Revealed preference and consumption behaviour at retirement

Peter Levell

Revealed preference and consumption behaviour at retirement*

Peter Levell

September 13, 2014

Abstract

This paper sets out revealed preference tests for different models of consumption behaviour over retirement that we applied to a Spanish consumption panel dataset. We reject the perfect foresight model both with separable preferences and allowing for preference change. The first order conditions for the life-cycle model allowing for uncertainty do not provide very strong restrictions on possible choices. In fact they are no stronger than those implied by the most basic revealed preference requirement: GARP. We then go on to investigate the patterns of deviations from a perfectly smoothed marginal utility of wealth and ask whether they fit the predictions of the life-cycle model. We find a tendency of these to increase over time, suggesting consumption falls more than we'd expect. After considering various possible explanations, we settle on non-rational behaviour as the most plausible.

1 Introduction

Economists have long been debating questions surrounding the consumption behaviour of households at retirement (see for instance Banks et al., 1998, Bernheim et al., 2001, Smith, 2006, Aguiar and Hurst, 2005, Luengo-Prado and Sevilla, 2013, etc.). This is to be expected as these are questions that are not only important but also difficult to answer. They are important as the behaviour of households at this important point in the life-cycle can help shed light both on significant policy questions (such as those concerning the adequacy of household's retirement savings), as well as on theoretical questions (such as the appropriateness of the commonly used life-cycle model). They are difficult to answer because changes in consumption at this time could be affected by coincident changes in a whole range of other variables. If preferences over consumption and leisure are non-separable, then we would expect the large change in labour supply that occurs when an individual retires to affect the allocation of consumption both across time and across periods. Retirement itself may also be associated with health shocks, changes in family size, or news concerning future earnings prospects - all of which could affect a household's consumption choices.

Significant progress has been made in addressing these issues, but many of the studies that have looked into these questions have made use of parametric methods which are limited in the questions they can answer. In this paper, we propose an alternative non-parametric approach, in the form of various revealed preference tests for different versions of life-cycle model, which we apply to households from a Spanish consumption panel over the period of retirement. This approach does not make ad-hoc functional form assumptions, and allows for complete heterogeneity in preferences across households. We also contribute to the revealed preference literature by allowing consumers to face uncertainty (which is frequently omitted in revealed preference tests), and considering the testable implications of this.

In what follows, we will set out revealed preference tests of the predictions of the life-cycle model where preferences over consumption goods are independent of labour force participation allowing for general uncertainty, perfect foresight, and a single revision in the marginal utility of wealth at retirement. We also set out a test of life-cycle model when we allow for preferences over consumption goods to change as the household retires (to be non-seperable with labour force participation) assuming perfect foresight in the periods before and after retirement. We find that only the weakest of these hypotheses, the basic life-cycle model with seperable preferences and uncertainty, can successfully rationalise the behaviour of a large number of households. However, the revealed preference test of this, which is derived from first order conditions for optimality and tested household by household, provides extremely weak evidence on the success of this model. Noting that this test does not exhaust the empirical implications of the life-cycle model, we then look at the pattern of changes in the marginal utility of wealth that would rationalise the data. Under the life-cycle model these are predicted to be martingale, but we find a systematic tendency for the marginal utility of wealth to increase over time (and

hence for consumption to fall too quickly relative to what we would expect). While such a result is possible to rationalise under the life-cycle model, we view non-rational behaviour (in particular time inconsistency) as a more likely explanation.

The plan of this paper is as follows. Section 2 describes the life-cycle model, parametric approaches, their findings and their various limitations. Section 3 then sets out the form of our battery of revealed preference tests. Section 4 describes the data we will be using. Section 5 presents our results. Section 6 discusses further empirical implications of the life-cycle model when we allow for uncertainty. Section 7 concludes.

2 The life-cycle model

At its most basic, economic theory predicts that consumers facing an intertemporal problem will choose the best allocation of goods over time that they can afford. We can represent the process of the consumer 'choosing the best' by having him or her maximise some lifetime utility function, that will depend on both how much and when individual goods are consumed. The aim of the researcher is to recover as much information as possible about this utility function using observations on consumer choices in different situations. This is typically done by first specifying some, parametric, functional form for the utility function which is then fitted to the data in some way. This framework allows us to carry out simple, statistical testing of important hypotheses such as consumption smoothing, separability restrictions between leisure and individual consumption demands, hyperbolic discounting and others.

A typical approach starts with a lifetime utility function that is separable across time, has a geometric discount factor and incorporates uncertainty by having the consumer maximise the *expected* sum of within-period utilities. That is the consumer will solve

$$\max_{q_1..q_T} U = E_t \left[\sum_{t=1}^{T} \beta^{t-1} u(q_t) \right] \tag{1}$$

subject to the squence of flow constraints

$$\Omega_{t+1} = \Omega_t + \omega_t l_t - \rho_t' q_t$$

where β is a discount factor, q_t is a vector of consumption commodities, l_t is labour suppl, ρ_t is a vector of discounted prices for consumption commodities, and ω_t is the discounted wage. Ω_t, ρ_t and ω_t are discounted values of actual wealth W_t, prices p_t, and wages w_t using the discount factor

$$1/(1 + r_t)(1 + r_{t-1}) \ldots (1 + r_2)$$

where r_t is the (nominal) interest rate. These models also typically propose that expectations are *rational* - that is they incorporate all information available at time t, are on average correct and are revised each period through a process of Bayesian updating. This is of course a hypothesis that can be tested like any other. We also assume that individuals can borrow or lend freely. Taken together, these assumptions form what we will call the life-cycle model.

Each period the consumer produces a new plan according to their current information set. Writing out the Lagrangian L_t of the problem in (1) being solved with period t information

$$L_t = \beta^{t-1} u(q_t) + E_t \left[\sum_{\tau=t+1}^{T} \beta^{\tau-t-1} u(q_\tau) \right] - \sum_{\tau=t}^{T} \lambda_t^\tau (\Omega_\tau + \rho_\tau' q_\tau - \Omega_{\tau+1}) - \sum_{\tau=t}^{T} \xi_\tau q_\tau$$

where the leads us the following vector of first order conditions for consumption in periods t and $t+1$

$$u'(q_t) \leq \frac{\lambda_t^t}{\beta^{t-1}} \rho_t \tag{2}$$

$$u'(q_{t+1}) \leq E_t \left[\frac{\lambda_t^{t+1}}{\beta^t} \rho_{t+1} \right] \tag{3}$$

This is an inequality rather than an equality as there may be zero consumption of some commodities in some periods. Here λ_t^τ is the Lagrange multiplier on the budget constraint for period τ when the problem is being solved in period t. This term can be interpreted as the (discounted) marginal utility of wealth: the utility value of relaxing the flow constraint period t. For simplicity we will denote λ_t^t (the λ term for the within period first order condition) as λ_t for the remainder of this paper. Similarly, λ_{t+1} will refer to the lagrange multiplier on the time $t+1$ constraint when the consumer is solving their lifetime optimsation problem in period $t+1$ (i.e. with period $t+1$ information). ξ_t is a vector of Kuhn-Tucker multipliers on the consumer's non-negativity constraints. Further manipulation of these conditions yields Euler conditions governing growth rates of consumption over time for each good i

$$u'(q_t^i) \geq \beta E_t \left[\frac{\rho_t^i}{\rho_{t+1}^i} u'(q_{t+1}^i) \right] \ \forall i \tag{4}$$

2.1 The life-cycle model and retirement

The life-cycle model makes a number of predictions about behaviour. The most important of these for individuals who are retiring is that consumer's consumption plans should only be revised from one period to the next if the consumer receives new information. Forseeable changes to income should not affect the consumer's optimal plan. When consumption is influenced by predictable changes in income this is referred to as 'excess sensivity'.

