AperTO - Archivio Istituzionale Open Access dell'Università di Torino

(Un)expected retirement and the consumption puzzle

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/148496> since 2015-12-10T10:58:20Z

Published version:

DOI:10.1007/s00181-013-0760-z

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

(Un)expected retirement and the consumption puzzle*

Margherita Borella**♦

Flavia Coda Moscarola*

Mariacristina Rossi**

Turin, June 2013

Abstract

In this work we revisit the retirement consumption puzzle using Italian panel data. As emphasised in the literature, the observed consumption drop might be due to unexpected wealth shocks at retirement which modify optimal consumption plans. Using an Euler equation approach, we test the impact of unexpected retirement on the consumption patterns of individuals around the age of retirement by using the panel component of the Survey of Household Income and Wealth (SHIW). This data set contains information on the expected age of retirement which can be used to distinguish between expected and unexpected retirement. Furthermore, we investigate the heterogeneous behaviour of individuals with different levels of education and wealth. We find evidence of a consumption drop at retirement especially for low-educated people and individuals with little wealth. The consumption drop at retirement, on average, does not seem to be a response to unexpected retirement. Disaggregating our sample, we find that the consumption drop persists among low-educated people with little wealth available, irrespective of whether retirement was expected or not. Highly educated people, conversely, do smooth their consumption, unless they have low wealth and are hit by an unexpected shock at retirement in which case they are forced to drop consumption.

Keywords: consumption, life cycle, retirement puzzle, unexpected retirement

JEL classification: D91, J26, D83, D84

This article is published as:

Empirical Economics (2014) 47:733–751

DOI 10.1007/s00181-013-0760-z

The final publication is available at link.springer.com

*University of Turin and CeRP-Collegio Carlo Alberto. ** University of Turin, CeRP-Collegio Carlo Alberto and NETSPAR.

Contacts: flavia.codamoscarola@carloalberto.org; mariacristina.rossi@unito.it

♦ Corresponding author: Margherita Borella, e-mail: margherita.borella@unito.it.

1. Introduction

One of the main predictions of the simple version of the Life Cycle Model (LCM) is that consumption should be kept stable over one's life span irrespective of income fluctuations¹. Consumption smoothing is accomplished through borrowing and dis/saving channels. One of the biggest fluctuations in income takes place at retirement when consumption is also observed to exhibit a sizeable drop which cannot be fully explained in a standard intertemporal utility optimisation framework (the 'retirement consumption puzzle', after the work by Banks, Blundell and Tanner, 1998). Understanding whether and why consumption drops at retirement is important not only as a way to test validity of the LCM, but also because it may signal a situation of vulnerability for the family at retirement, a matter in which policy-makers should perhaps intervene.

As first emphasized by Banks, Blundell and Tanner (1998), the consumption drop at retirement can be reconciled within the LCM to the extent that retirement and the related income drop – the pension income is commonly lower than the last wage – is unexpected: households should indeed react to unforeseen events by modifying their consumption rules. Accordingly, Smith (2006) found that food consumption in the UK dropped at retirement only for those households whose head had retired involuntarily while there was no drop for those who decided to retire. Using subjective retirement expectations, Haider and Stephens (2007) found a less pronounced decline in consumption in the US when retirement was expected by individuals. The degree of retirement expectedness could, in theory, be an important factor for better understanding whether the consumption drop is a response to an unexpected shock or something voluntarily planned by individuals. In this study, we test the impact of unexpected retirement on consumption using the Italian Survey of Household Income and Wealth (SHIW), a rich data set containing information on the expected age of retirement which can be used to distinguish between expected and unexpected retirement.

Other potential explanations for the consumption drop at retirement include the decrease in work-related expenditures, the non-separability of preferences about leisure and consumption and home production. Work-related expenditures (Banks, Blundell and Tanner, 1998; Miniaci, Monfardini and Weber, 2010) are expenses for goods and services that do not create utility per se for the individual. They are necessary while the individual works (travel, clothing and eating out expenses), but become superfluous during retirement. In the same vein, if workers compensate for the disutility of work with consumption (i.e. leisure is non-separable from consumption), they can reduce their consumption according to the increase in leisure occurring at retirement (French, 2005; Blau, 2008). Finally, as shown by Hurd and Rohwedder (2003) and Aguiar and Hurst (2005, 2008), individuals in the US have more spare time during retirement and can produce some goods (such as food) at home or search for the best price for quality goods. In this situation, they spend less but they attain the same level of utility they had before retirement².

An important and relatively little explored explanation, due to the lack of suitable data, for the retirement consumption puzzle relates to the information about retirement benefits and to the preparedness for retirement. If retirement occurs as a shock to workers, their unpreparedness for a reduction in income following early retirement could be the most sensible explanation for why

¹ In general, utility maximisation implies that the marginal utility of consumption is constant over time. If utility depends upon consumption only under equality between the interest rate and the subjective discount rate, the result converts into constant consumption.

² Outside the life-cycle framework, Bernheim, Skinner and Weinberg (2001) find for the US that the wealth accumulation behaviour of individuals responds more to a 'rule of thumb' than to the LCM. Angeletos et al (2001) demonstrate through simulation methods that hyperbolic (rather than geometric) discounting households have self-control problems that lead to dynamic-inconsistent behaviours and induce a planned fall in consumption at retirement.

consumption drop typically occurs. Households might not have adequately saved during their working life and therefore might have to reduce their expenditures when retirement occurs. Gustman and Stenmeier (2001) found that US workers often had misinformation or a lack of information about retirement benefits. Lusardi (1999, 2000) found that, *ceteris paribus*, households who had given little thought to retirement had far lower wealth than those that had given more thought to the subject. We argue that if a shock to retirement occurs, where individuals are forced to retire earlier than expected and reduce their expected retirement wealth, the drop in consumption will be higher than the average compared to households who gave little thought to retirement and who are less prepared to face an adverse wealth shock.

