CAN THE RETIREMENT CONSUMPTION PUZZLE BE RESOLVED? EVIDENCE FROM UK PANEL DATA

Sarah Smith

THE INSTITUTE FOR FISCAL STUDIES
WP04/07

Can the retirement-consumption puzzle be resolved?

Evidence from UK panel data

Sarah Smith*

This version: March 2004

Abstract

This paper uses UK panel data to shed further light on the fall in spending at retirement (the "retirement-consumption puzzle"). It compares the profiles of spending and well-being at retirement for different groups, defined according to whether retirement is voluntary or involuntary. Where retirement is voluntary, food spending and individual well-being are largely smoothed through retirement; where retirement is involuntary, both food spending and well-being fall. This is consistent with the retirement consumption puzzle being linked to negative wealth shocks. However, there remains one group for whom retirement appears to be voluntary, yet whose spending falls. Fully resolving the puzzle requires a better understanding of how the nature of retirement links to spending and of how different groups substitute leisure for consumption.

*London School of Economics and Institute for Fiscal Studies

Thanks to Richard Blundell and James Banks for their invaluable advice and help. Data from the British Household Panel Survey were made available by the Economic and Social Research Council. Any errors are the responsibility of the author

1. Introduction

A number of recent studies have looked at what happens to spending at retirement – and why (see Hamermesh (1984), Banks, Blundell and Tanner (1998), Bernheim, Skinner and Weinberg (2001), Ameriks, Caplan and Leahy (2002), Hurd and Rohwedder (2003), Miniaci, Monfardini and Weber (2003), Haider and Stephens (2004)).

This issue is important for a number of reasons. First, spending is likely to reflect an individual's well-being and can therefore give insights into how well off people are in retirement, compared to when they are working. Particularly if retired people hold substantial levels of (non-annuitised) wealth which they use to finance consumption, looking directly at spending may provide a better measure than income replacement rates. Of course, however, (changes in) spending may not exactly reflect (changes in) an individual's well-being when they retire if their spending needs change because of the end of work-related spending (eg commuting and suits) or because they choose to substitute increased leisure time for spending.

Secondly, there has been an interest in testing the predictions of the life-cycle model of consumption against what actually happens at retirement. The "stripped-down" version of the model predicts that consumption should be smoothed through anticipated changes in income, such as are likely to occur around retirement. In practice, the finding of all the studies listed above is that consumption falls significantly at retirement – this is common across a number of countries (US, UK and Italy), across different time periods and across different measures of spending. This fall has become known as the "retirement-consumption puzzle".

The puzzle has a number of possible interpretations. Bernheim, Skinner and Weinberg (2001) conclude that their "findings are difficult to interpret in the context of the life-cycle model" and that people, instead of rationally planning their retirement saving as the life-cycle model implies, tend to use simple rules of thumb. Banks, Blundell and Tanner (1998) suggest that there may be "unanticipated shocks occurring around the time of retirement" which could explain the fall in spending within the context of the life-cycle model, allowing for uncertainty. Hurd and Rohwedder (2003) argue that the fall in spending can be explained by an extended version of the life-cycle model which incorporates leisure in a household production function. In this case, spending falls at retirement because of the substitution of increased leisure time.

This paper uses information on the retirement of a sample of men in the British Household Panel Survey to shed further light on the retirement-consumption puzzle in

the UK. The panel data evidence confirms that spending (on food) does fall around retirement (see Figure 1), consistent with the earlier finding of Banks, Blundell and Tanner (1998) which was based on pseudo-cohort data from the Family Expenditure Survey. The observed fall is gradual but, as argued later, this is likely to be exaggerated by the nature of the survey.

Figure 1: Food spending around retirement

The paper compares the retirement experiences of different groups of retirees and in particular compares those for whom retirement is voluntary and those for whim it is involuntary. Previous studies have shown that the experience of retirement in the UK is very varied (see Tanner (1998), Blundell, Meghir and Smith (2002), Marmot et al (2004)) and, for a significant minority is linked to ill-health or redundancy.[1] For this group, retirement is more likely to be associated with a negative wealth shock and so comparing voluntary and involuntary retirements should provide evidence on the extent to which the fall in spending is linked to such unanticipated shocks.

The paper also looks at profiles of (a measure of) well-being in addition to looking at profiles of spending for further insights into the retirement-consumption puzzle. The intuition is straightforward. Assuming that retirement does not enter as a separate additive component in the utility function and that there is no substitution of leisure for consumption then, in the absence of shocks, both consumption and well-being should be smoothed through retirement. With a negative shock to wealth, however, both consumption and well-being would fall. By comparison, substitution of leisure

[1] Of course, redundancy does not necessarily lead to retirement (permanent labour market exit), but the wage cut someone would have to take in getting another job may be enough to make them stop working altogether.

for consumption at retirement (again assuming no additive components in the utility function), would imply a very different pattern at retirement; namely a fall in spending, but a smoother profile for well-being.

Of course retirement may enter the utility function as a separate additive component, implying a change in well-being at retirement. However, so long as this is the same across the groups, comparisons between the groups should allow some tentative conclusions to be drawn about the nature of the consumption-puzzle. The analysis controls for health which is likely to affect both the nature of retirement and well-being directly.

In practice, profiles of spending and well-being do vary across the different groups of retirees. Where retirement is voluntary, food spending and individual well-being are largely smoothed through retirement; where retirement is involuntary, both food spending and well-being fall. This is consistent with an explanation for the retirement consumption puzzle in terms of negative wealth shocks. However, there remains one group who retire at the state pension age (65) for whom retirement appears to be voluntary and whose well-being does not fall, yet whose spending falls.

The analysis suggests that negative wealth shocks play a role in resolving the retirement-consumption puzzle in the UK. But no single explanations is fully able to explain the patterns of consumption and well-being around retirement for all the groups. Fully resolving the puzzle is likely to require a better understanding of how the nature of retirement links to spending and of how different groups substitute leisure for consumption.

2. The puzzle and possible resolutions

The fact that observed spending falls at retirement is a challenge to the simple, one-consumption-good life cycle model. In its simplest form, with utility dependent only on consumption, no uncertainty and assuming that marginal utility is continuous and declining in consumption, the maximisation of lifetime utility implies that the marginal utility of consumption, and consumption itself, should be smoothed. Falling consumption at retirement contradicts this simple model and, if the model could not be extended, would imply irrational behaviour by consumers. This is the conclusion reached by Bernheim, Skinner and Weinberg (2001) who argue that the evidence of a fall in spending at retirement points to people using rules of thumb, rather than forward-looking optimising behaviour, to determine retirement saving.

One possible explanation, which would save the simple version of the life-cycle model is that the studies capture a fall in *spending* at retirement, which is not the same as a fall in utility-producing *consumption* at retirement. Households may stock up on

durables immediately prior to retirement and enjoy a higher flow of services from durables after retirement; thus while their observed spending may fall, their overall consumption remains the same. However, Miniaci, Monfardini and Weber (2003) find no evidence of pre-retirement stocking up of durables. Another possible explanation is that there is a necessary level of (non-utility-producing) spending associated with working, for example the cost of buying suits and travelling to work, that stops when people retire. Again, this would imply that, while observed spending falls, (utility-producing) consumption may be smoothed over retirement. This effect will be reinforced to the extent that the spending of the retired on certain items is subsidised (transport and prescription charges in the UK, health in the US). However, Banks, Blundell and Tanner (1998) take out obvious work-related spending items from total spending and look at sub-components of spending and still find evidence of a fall at retirement.

There are two possible extensions to the simple life cycle model which would make a fall in spending at retirement perfectly consistent with rational, optimising behaviour.

