

The British Household Panel Survey and its income data

Stephen P. Jenkins
Institute for Social and Economic Research
University of Essex

No. 2010-33
6 October 2010



INSTITUTE FOR SOCIAL
& ECONOMIC RESEARCH

Non-technical Summary

This paper is concerned with the construction of measures of household income for respondents to the British Household Panel Survey (BHPS). The focus is on a measure of net income. Sometimes also referred to as disposable income, net income refers to total money income from the labour market, cash benefits and tax credits including the state retirement pension, savings and investment income including private and occupational pensions, and private transfers, from which is deducted income tax payments, employee national insurance and occupational pension contributions, and local tax payments. (Gross income equals total income prior to the deductions referred to.)

The household income variable provided in the official release of BHPS data is a gross income measure. However it is widely agreed that analysis of the personal distribution of income, inequality and poverty – trends over time and longitudinal dynamics – should use net income rather than gross income measures. ISER researchers have derived net income measures using BHPS data and, on an ad hoc basis, made these available as unofficial supplements to the main BHPS release files.

BHPS net income variables are widely used. They also form part of the BHPS component of the Cross-National Equivalent File. (The CNEF contains comparable household panel data from the BHPS, Australian HILDA, Canadian SLID, German SOEP, US PSID, and the Swiss Household Panel Survey.)

It is important that users of the BHPS net income files are aware of the nature of the BHPS sample design and its other features, and also of how the net household income variables are constructed.

This paper addresses these needs. It provides a self-contained introduction to the BHPS, concentrating on aspects relevant to analysis of the distribution of household income. First, I discuss various BHPS design features drawing on the BHPS Quality Profile (Lynn 2006). Second, I discuss how data on net household income are derived. The BHPS net household income definition is modelled on that used in Britain's official personal income distribution statistics (*Households Below Average Income*, HBAI) based on the much larger and specialist cross-sectional income survey, the *Family Resources Survey*. The BHPS definitions are contrasted with those employed in the HBAI, and there are also comparisons of estimates of cross-sectional summaries of the income distribution. I show that the BHPS distributions track the HBAI ones relatively well over time.

The British Household Panel Survey and its income data

by

Stephen P. Jenkins

Abstract

This paper provides a self-contained introduction to the British Household Panel Survey (BHPS), concentrating on aspects relevant to analysis of the distribution of household income. I discuss BHPS design features and how data on net household income are derived. The BHPS net household income definition is modelled on that used in Britain's official personal income distribution statistics (*Households Below Average Income*, HBAI). I show that cross-sectional BHPS distributions track corresponding HBAI ones relatively well over time.

Key words: British Household Panel Survey, household income, net income, disposable income, Households Below Average Income, income distribution, inequality, poverty
JEL codes: C81, D31, I32

Acknowledgements

A version of this paper will appear as Chapter 4 of Jenkins (2011). The BHPS net income files releases have been developed by me in conjunction with Elena Bardasi, Sarah Jarvis, Horacio Levy, and John Rigg. I have drawn on the knowledge of Nick Buck, Peter Lynn, and Holly Sutherland. I thank all these colleagues but bear full responsibility for the material presented here.

Correspondence

Stephen P. Jenkins, Institute for Social and Economic Research, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, United Kingdom. Email: stephenj@essex.ac.uk.

The British Household Panel Survey and its income data

This paper provides a self-contained introduction to the British Household Panel Survey (BHPS), concentrating on aspects relevant to analysis of the distribution of household income. First, in Sections 1–6, I discuss various BHPS design features drawing on the BHPS Quality Profile (Lynn 2006). See also BHPS Documentation Team (2009) or, for a concise overview, Department for Work and Pensions (2008: Appendix 2).

Second, in Sections 7–10, I discuss how data on net household income are derived. The BHPS net household income definition is modelled on that used in Britain’s official personal income distribution statistics (*Households Below Average Income*, HBAI) based on the much larger and specialist cross-sectional income survey, the *Family Resources Survey* (FRS). The BHPS definitions are contrasted with those employed in the HBAI, and there are also comparisons of estimates of key cross-sectional summaries of the income distribution. I show that the BHPS distributions track the HBAI ones relatively well over time. (Analogous checks of longitudinal features such as poverty transition rates or income mobility are not possible because there is no comparable longitudinal data source.)

