American Economic Review 2009, 99:5, 2209–2226
http://www.aeaweb.org/articles.php?doi=10.1257/aer.99.5.2209

The Retirement Consumption Puzzle: Evidence from a Regression Discontinuity Approach

By ERICH BATTISTIN, AGAR BRUGIAVINI, ENRICO RETTORE, AND GUGLIELMO WEBER*

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 analyze 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 (Angus Deaton 1992; Martin Browning and Annamaria 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 lifetime optimizing behavior. This is documented for the United Kingdom (James Banks, Richard Blundell, and Sarah Tanner 1998), for the United States (B. Douglas Bernheim, Jonathan Skinner, and Steven Weinberg 2001), and for Italy (Raffaele Miniaci, Chiara Monfardini, and Weber 2003) 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 nonmarket time, unexpectedly low pensions or liquidity problems, as well as myopic or perhaps time-inconsistent behavior.

In this paper we investigate the size of the consumption drop associated with retirement in Italy by exploiting the exogenous variability in pension eligibility to identify the causal effect of retirement on consumption expenditures within a regression discontinuity framework (see Jinyong Hahn, Petra Todd, and Wilbert M. van der Klaauw 2001). To this end, we use data from the Bank of Italy Survey on Household Income and Wealth (henceforth SHIW) for the period 1993–2004. This survey collects information on food, nondurable and total household spending, as well as on the current/last job and the number of years of contributions toward the public retirement pension scheme. We evaluate the change in consumption caused by retirement by exploiting pension eligibility to correct for the endogenous nature of the retirement decision. Our identification strategy rests upon a standard assumption made in the program evaluation literature (see Guido W. Imbens and Thomas Lemieux 2008) which, in our context, amounts to

* Battistin: Department of Statistics, University of Padova, Via Cesare Battisti 241, 35121 Padova, Italy, and IRVAPP (e-mail: erich.battistin@unipd.it); Brugiavini: Department of Economics, University of "Ca Foscari" of Venice, S. Giobbe 873 30121, Venezia, Italy, SSAV, and IFS (e-mail: brugiavi@unive.it); Rettore: Department of Statistics, University of Padova, Via Cesare Battisti 241, 35121 Padova, Italy, and IRVAPP (e-mail: enrico.rettore@unipd.it); Weber: Department of Economics, University of Padova, Via del Santo 33, 35121 Padova, Italy, CEPR and IFS (e-mail: guglielmo.weber@unipd.it). An earlier version of this paper was circulated with the title "How Large is the Retirement Consumption Drop in Italy?" This was 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 (Milan), Ente Luigi Einaudi (Rome), Tsinghua University (Beijing), and CEMFI (Madrid). 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 and Lisa Callegaro for excellent research assistance.

assuming that spending behavior would remain smooth around the threshold for pension eligibility if no individual retired.

In our data, we show that a nonnegligible fraction of individuals retire as soon as they become eligible, and estimate at 9.8 percent the part of the nondurable consumption drop that is associated with retirement induced by eligibility. We find that such fall is not driven by liquidity problems for the less well-off in the population, and can be accounted for by drops in goods that are work-related expenses or leisure substitutes. We also show, however, that retirement induces a significant drop in the number of grown children living with their parents, and this can account for most of the retirement consumption drop. This finding is probably related not only to the fact that Italian youths live with their parents well into their twenties and thirties, but also to the receipt of a sizeable bonus payment at retirement for most Italian employees.

The remainder of this paper is organized as follows. Section I presents the identification strategy to deal with the endogeneity problem arising from self-selection of eligible individuals into retirement. Section II deals with data-related issues and, in particular, spells out the definition of pension eligibility. Section III presents our results and some robustness checks. The economic implications of our findings are discussed in Section IV, while Section V concludes.

I. Identification Strategy

In this paper we classify a household as retired if the head does not work and receives a pension. We focus on male heads of households, for reasons explained in the data section. Let R be a binary variable denoting the retirement status, with $R = 1$ for retired heads and $R = 0$ otherwise. Let S be the distance from the point in time the household head becomes eligible for retirement and let his *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 exhibit a value of the variable S not below zero. 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). Throughout this paper, S is assumed to be continuous.

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 household expenditure ("consumption") corresponding to the head being retired and not being retired, respectively. The causal effect of retirement on consumption 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 corresponding to a change in the retirement status of the household head, which is our quantity of interest.

Given our definition of retirement, the condition of pension eligibility does not necessarily imply that individuals are actually retired. On the other hand, individuals not yet eligible cannot be retired. As a result of the eligibility rule and of self-selection into retirement, the probability of retirement for those scoring negative values of S is zero by definition, since they are not yet eligible. The probability of retirement for those scoring positive values of S is less than one because retirement is not mandatory. This implies that the probability of retirement is discontinuous at the threshold for eligibility, but the discontinuity size is less than one. This discontinuity helps to solve the endogeneity problem due to the self-selection into retirement of eligible workers.[1]

[1] The decision to retire fits what is known as a *regression discontinuity design* in the programme evaluation literature (see Donald L. Thistlethwaite and Donald T. Campbell 1960; Hahn, Todd, and Van der Klauuw 2001). The potential of using eligibility rules to identify causal effects within such a design has been pointed out by several papers in the literature (see, among others, Battistin and Rettore 2008).

It can be shown that the average causal effect of retirement on consumption can be retrieved for those retirees around $S = 0$ (see Battistin and Rettore 2008 and the Web Appendix, available at http://www.aeaweb.org/articles.php?doi=10.1257/aer.99.5.2209, for further details) by considering the ratio

$$E\{\beta \mid R = 1, S = 0^{+}\} = \frac{E\{Y \mid S = 0^{+}\} - E\{Y \mid S = 0^{-}\}}{E\{R \mid S = 0^{+}\}}, \tag{1}$$

where values $S = 0^{+}$ and $S = 0^{-}$ denote individuals marginally above and marginally below $S = 0$, respectively, and the conditional expectations involved refer to consumption and the proportion of retired individuals in these two groups. The condition for this result to hold is the following, which will be maintained throughout the remainder of this paper.