A potential failure of the life-cycle model around retirement - known as the 'retirement consumption puzzle' - is the tendency of consumption expenditures to fall at the point of retirement. This phenomenon was observed in the UK by Banks et al. (1998) using household expenditure data from the UK, and it has since been observed in a number of other countries, including Italy (Battistin et al., 2009) and the United States (Aguiar and Hurst, 2007). The retirement consumption puzzle is 'puzzling' as the retirement event should be largely forseeable for many consumers, and so it should not result in a failure of consumption smoothing (a prediction of some variants of the life-cycle model). Some (e.g. Bernheim et al., 2001) have argued that it represents evidence of excess sensitivity of consumption to predictable income changes, which would challenge the life-cycle view that expectations are formed rationally. This would be troubling as it raises the possibility that households may not be saving adequately to provide for their old-age out of myopia or some behavioural bias.

Various papers have however proposed or defended alternative explanations for the puzzle that are consistent with the basic life-cycle model. Smith (2006) observes that in UK data food spending only falls for those who retire involuntarily - due to poor health or job loss. This suggests that falls in consumption may be due to negative wealth shocks associated with the retirement decision. There is also evidence to suggest that the declines in expenditure are mainly in food and work-related expenses (Bernheim et al. 2001, Aguila et al. 2011 and Battistin et al. 2009). A decline in work-related expenditures at the time of retirement would seem better explained by change in preferences over consumption goods following retirement (non-separabilites between consumption and leisure in the utility function) than by excess sensitivity. By this we do not mean that consumer has a utility function before and after retirement. Rather we mean that *preferences over consumption commodities* will depend on the amount of leisure the consumer enjoys, and that these will change at retirement. Aguiar and Hurst (2005, 2007) also show that in the United States there is good reason to believe that a reduction in food expenditures can be explained by an increase in home production of food and an increase in time spent shopping around (allowing retired households to pay lower prices for the same goods). Summarising these findings, Hurst (2008) deemed them sufficient to declare the "retirement of a consumption puzzle".[1]

Of special relevance to this paper is that both Luengo-Prado and Sevilla (2013) and Christensen (2008), find no evidence of a retirement consumption puzzle in the Spanish data covering the late 80s and early 90s that we will be using. Moreover neither paper finds any evidence for change in preferences over consumption goods. Christensen points out that income remains relatively stable over retirement in Spain, and argues that this means that any changes in the allocation of spending can be attributed solely to non-separabilities between consumption and leisure. However, she finds no decline in any spending group except for medical expenses (which are heavily subsidised for Spanish retirees). Somewhat surprisingly, this includes a specially constructed category of work-related expenditures (clothing, transportation, petrol and food out). Luenga-Prado and Sevilla similarly find no decrease in expenditures on any of the sub-categories they define, though they do find evidence for a decline in food spending in a successor survey covering later years. They argue with appeal to social attitude surveys that the difference between the two periods is linked to the development of more egalitarian attitudes towards women. The retirement of the (usually male) household head did not lead to an increase in home production in the earlier years as this was more often considered the role of the wife.

2.2 Parametric approaches

A common way of testing the predictions of the life-cycle model is to specify a parametric model of the utility function and to then take the Euler conditions in (4) to data. Studies looking in consumption behaviour will typically make a few additional assumptions. First we can group individual, non-durable, commodities into an aggregate c_t, and then we can assume that all consumers have within-period utility functions of the power form

$$u(c_t) = \exp(X_t' \delta) \frac{c_t^{1-\gamma}}{1-\gamma}$$

---

[1] As Hurst (2008) points out these findings do not necessarily imply that households save sufficiently to provide for an adequete retirement, only that there is no evidence that this is what lies behind a fall in expenditure at the point of retirement.

where X_t is a set of individual demographic and other characteristics (that may include terms such as employment status and hours worked to capture non-separabilities with leisure) and γ is the coefficient of relative risk aversion. Consumption will never be negative so long as the consumer has access to some resources whether in the current or future periods (as this is associated with an infinite marginal utility) so in this case we get the following Euler *equation*

$$c_t^{1-\gamma} = E_t \left[\beta(1 + r_t) \exp((X'_{t+1} - X_t)\delta)c_{t+1}^{1-\gamma} \right]$$

which can be 'log-linearised' to give

$$\frac{1}{\gamma}\Delta \ln c_{t+1} = \ln \beta + \ln(1 + r_t) + \Delta X'_t\delta + u_t \tag{5}$$

where u_t is a residual that captures individual heterogeneity, preference shocks and any expectational errors (see Carroll, 2005, for an explanation and a discussion of some of the problems involved in doing this). Suitably instrumented for the endogeneity of the interest rate (u_t will be higher when r_t is higher than the consumer expected), this model can be used to test for evidence of 'excess sensitivity of consumption to predictable changes in circumstances by including variables that should be in the consumer's information set on the right-hand side and checking for a positive coefficient as well as to estimate important preference parameters. An equation of the form in (5) has been used to test hypotheses about consumption behaviour at retirement in particular in a whole range of studies (Banks et al., 1998, Bernheim et al., 2001, Smith, 2006, Aguiar and Hurst, 2005, Luengo-Prado and Sevilla, 2013 among others).[2]

2.3 Limitations of parametric approaches

As we have seen, the standard parametric model makes a number of quite important assumptions. Firstly, the researcher must specify a functional form for within-period utility that will essentially be ad-hoc. Secondly, all consumers are assumed to have the same preferences with heterogeneity allowed in quite a restrictive manner. For instance, in power utility models, all consumers will have the same (constant) rate of relative risk aversion - a factor determining the curvature of the utility function. Thirdly, we assume that consumers' behaviour can be rationalised by an appropriate utility function at all.

The findings on consumption behaviour at retirement summarised above do not depend heavily on parametric assumptions: studies that look at the timing of changes in expenditure are not necessarily committed to a particular functional form for consumers' utility functions. However, parametric assumptions make it difficult to investigate important other, more general, questions about consumption behaviour. Firstly, we cannot use a parametric model to answer whether any individual consumer's choices can be rationalised by a stable utility function. Since this is precisely the claim that the retirement consumption puzzle is supposed to challenge, this is however certainly a question that is worthy of further investigation. Secondly, parametric models are not terribly useful for investigating questions about what is happening to the marginal utility of wealth. As we discuss below, a set of observations for a given individual may be rationalised by *both* an increasing or a decreasing path for the marginal utility of wealth over time - depending on the curvature of the utility function and size of the discount factor. This means that if we want to investigate the path of the marginal utility of wealth, we need to be open-minded with respect to possible function forms. A parametric approach specifies a functional form *a priori*.

Answering questions such as these will therefore require a nonparametric method. In an ideal world we would like to estimate demands flexibly, individual-by-individual, through an equation such as

$$q_t^i = m(q_t^{-i}, \rho_t, e_t)$$

where q_t^i is one element of the vector q_t, q_t^{-i} represents all the other elements, e_t represents temporal shocks, and m is some (potentially) non-linear function. We could then reject the hypothesis that a stable utility function rationalised the data if the resulting demands were not integrable (an hypothesis that we can test statistically).[3] If demands were integrable, we would then have a good deal of knowledge about the consumer's utility function and the possible paths of shocks they experienced.

