How Large is the Retirement Consumption Drop in Italy?*

Erich Battistin
University of Padova and IFS

Agar Brugiavini
University of Venice and IFS

Enrico Rettore
University of Padova

Guglielmo Weber
University of Padova, CEPR and IFS

first draft September 2005
this version 27th February 2006

Abstract

In this paper we investigate the size of the consumption drop at retirement in Italy. We use micro data on food, non-durable and total household spending covering the period 1991-2002, and evaluate the change in consumption that accompanies retirement by exploiting the exogenous variability in pension eligibility to correct for the endogenous nature of the retirement decision. We take a regression discontinuity design approach, and make the identifying assumption that consumption would be the same around the threshold for pension eligibility, if individuals would not retire. We check in our data that a non-negligible fraction of individuals retire as soon as they become eligible, and estimate at 9.2% the part of the consumption drop that is associated with retirement induced by eligibility.

Keywords: Consumption, Regression Discontinuity Design, Retirement
JEL Classification: D9; E2

*Paper presented at the RTN meeting on the Economics of Ageing, Venice, October 2005. We thank our discussant, Mike Hurd, and conference participants for their helpful comments. We are grateful to Alberto Fossaluzza for excellent research assistance, and to Lisa Callegaro for explanations on the codes that calculate pension eligibility for SHIW respondents.

1 Introduction

In most developed countries, consumption accounts for over two-thirds of GDP. In these countries a rising fraction of the population is past retirement age. The way consumers respond to retirement and the way they spend in their old age is thus a topic of great interest in the analysis of aggregate economic fluctuations and in the economic policy debate.

The standard model to analyse the consumption-saving choice by the household sector has been Modigliani's life cycle model. The model has been extended to cover uncertainty, leisure choice and a bequest motive (Deaton, 1992, Browning and Lusardi, 1996), but its key prediction can still be described as follows: consumers form intertemporal plans aimed at smoothing their standard of living (or marginal utility of consumption) over their life-cycle.

However, recent micro evidence has emphasized that there is a one-off drop in consumption at the time of retirement that may be hard to reconcile with life-time optimizing behavior. This is documented for the UK (Banks, Blundell and Tanner, 1998), for the US (Bernheim, Skinner and Weinberg, 2001), and for Italy (Miniaci, Monfardini and Weber, 2005, MMW in what follows) and is known as the retirement consumption puzzle (or retirement savings puzzle).

The literature mentions as possible reasons for this drop changes in preferences due to increased non-market time, unexpectedly low pensions or liquidity problems as well as myopic or perhaps time-inconsistent behavior (see Hamermesh, 1984, Hurd and Rohwedder, 2003, and Aguiar and Hurst, 2004).The Italian case is of particular interest because one can rule out explanations related to lack of resources. In fact, MMW document that actual and expected replacement rates were indeed close to each other independently of the type of job previously held by the newly retired. Also, liquidity problems are unlikely to play a role: Italian employees receive a large lump-sum payment upon retirement (technically, a severance pay worth three times the gross annual salary). If cash considerations matter, we would expect a surge in consumption at retirement rather than a drop. These two facts suggest that consumption falls at retirement cannot be attributed to unexpected income drops or liquidity problems.

The evidence provided in MMW is based on SFB diary-level data covering the 1985-96 period. The SFB (Survey on Family Budgets, run by ISTAT) contains excellent information on expenditures, but no information on past work activities. Thus MMW estimates the fall in non-durable consumption at retirement, and emphasizes the role that changes in leisure have on both the composition and the level of consumption, but cannot identify what part of the consumption fall at retirement is as planned and what other part is instead due to the realization of bad health outcomes or other shocks to the consumer's environment that affect the retirement decision.

In this paper we investigate the size of the consumption drop associated to retirement in Italy by exploiting the exogenous variability in eligibility to identify the causal effect in a regression discontinuity design framework. To this end, we use a different micro data set, the Bank of Italy Survey on Household Income and Wealth (SHIW), covering the 1991-2002 period, that has recall records on food, non-durable and total household spending (Battistin, Miniaci and Weber, 2003, compare consumption data across SFB and SHIW). We evaluate the change in consumption caused by retirement by exploiting the exogenous variability in pension eligibility to correct for the endogenous nature of the retirement decision: our identifying assumption is that consumption would be the same around the threshold for pension eligibility, if no individual retired.

The remainder of this paper is organized as follows. Section 2 presents how we deal with the endogeneity problem arising from self-selection of individuals into retirement. Section 3 deals with data-related issues, in particular with the definition of the pension eligibility. In Section 4 we show that pension eligibility is a variable that is measured with error, and we discuss the implications of this on our estimates. Section 5 presents our results. The implications of our findings for reconciling the retirement savings puzzle are presented in Section 6, while Section 7 concludes.

2 Identification

2.1 The regression discontinuity idea

This section presents the basic features of regression discontinuity analysis following the discussion in Hahn *et al.* (2001), to which the interested reader is referred for further details. The relationship with the literature on programme evaluation is established by comparing the retirement decision to that of being exposed to a treatment.

Following the notation of the potential outcome approach to causal inference, let (Y_1, Y_0) be the two potential outcomes one would experience by retiring and not retiring, respectively. In the context of this paper, Y_1 and Y_0 represent the household consumption expenditures corresponding to the head being retired and not being retired, respectively. The causal effect of retirement on expenditures is then defined as the difference between these outcomes, $\beta = Y_1 - Y_0$, which is not observable at the household level since being retired reveals Y_1 but conceals Y_0. Accordingly, though not observable β represents the change in consumption expenditures corresponding to a change in the retirement status of the household head, which is our quantity of interest.

Let R be the binary variable denoting the treatment status, with $R = 1$ for retired heads and $R = 0$ otherwise. A discontinuity design (Thistlethwaite and Campbell, 1960) arises when the treatment status R depends on an *observable* variable S and there exist a *known* point in the support of S where the probability of being treated changes discontinuously. Formally, if $\bar{s}$ is the discontinuity point, then a regression discontinuity is defined if

$$Pr\{R = 1|\bar{s}^+\} \neq Pr\{R = 1|\bar{s}^-\}. \tag{1}$$

Here and in the following $\bar{s}^+$ and $\bar{s}^-$ refer to those individuals *marginally* above and below $\bar{s}$, respectively. Throughout this paper, S is assumed to be continuous on the real line.

In the context of this paper, the expression in (1) implies that the probability of retirement varies discontinuously with an *observable* variable S. Throughout our analysis, S will denote the distance from the first time the household head becomes eligible for retirement. To fix ideas, let the *eligibility status* be established according to the deterministic rule $\mathbb{1}(S \geq 0)$. That is, individuals are eligible for retirement if and only if they present a value of the variable S above the threshold $\bar{s} = 0$. Of course, such a variable can take on negative values (if individuals are not yet eligible for retirement) as well as positive values (if individuals, regardless of their retirement status, already are).

The potential of using eligibility rules to overcome the selection problems arising in the study of causal effects has been already pointed out by several papers in the literature (see, amongst others, Battistin and Rettore, 2005). In the context of this paper, pension eligibility does not necessarily imply that individuals are actually retired; on the other hand, individuals not eligible can not be retired, thus inducing a discontinuity in the probability of retirement around the threshold for eligibility. As we will show in what follows, such a discontinuity can help solve the endogeneity problem arising from the analysis of the retirement status.

Following Trochim (1984), the distinction between *sharp* and *fuzzy* designs depends on the size of the discontinuity. The former design occurs when the probability of participating conditional on S steps from zero to one as S crosses the threshold $\bar{s}$. That is, the treatment status deterministically depends on whether individuals' values of S are above $\bar{s}$

$$R = \mathbb{1}(S \geq \bar{s}). \tag{2}$$

For example, a sharp design would correspond to the hypothetical situation in which retirement is mandatory conditional on eligibility: is such an extreme case, we would have that $S \geq 0$ implies $R = 1$ with probability one.

A fuzzy design occurs when the size of the discontinuity at $\bar{s}$ is smaller than one. For example, a fuzzy design can be thought as an instance in which R is a deterministic function of S for all subjects but this function is different across individuals (see Hahn *et al.*, 2001). It therefore follows that the retirement decision fits a fuzzy regression discontinuity conditional on S, as not all eligible individuals are actually retired and all ineligibles are in fact not retired.