CONDITION 1: *The mean value of Y_0 conditional on S is a continuous function of S at $S = 0$.*

According to this condition, discontinuities of Y around the threshold for eligibility can be given a causal interpretation only if in the absence of retirement no discontinuity would take place in the outcome Y around that threshold. This amounts to assuming a smooth consumption profile at $S = 0$ in a counterfactual world of no retirement.

Identification of causal effects is potentially undermined by measurement error in the variable S, an important issue in our data: a nonnegligible fraction of workers presenting negative values of S self-report to be retired from work (see, for example, Table 4). Although in principle this evidence is also consistent with measurement error in the reporting of the retirement status R, because of the questionnaire design and the way R is defined, measurement error in R appears most unlikely. For this reason, throughout the paper we maintain the assumption that R is not misreported and that all inconsistencies in the data between S and the observed retirement status are due to measurement error in S.[2]

In the Web Appendix we derive the conditions on the measurement error that allow us to retrieve the causal parameter in (1) from raw data. The main results can be summarized as follows. First, we show that the evidence that people are sometimes retired even though their S is negative (see Section IIC) is *not* consistent with the hypothesis of having classical measurement error in the eligibility variable. A more general model for measurement error is therefore called for. We do that by assuming that the observed distribution of the eligibility variable is a *mixture* of values that are measured correctly and values that are reported with error (not necessarily classical in form). In particular, we assume that

$$S_{obs} = S^{*}Z + S(1 - Z), \tag{2}$$

where S^{*} and S denote true and error-ridden measurements of the eligibility variable, respectively, Z is a binary indicator for the exact reporters, and S_{obs} is the survey measurement actually observed. This is known as the *contaminated sampling model* discussed by, among others, Joel L. Horowitz and Charles F. Manski (1995).

[2] 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 Andrew Chesher (2009). For the impact of misclassification (and/or misreporting) of R, see, for example, Battistin and Barbara Sianesi (forthcoming). In the context of this paper, having measurement error in S implies also that the eligibility status is potentially misrecorded, thus inducing a more complicated structure of the error.

Second, we show that, in the presence of the measurement error, the sample analogue of (1) is inconsistent for the parameter of interest. We therefore spell out the conditions under which the sample analogue of the following ratio:

$$\frac{E\{Y \mid S_{obs} = 0^{+}\} - E\{Y \mid S_{obs} = 0^{-}\}}{E\{R \mid S_{obs} = 0^{+}\} - E\{R \mid S_{obs} = 0^{-}\}}, \tag{3}$$

identifies the causal effect of retirement on consumption at $S^* = 0$ (see the Web Appendix for further details). The identification condition required is that, conditional on S^*, the process generating measurement error is orthogonal to the process of interest. This result 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 Joshua D. Angrist 1994; Hahn, Todd, and Van der Klaauw 2001).

To ease readability, in the remainder of this paper we will drop subscripts from S_{obs}. Hence S will be time to/from eligibility as measured in the data.

II. Data

In our analysis, information on consumption and pension status is obtained from SHIW for the period 1993–2004. In what follows, we describe how the variables of interest are defined from raw data. As it will be clear from the remainder of this section, we treat the data as repeated cross sections. Even though a panel component is available in SHIW, we do not exploit it in our empirical exercise, as identification of the causal discontinuity in consumption at retirement does not require panel data. In this sense, the approach suggested in this paper is much less data demanding than those that have been taken so far to study the retirement consumption puzzle. We will nonetheless make use of the longitudinal dimension of the survey to retrieve some of the key variables and/or to improve data quality.

A. *Information on Consumption*

Information on expenditure is derived from recall questions on consumption which underwent several changes over the period covered by our analysis. In the 1995 wave (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, and major house renovations, 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 nondurable consumption and food questions, but the wording and contents were otherwise identical. The 1993 wave, 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.[3]

[3] Battistin, Miniaci, and Weber (2003) compare consumption data across SHIW and diary-based data from the Italian statistical office (ISTAT) for 1995.

B. *The Definition of Retirement Status*

Retirement status is defined on the basis of two questions. In each survey wave, respondents were asked whether each household member was employed for most of the year. If the answer was negative, they were then asked whether the household member was a first-time work seeker, unemployed, homemaker, 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 a job-pensioner. This distinction contributes to the result that the sum of the percentages of workers and pensioners is, in some years, much below 100.

Table 1 provides a brief description of the number of workers and retired persons in SHIW data: the residual category includes other conditions such as disabled individuals or homemakers. In this table we have considered only household heads and their spouses/partners. It should be noted that the percentage of individuals who are active or have been active in the past is far larger for males than for females. This result is partly due to the labor market behavior of older cohorts: women were characterized by lower educational attainment and lower labor market participation. An indication of this is the fact that, over time, there is a nonnegligible increase in the percentage of females who are either working or retired. The low labor force participation by females explains why we focus on male retirement in our analysis of the retirement consumption drop.

C. *The Definition of Pension Eligibility*

The aim of this section is to summarize how we derived the variable S, which measures the time to (or from) pension eligibility. As we have shown in Section I, pension eligibility is a crucial variable in our analysis. It is measured on the basis of both age and seniority (accrued contribution years). The information on seniority is available since the 1995 wave, but the part of the 1993 sample that is reinterviewed in 1995 can also be used.

