

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 31st March 2007

Abstract

In this paper we investigate the size of the consumption drop at retirement in Italy. We use micro data on food and total non-durable household spending covering the period 1993-2004, 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 9.8% the part of the non-durable consumption drop that is associated with retirement induced by eligibility. We show that such fall is not driven by liquidity problems for the less well off in the population, and is compatible with a constant marginal utility of consumption if preferences are non-separable between consumption and leisure. We also produce evidence that this drop can be mostly accounted for by falls in goods that are typically considered to be work-related expenses or leisure substitutes.

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

*Paper presented at the RTN meeting on the Economics of Ageing, Venice, October 2005, the EALE 2006 conference, 2nd CSEF-IGIER Summer School, 2006 NBER Summer Institute Workshop on Aggregate Implications of Microeconomic Consumption Behavior, WPEG 2006, ESEM06, and at seminars at EUI, IGIER and Ente Luigi Einaudi, Rome. We thank conference and seminar participants for their helpful comments, and particularly Orazio Attanasio, Mike Hurd and Stephen Zeldes. 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 approaching or 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, 2003, 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. 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 on average 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 diary-level data covering the 1985-96 period. MMW estimate a fall in non-durable consumption at retirement of 5.4%. However, the data MMW use does not contain any retrospective information on work-histories, and MMW 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 pension eligibility to identify the causal effect in a regression

2 Literature Review

The life-cycle model of Modigliani and Brumberg (1954) predicts that individuals save to smooth consumption over time: in its simplest version, they save during their working lives to keep their consumption level constant once they retire. Hamermesh (1984) was the first paper to argue that consumers apparently do not save enough to achieve this aim. If households enter retirement with inadequate savings, they must cut their consumption level, contrary to the life-cycle model predictions.

The recent literature has focused on estimating how consumption levels change around retirement. The existence of a consumption fall around retirement is documented for the UK (Banks, Blundell and Tanner, 1998, BBT in what follows), for the US (Bernheim, Skinner and Weinberg, 2001, BSW), and for Italy (Miniaci, Monfardini and Weber, 2003, MMW) and is known as the retirement consumption puzzle (or retirement savings puzzle).

BBT use British cohort data and show that the standard Euler equation, in which consumption growth is a function of intertemporal prices and changes in demographics, underpredicts the level of consumption by as much as 1.5% on an annual basis for ages between 60 and 67. The cumulated consumption shortfall over this age band, where most people retire, is around 10%. BBT argue that only a fraction of this drop can be attributed to the increased leisure time that accompanies retirement. Later work by Smith (2006) uses food information on food for UK households who retired over the sample period, and stresses the importance of distinguishing between voluntary and involuntary retirement: a significant drop for food consumption is observed only for those who retire early because of poor health or job loss.

BSW use PSID data to estimate Euler equations for food consumption. Retirement status is instrumented by taking age-specific predicted probabilities conditional on demographics. The sample is split in groups: low wealth-to-income households drop their consumption most. BSW estimate a median drop of 14%, but higher drops for low wealth, low income replacement households. BSW conclude that "31% of the sample reduce their consumption by at least 35 percentage points". The evidence they provide is consistent with the notion that consumers do indeed enter retirement with inadequate savings. A number of papers have further investigate the issue on US data - Haider and Stephens (2004), who estimate a smaller consumption drop for those who retire at the expected time, Fisher et al (2005), who use CEX data and estimate a smaller drop (around 2.5%) for total expenditure than for food consumption (around 5.7%).

Recent papers by Aguiar (2005a and 2005b) and Hurd and Rohwedder (2006) stress that the drop in expenditure at retirement does not necessarily imply a drop in utility. For instance, work-related expenditure (transport to and from work, canteen meals and business clothing) is no longer needed - whether the account for a large enough part of pre-retirement consumption is an open issue. Also, home production of services (laundry, gardening, house-cleaning, cooking) may become advantageous, and the extra leisure time may allow consumers to shop more efficiently. This last channel has recently been stressed by Aguiar and Hurst, (2005a) and (2005b), in their careful analysis of food consumption around retirement, whilst the increase in home production of services by recent retirees has been documented by Hurd and Rohwedder (2006), who exploit time-use data.

Finally, the evidence provided in MMW is based on Italian diary-level data covering the 1985-96 period. MMW estimate a fall in non-durable consumption at retirement of 5.4%. MMW emphasize that this is a lower bound if there is heterogeneity in work-related expenses and those with higher expenses retire earlier. Also, the data MMW use does not contain any retrospective information on work-histories, and MMW 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.