As the decision to retire is entirely up to eligible individuals, it is crucial to discuss how the endogeneity problem arising from the retirement status can be accounted for in the context of discontinuity designs. We will consider the case of a sharp discontinuity first, and discuss the fuzzy case further below in this section. Let

$$Y = Y_0 + R(s)\beta$$

be the observed outcome as it results from taking or not taking part into the programme: it follows from the last expression that $Y \equiv Y_1$ or $Y \equiv Y_0$ depending on the retirement status of individuals ($R = 1$ and $R = 0$, respectively). The dependence of the retirement status R on S is stressed by writing $R(s)$. The difference of mean outcomes for individuals marginally above and below the threshold $\bar{s}$

$$E\{Y|\bar{s}^+\} - E\{Y|\bar{s}^-\} \tag{3}$$

can be written as

$$E\{Y_0|\bar{s}^+\} - E\{Y_0|\bar{s}^-\} + E\{R(s)\beta|\bar{s}^+\} - E\{R(s)\beta|\bar{s}^-\}, \tag{4}$$

which simplifies to

$$E\{Y_0|\bar{s}^+\} - E\{Y_0|\bar{s}^-\} + E\{\beta|\bar{s}^+\}$$

because of (2). The following condition is then sufficient for the mean impact of the treatment at $\bar{s}^+$ to be identified with a sharp discontinuity.

Condition 1. The mean value of Y_0 conditional on S is a continuous function of S at $\bar{s}$.

Accordingly, Condition 1 requires that in the counterfactual world of no retirement, no discontinuity would take place at the threshold for selection. In other words, this implies that the consumption profile under the no retirement alternative is smooth enough as S crosses $\bar{s}$. Intuitively, in order to give a causal interpretation to discontinuities of Y around the threshold for eligibility $\bar{s} = 0$, it has to be the case that in the absence of the treatment no discontinuity would be observed in the outcome Y around $\bar{s} = 0$. If this condition holds, we can write

$$E\{\beta|\bar{s}^+\} \equiv E\{Y|\bar{s}^+\} - E\{Y|\bar{s}^-\},$$

so that the difference in expected consumption expenditures above and below the threshold for eligibility identifies the causal effect of retirement on consumption.[1]

When the treatment status is not the result of a sharp assignment, the discontinuity in the probability to retire around the threshold is smaller than one. According to the current literature, such a discontinuity in the probability of retirement defines a fuzzy design. It follows that the mean impact at $\bar{s}$ cannot be identified by simply comparing the mean outcome for marginal retired to the mean outcome for marginal non-retired households. In general, additional conditions are required to recover meaningful causal parameters from (3), thus losing much of the simplicity of the design. Hahn *et al.*

[1] It is worth stressing again that to meaningfully define marginal units (with respect to $\bar{s}$) the selection variable S has to be continuous. Estimation of the right-hand side (left-hand side) of (3) makes use of data only in a neighborhood on the right (left) side of the discontinuity point. Unless one is willing to make some parametric assumptions about the regression curve away from $\bar{s}$, only data local to the discontinuity point help to estimate the jump. Asymptotically the neighborhood needs to shrink as with usual non-parametric estimation, implying a non-standard asymptotic theory for the resulting estimator of the mean impact (see Hahn *et al.*, 2001).

(2001) as well as many other authors in the literature point out that assumptions can be made to recover causal effects for a particular group of individuals around the threshold $\bar{s}$. Such assumptions qualify S as an instrumental variable for R around $\bar{s}$, so that a LATE (Angrist and Imbens, 1994) parameter can be estimated for the group of compliers. Heckman *et al.* (1999) emphasize this point by saying that much of the simplicity of the design is lost moving from a sharp design to a fuzzy design.

2.2 Discontinuities and pension eligibility

In the context of this paper, however, self selection of households into retirement fits the partially fuzzy design described by Battistin and Rettore (2005).[2] As a result of the eligibility rule and of self-selection, the probability of retirement for those heads scoring a value of S below the threshold $\bar{s}$ is zero by definition, since they are not eligible for retirement. The probability of retirement for those scoring above $\bar{s}$ is smaller than one because retirement is not mandatory. As a result, the probability of retirement is discontinuous at the threshold for eligibility and the size of the discontinuity is less than one (i.e. according to the terminology introduced in the previous section, a fuzzy design is defined). As pointed out by Battistin and Rettore (2005), despite the fuzziness of this design the existence of a sharp eligibility rule can help recover much of the simplicity of the design.

To recover the regularity conditions required for identification consider again the difference in (3). Since participation is precluded to marginally ineligibles ($R(\bar{s}^-) = 0$), the expression in (4) can be written as

$$E\{Y_0|\bar{s}^+\} - E\{Y_0|\bar{s}^-\} + E\{R(s)\beta|\bar{s}^+\}.$$

If Condition 1 holds, by using the law of iterated expectations and by noting that $E\{R(s)\beta|R = 0, \bar{s}^+\} = 0$ the previous expression equals

$$E\{R(s)\beta|\bar{s}^+\} = E\{\beta|R = 1, \bar{s}^+\}Pr\{R = 1|\bar{s}^+\},$$

so that the mean impact on participants in a right-neighborhood of $\bar{s}$ is identified by

$$E\{\beta|R = 1, \bar{s}^+\} = \frac{E\{Y|\bar{s}^+\} - E\{Y|\bar{s}^-\}}{E\{R|\bar{s}^+\}}. \tag{5}$$

In other words, Condition 1 is sufficient for the effect of the treatment on the treated to be identified locally at the threshold for eligibility $\bar{s}$.

It turns out that, despite the *prima facie* fuzzy nature of this set-up, the LATE (Imbens and Angrist, 1994) at the discontinuity point is identified under the same condition used to estimate the average treatment effect in the sharp design. The result rests on the fact that the probability of retirement on the left-hand side of $\bar{s}$ is zero by design, and this simplifies the expression in (4) without further assumptions on individuals' behavior.[3]

It also follows that (5) can be estimated from an instrumental variable procedure, where eligibility is used as an instrument for the actual status R conditional on S.

3 Data

3.1 The definition of pension eligibility

In our analysis, information on consumption expenditures and pension status is obtained from the Bank of Italy Survey on Household Income and Wealth (SHIW in what follows). The aim of this

[2]Heckman *et al.* (1999) as well as Van der Klaauw (2002) explicitly mentions the potential for using the dicontinuity arising from the eligibility criteria for a social programme.

[3]Results by Hahn *et al.* (2001) on non-parametric inference in regression discontinuity designs straightforwardly apply to the estimation of (5).

Table 1: Composition of the SHIW sample

	1991	1995	1998	2000	2002
Males					
Worker	47%	44%	44%	45%	44%
Retired	17%	20%	19%	22%	25%
Females					
Worker	24%	25%	26%	26%	27%
Retired	12%	14%	12%	13%	16%

section is to summarize how we derive the variable S that measures the time to (or from) pension eligibility (see also Boeri and Brugiavini, 2005).

As we have shown in the last section, pension eligibility is a crucial variable in our analysis. This is measured both on the basis of *age* and on the basis of *seniority* (accrued contributions years). The SHIW sample can be used to compute eligibility: it is a large cross-sectional sample and covers several years (1987 to 2002), though it has been run every two years. SHIW also has a small panel component (rotating panel) that can be used to study actual transitions from work to retirement.

Table 1 provides a brief description of the SHIW data available for workers and retired individuals (we neglected other conditions such as disability or being homemaker). It should be noted that a large percentage of individuals who are currently active or have been active in the past are men. This result is largely dominated by the labour market behaviour of older cohorts: women were characterized by lower educational attainments and lower labour market participation. An indication of this is the fact that over time there is a non negligible increase in the percentage of working women.