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 adopted in Italy until the 1992 pension reform, as well as during the transitional phase, which is not yet complete. 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 ten years' earnings. The 1995 pension reform changed this system radically, as benefits should now be computed according to a notionally defined contribution (NDC) method, that is, they 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: in the private sector, people could retire at age 60 (55 for women), 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 through this route before the reforms. Since the 1992 reform, normal retirement age was gradually increased and reached 65 for men and 60 for women in 2001. Starting in 1995, both age and seniority requirements for early retirement became more stringent over time (according to the type of job) as shown for males in Table 2, (These are the rules prevailing after 1998, according to Law 449/1997; these rules apply to white-collar employees, but differ only slightly for blue-collar employees). We can use these criteria to calculate eligibility for most observations in our data. A potential problem is that, with some interruptions, the Government imposed a delay window on retirees after they became eligible, which, depending on the month of birth, would

TABLE 1—COMPOSITION OF THE SAMPLE WITH RESPECT TO RETIREMENT STATUS

	1993	1995	1998	2000	2002	2004
Males						
Worker	62.3 percent	60.1 percent	60.5 percent	58.2 percent	54.9 percent	54.7 percent
Retired	31.8 percent	32.6 percent	31.6 percent	34.9 percent	38.9 percent	39.7 percent
Observations	6,512	6,559	5,844	6,430	6,175	6,126
Females						
Worker	31.3 percent	33.0 percent	34.9 percent	34.3 percent	33.5 percent	33.4 percent
Retired	14.3 percent	14.6 percent	12.7 percent	14.4 percent	17.0 percent	17.7 percent
Observations	5,900	5,910	5,238	5,578	5,323	5,171

Notes: Percentages and sample sizes of males and females in the sample classified as working or retired. The residual group comprises others, such as disabled or homemakers. The sample includes only household heads and their spouses/partners (thus excludes other household members) interviewed between 1993 and 2004.

TABLE 2—RETIREMENT ELIGIBILITY RULES FOR MALES: 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

Source: Brugiavini, Franco Peracchi, and David A. Wise (2002). Eligibility rules prevailing after enactment of Law 335/95 and Law 449/97.

postpone retirement by three or even six 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 a yearly frequency.

The Measurement of Years of Contributions.—In some waves, years of contribution have been explicitly recorded in SHIW for workers and for those currently retired; however, this question was *not* asked in 1993, and missing values are occasionally found in other waves. Because the age of the respondent does not provide enough information to measure eligibility, we have adopted a simple method to impute missing data on years of contribution by distinguishing subjects already retired at the interview time from those still at work. In fact, we make use of retrospective information on (i) self-reported age at entrance in the labor market and (ii) self-reported age of retirement (if retired). The imputation is carried out also taking into account 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 but to drop the observation. This more than halves the sample size for that year (see Table 3). Second, if the individual is observed in 1993 and belongs to the *panel* subsample of SHIW, an imputation of years of contribution is made on the basis of the

1995 recorded figure on contribution years. 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 questions on the age at retirement and the age of entry into the labor market. If the information is missing for a *worker* instead (who 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 labor market. Fourth, if a retiree or a worker belongs to the panel, missing values of contribution years can be recovered from the previous or subsequent waves.[4]

D. *The Working Sample*

In our empirical exercise, we take the head of the household to be the male. 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. We focus our analysis on males who are or have been active in the labor market, as they are the ones for whom job-pension eligibility can be defined. A detailed breakdown by survey year of the selection criteria adopted to derive the final sample used in the analysis is presented in Table 3.

Table 4 presents the percentage of individuals in our final sample who self-report to be retired by the values of the eligibility variable S. For brevity we report only those percentages for values of S between -4 and $+4$. 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. Observing retirees at negative values of S would 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 nonnegligible proportion of individuals whose imputed value of the eligibility variable is negative are retired. In particular, approximately 2.5 percent of males "marginally ineligible" for retirement self-report to be retired. This suggests adopting a measurement error robust estimation strategy, as explained in Section I.

III. Empirical Analysis

A. *Estimation Issues*

The causal effect of retirement on consumption can be estimated along the lines of what was discussed in Section I. Though this effect can be nonparametrically identified (see the Web Appendix for further details), the analysis presented in what follows builds upon a fully parametric approach, which we found to provide a reasonable fit to the data.

As the sample analogue of (3) coincides with an instrumental variable estimator, locally at $S = 0$, where the endogenous variable R is instrumented by the eligibility status $\mathbb{1}(S \geq 0)$ (see Imbens and Angrist 1994), we implemented an IV procedure that can be described as follows. First, for each survey year, we computed averages of household expenditure $Y_{s,t}$ and proportions of retired male heads $R_{s,t}$ by values of S. All variables are indexed by t and s to emphasize that they are defined as sample averages by calendar year and years to/from eligibility. The average sample size across cells for values of S between -10 and 10 is 90 observations, ranging from a minimum of 33 to a maximum of 192 observations.

[4] This imputation could result in 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). The distribution of years of contributions, as well as that of S resulting from this strategy, is reported in the Web Appendix.

Table 3—Sample Selection Criteria

	1993	1995	1998	2000	2002	2004	Total
Raw data	8,089	8,135	7,147	8,001	8,011	8,012	47,395
Step 0	3,329	7,566	6,714	7,377	7,340	7,383	39,709
Step 1	2,792	6,337	5,653	6,127	5,951	5,941	32,801
Step 2	2,655	5,876	5,219	5,733	5,599	5,621	30,703
With $S \in [-10, 10]$	1,085	2,353	1,890	2,042	1,998	1,910	11,278

Notes: Rows in the table refer to different selection steps. Numbers reported refer to observations left after each step. Description of sample selection procedure:
Raw data: Total number of households in SHIW;
Step 0: Keep all households where eligibility can be computed for head or spouse;
Step 1: Keep only households with male head or male spouse;
Step 2: Keep only households where male head or male spouse is working or retired.

Table 4—Percentage of Retired Males by Time to/from Eligibility

Distance	−4	−3	−2	−1	0	1	2	3	4
Percentage	1.3	1.5	1.5	2.5	33.3	62.6	61.3	65.2	74.0
Observations	318	442	467	372	697	527	501	550	500

Notes: Time to/from eligibility is measured as described in Section C, with positive (negative) values denoting the number of years from (to) eligibility. The working sample is obtained as described in Table 3.

Second, we specified the following regression of averages of household expenditure on proportions of retired heads and a quadratic polynomial in S:

$$Y_{s,t} = \beta_0 + \beta_1 R_{s,t} + \beta_2 S + \beta_3 S^2 + \varepsilon_{s,t}, \tag{4}$$

and added, as controls, year-specific intercepts to take into account changes in the SHIW survey questionnaire that occurred over the period covered by our data (see Section II). Implicit in this specification is the assumption that the leisure of the spouse is separable from that of the male head. We instrumented retirement with eligibility status by specifying the following regression:

$$R_{s,t} = \gamma_0 + \gamma_1 \mathbb{1}(S \geq 0) + \gamma_2 S + \beta_3 S^2 + \nu_{s,t}, \tag{5}$$

again allowing for year-specific intercepts.