3 Identification

3.1 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, $\tau = Y_1 - Y_0$, which is not observable at the household level since being retired reveals Y_1 but conceals Y_0 . Accordingly, though not observable τ 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 $\$$ is the discontinuity point, then a regression discontinuity is defined if

$$Pr\{R = 1|\$^+\} \neq Pr\{R = 1|\$^-\}: \tag{1}$$

Here and in the following $\$^+$ and $\$^-$ refer to those individuals *marginally* above and below $\$$, 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 an

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

(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 ξ . Such assumptions qualify S as an instrumental variable for R around ξ , 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.

3.2 Discontinuities and pension eligibility

In the context of this paper, however, self selection of households into retirement turns 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 ξ is zero by definition, since they are not eligible for retirement. The probability of retirement for those scoring above ξ 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 ineligible ($R(\xi^-) = 0$), the expression in (4) can be written as

$$E\{Y_0|\xi^+\} - E\{Y_0|\xi^-\} + E\{R(s) \mid \xi^+\}:$$

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

$$E\{R(s) \mid \xi^+\} = E\{\mathbb{1}\{R = 1; \xi^+\}\}Pr\{R = 1|\xi^+\};$$

so that the mean impact on participants in a right-neighborhood of ξ is identified by

$$E\{\mathbb{1}\{R = 1; \xi^+\}\} = \frac{E\{Y|\xi^+\} - E\{Y|\xi^-\}}{E\{R|\xi^+\}}. \quad (5)$$

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 ξ .

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 ξ 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 .

4 Data

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).

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

³Results 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

	1993	1995	1998	2000	2002	2004
Males						
Worker	66.8%	61.0%	60.1%	58.5%	54.9%	54.8%
Retired	28.4%	32.1%	31.8%	35.3%	39.0%	40.0 %
Females						
Worker	42.9%	43.4%	44.3%	42.9%	41.9%	43.8%
Retired	21.5%	26.3%	22.3%	24.4%	29.1%	29.7%

The survey contains a few recall questions on consumption. For instance, in the 1995 survey (collected in 1996) respondents were first asked "What was the monthly average spending of your household in 1995 on all consumer goods?" - they were instructed to exclude mortgage payments, rent, major house renovation as well as purchases of listed consumer durables (cars, furniture, appliances, jewels etc.). They were then asked "What instead is the monthly average figure for just food consumption? Consider spending on food products in supermarkets and the like and the spending on meals eaten regularly outside the home". Finally, they were asked questions on purchases and sales of consumer durables over the whole year. The same set of questions was asked in the 1998 wave.

In the 2000, 2002 and 2004 waves, the questions on durable purchases and sales were asked first, followed by the non-durable consumption and food questions, but the wording and contents was otherwise identical.

The 1993 survey, instead, asked for a different definition of food consumption: respondents were not instructed to include meals regularly consumed out of the home, but to consider expenditure at grocery stores and subtract home and personal cleaning products.⁴

Retirement status is defined on the basis of two questions. In each survey, respondents were asked whether each household member was employed for the most part of the year. If the answer was negative, they were then asked whether the household member was first-time work seeker, unemployed, home-maker, job pensioner, non-job pensioner, student, conscript or other. Non-job pensions were defined as disability, survivor and social pensions. In this paper, we consider a person as retired if she/he is classified as job-pensioner, this distinction contributes to the result that the sum of the percentages of workers and pensioners is in some years much below one hundred (see Table 1).

4.1 The definition of pension eligibility

The aim of this 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 2004), 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 - the residual category includes other conditions such as disable or homemaker. In the table, and in the rest of our analysis, we have considered only heads of households and in case their spouses/partners. 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.

⁴Battistin, Miniaci and Weber (2003) compare consumption data across SHIW and diary-based SFB for one particular year, 1995.

4.1.2 The measurement of years of contributions

In some years contributions have been explicitly recorded in SHIW for workers and 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 cases. First, if a worker or a retiree is interviewed *only once* in 1993, then we have no choice then to drop the observation. This halves the sample size of 1993. Second, if the individual is observed in 1993 and he belongs to the *panel* sub-sample of SHIW, an imputation of years of contribution is made on the basis of the 1995 recorded figure on contributions. Third, for a *retired* person who is observed only once between 1995 and 2004 and has a missing value for the contribution years an imputation is made on the basis of retrospective questions on the age of retirement and the age of entry into the labour market. If the information is missing for a *worker* instead (and he is observed only once between 1995 and 2004) the imputation is made on the basis of the difference between current age and age of entry into the labour market. Fourth, if a retiree or a worker belongs to the panel, missing values of contributions can be recovered from the previous or subsequent waves.

