kink = P_{MAX}$	
	$p = 1$	
	$h = 0.5$	$h = 1$
$\widehat{\gamma}_1^{MTAX}$	0.064	0.065
	(0.027)	(0.015)
Total nobs.	13,845	22,007

Notes: Standard errors between brackets. Models estimated using local linear regressions centered at the maximum pension with two bandwidths: 0.5 and 1.

5.2.2 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 7 shows that the probability of contributing to a pension fund experiences a discrete change at the maximum pension, and an increase in the derivative of the probability of contributing with respect to income (see Table 6 for a quantification). The increase in the maximum pension is at odds with a forward-looking behavior of saving, and could be consistent with some salience explanation of contributions: tax filers start contributing when their income is at the maximum level.

Figure 7: Probability to contribute to Pension Funds

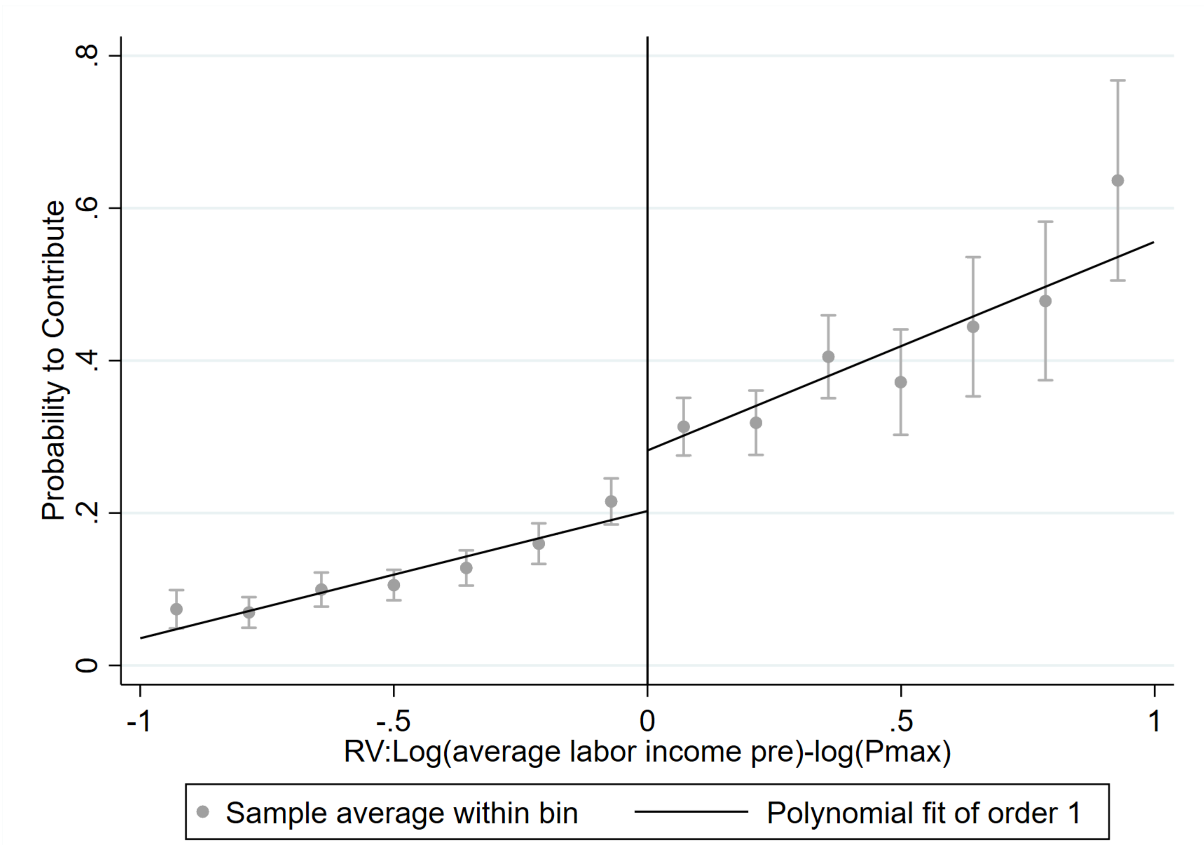

As to the fraction of income contributed (Figure 8) we obtain that above the cap, the relationship between contributions to pension funds and pre-retirement income increases by 0.47%, consistent with a change in saving behavior once the replacement rates start falling and the net interest rate increases -see Table 6 for the estimation results

