

HEDG Working Paper 05/04

Reporting bias and heterogeneity in self-
assessed health. Evidence from the British
Household Panel Survey

Cristina Hernández-Quevedo
Andrew M Jones
Nigel Rice

June 2005

ISSN 1751-1976

Reporting bias and heterogeneity in self-assessed health. Evidence from the British Household Panel Survey

Cristina Hernández-Quevedo¹, Andrew M Jones¹, Nigel Rice²

Abstract

This paper explores reporting bias and heterogeneity in the measure of self-assessed health (SAH) used in the British Household Panel Survey (BHPS). The ninth wave of the BHPS includes the SF-36 general health questionnaire, which incorporates a different wording to the self-assessed health variable used at other waves. Considerable attention has been devoted to the reliability of SAH and the scope for contamination by measurement error; the change in wording at wave 9 provides a form of natural experiment that allows us to assess the sensitivity of panel data analyses to a change in the measurement instrument. In particular, we investigate reporting bias due explicitly to the change in the question. We show how progressively more general specifications of reporting bias can be implemented using panel data ordered probit and generalised ordered probit models. Then we explore the sensitivity of measures of socioeconomic inequality and of mobility in health to changes in the measurement of SAH.

JEL codes I12 C23

Keywords: self-assessed health, reporting bias, ordered probit, generalised ordered probit, inequality in health

Data from the *British Household Panel Survey* (BHPS) were supplied by the ESRC Data Archive. Neither the original collectors of the data nor the Archive bear any responsibility for the analysis or interpretations presented here. This paper derives from the project “The dynamics of income, health and inequality over the lifecycle” (known as the ECuity III Project), which is funded in part by the European Community’s Quality of Life and Management of Living Resources programme (contract QLK6-CT-2002-02297). Nigel Rice is in part supported by the UK Department of Health programme of research at the Centre for Health Economics. The authors would like to thank Xander Koolman and Fabrice Etilé for their helpful advice and participants at the Health Economists’ Study Group, Glasgow 2004, and the York Seminars in Health Econometrics (YSHE) for their comments.

¹Department of Economics and Related Studies, University of York, Heslington, York, YO10 5DD, United Kingdom: Tel: +44-1904-433790; Fax: +44-1904-433759; Email: amj1@york.ac.uk

²Centre for Health Economics, University of York, United Kingdom

1. Introduction

Like many general population surveys, the British Household Panel Survey (BHPS) includes a measure of self-assessed health (SAH): respondents are asked to rate their own health on a five-point categorical scale. We use eleven years of the BHPS, 1991-2001; this provides the opportunity for longitudinal analysis of the relationship between self-assessed health and socio-economic status, allowing in-depth investigation of dynamics (state dependence) and individual heterogeneity. However longitudinal analyses of SAH face a challenge – the SF-36 questionnaire was included in wave 9 of the BHPS and, with it, the SAH question was re-worded and included a modification to the response categories. Does this allow satisfactory multivariate analyses of SAH that exploit the full span of the longitudinal data? With this challenge there is also an opportunity. Considerable attention has been devoted to the reliability of SAH and the scope for contamination by measurement error; the change in wording at wave 9 provides a form of natural experiment that allows us to assess the sensitivity of panel data analyses to a change in the measurement instrument. We focus on the extent to which the change in the SAH question induces reporting bias and what form any observed reporting bias takes. In particular, we investigate whether different groups of individuals respond to the change in the measurement instrument in different ways. Evidence that they do respond in different ways would be *prima facie* evidence of systematic differences in reporting behaviour. This would have implications for econometric analyses that use SAH as a dependent or independent variable and for analysis of inequalities in health based on SAH. For example, evidence of reporting bias may suggest that standard ordered probit models of SAH are overly restrictive precluding a more appropriate and flexible specification of the effect of regressors on health outcome probabilities. The structure of the paper is as follows. The self-assessed health variable in the BHPS is introduced in Section 2. Section 3 describes the issue of reporting bias and the estimation and testing strategy. The BHPS dataset is described in Section 4. Results are shown in Section 5, concluding the paper in Section 6.

2. Self-assessed health in the BHPS

Self-assessed health is often included in general socio-economic surveys, such as the BHPS and the European Community Household Panel (ECHP). SAH has been used widely in previous studies of the relationship between health and socioeconomic status (e.g., Adams et al. (2003), Benzeval et al. (2000), Deaton and Paxson (1998), Ettner (1996), Frijters et al. (2003), Salas (2002), Smith (1999)) and of the relationship between health and lifestyles (e.g., Kenkel (1995), Contoyannis and Jones (2004)). SAH is a simple subjective measure of health that provides an

ordinal ranking of perceived health status. However it has been shown to be a powerful predictor of subsequent mortality (see e.g., Idler and Kasl (1995), Idler and Benyamini (1997)) and its predictive power does not appear to vary across socioeconomic groups (see e.g., Burström and Fredlund (2001)). Categorical measures of SAH have been shown to be good predictors of subsequent use of medical care (see e.g., van Doorslaer et al. (2000), van Doorslaer, Koolman and Jones (2004)). Socioeconomic inequalities in SAH have been a focus of research (see e.g., van Doorslaer et al. (1997), van Doorslaer and Koolman (2004), van Ourti (2003)) and have been shown to predict inequalities in mortality (see e.g., van Doorslaer and Gerdtham (2003)).

Several studies have analysed the implications of including an age benchmark in the SAH question (Baron-Epel and Kaplan (2001), Kaplan and Baron-Epel (2003) and Eriksson *et al* (2004)). Baron-Epel & Kaplan (2001) analyse the “agreement” between two measures of SAH: a general question and an age-benchmarked question. They use a cross-sectional survey of adult Israeli residents (45-75 years old), based on 793 telephone interviews. The authors found “poor agreement” among older individuals (65-75 years old) with no reported diseases and those individuals with less than 12 years of education and no reported diseases. These two groups reported better health when they were asked to compare their health to people of their own age and sex. However, “excellent agreement” was found between the two questions, for respondents aged 55 – 65 with no disease. Having more years of education was related to more optimistic assessments of health, but only when using the age-related subjective health measure. Individuals provided different assessments of health, particularly for the older and less educated, even when they shared other characteristics such as language, culture, nationality and religion. In a later study, Kaplan and Baron-Epel (2003) analysed the inclusion of an age benchmark when asking individuals about their health. From their sample of 383 Israeli residents, they concluded that individuals report their health by comparing with people of their age, even when this age benchmark has not been included explicitly in the SAH question. The authors argue that the individuals interviewed tried to evaluate their health in an optimistic way. Hence, old and unhealthy individuals compared themselves to people of their own age. However, young, but unhealthy individuals avoided comparing themselves with those of their age, so as not to feel worse. Eriksson *et al* (2004) compare different measures of SAH, taking into account the inclusion of the age benchmark in the question. 8,200 individuals were provided a questionnaire, which included three measures of health: two non-comparative SAH measures one with seven (SAH-7) and one with five (SAH-5) response categories,, and a third measure that included a comparison with others of the same age (SAH-age). Although, they argue that the three measures provide “parallel assessments of subjective health”, that is, they can be applied equivalently, the authors recognize some differences between the measures. In particular, the high level of item

non-response for the oldest age group when answering SAH-7, together with the less skewed distribution for SAH-age, compared to the other two measures.

This subjective measure of health has caused debate in the literature concerning its validity. It has been argued that perceived health does not correspond with actual health (see Bound, 1991), while other researchers have regarded this variable as a valid indicator of true health (see Butler *et al.*, 1987). As a self-reported subjective measure of health, SAH may be prone to measurement error. General evidence of non-random measurement error in self-reported health is reviewed in Currie and Madrian (1999) and Lindeboom (2006). Crossley and Kennedy (2002) report evidence of measurement error in a 5-category SAH question. They exploit the fact that a random subsample of respondents to the 1995 Australian National Health Survey were asked the same version of the SAH question twice, before and after other morbidity questions. The first question was administered as part of the SF-36 questionnaire on a self-completion form, the second as part of a face-to-face interview on the main questionnaire. They find a statistically significant difference in the distribution of SAH between the two questions and evidence that these differences are related to age, income and occupation. This measurement error could be explained by a *mode of administration effect*, due to the use of self-completion and face-to-face interviews (Grootendorst *et al.* (1997) find evidence that self-completion questions reveal more morbidity); or a *framing or learning effect* by which SAH responses are influenced by the intervening morbidity questions.

It is sometimes argued that the mapping of “true health” into SAH categories may vary with respondent characteristics. This source of measurement error has been termed “state-dependent reporting bias” (Kerkhofs and Lindeboom (1995)), “scale of reference bias” (Groot (2000)) and “response category cut-point shift” (Sadana *et al.* (2000), Murray *et al.* (2001)). This occurs if subgroups of the population use systematically different cut-point levels when reporting their SAH, despite having the same level of ‘true health’.

Regression analysis of SAH can be achieved through specifying an ordered probability model, such as the ordered probit or logit. In the context of ordered probit models the symptoms of measurement error can be captured by making the cut-points dependent on some or all of the exogenous variables used in the model and estimating a generalised ordered probit. This requires strong *a priori* restrictions on which variables affect health and which affect reporting in order to separately identify the influence of variables on latent health and on measurement error. It is worth noting that allowing the scaling of SAH to vary across individual characteristics is equivalent to a heteroskedastic specification of the underlying latent variable equation (see e.g., van Doorslaer and Jones (2003)). This is because location and scale cannot be separately

identified in binary and ordered choice models and, in general, it is not possible to separate measurement error from heterogeneity. Attempts to surmount this fundamental identification problem include modelling the reporting bias based on more “objective” indicators of true health (Kerkhofs and Lindeboom (1995), Lindeboom and van Doorslaer (2004)) and the use of “vignettes” to fix the scale (Kapteyn *et al* 2004, Murray *et al.*, 2001)). Lindeboom and van Doorslaer (2004) analyse SAH in the Canadian National Population Health Survey and use the McMaster Health Utility Index (HUI-3) as their objective measure of health. They find evidence of cut-point shift with respect to age and gender, but not for income, education or linguistic group.