3.1.1 Working life and pension claims

A crucial feature of many pension systems is the design of pension benefits: in most European countries this is the defined benefit (DB) variety and it is related to some average of lifetime earnings. An extreme version of this is a final salary computation method: this was basically adopted in Italy until the 1992 pension reform as well as during the transitional phase which Italy is going through. In particular, until 1992 the pension benefit was based on the average of the last five years earnings, during the transitional phase these became the last 10 years earnings. The 1995 Pension reform changes radically this system as benefit should be computed according to a Notionally Defined Contribution (NDC) method. In the latter case, pension benefits are automatically linked to an average of lifetime earnings, adjusted by some actuarial coefficients.

In this paper we are particularly interested in eligibility conditions: until 1992-1993 they were quite simple as people could retire at age 60 (55 women) in the private sector, or any age if they had completed 35 years of contributions. The early retirement option was quite generous because it did not attract any actuarial penalty and a large fraction of workers retired before the reforms through this route. After the 1992 reform the normal retirement age were set at 65 for men and 60 for women (to be reached gradually by the year 2001). Both age and seniority requirements for early retirement grew over time, starting essentially in 1995 (according to the sector of employment) as shown in Table 2 (rules prevailing after 1998 according to the Law 449/1997; these rules apply to white- collar employees, they differ only slightly for blue-collar employees). There was enough flexibility offered by these criteria which we can model explicitly in our data. One final problem to be mentioned is that, with some interruptions, the Government has imposed a delay-window on retirees after they became eligible, which, depending on the month of birth, could postpone retirement by 3 or even 6 months. While we cannot observe the effects of these windows directly, we do not think they would introduce too much noise in our data because we measure consumption and eligibility at yearly frequency.

Table 2: Retirement eligibility rules: age and years of contributions

	Private Sector		Public Sector		Self-employed	
	age and years	only years	age and years	only years	age and years	only years
1996	54 and 35	36	53 and 35	36	57 and 35	40
1997	54 and 35	36	53 and 35	36	57 and 35	40
1998	54 and 35	36	53 and 35	36	57 and 35	40
1999	55 and 35	37	53 and 35	37	57 and 35	40
2000	55 and 35	37	54 and 35	37	57 and 35	40
2001	56 and 35	37	55 and 35	37	58 and 35	40
2002	57 and 35	37	55 and 35	37	58 and 35	40
2003	57 and 35	37	56 and 35	37	58 and 35	40
2004	57 and 35	38	57 and 35	38	58 and 35	40
2005	57 and 35	38	57 and 35	38	58 and 35	40
2006	57 and 35	39	57 and 35	39	58 and 35	40
2007	57 and 35	39	57 and 35	39	58 and 35	40
2008	57 and 35	40	57 and 35	40	58 and 35	40

3.1.2 The measurement of years of contributions

In some years contributions have been explicitly recorded in SHIW for workers (not for people currently retired); however this question was not asked in 1993 and missing values are occasionally found also in other years. Because the age of the respondent does not provide enough information to measure eligibility, we have adopted a simple imputation method, by distinguishing retirees from workers. In fact, we make use of retrospective information on (i) the age that the respondent reported as age entering the labour market, and (ii) the self-reported age of retirement (if retired). The imputation is carried out also taking account of whether the individual belongs to the panel component or not.

We can distinguish the following *five* cases. First, if the *worker* is interviewed *only once* in 1993, then we have no choice than to use current age minus the self-reported age of entry in the labour market. Second, if the *worker* belongs to the *panel* sub-sample of SHIW, such a measure of years of contributions is corrected to account for the contributions recorded in successive years. Third, for a *retired* person who is *only* observed in 1993, use is made of the retirement age minus entry in the labour market. Fourth, if a *retired* person belongs to the *panel*, a coherency check is carried out. Finally, if a retired person belongs to the panel and worked in a previous wave, contributions can be recovered from the previous wave.

It is of course worth looking at the difference between the measure of years contributions as self-reported by individuals and the measure imputed by our procedure. This imputation could be a noisy measure because the respondent can have a vague recollection (especially if far in the past) of the events. However as Figure 1 shows, if we exclude the year 1993 there is basically complete coincidence between the original years of contributions and the final sample (including imputations based on potential experience). When we include the year 1993 there is a more pronounced heaping at age 35 for men, while for women the distribution shows an equally disperse pattern (see Figure 2).

It follows that the variable which measures the distance from the first eligibility year (S) is computed by first establishing on the basis of recorded age and years of contributions the eligibility year and taking the difference with the survey year. The distribution of this variable for males and females is reported in Figure 3. Negative values of this variable imply that eligibility for retirement has not yet been attained. Positive values measure the time from the first year of eligibility.