Unfortunately such methods are typically highly data intensive (they will be subject to the curse of dimensionality) and so cannot by applied in most datasets. An alternative, *non-statistical,* approach set out by Varian (1982) and in the context of the life-cycle model by Browning (1989) is the method of revealed preference. As we will show, this method can be taken to data covering only a few periods but will nonetheless allow us to contribute to the literature on consumption behaviour around retirement in a number of ways. In particular, it will allow us to directly test whether choices over time can be rationalised by *any* utility function. Secondly, as we will show the revealed preference approach will also allow us to directly investigate whether there is evidence for preference change when leisure time increases at the point of retirement

---

[2]Sometimes consumption is separated into different components, and sometimes the interest rate is omitted.

[3]For instance, Haag, Hoderlein and Pendakur (2009) show how to impose and test Slustky symmetry in a nonparametric demand system.

(that preferences between consumption and leisure are non-separable). Thirdly, the revealed preference procedure can be used to investigate different hypotheses regarding the evolution of the marginal utility of wealth.

3   Revealed preference

The revealed preference approach starts with the vector of first order conditions from the basic life-cycle model given by (2). A consumer maximising a lifetime utility function of the form in (1) must satisfy these, and so these gives us a way of telling whether the consumer's observed choices are generated from utility maximising behaviour. If so, then we say the consumer's choices can be *rationalised* by the life-cycle model (or are 'life-cycle consistent').

Definition 1. We say that the life-cycle model rationalises some data (q_t, ρ_t) if $\exists$ a real, concave, non-satiated, real-valued, differentiable function $u(.)$ and a discount rate $\beta \in [0, 1]$ such that $\beta^{t-1} u'(q_t) \leq \lambda_t \rho_t$ for all t.

To find a within-period utility function with the required properties, we need a condition that can be taken to data. The definition of concavity means that the utility function should satisfy the following

$$u(q_s) \leq u(q_t) + u'(q_t)(q_s - q_t)$$

$$\implies u(q_s) \leq u(q_t) + \frac{\lambda_t}{\beta^{t-1}} \rho'_t(q_s - q_t)$$

Furthermore, since $u(.)$ is continuous, there would have to be real numbers u_s and u_t such that

$$u_s \leq u_t + \frac{\lambda_t}{\beta^{t-1}} \rho'_t(q_s - q_t) \; \forall s, t \tag{6}$$

Finally, we know from monotonicity that λ_t must be positive (and the restrictions on β imply that β^{t-1} must also be positive). This means that for a given β we will have $2T$ unknowns $u_1..u_T$ and $\lambda_1..\lambda_T$ and $T(T-1)$ inequalities that may or may not be satisfied for a given set of quantities. By grid searching among possible values of β, we will be able to check if values of $u_1..u_T$ and $\lambda_1..\lambda_T$ exist such that this condition is satisfied for an individual over some period of time, and hence whether their behaviour can be rationalised by the life-cycle model.

It turns out that (6) is equivalent to the most basic revealed preference requirement - the generalised axiom of revealed preference or GARP - which states that if any bundle i is revealed preferred to j (directly or indirectly), then i must be unaffordable when j is chosen

$$x^i R x^j \implies p^j x^j \leq p^j x^i \; \forall i, j$$

This is due to Afriat's theorem, which implies that if GARP holds for current prices, then it will hold for a set of transformed prices (in our case prices transformed by $\frac{\lambda_t}{\beta^{t-1}} \times \frac{1}{(1+r_t)(1+r_{t-1})...(1+r_2)}$).

Proposition 1. *The following two statements are equivalent. 1) The data (q_t, ρ_t) are rationalised by the life-cycle model. 2) The data (q_t, ρ_t) satisfy GARP.*

Proof. See appendix.   □

As we said GARP is the most basic revealed preference test. The reason that the life-cycle model does not make any stronger predictions about behaviour is that we are allowing for uncertainty and the possibility of the consumer replanning in every period. In an uncertain environment, the consumer acts as though they face different lifetime constraints each period according to the news they received. Without access to this news, we will not be able impose any restrictions on the evolution of λ_t, and we are left with no more restrictions on behaviour than the simple consistency requirements of GARP. This is not to say however that a test of whether or not the consumer passes a test of GARP will exhaust the implications of the life-cycle model. In addition, in standard formulations of the life-cycle model innovations to λ_t will be martingale (Hall,1978): such that $E_t[\lambda_{t+1}] = \lambda_t$. We will return to the importance of this below when we look for nonrandom patterns in changes in the marginal utility of wealth across consumers.

In addition, we can employ a stricter test by imposing some restrictions on λ_t for any given consumer. For instance we could impose the condition that λ_t must be constant in condition (6). This would give us a condition of the form

$$u_s \leq u_t + \frac{\lambda}{\beta^{t-1}} \rho'_t(q_s - q_t) \; \forall s, t \tag{7}$$

which Browning (1989) calls cyclical monotonicity (CM). It turns out (in a result due to Browning, 1989) that this is a test of the life-cycle model with the additional assumption of perfect foresight. λ_t is constant as consumers are able to

perfectly smooth their marginal utilties. The definition of rationalisability with perfect foresight is the same as definition 1 except that λ is not longer allowed to vary with time (and so has no t subscript).

Definition 2. We say that the life-cycle model with perfect foresight rationalises some data (q_t, ρ_t) if $\exists$ a real, concave, non-satiated, continuous, differentiable function $u(.)$ and a discount rate $\beta \in [0, 1]$ such that $\beta^{t-1} u'(q_t) \leq \lambda \rho_t$.

Proposition 2. *The following two statements are equivalent. 1) The data (q_t, ρ_t) are rationalised by the life-cycle model under the assumption of perfect foresight. 2) The data (q_t, ρ_t) satisfy CM.*

Straightforwardly it can be seen that for any subset of periods (6) implies (7) and hence that $CM \implies GARP$.

The assumption of perfect foresight may well seem unreasonably strong. When we speak of perfect foresight however, we only mean over the periods when we observe the consumer. More precisely we might call this accurate foresight, since all it means is that the consumer is able to smooth their discounted marginal utility of wealth for a few periods without their plans being upset by unpredicted shocks. For short time horizons, this need not be a (completely) unreasonable assumption.

We can test GARP and CM if we have data on quantities and (discounted) prices for a sufficient number of periods. Conducting GARP and CM individual by individual for a dataset would tell us

1. What proportion of consumers' behaviour can be rationalised by a stable utility function of *any* form (GARP)?

2. What proportion of consumers' behaviour can be rationalised by the life-cycle model with perfect foresight (CM)?

A test could in principle including consumption of leisure as a commodity and the wage as a price. It may not be practical to take such a test to data however, as it is difficult to identify the price of leisure (the wage rate) in most household surveys. Information on hourly wages is often not included and calculating wages by for instance dividing earnings by hours may be misleading as wages can themselves vary with hours. Instead we can try to learn as much as possible with the limited information we have.

Without leisure, conditions (6) and (7) are necessary (but not sufficient) conditions for a version of the life-cycle model where preferences over leisure and consumption are weakly separable (Varian, 1983).[4]. Since they are insufficient, if the data pass these tests this does not necessarily vindicate the model, though a failure can still be considered a violation. These are the first tests we will take to our data.

Since GARP implies CM, these tests can have three possible outcomes. These are shown in table 1.