In general, the SAH question is measured following an ordinal scale, with possible responses from “very poor” or “poor” to “very good” or “excellent”. The SAH variable that is included in the BHPS has two wordings, depending on the wave. For waves 1-8 and 10 onwards, the SAH variable represents “health status over the last 12 months”. Respondents are asked: “Compared to people of your own age, would you say your health over the last 12 months on the whole has been: excellent, good, fair, poor, very poor?”. However, the SF-36 questionnaire was included in wave 9. The SF-36 questionnaire includes 36 items that measure health across eight dimensions of physical functioning, social functioning, role limitations due to physical problems, role limitations due to emotional problems, mental health, vitality, pain, and general health perceptions (Jenkinson *et al.*, 1996). In this questionnaire, the SAH variable for wave 9 represents “general state of health”, using the question: “In general, would you say your health is: excellent, very good, good, fair, poor?”.

If the two wordings of the question are compared, it is possible to distinguish three main differences. Firstly, the self-assessed health question in wave 9 does not include the age benchmark that is present in the rest of the waves, in which individuals are asked to assess their level of health “compared to people of your own age”. The second is the frame of reference: respondents in waves other than 9 are given a reference period of 12 months, while those at wave 9 are not. The third difference is the way in which the categories have been labelled. Although both questions provide five possible answers to the respondents, the category “very poor” is not available in wave 9, but “very good” is incorporated between “good” and “excellent”.

Figure 1 presents the distribution of SAH for each wave, using a balanced panel of individuals who are observed for all 11 waves. Accordingly, all wave-specific distributions are based on the same sample of individuals followed over time. The different categories are shown on the horizontal axis of each graph, with “1” representing the lowest level of health¹, and “5”, the

¹ This corresponds to “very poor” health at waves 1-8, 10-11 and “poor” health at wave 9.

highest level². The histograms for waves 1-8 and 10-11 follow a similar pattern; a skewed distribution is clear, with the majority of respondents reporting their health as “good” or “excellent”. Around 50 per cent of individuals report a “good” level of health and around 70 per cent, either “good” or “excellent” levels of health. These percentages vary little across waves. Note that although the question is framed in terms of self-rating of health compared to people of ones own age, the distribution of responses across the ten waves indicates a worsening of reported health as the cohort ages. For example, 5% of individuals report “poor” health in wave 1 compared to 7.5% in wave 11; 30% of respondents report “excellent” health in wave 1 compared to 21% in wave 11.

Insert Figure 1 around here.

At wave 9, where the SF-36 version of SAH was used there is a notable change in the distribution of responses. Fewer individuals report their health status within the two highest categories (“excellent” and “very good” health), while a greater proportion of respondents report within the bottom two categories (“poor” and “fair” health). In general, there appears to be a shift to the left of the distribution together with a flattening of its mass points. There is also evidence of a floor effect: individuals who report their health as “poor” in wave 9, may have assessed their health as “very poor” if this category had been available.

3. Multivariate models for ordered responses in longitudinal data

3.1 Reporting bias

Given the change in the distribution of the SAH variable at wave 9, reporting bias could be considered as an explanation for this observation. Reporting bias has been a concern in the literature and can be defined as the differential reporting of health across individuals or groups of individuals with the same health status. The systematic use of different threshold levels by sub-groups of a population reflect the existence of reporting bias (Lindeboom and van Doorslaer, 2004; Murray *et al*, 2001; Groot, 2000; Sadana *et al*, 2000; Kerkhofs and Lindeboom, 1995). These differences may be influenced by, among other things, age, gender, education, income, language and personal experience of illness. Basically, it means that different groups appear to interpret the question within their own specific context and therefore use different reference points when they are responding to the same question.

² Corresponding to “excellent” health at all waves.

In their analysis of reporting bias Lindeboom and van Doorslaer (2004) distinguish between *index shift* and *cut-point shift*. *Index shift* occurs if the shape of the distribution of SAH remains the same, but there is a change in its location such that there is a parallel shift in all of the reporting thresholds for particular sub-groups of the population. This is illustrated in Figure 2:




Figure 2

It should be stressed that the term “index shift” may be misleading. The parallel shift in the distribution may be due to a shift in the cut-points (reporting bias) or due to a shift in the underlying measure of “true health”. In general, it is not possible to separately identify the two reasons for index shift.

If the reporting bias is due to *cut-point shift*, this implies that there is a change in the relative positions of the reporting thresholds for particular sub-groups of the population (see Figure 3 for an example). This will result in a change in the overall distribution of SAH.




Figure 3

A comparison of the two versions of the SAH question in the BHPS has the added complexity of the wave 9 SF-36 version of SAH having a different wording and a different set of response categories compared to the wave 1-8 and 10-11 question. The resulting differences in the distributions of SAH at wave 9 compared to other waves could be due to either an index shift in the mapping of “true” health to the SAH categories or a change in the relative location of the cut-points. Figure 4 illustrates a cut-point shift at wave 9.




Figure 4

Figure 1 shows evidence that, at least at an aggregate level, individuals appear to retain a reasonably consistent interpretation of the original SAH question. That is, the distribution of responses to the SAH question over waves 1-8 and 10-11 are very similar. However, there is a trend towards reporting worsening health over the eleven waves. This may be indicative of “true” health deteriorating as the cohort ages or it could be due to a shift in reporting thresholds used by respondents. The distribution of SAH at wave 9 clearly differs from the distributions observed at other waves. The joint effect of a change in the framing of the question together with the addition of the “very good” category and the removal of the “very poor” category is the likely source of the majority of the differences observed.³ It is of interest to ascertain whether these effects have resulted in either an index shift or a cut-point shift in the relationship mapping latent “true” health to the observed categories of SAH.

3.2 Pooled ordered probit and random effects ordered probit

One of the aims of this paper is to test whether the change in the measurement of SAH at wave 9 of the BHPS influences the estimates of panel data models for ordered responses and what this can tell us about measurement error. For that purpose, pooled ordered probit (POP) models have been applied. The literature suggests that ordered probit or ordered logit models can be used when considering an ordered categorical variable, such as SAH (McElvey and Zavoina, 1975), and if the scaling of the variable is available, interval regression is an alternative (van Doorslaer and Jones, 2003). The pooled model is estimated with robust standard errors, allowing for clustering of observations within individuals (see e.g., Wooldridge, 2002).

Ordered probit (and logit) models are typically motivated by a latent variable specification: this makes an explicit distinction between the process that determines an individual’s “true” health – the *latent variable model* – and the thresholds that determine which category of SAH they report,

³ It is possible that the distribution of “true” health at wave 9 differs from the distribution of “true” health at other waves and that this is reflected in self-assessed health status irrespective of the framing of the question. However in light of the evidence on the consistency in the reported distribution of SAH across the other waves this would appear highly unlikely.

given their true health – the *measurement model*. The *latent variable* specification for the models considered in this study is given by:

$$h_{it}^* = \beta'x_{it} + \varepsilon_{it} \quad i=1,\dots,N; t=1,\dots,T \quad (1)$$

where ε_{it} is an idiosyncratic error term uncorrelated with the set of regressors x_{it} . h_{it}^* represents the “true” level of health for individual i at time t . As “true” health is not observable we rely on the use of SAH as an appropriate dependent variable. It is useful to note that consistency of the POP estimator relies on the assumption that x_{it} is independent of the latent error, ε_{it} , in each time period. The POP gives *population-averaged* coefficients (Zeger et al., 1988). The POP assumes that the error term, ε_{it} , is distributed $N(0,1)$ for identification of β (see e.g., Wooldridge, 2002, p.486). The *measurement model* relates values of h_{it}^* to reported categories of SAH. If the re-wording of the SAH variable at wave 9 is taken into account, the latent variable specification would be the following⁴:

- For the original wording of the question, i.e., for waves 1-8 and 10-11,

$$h_{it} = j \text{ if } \mu_{j-1}^0 < h_{it}^* < \mu_j^0, \text{ with } j=1 \text{ (very poor), } 2 \text{ (poor), } 3 \text{ (fair), } 4 \text{ (good), } 5 \text{ (excellent)}$$

- In the case of the re-wording of the question (wave 9),

$$h_{it} = j \text{ if } \mu_{j-1}^1 < h_{it}^* < \mu_j^1, \text{ with } j = 1 \text{ (poor), } 2 \text{ (fair), } 3 \text{ (good), } 4 \text{ (very good), } 5 \text{ (excellent)}$$

Given this change in the phrasing of the question, the probability of an individual i choosing the j^{th} category will be given by the following:

$$P(h_{it} = j \mid x_{it}) = \begin{aligned} & \Phi(\mu_j^0 - \beta'x_{it}) - \Phi(\mu_{j-1}^0 - \beta'x_{it}) && t = 1-8, 10-11 \\ & \Phi(\mu_j^1 - \beta'x_{it}) - \Phi(\mu_{j-1}^1 - \beta'x_{it}) && t = 9 \end{aligned} \quad (2)$$

It can be seen that the β 's coefficients are common for all waves; given the latent variable specification, these parameters should be unaffected by changes in the measurement model. This

⁴ To identify the model, we follow STATA in using the conventional normalisation, that is: $\mu_0^0 = -\infty$, $\mu_5^0 = +\infty$ and there is no constant term in the linear index given by equation (1).

assumption provides the basis of our tests for reporting bias. Interest focuses on the estimates of the cut-points μ^0_j and μ^1_j .

3.3. Empirical Tests

The simplest form of index shift is where there is a parallel shift in the reporting thresholds that is common to all individuals. This implies that the original cut-points (μ^0_j) can be obtained by adding a constant, γ , to the cut-points corresponding to wave 9 (μ^1_j):

$$\mu^0_j = \mu^1_j + \gamma, \quad \text{for all } j \quad (3)$$

In this particular case, the probability that an individual i chooses category j to assess their health in time t , can be written as:

$$P(h_{it} = j \mid x_{it}) =$$

$$\begin{aligned} & \Phi(\mu^0_j - \beta'x_{it}) - \Phi(\mu^0_{j-1} - \beta'x_{it}) && t = 1-8, 10-11 \\ & \Phi(\mu^0_j - \gamma - \beta'x_{it}) - \Phi(\mu^0_{j-1} - \gamma - \beta'x_{it}) && t = 9 \end{aligned} \quad (4)$$

This shows that the parallel shift in wave 9 can be captured by adding a dummy variable for wave 9 to the latent variable equation – the parallel shift in the cut-points in (3) is equivalent to a shift in the intercept of the latent variable equation (1) at wave 9, hence the label “index shift” used by Lindeboom and van Doorslaer (2004). However, in the context of our application to the BHPS we are maintaining the assumption that the true health equation, (1), does not change at wave 9 but the measurement model, (2), does. So a more accurate label would be “parallel cut-point shift”. In practice, time dummies are included for each wave of the panel and the focus is on whether the variable for wave 9 stands out from the underlying trend.