While potentially due to a number of different causes, the consumption drop at retirement is likely to present a large amount of heterogeneity across households. In this paper, we are able to disentangle the effect of consumption drop as a response to shock rather than a planned behaviour, unaffected by any surprise. As the ability to plan, even in the absence of shock, is likely to be determined by educational level, we also explore in-depth the impact of education, which we consider as a proxy for planning ability. In particular, we analyse whether potentially better planners behave differently from their peers both in the presence and in the absence of a retirement shock. Differences in planning effort and in the propensity to plan are likely to be strongly associated with differences in wealth accumulation (Ameriks, Caplin and Leahy, 2003) and with different variations in consumption after retirement (Bernheim, Skinner and Weinberg, 2001; Aguiar and Hurst, 2005). To the extent that the retirement income decline is partly unanticipated, we expect those households with a more consistent level of net worth to be more protected against negative income shocks.

In this paper, we study the consumption dynamics of Italian households by exploring the panel dimension of the SHIW database. We make use of a unique feature of the SHIW data: the expected age of retirement for each worker in each wave. In addition, we characterise the behaviour of each household based on some observable features that can reveal interesting differences between individuals. Primarily, individuals' educational background is considered as it is likely to influence their wealth, hence their capability to self-insure against risks.

Higher level of education are likely to be associated with better planning and to a higher propensity to achieve consumption smoothing, thus avoiding consumption drop. An additional component we think it is relevant for our analysis is the buffer stock of wealth, which might ease the consumption smoothing mechanism. We look at accumulated wealth level to account explicitly for the role of buffer stock in preventing a consumption drop at retirement.

Our analysis detects a significant consumption drop at retirement on expenses for non-durables. Adding to the literature, our study sheds light on whether households reduce their consumption at retirement as a response to a shock in the time of retirement. In our analysis, we isolate retired individuals according to whether they expected to retire or not to clearly identify the extent to which households are forced to reduce their welfare (by reducing expenditures) when retirement occurs unexpectedly. Our findings suggest that the consumption drop is not a response to retirement unexpectedness. Both surprised and unsurprised households, on average, drop consumption at retirement with undistinguishable intensity. However, the consumption drop is different across educational levels. Disaggregating our sample by education and wealth, in fact, we find that low educated households with little available wealth drop their consumption at retirement, irrespective of whether the event was expected or not. Highly educated households, conversely, do smooth their consumption, unless they have low wealth and are hit by an unexpected shock at retirement where they are forced to drop consumption.

The rest of the paper is laid out as follows: Section 2 shows the empirical strategy we apply to test the presence and the determinants of the consumption drop; Section 3 illustrates the data; Section 4 discusses the results; and Section 5 reports the conclusions.

2. Empirical strategy

In order to analyse the consumption patterns of individuals around retirement, we estimate an Euler equation derived on the assumption of intertemporal separable lifetime preferences and constant relative risk aversion (CRRA) within-period utility. To take into account the role of demographic variables, we specify the within-period utility function as:

$$u(C_{i,t}) = \frac{\exp(\beta_1 Z_{i,t})}{1 - \rho} C_{i,t}^{1-\rho} \quad (1)$$

where ρ is the inverse of the elasticity of intertemporal substitution – assumed constant across individuals – and Z is a set of demographic characteristics acting as taste shifters.

The resulting Euler equation shows the consumption evolution over time as a function of the parameters of the utility function, of the intertemporal rate of time preference – which, following Banks, Blundell and Tanner (1998) is allowed to depend on age – and of the market interest rate:

$$\ln C_{i,t} - \ln C_{i,t-1} = k + \frac{1}{\rho} r_t + \beta_1 \Delta Z_{i,t} + \beta_2 age_{it} + e_{i,t} \quad (2)$$

The constant k captures both the (constant across households part of the) discount factor and conditional higher moments of consumption growth and of the interest rate, and $e_{i,t}$ represents all the unexpected news received in year t . Shocks may be related to individual or aggregate factors, such as an unforeseen unemployment spell or an unexpected recession, that cause a revision to lifetime resources and hence to consumption.

As long as retirement and the consequent change in income are expected – and therefore households are not caught unprepared – permanent income should not change and the consumption growth should not be affected at all: to test this implication we add to equation (2) an indicator variable equal to one at the time of retirement:

$$\Delta \ln C_{it} = k + \frac{1}{\rho} r_t + \beta_1 \Delta Z_{it} + \beta_2 age_{it} + \gamma * retire_{it} + e_{it} \quad (3)$$

According to this specification of the model, if retirement is expected, the coefficient γ should be equal to zero.³

If, however, retirement is unexpected – that is, if retirement occurs earlier than expected as a consequence of early dismissal from work or redundancy, for example – then it is accompanied by an unexpected wealth shock, which causes, in turn, a negative revision of consumption (Banks, Blundell and Tanner, 1998). In order to distinguish between expected and unexpected decisions to retire, we use a unique source of information present in our dataset: the expected age of retirement. The data provide information on the expected retirement age gathered at each wave for each working respondent. By using the panel dimension of the data we are also able to compare the actual age of retirement of the newly retired individuals with the expected age they stated in the previous wave, hence determining if their retirement was expected or not. More precisely, by comparing actual and expected retirement ages, we distinguish among four cases:

- i) individuals retiring when expected,
- ii) individuals retiring unexpectedly,
- iii) individuals who expected to retire but did not,

³ The assumptions underlying the model we use rely on the separability of consumption and leisure and do not account for home production or for the effect of precautionary saving arising from income uncertainty.

iv) individuals who did not retire and did not expect to do so.