First, allowing for uncertainty, the prediction of the life-cycle model is that consumption should not fall only in the absence of any negative shocks to lifetime wealth. The evidence from the UK suggests that for many, retirement may indeed be associated with such a negative shock. Disney and Tanner (1999) show that more people tend to retire earlier than expected than later and, while this may be associated with a positive wealth shock which causes earlier than expected retirement, Tanner (1998) and Marmot et al (2004) find a significant number of people citing ill health and compulsory early redundancy as the main reason for early retirement. In a recent US study, Haider and Stephens (2004) look at the extent to which consumption falls when retirement is expected (on the basis of previously reported subjective retirement expectations) and find that the magnitude of the fall is far smaller than when looking at the fall associated with actual retirement. This suggests that shocks play a role in explaining some of the observed fall in spending at retirement, although, as noted, actual and expected retirement may differ for a number of different reasons.

Second, utility may depend not just on consumption but also on household demographics, including leisure and health. In this case, there is no reason to expect consumption itself to be smoothed; instead the optimal path of consumption will depend on how the marginal utility of consumption varies with changes in demographics. Börsch-Supan and Stahl (1991) argue that people's ability to derive utility from consumption may fall with old age, but this explanation is unlikely to explain the fall in spending among the "younger old" at retirement. Clearly, however, retirement is associated with a change in leisure, and, for some, a change in health that may affect consumption. If, for example, consumption and leisure are substitutes,

the increase in leisure will cause a reduction in consumption. Moreover, if people expect leisure to increase when they retire, they will fully anticipate this decline in consumption.

In practice, the change in leisure at retirement in the UK tends to be discontinuous. Most people move straight from full-time employment into retirement and do not gradually reduce the number of hours worked (see Tanner (1998)). There is clearly an issue about whether this is optimal from the individual's point of view given diminishing marginal returns to leisure. There are possible reasons why individuals may not want to reduce their hours gradually, including fixed costs associated with working and/or economies of scale in converting time into utility-producing leisure. More likely, they may face constraints in their choice of the number of hours to work as a result of the fixed costs of employment to the employer and, for people with a defined benefit occupational pension in the UK, current legal restrictions on drawing any pension income while still working for the same employer. These issues are not the focus of this paper; here it will simply be assumed that there is a discontinuous change in leisure at retirement, which would imply a (largely anticipated) one-off fall in consumption at retirement if leisure and consumption are substitutes.

Hurd and Rohwedder (2003) explore these issues using information on anticipated and actual spending changes at retirement from the 2001 Consumption and Activities Mail Survey, a supplementary survey to the longitudinal Health and Retirement Survey. Their main finding is that, among the pre-retired, there is an anticipated decline in spending at retirement that is, if anything, greater than the fall in spending that actually occurred among (a different group of) those who had already retired (20% compared to 12% among married couples, for example). Ameriks, Caplin and Leahy (2002) also find that many people expect to spend less in retirement.

The finding that many people expect spending to fall is extremely interesting, but it is not conclusive about the mechanism that causes actual consumption to fall. If panel data showed a strong link between anticipated and actual spending this would imply the absence of shocks, but could be explained by consumers applying rules of thumb to their retirement spending (for example simply allowing spending to track income). The evidence is less convincing for being based on cross-section analysis and there are important differences between sub-groups that suggest possible wealth shocks, at least for some. For example, anticipated declines in spending at retirement vary little with income, wealth and health status, but the actual falls in spending are far greater for those who, post-retirement, are in the bottom income and wealth quartiles and self-report poor health. Using data from the earlier Retirement History Survey, which does link expected and actual changes in spending for the same people, Haider and

Stephens (2003) show there is little correlation between the two – the fall in spending that occurs in retirement is broadly the same whatever people's prior expectations.

Miniaci, Monfardini and Weber (2003) investigate whether the fall in spending at retirement in Italy can be explained by shocks to retirement wealth or by a leisure-substitution hypothesis. They rule out negative wealth shocks by presenting evidence that shows that the pre-retired were able to predict their income replacement rate at retirement very accurately.[2] They find that inclusion of controls for leisure (number of workers in the household, employment and occupation of the head) considerably reduces the size of the retirement consumption fall – although a small and gradual reduction in consumption remains. It is, however, debatable whether other (typically temporary) non-employed states (also used by Banks, Blundell and Tanner (1998) as a control for leisure) are really comparable to (typically permanent) retirement.

This paper uses data from the British Household Panel Survey to shed further light on the retirement consumption puzzle in two ways.

- First, recognising the varied experience of retirement, it compares the spending profiles for different groups of people – in particular those for whom retirement can be largely characterised as involuntary (and hence likely to be accompanied by a negative wealth shock) and those for whom retirement can be largely characterised as voluntary (and hence unlikely to be accompanied by a negative wealth shock).
- Second, as well as looking at what happens to spending, the paper compares profiles of (a measure of) well-being for the different groups for additional insights into the nature of the retirement-consumption puzzle.

Each of these ideas is discussed in turn.

Voluntary/ involuntary retirement

Previous studies of retirement in the UK have highlighted the very different retirement experiences that people have (see Tanner (1998) and Marmot et al (2004)). There is considerable variation in the timing of retirement, and in particular whether people retire before, at or after the state pension ages (currently 65 for men and 60 for women). There is also variation in the reasons why people retire. When asked to give the main reason why they left work before the state pension age (or the normal retirement age in their job), for example, people give reasons that broadly reflect involuntary retirement ('ill-health' or 'compulsory redundancy') or voluntary

[2] Of course, accurate prediction of the replacement rate at retirement does not completely preclude negative wealth shocks since retirement wealth depends also on annual uprating as well as initial level of income. Over this period there was a change in uprating from indexation to earnings to indexation to prices.

retirement ('redundancy/retirement with reasonable financial terms' or 'to enjoy life while young and fit'). On the basis of the reasons given, it is reasonable to assume that those who retire involuntarily are more likely to experience a negative wealth shock than those for whom retirement is voluntary.

In turn, these different retirement experiences tend to be reflected in different pathways into retirement and, in particular, the way individuals report their non-working status: Those who leave work involuntarily are much less likely to describe themselves as retired immediately they have left work.

There are a number of possible ways to define retirement. One is to take a purely objective approach and define retirement as the point of permanent exit from employment. A second is to use individuals' self-reported status, conditional on them not working. Using this definition, there will be some people who have permanently left employment, but who don't self-report themselves as retired. For those who retire involuntarily, the pathway into retirement often takes them through a spell when they describe themselves as unemployed or, more usually, sick,[3] before they actually report that they are retired. A third is to use individuals' self-reported retirement status, irrespective of whether they are working or not. In this case, as well as some people who have permanently left employment who are not retired, there will be some people who are retired who have not yet permanently left employment.

Figure 2 illustrates the different pathways into retirement for men using data from the BHPS. In the survey, self-reported retirement is defined conditional on not-working (ie definition 2 above). As the figure shows, in the run-up to 'retirement', an increasing proportion of men report themselves as unemployed or sick.

Tanner (1998) and Marmot et al (2004) show that those for whom retirement is largely involuntary are more likely to self-report themselves as unemployed or sick before they report themselves as retired. If retirement is defined to mean permanent exit from employment, then, very broadly, those for whom retirement is involuntary are more likely to begin their retirement with a spell being unemployed/sick. By contrast, those for whom retirement is voluntary are more likely to begin their retirement by self-reporting themselves as retired.

[3] In the UK, the tendency to self-report as sick may be exaggerated by the use of incapacity benefit as an early retirement vehicle (see Blundell, Meghir and Smith (2002) and Disney, Meghir and Wakefield (2003)).