Table 1: Possible outcomes

	GARP	Perfect foresight
1.	Pass	Pass
2.	Pass	Fail
3.	Fail	Fail

The interpretation of a high proportion of passes for the perfect foresight test is straightforward. We would say we have no reason to reject the hypothesis that the life-cycle model with perfect foresight and separable preferences successfully rationalise behaviour. Various explanations might account for a high proportion of households in case 2. The first is that simply that there is some uncertainty leading to a failure of perfect foresight. This might be because individuals experienced a wealth shock (or news shock) at retirement or at some other point. If these shocks were coincident with the retirement decision (such as unexpected job loss), then they would be consistent with the explanation offered by Smith (2006) for the retirement consumption puzzle in the UK. One way to investigate this would be to allow one single change in λ_t in the retirement period (and surrounding periods) and to then repeat the test. A second explanation is that preferences over consumption goods change at the point of retirement, due to the fact that preferences between consumption and leisure are non-separable and consumers now have more leisure time, but that there is not enough variation in prices for this change to be detected as a violation by the GARP test. One way to examine this hypothesis is to conduct separate tests of the perfect foresight case for the periods before and after retirement. If it turns out that while there are many failures over the whole period, there are a substantial number of passes for the two sub-periods, we could interpret this as evidence for a depedence of preferences over consumption goods on labour force participation.[5]

---

[4]These conditions are not sufficient as the consumer would also have to satisfy additional restrictions on their choices over the omitted good (in this case leisure).

[5]This test is in the spirit of Varian (1988)'s test for cases where we do not observe prices for one of our goods: that revealed preference conditions for the remaining goods should still hold when consumption of the omitted good is constant.

In the event that none of these explanations work, then we will have to consider alternatives including more general uncertainty (multiple unexpected shocks), or simple failures of the life-cycle model (behavioural biases, credit constraints and so on). We will discuss how to investigate these in section 6.

A high proportion of of individuals in situation 3 would suggests that consumption behaviour cannot be rationalised by a stable, time separable subutility function. This would point to evidence of a dependence on consumption preferences on labour force status, which we can further investigate as discussed above. If tests carried out in individual sub-periods also failed then we would have reason to question some of the fundamental assumptions of the life-cycle model.

These tests require panel data with detailed information of on expenditures and labour force status for a sufficient period of time. For this we make us of the the Spanish Continuous Household Budget Survey (*Encuesta Continua de Presupuestos Familiares*, ECPF), which we now describe.

4   Data and sample

The ECPF is a household budget survey conducted by the Spanish National Institute of Statistics (INE). The survey is a rotating quarterly panel, with individual households followed for up to 8 quarters. The survey went through two different incarnations ECPF-85, which ran from 1985-1997, and ECPF-97, which ran from 1997-2005. The latter retained many features of the old survey - including the panel component - but differs substantially in the way information is collected. We only make use of ECPF-85. Both surveys were used to construct weights for the Spanish CPI.

ECPF-85 covered roughly 3,200 households each quarter, and collected detailed information on a broad range of different expenditures, demographic information (including the labour force status of both husband and wife) and data on income from a variety of sources (salary, self-employment, capital, pension, transfers and other). Expenditure data is quarterly and was collected through a combination of diary and recall questions.

We group expenditures into 10 non-durable commodity groups: food at home, utilities, food out, adult clothing, child clothing, transport, communications, recreation, personal care, household goods/services. This division is intended to separate out expenditures which are likely to be substitutes or complements for leisure (such as transport or adult clothing) from those which less likely to be (such as child clothing). We exclude medical goods as these are not pure consumption goods and were subsidised for Spanish retirees during this period (Boldrin, 1997).[6] Including them does not significantly affect the results.

Prices are drawn from the Spanish CPI. Where no price was available for a subset of goods, a price was calculated using a stone price index. Prices are discounted using interest rates on Spanish Treasury Bills taken from the International Monetary Fund. They are converted to quarterly rates by taking them to the power 0.25 to match the period of expenditure.

Our sample consists of households where the principal earner in the first quarter the household was surveyed retires and does not return to work in the period we observe them. We also restrict our sample to cases where households are observed at least two quarters before and after the retirement date. This gives us a sample of 312 households who we are able to follow over the retirement threshold. Some descriptive statistics for these households are presented in table 2.

Table 2: Summary statistics

	Mean	SD	Min	Max	Obs
Periods in Survey	7.23	1.18	4	8	312
Age at retirement	60.02	6.6	31	84	312
# HH members	3.47	1.55	1	11	312
Spouse in HH	0.94	0.24	0	1	312
Spouse always employed	0.11	0.31	0	1	293
Spouse never employed	0.67	0.47	0	1	293

The average age of retirement is 60 which is quite far below the official retirement age in this period of 65, though many pension schemes allowed retirement before this (Boldrin et al. 1997). One concern in the exercises that follow might be that the labour supply of spouses may change as the principal earner retires. The information on spouse's employment in table 1 suggests that this might not be too much of a problem. In most (67%) of our households, spouses are out of the labour force for the whole period where the household is observed, and an additional 11% remain employed even as the head retires.

Figures 1 and 2 show what happens to income and expenditure in the periods leading up to and following the point of retirement. Both suggest that these variables are relatively constant over the retirement period in accordance with the

---

[6] For instance public health care assistance was provided to those receiving non-contributory pensions administered by the *Instituto Nacional de Servicios Sociales*. In the later years, most regional government subsidised some medical costs for pensioners (as well as holidays and public transportation costs).

findings of Christensen (2008).
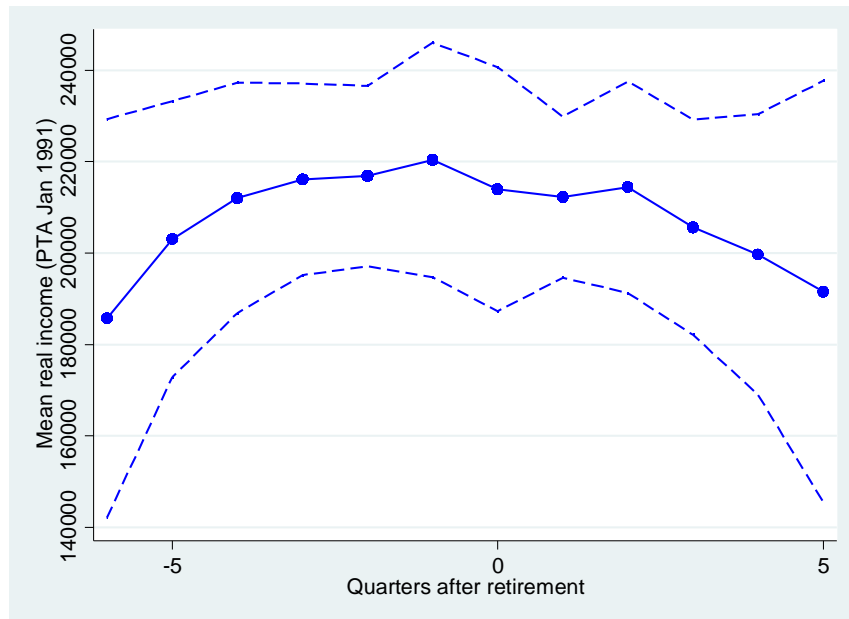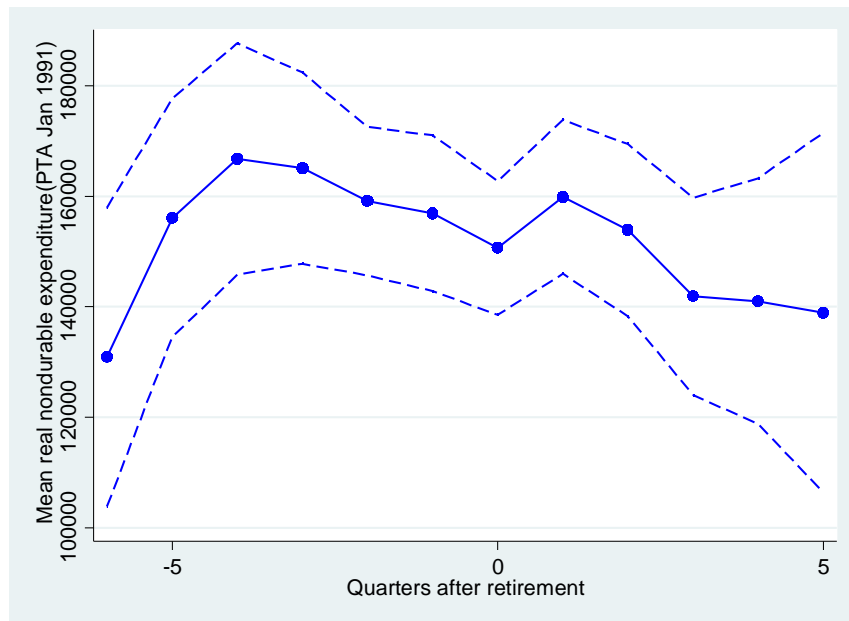
Figure 1: Total income over retirement