Figure 8: Mean Contributions to Pension Funds

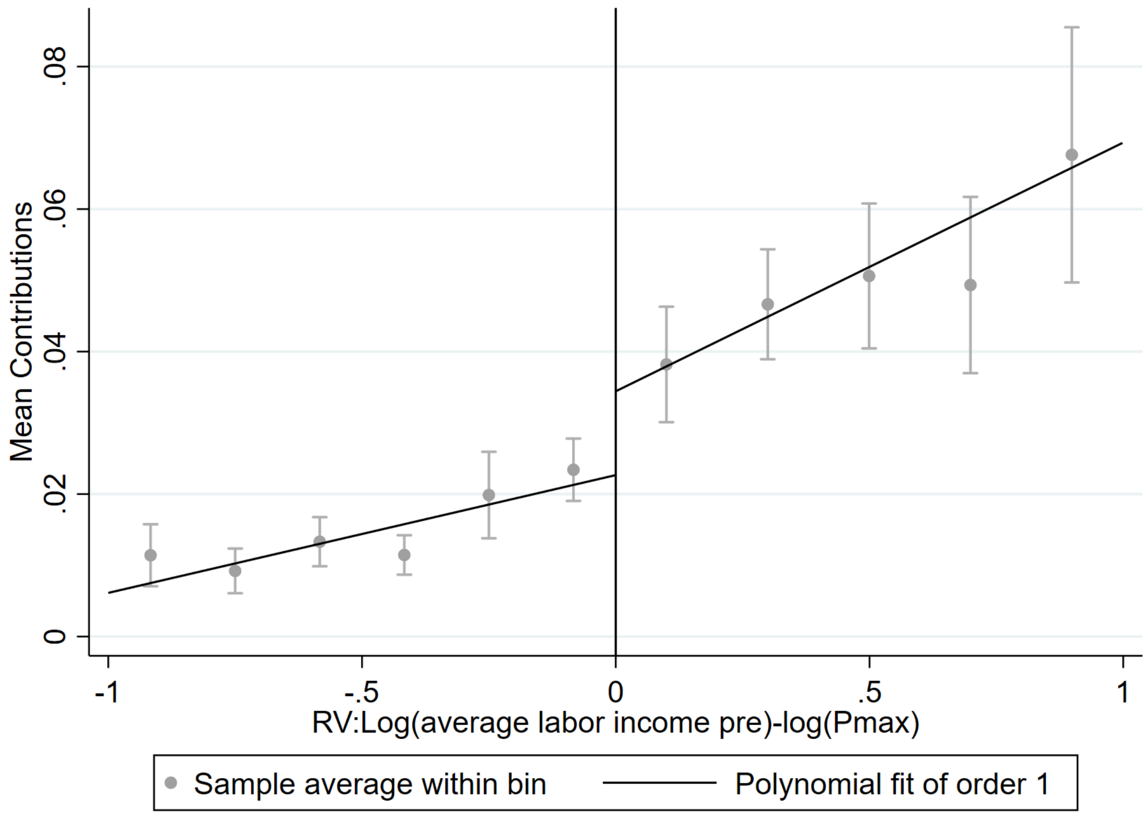

Table 6: Kink Prob. Contribute and Mean Contributions

	$Kink = P_{MAX}$	
	$p = 1$	
	$h = 0.5$	$h = 1$
$\widehat{\gamma}_1^{\text{Pr\_}Cont}$	.0491	.0596
	(.054)	(.027)
$\widehat{\gamma}_1^{Mean\_Cont}$	.0079	.0046
	(.0058)	(.0027)
Total nobs.	34,768	56,366

Notes: St. errors between brackets. Each tax filler-year contributes one observation. First row: uses as dependent variable an indicator of the fraction of tax filers who contribute to pension funds. The second row uses as the dependent variable of the contribution of pension funds to gross yearly income.

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

To examine the effectiveness of tax incentives for retirement savings in tilting the income trajectory over the life cycle toward later ages, our analysis focuses on the degree to which individuals adjust their income profiles to include higher levels of income during retirement. To assess this, we compute the total income at retirement, encompassing public pensions as well as income from private investments (such as financial income, rents, and self-employment income), and measure the probability that the total income replacement rate exceeds 70%.

The horizontal axis of Figure 9 is the deviation of the (log of) average labor income in the eight years preceding retirement from the contribution cap on the horizontal axis. On the vertical axis, we display the probability of the *total* replacement rate exceeding 70%. As discussed in Section 2, the relationship changes around the cap. Namely, consider tax filers whose income is below the maximum. Among those right above the cap, 65% have total post-retirement income higher than 70%. Among those whose income is .5 below the cap, the fraction is about 70%, not very different. The relatively flat relationship between total replacement rates and income was already discussed in Section 2 and is the result of constant public pension replacement rates and weak incentives to save due to a substitution effect. Instead, among filers whose pre-retirement income is .5 above the maximum, 55% have a total replacement rate above 70%. As mentioned in Section 2, the fall in replacement rates with income is to be expected, as savings cannot fully replace the fall in public pension replacement rates. However, as we mentioned momentarily, the relevant comparison relates to the change in the public and total replacement rates.