Figure 2: Employment status prior to self-reported retirement

In this paper, retirement is defined as the point of permanent exit from employment (definition 1 above). But, differences in the timing of retirement and how individuals self-report their employment status when they retire, are used to allocate people into groups that, broadly, reflect the extent to which retirement is voluntary or involuntary. Four groups are defined as follows (summarised in Table 1):

- Group 1 – individuals who leave employment at the state pension age and immediately self-report themselves as being retired. For this group, retirement is largely assumed to be voluntary. However, it is possible that some people in this group may have been forced into leaving work as a result of age discrimination (and because of the absence of anti age discrimination laws consider themselves unlikely to find another job).
- Group 2 – individuals who leave employment before the state pension age and immediately self-report themselves as being retired. For this group, retirement is largely assumed to be voluntary. They could have carried on working until state pension age, but chose to self-report themselves as retired.
- Group 3 – individuals who leave employment after the state pension age and immediately self-report themselves as being retired. For this group, retirement is largely assumed to be voluntary since they were able to carry on working past the state pension age.
- Group 4 – individuals who leave employment before the state pension age and report themselves as being unemployed or sick before later self-reporting themselves as retired. For this group, retirement is largely assumed to be involuntary.

Table 1: Summary definition of groups

	Age of retirement	Self-reported employment status after retirement	Retirement largely voluntary or involuntary?
Group 1	65	Retired	Voluntary
Group 2	Less than 65	Retired	Voluntary
Group 3	More than 65	Retired	Voluntary
Group 4	Less than 65	Unemployed or sick	Involuntary

An alternative way to allocate individuals into groups would be to use their own self-reporting of the reason for retiring. This information is collected in a detailed retirement module in wave 11 of the BHPS. However, the reasons given may be subject to recall error (since the questions are asked several years after retirement) and/or post-hoc rationalisation. Assignment of individuals into the four groups is therefore akin to an instrumental variable solution to these measurement error problems. This issue is returned to in more detail in the discussion of estimation in section 3 below.

Assuming that there is systematic variation between the groups in the extent to which retirement is voluntary or involuntary (especially between groups 1 – 3 and group 4), this paper investigates whether this is reflected in differences in the path of spending around retirement and, hence, whether negative wealth shocks play a role in explaining the retirement-consumption puzzle.

Well-being and consumption

As well as looking at what happens to spending at retirement, this paper also looks at what happens to (a measure of) well-being for further insights into the nature of the retirement-consumption puzzle. Very broadly, if there are negative wealth shocks it is likely that both spending and well-being will fall; if leisure is substituted for spending, then spending will fall, but well-being should be far smoother.

While there has been some recent advocacy of the view that economists should start focusing more on the ultimate goal of happiness itself (see Oswald (1997) and Layard (2003)), this approach is open to challenge on two grounds. First, that there is no universally agreed definition of what well-being is (or how to measure it in surveys) and second, that the life-cycle model implies that (expected) marginal utilities, not the level of utility or well-being, should be smoothed.

The response to the first challenge is to appeal to psychologists' view of well-being. Most importantly, they believe that it is something you can ask people about and

receive reliable responses. As to what kinds of questions to ask, Argyle (1989) proposes that well-being or happiness has three key elements – satisfaction, joy and psychological distress (or disutility).[4] This paper focuses on the third of these – psychological distress (or disutility) – for which the General Health Questionnaire (GHQ) assessment contained in the BHPS is seen as being "one of the most reliable indicators" (Argyle (1989)). Clearly this is only a partial measure and, arguably, only a partial proxy (although, to a lesser extent, the same argument applies to non-durable spending as a measure of consumption). However, since the aim in this paper is to compare profiles of well-being of different groups of retirees then, so long as the measurement error problems do not vary systematically across the groups, the comparisons themselves should be valid, even if it is not possible to draw any firm conclusions about the level of well-being.

The second challenge is that lifetime utility maximisation implies that marginal utilities should be smoothed, not that the level of utility itself should be smoothed. If retirement enters the utility function as an additive component, for example, there is no expectation that the level of well-being should be smoothed through retirement. However, in order to draw inferences about the nature of the retirement-consumption puzzle from a comparison between the groups, it is not necessary to assume that retirement has no separate effect on well-being, only that the effect of retirement is the same for all groups. In other words, it is necessary to assume that any differential effect of involuntary retirement on well-being occurs as a result of the associated negative wealth shock, not as a result of the retirement. The analysis controls for health which is likely to affect both the nature of retirement and well-being.

Previous empirical studies looking directly at the effect of retirement on measures of well-being are fairly inconclusive, but tend to point to retirement having a positive effect. In his general survey of the factors associated with differences in well-being, Oswald (1997) reports that retirement has been found to be a positive risk factor, but many of the studies on which this conclusion is based use cross-section data and fail to control for the possible endogeneity of retirement with respect to well-being.

In a more detailed study, Charles (1999) used panel data from the US Health and Retirement Survey to look at the impact of retirement on two measures of well-being – whether someone had recently been depressed and whether they recently felt lonely. The use of panel data controls for any fixed effects (ie people with permanently high well-being may be more or less likely to retire); the paper also controls for shocks to well-being that may be correlated with retirement by instrumenting retirement with policy changes to the social security system and the individual's exposure to a

[4] Health is seen either as a possible fourth component of happiness, or as factor that is strongly correlated with the other three.

mandatory retirement rule in their job. The paper finds a small, but significant fall in the probability of being depressed associated with retirement. In this paper, the use of panel data will similarly control for fixed effects. The next section discusses in more detail the proposed estimation strategy.

3. Estimation

A fixed effects estimator is used to estimate the differential impact of retirement on consumption for the different groups. This is derived from a marginal-utility-of-wealth-constant consumption demand function, or Frisch function (see Browning, Deaton and Irish (1985) and Blundell and Macurdy (1999))

Consumers are assumed to choose consumption and leisure according to the value function:

$$V(A_t, t) = \max\{U(C_t, L_t, X_t) + \delta E[V(A_{t+1}, t+1)]\}$$

subject to the following budget constraint:

$$A_{t+1} = (1+r)(A_t + B_t + W_t H_t - C_t)$$

where δ is the consumer's discount rate, A_t is total wealth, C_t is consumption, L_t is leisure, X_t is a vector of demographics, r is the (constant) interest rate, B_t is unearned income, W_t is the wage rate and H_t is number of hours worked.

This yields the following first-order-condition for the marginal utility of consumption and the marginal utility of wealth, λt ($=\partial V / \partial A_t$):

$$U_C(C_t, L_t X_t) = \lambda_t$$

$$\lambda_t = \delta E_t[\lambda_{t+1}(1+r)]$$

implying a consumption demand function of the form, $C_t = C(\lambda_t, W_t, X_t)$.

This allows consumption demand to be expressed as a function of an individual's current characteristics (including wages) and a single statistic – the marginal utility of wealth – capturing all other (expected) future information that determines the level of consumption today.

With uncertainty, shocks will be reflected in changes in the marginal utility of wealth from one period to the next. It is possible to express the stochastic process for the marginal utility of wealth as follows:

$$\ln \lambda_t = b_t^* + \ln \lambda_{t-1} + \varepsilon_t^* = \sum_{j=1}^{t} b_j^* + \ln \lambda_0 + \sum_{j=0}^{t} \varepsilon_j^*$$

(where b^*_t depends on the discount factor, the interest rate and the moments of the forecast error ε^*_t).

With this specification, the marginal utility of wealth can be captured by an individual fixed effect, λ_0, plus a function of age plus a random error term, reflecting expectational error in the current period. Thus the level of consumption demand can be modelled as a function of the individual's current characteristics (including wages), age and a fixed effect.