Figure 2: Total non-durable expenditure over retirement

As an additional indication of the impact of retirement on income and spending, we run two panel regressions of the logs of these variables on an individual fixed effect and a retirement dummy with no additional controls. The results are reported in table 3. There is a small statistically significant 3% decline in log expenditure and a statistically insignificant increase in log income, but this disappears once household size is included as an additional control. Changing is not something which our revealed preference tests take account of, but it could plausibly account for apparent preference changes at the household level. We will return to this issue in what follows.

Table 3: Regressions: income and expenditure at retirement

|	Log Total Non-durable Expenditure		Log Total Income	
Retired Dummy	-0.038*	-0.030	0.018	0.023
	(0.016)	(0.016)	(0.022)	(0.021)
Household size		0.136**		0.081**
		(0.020)		(0.027)
# Obs	2256	2256	2254	2254
#Households	312	312	312	312

Notes: * indicates significant at the 5% level ** indicates significant at the 1% level. Regressions are fixed-effects (within) models.

5 Empirical Results

We begin by testing the life-cycle model by implementing a test of GARP, for the whole period when households are observed, followed by the strictest of our tests (perfect foresight). We then go on to consider an intermediate case - allowing a single revision to the discounted marginal utility of wealth - before carrying out a test allowing for a change in preferences over consumption goods at the point of retirement (non-separable preferences). Each of these tests is run on each household separately and independently, and so each allows for *complete* heterogeneity in the utility functions of each individual (both in terms of whether they pass or not and in terms of the form of their preferences). In tests where a discount factor is involved, we grid search for a solution over 100 possible values of β between 1 and 0.

5.1 GARP/Life-cycle model

As we said, this is equivalent to test of the life-cycle model allowing for uncertainty in a very general way but where consumption and leisure are separable. The results of this test are summarised in figure 3.

Figure 3: GARP
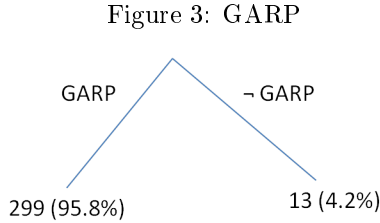
96% of households satisfy GARP over the whole period. This means that we have no evidence for any change in preferences over consumption goods over the period of retirement (consistent with the findings of Christensen, 2008 and Luenga-Prado and Sevilla, 2013). However, absence of evidence is not always good evidence of absence. The informational content of this result depends on the difficulty of passing GARP if individual tastes were *not* stable over the period, as is clear if we formulate this question in Bayesian terms. Let U be the probability that an individual acts as a maximiser of a stable utility function and let G be the probability that the individual passes a GARP test (or any other revealed preference test). Bayes theorem implies that

$$P(U|G) = \frac{P(G|U)P(U)}{P(G|U)P(U) + P(G|\neg U)[1 - P(U)]}$$

Since GARP is necessary and sufficient for utility maximising behaviour $P(G|U) = 1$, so

$$P(U|G) = \frac{P(U)}{P(U) + P(G|\neg U)[1 - P(U)]}$$

which depends inversely on $P(G|\neg U)$ - the probability that an individual passes GARP over the period of retirement even if their preferences did change (or if they do not in general act like utility maximisers). When $P(G|\neg U) = 1$ then we will have no reason to update our prior. Conversely, as $P(G|\neg U) \to 0$, $P(U|G) \to 1$. As this term gets smaller, the more confident we can be that a successful GARP test is evidence of utility maximising behaviour.

To compute $P(G|\neg U)$ we need some alternative data-generating process. Unfortunately, there are many forms of irrational behaviour that could serve this purpose, and so we will have to select one among the possible contenders to utility maximisation. Following Bronars (1987) we will adopt Becker's (1962) concept of irrational behaviour - uniformly

random choices across bundles on the consumer's budget hyperplane. Under this alternative, $P(G|\neg U)$ will be the probability of a consumer passing GARP if they had chosen which bundle to consume by rolling a multi-sided die (with one side for each possible bundle). We calculate this by first observing how often a household passes GARP when their budget shares are drawn randomly from a uniform distribution, and then averaging these probabilities across individuals.[7] In our case the probability of a pass was 98.2% across the sample. This is actually greater than the pass rate of our data, and is of course very close to 100% (which would give us no reason to update any of our priors at all).[8] Thus, this exercise does essentially nothing to make the hypothesis of "no taste change" any more likely. We consider further testable implications of the life-cycle model in section 6 below.

5.2 Perfect Foresight

The GARP test turns out not to be too informative. A stricter test is of whether or not the consumer passes a test of the life-cycle model with perfect foresight. In this test we keep the discounted marginal utility of wealth constant across periods. Passing GARP over the whole period is a necessary condition for passing this test. Since we conduct this test on all periods (before and after retirement) this test also assumes no preference change at the point of retirement. The results are shown in figure 4. Of those households who passed GARP only 3 (1%) also pass the perfect foresight test.

Figure 4: Perfect Foresight

GARP ¬ GARP

299 (95.8%)   13 (4.2%)

Perfect
foresight

3 (1%)   296 (99%)
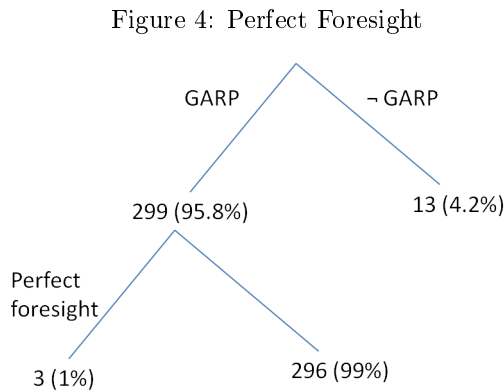
5.3 Single change in λ

Now we allow for a single change to λ_t at the time of retirement. This is an intermediate case between perfect foresight and GARP. This test allows for some uncertainty in that we allow for shocks (both positive and negative) that are coincident with the retirement decision such as job loss or health shocks. The results of this test are shown in figure 5. Only an additional 7 individuals pass when we allow a single adjustment to λ_t in this way.