3. Data

We use ten waves of the SHIW for the period 1991–2010. The survey began in the 1960s with the aim of gathering data on the incomes and savings of Italian households. Over the years, the scope of the survey has grown and it now contains detailed information on Italian households' consumption and household members' demographics, labour supply, income and real and financial wealth.

The data are representative of the Italian resident population and are collected (since 1989) every 2 years⁴. Each wave approximately covers 8,000 households and 50 % of the sample is re-interviewed in order to build up a rotating panel component. The unit of observation is the family, defined as all the persons residing in the same dwelling who are related by blood, marriage, common-law marriage or adoption. Brandolini and Cannari (1994) and Brandolini (1999) describe in detail the set-up of this dataset and its quality.

The availability of data on households' expenditure and on individuals' characteristics allows us to estimate the Euler equation (equation 3) for non-durable consumption. The utility changes associated with the consumption of durables are indeed difficult to measure as households' current utility can depend on service flows from past purchases of these goods.

The data on individuals' expectations about retirement age constitute a special asset for our analysis, as we are able to observe whether retirement was truly expected by individuals. In addition, knowing the educational attainment for each household component allows us to detect whether individuals with a low education level behave differently from individuals with higher educational attainment. Finally, the survey contains information on financial and real wealth which we utilise to explain the heterogeneity of individual behaviour.

To carry out the empirical analysis, we select a sample of households whose head is around the age of retirement (in the age range 50–72). In particular, as we are interested in time variations, we focus on the group of households that are observed for at least 2 consecutive waves. In order to avoid exceptional and unrepresentative situations, we restrict our sample to male-headed households (the vast majority in Italy, about 73% of the sample in the selected age range). Our initial sample includes 18,524 observations. Moreover, we select households whose head is an employee, excluding those who exhibited any time out of the labour market other than retirement (8,079 observations out of the initial 18,524). We also exclude those individuals who returned to work after retirement (36 individuals corresponding to 144 observations) or are always in retirement (117 additional observations dropped), and the individuals for which information about expected retirement age is missing (3863 observations dropped). The final selected sample is 3,955 observations for 1,572 heads of households.

Our main variable of interest is yearly non-durable expenditure, defined on the basis of recalled average monthly expenditure on food and non-durable items, hence excluding expenditure on durables, rents and mortgages⁵. Therefore, non-durable consumption comprises work-related expenditures, expenditures for meals outside home and commuting costs.

Descriptive statistics of the variables of interest for the selected sample are presented in Table 1. The average family spends about 28 thousand euro per year on non-durables. The average consumption growth, given by the first difference of the logarithm of non-durable expenditure, is close to zero. The average age of the heads of households is around 57 while the age at retirement is about 59. In the sample, we observe 489 retiring heads. The dummy variable "retire" is assigned a

⁴ With the exception of the 1997 wave, which was collected in 1998.

⁵ Interviewed people have been asked about the average monthly expenditures for all goods with the exception of: jewellery, vehicles, furniture, electric domestics and other devices, donations, transfers, child support, alimonies, house emergency maintenance services, rent, mortgages, life insurance premiums and pension fund contributions.

value of one in the wave in which the head is observed retiring, or otherwise assigned a value of zero. The overall mean is equal to 0.12, while the percentage of household heads retiring during the sample period is 31%. An individual is considered retired if his response to the current labour market status question is “retired from work”⁶.

The dichotomous variable “low educated” is based on the educational level of the head of household and is assigned a value of one if the head completed at most compulsory education; 55% of the household heads in the selected sample belongs to this group. The average family in our sample is made up of about 3 components, of which 2 are income earners; in many cases the second earner is the spouse, in others a grown-up child. The number of income earners other than the spouse, which is equal to 1.6, and the number of income earners other than the spouse and children, which shrinks to 1.36, are reported in the table. The spouse is present in 3,023 households and working in 40% of household/year observations. Among the spouses, only 82 retire in the sample period (with the relative dummy variable having an overall average equal to 0.03).

Household net wealth is defined as the sum of real wealth (i.e. real assets) and financial assets minus financial liabilities. In our sample of households whose head is close to retirement, the average wealth is about 275 thousand euro. Low wealth individuals are those whose wealth is below the median wealth calculated for each wave.

Table 1 – Descriptive statistics

	Number of Obs.	Mean	Std Dev.
Non-durable expenditure	3,955	27,891	13,052
$\Delta \log$ (non-durable expenditure)	3,955	-0.0007	0.3
Age	3,955	57.1	5.3
Age at retirement	489	59.1	4.0
Retire (dummy)	3,955	0.12	0.3
Low educated (dummy)	3,955	0.55	0.5
Household members	3,955	3.25	1.2
Δ Household members	3,955	-0.16	0.5
Household income recipients	3,955	1.99	0.8
Δ Household income recipients	3,955	0.02	0.7
HH income recipients other than the spouse	3,955	1.58	0.8
Δ HH income recipients other than the spouse	3,955	0.02	0.6
HH income rec. other than spouse and children	3,955	1.36	0.6
Δ HH income rec. other than spouse and children	3,955	0.03	0.5
Number of children working	3,955	0.22	0.5
Δ Number of children working	3,955	-0.03	0.4
Working spouse (dummy)	3,023	0.40	0.5
Spouse retire (dummy)	3,023	0.03	0.2
Wealth (total)	3,955	274,621	560,926

Note: Source: SHIW 1991–2010, pooled sample of 3,955 observations related to 1,572 male household heads observed retiring in the period 1992–2010. Non-durable expenditure and wealth are reported in euro at 2010 prices.⁷

⁶ The variable providing this information is “apqual” in the SHIW questionnaire.