We limited the estimation sample to values of S within a ten-year band from eligibility. We excluded the value $S = 0$, as the recall question on consumption could, for those who do retire, cover both pre- and postretirement periods.[5] As a result of this procedure, we ended up with 120 cells resulting from 20 values of S and 6 survey waves. The first-stage regression yielded an R^2 of 0.92, and the coefficient on eligibility estimated at 0.435 with a standard error of 0.038. As shown in Figure 1, the relationship between retirement status and eligibility is relatively stable over the years.

[5] We experimented with different selection criteria: though the estimated key parameter is not much affected by the bandwidth choice, we find that with smaller bands (such as five-year bands) estimates are affected by outliers in the age profile; with larger bands (15-year and more) composition effects start playing an important role, because of mortality.

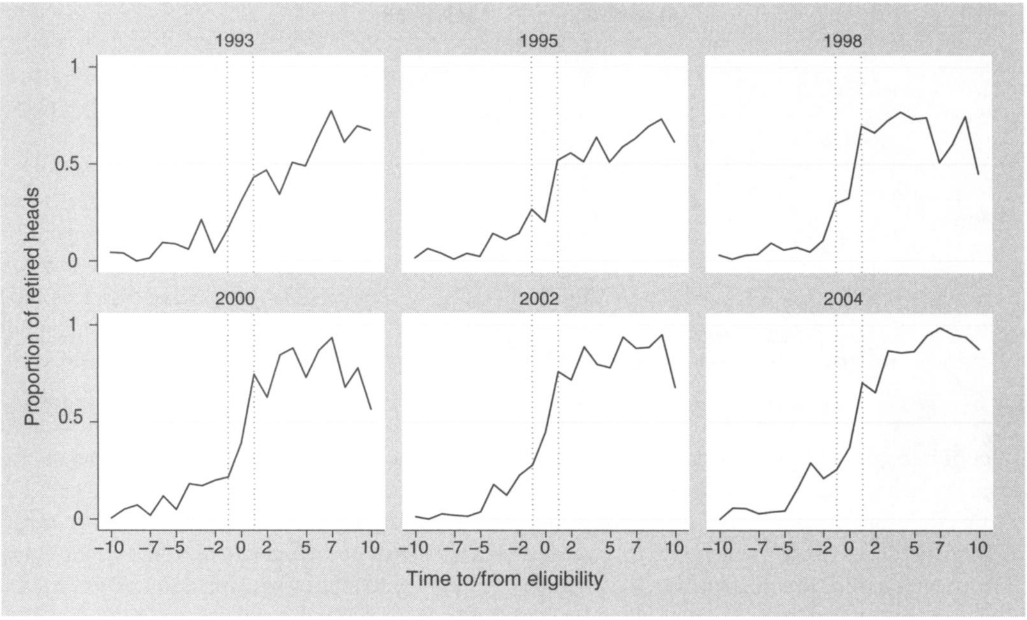

FIGURE 1. PROPORTION OF RETIRED MALES BY TIME TO/FROM ELIGIBILITY AND SURVEY YEAR

B. *Results*

Two sets of results are presented in what follows, for nondurable and for food consumption. Results for nondurable consumption are reported in the top panel of Table 5, suggesting a consumption drop of around 9.8 percent, which is significantly different from zero at the 10 percent level. Results for food are reported in the bottom panel, suggesting a larger consumption drop (14.1 percent), which is significantly different from zero at the 5 percent level.

The lack of precision in the estimation of the drop for nondurable consumption is likely to reflect the noise-to-signal ratio to be expected in this type of recall question (see Martin Browning, Thomas F. Crossley, and Weber (2003) for an appraisal). Recall that questions on food consumption are less heavily affected by memory problems, and therefore more informative (Battistin, Miniaci, and Weber 2003). But if we take point estimates at face value, our finding that food expenditure falls more at retirement than nondurable expenditure may appear at odds with expectations. However, for all years but one (1993), 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. Also, Mark Aguiar and Erik Hurst (2005) make a convincing case that drops in food expenditure at retirement do not translate into drops in food consumption, because of better shopping opportunities associated with more leisure time.

C. *Robustness Checks*

To gather evidence on the validity of Condition 1 (on which our identification strategy crucially relies), we implemented an overidentification test following David S. Lee (2008). The idea behind this test is to exploit outcomes satisfying the following two conditions. First, it has to be known a priori that they are not affected by the eligibility status. Thus, there must be no causal

TABLE 5—ESTIMATION RESULTS

	Coefficient	Standard error	t	p-value
Nondurable expenditure (logs)				
Retired	−0.0983	0.0567	−1.74	0.085
S	−0.0055	0.0027	−2.05	0.043
S^2	−0.0003	0.0002	−1.91	0.059
Food expenditure (logs)				
Retired	−0.1409	0.05442	−2.59	0.011
S	−0.0028	0.0026	−1.07	0.287
S^2	−0.0001	0.0001	−0.58	0.561

Notes: Instrumental variables estimates based on 120 cell means obtained from micro data as described in Section A. The estimated equation relates expenditure to a dummy for retirement, controlling for time to/from eligibility and survey year dummies. Retirement is instrumented by eligibility status. Standard errors are robust to heteroskedasticity.

effect of becoming eligible on these outcomes. Second, they have to be correlated with the unobservables, which are likely to affect the level of consumption.

Consider, for instance, the case of education. It is well known that education is a good proxy for lifetime access to economic resources, and should therefore affect consumption. However, educational choices are determined long before eligibility for retirement. Thus, the causal effect of retirement on education must be zero. It therefore follows that if household heads marginally below and above the threshold $S = 0$ presented significant differences in their education level, we would have evidence counter to the validity of our identification strategy. Discontinuities in outcome variables determined prior to eligibility are consistent with the hypothesis that households at $S = 0^-$ are not fully comparable to households at $S = 0^+$, so any discontinuity detected in the outcome of interest (consumption) may signal other sources of selection.