In this paper, log spending is food is modelled as a function of time-varying characteristics, X_1, including health and marital status, fixed characteristics, X_2, including including education, occupation and pension status (ie whether the individual has an occupational or personal pension), age (A) and a time trend. The error term has two components reflecting unobserved fixed characteristics that may affect the marginal utility of wealth and time-varying shocks. Wages are not included directly, but are assumed to be determined by the characteristics X_1, X_2 and age.

$$(1)\ \ln C_{igt} = \sum_{g=1}^{4} \alpha_{11g} G_{ig} + \sum_{g=1}^{4} \alpha_{12g} G_{ig} R_{igt} + \beta_{11}' X_{1igt} + \beta_{12}' X_{2ig} + \gamma_1 A_{igt} + \delta_1 T + \omega_{1ig} + u_{1igt}$$

G is an identifier denoting which group the individual belongs to (and hence whether retirement is voluntary or involuntary), while variable *R* denotes retirement (defined as permanent exit from employment). In the regressions, this is included in two different ways – by the transition into retirement (ie R = 1 only during the first period in which the individual leaves employment), and by the state of retirement (ie R = 1 if the individual is retired). The former would capture one-off changes in consumption at retirement, the latter would capture longer-term changes in consumption associated with retirement if, for example, it takes time for people to adjust spending after retirement.

The key parameters of interest are the coefficients on the set of group-retirement interaction terms (i.e. $\sum_{g=1}^{4} \alpha_{igt}$). As discussed, within the context of the life-cycle model there are two potential explanations for why consumption may fall at retirement, with different implications for the spending of the different groups at retirement. One is because consumers substitute increased leisure time for spending, the other is because of negative wealth shocks. Changes in health associated with retirement may also cause consumption to change, but health status is included directly as a control variable. In the case of leisure substitution, retirement should have a similar effect on all groups, and the coefficients should not be significantly different. Negative wealth shocks, however, should only affect consumption where

retirement is involuntary, and there should be a significant difference between the coefficients for groups 1 – 3 and group 4.

The impact of retirement on well-being is also modelled using a similar empirical specification. Rather than modelling well-being directly as a function of consumption, a reduced form model is used, ie including the same set of characteristics that are assumed to affect consumption, but allowing the coefficients to vary:

(2)

$$WB_{igt} = \sum_{g=1}^{4} \alpha_{12g} G_{ig} + \sum_{g=1}^{4} \alpha_{22g} G_{ig} R_{igt} + \beta_{12}' X_{1igt} + \beta_{22}' X_{2ig} + \gamma_2 A_{igt} + \delta_2 T + \omega_{2ig} + u_{2igt}$$

Again, the key parameters of interest are the coefficients on the group-retirement interaction terms. The assumption is that involuntary retirement will have a different impact on well-being to voluntary retirement because of the associated negative wealth shock, causing a significant difference between the coefficients for groups 1 – 3 and group 4.

As discussed in the previous section, assignment of individuals to the different groups depends when they are observed to retire (relative to the state pension age) and how they self-report their employment state post-retirement. It is assumed that which group an individual is in will depend on their time-varying, and fixed observable characteristics, unobservable characteristics and a random error term which is assumed to reflect other events which result in unexpected/ involuntary retirements (such as compulsory redundancy).

(3) $G_{ig} = f(Z_{1ig}, Z_{2igt}, T, \omega_{ig}, u_{igt})$

In order for the groups to act as a valid instrument for voluntary/ involuntary retirement, two conditions must be satisfied. First, there must be correlation between group membership and voluntary/ involuntary retirement (instrument relevance) and, secondly, group membership must be uncorrelated with the error terms in equations (1) and (2) (instrument exogeneity).

Data from wave 11 of the BHPS support the relevance of the groups as instruments. In wave 11, respondents were asked detailed questions about whether they retired early/late/at the normal age, and at their main reason for retiring early. These responses are summarised in Table 2. As discussed above, the reporting may be subject to error because of recall or post-hoc rationalisation. But, there are clear differences across the groups. As would be expected, of the two groups of early retirees (groups 2 and 4), those in group 4 are far more likely than those in group 2 to

report early retirement because of ill-health, while those in group 2 are far more likely than those in group 4 to report early retirement to "enjoy life while young and fit".

Table 2: Main reason for early retirement

	Own ill-health	Redundant – no choice	Offered reasonable financial terms	To enjoy life while young and fit
Group 2	10%	16%	36%	28%
Group 4	30%	16%	19%	6%

What about instrument endogeneity? Using fixed effects estimation, the problem of correlation between unobservable characteristics ω_{ig} in equation (3) and ω_{1ig} and ω_{2ig} equations (1) and (2) disappears. But there will be an endogeneity problem if the random error term u_{igt} is correlated with u_{1igt} or u_{2igt}. This will be the case if there are factors that cause unexpected/ involuntary retirement that also result in changes in psychological well-being and/or consumption that cannot be explicitly controlled for. Clearly, a change in health is likely to be related both to the nature of retirement and to well-being and consumption, but this is controlled for explicitly. It is assumed that other factors that may result in involuntary retirements (such as redundancy or being dismissed as a result of age discrimination) do not have a direct effect on well-being or consumption other than through their effect on being retired.

4. Data

The data are drawn from the British Household Panel Survey (BHPS). This is a panel dataset which has been collecting information on the same sample of approximately 10,000 individuals each year since 1991. The BHPS covers all ages and, compared to the US Health and Retirement Survey for example, has a smaller number of individuals in the relevant age range for studying retirement. This paper selects the cohort of men aged 45 – 64 in the first year of the survey, a total sample of around 2,000.

In the longer term, The English Longitudinal Survey of Ageing, the new panel dataset focusing on older individuals, will provide richer information on a larger sample for studying retirement in the UK (see Marmot et al (2004)). For the time being, however, the BHPS has a number of strengths:

- after running for more than ten years, the dataset contains a reasonable number of retirements – more than 500 men retire over the entire period – and is increasingly being used to study retirement;[5]
- nearly half the sample has been in the survey for the entire period, giving ten waves of information to use in the analysis.[6] This is particularly important for assigning individuals into the different groups, a process which relies on information on employment before and after retirement; and
- the BHPS contains a wide number of variables, including information on spending, well-being, income and health. As discussed further below, the variables are often not as detailed as we would like (the information on spending, for example is very limited compared to the Family Expenditure Survey). However, given that the main purpose of this paper is to compare behaviour across different groups, this is arguably less of a problem than it otherwise would be.

Measures of retirement in the BHPS

There are two possible ways to define retirement in the BHPS. One is to use the point at which individuals self-report themselves to be retired. The other is to use the point at which they permanently leave employment. For the analysis here, the second definition is used, with the self-reported employment state used to allocate individuals to different groups.

Figure 3 shows the distribution of retirement ages for men. As previous studies of retirement in the UK have shown, there is a spike corresponding to the state pension age, but the majority of men retire before this age. Using individuals' self-reported status to define the age of retirement, the spike at 65 is bigger, reflecting the importance of the state pension age for people's definition of retirement.

[5] See for example Bardasi, Rigg and Jenkins (2000) for an analysis of retirement and poverty using the BHPS and Disney, Emerson and Wakefield (2003) for an analysis of ill-health and labour market exits.
[6] Wave 11 available, but not used because of no information on income.