⁷ We use the ISTAT consumer price index for blue- and white-collar worker households (FOI).

The average (log) consumption across the years around retirement for the sample considered in our analysis is illustrated in Figure 1. At wave zero, the household head is experiencing retirement. The consumption behaviour around the wave (wave 0 in Figure 1) when retirement occurs is the focus of our analysis. As a first snapshot of our variable of interest, we observe a general decline in consumption occurring in this part of the life cycle with a more substantial drop at the time of retirement.

Figure 1. Consumption across waves centred at retirement

Note: Negative (positive) waves are associated with waves before (after) retirement occurs.

Table 2 – Effective and expected age of retirement

	Num. of Obs.	All	Education		Wealth	
			Low	High	Low	High
Effective age	489	59.05	58.48	60.07	59.02	59.08
Expected age wave-1	489	59.42	59.00	60.15	59.77	59.04
Expected age wave-2	299	60.01	59.11	61.21	60.27	59.78
Expected age wave-3	163	60.48	60.13	60.85	60.87	60.15

Note: expected age wave-1 is the expected age of retirement as declared one wave prior to retirement; the other variables are defined similarly.

In order to shed some light on the retirement expectations and their evolution, Table 2 reports the average effective retirement age in our sample, as well as the average of the self-predicted age of retirement reported one, two or three waves prior to retirement. As our sample includes individuals observed for at least two waves, the number of observations for the effective age of retirement and the expected age one wave prior to retirement is the same, while going back further, we have fewer observations. The table shows that self-prediction about the age of retirement is more accurate the closer the prediction is made to retirement for the overall sample as well as for the subsamples considered. The F-tests (not shown) reveal that the difference between the effective age and the predictions is always statistically significant at any standard confidence level; only high educated individuals with wealth above the median the wave before retirement predict their retirement age correctly on average. Hence, low educated and/or low wealth individuals appear less able to make accurate predictions about their age of retirement, even one wave prior to the event compared to high education/wealth individuals. This pattern could be the result of shocks (to health, or to employment) hitting with higher probability in these groups or might reflect less ability to forecast the future.

Finally, it is worth noting that when splitting the sample by education, the average age of retirement (both effective and predicted) is statistically different in the two groups with the low educated heads retiring, on average, about one year and a half before high educated individuals. This difference is not present when splitting the sample by wealth.

Based on the information on effective and expected retirement age, we are able to distinguish among expected and unexpected retirements. An unexpected retirement is defined as a retirement taking place in wave t at an age which is lower than expected in wave $t-1$. Table 3 reports the distribution of retiring people by educational and wealth level. Overall, 138 out of 489 retirement episodes (about 28%) are unexpected. The incidence of unexpected retirement is lower for people with high education levels than for those with low education levels (26% and 29% respectively) and for high wealth than for low wealth households (22% and 34% respectively).

Table 3 – Distribution of retiring people by educational level and wealth

	Retire when expected	Retire when not expected	Total
<i>By education:</i>			
low educated	221	92	313
high educated	130	46	176
<i>By wealth levels:</i>			
below the median	167	87	254
above the median	184	51	235
<i>Total</i>	351	138	489

Note: A low educated individual has completed at most the compulsory education. A low wealth individual has wealth below the median level.

4. Results

As discussed in section 2, we write our basic specification for consumption growth as⁸:

$$\Delta \ln C_{it} = \ln C_{it} - \ln C_{it-2} = \alpha + \gamma * retire_{it} + X_{it}\beta + T_t\theta + e_{it} \quad (4)$$

where our dependent variable is the variation in the logarithm of non-durable household expenditures that occurred between two consecutive waves of the survey.

In our baseline specification, the set of regressors includes the dummy variable for retiring in the current wave,⁹ the coefficient γ which measures the consumption drop at retirement, and the time dummy variables T_t , to capture the effect of time-varying interest rates. In addition, we control for a set of individual specific characteristics (X) including: the age of the household head, capturing the changes in the intertemporal rate of time preferences; the educational level of the household head, as it normally shapes individual behaviour; the change in the number of household components, as it is expected to affect consumption growth substantially¹⁰; the change in the number of income earners, to capture the potential labour supply in the household which could act as an insurance device against the income drop at retirement; and the area of residence (north, centre and south), to allow for macro-regional differences in discount rates.

The results of our baseline models are reported in Table 4. Given the estimations from the panel dimension of the data, we compute clustered robust standard errors in order to account for heteroscedasticity and serial correlation of individual errors over time.

For brevity's sake, we provide comments on the coefficient estimates for only the main variables of interest. In the first and simplest specification considered (see column 1 of Table 4), the OLS coefficient on the retire dummy variable reveals an average consumption drop of about 4% at retirement.