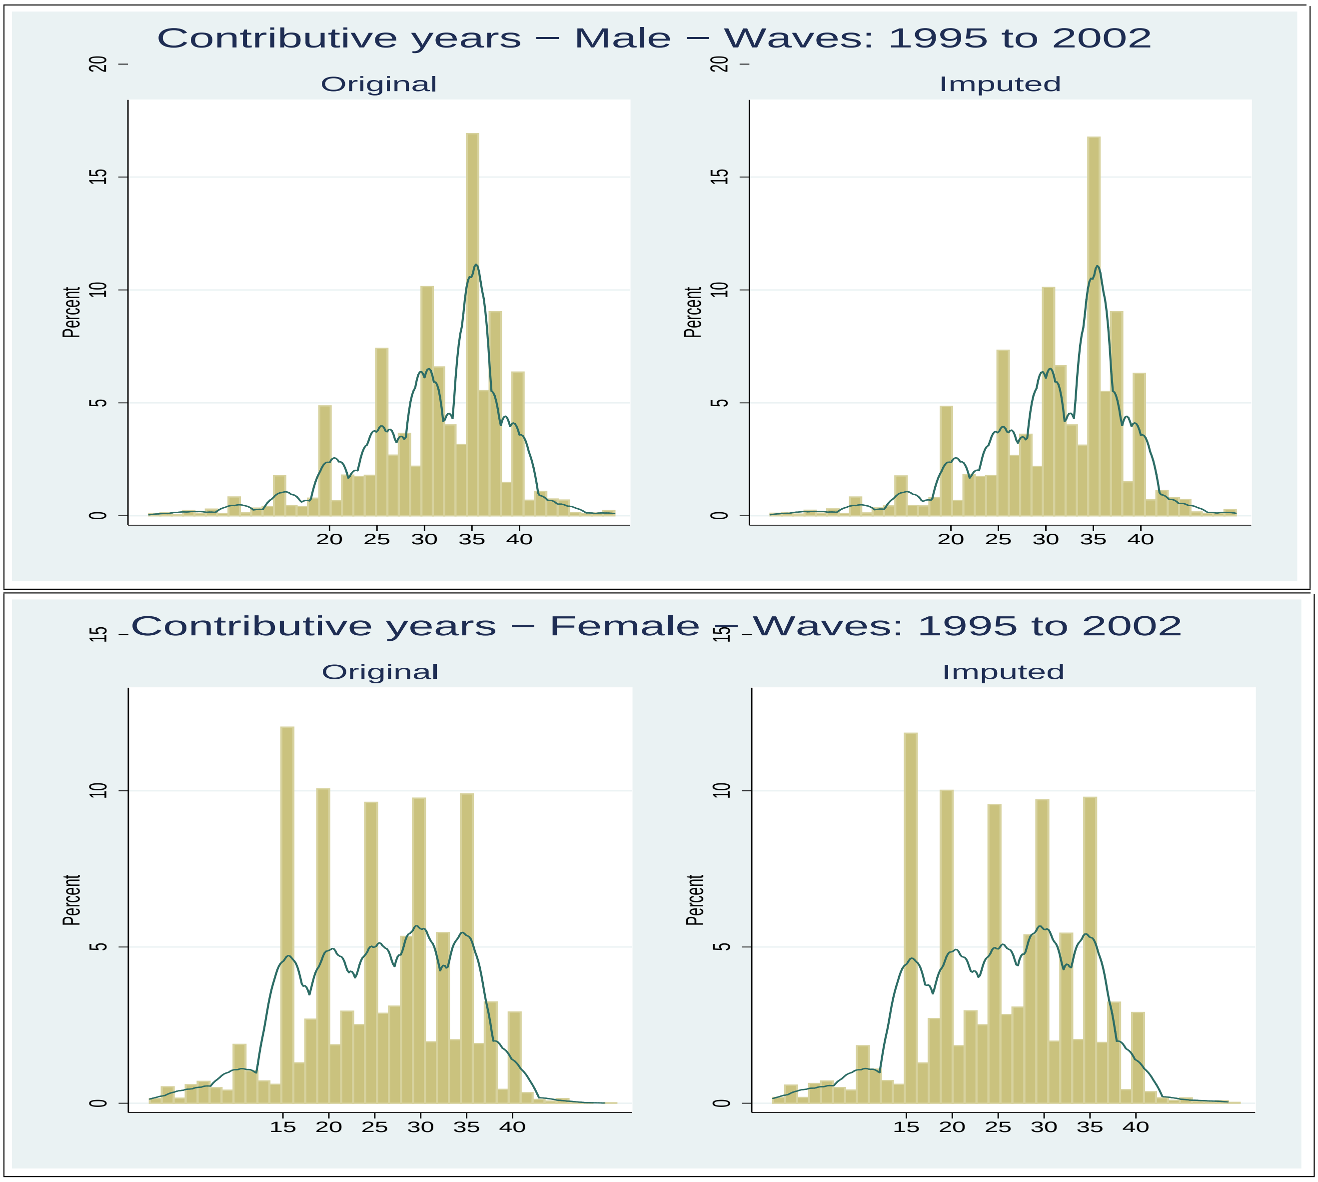

Figure 1: Contributive years - waves 1995 to 2002

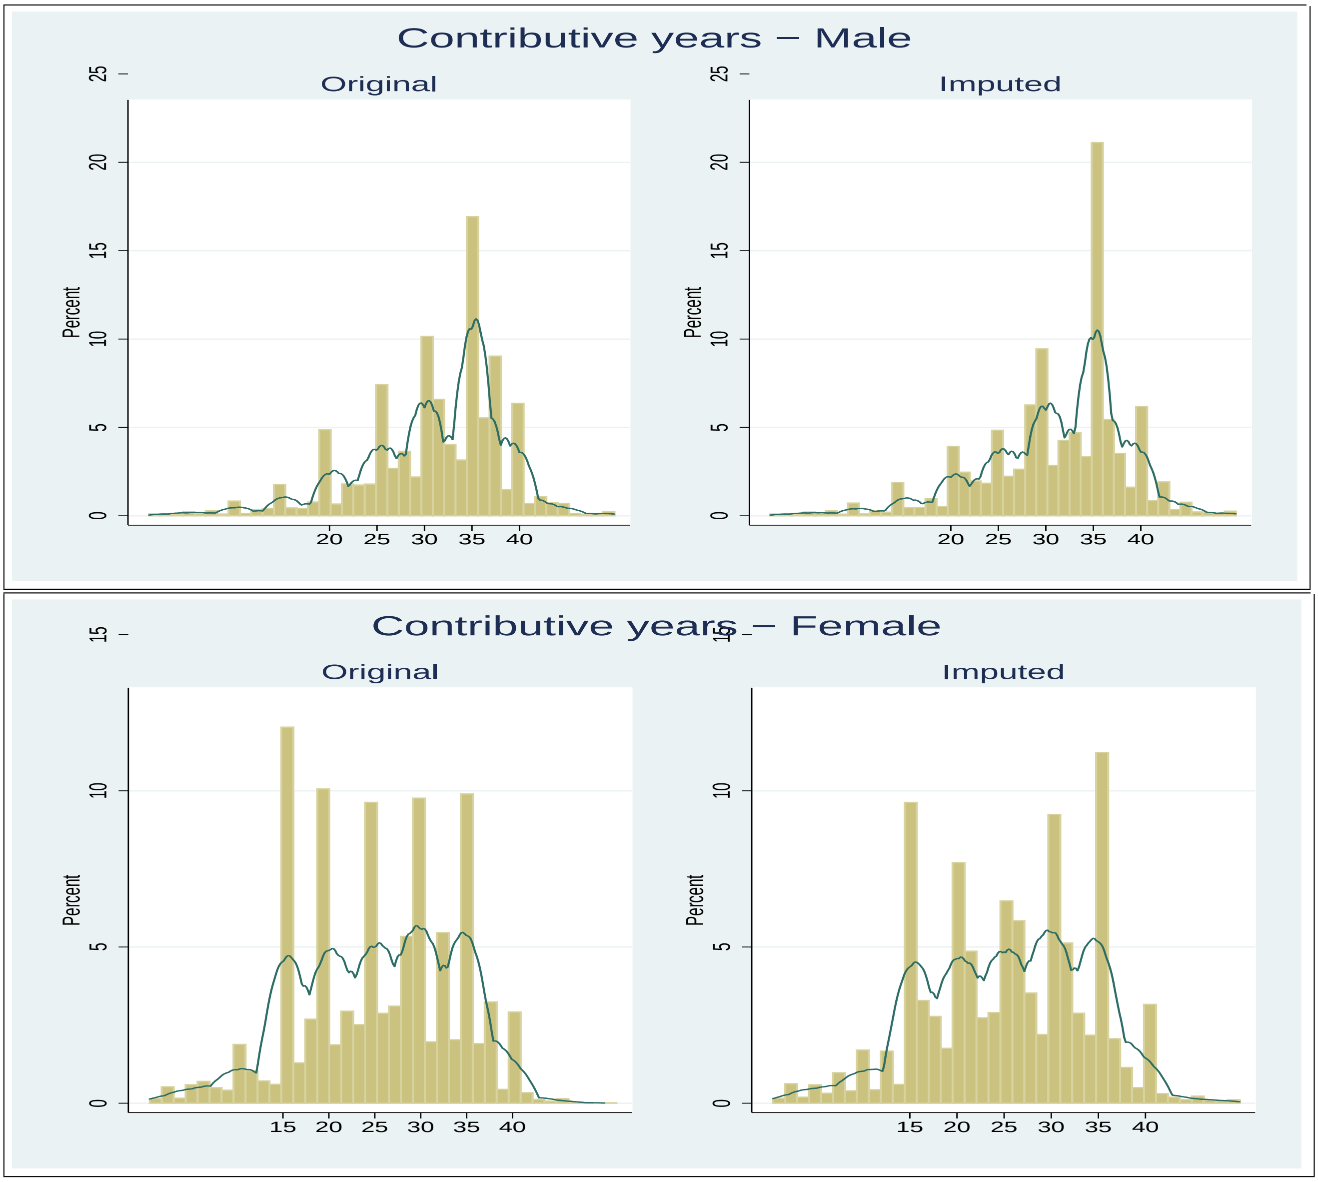

Figure 2: Contributive years - waves 1993 to 2002

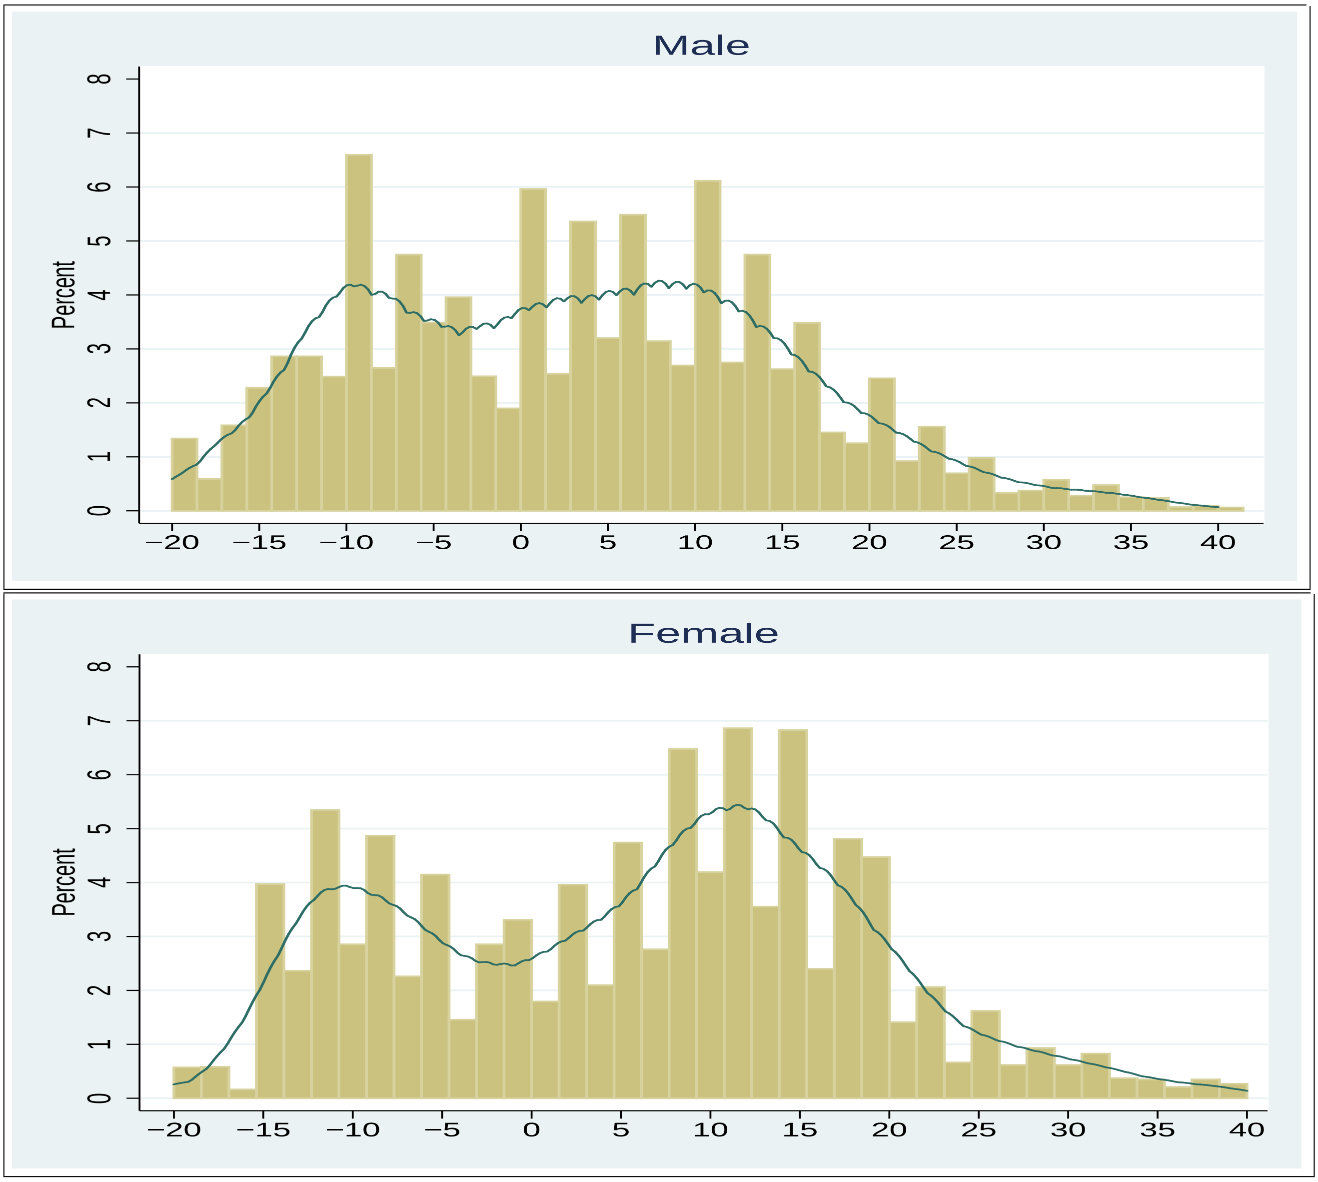

Figure 3: Distribution of S

Table 3: Percentage of retired individuals by eligibility status

eligibility	**males**	**females**
-5	0.0441426	0.0269058
-4	0.1073826	0.0681818
-3	0.0893855	0.0431655
-2	0.0677966	0.0854701
-1	0.1645161	0.2121212
0	0.2810345	0.4880952
1	0.5779377	0.7289157
2	0.5747663	0.7167630
3	0.6376812	0.7113402
4	0.7164179	0.7927461
5	0.6564885	0.6898148

3.1.3 Years of contributions and the retirement status

To conclude this section, in Table 3 we present the percentage of individuals self-reporting to be retired by the values of the eligibility variable S. As already discussed above, since retirement can only be entered conditional on eligibility, it has to be the case that no retired individuals are observed for negative values of the eligibility variable. If this were the case, it could be symptomatic of errors in the recording of the eligibility variable or, possibly, misreporting of the retirement status (or both).

Despite the sharp design implied by the eligibility rule, we find that a non-negligible proportion of individuals whose imputed value of the eligibility variable is negative are in fact retired. Notably, something like 16 percent of males and 21 percent of females marginally ineligible for retirement self-report to be retired. These figures decrease as the time to the first eligibility year increases, though proportions are still non-negligible when $S = -5$.

Sadly enough, it is somehow immediate to conclude that the identification strategy outlined in Section 2 no longer applies. More precisely, estimators of the causal effect of retirement on consumption based on discontinuities around $S = 0$ are in general biased for the parameter of interest. In the next section, we will address this problem and propose and estimation strategy which - under certain assumptions on the data generating process of the errors - provides consistent estimates of the causal parameter.

3.2 Sample selection

In estimation we take all observations that are within a 10-year band from eligibility. However we drop those that are exactly at eligibility, because the recall question on consumption could cover both pre and post-retirement periods. We investigated different selection criteria: smaller bands (such as 5-years bands) are heavily affected by outliers in the age profile; with larger bands (15-year and more) composition effects start playing an important role, because of mortality.

In our empirical exercise, we take the head of the household to be the man. We select couples and single males, and do not use information on the wife's employment position to classify a household as retired or otherwise.