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 the number of cases where an imputation had to be carried out is not large (apart from the year 1993). Figure 1 shows the distribution of years of contributions separately for men and women. For men there is high percentage of cases with contribution spells above 30 years - with a relevant spike at 35 years. For women the distributions is much more dispersed and in many cases there are only 5 or 10 years of contributions completed.

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 2. 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. While for men the distribution of S is rather even over the range of negative and positive values, for women there is a prevalence of positive values of S , indicating that a large percentage of women have past their first eligibility year (and they are presumably retired).

4.1.3 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 ; for brevity we only show these percentages for values of S between -5 and $+5$: 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 11 percent of males and 25 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 3 no longer applies. More precisely, estimators of the causal effect of retirement on consumption




Figure 1: Contributive years - waves 1993 to 2004


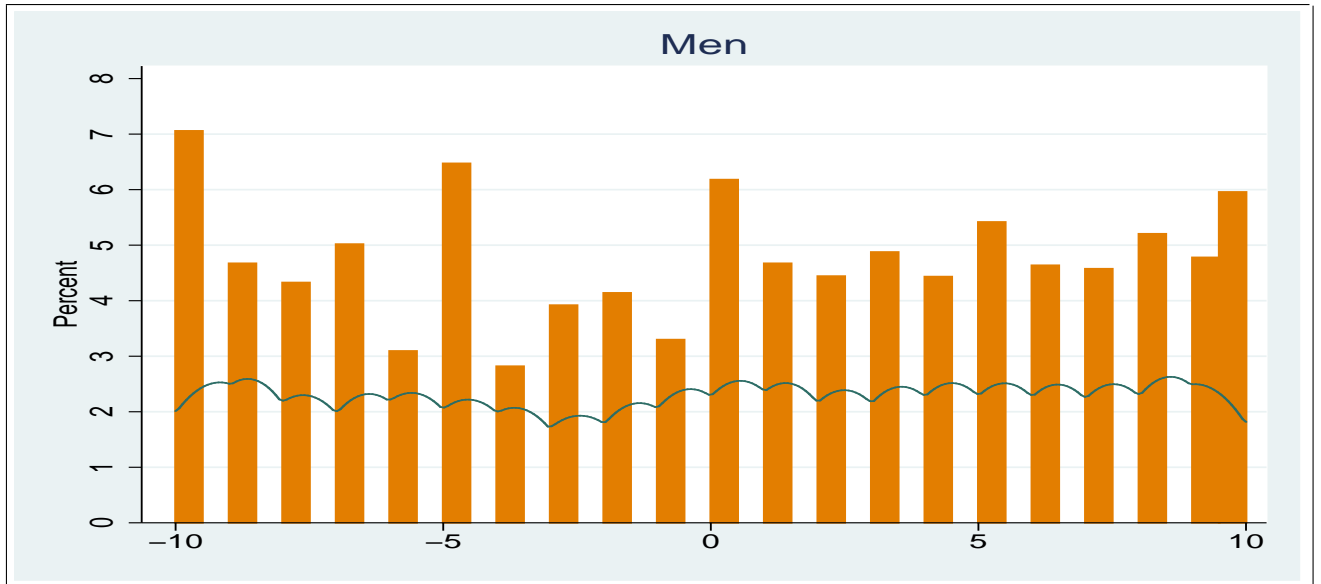


Figure 2: Distribution of S

Table 3: Percentage of retired individuals by eligibility status

eligibility	males	females
-5	0.045	0.013
-4	0.013	0.023
-3	0.015	0.024
-2	0.015	0.035
-1	0.025	0.140
0	0.333	0.568
1	0.626	0.704
2	0.613	0.800
3	0.652	0.627
4	0.740	0.781
5	0.666	0.605

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 an estimation strategy which - under certain assumptions on the data generating process of the errors - provides consistent estimates of the causal parameter.

4.2 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.

5 Measurement errors

Throughout this section we will allow for measurement error in the variable S , but we will maintain the assumption that R is 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 4.1, 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 4.1 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 ζ are in fact a *mixture* of individuals whose true value, S^* say, is $S^* = \zeta$ and individuals whose reported value $S = \zeta$ 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.⁵ 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).

5.1 A formal setup

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

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

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

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

as $E\{R|S^* = \zeta\} = 0$ when $\zeta < \bar{s}$. 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).