This argument can be extended to the case where the index shift differs across different types of individual, defined by their characteristics x measured at wave 9:

$$\mu^0_j = \mu^1_j + \delta'x_{i9}, \quad \text{for all } j \quad (5)$$

Following the same argument as above, the model can be reformulated so that the latent variable is expressed as follows:

$$\begin{aligned}
h_{it}^* &= \beta'x_{it} + \varepsilon_{it}, & t = 1-8, 10-11 \\
h_{it}^* &= (\beta + \delta)'x_{it} + \varepsilon_{it}, & t = 9
\end{aligned} \tag{6}$$

In this case, differences in reporting at wave 9 are captured by the inclusion of interaction terms between the wave 9 dummy variable and the explanatory variables. Whilst the shift in reporting thresholds at wave 9 is a function of individual characteristics, the same shift applies to each of the thresholds.

Our specification relies on the distinction between the latent variable model (1) and the measurement model. In particular we have assumed that there is a single index that describes “true” health. This could be relaxed by allowing for a generalised ordered probit specification in which the β s are allowed to vary across the reported categories of SAH (β_j , $j=1\dots5$). This generalised specification makes intuitive sense if the ordered categories reflect choices available to the individual, but this is not the case with self-assessed health. The categories of SAH are an artefact of the design of the survey question and it is hard to construct an argument that the relationship between someone’s true health and their socio-economic characteristics will be contingent on the number and labelling of these arbitrary categories. Nevertheless, the existence of individual heterogeneity in the β s (i.e., allowing for β_j to differ across individuals) such that, for example, the association between income and health is lower for those in better health, could lead to heterogeneity across the categories of SAH (see van Doorslaer and Jones, 2003). To assess the *single index* assumption – sometimes referred to as the proportionality assumption (see Long, 1997, p.140), we estimate the POP specifications with different numbers of categories of SAH. As well as 5 categories, we also estimate models with SAH collapsed to 4 and 3 categories. Interest focuses on whether the β ’s remain stable with different categorisations of the dependent variable. If they do, then this provides prima facia evidence that our underlying latent variable model is consistent with the data.

This approach also sheds light on *cut-point shift*. This occurs when the shift in specific thresholds is allowed to differ. In general it is not possible to distinguish between heterogeneity and measurement error (see e.g., van Doorslaer and Jones, 2003). However, a comparison of specifications with 5, 4 or 3 categories will provide evidence whether the coefficients remain stable across categories. If they are stable then we can reasonably assume that any cut-point shift observed is due to measurement error rather than heterogeneity in latent health.

The previous analysis can be extended, in order to test for *cut-point shift*. For that purpose, we estimate a model allowing parameters to vary across categories of y for every wave, following a Generalised Ordered Probit model (GOP). The GOP allows the relaxation of one of the

restrictive characteristics of the ordered probit model: the constancy of threshold parameters (Terza, 1985). In order to allow cut-points to vary with the regressors, μ^{θ_j} needs to be reformulated as follows:

$$\mu^{\theta_{ij}} = \mu^{\theta_j} + \delta^{\theta_j} x_{it} \quad \text{for all } j \quad (7)$$

If the impact of socio-economic characteristics on the thresholds at wave 9 differs across thresholds, (5) becomes:

$$\mu^{\theta_{ij}} = \mu^{\theta_j} + \delta^{\theta_j} x_{it}, \quad \text{for } t = 9 \quad (8)$$

Then, δ^{θ_j} in equation (8) represents the additional effect of the regressors in the cut-points. Given (1), (7) and (8):

$$P(h_{it} = j \mid x_{it}) =$$

$$\frac{\Phi(\mu^{\theta_j} - (\beta + \delta^{\theta_j})' x_{it}) - \Phi(\mu^{\theta_{j-1}} - (\beta + \delta^{\theta_{j-1}})' x_{it})}{\Phi(\mu^{\theta_j} - (\beta + \delta^{\theta_j} - \delta^{\theta_{j-1}})' x_{it}) - \Phi(\mu^{\theta_{j-1}} - (\beta + \delta^{\theta_{j-1}})' x_{it})} \quad t = 1-8, 10-11 \quad (9)$$

$$\frac{\Phi(\mu^{\theta_j} - (\beta + \delta^{\theta_j} - \delta^{\theta_{j-1}})' x_{it}) - \Phi(\mu^{\theta_{j-1}} - (\beta + \delta^{\theta_{j-1}})' x_{it})}{\Phi(\mu^{\theta_j} - (\beta + \delta^{\theta_j} - \delta^{\theta_{j-1}})' x_{it}) - \Phi(\mu^{\theta_{j-1}} - (\beta + \delta^{\theta_{j-1}})' x_{it})} \quad t = 9$$

This gives a specification in which the coefficients vary across categories of SAH, as a result of cut-point shift. The pooled GOP is estimated directly by maximum likelihood, with robust standard errors corrected for the clustering within-individuals. For identification, the normalisation $\mu^{\theta_{i4}} = \mu^{\theta_4}$ is imposed, so the cut-point shift associated with the x 's is measured relative to μ^{θ_4} .

4. The BHPS data

The BHPS is a longitudinal survey of private households in Great Britain (England, Wales and Scotland south of the Caledonian Canal), and was designed as an annual survey of each adult (16+) member of a nationally representative sample of more than 5,000 households, with a total of approximately 10,000 individual interviews. The first wave of the survey was conducted between 1st September 1990 and 30th April 1991. The initial selection of households for inclusion in the survey was performed using a two-stage stratified systematic sampling procedure designed

to give each address an approximately equal probability of selection⁵. The same individuals are re-interviewed in successive waves and, if they split off from their original households are also re-interviewed along with all adult members of their new households.

A balanced panel is used, which means that only individuals from the first wave who were interviewed in each subsequent wave are included. Following Contoyannis et al. (2004), we include individuals who gave a full interview at each wave.⁶ Our working sample consists of 2,255 men and 2,841 women. All models are estimated separately for men and women.

Self-assessed health

The two versions of SAH available in the BHPS have been described in section 2. For both questions the responses are coded in increasing order of health. For example, for waves 1-8 and 10-11, “very poor” health is coded as 1, whilst “excellent” health is coded as 5. At wave 9, “poor” health is coded as 1 and “excellent” health as 5.

Insert Table 1 about here

To assess the single index assumption we collapse the raw categories of SAH so that there is common support within each collapsed category across the two versions of the question. For example, we first collapse SAH into four categories creating a new SAH variable with the following codes: 1: individuals who reported either “very poor” or “poor” health at waves 1-8 and 10-11, or individuals who reported “poor” health at wave 9; 2: individuals who reported “fair” health; 3: individuals who reported “good” health at waves 1-8 and 10-11, or “good” or “very good” health at wave 9; 4: individuals who reported “excellent” health. The original SAH categories were also further collapsed to a three-category variable. The collapsed categorisations are summarised in Table 1.

A further construction of SAH containing 4 categories is performed. This follows an approach suggested by Lindley and Lorgelly (2003). They estimate the relationship between income inequality and self-assessed health using the first nine waves of the BHPS. The interest of their paper for the purpose of this study is the way in which wave 9 data is included in the analysis. In particular, SAH has been constructed as a categorical variable with 4 categories, recoded as “excellent”, “good”, “fair” and “poor”. “Poor” and “very poor” were merged into the new category of “poor”. At wave 9 respondents reporting “very good” health were re-assigned to either “good” or “excellent” health on the following basis. A random sample consisting of 33%

⁵ For further details see Taylor (1998).

⁶ This excludes individuals for whom proxy interviews were provided together with those who responded to telephone interviews.

of individuals was recoded as “excellent”. The remaining individuals were recoded as “good”. A random sample of 20% of individuals reporting “good” health at wave 9 was recoded as “fair”. Further, a random sample of 20% of individuals reporting “fair” health at wave 9 was recoded as “poor”. It is argued by the authors that these re-codings and weights were chosen in order to maintain averages for each category of self-assessed health, for all waves (Lindley and Lorgelly, 2003). Table 2, presents the proportion of respondents in each category for the various constructions of the SAH variable. A comparison of the observed cell frequencies across the eleven waves for the four classifications of the SAH variable shows that the approach adopted by Lindley and Lorgelly (2003) provides greatest consistency. In particular, the frequencies for each of the categories observed at wave 9 are close to those obtained for the same categories at other waves.

Insert Table 2 about here

Socio-economic variables

Our empirical models are intended to capture the association between SAH and a range of socio-economic variables. There are three variables representing marital status (Widow, Single, Div/Sep), with Married as the reference category. An indicator of ethnicity is introduced (Non-White). Four dummy variables have been constructed to represent maximum level of education achieved, focusing on the human capital of the respondents. These are: Degree (higher degree/first degree), HND/A (Higher National Certificate, Higher National Diploma, A-level), O/CSE (O-level, Certificate of Secondary Education). No-Qual (no qualifications) is the base case for the education variables. The size of the household (HHSize) and the number of children in the household, by age, are also included in the analysis to capture the stage of life of the respondent. The income variable is the logarithm of equivalised real income, adjusted using the Retail Price Index and equivalised by the McClement’s scale to adjust for household size and composition. A quartic polynomial function of age is included (Age , $Age2=Age^2/100$, $Age3=Age^3/10,000$, $Age4 = Age^4/1,000,000$). A vector of time dummies is also included in the analysis. The variables are summarised in Table 3.

Insert Table 3 about here

A check for consistency

There is a useful additional question only available in wave 9 of the BHPS (included in the SF-36 questionnaire) regarding the individual’s health compared to one year ago. Individuals are asked: “Compared to one year ago, how would you rate your health in general now?”, with 5 possible answers: “much better”, “better”, “same”, “worse” and “much worse”. This variable links waves

8 and 9 and can be used to examine the consistency of the answers obtained in wave 9 and allows us to focus on the change in the labelling of the five categories (rather than the framing of the questions). The “same” health compared to one year ago is reported by 70.33% of the sample, with the next largest percentage reporting a higher level of health than the previous year (16.22% of individuals reported “much better” or “better” health) and 13.45% reporting “worse” or “much worse” health than a year ago. If we focus only on those individuals reporting the “same” health compared to one year ago, we can analyse the distribution of self-assessed health in wave 9, given each of the answers available for respondents in wave 8. Results are summarised in Tables 4 and 5.