Figure 9: RV vs.Total Income Replacement Rate

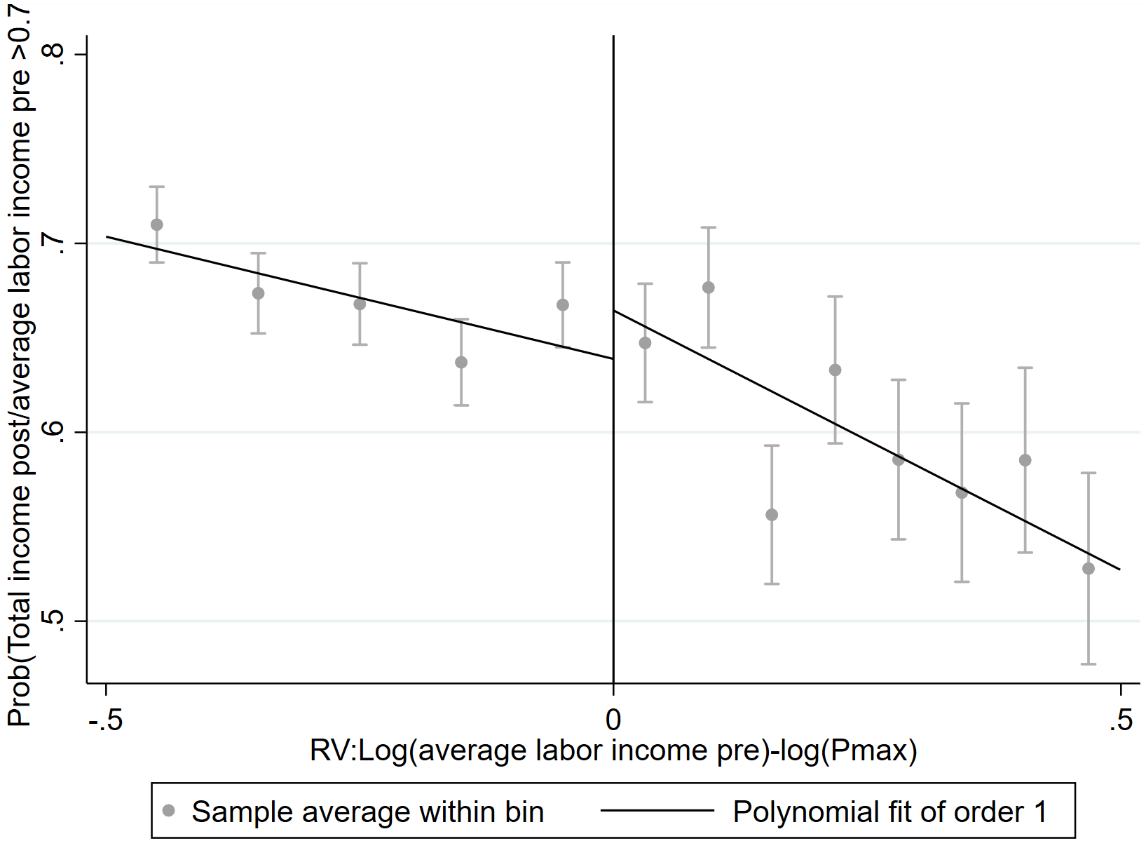

Figure 10 illustrates the changes in slope observed when analyzing the public pension replacement rate (left panel) versus the total replacement rate (right panel). It is noteworthy that, as mentioned earlier, there is a more pronounced change in slope for the former, in absolute value, compared to the latter. Table 7 presents the estimation results, with our primary focus on the effect measured by the ratio $\frac{\hat{\gamma}_1^T}{\hat{\gamma}_1^L}$. The estimate is .40, suggesting that pension funds absorb 60% of the fall in post-retirement income due to the public pension replacement rate.