In what follows, we implement the general idea in Lee (2008) by applying the same estimation procedure described in Section IIIA to detect discontinuities on a battery of outcomes determined prior to eligibility. They include dummies for educational qualifications, age of the head, years of pension contributions of the head, and size of the main residence. Results are reported in the top panel of Table 6, and in all cases we considered they are consistent with our identifying restriction.

IV. Economic Interpretation

We have found that nondurable consumption drops at retirement by a relatively large amount: 9.8 percent. The question is whether this should be taken as evidence against lifetime optimizing behavior.

International evidence on the size of the drop is in line with our findings. Banks, Blundell, and Tanner (1998) 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, overpredicts the level of consumption by as much as 1.5 percent on an annual basis for those between the ages of 60 and 67. The cumulated consumption shortfall over this age band, where most people retire, is around 10 percent. They argue that only a fraction of this drop can be attributed to the increased leisure time that accompanies retirement. Later work by Sarah Smith (2006) uses information on food for UK households that retired over the sample period, and stresses the importance of distinguishing between voluntary and involuntary retirement: a significant drop in food consumption is observed only for those who retire early because of poor health or job loss. Indeed, David M. Blau (2008) stresses that consumption drops at retirement can be reconciled with lifetime optimization if there is uncertainty over layoffs, job offers, health, and mortality,

TABLE 6—OVERIDENTIFICATION TESTS

	Coefficient	Standard error	t	p-value
Variables determined before retirement plans (see Section IIIC)				
College degree	0.0219	0.0238	0.92	0.358
High school diploma	−0.0380	0.0469	−0.81	0.419
Middle school	0.0300	0.0503	0.60	0.551
Primary school	0.0075	0.0607	0.12	0.902
Years of contributions	−1.9376	1.3585	−1.43	0.157
Age	−0.5662	1.5568	−0.36	0.717
Age^2	−8.9669	169.65	−0.05	0.958
Homesize	0.1138	5.7213	0.02	0.984
Wealth variables (see Section IVB)				
Wealth	2,170.5	17,999.7	0.12	0.904
Poverty dummy	0.0033	0.0471	0.07	0.946

Notes: Instrumental variable estimates based on 120 cell means. The estimated equation relates each variable in the table to a dummy for retirement controlling for time to/from eligibility and survey year dummies. Retirement is instrumented by eligibility status. Wealth has been imputed using the procedure described in Section IVB. The poverty indicator takes value 1 if imputed wealth is in the bottom third of the distribution. Standard errors are robust to heteroskedasticity.

and if retirement is a discrete event that is freely chosen by the household. However, in Blau's model the causal effect of retirement on consumption is zero.

Bernheim, Skinner, and Weinberg (2001, henceforth BSW) use Panel Study of Income Dynamics (PSID) data to estimate Euler equations for food consumption. The retirement status is instrumented by taking age-specific predicted probabilities conditional on demographics. The sample is split into groups: low wealth-to-income households drop their consumption most. They estimate a median drop of 14 percent, but higher drops for low wealth ratio, 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 retire with inadequate savings (as originally suggested by Daniel S. Hamermesh 1984). Recently, smaller drops have been estimated by Steven J. Haider and Melvin Stephens Jr. (2007) (for those who retire at the expected time) and by Jonathan Fisher et al. (2005) who use Consumer Expenditure Survey (CEX) data, deflate expenditure by the squared root of household size, and report a smaller drop (around 2.5 percent) for total expenditure than for food consumption (around 5.7 percent).

The only available evidence for Italy is provided in Miniaci, Monfardini, and Weber (forthcoming) and is based on diary-level data covering the 1985–1996 period. They estimate a fall in nondurable consumption at retirement of 5.4 percent and argue that this is a lower bound if there is heterogeneity in work-related expenses and individuals with higher expenses retire earlier. The data they use do not contain any information on work histories, and they cannot identify what part of the consumption fall at retirement is planned and what part is instead due to the realization of shocks that affect the retirement decision. Miniaci, Monfardini, and Weber stress that the Italian case is of particular interest because liquidity problems are unlikely to play a role: Italian employees receive a large bonus payment upon retirement (technically, severance pay). If cash considerations mattered, one would expect a surge in consumption at retirement rather than a drop.[6]

[6] Private sector employees contribute a nonnegligible fraction of gross annual earnings (6.91 percent) to a severance pay fund, which earns 1.5 percent return plus 75 percent of inflation for each year, and is paid as a lump sum when the job is terminated. Public sector employees contribute the same, but their final payment is based on their last pay

Recent papers stress that a drop in expenditure after retirement does not necessarily imply a drop in utility. Work-related expenditure (transport to and from work, meals, and business clothing) is no longer needed (Miniaci, Monfardini, and Weber 2003), home production of services (laundry, gardening, house-cleaning, cooking) may become advantageous (Michael D. Hurd and Susann Rohwedder 2006), and the extra leisure time may allow consumers to shop more efficiently (Aguiar and Hurst 2005, 2007). Erik Hurst (2008) argues that the available evidence is fully compatible with the life-cycle model, and does not imply time-inconsistent behavior (George-Marios Angeletos et al. 2001).

In what follows, we produce a back-of-the-envelope exercise that may shed light on whether the estimated drop in consumption induced by retirement is due to increased consumption of leisure. We then explore alternative explanations by investigating whether the drop is larger for people who enter retirement with little wealth, whether family size is affected by retirement of the head, and how the composition of expenditures varies after retirement.