However, as shown in column 2, retiring does not have the same impact across households. High educated people seem not to be significantly affected by retiring. Highly educated individuals – here defined as individuals with more than compulsory education – actually show a drop in consumption of about 0.7 percentage points, not statistically different from zero. Conversely, people with a lower level of education are characterised by a significant decrease in consumption of 6 percentage points. The two groups statistically differ, albeit marginally (at 11%). One possible interpretation of these results could be tied to home production which could be more pronounced among less educated people. Different educational levels are likely to be associated with different preferences for leisure and consumption. In particular, the low educated are likely to better know how to produce a variety of goods at home and finally have the time to do so during retirement; therefore they are able to reduce their consumption of market goods (Hurd and Rohwedder, 2003).

Alternatively, it could be argued that education acts as proxy wealth, with low educated individuals having low wealth being more likely to drop consumption at retirement if hit by a negative shock. However, in our sample 70% of the highly educated individuals accumulated wealth above the median (almost 200,000 euro), as opposed to 40% of individuals with low education levels. Hence,

⁸ See Banks, Blundell and Tanner (1998) for a similar specification.

⁹ As there is a two-year lag between each wave, an individual observed retiring in wave t could actually have retired one year previously.

¹⁰ See Attanasio and Weber (1995). Alternatively, as in Miniaci et al (2010), we could have attributed shares of the household consumption to the household members according to an equivalence scale and we could have regressed the individual consumption on each household member's characteristics. However, this would have been at the cost of introducing a substantial measurement error.

we actually find that education cannot be considered a good proxy for wealth. To test whether wealthier individuals manage to smooth consumption more easily, we run a regression disentangling the retiring population into two groups: those with wealth above or below the median of the distribution (see column 3).

We find that the consumption drop only characterises the group of individuals whose wealth is below the median, while those with wealth above the median exhibit smooth consumption between work and retirement. For the former group, the coefficient of the dummy variable retire in the estimated equation is -5.3%, significant at the .05 level. For the latter, it is about -2.5% and not statistically different from zero. However, the two groups are not statistically different. Differences between wealth groups do not seem to be responsible for consistent differences in behaviour.

Our results lead us to explore the role of education interacted with wealth at retirement (column 4, table 4). The results show that education has an impact only within the low wealth group where high and low educated households significantly differ in their consumption drop (at 6% level).

Table 4 – OLS by alternative specifications

	(1)	(2)	(3)	(4)
Retire	-0.0401** (0.0165)			
Retire*low educated		-0.0595*** (0.0207)		
Retire*high educated		-0.0068 (0.0259)		
Retire*low wealth			-0.0534** (0.0223)	
Retire*high wealth			-0.0254 (0.0229)	
Retire*low educated*low wealth				-0.0746*** (0.0252)
Retire*low educated*high wealth				-0.0334 (0.0323)
Retire*high educated*low wealth				0.0195 (0.0435)
Retire*high educated*high wealth				-0.0193 (0.0306)
Low educated	-0.0017 (0.0082)	0.0044 (0.0094)	-0.0018 (0.0082)	0.0039 (0.0095)
Age	-0.0033*** (0.0011)	-0.0034*** (0.0011)	-0.0033*** (0.0011)	-0.0034*** (0.0011)
HH members (Δ)	0.0427*** (0.0116)	0.0429*** (0.0116)	0.0426*** (0.0116)	0.0430*** (0.0117)
HH recipients (Δ)	0.0609*** (0.0098)	0.0606*** (0.0098)	0.0610*** (0.0098)	0.0606*** (0.0098)
Always retired	0.0049 (0.0140)	0.0048 (0.0140)	0.0051 (0.0140)	0.0045 (0.0141)
High wealth	0.0266*** (0.0096)	0.0266*** (0.0096)	0.0230** (0.0103)	0.0249** (0.0105)
Constant	0.0829 (0.0634)	0.0850 (0.0634)	0.0849 (0.0635)	0.0832 (0.0637)
R-squared	0.045	0.045	0.045	0.045
N	3955	3955	3955	3955
F-test of joint significance (P-value)		0.017	0.023	0.044

Note: ***1% significance level; **5% significance level; *10% significance level.

Clustered standard errors. Standard errors in parentheses. In the last row, we report the p-value of the F-test of joint significance of the interactions of the variable “retire”.

After investigating the consumption drop for various subgroups of the population, we analyse whether expectations about retirement play a role: retirement can be unexpected and, if so, it is likely to be associated with a wealth shock having sizeable effects on consumption. Permanent

income, the main determinant of current consumption, is negatively affected by earlier retirement, as retirement coincides with a flow of pensions that are lower than labour incomes. In addition, unexpected retirement is associated with a lower than expected replacement rate, further exacerbating the income drop when retired. This leads to an estimation of the consumption drop at retirement that for some individuals mixes up the effects of the unexpected income shock as a consequence of retirement with all the other potential explanations for the reduction in consumption at retirement, namely preferences, non-separability between consumption and leisure, etc.

As discussed in section 2, we are interested in unexpected retirement and we use the information on the expected age of retirement to create a dummy coded variable. At each time t , we observe whether an individual is currently retired and whether he expected to be so. A value of 1 corresponds to an individual in wave t who just retired and stated in the previous wave that he expected to do so. This allows us to examine our sample according to the unexpectedness of retirement which could potentially drive the consumption drop. By creating an interaction between these two dichotomous variables, we obtain four dummy variables to cover the four possible cases: (i) in period t an individual is not retired and did not expect to be retired; (ii) in period t an individual is not retired but expected to be retired; (iii) in period t an individual is retired but did not expect to be retired and (iv) in period t an individual is retired and expected to retire. Case iii is labelled ‘retired when not expected’ and case iv is labelled ‘retired when expected’. In our sample, about 72% of the individuals who retire correctly predict retirement. Case i is our base case and we exclude the corresponding dummy variable from the regressions.