Figure 10: Public Pension RR and Total RR

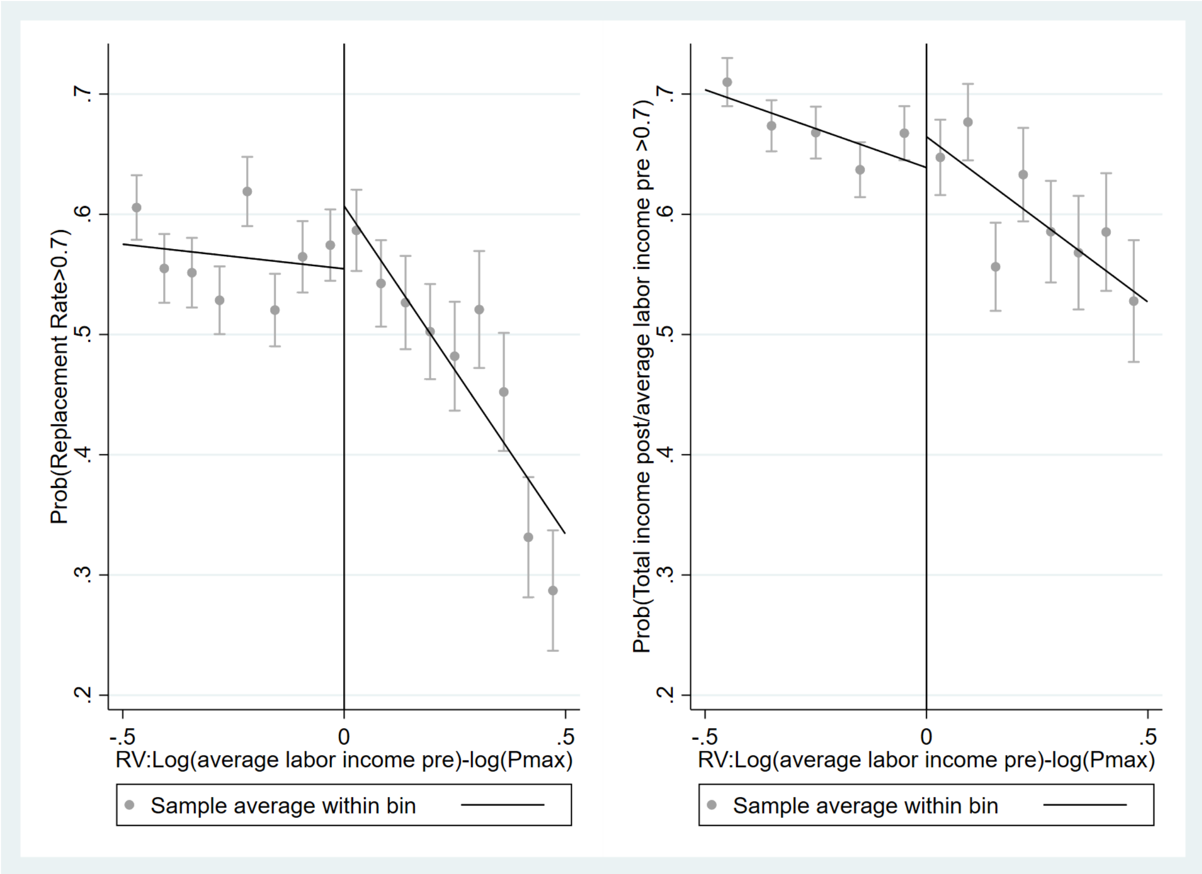

Table 7: Estimation Results

	$Kink = P_{MAX}$	
	$p = 1$	
	$h = 0.5$	$h = 1$
$\widehat{\gamma}_1^T$	-0.200	-0.122
	(0,134)	(0.064)
$\widehat{\gamma}_1^L$	-0.486	-0.284
	(0.144)	(0.067)
$\frac{\widehat{\gamma}_1^T}{\widehat{\gamma}_1^L}$	0.412	0.430
Total nobs.	13,845	22,007

Notes: St. errors between brackets.

5.3 By income type

TO DO

6 A benchmark: cohorts retiring before 1988

As a benchmark, we perform a similar analysis but with the individuals retired before 1988. Before 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 the 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 and present our graphical results. Although not using the full eight years to compute the running variable results in some noise, Figure shows an elasticity of public income (only public pension) to pre-retirement income of $-.376$, while according to the Figure 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.

INCOMPLETE

References

Attanasio, O., and A. Brugiavini (2003), “Social Security and Households’ Saving”, *The Quarterly Journal of Economics,* 118(3), 1075-1119.

Card, D., Lee, D.S., Pei, Z., and A.Weber (2016), “Regression Kink Design: Theory and Practice”, NBER Working Papers 22781, National Bureau of Economic Research, Inc.

Chetty, R., Friedman, J.N., Leth-Petersen, S., Nielsen, T.H., and Toren Olsen (2014), “Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark”, *The Quarterly Journal of Economics*, 129(3), 1141–1219.

Engen, E., and J. Gruber (2001), “Unemployment Insurance and Precautionary Saving”, *Journal of Monetary Economics,* 47, 545–579.

Heckman, J. (1979), “Sample selection bias as a specification error”, *Econometrica*, 47, 153-161.

Moser, C., and P. Olea de Souza e Silva (2019), “Optimal Paternalistic Savings Policies”, Columbia Business School Research Paper No. 17-51.