Figure 3: Distribution of retirement ages

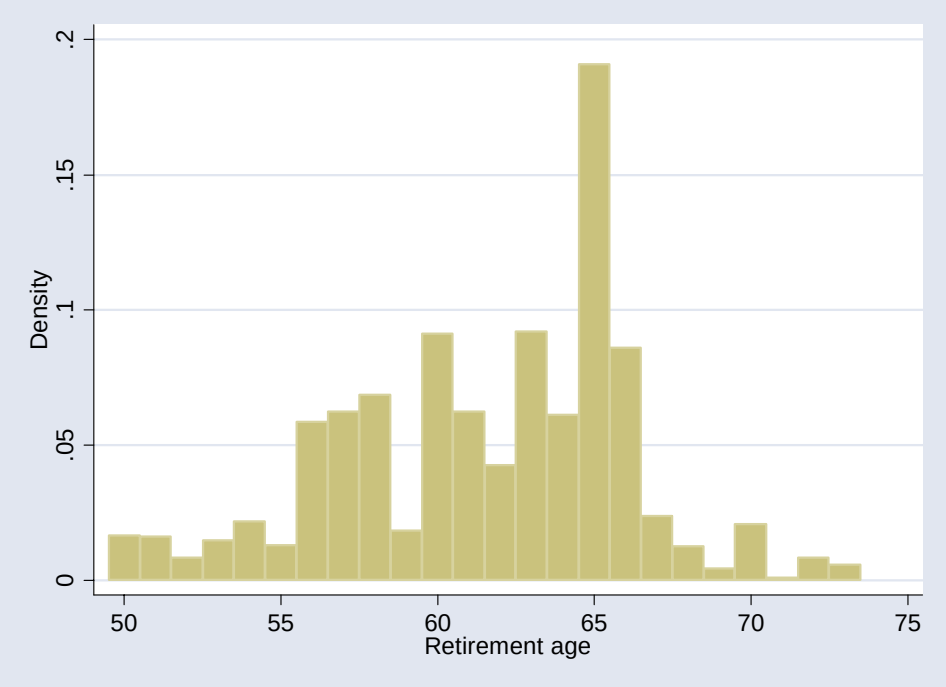

On the basis of the age at which individuals leave employment and their self-reported employment state, the individuals in the BHPS are assigned to one of four groups as follows:

Table 3: Group assignment criteria and sample sizes

Group	Assignment criteria	N
1	Observed working for at least two periods prior to retirement; retires at the state pension age (65); self-reports retired status after leaving work; not observed to re-enter employment.	51
2	Observed working for at least two periods prior to retirement; retires before the state pension age (65); self-reports retired status after leaving work; not observed to re-enter employment.	111
3	Observed working for at least two periods prior to retirement; retires after the state pension age (65); self-reports retired status after leaving work; not observed to re-enter employment.	42
4	Observed working for at least two periods prior to retirement; retires before the state pension age (65); self-reports unemployed/ sick status after leaving work; not observed to re-enter employment.	65

Taken together, the individuals allocated to the four groups make up a relatively small proportion of the total retirements observed in the BHPS. In other cases, there are two few observations before or after retirement to be able to make an assignment (eg

some retire right at the beginning or end of the survey period, some retire from unemployment without being observed to work). This is likely particularly to affect group 4, which requires someone to be observed leaving work and subsequently self-reporting themselves as retired. In practice, the number of men following this pathway into retirement is likely to be greater than the numbers here suggest. However, while the fairly tight criteria for allocating individuals to the different groups may reduce sample size, the advantage is that membership of the groups should be more homogeneous, making any differences in behaviour more pronounced.

The characteristics of the groups are fairly distinct, as shown in Table 4.

Table 4: Characteristics of the four groups

	Professional occupation	Manual occupation	Occup'al pension	Public sector	No educ quals	Divorced/ widowed/ separated
1	0.18	0.78	0.49	0.06	0.46	0.12
2	0.37	0.40	0.85	0.32	0.28	0.06
3	0.19	0.58	0.21	0.19	0.33	0.13
4	0.17	0.61	0.57	0.20	0.31	0.06

- Individuals in group 1 are predominantly manual workers, overwhelmingly in the private sector, with the lowest average level of educational qualifications. Around half have an occupational pension.
- Individuals in group 2 are predominantly non-manual workers (and include a large number of people in professional/ managerial occupations) with an occupational pension. They are the most likely to work in the public sector and have the highest average level of educational qualifications.
- Individuals in group 3 tend to be in manual occupations, with the overwhelming majority not having an occupational pension.
- Interestingly, individuals in group 4 (for whom retirement is assumed to be largely involuntary) tend to fall somewhere in between these three groups. The proportion in manual jobs is not as low as in group 2, but not as high as in group 1. Similarly the proportion with an occupational pension is not as high as group 2, but not as low as group 3. The level of educational qualifications is average, and they are neither under- nor over-represented in the public sector.

Measures of spending in the BHPS

The BHPS collects limited information on spending on a small number of goods – food in, gas and electricity and durables – in all waves. This paper focuses on the information on food spending. Households are asked "approximately how much does your household usually spend each week in total on food and groceries." In the first wave, they are asked to give a continuous answer; in subsequent waves, they are asked to say in which band (out of 12) their weekly food spending lies. They are told to include all food, bread, milk, soft drinks etc, but asked to exclude pet food, alcohol, cigarettes and meals out. Take-aways eaten in the home are, however, included.

To obtain a weekly spending figure, each individual is assigned the mid-point of their reported band each year, adjusted for inflation in food prices.[7] Comparisons with the more detailed spending information in the Family Expenditure Survey shows that mean food spending in the BHPS is slightly higher than in the FES.[8] In part this may reflect the fact that there are fewer observations in the lowest bands in the BHPS (respondents may ignore atypical weeks when they spend very little). Alternatively, respondents may include other grocery items that they regularly buy at the supermarket such as washing powder, toilet roll etc. When these items are included in the FES spending figures, the two sets of numbers are very similar.

If people do substitute leisure for spending when they retire, what would happen to this measure of food spending? To the extent that people eat out less and produce more meals at home, spending on food may rise. However, to the extent that people substitute home-cooked meals for more expensive take-aways, this measure of food spending may fall. A priori, there is no clear direction for any change at retirement.

Measures of well-being in the BHPS

The BHPS contains a number of measures of well-being, including life satisfaction, but only one measure – the General Health Questionnaire (GHQ) – in all waves.

The GHQ is designed to measure psychological distress or disutility which is seen as one of three key elements contributing to an individual's overall well-being or happiness (see Argyle (1989)). Clearly, it is not an ideal measure of what economists may think of as utility. However, it has been used in previous studies (for example Oswald (1997)) and shown to vary systematically with risk factors that may be thought to increase or reduce utility (such as unemployment).

The GHQ has 12 individual measures covering concentration, loss of sleep, whether the individual feels they play a useful role, whether they are capable of making

[7] For wave 1, the continuous answers are first banded, and then the midpoints are assigned.
[8] To calculate the FES figures, the continuous weekly spending figures are converted into bands and then mid-points as in the BHPS.

decisions, whether they are constantly under strain, whether they have problems overcoming difficulties, whether they enjoy day-to-day activities, their ability to face problems, whether they are unhappy/depressed, whether they are losing confidence, their belief in their self-worth and their general happiness.

In all cases, questions are asked relative to the individual's usual state. So, for example, individuals are asked "Have you recently been feeling reasonably happy, all things considered" and given the following four options: 1 = more than usual; 2 = same as usual; 3 = less so than usual; 4 = much less than usual. The responses are always numbered such that a higher number reflects greater disutility. This paper uses the aggregate Likert index which recodes the responses from 0 – 3 and sums the twelve measures to produce a single index with a range of 0 – 36.[9]

Measures of health in the BHPS

As with well-being, the BHPS contains a large number of variables measuring individuals' health, but only a limited number of health variables in all ten waves (see Disney, Emerson and Wakefield (2003) for a more detailed analysis of health and labour market exit using the BHPS data). In this paper, the main health variable is number of health problems reported by the individual in each year (out of a maximum of 13, including arms, legs and hands; sight; hearing; skin conditions/ allergy; chest/ breathing; heart/ blood pressure; stomache/ digestion; diabetes; anxiety/ depression; alcohol & drugs; epilepsy; migraine and other). Information on whether the individual is registered as disabled is also available in all waves.