Table 5 reports the core results of our paper which include the predictions among the regressors. Correctly expecting retirement should make people behave under optimal conditions, while an unexpected retirement, or shock, should be more likely to change consumption behaviour to adapt to the new environment. We would thus expect no consumption drop for respondents who retired and correctly predicted it. Our baseline specification (column 1) shows instead that both types of households reduced their consumption. Those not affected by a surprise reduced their consumption by about 4%; the same result is obtained for surprised households and the two coefficients do not significantly differ. Such a finding sheds important light on the willingness to drop consumption at retirement. Households are in fact aware that retirement is happening and they reduce their consumption when retirement occurs voluntarily. On average, retirement, whether unexpected or not, does not relate to drop in consumption. In other words, a surprise retirement is not driving the average consumption drop.

In order to explore, with greater detail, the characteristics that may be driving this drop, we further examine our retirement variable by the surprise intensity degree and the educational level of the head of the household, as reported in column 2 and 3, respectively.

In column 2, we use the time lag between the stated retirement age and the actual one by assigning a value of one to the dummy variable if the lag is longer than two years (variable “Retire when not expected (more than two years)”). Results do not show any significant relationships with surprise intensity, but actually show a reduced significant consumption drop for all surprised respondents regardless of the intensity of the surprise.

The education of the household head drives important differences in consumption behaviour. As shown in column 3, consumption drops at retirement only for low educated respondents, irrespective of whether retirement is expected or not. Less educated people smooth their consumption less than more educated peers, albeit the difference in the coefficients is not statistically different.

Table 5 – OLS by alternative specifications disentangling expected and unexpected retirement

	(1)	(2)	(3)	(4)
Retire when expected	-0.0429** (0.0195)	-0.0430** (0.0195)		
Retire when not expected	-0.0426* (0.0258)	-0.0534 (0.0430)		
Retire when not expected (more than two years)		0.0182 (0.0522)		
Not retired when expected	-0.0209 (0.0255)	-0.0210 (0.0255)		
Retire when expected*low education			-0.0613** (0.0244)	
Retire when not expected*low education			-0.0573* (0.0331)	
Not retired when expected*low education			0.0175 (0.0337)	
Retire when expected*high education			-0.0120 (0.0309)	
Retire when not expected*high education			-0.0139 (0.0397)	
Not retired when expected*high education			-0.0705* (0.0379)	
Retire when expected*low wealth				-0.0448 (0.0285)
Retire when not expected*low wealth				-0.0768** (0.0309)
Not retired when expected*low wealth				-0.0171 (0.0334)
Retire when expected*high wealth				-0.0407 (0.0251)
Retire when not expected*high wealth				0.0153 (0.0446)
Not retired when expected*high wealth				-0.0235 (0.0387)
Low education	-0.0011 (0.0082)	-0.0012 (0.0082)	0.0001 (0.0096)	-0.0014 (0.0082)
Age	-0.0031*** (0.0011)	-0.0031*** (0.0011)	-0.0031*** (0.0011)	-0.0031*** (0.0011)
Hh members (Δ)	0.0426*** (0.0116)	0.0426*** (0.0116)	0.0428*** (0.0116)	0.0426*** (0.0116)
HH recipients (Δ)	0.0609*** (0.0098)	0.0609*** (0.0098)	0.0611*** (0.0098)	0.0613*** (0.0098)
Always retired	0.0012 (0.0141)	0.0010 (0.0141)	0.0007 (0.0141)	0.0014 (0.0141)
High wealth	0.0271*** (0.0095)	0.0270*** (0.0095)	0.0270*** (0.0095)	0.0239** (0.0105)
Constant	0.0717 (0.0629)	0.0704 (0.0629)	0.0690 (0.0631)	0.0730 (0.0629)
R-squared	0.045	0.045	0.045	0.045
N	3955	3955	3955	3955
F-test of joint significance (P-value)	0.02	0.04	0.057	0.027

Note: ***1% significance level; **5% significance level; *10% significance level. Clustered standard errors. Standard errors in parentheses. In the last row, we report the p-value of the F-test of joint significance of the interactions with the variable “retire”.

We then break down our sample of retiring individuals according to whether their wealth at retirement is above or below the median (column 4). We find evidence of a consumption drop at retirement significantly different from zero only for low wealth individuals. The consumption drop is sizeable, equal to about 7.7% if retirement was unexpected, while the drop is around 4% when expected, albeit not statistically significant (the p-value being equal to 42%). On the other hand, for high wealth individuals, we find no significant consumption drop irrespective of whether retirement is expected or not.

In summary, we find that low educated individuals, on average, drop consumption at retirement, while high educated individuals do not, regardless of expectedness. When we consider wealth, we find evidence that low wealth individuals drop consumption regardless of whether retirement is expected or not, while high wealth individuals do not drop consumption.

To investigate these results further and to shed more light on the effect of education and wealth, we break down our sample using interactions between wealth and education by the degree of preparedness for retirement, as shown in Table 6.

First, consumption still drops among low educated respondents—irrespective of whether their retirement was unexpected or not, albeit only for a restricted sub-sample: those who are not well-off. Their consumption drops by about 8%.

Second, among low wealth households who retire when expected, the highly educated do not drop consumption while the low educated do reduce consumption by about 7%: the consumption drop of the low educated and high educated is significantly different (the p-value of the F-test is 3.2%)

Third, highly educated individuals with low wealth drop their consumption by 9.7% when retirement is unexpected. On the other hand, highly educated individuals with low wealth who retire when expected do not drop their consumption (in this case the two coefficients are different at the 3% level).