4 Measurement errors

Throughout this section we will allow for measurement error in the variable S, but we will maintain the assumption that R in not mismeasured. For a detailed analysis of the impact of measurement error in S for the identification of the causal effects of a binary treatment see Battistin and Chesher

(2004). For the impact of misclassification (and/or misreporting) of R see Battistin and Sianesi (2005). Note, however, that in the context of this paper having measurement error in S implies that also the eligibility status is potentially misrecorded, thus inducing a more complicated structure of the error.

The motivation of allowing for measurement error in the time-to(from)-eligibility variable S builds on the evidence presented in Section 3.1, in that a non-negligible fraction of non-eligibles for pension indeed self-report to be pensioners. Under the maintained assumption that the pension status is not misreported, such evidence is consistent with the problem of having measurement error in S. It is worth noting that the crucial assumption of correct classification of pensioners can be motivated by the structure of the SHIW questionnaire, in that those individuals who self-report to receive pension benefits are actually forced to answer a very detailed set of questions in this respect.

The results in this section can be summarized as follows. First, we show that the evidence provided in Section 3.1 is *not* consistent with the hypothesis of having classical measurement error in S. A more general model for measurement error is therefore needed. Second, we assume that individuals whose observed value of the eligibility variable is τ are in fact a *mixture* of individuals whose true value, S^* say, is $S^* = \tau$ and individuals whose reported value $S = \tau$ is affected by measurement error

$$S_{obs} = S^* Z + S(1 - Z),$$

where Z denotes a dummy taking value one for the exact reporters and $S \neq S^*$ because of measurement error. This is known as the *contaminated sampling model* discussed, amongst others, by Horowitz and Manski (1995). Third, we show that if the mixture groups indexed by Z are *not* systematically different with respect to (Y, R, S^*), the sample analogues of the quantities in (5) would define a biased estimator for the parameter of interest.[4] Finally, we show that the sample analogue of the fuzzy regression discontinuity estimand is in fact consistent for the parameter defined in (5). This results will be heavily used in the estimation section below, as it implies that consistent estimates of the causal effect of retirement on consumption can be recovered by a simple instrumental variable strategy where the eligibility status is used to solve for the endogeneity of the retirement status (see Imbens and Angrist, 1994, and Hanh *et al.*, 2001).