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

$$E\{R|S_{obs} = \bar{s}^+\} - E\{R|S_{obs} = \bar{s}^-\} = E\{R|S^* = \bar{s}^+\}E\{Z|S_{obs} = \bar{s}^-\}; \quad (6)$$

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} , $P\{R = 1|S_{obs} = \bar{s}^-\}$.

⁵In future research we aim at relaxing this assumption by allowing for non-zero correlation between Z and (Y, R, S^*) .

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

$$E\{Y|S_{obs} = \mathfrak{s}^+\} - E\{Y|S_{obs} = \mathfrak{s}^-\} = \frac{(E\{Y|S^* = \mathfrak{s}^+\} - E\{Y|S^* = \mathfrak{s}^-\})}{E\{Z|S_{obs} = \mathfrak{s}^-\}}; \quad (7)$$

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

5.2 Identification

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

$$E\{R|S^* = \mathfrak{s}^+\} = \frac{E\{Y|S^* = \mathfrak{s}^+\} - E\{Y|S^* = \mathfrak{s}^-\}}{E\{R|S^* = \mathfrak{s}^+\}};$$

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

$$\frac{E\{Y|S_{obs} = \mathfrak{s}^+\} - E\{Y|S_{obs} = \mathfrak{s}^-\}}{E\{R|S_{obs} = \mathfrak{s}^+\} - E\{R|S_{obs} = \mathfrak{s}^-\}}; \quad (8)$$

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 .

6 Empirical analysis

6.1 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 3 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 $\mathbb{1}(S_{obs} \geq \mathfrak{s})$ (see Imbens and Angrist, 1994), the regressions presented in the next section will all take the following form:

$$Y_{s,t} = \alpha_0 t + \alpha_1 R_{s,t} + \alpha_2 S_{obs,t} + \alpha_3 S_{obs,t}^2 + \epsilon_{s,t}$$

where the first-stage regression is run on $\mathbb{1}(S_{obs} \geq \mathfrak{s})$. All variables are indexed by t and s to emphasize that we take sample averages by calendar year, as well as by years to/from eligibility. We allow for year-specific intercepts to take into account the changes in the wording of consumption questions that are discussed above.