BHPS net income variables are widely used. They also form part of the BHPS component of the Cross-National Equivalent File. (The CNEF contains comparable household panel data from the BHPS, Canadian SLID, German SOEP, US PSID, and the Swiss Household Panel Survey: see Frick et al. 2007.) It is important that users of the BHPS net income files are aware of the nature of the BHPS sample design and its other features, and also of how the net household income variables are constructed. By the end of the paper, readers should have a good appreciation of the strengths and weaknesses of the BHPS and its net household income data.

1. BHPS: its design and other features

The BHPS is a classic example of a household panel survey designed to address a wide range of research topics (Jenkins 2011: Chapter 3). The dynamics of household income in general, and poverty dynamics in particular, were among the core research topics initially envisaged. Other topics include labour market behaviour, education and training, housing, household formation, dissolution and fertility, social and political attitudes and values, health. The general purpose nature of the survey means that there is inevitably some degree of

compromise in the specification of measures relating any particular topic area, including household income.

The BHPS was originally designed as an indefinite life panel but has now ended, at least in its current form. There have been 18 waves of annual interviewing, with the last wave completed in survey year 2008. The BHPS sample is now incorporated into Understanding Society – the UK Household Longitudinal Study, providing a sample with a long run of panel data that will supplement data for new samples of respondents with whom interviewing began in 2009 (<http://www.understandingsociety.org.uk>).

The first wave of the BHPS was intended to represent the private household population of Great Britain south of the Caledonian Canal. Great Britain consists of England, Wales and Scotland. The United Kingdom is Great Britain plus Northern Ireland. (The Caledonian Canal traverses northern Scotland and, to its north, population density is very low.) Residential addresses were selected using a equal-probability clustered and stratified design from the Postal Address File (PAF), the source also used to select samples for major national cross-sectional surveys, and then all households at each address (with a selection of households for the 3 per cent of addresses with more than three households). As explained in earlier sections, additional samples drawn from Scotland, Wales, and Northern Ireland were added to the original sample in the mid- to late-1990s. I do not discuss the nature of these samples further: see Lynn (2006) for details.

The BHPS design means that individuals residing in institutions were not eligible for selection, for example residents of nursing homes, military barracks, or student halls of residence. The National Equality Panel (2010: Appendix 3) discusses the size and composition of the non-household population in the UK, and estimates its size to be around 2 per cent of the total population. The BHPS's design also means that people without a residential address are excluded from sample coverage. Since most homeless people are also destitute, it is clear that the BHPS – like all other national surveys in Britain – undercounts the number of people who are poor but the numbers are very small. The National Equality Panel's conclusion is that, because of the wide range of incomes in the non-household population, 'the data ... on the household population, while incomplete, can still present a fair picture of the circumstances of the population as a whole' (2010: 411). I assume that this is the case here as well.

The BHPS definition of a 'household', and the unit to which 'household income' refers, is the same as that used in the UK's national statistical practice (Lynn 2006: 16), that is 'one person living alone or a group of people who either share living accommodation, or

share one meal a day and who have the address as their only or main residence'. Living together requires six months continuous residence. This means that students are included if their term-time address was selected unless they were living in a hall of residence.

The BHPS definition of the 'family' coincides with the definition in the British tax-benefit system (also known as the 'benefit unit'), that is, a single person or a couple living together with or without dependent children. A dependent child is aged less than 16 years, or more than 16 years but under 19 years and unmarried, in full-time non-advanced education and living with his/her parent or parents. Parent status is defined by blood, adoption, or guardianship. A household may contain several benefit units. Examples of this are a non-dependent child living with his parents (two benefit units), or three single adults sharing a house (three benefit units). The choice of the household versus family as the income recipient unit can have marked differences on estimates of statistics such as poverty rates (Jenkins: 2011, Chapter 2).