4.1 A formal setup

From the definition of S_{obs} given above and by using the law of iterated expectations we have

$$\begin{aligned} E\{R|S_{obs} = s_{obs}\} &= E\{R|S^* = s_{obs}, Z = 1\}E\{Z|S_{obs} = s_{obs}\} \\ &+ E\{R|S = s_{obs}, Z = 0\}(1 - E\{Z|S_{obs} = s_{obs}\}) \\ &= E\{R|S^* = s_{obs}, Z = 1\}E\{Z|S_{obs} = s_{obs}\} \\ &+ (1 - E\{Z|S_{obs} = s_{obs}\}) \\ &\int E\{R|S = s_{obs}, S^* = \tau, Z = 0\} f_{S^*|S,Z}(\tau|s_{obs}, 0)d\tau. \end{aligned}$$

Under the assumption $Z \perp (Y, R, S^*, S)$ which states the irrelevance of Z for the sampling process, and the assumption of non-differential measurement error, $(Y, R) \perp S|S^*$, the last expression becomes

$$\begin{aligned} E\{R|S_{obs} = s_{obs}\} &= E\{R|S^* = s_{obs}\}E\{Z|S_{obs} = s_{obs}\} \\ &+ (1 - E\{Z|S_{obs} = s_{obs}\}) \\ &\int_{\bar{s}}^{+\infty} E\{R|S^* = \tau\} f_{S^*|S}(\tau|s_{obs})d\tau, \end{aligned}$$

as $E\{R|S^* = \tau\} = 0$ when $\tau < \bar{s}$. Note that, in general, the measurement error in S does not need to be classical, though it has to be non-differential, i.e. it must contain no information on (Y, R) once the true value S^* has been controlled for (see Bound *et al.*, 2001).

[4] In future research we aim at relaxing this assumption by allowing for no-zero correlation between Z and (Y, R, S^*).

Under smoothness conditions of the distribution of (S, S^*) around $(\bar{s}, \bar{s})$, it follows from the last expression that

$$E\{R|S_{obs} = \bar{s}^+\} - E\{R|S_{obs} = \bar{s}^-\} = E\{R|S^* = \bar{s}^+\}E\{Z|S_{obs} = \bar{s}^-\}, \tag{6}$$

implying that the discontinuity in the retirement probability observed around the threshold for eligibility understates the true discontinuity by means of the term $E\{Z|S_{obs} = \bar{s}^-\}$. It therefore follows that the estimated discontinuity is downward biased for the true discontinuity. The bias term can be estimated from the proportion of heads who self-report being retired though marginally ineligible according to S_{obs}, $P\{R = 1|S_{obs} = \bar{s}^-\}$.

By applying a similar argument to the regression function of Y on S_{obs} we obtain

$$\begin{aligned} E\{Y|S_{obs} = \bar{s}^+\} - E\{Y|S_{obs} = \bar{s}^-\} \quad = \quad & (E\{Y|S^* = \bar{s}^+\} - E\{Y|S^* = \bar{s}^-\}) \\ & E\{Z|S_{obs} = \bar{s}^-\}, \end{aligned} \tag{7}$$

which implies that the discontinuity in consumption expenditures estimated around $S_{obs} = \bar{s}$ is still downward biased for the true discontinuity.

4.2 Identification

In the notation of the previous section, the parameter in (5) can be written as

$$E\{\beta|R = 1, S^* = \bar{s}^+\} = \frac{E\{Y|S^* = \bar{s}^+\} - E\{Y|S^* = \bar{s}^-\}}{E\{R|S^* = \bar{s}^+\}},$$

which depends on the joint distribution of (Y, R, S^*). Because of measurement error in S_{obs}, (6) and (7) imply that the estimator constructed by using the empirical analogues of the quantities in the last expression from raw data (Y, R, S) is *not* consistent for the parameter of interest.

However, it is immediate to see how the following estimator

$$\frac{E\{Y|S_{obs} = \bar{s}^+\} - E\{Y|S_{obs} = \bar{s}^-\}}{E\{R|S_{obs} = \bar{s}^+\} - E\{R|S_{obs} = \bar{s}^-\}}, \tag{8}$$

is consistent for the causal effect of retirement on consumption under the assumptions made on the measurement error in S_{obs}. As pointed out by Imbens and Angrist (1994), the latter expression can be interpreted as an instrumental variable estimand, where the eligibility status is used to correct for the endogeneity of R.

5 Empirical analysis

5.1 Estimation issues

The causal effects of retirement on consumption can be estimated along the lines of what discussed in the last section. The sample analogue of (8) can be obtained by taking the ratio of the discontinuity pictured in the top panel of Figure 4 to the discontinuity in the bottom panel of the same figure.

Though the effect of interest can be non-parametrically identified, the analysis presented in what follows builds on a fully parametric approach, which - we think - represents a better framework to use while communicating our results. As the sample analogue of (8) coincides with an instrumental variable estimator where the endogenous variable R is instrumented with the eligibility status $\mathbb{1}(S_{obs} \geq \bar{s})$ (see Imbens and Angrist, 1994), the regressions presented in the next section will all take the following form

$$Y = \beta_0 + \beta_1 R + \beta_2 S_{obs} + \beta_3 S_{obs}^2 + \varepsilon,$$

where the first-stage regression is run on $\mathbb{1}(S_{obs} \geq \bar{s})$.