Table 4 shows the percentage of individuals that reported either “excellent” or “good” health in wave 8, who assessed their health in wave 9 as “excellent” or “very good” (if they reported “excellent” in wave 8), and “very good” or “good” health (if they reported “good” health in wave 8). This provides evidence on the effect of introducing the category “very good” between “excellent” and “good”, in wave 9. It is possible to check for patterns in the mapping from the wave 8 responses to the new “very good” category with respect to some of the socioeconomic variables, such as gender, highest level of education achieved, income quantiles and age. 42.04% of individuals who reported “excellent” health in wave 8, assessed their health as “very good” when this option became available in wave 9. This percentage is similar for men and women (42.02% and 42.05%, respectively). By socioeconomic groups, it can be seen that the percentage of individuals reporting “very good” health in wave 9 follows an increasing trend with higher levels of education, i.e. those who are more educated are more likely to change their response from wave 8 and assess their health as “very good” in wave 9. For income quantiles, those individuals in the fourth quantile represent the smallest percentage of individuals who report “very good” health in wave 9. A positive gradient can be seen for the age groups considered. In fact, the percentage of individuals reporting “very good” health in wave 9 increases as individuals get older. The same analysis could be applied for those individuals who reported “good” health in wave 8. In this case, 42.13% of individuals report “very good” health in wave 9. By gender, it can be seen that more women (44.13%) than men (39.73%) report “very good” health in wave 9. The percentage of those individuals reporting “very good” health increases with levels of education. If the richest group of the sample (Q5) is compared with the poorest one (Q1), it can be seen that the percentage of those reporting “very good” health is greater for those in the highest income group. Older individuals are less likely to report “very good” health in wave 9, once they have reported “good” health in wave 8.

Insert Table 4 about here

The changes in SAH responses from wave 8 to wave 9 reported in Table 4 do not necessarily imply inconsistencies in reporting, for example, someone who reports “excellent” health in wave 8 and says that their health is still the same may legitimately report “very good” when this option becomes available at wave 9. In contrast, Table 5 provides evidence of inconsistencies of the self-assessments for those individuals who claimed to have the “same” level of health in wave 9, compared to last year. From these results it is possible to check how this inconsistency varies across socioeconomic groups. For that purpose, we distinguish between those individuals whose assessments reflect either a *decline* or an *improvement* in health, comparing their responses for wave 8 with the answers provided in wave 9. The label *decline* in Table 5 represents those individuals that make an inconsistent movement from one category in wave 8 to a lower category in wave 9, being inconsistent because these individuals reported having the “same” health in wave 9 compared to last year. We assume that *decline* includes the following possibilities: moving from “excellent” in wave 8 to either “good”, “fair” or “poor” in wave 9; moving from “good” in wave 8 to “fair” or “poor” in wave 9; moving from “fair” in wave 8 to “poor” in wave 9. On the other hand, the label *improvement* represents those individuals that report higher categories in wave 9, having reported that their health is the “same” compared to last year. These are the choices considered: moving from “good” in wave 8 to “excellent” in wave 9; moving from “fair” in wave 8 and to either “excellent”, “very good” or “good” in wave 9; moving from “poor” in wave 8 to either “excellent”, “very good”, “good” or “fair” in wave 9; moving from “very poor” in wave 8 and to either “poor”, “fair”, “good”, “very good” or “excellent” in wave 9.

From these results, it can be seen that a higher percentage of individuals who reported having the *same* health in both waves, improve their assessment of health (25.98%) in wave 9, while 9.88% report a decline in their own health. The percentage of women reporting an improvement is slightly higher (23.87%) than for men (22.46%), with men being more likely to be inconsistent in their answers when reporting a decline in health (10.33%). The percentage of individuals reporting either an improvement or a decline in their health compared to wave 8 decreases with level of education. Hence, individuals with no qualifications tend to be less consistent. There is a negative gradient with income, as individuals in the lowest income group tend to be less consistent and report an improvement in their health, compared to individuals in the highest quantile. It is not possible to distinguish a clear pattern of those reporting a decline in health by income groups. However, if the lowest income group (Q1) is compared to the highest one (Q5), it can be seen that richer individuals are less likely to report a decline in their health. There is not a clear pattern by age groups either. However, those older than 70 are more likely to report either an improvement in health (30.28%) or a decline in health (10.24%), while those aged 50 to 70, are less likely to be inconsistent in their answers, as 7.77% report a decline in their health and 19.55% report an improvement in their assessment of health.

Insert Table 5 about here

The analysis of the introduction of very good health as a response category at wave 9, as shown in Table 4, indicates, on the whole, that respondents are reasonably consistent in their reporting of SAH, given that they stated their health had remained the same across waves 8 and 9 and that they had reported excellent or good health in wave 8. However, when considering all possible categories of SAH at wave 8, there is a larger degree of inconsistency in the reporting behaviour between waves 8 and 9 with up to 30% of individuals providing a response at wave 9 that is inconsistent with reporting that their health had remained the same. If we are prepared to assume that individuals can better recall their underlying health status a year ago with greater accuracy than they recall the actual response category they used when responding to the BHPS survey, then these results provide prima facie evidence of measurement error at wave 9 that appears to have been induced by the change in the labelling of SAH response categories when moving to the SF 36 questionnaire.

5. Results

The latent variable model

Results for POP models for SAH with 5 categories, estimated for men and women separately on all data excluding wave 9, are presented in Table 6. These results represent a “benchmark” against which the results of incorporating the wave 9 SF-36 version of SAH can be compared. A number of features are worth noting. There is a clear gradient in the relationship between educational attainment and self-assessed health with individuals with greater qualifications reporting better health. There is also the expected positive relationship between income and health, although for both men and women.

Insert Table 6 about here

We observed a general decline in the reporting of better health states when we compared the raw distributions of SAH across the eleven waves as shown in Figure 1. When we condition on relevant regressors such as age, income and educational attainment we still observe a clear gradient across waves, this is evident from the coefficients on the year dummy variables. These become more negative as the cohort ages indicating the reporting of worsening health. This may be due to an age effect not fully removed by the wording of the SAH question or not fully controlled for by conditioning on a polynomial in age. However, it may also be due to changes in

reporting as individuals age, reflected in an index shift over time. This would imply that the context within which individuals rate their own health changes with increasing age. However this second hypothesis seems unlikely as the direction of the trend contradicts the notion of *adaptation*: that individuals adapt to worsening health, becoming more optimistic in their assessment of self-assessed health (Groot, 2000).

The measurement model

Table 7 presents the results of POP specifications for men and women when incorporating wave 9 data on SAH. The columns of the tables represent the alternative classifications of SAH as outlined in Tables 1 and 2 above. A comparison of the SAH 5 coefficients to their respective coefficients reported in Table 6 reveals little change in the estimated magnitude of the effects of socio-economic characteristics on health, although, in general these effects increase slightly in their absolute magnitudes from zero. This is evident for both men and women. A further comparison of the coefficient estimates across the four different categorisations of SAH reveals a large degree of consistency with perhaps SAH 3 being an exception: collapsing the original five-category SAH does not appear to affect substantially the estimated relationship between socio-economic characteristics and health. This finding supports the notion that the underlying latent variable model specification is appropriate. Larger differences in coefficients across the four models are observed for the estimated cut-points indicating, as one would expect, that collapsing the original five-point categorical scale into four or three categories has most effect on the measurement model that determines which category of SAH individuals report given their true latent health.

Insert Table 7 about here

A Common index shift

Including wave 9 data has a marked effect on the estimated year dummy coefficients. The results presented in Table 6 that excluded wave 9 information showed a gradient decreasing in reported health as the cohort aged. This gradient remains after the inclusion of wave 9 data. However, the estimated coefficient for the dummy variable for wave 9 (yr9900) in general does not conform to the trend observed across the other years (see Table 7). This is particularly true for the SAH 5 model where the estimated coefficient for yr9900 is substantially larger than the coefficients for other years. This is not surprising given that the construction and coding of the SAH variable for wave 9 (based on the SF-36 version of SAH) is different to the construction and coding of the SAH variable at other waves. The large negative estimated coefficient on the

wave 9 dummy variable appears as a direct artefact of the way in which the variable has been constructed.

If we consider the coefficients at wave 9 for the models adopting the SAH 4 categorisation and the approach based on the methodology of Lindley and Lorgelly (L & L) then we can see that for men the estimates are close to being consistent with the trends in the gradients observed across the other years. However, the estimates are closer to zero than true consistency with the time trends would predict. For women, the estimates do not conform to the trends observed across other years and are much closer to zero than would have been predicted on the basis of the trends. The smaller absolute value of the wave 9 coefficient obtained for both men and women would appear to reflect two things: firstly, a floor effect at wave 9 for the lowest health category of “poor” health and secondly, the larger proportion of respondents rating their health in the combined category of “good, very good” health compared to the proportion in the equivalent category of “good” health at other waves. The differences in proportions can be seen in Table 2. The combined effect of this is to induce an index shift towards reporting better health status at wave 9.

The coefficient estimate for wave 9 when considering the SAH 3 categorisation is positive for men and large and positive for women. These do not conform to the trends observed in the estimated coefficients across the other years of data. A cursory glance at Table 2 shows that the distribution of responses at wave 9 when using the SAH 3 categorisation is skewed towards reporting better health status. Again, this is consistent with an index shift towards reporting better health at wave 9.

Reporting bias

The results reported above indicate that the data are consistent with an overall index shift at wave 9. To assess whether the observed shift in cut-points is systematically related to population subgroup characteristics we interact the wave 9 dummy variable with the socio-economic variables. Table 8 presents the results for interactions with educational attainment, income and age. Again, results are presented for POP specifications only. Models were estimated for the original five-category SAH variable together with the SAH 4 categorisation and the approach suggested by Lindley and Lorgelly. Log-likelihood ratio tests of restricted (without interactions terms) versus unrestricted (with interaction terms) models reject the restricted model for the five-category SAH variable and the Lindley and Lorgelly categorisation of SAH for women only. However, a comparison of the individual coefficients to their standard errors reveals no significant interaction terms between the socio-economic variables and the wave 9 dummy variable for any

of the model specifications tested⁷. This latter result suggests that the cut-points are not systematically related to identifiable population sub-groups.