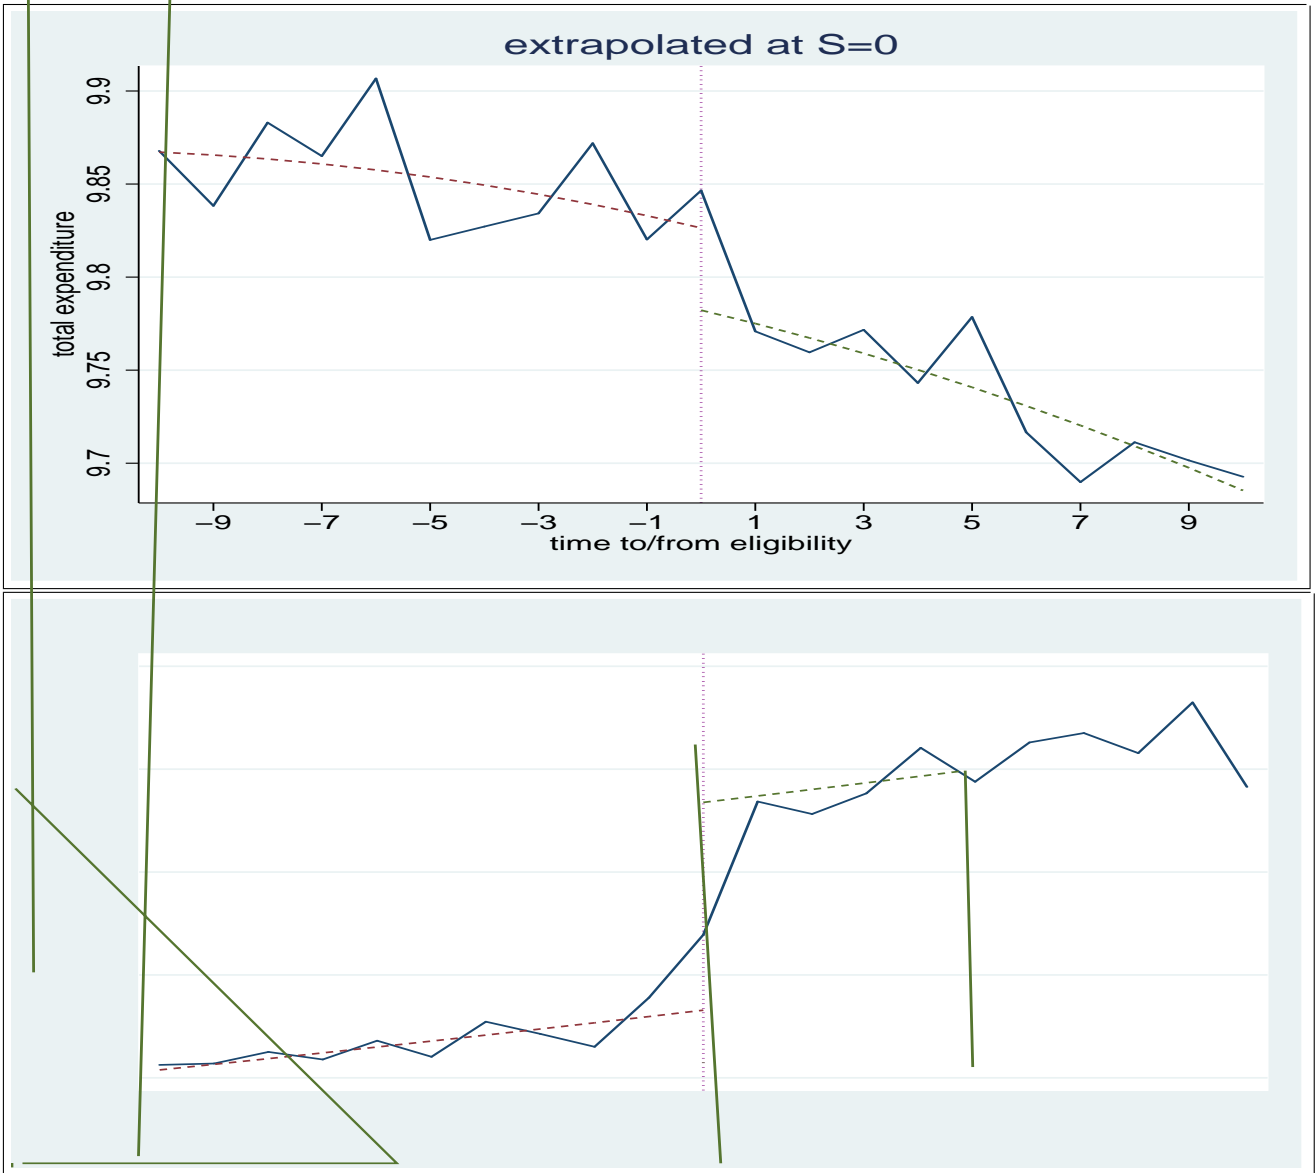


Figure 3: Non-parametric estimation of the causal effect of retirement on non-durable consumption expenditures