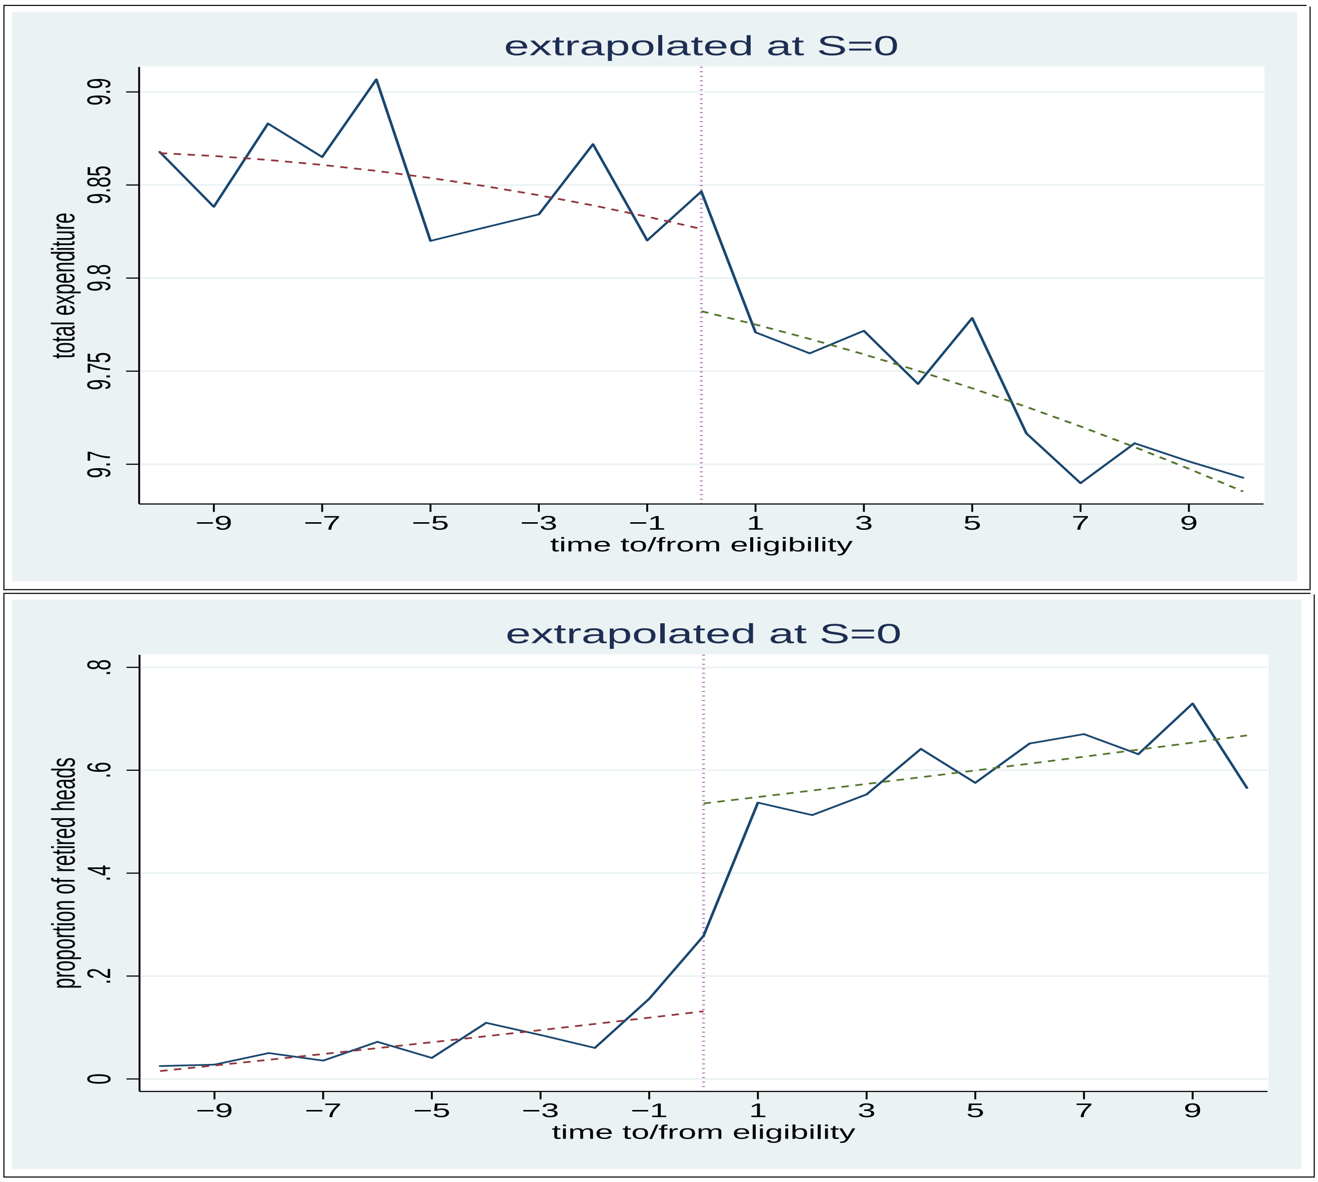

Figure 4: Non-parametric estimation of the causal effect of retirement on non-durable consumption expenditures

Table 4: Estimation results

Non-durable expenditure				
	coefficient	**std. err.**	**t**	**p-value**
Male retired	-0.09217	0.04111	-2.24	0.039
S_{obs} male	-0.00485	0.00145	-3.35	0.004
S_{obs}^2 male	-0.00026	0.00008	-2.96	0.009
Food expenditure				
	coefficient	**std. err.**	**t**	**p-value**
Male retired	-0.06882	0.05108	-1.35	0.197
S_{obs} male	-0.00480	0.00191	-2.51	0.023
S_{obs}^2 male	-0.00024	0.00001	-2.01	0.062

5.2 Results

The estimation procedure that we take can be described as follows. First, averages of household non-durable expenditure and proportions of retired male heads by values of S_{obs} (between -10 and 10) are considered, that is the twenty observations corresponding to the top and the bottom panels of Figure 4, respectively. Second, we regress averages of household non-durable expenditure on proportions of retired heads and a quadratic polynomial in S_{obs}, restricting the sample to "husband and wife" and "single males" household only (the inclusion of additional regressors such as time dummies leads to qualitatively similar results, so we have decided to report the simplest specification). Implicit in this specification is the assumption that the leisure of the spouse is separable from that of the male head. Retirement is instrumented with the eligibility status, yielding good results for the first stage. Clustering is explicitly taken into account in estimation, by computing a weighted IV estimator with weights reflecting the cell size.

Two sets of results are presented: for non-durable consumption and for food at home consumption. Results for non-durable consumption are reported in the top panel of Table 4, suggesting a consumption drop of around 9.2 percent, that is significantly different from zero at the 95 percent level. Results for food at home are reported in the bottom panel, suggesting a smaller consumption drop (6.9 percent), that is not significantly different from zero. We also estimated these equations using per-equivalent adult consumption measures, with almost identical results.

6 Some back of the envelope stuff

We have found that consumption drops at retirement by a relatively large amount: 9.2%. The question is whether this drop should be taken as evidence against life-time optimizing behaviour.

A number of papers have emphasized that a consumption drop at retirement is compatible with the life-cycle model if leisure affects the marginal utility of consumption. Specific ways in which this non-separability may come into play at retirement are that work-related expenditures are no longer needed (Banks, Blundell and Tanner, 1998, emphasize this channel, but their estimate of the consumption drop is much smaller, around 3%, two thirds of which is anticipated) or some goods and services may start being produced at home (as argued in Aguiar and Hurst, 2004, and Hurd and Rohwedder, 2003, who report anticipated drops in the $15-20\%$ region).

A simple way to find out whether our estimated consumption drop is consistent with the life cycle model is to perform some back-of-the-envelope calculations. Let us take the simplest possible utility function, a power utility defined over a Cobb-Douglas composite good made of non-durable consumption, C, and (male) leisure, l:

$$U_t = \frac{\left(C_t^{\alpha} l_t^{1-\alpha}\right)^{1-\gamma}}{1-\gamma}$$

where $\gamma > 0$ is the reciprocal of the elasticity of intertemporal substitution (EIS) and α measures the within period consumption share for periods when leisure is not at a corner (hence $0 < \alpha < 1$):

$$\alpha = \frac{p_t C_t}{p_t C_t + w_t l_t}$$

where p_t is the price of the consumption good at time t and w_t the nominal wage at time t.

Life-time optimization implies that households should keep the marginal utility of consumption, U_c, constant:

$$U_c = \alpha C^{\alpha-1-\alpha\gamma} l^{(1-\alpha)(1-\gamma)}$$

This implies that consumption must react to changes in leisure according to the following relation:

$$\frac{dC}{dl}|_{U_c=\overline{U}} = \frac{C(1-\alpha)(1-\gamma)}{l(1-\alpha+\alpha\gamma)}.$$

We see that consumption should drop as a result of an increase in leisure if $\gamma > 1$, it should increase otherwise. A great attraction of the Cobb-Douglas lies in its analytical tractability: the formula above can be used to derive an elasticity as a function of just two preference parameters, α and γ.

We can estimate α in our data as follows. We take a sample of prime-age workers and assume that l is defined as $(T-h)$, where T is the maximum number of hours a person could possibly work (16 a day, say) and h is hours of work. In the case where $h = 8$, for instance, $l = 8$, and $wl = wh =$ earnings. The average propensity to consume out of earnings is roughly 80%, so $\alpha = 0.44$.No estimate for the elasticity of intertemporal substitution exists in Italy (to our knowledge), but studies on micro data run in other countries suggest that the elasticity should lie in the $0.5 - 0.8$ range. The real business cycle literature as well as recent studies on aggregate consumption and asset returns (Bansal and Yaron, 2004) emphasize the elasticity should be in excess of unity. We shall therefore consider a range for γ between 0.5 and 1.5.

We have estimated the average consumption drop associated to retirement to be 9.2%, with a standard error of 4.1%. In our calculations of the α parameter, retirement doubles leisure. Thus the estimated elasticity is indeed -0.092.

A 0.092 consumption drop is consistent with utility maximization if $\gamma = 1.1775$ ($EIS = 0.85$). If we consider a one-standard error range around the point estimate, we find that a 0.133 drop is consistent with $\gamma = 1.265$ ($EIS = 0.79$) and a 0.051 drop is consistent with $\gamma = 1.095$ ($EIS = 0.91$). In this context, a zero drop obtains if $\gamma = 1$ (EIS =1), whereas consumption should *increase* at retirement for smaller values of γ.

This simple, one-good case hides the fact that the impact of retirement on total (non-durable) consumption is ambiguous, because some goods may be leisure substitutes and some other leisure complements. MMW show examples of both, notably food out and transport as substitutes, food at home and heating fuel as complements.

We do not have detailed expenditure information in our main data set, SHIW, but can access diary-level data on consumer spending for the year 2002. This large data set, collected by the Italian statistical office (ISTAT), contains records of current employment, household composition, size of the main residence and a few other household characteristics, but no information on years of contributions, or past employment histories. Thus years to and from pension eligibility are not known, and our identification strategy cannot be applied.