Insert Table 8 about here

Cut-point shift

In general it is not possible to distinguish between cut-point shift and heterogeneity in the latent variable equation. However if it is assumed that equations (1) and (9) are correctly specified – that is, there is no heterogeneity in the β s and reporting bias leads to cut-point shift – then the results of the GOP can be interpreted as separating the effects of the regressors on true health from reporting bias. The former are captured by the β s and the latter by the δ s. In order to interpret the results from the GOP model estimates, it is necessary to take into account that the top cut-point (μ_4) is fixed. We estimate two versions of the GOP model. In the first version the cut-points are modelled as a function of the explanatory variables corresponding to equation (7). We term this the restricted GOP model. The second version of the GOP model extends the set of explanatory variables used to test the constancy of the threshold parameters to the interaction terms between the initial set of explanatory variables used in the restricted model and the wave 9 dummy variable. We term this model the unrestricted GOP model. The unrestricted model corresponds to equations (7) and (8).

Table 9 shows GOP results for the restricted model for both men and women respondents, in particular, the coefficients obtained for the education variables, the income variable and the dummy variable corresponding to wave 9. It can be seen that for education and income variables the corresponding coefficients in the linear index model present the expected positive sign, for both men and women. The magnitudes of the β s are similar but they tend to be smaller in the GOP models compared to the POP. When the thresholds are allowed to vary with the explanatory variables, the δ_j coefficients are negative. This result implies that, as respondents increase either their level of education or their income, their assessment of health improves. Hence, individuals with higher income or level of education tend to be more optimistic about their health.

Comparing the coefficients corresponding to level of education for men, it can be seen that the absolute value of the δ_j coefficients decrease with level of education, implying that individuals with higher level of education are more optimistic about their health than those with O/CSE or HND/A level. The same finding applies to the coefficients for the income variable. Similar results are found for women. The coefficients for level of education and income variables have

⁷ This is also true for the interaction terms not presented in Table 8.

the expected sign. However, in absolute terms, it can be seen that the coefficients are lower for women than for men. Hence, the shift in the cut-points is smaller for women.

Insert Table 9 about here

Table 10 presents the GOP results for the unrestricted model. Here all interaction terms between the set of explanatory variables and the wave 9 dummy variable are used to model both the latent health index and the cut-point thresholds. It can be seen that the coefficient estimates for the main effects of educational attainment and income are very similar to the results obtained in the restricted model. This holds for both men and women. However, the estimated coefficient on the wave 9 dummy variable differs dramatically and for men reflects closely that obtained for the POP model including interaction terms, reported in Table 8.

Few of the additional parameter estimates in the unrestricted GOP model achieve statistical significance at the 5% level. Of the additional (over the restricted model) 48 parameters estimated 3 achieved statistical significance for women and only 1 for men.⁸ A log-likelihood ratio test of unrestricted versus restricted model fails to reject the null hypothesis that the restrictions are valid suggesting that the restricted GOP model is an adequate specification for these data.⁹ Again this suggests that the change in wording of the SAH variable at wave 9 is not associated with any cut-point shift in the measurement model.

Insert Table 10 about here

Concentration and mobility indices

To complete our analysis we present estimates of income-related health inequality and mobility (van Doorslaer and Koolman, 2004, Jones and López, 2004) based on the different measures of SAH. In Table 11, short-term and long-term concentration indices are presented, together with mobility indices for self-assessed health (SAH) for all individuals. Given the change of wording in the SAH variable at wave 9, we have used SAH measures with 5, 4 and 3 categories for all waves as outlined in Table 1. Lindley and Lorgelly's approach has also been implemented to obtain the SAH variable with 4 categories for all waves. The income measure used to construct the indices is equivalised real income that has been adjusted using the Retail Price Index and equivalised by the McClement's scale to adjust for household size and composition. A balanced panel is used, which means that only individuals from the first wave who were interviewed in each subsequent wave

⁸ For women these are: yr9900 x HND/A, yr9900 x O/CSE and yr9900 x Div/Sep as parameters for cut-point 3 (Cut 3). For men yr9900 x HND/A for cut-point 3 is significant.

⁹ For men: $\chi^2_{47} = 42.2$; $p = 0.67$, for women: $\chi^2_{48} = 57.8$; $p = 0.16$.

are included. Following Contoyannis et al. (2003), we include individuals who gave a full interview at each wave.

In Table 11, it can be seen that the short-term concentration indices for all waves are positive, indicating that there is “pro-rich” inequality in health in all periods. It is also possible to see an increase in the size of the short-term concentration indices with time. However, this trend breaks at wave 9 for all constructions of SAH considered, excepting for the SAH variable constructed by following Lindley & Lorgelly’s approach (column headed L&L). In particular, for SAH with 5 categories there is an upward shift in the short-term concentration index for wave 9. For SAH with 3 and 4 categories, the estimate for wave 9 shifts downwards. Long-run concentration indices are also shown, being positive for all the variables considered. The negative value of the mobility index implies that the level of long-run income-related health inequality is greater than we could infer from cross-sectional information. If we compare the estimates across constructions of SAH, it is possible to see that the short-term and long-term concentration indices together with the mobility indices, are very similar between SAH with 4 categories and Lindley and Lorgelly’s approach. However, estimates for SAH with 5 or 3 categories are smaller than the other constructions.

Tables 12 and 13 show the results of the concentration indices and mobility index, differentiating between male and female respondents. The short-term concentration index for both men and women are positive, indicating that there is “pro-rich” inequality in health in all periods. There is also an increase in the size of the estimates with time. Although men start with a lower concentration index than women for wave 1, their estimates reach higher values in the last waves of the panel, compared to women. This result applies to all the constructions of SAH that have been applied. Considering results by gender, it can be seen that the increasing trend of the short-term concentration index with time also breaks in wave 9; the shift being evident for all constructions of SAH, with the exception of Lindley & Lorgelly’s approach. Again we can see an upward shift at wave 9 for the 5-category SAH variable and a downward shift for SAH 4 and SAH 3. This holds for both men and women. The long-term concentration index is also positive for both men and women, being higher for women than for men. These estimates are greater for SAH with 4 categories and following Lindley & Lorgelly’s approach than for the other constructions.

The mobility indices are negative for both sub-samples, implying that the level of long-run income related health inequality is greater than we could infer from cross-sectional information. The size of the mobility indices is slightly greater for women than for men, in absolute terms.

6. Conclusion

Several indices and instruments are available to obtain valid information on an individual's general level of health. The reliability of subjective information such as self-reported general health status raises concern due to it being vulnerable to several biasing factors (Knäuper and Turner, 2000). However, objective measures of health status are rare in survey data and where they do exist they are often too specific to particular health conditions. Accordingly their applicability as an overall measure of an individual's health status is often limited. An appeal of measures of general health status is their ability to encapsulate and summarise a multitude of health conditions, but this is at the cost of them largely being based on subjective assessment and self-report. This is the case in the BHPS that offers a rich and diverse source of information on individual behaviour and socio-economic situation. The issue surrounding the use of self-report general health status is confounded further in the BHPS due to a change in the SAH question at wave 9. The inclusion of the SF-36 questionnaire in the ninth wave and the consequent change in the phrasing of the SAH question together with a change to the response categories made available raises challenges for the analyst wishing to conduct longitudinal work where the focus of interest is health status. However, and importantly, the change in the SAH question provides an opportunity to investigate the extent and nature of reporting bias and more specifically whether groups of individuals, identified through socio-economic characteristics, respond in different ways to the change in the measurement instrument. Should this be the case, then the results have implications for the modelling of SAH in the BHPS and more generally. Evidence of reporting bias in the form of cut-point shift would suggest that standard ordered probit models of SAH are overly restrictive and preclude flexible estimation of the effect of regressors on the outcome probabilities.

Considerable attention has been devoted to the reliability of SAH and the scope for contamination by measurement error; the change in wording at wave 9 provides a form of natural experiment that allows us to assess the sensitivity of panel data analyses to a change in the measurement instrument. First, we show the effect that the change of wording of the SAH question has on longitudinal models of SAH and secondly, how this information can be exploited to inform about the nature of reporting bias. Progressively more general specifications of reporting bias are implemented using panel data ordered probit and generalised ordered probit models. These are designed specifically to test hypotheses about the role of heterogeneity and reporting bias in individuals' responses to the SAH question. Further, given the increased reliance on SAH in studies on inequality in health, we explore the sensitivity of the concentration index of health inequality and of mobility in health to changes in the measurement of SAH.

Our findings suggest that collapsing the original SAH variable to create a categorisation that has common support over the two versions of the question does not affect substantially the estimated relationship between socio-economic characteristics and self-assessed health. Collapsing the categories of SAH is best achieved by collapsing to a four-category version of SAH, either by combining “very poor” and “poor” health in waves 1-8 and 10-11 and “good” and “very good” in wave 9 or by adopting the methodology proposed by Lindley and Lorgelly (2003). On a practical level, this is useful for longitudinal studies of the BHPS, or other panel datasets that incorporate a change in the wording of SAH, that wish to make full use of the available data, but is also relevant to comparative studies that draw on sources of data that contain the different versions of the self-assessed health question

The assumption that our underlying latent variable model of SAH is an appropriate specification is supported by the finding that collapsing the categories of SAH does not substantively alter the relationship between socioeconomic characteristics and health. Building on this result, our investigations of reporting bias suggest that while there is evidence of an index (parallel) shift in respondents’ use of thresholds to map underlying health status to the categories of SAH, the magnitude of this parallel shift does not vary by socio-economic characteristics. Further, we find little evidence to suggest that reporting bias induced by the change in wording is characterised by cut-point shift; our results do not indicate that the relative positions of the cut-point thresholds vary in a way determined by individual socio-economic characteristics.