5. Results

This section presents the main results on income, spending, well-being and health around retirement for the four different groups. Table 5 summarizes means for each of the variables before and after retirement, while Figures 4-7 present the same information slightly differently, showing the paths of the variables in each of the four years before, and the four years after, retirement. In the figures, year 0 represents the first year in which the individual is retired.

While the numbers presented in Table 5 and Figures 4-7 clearly highlight the key differences between the groups, they are not conclusive – firstly because they fail to control for other factors (age, for example is related to well-being and varies systematically across the groups) and secondly, because of compositional changes (i.e. the sample one year after retirement is not necessarily the same as the sample two years after retirement and so on).

[9] The main alternative is the Caseness index, where scores of 0 or 1 are re-coded as 0, and scores of 2 or 3 are re-coded as 1. This produces a narrower aggregate index from 1 – 12.

Tables 6 and 7 therefore presents the results of fixed effects regressions on the two key variables, log real food spending and GHQ score. In the regression analysis, the retirement variable is interacted with group dummies to capture any significant differences across the groups. Because retirement for individuals in group 4 is thought largely to be involuntary (compared to the other three groups) this is chosen as the reference group. As discussed above, retirement is specified in two ways – first as the transition into retirement (ie R=1 only during the first period in which the individual leaves employment), and second, by the state of retirement (ie R = 1 if the individual is retired). The former picks up one-off changes in spending and well-being associated with retirement; the latter, longer-term adjustment. All regressions include a full set of age- and time dummies and controls for marital status and health (measured by number of health problems and whether or not the individual is officially registered as disabled).

Table 5: Mean income, spending, disutility and health problems

	Before retirement	After retirement
Equivalised real weekly income		
Group 1	£322	£238
Group 2	£505	£358
Group 3	£395	£283
Group 4	£333	£287
Real weekly food spending		
Group 1	£53	£44
Group 2	£57	£56
Group 3	£44	£45
Group 4	£56	£51
Disutility		
Group 1	9.45	8.71
Group 2	10.23	9.73
Group 3	9.33	8.85
Group 4	10.50	11.05
Health problems		
Group 1	1.00	1.32
Group 2	0.89	1.40
Group 3	1.10	1.53
Group 4	1.15	1.86

Income

As would be expected from the characteristics of the four groups shown in Table 4, they have very different levels of income before and after retirement. Individuals in group 2 have the highest average income pre- and post-retirement, individuals in group 1, the lowest.

Figure 4: Mean equivalised weekly real household income (£)

Group 1

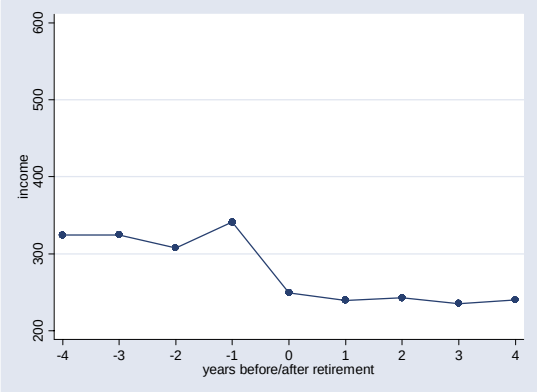

Group 2

Group 3

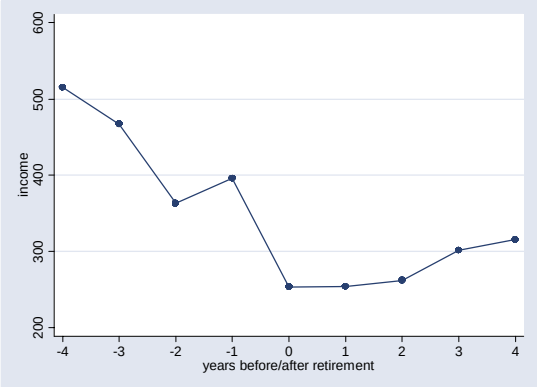

Group 4

For groups 1, 2 and 3, the replacement rate is broadly the same (average post-retirement income is, respectively, 74%, 71% and 72% of average pre-retirement income). For group 4, the replacement rate is higher (86%). For all four groups, Figure 4 shows a clear one-off change in income at retirement. For group 2, income rises slightly in the years after retirement which may reflect post-retirement eligibility for the state pension (from age 65).

Food spending

The groups also have quite different profiles for food spending around retirement. The mean spending figures in Table 5 show that groups 2 and 3 broadly maintain their level of spending on food before and after retirement, albeit at very different levels.

In the case of groups 1 and 4, average food spending after retirement is, respectively 17% and 10% lower than the average pre-retirement level. Figure 5 shows a fairly steep decline in spending on food for group 1, beginning in the year of retirement, and a more gradual decline for group 4. The nature of the question in the BHPS (which asks about usual spending) may be more likely to produce a smoothed profile.

Figure 5: Mean weekly real spending on food (£)

Group 1

Group 2

Group 3

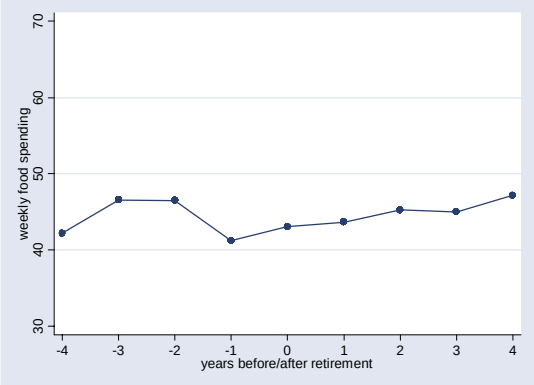

Group 4

These patterns across the groups are confirmed by the results of the regression analysis reported in Table 6. There are two columns of results reflecting inclusion of retirement as a transition variable and as a state variable. Because the BHPS asks about "usual spending on food", it is more likely that any change in food spending will be picked up by the state variable than by the transition variable, particularly if food spending adjusts gradually, and this is indeed the case. There is no significant change in reported usual food spending associated with the transition into retirement for any of the groups. When retirement is included as a state variable, however, there is a significant fall in food spending associated with retirement of around 7% for group 4. The interaction term for group 1 is insignificant, implying a fall in food spending of a similar magnitude. For both groups 2 and 3, the interaction terms are

significant and positive and similar in absolute magnitude to the decline for group 4, implying that both groups smooth their food spending.

Table 6: Regression results for (log) food spending – fixed effects estimation

Dependent variable = (log) weekly real spending on food

	Retirement as a transition variable		Retirement as a state variable	
	Coefficient	SE	Coefficient	SE
Retired (0/1)	0.3192	0.3796	-0.0739	0.0279
Retired*Group1	0.2346	0.5665	0.0372	0.0415
Retired*Group2	-0.2169	0.4316	0.0809	0.0319
Retired*Group3	-0.2295	0.5948	0.0924	0.0439
Divorced/ widowed/ separated (0/1)	-0.5907	0.6851	-0.4542	0.0499
Number of health problems	0.4366	0.0998	0.0176	0.0072
Disabled (0/1)	0.5868	0.5614	-0.0589	0.0399
Age dummies	Yes		Yes	
Time dummies	Yes		Yes	
No. individuals	269		269	

The regression results confirm that there are significant differences in the spending profiles of the groups, a finding which is hard to reconcile with the leisure-substitution hypothesis. The patterns for groups 2, 3 and 4 are consistent with a fall in spending linked to a negative wealth shock – ie a fall in spending when retirement is involuntary, but smoothed spending when retirement is voluntary. Group 1, however, does not fit this pattern; some possible explanations are explored in the conclusions.