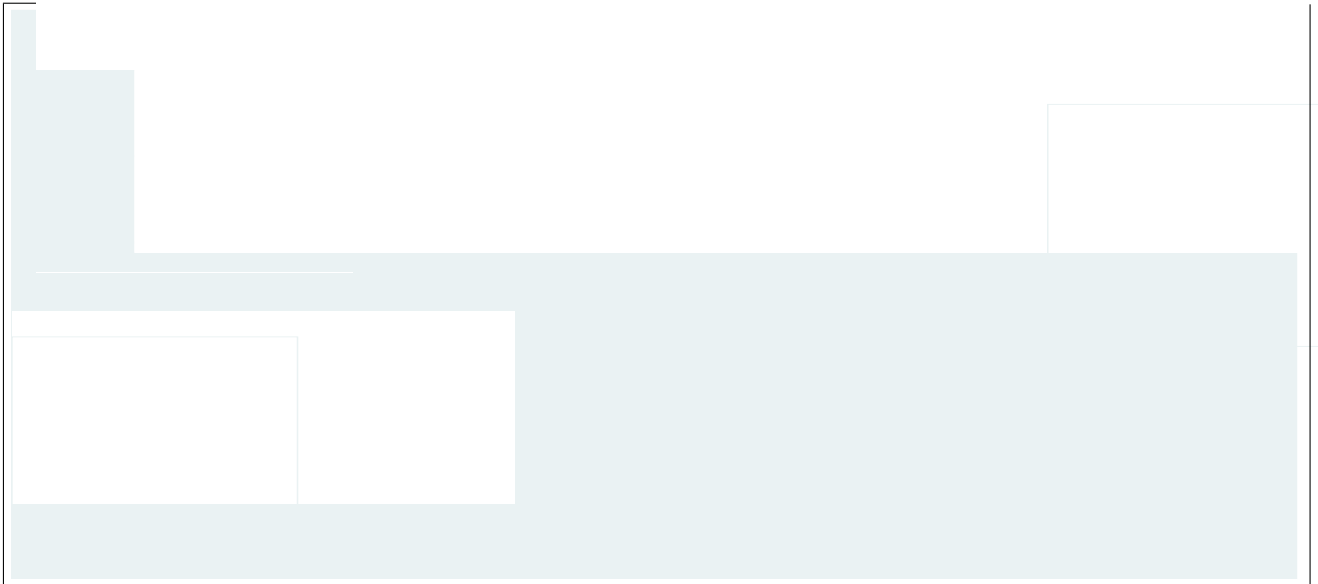


Figure 4: First stage - by year

6.2 Results

The estimation procedure that we take can be described as follows. First, for each survey year, we compute averages of household non-durable expenditure and proportions of retired male heads by values of S_{obs} (between -10 and 10), that is the twenty observations corresponding to the top and the bottom panels of Figure 3, 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. In all specifications we add time dummies as regressors. 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 the computation of the standard errors.

The first stage regression of retirement status on pension eligibility, a second order polynomial in S_{obs} , and time dummies, presents an R^2 of .92 - the coefficient on eligibility is estimated at .435, with a standard error of .038. The relation between retirement status and eligibility is relatively stable over the years, as graphically shown in Figure 4.

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 9.8 percent, that is significantly different from zero at the 90 percent level. Results for food are reported in the bottom panel, suggesting a larger consumption drop (14.1 percent), that is significantly different from zero at the 95 per cent level.

The lack of precision in the estimation of the non-durable consumption drop likely reflects the high noise to signal ratio to be expected in this type of recall question (see Browning, Crossley and Weber, 2003, for an appraisal). Recall questions on food consumption are less heavily affected by memory problems, and therefore more informative (Battistin, Miniaci and Weber, 2003). But if we take the point estimates at face value, our finding that food expenditure falls more at retirement than non-durable expenditure may appear at odds with expectations. However, for all years but one food consumption is defined to include meals regularly consumed out of the home, and this is a typical example of work-related expenses that are likely to vanish when an individual retires.

6.3 Robustness Checks

Table 4: Estimation results

Non-durable expenditure				
	coef	std. err.	t	p-value
Male retired	-0.0983	0.0567	-1.74	0.085
S_{obs} male	-0.0055	0.0027	-2.05	0.043
S_{obs}^2 male	-0.00029	0.00015	-1.91	0.059
Food expenditure				
	coef	std. err.	t	p-value
Male retired	-0.1409	0.05442	-2.59	0.011
S_{obs} male	-0.0028	0.0026	-1.07	0.287
S_{obs}^2 male	-0.00008	0.00014	-0.58	0.561

To gather evidence on the validity of condition 1 on which our identification strategy relies, we implement an over-identification test following Lee (2006). Consider the set of *pre-intervention* outcomes that meet the following two conditions: they should not be affected by the eligibility status but they should be correlated to the unobservables likely to affect the level of consumption.

Consider the case of education. It is well known that education is a good proxy for life-time access to economic resources, and should therefore affect consumption. But it is clear that there is no room for a causal effect of the eligibility status on the education of males in a neighbourhoods of eligibility for retirement. As a consequence, upon finding that households on the two sides of the discontinuity point $S = 0$ differ with respect to education of their bread-winner, we would have to conclude that our identification strategy fails since households assigned to S^- are not comparable to households assigned to S^+ with respect to an observable known to be related to several unobservables relevant for the consumption pattern.

A formal test confirming this evidence is implemented by running the same parametric IV regression presented in the previous subsection using as a dependent variable a battery of *pre-intervention* outcomes. They include four education variables, the age of the head and the size of the main residence. The evidence is reported in the left columns of Table 5 and in all cases we consider it does not lead to rejecting our identifying restriction.

Due to the rules of the Italian pension system, an additional likely causal effect of the household head being eligible for retirement is a change in household composition. Employees receive a severance pay upon retirement, known as *Liquidazione*, that is proportional to the number of years on the job: for a person who retires after forty years with the same firm, this lump-sum pay is normally worth three years' gross annual salary (and is taxed at an 11% reduced rate). The anecdotal evidence is that as fathers (and mothers) retire they use the *Liquidazione* to buy a house for their sons and daughters which this way leave the household to set a new one (Guiso and Jappelli, 2002, document the role that inter-vivos transfers play in Italy on home-ownership). In this sense the decision of the household head to retire and the decision of the son/daughter to leave the household may be co-determined and affected by the eligibility status of the household head. Since the consumption pattern of the household depends on the size of the household itself it is clear that the instrument we use to identify the causal effect of retirement on consumption - the eligibility for retirement status - potentially affects consumption also via the size of the household.