---

[7] To obtain these budget shares, we first draw fractions randomly from a Dirichlet distribution and then divide them by their sum to obtain a set of budget shares that sum to one. This procedure ensures that the resulting budget shares will be uniformly distributed. We then calculate quantities as $q_t^i = w_t^i \times \frac{x_t}{p_t^i}$ where w_t^i is the budget share of good i in period t and x_t is the individual's expenditure and then use these to rerun the GARP test.

[8] An alternative means of demonstrating the predictive power of GARP proposed by Selten (1991) and axiomatised by Crawford and Beatty (2011) is the Selten score. This is simply the pass rate less the Selten area. This gives a value which lies between -1 and 1. A score approaching 1 represents an ideal case where the choices seen in the data all pass GARP but individuals choosing randomly would hardly ever pass. A score approaching -1 would imply the converse: the actual choices seen in the data fail GARP much more often than random choices. In our case the probability of a pass was 98.2% across the sample, giving a selten score of 0.043.

Figure 5: Single change in λ
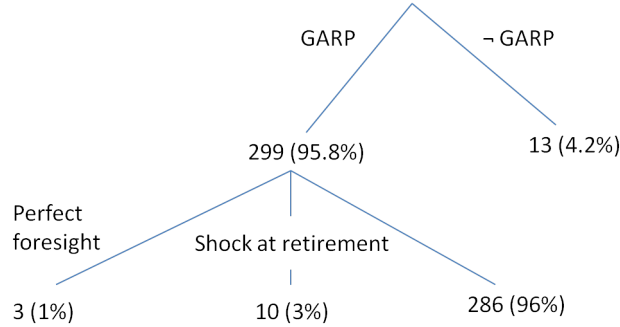
One concern may be that the decision to retire may not be exactly coincident with any shocks the individual experiences. For instance, an individual might retire only after a shock has been realised, or else might retire in anticipation of some shock. To allow for this possibility we can simply take the whole period see if a single revision to λ *at any time* can rationalise the data. For this exercise we need to have sufficient periods before and after a supposed shock, so we restrict our sample to those individuals observed for the maximum 8 periods. The results of the various tests are shown in table 4 below. Allowing for a single change at any date does not give us any additional passes.

Table 4: Pass rates with single change to λ

	No change to λ	Single change (at retirement)	Single change (any date including retirement)
Passes	0	1	1
Total Obs	197	197	197

5.4 Allowing for non-separable preferences

Here we allow for preference change over consumption goods at retirement (due to the change in labour force status) by carrying out the perfect foresight test in the periods before and then after retirement separately. The pass rate on our revealed preference tests can only increase when we reduce the number of periods we are testing (since this merely reduces the chances of encountering a violation). Nonetheless, an substantial increase in the pass rate might still give us reason to believe that a depedence on preferences over consumption goods on labour force participation is a likely explanation for the consumption behaviour we see around retirement. The results are shown in table 5.

Table 5: Pass rates in periods before and after retirement

	Passes			Total
	Before	After	Before and After	
GARP	311	310	309	312
CM	135	120	32	312

309/312 individuals pass GARP when we test separately before and after retirement. Of these, just 32 - or 10% - pass the life-cycle model both in the periods before and in the periods after retirement. Simple preference change does not do well at explaining violations of the perfect foresight model.

What can we conclude from this battery of tests? Firstly, the restrictions imposed by GARP do not provide a very strong test of the life-cycle model. Secondly, if preferences between consumption and leisure are separable then over this period something appears to happen to the marginal utility of wealth: a test of the perfect foresight model with this assumption fails. Moreover this is more systematic than a simple one-off revision (whether at the time of retirement or any other time). Finally, a test of the perfect foresight model allowing for separable preferences also does poorly. In the next section we consider further implications of the life-cycle model with more general uncertainty.

6 Violations

As we noted previously, GARP does not exhaust the empirical implications of the life-cycle model with uncertainty. To investigate the plaubsibility of this model further we can consider deviations e_t from the perfect foresight case

$$u_s \leq u_t + \frac{(\lambda + e_t)}{\beta^{t-1}} \rho'_{\mathbf{t}}(\mathbf{q_s} - \mathbf{q_t}) \tag{8}$$

where $(\lambda + e_t) > 0$. Under the life-cycle model, these deviations should be martingale, that is with $E_t[e_{t+1}] = e_t$. We can now ask what the paths of $\{e_1 \ldots e_T\}$ which rationalise our data look like. Multiple paths may rationalise the behaviour of any given household, and the data cannot tell us which of these is right. We will select the path which minimises the sum of squared deviations $\sum e'e$, which will in turn isolate a particular path for λ_t. This is the path with the smallest deviations from the pefect foresight life-cycle model and so is as favourable as possible to the maintained hypothesis of rational, forward-looking consumers. We can now ask whether the pattern of deviations across individuals are plausible given the predicted martingale property of the deviations.

These violations are terms in a utilty function and so do not have a cardinal interpretation. We cannot compare the size of violations across individuals, and there is no sense in saying that one violation is for instance 'twice as big' as another. However, the *change* in these terms from one period to another does have an interpretation. A systematic tendency to get decreasing or increasing e's indicates that something systematic is happening to the discounted marginal utility of wealth. Consumption would be increasing or decreasing faster than it 'should' do under the life-cycle model, and this would point to a rejection of this model.

Figure 6 compares λ_t to λ at the point of retirement for households who we observe for all 8 quarters. It is clear that, in each period before retirement, a majority households have λ_t's which lie below those in the first period of retirement. Moreover, in each period after retirement, a majority of households have λ_t's which exceed it. Moreover the percentage of households with lower λ_t than the retirement date seems to monotonically increase with time. This points to a general tendency for the marginal utility of wealth to be increasing over time (suggesting that consumption is falling faster than we might expect). This may be a phemenon related to retirement in particular or it may be a more general phenomenon. Figure 7 plots the proportions of households with marginal utilities of wealth below the marginal utility seen in the 4th quarter we observe them for working age households (where the head remains employed throughout the 8 quarters we observe them and is aged 55-60). A very similar pattern is evident, pointing towards a general problem with the life-cycle model (although the extent of the increase is smaller).

Figure 6: Proportion λ_t higher than λ_{retire}
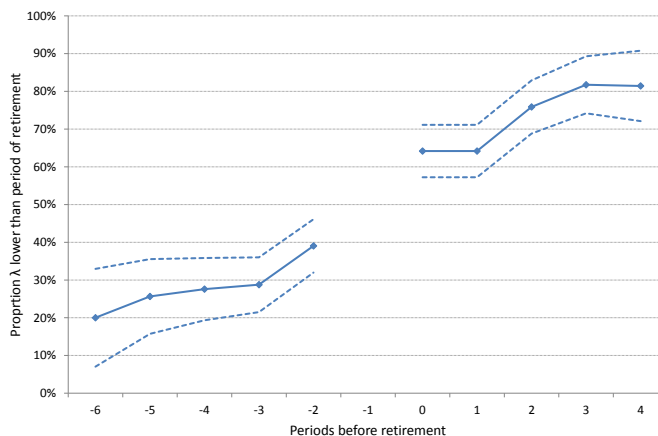
Note: Dotted lines represent 95% confidence intervals

Figure 7: Proportion λ_t higher than λ_4, working age households
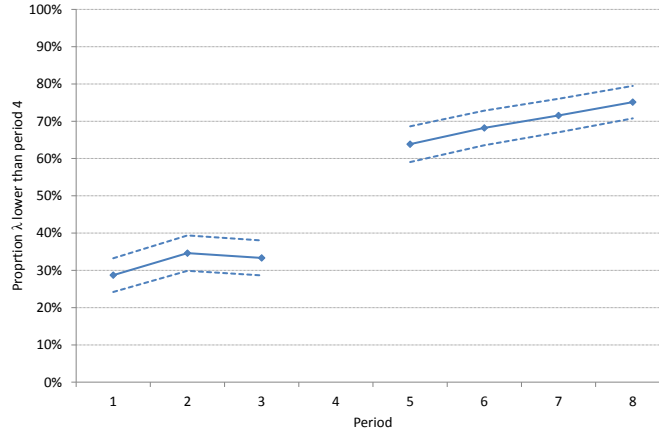
Note: Dotted lines represent 95% confidence intervals

The fact that it we also find increasing λ's for households who do not change their labour supply suggests that it is not something which can be dismissed as being due to non-separable preferences between consumption and labour force participation. As we saw from table 2 the proportion of spouse's who change their labour supply in our retiring households is also low, making it unlikely that this explains the result either. So what might explain these patterns? In the rest of this section we will discuss some possibilities.