References

- Adams, P., Hurd, M.D, McFadden, D., Merrill, A., and Ribeiro, T. (2003). Healthy, wealthy and wise? Tests for direct causal paths between health and socioeconomic status. *Journal of Econometrics* 112: 3-56.
- Benzeval, M., Taylor, J., and Judge, K. (2000). Evidence on the relationship between low income and poor health: Is the Government doing enough?. *Fiscal Studies*, 21(3): 375-399.
- Bound, J (1991). Self-Reported versus objective measures of health in retirement models. *The Journal of Human Resources*, 26, 106-137.
- Buck N., Gershuny J., Rose D., Scott J. (1994). *Changing Households. The BHPS 1990-1992*. ESRC Research Centre on Micro-Social Change.
- Burström, B. & Fredlund, P. (2001). Self rated health: is it as good a predictor of subsequent mortality among adults in lower as well as in higher social classes?. *Journal of Epidemiology and Community Health*, 55: 836-840.
- Butler J.S., Richard V. Burkhauser, Jena M Mitchel, Theodore P. Pincus (1987). Measurement error in self-reported health variables. *The Review of Economics and Statistics*, 69 (4): 644-650.
- Contoyannis, P. and Jones, A. (2004). Socio-economic status, health and lifestyle. *Journal of Health Economics*, 23: 965-995.
- Contoyannis P., Jones A. M., and Rice N. (2004). The dynamics of health in the British Household Panel Survey. *Journal of Applied Econometrics*., 19: 473-503.
- Crossley T.F. and Kennedy S. (2002). The reliability of self-assessed health status. *Journal of Health Economics*, 21: 643-658.
- Currie, J. and Madrian, B.C. (1999). Health, health insurance and the labor market. In Ashenfelter, O. and card, D. (eds.) *Handbook of Labour Economics, Vol.3*, Elsevier, Amsterdam.
- Deaton, A..S. and Paxson, C.H. (1998). Ageing and inequality in income and health. *American Economic Review, Papers and Proceedings* 88: 248-253.
- Ettner, S. (1996). New evidence on the relationship between income and health. *Journal of Health Economics*, 15: 67-85.
- Frijters, P., Haisken-DeNew, J.P. and Shields, M.A. (2003). Estimating the causal effect of income on health: evidence from post reunification East Germany, *Centre for Economic Policy Discussion Paper No. 465*, Australian National University.
- Groot, W. (2000). Adaptation and scale of reference bias in self-assessments of quality of life. *Journal of Health Economics* 19(3), 403-420.
- Grootendorst, P., Feeny, D., and Furlong, W. (1997). Does it matter whom and how you ask? Inter and intra-rater agreement in the Ontario health survey. *Journal of Clinical Epidemiology* 50: 127-136.
- Idler, E. L. and Y. Benyamini (1997). Self-rated health and mortality: a review of twenty-seven community studies. *Journal of Health and Social Behavior* 38(1): 21-37.
- Idler, E.L. and Kasl, S.V. (1995). Self-ratings of health: do they also predict change in functional ability?. *Journal of Gerontology* 50B: S344-S353.

- Jenkinson, C. J. et al. (1996). *The U.K. SF-36: an analysis and interpretation manual. A guide to health status measurement with particular reference to the Short Form 36 Health Survey*. Oxford University, Department of Public Health and Primary Care.
- Kaptein, A, JP Smith, A van Soest (2004). Self-reported work disability in the US and The Netherlands. *RAND Labor and Population Working Paper*.
- Kenkel, D. (1995). Should you eat breakfast? Estimates form health production functions. *Health Economics*, 4: 15-29.
- Kerkhofs, M., Lindeboom, M. (1995). Subjective health measures and state dependent reporting errors. *Health Economics* 4: 221-235.
- Knäuper Barbel and Patricia A Turner (2000). Measuring health: improving the validity of health assessments. In J. Dunbar-Jacob (Ed.), *Quality of Life: In the Eye of the Beholder*. University of Pittsburgh: CRC.
- Lindeboom M (2006). Health and work of older workers. In A.M.Jones (Ed.) *Elgar Companion to Health Economics*: Edward Elgar, Cheltenham.
- Lindeboom M. and van Doorslaer E. (2004). Cut-point shift and index shift in self-reported health. *Journal of Health Economics*, 23: 1083-1099.
- Lindley, J., Lorgelly, P.K. (2003). *The relative income hypothesis: does it exist over time? Evidence from the BHPS*. Paper prepared for the Health Economists' Study Group. Winter 2003, Leeds.
- Long, J. Scott (1997). *Regression models for categorical and limited dependent variables*. Sage Publications.
- McKelvey R.D. and Zavoina W. (1975). A statistical model for the analysis of ordinal level dependent variables. *Journal of Mathematical Sociology*, 4: 103-120.
- Murray, CJL, A Tandon, J Salomon and CD Mathers (2001). *Enhancing cross-population comparability of survey results*, GPE Discussion Paper Nr 35, WHO/EIP, Geneva.
- Sadana, R, CD Mathers, AD Lopez, CJL Murray and K Iburg (2000), *Comparative analysis of more than 50 household surveys on health status*, GPE Discussion Paper No 15, EIP/GPE/EBD, World Health organisation, Geneva.
- Salas, C. (2002). On the empirical association between poor health and low socioeconomic status at old age. *Health Economics*, 11: 207-220.
- Smith, J.P. (1999). Healthy bodies and thick wallets: the dual relationship between health and economic status. *Journal of Economic Perspectives* 13: 145-166.
- Terza, J.V. (1985). Ordinal Probit: A Generalization. *Communications on Statistics – Theory and Methods*, 14 (1): 1-11.
- Taylor, M., Brice, J., Buck, N., and Prentice, E. (1998). *British Household Panel Survey User Manual Volume A: Introduction, Technical Report and Appendices*. University of Essex, Colchester.
- van Doorslaer, E. (1987). *Health, knowledge and the demand for medical care*. Assen, Maastricht.
- van Doorslaer, E. and Gerdtham, U-G.(2003). Does inequality in self-assessed health predict inequality in survival by income? Evidence from Swedish data, *Social Science and Medicine*, 57: 1621-1629.

- van Doorslaer, E. and Koolman, X. (2004). Explaining the differences in income-related health inequalities across European countries. *Health Economics*, 13: 609-628.
- van Doorslaer, E. and Jones, A. M. (2003). Inequalities in self-reported health: validation of a new approach to measurement. *Journal of Health Economics*, 22(1): 61-87.
- van Doorslaer, E., Jones, A.M. and Koolman, X. (2004). Explaining income-related inequalities in doctor utilisation in Europe. *Health Economics*, 13: 629-647.
- van Doorslaer, E. et al. (1997). Income-related inequalities in health: some international comparisons. *Journal of Health Economics*, 16 (1): 93-112.
- van Doorslaer, E., et al. (2000). Equity in the delivery of health care in Europe and the US. *Journal of Health Economics*, 19: 553-583.
- van Ourti, T. (2003). Socio-economic inequality in ill-health amongst the elderly. Should one use current income or permanent income? *Journal of Health Economics*, 22: 187-217.
- Verbeek M. (2000). *Modern econometrics*. Wiley.
- Wooldridge, J. (2002). *Econometric Analysis of cross section and panel data*. MIT Press.
- Zeger, S. L., K.-Y. Liang and P.A. Albert (1988). Models for longitudinal data: a generalized estimating equation approach. *Biometrics*, 44: 1049-60.

Figure 1. Distribution of SAH for each wave

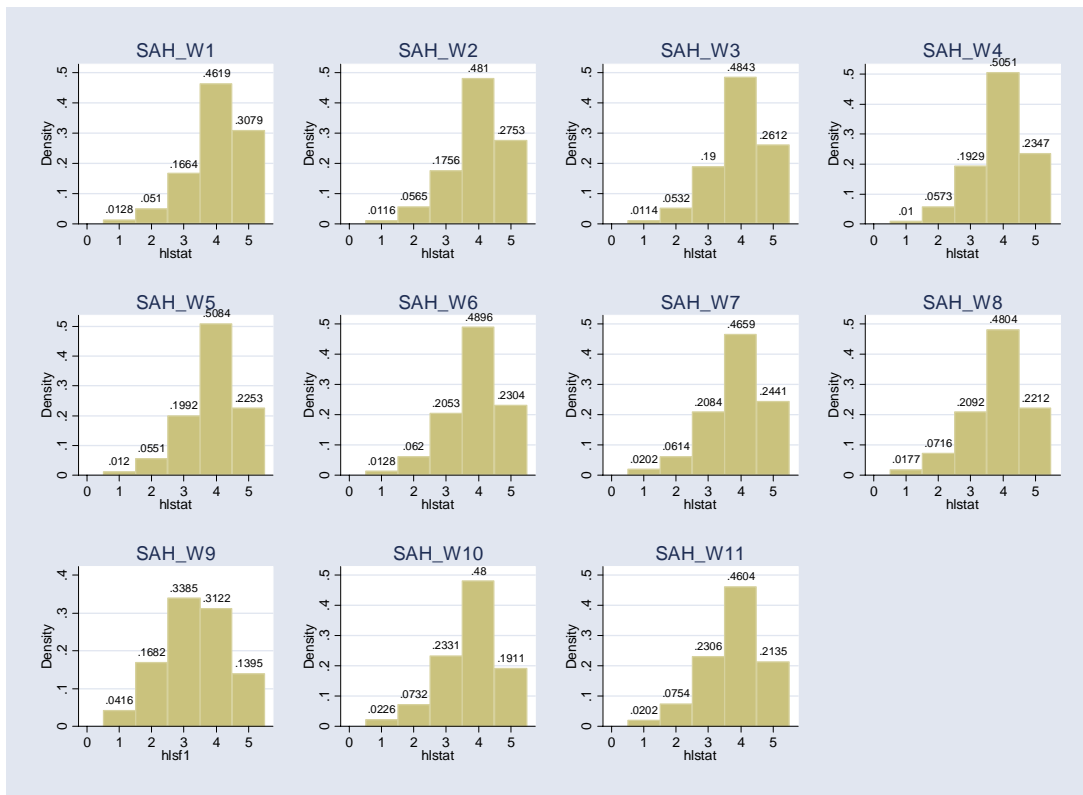


Table 1. Construction of SAH variable

Wave	Category	SAH 5 categories	SAH 4 categories	SAH 3 categories
1-8, 10-11	1	very poor	very poor, poor	very poor, poor
	2	poor	fair	fair
	3	fair	good	good, excellent
	4	good	excellent	
	5	excellent		
9	1	poor	poor	poor
	2	fair	fair	fair
	3	good	good, very good	good, very good, excellent
	4	very good	excellent	
	5	excellent		