The BHPS definitions differ from those used in other panel surveys. In the PSID for instance, the unit of focus is the 'family', defined to be 'a group of people living together as a family. They are generally related by blood, marriage, or adoption, but unrelated persons can be part of a FU if they are permanently living together and share both income and expenses'. (<http://psidonline.isr.umich.edu/Guide/FAQ.aspx#90>. See also Hill 1992.) Stated thus, the definition is close to the BHPS's definition of a household. But income distribution researchers using the PSID often focus on family income definitions using a narrower, US Census Bureau, definition of the family that excludes unrelated individuals – who are treated as one-person families. (See for example Gottschalk and Danziger 2005.) The definition of the income-receiving unit used in most UK income distribution research is therefore wider than in much US research on income distribution.

All individuals, adults and children, enumerated in BHPS respondent sample households at wave 1 became part of the longitudinal sample, and have been followed over time. Each person in this group is an Original Sample Member (OSM). New permanent members of the sample joining the longitudinal sample after the initial BHPS wave are either babies born to or adopted by OSMs after the initial wave, or the parent of a longitudinal sample member who joins the household of an OSM. So, if an OSM got married in 1994 and the couple had a child together, the spouse and the baby would both become permanent sample members (PSMs). If the partners subsequently divorced and lived separately, they and their baby would each be followed as part of the longitudinal sample. Otherwise, all persons joining the household of a longitudinal sample member are interviewed in the waves at which

they are present in the sample household, but they are not followed if they leave that household – they are temporary sample members (TSMs).

The fieldwork for wave 1 was carried out between 1 September and 1 December in 1991. In subsequent waves, the fieldwork period was broadly the same, except that there was an extension running, in principle, through to the following May, in order to try and re-interview respondents who were difficult to trace or contact or to secure response from. But fieldwork remained heavily concentrated in the Autumn: the modal interview month was October for waves 1–5, and September thereafter, with at least 80 per cent of interviews undertaken in either September or October (Lynn 2006: Table 21). Wave 9 was the only exception when around 12 per cent of interviews took place in the January–May period, arising because of the difficulties associated with the introduction of CAPI. The concentration of fieldwork has the advantage of helping to control for seasonal effects on response, including effects associated with Christmas. But, equally, these aspects cannot be studied.

The main survey instruments are an Individual Questionnaire answered by each adult member of a sampled household (lasting around 45 minutes on average), and a Household Questionnaire answered by one of these persons on behalf of the household (a further 15 minutes on average). There is also an adult self-completion questionnaire and, from wave 4 onwards, a self-completion questionnaire for children aged 11–15. (The repeated responses to this youth questionnaire are sometimes referred to collectively as the ‘British Youth Panel’.) When children reach the age of 16, they become full sample members in their own right, and interviews are based on the instruments for adults. In addition, at the first three waves of the BHPS, there were a number of additional modules focusing on respondent life histories prior to the initial wave, using respondent retrospective recall to collect data about work and jobs, partnerships (legal and cohabiting) and fertility. All the information used to collect the various components of household income are derived, however, from the two main instruments (the Individual and Household questionnaires) and additional data about the household derived as part of the survey process (such as the enumeration of its members).

2. Sample size

As Lynn (2006: 17–8) documents, the BHPS initial sample selection process yielded 8,167 addresses, with fieldwork identifying 13,840 persons at those addresses, including 10,751 aged 16 or older eligible for personal interviews. The number of personal interviews achieved

at wave 1 (including proxy interviews) was 10,264 spread across 5,505 households. As the panel has matured, the number of achieved interviews with main sample OSMs has fallen gradually, reaching just under 8,155 at Wave 7, and 7,120 by wave 13 (Lynn 2006: Table 4). But, at the same time, the number of personal interviews achieved with PSMs has also increased gradually, from 10 at wave 2 to 240 at wave 7 and 299 at wave 13. The corresponding numbers of TSM personal interviews are 484, 1,071, and 1,236. Thus, the total number of achieved individual interviews went from 9,845 at wave 2, to 9,466 at wave 7 to 8,655 at wave 13 (Lynn 2006: Table 4). These trends reflect attrition from the original sample, but the numbers themselves cannot be used to infer response rates. On these, see below.

Longitudinal sample sizes are more difficult to derive than cross-sectional ones because numbers depend on the particular research issue addressed, and because there are many ways of looking at the data longitudinally, including for example using long sequences of repeated observations on individuals or pooling year-on-year transitions from successive years. Numbers depend on initial sample sizes and subsequent attrition. (Lynn 2006: 18).