6.1   Adverse shocks

Could a series of adverse shocks account for an increasing discounted marginal utility of wealth? To see how uncertainty affects the λ_t we need to return to the first order conditions for consumption commodities from the problem in (1).

Given the estimates of lifetime wealth in period t, the first order conditions for consumption of the goods q in period t (ignoring non-negativity constraints) are

$$\lambda_t = \frac{\beta^{t-1} u'(q_t)}{\rho_t} = E_t \left[\frac{\beta^t u'(q_{t+1})}{\rho_{t+1}} \right] \tag{9}$$

Now we know that based the assessment of lifetime wealth made in period $t+1$

$$\lambda_{t+1} = \frac{\beta^t u'(q_{t+1})}{\rho_{t+1}} \tag{10}$$

Substituting (10) into (9) gives

$$\lambda_t = E_t \left[\lambda_{t+1} \right]$$

or the martingale result. Given some regularity conditions we can decompose into the value of $E_t \left[\lambda_{t+1} \right]$ (which is λ_{t+1} if expectations are rational) and an additive, mean-zero error

$$\lambda_t = \lambda_{t+1} + \varepsilon_t$$

Since the utility function is concave, a positive value of ε_t will mean that the consumer has less wealth than they had expected in the second period: the marginal utility of wealth increases when the consumer experiences a negative wealth shock, and decreases when the consumer experiences a positive wealth shock.

Rational expectations imply that idiosyncratic shocks should be random, not systematic across individuals, and so these would not account for the pattern of violations that we saw in figure 6. A systematic tendency across a fairly large group of individuals could be the result of common, or aggregate shocks, such as unanticipated decreases in pensioner benefits or productivity shocks. But this seems an unlikely explanation for a pattern observed consistently for many individuals over a period of 12 years (apparently including those of working age). One would have to appeal to quite a long sequence of unanticipated, negative shocks to generate this.

Figure 8: Proportion λ_t higher than λ_{retire}, no change in HH size
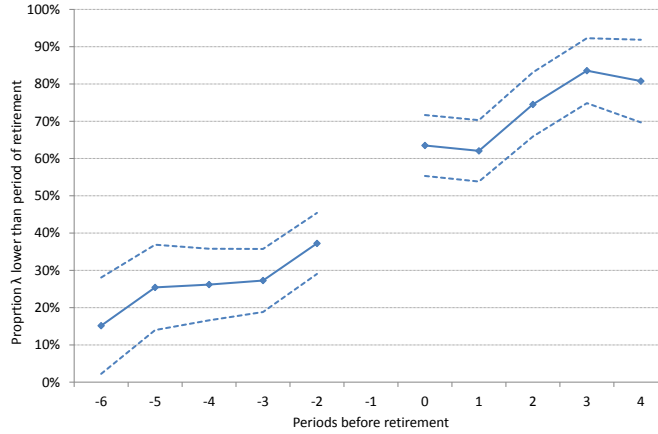
Note: Dotted lines represent 95% confidence intervals

It should be noted that increasing the martingale predicts a constant *mean* value of the discounted marginal utility of wealth, while in figure 6 we saw evidence of an increasing proportion of households who had λ_t greater than λ at the point of retirement. This could be rationalised with the life-cycle model by a low probability possibility of a large positive wealth shock leading households to spend more than they otherwise would. Such an explanation seems unlikely however, as it is difficult to imagine what such a positive wealth shock could be.

6.2   Changing household composition

One possibility suggested by the results in table 4 is that the composition of households is changing over time. The evidence for a 'retirement consumption puzzle' in this data disappeared once household size was controlled for. We can check for this by looking at households where household size remains constant. Figure 8 is the same as figure 6 except that we restrict the sample to households where there is no change in household size. It shows much the same pattern.

6.3   Measurement error in the interest rate

Another possibility is that we have simply mismeasured the real interest rate we use to discount prices. If for instance, we used an interest rate that was too large relative to the interest rate individuals could obtain on their savings, then consumption would appear to decrease more than it should do (and the discounted marginal utility of wealth would appear to increase). However, the interest rate data we use is likely to be conservatively low. We use interest rates on Spanish Treasury bills (made quarterly) which may be lower than the interest rates savers actually receive.

6.4   Credit/saving constraints

The life-cycle model assumes perfect capital markets, which is what allows the consumer to smooth their marginal utilities. If we relax this assumption, then the λ_t need no longer be constant even in the perfect foresight case. To see how, we can introduce a credit constraint into the problem in (1): consumption in period t cannot exceed some limit b_t. Abstracting from uncertainty (and dropping the non-negativity constraints) this gives us the Lagrangian

$$L = \sum \beta^{t-1} u(q_t) - \lambda(\sum \rho'_t q_t - W) - \mu_t(\rho'_t q_t - b_t)$$

where W is *total* lifetime wealth. This is associated with the first order condition

$$u'(q_t) = \frac{\rho_t(\lambda + \mu_t)}{\beta^{t-1}}$$

and hence the cyclical monotonicity condition

$$u_t \leq u_s + \frac{(\lambda + \mu_t)}{\beta^{t-1}} \rho_t'(q_s - q_t)$$

μ_t is the Kuhn-Tucker multiplier on the credit constraint (or the marginal value of relaxing it). When the constraint is binding $\mu_t > 0$. For a saving constraint, we would have $\mu_t < 0$. A path of increasing λ_t could therefore be rationalised if we thought consumers were unable to save in earlier periods. This would mean that while consumers knew that their future resources were going to decrease, they were not be able to smooth their consumption in they way they would like because they were unable to save. This seems unlikely. Indeed, a fair number (54%) of households saw their real incomes increase on average over the period of retirement.

6.5   Time inconsistency

A final possibility is that consumers are non-rational. As we noted above this was the preferred explanation of Bernheim et al. (2001) for the retirement consumption puzzle. A tendency to consume too much at younger ages (due to for instance hyperbolic preferences) would explain the patterns we see in the data,and would help to explain why the same pattern is also observed among households who are not retiring. Further tests of these models (such as those proposed in Blow et al., 2014) using a revealed preference framework could lend further support to this hypothesis.

7   Conclusion

In this paper we set out revealed preference tests for different models of consumption behaviour over retirement that we applied to a Spanish consumption panel dataset. We reject the perfect foresight model both with separable preferences and allowing for preference change. The first order conditions for the life-cycle model allowing for uncertainty do not provide very strong restrictions on possible choices. In fact they are no stronger than those implied by the most basic revealed preference requirement: GARP. We then go on to investigate the patterns of deviations from a perfectly smoothed marginal utility of wealth and ask whether they fit the predictions of the life-cycle model. We find a tendency of these to increase over time, suggesting consumption falls more than we'd expect. After considering various possible explanations, we settle on non-rational behaviour as the most plausible.