The upper right portion of Table 5 presents the evidence from our data. The estimated causal effect of the eligibility status on the household size is as large as $-.30$, statistically significant. By breaking down this causal effect by type of household membership it is clear that it is driven by the negative causal effect of the household head eligibility for retirement status on the number of children older than 18, that is the kind of effect we mentioned above. No effect is instead found on the proportion

Table 5: Override identification Tests

variable	coefficient	se	variable	coefficient	se
college	0.0219	0.0238	family size	-0.3041*	0.1259
high school	-.0380	0.0469	kids 18+	-0.2539*	0.0926
middle school	0.0300	0.0503	couple	-0.0222	0.0457
primary school	0.0075	0.0607	<i>Poverty sample variables</i>		
age	-.5662	1.5568	w _{it}	2170.5	17999.7
home size	0.1138	5.7213	w _{poor}	0.0033	0.0471

of couples in the sample.

Despite this violation of this exclusion restriction we believe that our interpretation of the results as compatible with the estimated retirement consumption drop is unlikely to be negatively affected. This is because the additional channel through which the eligibility status affects consumption bears as an implication an overstatement of the size of the causal effect of retirement on consumption (a reduction of household size other things being equal causes a reduction of the household consumption). Since in the next section we argue that our (overstated) estimate is consistent with a life-cycle optimising behaviour, *a fortiori* it must be the case that the true causal effect of retirement on consumption is consistent with it.

7 Economic Interpretation

We have found that consumption drops at retirement by a relatively large amount: 9.8%. The question is whether this drop should be taken as evidence against life-time optimizing behaviour.

7.1 Poverty subsample

In an influential paper, Bernheim, Skinner and Weinberg (2001), BSW, relate the size of the consumption drop to wealth prior to retirement. They find that the drop is largest amongst the relatively poor, and this strongly supports their conclusion that the retirement consumption drop is due to inadequate provision for old age by "more impatient" or less informed consumers. We have argued in the introduction that we do not expect lack of resources to be a problem for individuals who retired during our sample period, partly because the pension system provisions were generous and well understood for these cohorts, and partly because employees receive a large lump-sum payment upon retirement (*Liquidazione*).

We have run a formal test similar to that proposed by BSW as follows. We have taken a sub-sample of individuals not yet eligible for retirement and regressed their total wealth on variables that correlate with life-time wealth, but do not change abruptly as people become eligible. The specification includes a set of dummy variables for education and for couples, plus a second order polynomial in the size of the home and a set of zero-sum year dummies. On the basis of this equation (that explains 15.38% of the variance of the dependent variable) we can assign a predicted wealth value (labelled w_{it}) for each sample observation, whether eligible or not.

We then select the lowest third of the w_{it} distribution - and call this is the poverty sub-sample. If BSW mechanism is at play, we expect households in the poverty sub-sample to drop their consumption most at retirement. Our findings go in the opposite direction. Despite a good fit of the first stage regressions ($R^2 = .935$; coefficient on eligibility of .407 with a s.e. of .0382), the effect of retirement on non-durable consumption is effectively zero, on food consumption is $-.027$, with a s.e. of .088.

Thus for the sample of relatively poor households, retirement is strongly associated to eligibility, but does not lead to reductions in consumption. This corroborates our claim, that liquidity problems

at retirement are not an issue in Italy, but is at first surprising. A possible explanation for the negligible impact of retirement on consumption of the relatively poor is that work-related expenses are much less important for this group, that is mostly made of blue collar workers. Blue collar workers normally eat meals at factory canteens for free, use heavily subsidized public transport to go to work, and do not pay for their work clothes. White collar workers, on the contrary, tend to eat in bars and restaurants, drive to work and buy expensive suits and dresses to wear at the office. Thus retirement has much more of an impact on the latter than on the former group's consumption.

7.2 Back-of-the-envelope calculations

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, 2005, and Hurd and Rohwedder, 2006, who report anticipated drops around 13%).

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 :

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

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

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

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:

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

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

$$\frac{dC}{dl} \Big|_{U_c=\bar{U}} = \frac{C(1-\theta)(1-\sigma)}{l(1-\theta+\theta\sigma)}$$

We see that consumption should drop as a result of an increase in leisure if $\sigma > 1$, 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, θ and σ .