Table 2. Frequencies for the classifications of SAH

	Wave										
	1	2	3	4	5	6	7	8	9	10	11
SAH 5											
1	1.28	1.16	1.14	1.00	1.20	1.28	2.02	1.77	4.16	2.26	2.02
2	5.10	5.65	5.32	5.73	5.51	6.20	6.14	7.16	16.8	7.32	7.54
3	16.6	17.6	19.0	19.3	19.9	20.5	20.8	20.9	33.9	23.3	23.1
4	46.2	48.1	48.4	50.5	50.8	49.0	46.6	48.0	31.2	48.0	46.0
5	30.8	27.5	26.1	23.4	22.5	23.0	24.4	22.1	14.0	19.1	21.4
SAH 4											
1	6.38	6.81	6.46	6.73	6.71	7.48	8.16	8.93	4.16	9.58	9.56
2	16.6	17.6	19.0	19.3	19.9	20.5	20.8	20.9	16.8	23.3	23.1
3	46.2	48.1	48.4	50.5	50.8	49.0	46.6	48.0	65.1	48.0	46.0
4	30.8	27.5	26.1	23.5	22.5	23.0	24.4	22.1	14.0	19.1	21.4
SAH 3											
1	6.38	6.81	6.46	6.73	6.71	7.48	8.16	8.93	4.16	9.58	9.56
2	16.6	17.6	19.0	19.3	19.9	20.5	20.8	20.9	16.8	23.3	23.1
3	77.0	75.6	74.6	74.0	73.4	72.0	71.0	70.2	79.0	67.1	67.4
L & L*											
1	6.38	6.81	6.46	6.73	6.71	7.48	8.16	8.93	7.69	9.58	9.56
2	16.6	17.6	19.0	19.3	19.9	20.5	20.8	20.9	19.8	23.3	23.1
3	46.2	48.1	48.4	50.5	50.8	49.0	46.6	48.0	48.2	48.0	46.0
4	30.8	27.5	26.1	23.5	22.5	23.0	24.4	22.1	24.3	19.1	21.4

*Following approach of Lindley and Lorgelly (2003)

Table 3. Definition of the regressors

Widow	1 if widowed, 0 otherwise
Single	1 if never married, 0 otherwise
Div/Sep	1 if divorced or separated, 0 otherwise
Non-White	1 if a member of ethnic group other than white, 0 otherwise
Degree	1 if highest academic qualification is a degree or higher degree, 0 otherwise
HND/A	1 if highest academic qualification is HND or A Level, 0 otherwise
O/CSE	1 if highest academic qualification is O level or CSE, 0 otherwise
HHSIZE	Number of people in household including respondent
NCH04	Number of children in household aged 0-4
NCH511	Number of children in household aged 5-11
NCH1218	Number of children in household aged 12-18
Income	Equivalised annual real household income in pounds
Age	Age in years at 1st December of current wave

Table 4. Introduction of “very good” at wave 9

	Ex (w8)		G (w8)	
	Ex (w9)	VG (w9)	VG (w9)	G (w9)
All	42.12	42.04	42.13	42.48
Men	40.82	42.02	39.73	43.88
Women	43.51	42.05	44.13	41.32
Degree	51.37	37.7	48.85	37.38
HND	40.39	44.95	43.01	42.65
OCSE	43.51	41.08	43.67	41.39
None	35.69	43.12	37.22	45.65
Income quantiles:				
Q1	40.68	44.07	36.96	45.54
Q2	38.89	40.74	40.87	42.61
Q3	39.78	45.3	43.81	41.9
Q4	42.86	35.29	40.69	45.41
Q5	44.44	44.05	47.33	37.62
Age categories:				
<30	41.82	38.18	43.3	37.63
30-50	46.84	36.92	45.75	39.2
50-70	36.75	46.64	39.29	48.05
70+	29.76	54.76	37.35	42.97

Table 5. Evidence of inconsistent responses

	Decline	Improvement
All	354 (9.88%)	931 (25.98%)
Men	172 (10.33%)	374 (22.46%)
Women	169 (8.81%)	458 (23.87%)
Degree	32 (6.93%)	88 (19.04%)
HND	71 (8.33%)	181 (21.24%)
OCSE	98 (8.85%)	246 (22.22%)
No qual	139 (12.01%)	312 (26.97%)
Income quantiles:		
Q1	51 (8.5%)	173 (28.83%)
Q2	66 (9.78%)	158 (23.41%)
Q3	57 (7.45%)	176 (23.01%)
Q4	67 (8.85%)	158 (20.87%)
Q5	46 (5.84%)	167 (21.22%)
Age categories:		
<30	30 (8.09%)	91 (24.53%)
30-50	121 (7.52%)	340 (21.14%)
50-70	89 (7.77%)	224 (19.55%)
70+	47 (10.24%)	139 (30.28%)

Table 6. POP models excluding wave 9

	MEN NT = 2,2539	WOMEN NT = 2,8386
Degree	.377 (.058)	.381 (.052)
HND/A	.260 (.048)	.279 (.045)
O/CSE	.218 (.049)	.289 (.041)
Age	.071 (.060)	.031 (.059)
Age2	-.236 (.193)	-.123 (.187)
Age3	.300 (.262)	.209 (.249)
Age4	-.129 (.126)	-.126 (.118)
Ln(income)	.244 (.025)	.206 (.019)
yr9293	-.087 (.023)	-.067 (.020)
yr9394	-.133 (.024)	-.095 (.021)
yr9495	-.170 (.025)	-.147 (.022)
yr9596	-.215 (.025)	-.160 (.023)
yr9697	-.212 (.027)	-.202 (.024)
yr9798	-.237 (.028)	-.197 (.025)
yr9899	-.285 (.028)	-.238 (.025)
yr9900		
yr0001	-.368 (.029)	-.336 (.026)
yr0102	-.336 (.030)	-.276 (.027)
Cut1	.747	.117
Cut2	1.514	.889
Cut3	2.396	1.764
Cut4	3.736	3.119
Intra-class correlation		
Log Likelihood	-26874.1	-35101.4
RESET (p value)	.04 (.85)	1.49 (.22)

Notes:

1. SAH is based on the categorisation: 1: “very poor”, 2: “poor”; 3 “fair”; 4 “good”; 5: “excellent” health.
2. Estimated coefficients for the variables Widow, Single, Div/Sep, Non-White, HHSIZE, NCH04, NCH511, NCH1218 have been suppressed to conserve space. Full results are available from the authors upon request.
3. Standard errors in parentheses.

Table 7. POP models including wave 9 data

	Men NT = 24791				Women NT = 31225			
	SAH 5	SAH 4	SAH 3	L & L	SAH 5	SAH 4	SAH 3	L & L
Degree	.385 (.058)	.381 (.058)	.454 (.068)	.377 (.057)	.388 (.053)	.382 (.052)	.477 (.060)	.383 (.052)
HND/A	.266 (.047)	.262 (.048)	.324 (.054)	.263 (.047)	.287 (.045)	.280 (.045)	.338 (.051)	.286 (.045)
O/CSE	.222 (.048)	.219 (.049)	.268 (.055)	.219 (.048)	.292 (.041)	.289 (.041)	.326 (.045)	.290 (.040)
Age	.080 (.059)	.069 (.060)	.086 (.070)	.072 (.059)	.043 (.059)	.044 (.060)	.035 (.066)	.038 (.059)
Age2	-.266 (.192)	-.229 (.192)	-.287 (.226)	-.234 (.190)	-.163 (.186)	-.163 (.189)	-.140 (.207)	-.143 (.187)
Age3	.341 (.259)	.288 (.259)	.359 (.304)	.294 (.246)	.264 (.246)	.259 (.250)	.244 (.274)	.233 (.248)
Age4	-.148 (.125)	-.122 (.124)	-.147 (.144)	-.124 (.123)	-.152 (.116)	-.148 (.118)	-.148 (.129)	-.136 (.118)
Ln(income)	.236 (.025)	.241 (.025)	.274 (.030)	.238 (.025)	.207 (.019)	.207 (.019)	.212 (.021)	.208 (.019)
yr9293	-.085 (.022)	-.091 (.023)	-.070 (.030)	-.089 (.022)	-.067 (.020)	-.071 (.020)	-.037 (.024)	-.070 (.020)
yr9394	-.130 (.024)	-.139 (.024)	-.121 (.031)	-.136 (.024)	-.094 (.021)	-.097 (.022)	-.043 (.026)	-.096 (.022)
yr9495	-.167 (.025)	-.179 (.026)	-.136 (.033)	-.175 (.025)	-.146 (.022)	-.152 (.023)	-.057 (.027)	-.149 (.022)
yr9596	-.210 (.025)	-.222 (.026)	-.145 (.034)	-.217 (.025)	-.158 (.023)	-.164 (.024)	-.088 (.028)	-.161 (.023)
yr9697	-.207 (.027)	-.221 (.028)	-.196 (.035)	-.216 (.027)	-.200 (.023)	-.207 (.024)	-.143 (.028)	-.203 (.024)
yr9798	-.232 (.027)	-.238 (.028)	-.240 (.036)	-.233 (.027)	-.195 (.025)	-.196 (.025)	-.173 (.029)	-.193 (.025)
yr9899	-.280 (.028)	-.294 (.029)	-.278 (.036)	-.287 (.028)	-.235 (.025)	-.242 (.025)	-.194 (.029)	-.238 (.025)
yr9900	-.833 (.030)	-.288 (.027)	.183 (.038)	-.242 (.029)	-.745 (.026)	-.164 (.024)	.348 (.032)	-.128 (.026)
yr0001	-.361 (.029)	-.374 (.029)	-.356 (.035)	-.365 (.029)	-.332 (.025)	-.340 (.026)	-.274 (.030)	-.333 (.026)
yr0102	-.329 (.029)	-.344 (.030)	-.346 (.036)	-.335 (.030)	-.272 (.026)	-.280 (.027)	-.257 (.031)	-.275 (.027)
Cut1	.718	1.418	1.976	1.458	.22	1.005	1.086	.999
Cut2	1.53	2.306	2.892	2.339	1.028	1.884	1.992	1.874
Cut3	2.423	3.696		3.681	1.915	3.291		3.228
Cut4	3.733				3.245			
Log Likelihood	-30074.2	-28436.4	-16197.9	-28797.5	-39110.8	-36730.8	-23046.3	-37223.2
RESET (p value)	.29 (.59)	.00 (.98)	.03 (.87)	.02 (.88)	.17 (.68)	2.35 (.13)	3.35 (.07)	2.92 (0.09)

Notes:

1. Estimated coefficients for the variables Widow, Single, Div/Sep, Non-White, HHSIZE, NCH04, NCH511, NCH1218 have been suppressed to conserve space. Full results are available from the authors upon request.
2. Standard errors in parentheses.