Appendix

Proof of proposition 1.

Proof. To prove this proposition we need to show that firstly the test in (6) is equivalent to GARP. Secondly, we need to show that data that satisfy the test in (6) then this implies that they will be rationalised by the life-cycle model in the sense given in definition 1.

Equivalence of (6) and GARP: Afriat's theorem states that the following two statements are equivalent (see Afriat, 1967, or Varian,1982, for a proof):

(i) The data satisfy GARP: if $q_s R q_t$, then $p_t q_t \leq p_t q_s$ for all s, t
(ii) There exist 'Afriat numbers' $U_s, U_t, \lambda_t > 0$ such that $U_s \leq U_t + \lambda_t p_t'(q_s - q_t)$ for all s, t

Now if condition (ii) is satisfied for some data p_t, q_t, then it will also be satisfied for the data $\sigma_t p_t, q_t$ where $\sigma_1, .., \sigma_T$ are any positive numbers. This is because if $U_1...U_T$ and $\{\lambda_1..\lambda_T\}$ satisfied (ii) in the old data, then $U_1...U_T$ and $\{\hat{\lambda}_1..\hat{\lambda}_T\}$, where $\hat{\lambda}_t = \lambda_t/\sigma_t$, will satisfy it in the new.

Conversely, if (ii) is satisfied for some $\sigma_t p_t, q_t$, then it will also be satisfied for p_t, q_t. Hence, by Afriat's theorem p_t, q_t will satisfy GARP if $\sigma_t p_t, q_t$ satisfy (ii).

If we set $\sigma_t = \beta^{t-1}(1 + r_t)(1 + r_{t-1}) \ldots (1 + r_2)$, then (ii) becomes the test of the life-cycle model in (6). Since this both implies and is implied by GARP, we have that $GARP \Longleftrightarrow (6)$.

Equivalence of the data satisfying (6) and the data being rationalised by the life-cycle model (LCM): That LCM implies (6) is demonstrated in the text. Thus, (6) is a necessary condition for LCM. It remains to be demonstrated that (6) is sufficient for LCM. We can rearrange our test to get

$$\frac{\lambda_t}{\beta^{t-1}} \rho_t'(q_s - q_t) \geq u(q_s) - u(q_t)$$

Let's call $\frac{\lambda_t}{\beta^{t-1}} \rho_t' = \Xi_t$. Across periods this gives us

$$\Xi_t(q_s - q_t) \geq u(q_s) - u(q_t)$$

$$\Xi_s(q_u - q_t) \geq u(q_u) - u(q_s)$$

$$\vdots$$

$$\Xi_u(q_z - q_u) \geq u(q_z) - u(q_u)$$

Adding these up gives

$$\Xi_t(q_s - q_t) + \Xi_s(q_u - q_s) + \ldots + \Xi_u(q_z - q_u) \geq 0$$

This matches what Rockafellar (1997) and Browning (1989) call "cyclical monotonicty". Following Browning (1989) we can then apply Theorem 24.8 in Rockafellar to demonstrate the existence of a concave utility function $u(.)$ such that $u'(q_t) = \Xi_t$ for all t. This satisfies the requirement in definition 1, completing the proof. $\square$

References

[1] Afriat, S.N. (1967) "The Construction of Utility Functions from Expenditure Data," *International Economic Review*, 8, 67–77.

[2] Aguiar, M., Hurst, E. (2005) "Consumption versus Expenditure," *Journal of Political Economy*, 113, 919–948.

[3] Aguiar, M., Hurst, E. (2007) "Life-Cycle Prices and Production," *American Economic Review*, 97, 1533–1559.

[4] Aguila, E., Attanasio, O., Meghir, C. (2010) "Changes in Consumption at Retirement: Evidence from Panel Data," *Review of Economics and Statistics*, 93, 1094–1099.

[5] Banks, J., Blundell, R., Tanner, S. (1998) "Is There a Retirement-Savings Puzzle?" *American Economic Review*, 88, 769–788.

[6] Battistin, E., Brugiavini, A., Rettore, E., Weber, G. (2009) "The Retirement Consumption Puzzle: Evidence from a Regression Discontinuity Approach," *American Economic Review,* 99, 2209–2226.

[7] Beatty, T.K.M., Crawford, I.A. (2011) "How Demanding Is the Revealed Preference Approach to Demand?" *American Economic Review,* 101, 2782–2795.

[8] Becker, G.S. (1962) "Irrational Behavior and Economic Theory," *Journal of Political Economy,* 70, 1–13.

[9] Bernheim, B.D., Skinner, J., Weinberg, S. (2001) "What Accounts for the Variation in Retirement Wealth among U.S. Households?," *American Economic Review,* 91, 832–857.

[10] Boldrin, M., Jimenez-Martni, S., Peracchi, F. (1997) "Social Security and Retirement in Spain," Working Paper No. 6136, National Bureau of Economic Research.

[11] Blow, L., Browning, M., and Crawford, I. (2014) "Never mind the hyperbolics: nonparametric analysis of time-inconsistent preferences," IFS Working Papers (W14/17)

[12] Bronars, S.G. (1987) "The Power of Nonparametric Tests of Preference Maximization," *Econometrica* 55, 693–698.

[13] Browning, M. (1989) "A Nonparametric Test of the Life-Cycle Rational Expectations Hypothesis," *International Economic Review* 30, 979–992.

[14] Carroll, C.D., 2001, "Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation)," *Advances in Macroeconomics,* 1.

[15] Christensen, M. (2008) "Demand patterns around retirement: Evidence from Spanish panel data," The School of Economics Discussion Paper Series No. 0809, Economics, The University of Manchester.

[16] Crawford, I., Polisson, M. (2013) "Testing for Intertemporal Nonseparability" Discussion Papers in Economics No. 13/08, Department of Economics, University of Leicester.

[17] Haag, B.R., Hoderlein, S., Pendakur, K. (2009) "Testing and imposing Slutsky symmetry in nonparametric demand systems," *Journal of Econometrics*, 153, 33–50.

[18] Hall, R. (1978) "Stochastic implications of the life cycle-permanent income hypothesis," *Journal of Political Economy*, 86(6), pp. 971-87.

[19] Hurst, E. (2008) "The Retirement of a Consumption Puzzle," Working Paper No. 13789, National Bureau of Economic Research.

[20] Kimball, M.S. (1990) "Precautionary Saving and the Marginal Propensity to Consume," NBER Working Paper No. 3403, National Bureau of Economic Research, Inc.

[21] Luengo-Prado, M.J., Sevilla, A. (2013) "Time to Cook: Expenditure at Retirement in Spain," *The Economic Journal*, 123, 764–789.

[22] Rockafellar, R.T. (1997) "Convex Analysis" Princeton University Press.

[23] Selten, R. (1991) "Properties of a measure of predictive success," *Mathematical Social Sciences*, 21, 153–167.

[24] Smith, S. (2006) "The Retirement-Consumption Puzzle and Involuntary Early Retirement: Evidence from the British Household Panel Survey," *The Economic Journal*, 116, C130–C148.

[25] Varian, H.R. (1982) "The Nonparametric Approach to Demand Analysis," *Econometrica*, 50, 945–973.

[26] Varian, H.R. (1983) "Non-Parametric Tests of Consumer Behaviour," *The Review of Economic Studies*, 50, 99–110.

[27] Varian, H.R. (1988) "Revealed preference with a subset of goods," *Journal of Economic Theory*, 46, 179–185.