A. *Back-of-the Envelope Calculations*

A simple way to check whether our estimated consumption drop is consistent with the life-cycle model without uncertainty 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 nondurable consumption, C, and (male) leisure, l, both at time t (we drop time subscripts for convenience):

$$U = \frac{(C^{\alpha} l^{1-\alpha})^{1-\gamma}}{1-\gamma}, \tag{6}$$

where $\gamma > 0$ is the reciprocal of the elasticity of intertemporal substitution (EIS), and

$$\alpha = \frac{pC}{pC + wl} \tag{7}$$

measures the within-period consumption share for periods when leisure is not at a corner (hence $0 < \alpha < 1$). Moreover, p is the price of the consumption good at time t, and w is the nominal wage at time t. Life-time optimization implies that households should keep the marginal utility of consumption, U_c, constant:

$$U_c = \alpha C^{\alpha - 1 - \alpha\gamma} l^{(1-\alpha)(1-\gamma)}. \tag{8}$$

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

$$\frac{dC}{dl}\Big|_{U_c = \bar{U}} = \frac{C(1-\alpha)(1-\gamma)}{l(1-\alpha+\alpha\gamma)}. \tag{9}$$

We see that consumption should drop as a result of an increase in leisure if $\gamma > 1$, and it should increase otherwise. A great attraction of the Cobb-Douglas functional form lies in its analytical

and tenure (it can be worth up to 40 times their gross monthly salary). Both types of employees can borrow against their severance fund payment an amount that cannot exceed 70 percent of its current value, but only for exceptional health expenses or for the purchase of their own home or that of their children. There is no such fund for self-employed workers.

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 percent, so $\alpha = 0.44$. No estimate for the EIS 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 (Ravi Bansal and Amir Yaron 2004), emphasize that 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 with retirement to be 9.8 percent. In our calculations of the α parameter, retirement doubles leisure. Thus, the estimated elasticity is indeed -0.098. Such a drop is consistent with utility maximization if $\gamma = 1.19$ (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 $\gamma = 1.315$ ($EIS = 0.76$) and a 0.041 drop is consistent with $\gamma = 1.076$ (EIS $= 0.93$). In this context, a zero drop is obtained if $\gamma = 1$ (EIS $= 1$), whereas consumption should *increase* at retirement for smaller values of γ.

B. *Poverty Subsample*

In an influential paper, BSW relate the size of the consumption drop to wealth prior to retirement. They find that the drop is largest among 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 do not expect lack of resources to be a problem in Italy 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 received a large lump-sum payment upon retirement.

However, we have also run a formal test similar to that proposed by BSW, who have panel data, so can condition on preretirement wealth. Since the panel component of SHIW comprises too few individuals, in what follows we exploit only repeated cross sections of data, and impute preeligibility wealth to all individuals in the final sample (workers and retired)—this is preferable to computing preretirement wealth, given that we treat retirement as endogenous. We select individuals not yet eligible for retirement and regress their total wealth on variables that correlate with lifetime 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 (which explains 15.38 percent of the variance of the dependent variable), we can assign a predicted wealth value for each sample observation, whether eligible or not.

We then select the lowest third of the imputed wealth distribution and call it the poverty subsample. We check that the variables we use to split the sample can be considered to be genuinely predetermined with respect to eligibility for retirement variables by running the same overidentifying restrictions test that we presented for education, age, and home size in Section IIIC. In this case, we ran instrumental variables regressions for both the continuous wealth variable and the discrete zero-one indicator on retirement and found insignificant coefficients, as shown in the bottom panel of Table 6.

If the BSW mechanism is at play, we expect households in the poverty subsample to drop their consumption most at retirement. Our findings are in the opposite direction. Despite a good fit of the first-stage regressions ($R^2 = 0.935$; coefficient on eligibility of 0.407 with a standard error of 0.0382), the effect of retirement on nondurable consumption is effectively zero, and on food consumption is −0.027, with a standard error of 0.088.

Thus, for the sample of relatively poor households, retirement is strongly associated with 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 calls for an explanation. A possible reason for the negligible impact of retirement on consumption of the relatively poor is that work-related expenses are much less important for this group, which is primarily 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.

C. *Cohabitation of Parents and Children*

Italian youths live with their parents well into their twenties and thirties. This remarkable feature of Italian society is at least partly explained by economic considerations: Sascha O. Becker et al. (2004) have focused on the inability of young Italians to find secure jobs, Marco Manacorda and Enrico Moretti (2006) have emphasized parents' desire for cohabitation, while Rob Alessie, Brugiavini, and Weber (2006) have stressed that young people often do not have adequate financial resources to move out and enter the housing market.

This wide-spread cohabitation, together with the rules of the Italian pension system, suggest that an additional likely causal effect of the household head being eligible for retirement may be a change in household composition. In fact, as we have seen, employees receive severance pay upon retirement, which is related to the number of years on the job. The anecdotal evidence is that as fathers (and mothers) retire, they use it to (help) buy a house for their sons and daughters, who then leave the parental home (Luigi Guiso and Tullio Jappelli (2002) document the role that inter vivos transfers play in Italy as regards home-ownership). In this sense the decision of the household head to retire and the decision of the son/daughter to leave home may be codetermined 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 retirement induced by eligibility affects consumption both directly and via its effect on household size.[7]

In this case, we have run instrumental variables regressions for overall household size and the number of cohabiting grown children on retirement and found highly significant coefficients (see the top panel of Table 7). The estimated causal effect of the retirement status on the household size is as large as −0.30, 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's retirement status on the number of children older than 18, which is the kind of effect we mentioned above. However, no effect is found regarding the number of younger children or the proportion of couples in the sample.

The immediate consequence of these results is that the estimated retirement consumption drop does not translate into a similar change in marginal utility of consumption, if utility depends on family size. The simplest way to allow for the effects of family size on utility is to express consumption in per capita terms, either by dividing it directly by the number of

[7] Manacorda and Moretti (2006) also stress the importance of parents' pension eligibility on the home-leaving decision by the children. They claim that pension reforms of the 1990s forced parents to stay at work longer and hence raised their income. In their model, parents have a taste for coresidence: the additional income due to postponed retirement would have been used to bribe the children to stay home.