Well-being

For groups 1, 2 and 3, the mean GHQ scores shown in Table 5 show a lower level of disutility (hence a higher level of well-being) after retirement compared to before. This may reflect an age effect (in general, GHQ scores tend to fall with age), as well as/ instead of a pure retirement effect. For Group 4, however, the level of disutility is higher post-retirement than before.

Figure 6 shows that groups 2 and 4 experience sharp changes in GHQ scores coinciding with retirement. For group 2, there is a sharp fall in GHQ score at retirement – ie a fall in disutility – followed by a slight increase. For group 4 there is a sharp increase in GHQ score, followed almost immediately by a fall of around the same magnitude. It is worth pointing out that the GHQ asks about indicators of psychological distress relative to usual. It is hard to know exactly how people will

interpret “usual”, but if their reference point is the previous year, such a one-off change followed by a return to (roughly) the same level would imply a one-off decline, with no further deterioration. By the same measure, the profile of group 2 would imply a one-off improvement, followed by further (smaller) improvements.

Figure 6: Mean disutility score

Group 1

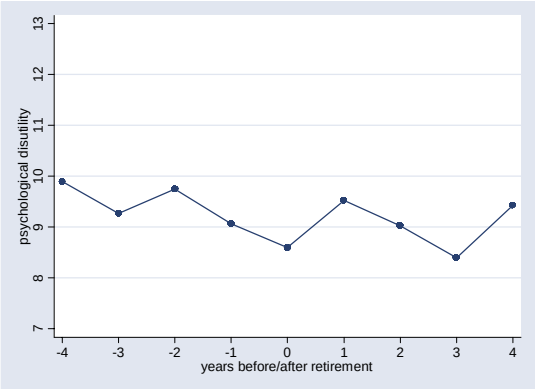

Group 2

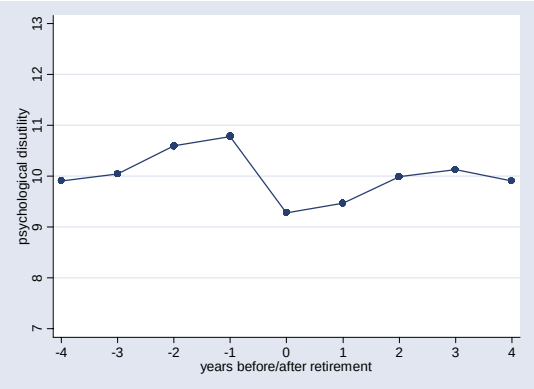

Group 3

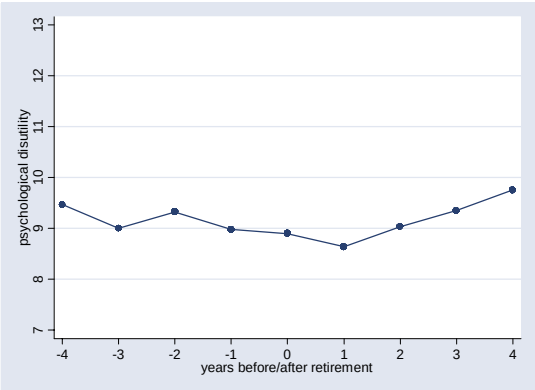

Group 4

Of course, the changes in GHQ scores at retirement, particularly for group 4, may simply reflect changes in health that are linked to labour market exit, rather than an effect of retirement per se. As shown in Table 2, ill-health is one of the main reported reasons for early retirement among this group. Figure 7 therefore shows the profiles of the average number health of problems for each of the groups. The profile of group 4 does show an increase in average number of health problems around retirement, but this does not coincide exactly with the increase in disutility. Moreover, group 1 shows a similar increase in average health problems of a similar magnitude, at a similar point in time, but the profile of GHQ scores for group 1 is quite different to that of group 4. For group 2 there is no fall in average number of health problems that might account for the fall in disutility. Clearly health is likely to be strongly linked to GHQ scores, and is controlled for in the regressions reported in

Table 6, but the group profiles presented in Figure 7 suggest that health is not the only explanation.

Figure 7: Mean Health problems

Group 1 Group 2

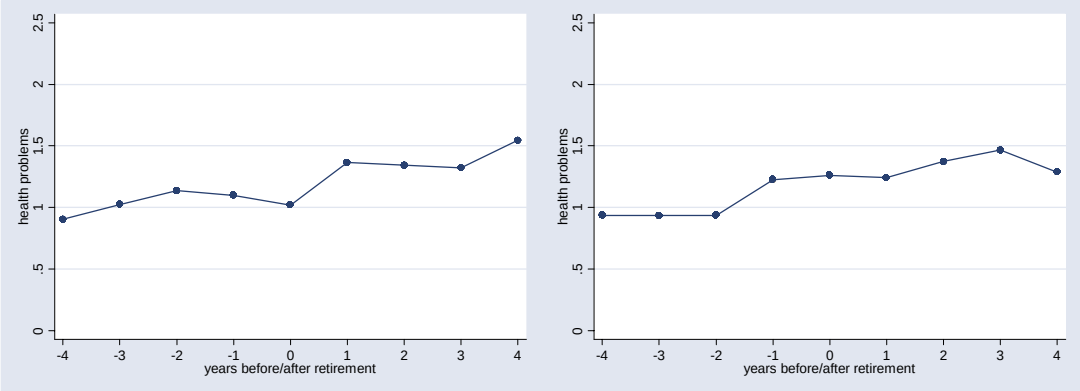

Group 3 Group 4

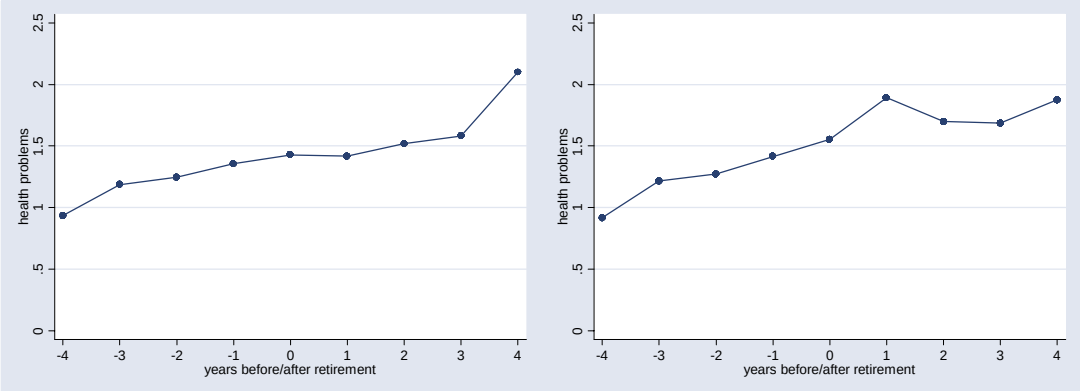

Table 7 reports the results from the regression. Again, two sets of results are reported corresponding to the inclusion of retirement as a transition and state variable respectively. Since the questions ask respondents to report their level of disutility "compared to usual", it is more likely that a change in disutility would be picked up by the transition variable than by the state variable, and this is indeed the case. The results show that there is a significant fall in well-being at retirement for individuals in group 4 (a significant increase in GHQ score of 1.6 points), controlling for health. Compared to this reference group, the coefficients on the interaction terms for the three other groups are negative and significant – with the magnitudes implying a slight increase in well-being. Controlling for other factors, the change for group 2 does not appear to be different to that for the other two groups. When retirement is included as a state variable, there is no significant change for any of the groups.