What we can do is to compare two groups of households, those whose head's age lies between 50 and 54 and 65 to 69. Heads of household in the former group are mostly employed (81.8% are

Table 5: Consumption drop in SFB diary data

	unadjusted		p-score adjusted	
Category	**difference**	s.e.	**difference**	s.e.
Non-durable	-510.35	22.70	-241.02	29.74
Food at home	-89.70	6.74	-5.96	8.90
Meals out	-44.37	2.95	-35.68	3.46
Alcohol	-5.24	0.96	-2.75	1.13
Tobacco	-13.35	0.95	-7.98	1.26
Clothing	-99.09	7.26	-58.05	9.28
Personal services	-7.96	2.73	-5.29	3.01
Transport	-153.67	7.56	-76.04	9.91
Heating	-19.75	2.56	-8.73	3.19
Phones	-15.95	0.96	-8.99	1.14
Housing services	-7.82	4.34	-2.55	4.47
Other	-53.43	4.02	-28.99	4.93

employed, 9.6% are retired, the others are either unemployed or out of the labour force), in the latter they are mostly retired (82.7% are retired, 8.0% are employed, all the others are out of the labour force).

In Table 5 we report the difference between average spending of the older group and average spending of the younger group, and its standard error. The first column lists the various commodities considered, the second and third columns present a straight comparison, whereas the fourth and fifth columns refer to a comparison that corrects for composition effects in terms of region of residence, number of equivalent adults and size of the main residence.[5] This correction is meant to remove that part of difference that can be attributed to age-related changes in household composition as well as (to some extent) cohort effects, under the assumption that the size of the main residence for a given household size correlates with life-time wealth.

The first row of numbers tells us that non-durable consumption falls by 510 euros a month (-31.1%) between the early fifties and late sixties. However, once composition effects are taken into account, this drop is reduced to 241 euros a month (-15.6%). This is larger than our estimate of the consumption drop at retirement, suggesting that age and composition effects play a role that is not fully accounted for in our adjustment procedure. One possible interpretation of our estimates is that almost 60% of the overall consumption drop over this period of the life cycle is due to retirement, the remaining 40% reflects changes in preferences due to poorer health and other unobservable age-related characteristics.

If we look at the adjusted drops, we see that the largest items are meals out, clothing and transport, that account together for 169.77 euros, that is for over two thirds of the overall drop. At least two of these items, clothing and transport, are typically considered work-related expenses, whereas the drop in meals out is consistent with the home production hypothesis. In fact, these is a switch from meals out to food at home in relative terms: food at home stays constant once composition effects are taken into account (and increases its budget share considerably), whereas meals out are reduced by 35.68 euros - around 41% in relative terms! - and so does its budget share.

The conclusion that we draw from this exercise is that the our estimated retirement consumption drop could well be due entirely to a reduction of work-related expenses and a substitution away from market goods to home-production of food.

[5]The adjustment was made so to make the distribution of these characteristics for the younger and the older groups be equal to the distribution of households whose head is aged between 59 and 64.

7 Conclusions

In this paper we have investigated the size of the consumption drop due to retirement in Italy. We have used micro data covering the 1991-2002 period on food, non-durable and total household spending, and evaluated the change in consumption that accompanies retirement by exploiting the exogenous variability in pension eligibility to correct for the endogenous nature of the retirement decision. We have taken a regression discontinuity design approach, and made the identifying assumption that consumption would be the same around the threshold for pension eligibility, if the individual could not retire. We have shown that a non-negligible fraction of individuals retire as soon as they become eligible, and estimated the part of the consumption drop that is associated with retirement induced by eligibility. Given that pension eligibility is a variable that is measured with error, we have also evaluated the impact of measurement error on our estimates.

Our key result is that non-durable consumption drops by 9.2% because of (male) retirement. We have shown that such fall is compatible with a constant marginal utility of consumption if preferences are non-separable between consumption and leisure. In particular, if preferences are Cobb-Douglas in consumption and male leisure, we have shown that we would require the elasticity of intertemporal substitution to be between 0.8 and 0.9, in line with recent micro evidence for the US and the UK. We have also produced evidence that this drop can be mostly accounted for by falls in goods that are typicaly considered to be work-related expenses or leisure substitutes (clothing, transport, meals out).

References

[1] Aguiar Mark and Erik Hurst (2004), "Consumption vs. Expenditure", NBER WP 10307

[2] Angrist, Joshua D., and Guido W. Imbens (1994), "Identification and Estimation of Local Average Treatment Effect*s"*, *Econometrica*, **62(2)**, 467-476

[3] Banks, James, Richard Blundell and Sarah Tanner (1998) "Is There a Retirement-Savings Puzzle?", *American Economic Review,* **88(4)**, 769-788

[4] Battistin, Erich and Andrew Chesher (2004), "The Impact of Measurement Error on Evaluation Methods based on Strong Ignorability", mimeo, Institute for Fiscal Studies, London

[5] Battistin, Erich, Raffaele Miniaci and Guglielmo Weber (2003) "What Do We Learn from Recall Consumption Data?", *Journal of Human Resources*, **38(2)**, 354-385

[6] Battistin, Erich and Enrico Rettore (2005), "Ineligibles and Eligible Non-Participants as a Double Comparison Group in Regression Discontinuity Designs", *Journal of Econometrics*, forthcoming

[7] Battistin, Erich and Barbara Sianesi (2005), "Misreported Schooling and Returns to Education: Evidence from the UK", mimeo, Institute for Fiscal Studies, London

[8] Bernheim, B. Douglas, Jonathan Skinner and Steven Weinberg (2001) "What Accounts for the Variation in Retirement Wealth Among U.S. Households?", *American Economic Review*, **91(4)**, 832-857

[9] Boeri, Tito and Agar Brugiavini (2005), "Pension Reforms and Women at Work", paper presented at IZA conference on *The Older Worker in the Labor Market: Work Time and Work Incentives*, Lisbon.

[10] Bound, John, Brown, Charlie and Nancy A. Mathiowetz, N. (2001), "Measurement error in survey data", in J.J. Heckman and E. Leamer (eds.), *Handbook of Econometrics. Vol. 5*, Amsterdam: North-Holland, 3705-3843

[11] Browning, Martin and Anna Maria Lusardi (1996) "Household Saving: Micro Theories and Micro Facts", *Journal of Economic Literature*, **34(4)**, 1797-1855

[12] Deaton, Angus (1992) *Understanding Consumption*, Oxford: Oxford University Press

[13] Hahn, Jinyong, Petra Todd. and Wilbert van der Klaauw(2001), "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design", *Econometrica*, **69(1)**, 201-209

[14] Hamermesh, Daniel S. (1984) "Consumption During Retirement: The Missing Link in the Life-Cycle", *Review of Economics and Statistics*, **66(1)**, 1-7.

[15] Heckman, James J., Robert Lalonde and Jeffrey Smith (1999), "The Economics and Econometrics of Active Labor Market Programs", *Handbook of Labor Economics*, Volume 3, Ashenfelter, A. and Card, D. (eds.), Amsterdam: Elsevier Science

[16] Horowitz, Joel L. and Charles F. Manski (1995), "Identification and Robustness with Contaminated and Corrupted Data", *Econometrica*, **63(2)**, 281-302

[17] Hurd, Michael and Susann Rohwedder (2003), "The Retirement-Consumption Puzzle: Anticipated and Actual Declines in Spending at Retirement", NBER WP 9586

[18] Johnson, Norman L., Samuel Kotz and N. Balakrishnan (1994), *Continuous Univariate Distributions*, New York: Wiley

[19] van der Klaauw, Wilbert (2002), "Estimating the Effect of Financial Aid Offers on College Enrollment: A Regression-Discontinuity Approach", *International Economic Review*, **43(4),** 1249-1287

[20] Miniaci, Raffaele, Chiara Monfardini and Guglielmo Weber (2005) "Is There a Retirement Consumption Puzzle in Italy?", mimeo, University of Padua

[21] Thistlethwaite, Donald L. and Donald T. Campbell (1960), "Regression discontinuity analysis: an alternative to the ex post facto experiment", *Journal of Educational Psychology*, **51(6)**, 309-317

[22] Trochim, William M. K. (1984), *Research Design for Program Evaluation: the Regression-Discontinuity Approach*, Beverly Hills: Sage Publications