TABLE 7—EFFECT OF RETIREMENT ON DEMOGRAPHICS AND EQUIVALIZED EXPENDITURES

	Coefficient	Standard error	t	p-value
Demographics (see Section IVC)				
Household size	−0.3041	0.1259	−2.42	0.017
Number of children under 18	0.0372	0.0249	1.49	0.136
Number of children 18+	−0.2539	0.0926	−2.74	0.007
Dummy for couple	−0.0222	0.0457	−0.49	0.628
Equivalized expenditures using family size (see Section IVC)				
Nondurable (logs)	0.0165	0.0583	0.28	0.777
Food (logs)	−0.0271	0.0526	−0.77	0.443
Equivalized expenditures using square root of family size (see Section IVC)				
Nondurable (logs)	−0.0411	0.0536	−0.52	0.606
Food (logs)	−0.0840	0.0512	−1.64	0.100

Notes: Instrumental variables estimates based on 120 cell means. The estimated equation relates each variable in the table to a dummy for retirement controlling for time to/from eligibility and survey year dummies. Retirement is instrumented by eligibility status. Standard errors are robust to heteroskedasticity.

household members, or, more sensibly, by some function (the square root for instance) of the number of household members that takes into account economies of scale. In the two bottom panels of Table 7 we provide point estimates and standard errors of the coefficient of interest when the dependent variable is redefined to take into account family size (we do not report the coefficients on the constant, S, its square, or year dummies). The central panel corresponds to the case where consumption is defined in per capita terms; the bottom panel corresponds to the case where it is deflated by the square root of family size (as in OECD 2008 and in Fisher et al. 2005). We see that the retirement consumption drop disappears in the former case, and is halved in the latter case.

Given that there is a strong case for taking economies of scale into consideration, we shall consider the estimates presented in the bottom panel as our preferred estimates. We can thus conclude that retirement induces a 4.1 percent drop in nondurable consumption (insignificantly different from zero), and an 8.4 percent drop in total food (borderline significant).

D. *Work-Related Expenses*

The simple, one-good case presented in our back-of-the-envelope calculations hides the fact that the impact of retirement on total (nondurable) consumption is ambiguous, because some goods may be leisure substitutes and some other leisure complements. Miniaci, Monfardini, and Weber (2003) show examples of both, notably food out and transport as substitutes, food at home and heating fuel as complements. Aguiar and Hurst (2008) show that entertainment expenses also increase after retirement, in agreement with their nature of leisure complements.

We do not have detailed expenditure information in our main dataset, SHIW, as we know only nondurable consumption and a food item that is the sum of food at home and meals regularly consumed out of the home. However, we gained access to diary-level data on consumer spending for 2002 for a different sample of households. This large dataset, collected by 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 (Battistin, Miniaci, and Weber (2003, forthcoming), compare information on food spending from this dataset and from SHIW for 1995). Thus, years to and from pension eligibility are not known, and our identification strategy cannot be applied.

What we can do is compare two groups of households, one whose head's age lies between 50 and 54 and the other between 65 to 69, respectively. Heads of household in the former group are mostly employed (81.8 percent are employed, 9.6 percent are retired, the others are either unemployed or out of the labor force); in the latter they are mostly retired (82.7 percent are retired, 8.0 percent are employed, all the others are out of the labor force).

In Table 8 we report the difference between average spending of the older group and average spending of the younger group, and standard errors. The first column lists the various commodities considered, the second and third columns present a straight comparison, and 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.[8]

The first row of the table depicts that nondurable consumption falls by €510 a month (−31.1 percent) between the early fifties and late sixties. However, once composition effects are taken into account, this drop is reduced to €241 a month (−15.6 percent). 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 percent of the overall consumption drop over this period of the life cycle is due to retirement; the remaining 40 percent 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 together account for €169, 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 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—around 41 percent in relative terms!—and so does its budget share.

As we know from the previous section, eligibility-induced retirement causes a significant drop in the number of grown children living in the household. Thus far we have neglected this drop, on the assumption that children who are induced to leave home by retirement would leave home at some later stage anyway. However, we can check what would happen if this assumption were incorrect, and all such children stayed home, by evaluating the p-score correction for an increased number of cohabiting children in the older age group (in a sense, in doing so we partially adjust the number of equivalent adults). Given the estimates presented in Table 7, we attribute to all households age 65–69 a fourth of an adult more, and obtain the results shown in the last two columns of Table 8. These are reasonably close to the ones presented in the previous two columns.

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.

V. Conclusions

In this paper we have investigated the size of the consumption drop in Italy due to retirement. We have used micro data covering the 1993–2004 period on food and nondurable household spending, and evaluated the causal effect of retirement on consumption 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

[8] The adjustment was made via propensity score weighting so as to make the distribution of these characteristics for the younger and the older groups equal to the distribution of households whose head is between the ages of 59 and 64.

TABLE 8—CONSUMPTION DROP IN DIARY DATA

	Unadjusted		*p*-score adjusted		Partially adjusted	
Category	Difference	Standard error	Difference	Standard error	Difference	Standard error
Nondurable	−510.35	22.70	−241.02	29.74	−305.91	24.45
Food at home	−89.70	6.74	−5.96	8.90	−25.57	7.20
Meals out	−44.37	2.95	−35.68	3.46	−39.38	3.26
Alcohol	−5.24	0.96	−2.75	1.13	−3.46	1.05
Tobacco	−13.35	0.95	−7.98	1.26	−9.56	0.03
Clothing	−99.09	7.26	−58.05	9.28	−68.48	7.64
Personal services	−7.96	2.73	−5.29	3.01	−5.65	2.81
Transport	−153.67	7.56	−76.04	9.91	−96.27	8.09
Heating	−19.75	2.56	−8.73	3.19	−10.31	2.72
Phones	−15.95	0.96	−8.99	1.14	−10.20	1.03
Housing services	−7.82	4.34	−2.55	4.47	−3.01	4.39
Other	−53.43	4.02	−28.99	4.93	−34.01	4.32

Notes: Difference in mean spending of persons 65–69 to 50–54 years old on various categories and their standard errors using 2002 diary data from ISTAT. Number of observations: 2,720 for 50–54 age band; 2,460 for 65–69 age band.

assumption that consumption would be the same around the threshold for pension eligibility if the individual would not retire. We have shown that a nonnegligible 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 time to/from 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 nondurable consumption drops by 9.8 percent because of (male) retirement. We show that such a fall is not driven by liquidity problems for the less well-off in the population, and can be accounted for by drops in goods that are work-related expenses or leisure substitutes, such as clothing, meals out, and transport. However, we also show that retirement induces a significant drop in the number of grown children living with their parents, and this can account for a large fraction of the retirement consumption drop: if expenditure is deflated by a standard family size measure that takes into account economies of scale in consumption, the estimated drop is more than halved, and its point estimate is not significantly different from zero.