The regression results confirm that there is a significant difference between the profile of well-being of group 4 and the other three groups. The change in well-being associated with involuntary retirement is significantly different to (and worse than) the change associated with voluntary retirement, which is broadly consistent with a negative wealth shock hypothesis. Of course, the difference may directly reflect the factors that cause people to retire involuntarily – ie ill-health or redundancy. While a number of health variables have been included in the regression, they may not be adequate controls for health. It is also possible that redundancy may have a direct effect on well-being other than through the associated negative wealth shock – this is not something that can be controlled for explicitly. Nevertheless, as with the results on food spending, these findings on well-being imply that understanding the voluntary/ involuntary nature of retirement is important for resolving the puzzle.

Table 7: Regression results for well-being – fixed effects estimation

Dependent variable = Aggregate GHQ (disutility) score

	Retirement as a transition variable		Retirement as a state variable	
	Coefficient	SE	Coefficient	SE
Retired (0/1)	1.5850	0.4535	-0.0059	0.0334
Retired*Group1	-2.0198	0.8099	0.0597	0.0599
Retired*Group2	-2.0000	0.5939	-0.0065	0.0438
Retired*Group3	-1.7152	0.7765	-0.0059	0.0572
Divorced/ widowed/ separated (0/1)	-0.5900	0.6833	-0.4480	0.0499
Number of health problems	0.4503	0.0995	0.0168	0.0072
Disabled (0/1)	0.5885	0.5567	-0.0751	0.0396
Age dummies	Yes		Yes	
Time dummies	Yes		Yes	
No. individuals	269		269	

6. Conclusions

This paper has looked at what happens to food spending and well-being at retirement for different groups of retirees. In particular, it has examined whether there are differences between those for whom retirement is largely voluntary (who are therefore assumed to experience no negative wealth shocks) and those for whom retirement is largely involuntary (who are therefore assumed to be more likely to experience negative wealth shocks).

There are two groups – one retiring voluntarily before the state pension age (65) and another retiring voluntarily after the state pension age (65) – whose profiles of spending and well-being are consistent with the life-cycle model in the absence of any negative wealth shock. The two groups are very different in their characteristics (occupation, sector of employment and pension) and in their levels of income and spending before and after retirement. Yet both groups smooth their food spending through retirement, and retirement is not associated with any significant fall in well-being. If anything, both groups experience a slight increase in well-being at retirement.

For another group, retirement is assumed to be involuntary, largely reflecting ill-health and redundancy, and likely to be associated with a negative wealth shock. The profiles of food spending and well-being around retirement reflect this. There is a significant one-off reduction in well-being at retirement (controlling for health) and a reduction in food spending. For this group, resolving the retirement-consumption puzzle appears to lie in the involuntary nature of their retirement (and associated negative wealth shock).

However, there is another group – who retire apparently voluntarily at the state pension age (65) – whose profiles of food spending and well-being do not fit with this pattern, showing a significant fall in food spending after retirement, but no significant change in well-being. So, what possible explanations could account for the observed spending and well-being profiles of this group?

One possibility is that individuals in this group may have wanted to work for longer, but were prevented from doing so by age discrimination.[10] However, this would explain why food spending fell, but not why there was no change in well-being. It is plausible if well-being is defined relative to peers,[11] but if so, the fall in well-being among the other group of retirees would have to be explained.

Another possibility is that, among this group, there is substitution of leisure for consumption – this would account for the fall in food spending with no decline in well-being. The obvious question then, however, is why this would happen uniquely among this group? It is possible to reconcile with the profiles of group 4 (involuntary retirees) – which are consistent with both some substitution of leisure for consumption *and* a negative wealth shock. For those in group 3 who retire after the state pension age, one explanation may lie in slightly higher levels of part-time work (ie some substitution of leisure for consumption has already occurred), although the overwhelming majority continue to work full-time right up until retirement. For

[10] Self-reported retirement may reflect strong social norm associated with state pension age
[11] See Fox and Kahneman (1992)

group 2, who retire voluntarily before age 65 their higher level of post-retirement income may allow them to maintain their level of food spending, but this is slightly at odds with the evidence presented in Hurd and Rohwedder (2003) showing similar anticipated falls in spending across all income and wealth groups.

This paper has taken "two steps forward and one step backwards" towards resolving the retirement-consumption puzzle. No single, simple explanation appears to fit all the facts in the BHPS data. By comparing different groups of retirees, the paper shows that understanding the nature of retirement, and in particular the extent to which retirement is likely to be associated with a negative wealth shock, is key to understanding why spending falls – at least for some. However, there is a group of retirees whose retirement appears to be voluntary, yet who nevertheless experience a fall in spending. To fully resolve the puzzle, further analysis is needed to get a better understanding of the link between the nature of retirement and spending around retirement and/or how different groups engage in the substitution of leisure for consumption.

References

Ameriks, J., Caplin, A. and J. Leahy (2002) "Retirement consumption: Insights from a survey" National Bureau of Economic Research Working Paper No. 8735

Argyle (1989) The psychology of happiness

Banks, J., Blundell, R. and S. Tanner, S. (1998) "Is there a retirement-savings puzzle?" American Economic Review 88:4 769–88.

Bardasi, E., Rigg, J. and S. Jenkins (2000) Retirement and the economic well-being of the elderly: A British perspective, ISER working paper 2000-33

Bernheim, D., Skinner, J. and S. Weinberg (2001) "What accounts for the variation in retirement wealth among US households?" American Economic Review 91:4 832–57

Blundell, R. and T. Macurdy (1999) "Labor supply: A review of alternative approaches" in Ashenfelter, O. and D. Card (eds) Handbook of Labor Economics

Blundell, R., Meghir, C. and S. Smith (2002) "Retirement incentives and labour supply in the UK" Economic Journal

Borsch-Supan, A. and K. Stahl (1991) "Life-cycle savings and consumption constraints: Theory, empirical evidence and fiscal implications" Journal of Population Economics 4(3) 233–55

Browning, M., Deaton, A. and M. Irish (1985) "A profitable approach to labor supply and commodity demands over the life-cycle" Econometrica 53(3) 503–43

Charles, K. (1999) "Is retirement depressing? Labor force inactivity and psychological well-being in later life" University of Michigan mimeo

Disney, R., Emerson, C. and M. Wakefield (2003) Ill-health and retirement in Britain: A panel data based analysis IFS Working paper No. 03/02

Disney, R. and S. Tanner (1999) "What can we learn from retirement expectations data?" Institute for Fiscal Studies Working Paper No. 99/17

Fox and Kaheman (1992) Correlations, causes and heuristics in surveys of life satisfaction, Social Indicators Research 27 221–34

Haider, S. and M. Stephens (2004) "Is there a retirement-consumption puzzle? Evidence using subjective retirement expectations" National Bureau of Economic Research Working Paper No. 10257

Hamermesh, D. (1984) "Life-cycle effects on consumption and retirement" Journal of Labor Economics 2 353–370

Hurd, M. and S. Rohwedder (2003) "The retirement-consumption puzzle: Anticipated and actual declines in spending and retirement" National Bureau of Economic Research Working Paper No. 9586

Layard, R. (2003) "Happiness: Has social science a clue?" Lionel Robbins Memorial Lectures 2003/03

Marmot et al (2004) Health, wealth and lifestyles of the older population in England: The 2002 English Longitudinal Study of Ageing

Miniaci, R., Monfardini, C. and G. Weber (2003) "Is there a retirement-consumption puzzle in Italy?" Institute for Fiscal Studies Working Paper No. 03/14

Oswald, A. (1997) "Happiness and economic performance" Economic Journal 107 1815–1831.

Tanner, S. (1998) "The dynamics of male retirement behaviour" Fiscal Studies