We can estimate θ 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 $\theta = 0.44$. 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 σ between 0.5 and 1.5.

We have estimated the average consumption drop associated to retirement to be 9.8%, with a standard error of 5.7%. In our calculations of the σ parameter, retirement doubles leisure. Thus the estimated elasticity is indeed -0.098 .

A 0.098 consumption drop is consistent with utility maximization if $\sigma = 1.19$ ($EIS = 0.84$). If we consider a one-standard error range around the point estimate, we find that a 0.155 drop is consistent with $\sigma = 1.315$ ($EIS = 0.76$) and a 0.041 drop is consistent with $\sigma = 1.076$ ($EIS = 0.93$). In this context, a zero drop obtains if $\sigma = 1$ ($EIS = 1$), whereas consumption should *increase* at retirement for smaller values of σ .

7.3 Work-related expenses

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, we only know non-durable consumption and a food item that is the sum of food at home and meals regularly consumed out of the home. However, we 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 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 6 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.⁶ 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

⁶The 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.

Table 6: Consumption drop in SFB diary data

Category	unadjusted		p-score adjusted	
	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

in meals out is consistent with both a reduction in work-related expenses (meals regularly consumed out of the home) and the home production hypothesis. In fact, there 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 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.

8 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 1993-2004 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 9.8% because of (male) retirement. We have shown that such fall is not driven by liquidity problems for the less well off in the population, and that is compatible with a constant marginal utility of consumption if preferences are non-separable between consumption and leisure. We have also produced evidence that this drop can be mostly accounted for by falls in goods that are typically considered to be work-related expenses or leisure substitutes (clothing, transport, meals out).

References

- [1] Aguiar Mark and Erik Hurst (2005), "Consumption vs. Expenditure", *Journal of Political Economy*, **113(5)**, 919-948.
- [2] Angrist, Joshua D., and Guido W. Imbens (1994), "Identification and Estimation of Local Average Treatment Effects", *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, Ra@aele 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] Browning, Martin, Thomas Crossley and Guglielmo Weber (2003) "Asking Consumption Questions in General Purpose Surveys", *Economic Journal*, 113(491), F540-F587
- [13] Deaton, Angus (1992) *Understanding Consumption*, Oxford: Oxford University Press
- [14] Fisher, Jonathan, David S. Johnson, Joseph Marchand, Timothy M. Smeeding, and Barbara Boyle Torrey (2005) "The Retirement Consumption Conundrum: Evidence from a Consumption Survey", The Center for Retirement Research WP #2005-14
- [15] Guiso, Luigi and Jappelli, Tullio, 2002. "Private Transfers, Borrowing Constraints and the Timing of Homeownership," *Journal of Money, Credit and Banking*, **34(2)**, 315-39.
- [16] 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
- [17] Haider, Steven J., and Melvin Stephens Jr. (2004). "Is There a Retirement-Consumption Puzzle? Evidence Using Subjective Retirement Expectations." NBER Working Paper No. W10257, forthcoming: *Review of Economics and Statistics*.

- [18] Hamermesh, Daniel S. (1984) "Consumption During Retirement: The Missing Link in the Life-Cycle", *Review of Economics and Statistics*, **66(1)**, 1-7.
- [19] 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
- [20] Horowitz, Joel L. and Charles F. Manski (1995), "Identification and Robustness with Contaminated and Corrupted Data", *Econometrica*, **63(2)**, 281-302
- [21] Hurd, Michael and Susann Rohwedder (2006), "Some Answers to the Retirement - Consumption Puzzle", NBER WP 12057
- [22] Johnson, Norman L., Samuel Kotz and N. Balakrishnan (1994), *Continuous Univariate Distributions*, New York: Wiley
- [23] 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
- [24] Miniaci, Raëale, Chiara Monfardini and Guglielmo Weber (2003) "Is There a Retirement Consumption Puzzle in Italy?", IFS WP03/14, revised March 2007
- [25] Modigliani, Franco and Richard Brumberg (1954), "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data". *Post-Keynesian Economics*, Kenneth K. Kurihara (ed.), New Brunswick: Rutgers University Press.
- [26] Smith, Sarah (2006), "Can the Retirement Puzzle Be Resolved? Evidence from UK Panel Data", *Economic Journal*, **116(510)**, C130-C148.
- [27] 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
- [28] Trochim, William M. K. (1984), *Research Design for Program Evaluation: the Regression-Discontinuity Approach*, Beverly Hills: Sage Publications