We conclude that differences across wealth levels alone do not consistently shape consumption growth at retirement. Rather, it is within wealth categories where we find the most interesting differences. Particularly, within the low wealth group, the interactions between education and expectedness are responsible for statistically significant differences in consumption drop at retirement.

These results seem to indicate that highly educated people drop their consumption as a last resort. This suggests that more educated individuals are likely to be better planners and thus more oriented to smooth their consumption, even when retirement occurs. Only under unforeseen circumstances, such as lack of sufficient buffer stock when a shock occurs, would they abandon their planned behaviour. Indeed, educated households exhibit a drop in their consumption when retirement comes as a shock and when they have little wealth. When retirement is expected, regardless wealth, highly educated respondents seem to smooth their consumption. Conversely, low educated people with little wealth drop their consumption equally, irrespective of whether they are expecting to retire or not. In other words, low educated people with low wealth seem not to be interested in planning consumption smoothing.

Table 6 – OLS by alternative specifications interacting expected retirement, education and wealth

	(1)
Retire when expected*low education*low wealth	-0.0732** (0.0316)
Retire when not expected* low education*low wealth	-0.0799** (0.0362)
Not retired when expected* low education*low wealth	-0.0066 (0.0412)
Retire when expected*low education*high wealth	-0.0412 (0.0347)
Retire when not expected* low education*high wealth	-0.0014 (0.0693)
Not retired when expected* low education*high wealth	0.0542 (0.0611)
Retire when expected*high education*low wealth	0.0659 (0.0580)
Retire when not expected* high education*low wealth	-0.0969* (0.0514)
Not retired when expected* high education*low wealth	-0.0456 (0.0507)
Retire when expected*high education*high wealth	-0.0387 (0.0347)
Retire when not expected* high education*high wealth	0.0359 (0.0528)
Not retired when expected* high education*high wealth	-0.0647 (0.0464)
Low education	-0.0014 (0.0096)
High wealth	0.0235** (0.0106)
R-squared	0.045
N	3955

Note: ***1% significance level; **5% significance level; *10% significance level.

Clustered standard errors. Standard errors in parentheses. The regression includes all the controls as in table 4 and 5.

5. Sensitivity analysis

As a robustness check, we tried several other specifications. We particularly concentrate on the possible role of labour supply of the spouse and of other members of the household. If retirement surprised the household, the entity of the shock might be less pronounced if other persons in the household work. Reducing consumption might be less of a necessity when there are other family members who can bring a salary to the household.

Particularly when the shock is unexpected we expect that having a working family member can act as a smoothing consumption device. We first check whether consumption drop at retirement is mitigated by the presence of a working spouse (Table 7, column 1) by adding an interaction variable: the product of the dummy variable “working spouse in the family” with the dummy variable retiring with or without surprise. Despite the hypothesised direction, while retirees do drop their consumption less when they have a working spouse to rely on, the results are not statistically significant. Hence, adding the dimension of a working partner does not change the prediction or the

sign of consumption drop, which persists in the household, irrespective of whether the spouse works.

We also try adding working children living with their parents (column 2) in order to check whether a similar relationship can be found. However, we find no significant evidence that having working children act as a smoothing device.

Table 7 – OLS by alternative specifications, adding working spouse and children

	(1) b/se	(2) b/se
Retire when expected	-0.0475** (0.0227)	-0.0294 (0.0217)
Retire when not expected	-0.0502* (0.0301)	-0.0396 (0.0260)
Not Retired when expected	-0.0122 (0.0336)	-0.0116 (0.0279)
Retirement when expected * Working spouse	0.0199 (0.0441)	
Retirement when not expected* Working spouse	0.0372 (0.0563)	
Not Retired when expected* Working spouse	-0.0257 (0.0494)	
Working spouse	0.0056 (0.0105)	
Working children		0.0117 (0.0148)
Retire when expected * Working children		-0.0551 (0.0464)
Retire when not expected * Working children		-0.0148 (0.0766)
Not Retired when expected * Working children		-0.0413 (0.0636)
Low Education	-0.0005 (0.0082)	-0.0017 (0.0082)
Age	-0.0030*** (0.0011)	-0.0032*** (0.0011)
Hh members (Δ)	0.0424*** (0.0116)	0.0427*** (0.0117)
HH recipients (Δ)	0.0608*** (0.0098)	0.0604*** (0.0101)
Always retired	0.0016 (0.0143)	0.0017 (0.0142)
High wealth	0.0262*** (0.0098)	0.0271*** (0.0095)
Constant	0.0636 (0.0638)	0.0728 (0.0632)
R-squared	0.044	0.044
N	3955	3955

Note: ***1% significance level; **5% significance level; *10% significance level.

Clustered standard errors. Standard errors in parentheses. The regression includes all the controls as in table 4 and 5.

6. Conclusions

In this paper, we utilise the panel dimension of the SHIW dataset to estimate the size of the consumption drop at retirement in Italy. Our results show that, on average, the non-durable consumption drop at retirement in Italy is about 4%. This finding is in line with previous research on Italian data (Battistin et al., 2009; Miniaci, Monfardini and Weber, 2010).

We then use information on the expected retirement age to distinguish among households whose head retires when expected from those for whom retirement comes as a surprise. Retiring when expected should make people behave under the optimal conditions. Conversely, households affected by a negative shock at retirement, retiring earlier than expected, should be more likely to change their consumption behaviour to adapt to the new environment. Hence, we expect a bigger drop in consumption for respondents who retired unexpectedly. Our baseline specification shows instead that both types of households reduced their consumption at retirement of about 4%. Such a finding sheds important light on the consumption drop at retirement. On average, whether retirement is unexpected or not does not affect the respondents. In other words, the surprise in retirement is not driving the average consumption drop.