REFERENCES

Aguiar, Mark, and Erik Hurst. 2005. "Consumption versus Expenditure." *Journal of Political Economy*, 113(5): 919–48.

Aguiar, Mark, and Erik Hurst. 2007. "Life-Cycle Prices and Production." *American Economic Review*, 97(5): 1533–59.

Aguiar, Mark, and Erik Hurst. 2008. "Deconstructing Lifecycle Expenditure." National Bureau of Economic Research Working Paper 13893.

Alessie, Rob, Agar Brugiavini, and Guglielmo Weber. 2006. "Saving and Cohabitation: The Economic Consequences of Living with One's Parents in Italy and the Netherlands." In *NBER International Seminar on Macroeconomics 2004*, ed. Richard H. Clarida, Jeffrey A. Frankel, Francesco Giavazzi, and Kenneth D. West, 413–58. Cambridge, MA: MIT Press.

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg. 2001. "The Hyberbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation." *Journal of Economic Perspectives*, 15(3): 47–68.

Banks, James, Richard Blundell, and Sarah Tanner. 1998. "Is There a Retirement-Savings Puzzle?" *American Economic Review*, 88(4): 769–88.

Bansal, Ravi, and Amir Yaron. 2004. "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles." *Journal of Finance*, 59(4): 1481–1509.

Battistin, Erich, and Andrew Chesher. 2009. "Treatment Effect Estimation with Covariate Measurement Error." CeNMAP Working Paper CW25/09.

Battistin, Erich, and Enrico Rettore. 2008. "Ineligibles and Eligible Non-participants as a Double Comparison Group in Regression-Discontinuity Designs." *Journal of Econometrics*, 142(2): 715–30.

Battistin, Erich, and Barbara Sianesi. Forthcoming. "Misclassified Treatment Status and Treatment Effects: An Application to Returns to Education in the UK." *The Review of Economics and Statistics*.

Battistin, Erich, Raffaele Miniaci, and Guglielmo Weber. 2003. "What Do We Learn from Recall Consumption Data?" *Journal of Human Resources*, 38(2): 354–85.

Becker, Sascha O., Samuel Bentolila, Ana Fernadez, and Andrea Ichino. 2004. "Job Insecurity and Children's Emancipation: The Italian Puzzle." CESifo Working Paper WP1144.

Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg. 2001. "What Accounts for the Variation in Retirement Wealth among US Households?" *American Economic Review*, 91(4): 832–57.

Blau, David M. 2008. "Retirement and Consumption in a Life Cycle Model." *Journal of Labor Economics*, 26(1): 35–71.

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

Browning, Martin, Thomas F. Crossley, and Guglielmo Weber. 2003. "Asking Consumption Questions in General Purpose Surveys." *Economic Journal*, 113(491): F540–67.

Brugiavini, Agar, Franco Peracchi, and David A. Wise. 2002. "Pensions and Retirement Incentives: A Tale of Three Countries: Italy, Spain and the USA." *Giornale degli Economisti e Annali di Economia*, 61(2): 131–69.

Deaton, Angus. 1992. *Understanding Consumption*. Oxford: Oxford University Press.

Fisher, Jonathan, David Johnson, Joseph Marchand, Timothy M. Smeeding, and Barbara Boyle Torrey. 2005. "The Retirement Consumption Conundrum: Evidence from a Consumption Survey." Boston College Center for Retirement Research Working Paper 2005–14.

Guiso, Luigi, and Tullio Jappelli. 2002. "Private Transfers, Borrowing Constraints and the Timing of Homeownership." *Journal of Money, Credit, and Banking*, 34(2): 315–39.

Hahn, Jinyong, Petra Todd, and Wilbert van der Klaauw. 2001. "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design." *Econometrica*, 69(1): 201–09.

Haider, Steven J., and Melvin Stephens, Jr. 2007. "Is There a Retirement-Consumption Puzzle? Evidence Using Subjective Retirement Expectations." *Review of Economics and Statistics*, 89(2): 247–64.

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

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

Hurd, Michael D., and Susann Rohwedder. 2006. "Some Answers to the Retirement-Consumption Puzzle." National Bureau of Economic Research Working Paper 12057.

Hurst, Erik. 2008. "The Retirement of a Consumption Puzzle." National Bureau of Economic Research Working Paper 13789.

Imbens, Guido W., and Joshua D. Angrist. 1994. "Identification and Estimation of Local Average Treatment Effects." *Econometrica*, 62(2): 467–75.

Imbens, Guido W., and Thomas Lemieux. 2008. "Regression Discontinuity Designs: A Guide to Practice." *Journal of Econometrics*, 142(2): 615–35.

Lee, David S. 2008. "Randomized Experiments from Non-random Selection in U.S. House Elections." *Journal of Econometrics*, 142(2): 675–97.

Manacorda, Marco, and Enrico Moretti. 2006. "Why Do Most Italian Youths Live with Their Parents? Intergenerational Transfers and Household Structure." *Journal of the European Economic Association*, 4(4): 800–29.

Miniaci, Raffaele, Chiara Monfardini, and Guglielmo Weber. 2003. "Is There a Retirement Consumption Puzzle in Italy?" Institute for Fiscal Studies Working Paper W03/14.

Miniaci, Raffaele, Chiara Monfardini, and Guglielmo Weber. Forthcoming. "How Does Consumption Change upon Retirement?" *Empirical Economics*.

OECD. 2008. *Growing Unequal? Income Distribution and Poverty in OECD Countries*. Paris, http://www.oecd.org/els/social/inequality.

Smith, Sarah. 2006. "The Retirement-Consumption Puzzle and Involuntary Early Retirement: Evidence from the British Household Panel Survey." *Economic Journal*, 116(510): C130–48.

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): 309–17.