Table 8. POP models with interaction terms, including wave 9 data

	Men NT = 24791			Women NT = 31225		
	SAH 5	SAH 4	L & L	SAH 5	SAH 4	L & L
yr9900	2.947 (1.969)	3.237 (1.525)	3.493 (1.907)	-.154 (1.677)	2.024 (1.369)	2.392 (1.737)
Interactions with yr9900:						
Degree	.111 (.068)	-.028 (.058)	-.023 (.069)	.085 (.062)	-.067 (.054)	.041 (.069)
HND/A	.092 (.051)	-.019 (.044)	.057 (.055)	.088 (.048)	-.034 (.042)	.094 (.052)
O/CSE	.071 (.052)	-.018 (.044)	.032 (.055)	.045 (.041)	-.056 (.036)	.030 (.045)
Age	-.227 (.156)	-.180 (.121)	-.221 (.150)	-.038 (.130)	-.109 (.108)	-.200 (.136)
Age2	.590 (.456)	.481 (.357)	.603 (.440)	.033 (.367)	.253 (.308)	.547 (.388)
Age3	-.665 (.566)	-.549 (.449)	-.711 (.546)	.027 (.439)	-.267 (.374)	-.645 (.468)
Age4	.270 (.254)	.225 (.204)	.301 (.245)	-.032 (.189)	.110 (.164)	.278 (.204)
Ln(income)	-.067 (.043)	-.102 (.040)	-.078 (.045)	.022 (.038)	-.050 (.032)	-.003 (.040)
Cut1	.700	1.500	1.472	.075	1.012	.991
Cut2	1.513	2.388	2.354	.885	1.891	1.866
Cut3	2.407	3.779	3.696	1.774	3.298	3.222
Cut4	3.718			3.103		
Log Likelihood	-30062.6	-28429.0	-28786.6	-39090.654	-36722.1	--37205.155
RESET (p value)	.05 (.820)	.06 (.804)	0.05 (.826)	3.19 (.074)	1.86 (.172)	2.20 (.138)

Notes:

1. Estimated coefficients for the variables Widow, Single, Div/Sep, Non-White, HHSIZE, NCH04, NCH511, NCH1218 and their interactions with the wave 9 dummy variable have been suppressed to conserve space. Full results are available from authors upon request.
2. Standard errors in parentheses.

Table 9. Generalized Ordered Probit (GOP) – restricted models

		SAH (not w9) NT = 22539	Men SAH 4 (not w9) NT = 2259	SAH 4 NT = 24791	SAH 5 (not w9) NT = 28386	Women SAH 4 (not w9) NT = 28386	SAH 4 NT = 31225
Degree	β	.305 (.065)	.305 (.065)	.309 (.065)	.270 (.064)	.269 (.064)	.265 (.064)
	Cut1	-.121 (.258)	-.191 (.114)	-.191 (.117)	-.221 (.133)	-.256 (.086)	-.268 (.087)
	Cut2	-.191 (.115)	-.130 (.065)	-.141 (.065)	-.257 (.086)	-.193 (.062)	-.207 (.062)
	Cut3	-.130 (.065)			-.193 (.062)		
HND/A	β	.192 (.054)	.192 (.054)	.191 (.054)	.203 (.054)	.203 (.054)	.198 (.054)
	Cut1	-.079 (.121)	-.037 (.079)	-.039 (.080)	-.149 (.107)	-.115 (.068)	-.124 (.068)
	Cut2	-.037 (.079)	-.141 (.052)	-.154 (.052)	-.114 (.068)	-.128 (.052)	-.146 (.052)
	Cut3	-.141 (.052)			-.128 (.052)		
O/CSE	β	.147 (.054)	.147 (.054)	.151 (.054)	.237 (.048)	.237 (.048)	.229 (.048)
	Cut1	-.105 (.117)	-.096 (.080)	-.081 (.081)	-.061 (.086)	-.069 (.062)	-.077 (.063)
	Cut2	-.096 (.081)	-.116 (.053)	-.119 (.052)	-.070 (.062)	-.085 (.046)	-.103 (.047)
	Cut3	-.116 (.053)			-.085 (.046)		
Ln(income)	β	.191 (.026)	.191 (.026)	.183 (.026)	.195 (.024)	.195 (.024)	.195 (.024)
	Cut1	-.066 (.059)	-.150 (.039)	-.164 (.038)	.051 (.049)	-.001 (.031)	-.007 (.032)
	Cut2	-.149 (.039)	-.075 (.028)	-.080 (.027)	-.001 (.031)	-.028 (.024)	-.027 (.024)
	Cut3	-.075 (.028)			-.028 (.024)		
yr9900	β			-.631 (.039)			-.583 (.037)
	Cut1			-.786 (.067)			-.860 (.060)
	Cut2			-.624 (.050)			-.702 (.044)
	Cut3						
Cut1	2.98 (1.55)	3.13 (1.15)	2.93 (.003)	-4.18 (1.75)	.637 (.993)	.79 (.429)	
Cut2	3.12 (1.14)	2.84 (.830)	3.37 (.001)	.653 (.988)	2.06 (.789)	2.69 (.007)	
Cut3	2.84 (.829)	2.66 (.836)	3.18 (.001)	2.072 (.787)	2.63 (.863)	3.20 (.001)	
Cut4	2.65 (.838)			2.63 (.864)			
Log Likelihood	-26727.2	-26048.4	-28167.8	-34971.8	-33758.2	-36439.1	

Notes:

1. Estimated coefficients for the variables Widow, Single, Div/Sep, Non-White, HHSIZE, NCH04, NCH511, NCH1218 have been suppressed to conserve space. Full results are available from authors upon request.
2. Standard errors in parentheses.

Table 10. Generalized Ordered Probit (GOP) - unrestricted model

		Men SAH 4 NT = 24791	Women SAH 4 NT = 31225
Degree	β	.308 (.066)	.270 (.065)
	Cut1	-.183 (.115)	-.256 (.087)
	Cut2	-.127 (.065)	-.193 (.063)
HND/A	β	.194 (.054)	.204 (.055)
	Cut1	-.033 (.080)	-.115 (.068)
	Cut2	-.140 (.052)	-.129 (.053)
O/CSE	β	.149 (.054)	.238 (.049)
	Cut1	-.093 (.081)	-.069 (.063)
	Cut2	-.115 (.053)	-.086 (.047)
Ln(income)	β	.188 (.027)	.196 (.025)
	Cut1	-.165 (.039)	-.002 (.032)
	Cut2	-.081 (.028)	-.029 (.024)
Yr9900	β	2.365 (3.043)	.283 (2.697)
	Cut1	1.394 (6.457)	-.461 (4.878)
	Cut2	-4.010 (3.972)	-3.315 (3.561)

Interactions			
With yr9900:			
Degree	β	.002 (.104)	-.080 (.101)
	Cut1	-.223 (.251)	-.196 (.235)
	Cut2	-.208 (.142)	-.227 (.138)
HND/A	β	-.052 (.087)	-.109 (.084)
	Cut1	-.106 (.141)	-.124 (.153)
	Cut2	-.226 (.109)	-.281 (.107)
O/CSE	β	.030 (.088)	-.140 (.076)
	Cut1	.208 (.139)	-.120 (.123)
	Cut2	-.060 (.108)	-.289 (.094)
Ln(Income)	β	-.055 (.053)	.002 (.051)
	Cut1	†	-.118 (.091)
	Cut2	-.003 (.066)	.021 (.066)
Cut1		2.78 (.005)	.639 (.994)
Cut2		3.44 (.001)	2.069 (.789)
Cut3		3.14 (.002)	2.636 (.864)
Log Likelihood		-28146.7	-36410.2

Notes:

1. Estimated coefficients for the variables Widow, Single, Div/Sep, Non-White, HHSIZE, NCH04, NCH511, NCH1218 and their interactions with the wave 9 dummy variable have been suppressed to conserve space. Full results are available from authors upon request.
2. Standard errors in parentheses.
3. † Maximum likelihood estimation of this model encountered numerical optimisation problems. An empirical fix to this was achieved by imposing the restriction that the coefficient on the variable Yr9900 x Ln(Income) for the second cut-point was equal to zero.

Table 11. Results for the Concentration Indices and Mobility Index - all individuals

	SAH 5	SAH 4	L&L	SAH 3	
Short-term CI	wave 1	0.020	0.025	0.025	0.017
	wave 2	0.021	0.027	0.027	0.019
	wave 3	0.023	0.030	0.030	0.020
	wave 4	0.024	0.031	0.031	0.023
	wave 5	0.025	0.032	0.032	0.022
	wave 6	0.025	0.032	0.032	0.023
	wave 7	0.027	0.034	0.034	0.026
	wave 8	0.028	0.035	0.035	0.026
	wave 9	0.039	0.028	0.035	0.023
	wave 10	0.029	0.037	0.037	0.028
	wave 11	0.029	0.038	0.038	0.028
Long - term CI	0.030	0.036	0.037	0.026	
Mobility Index	-0.135	-0.134	-0.133	-0.129	

Table 12. Results for the Concentration Indices and Mobility Index - men

	SAH 5	SAH 4	L&L	SAH 3	
Short-term CI	wave 1	0.016	0.019	0.019	0.014
	wave 2	0.017	0.022	0.022	0.015
	wave 3	0.020	0.026	0.026	0.017
	wave 4	0.023	0.030	0.030	0.025
	wave 5	0.024	0.031	0.031	0.020
	wave 6	0.024	0.031	0.031	0.022
	wave 7	0.026	0.033	0.033	0.026
	wave 8	0.027	0.034	0.034	0.026
	wave 9	0.034	0.026	0.032	0.022
	wave 10	0.031	0.039	0.039	0.030
	wave 11	0.030	0.039	0.039	0.029
Long - term CI	0.028	0.346	0.035	0.025	
Mobility Index	-0.131	-0.133	-0.130	-0.122	

Table 13. Results for the Concentration Indices and Mobility Index - women

	SAH 5	SAH 4	L&L	SAH 3	
Short-term CI	wave 1	0.022	0.028	0.028	0.0200
	wave 2	0.023	0.029	0.029	0.021
	wave 3	0.024	0.031	0.031	0.021
	wave 4	0.023	0.030	0.030	0.021
	wave 5	0.024	0.032	0.032	0.022
	wave 6	0.024	0.031	0.031	0.023
	wave 7	0.026	0.033	0.003	0.024
	wave 8	0.027	0.031	0.034	0.026
	wave 9	0.041	0.029	0.035	0.022
	wave 10	0.026	0.033	0.033	0.025
	wave 11	0.027	0.035	0.035	0.027
Long - term CI	0.029	0.036	0.036	0.026	
Mobility Index	-0.14071636	-0.136	-0.138	-0.136	