We then investigate whether differences in education and/or wealth affect our results. We find that low educated individuals, on average, drop consumption at retirement, while high educated individuals do not, regardless of expectedness. When we consider wealth, we find evidence that low wealth individuals drop consumption, whether retirement is expected or not, while high wealth individuals do not.

When we break down our sample using interactions between wealth and education given the degree of preparedness for retirement, we find the most interesting results. The results show that low educated respondents – irrespective of whether their retirement was unexpected or not – drop their consumption but only for those with low wealth. Their consumption drops by almost 8%.

In addition, we find that a subsample of highly educated individuals drop their consumption, by almost 10%, when retirement is unexpected only among those with low wealth. On the other hand, high educated households with low wealth who retire when expected do not drop their consumption.

These results suggest that highly educated people drop their consumption as a last resort. In fact, this occurs when retirement comes as a shock and they possess little wealth. When retirement is expected, irrespective of their wealth, highly educated respondents smooth their consumption according to their plans. Conversely, low educated people, with little wealth, drop their consumption in the case of both expected and unexpected retirement.

Acknowledgements

We wish to thank two anonymous referees, Rob Alessie, Ainhoa Aparicio, Adrian Kalwij, David Laibson, Raffele Miniaci, Nicola Pavoni and participants at the International Pension Workshop NETSPAR (Turin, 2011), the SIE (Società Italiana degli Economisti) 2011 meeting, the Economics Department of the University of Verona seminar series and the 2nd International Workshop on the Socio-Economics of Ageing (Lisbon, 2011) for their useful comments. We are grateful to AGING – Regione Piemonte and MIUR for their financial support.

References

- Aguiar, M. and Hurst, E. (2005) Consumption versus expenditure. *Journal of Political Economy*, 113 (5), 919-948.
- Aguiar, M. and Hurst, E. (2008) *Deconstructing Lifecycle Expenditure*. National Bureau of Economic Research, Inc. NBER Working Papers 13893.
- Ameriks, J., Caplin, A. and Leahy, J. (2003) Wealth accumulation and the propensity to plan. *The Quarterly Journal of Economics*, MIT Press, 118 (3) (August), 1007-1047.
- Angeletos, G., Laibson, D., Repetto, A., Tobacman, J. and Weinberg, S. (2001) The hyperbolic consumption model: Calibration, simulation, and empirical evaluation. *Journal of Economic Perspectives*, 15 (3), 47-68.
- Attanasio, O. P. and Weber G., (1995) Is consumption growth consistent with intertemporal optimization? Evidence from the Consumer Expenditure Survey. *Journal of Political Economy*, University of Chicago Press, 103 (6), 1121-57.
- Banks, J., Blundell, R. and Tanner, S. (1998) Is there a retirement-savings puzzle? *The American Economic Review*, 88 (4), 769-788.
- Battistin, E., Brugiavini, A., Rettore, E. and Weber, G. (2009) The retirement consumption puzzle: Evidence from a regression discontinuity approach. *The American Economic Review*, American Economic Association, 99 (5), 2209-2226.
- Bernheim, D., Skinner, J. and Weinberg, S. (2001) What accounts for the variation in retirement wealth among U.S. households? *The American Economic Review*, American Economic Association, 91 (4) (September), 832-857.
- Blau, D. M. (2008) Retirement and consumption in a life cycle model. *Journal of Labor Economics*, University of Chicago Press, 26, 35-71.
- Brandolini, A. (1999) The distribution of personal income in post-war Italy: Source description, data quality, and the time pattern of income inequality. *Giornale degli Economisti e Annali di Economia*, 58, 183-239.
- Brandolini, A. and Cannari, L. (1994) Methodological appendix: The Bank of Italy's survey of household income and wealth. In: Ando, A., Guiso, L. and Visco, I. (eds.) *Saving and the Accumulation of Wealth. Essays on Italian Household and Government Saving Behavior*. Cambridge, Cambridge University Press.
- French, E. (2005) The effects of health, wealth, and wages on labor supply and retirement behavior. *Review of Economic Studies*, 72 (2), 395-427.
- Gustman, A. L. and Steinmeier, T. L. (2001) *Imperfect Knowledge, Retirement and Saving*. Michigan Retirement Research Center Research Paper No. WP 2001-012.
- Haider, S. J. and Stephens, M. (2007) Is there a retirement consumption puzzle? Evidence from subjective retirement expectations. *The Review of Economics and Statistics*, 89 (2), 247-264.
- Hurd, M. and Rohwedder, S. (2003) *The Retirement-Consumption Puzzle: Anticipated and Actual Declines in Spending at Retirement*. NBER Working Papers 9586.
- Lusardi, A. (1999) Information, expectations, and savings. In: Aaron, H. (ed.) *Behavioral Dimensions of Retirement Economics*. New York, Brookings Institution/Russell Sage Foundation, ch. 3.
- Lusardi, A. (2000) *Explaining Why So Many Households Do Not Save*. Working Paper, University of Chicago, Harris School of Public Policy.

Miniaci, R., Monfardini, C. and Weber, G. (2010) Is there a retirement puzzle in Italy? *Empirical Economics*, 38, 257-280.

Smith, S. (2006) The retirement consumption puzzle and involuntary early retirement: Evidence from the British Household Panel Survey. *The Economic Journal*, 116 (March), C130-C148.