Lynn (2006) provides six tables illustrating these points, including breakdowns by age. His Table 5 shows, for instance, that of the 6801 continuing OSMs with achieved interviews at wave 13, 5,481 provided a wave 1 response, and 4,648 provided interviews at all waves from 1 to 13. Lynn's Table 6 provides information about sequential response from wave 1 onwards. There were 4,653 respondents present at every wave from 1 to 13, but more than twice that who respondents at 1 or more waves (9,912). Table 7 repeats the analysis except that the calculations are for sequential wave response for those present at wave 5. 5,481 provided interviews at every wave, wave 3 through 13, but 8,162 provided 1 or more interviews. Table 8 summarizes the number of pairs of successive waves at which respondents gave a full interview – the sample size relevant to estimation of (average) transition rates such as proportions moving into and out of poverty. In this case sample sizes are very large, over 110,000 (these numbers include the extension samples). Tables 9 and 10 show sample sizes for numbers of events. Table 9 shows that, over waves 1–13, the number of employment to employment transitions is very large (more than 60,000) but, for some transitions of particular policy interest such as those from employment to unemployment, the numbers are much smaller (around 1,400). There is a similar issue with the numbers of respondents moving into or out of poverty being small relative to the number staying non-poor. Finally, Lynn's Table 10 shows numbers of demographic events experienced by respondents over the 13 interview waves, referring to partnership formation and dissolution,


and arrival and departure of children. Here the number of events is of the order of one to two thousand, that is, relatively small, especially once breakdowns by other characteristics are undertaken.

3. Response rates, including attrition

Response rates can be calculated in many ways. A first approach is similar to that used for cross-section surveys, documenting wave by wave, data about field outcomes and response rates. Lynn (2006: Tables 25–37) provides this type of information for waves 1–13 for the original BHPS sample. For example, at wave 1, there was complete coverage within 69 per cent of the 7,491 eligible households including proxies, and partial coverage with 74 per cent. In terms of individual adults ($n = 10,751$), 92 per cent provided full interviews, and a further 2 per cent provided proxy interviews. The most reason for non-response was refusal (4 per cent) with reasons such as non-contact or absence, and age, infirmity, disability or language difficulty being relatively unimportant. At wave 13 (individual adults $n = 9,956$), 87 per cent provided full interviews, and the refusal rate was 10 per cent.

The full 13 wave pattern is summarized in Figure 1, which shows that the cross-sectional response rate for individual interviews has hovered around 90 per cent after an initial fall and recovery as the panel settled in. A small downward trend in response rates is perhaps discernable towards the end of wave 13. Correspondingly refusal rates typically fluctuate at around 10 per cent, with perhaps a slight upward trend towards the end of the period. Observe that there is no apparent change in response rates around wave 9 when CAPI was introduced. These rates (and trends) are in line with other leading household panel surveys such as the German SOEP and the Australian HILDA. For details of response in these surveys, see for example Kroh (2009) and Watson and Wooden (2006).

Figure 1. Individual interview response rate (%) and refusal rate (%), BHPS Waves 1–13



Note: Calculations based on BHPS original sample.

Source: Author's derivation from Lynn (2006: Tables 25–37).

One-wave-at-a-time response rates are less useful for assessing household panel surveys because, for most analysis purposes, it is longitudinal response or non-response and its cumulation over time that is relevant, whether for longitudinal statistics such as poverty transition rates, or cross-sectional statistics (after wave 1) such as poverty rates (Lynn 2006: 75). But, in this case, there is no single response rate calculation, as it depends on the combination of waves that the analyst wishes to use, and the number of possibilities is very large (Lynn 2006: 75). Inevitably, therefore, Lynn focuses on a small number of summary measures of response rates, which I now review.

Table 1 summarizes interview outcomes for individuals who were full respondents at wave 1. Ineligibility arises because for example sample members move abroad or die, and the rates have risen as the panel matures (remember the rates refer to wave 1 respondents only – an ageing cohort, by definition). The second column shows the fraction of those eligible (those alive, and in Britain south of the Caledonian Canal) who responded, and this rate has fallen as the panel matures from almost 88 per cent at wave 2 to around 65 per cent at wave 13. The final two columns summarize different types of longitudinal response rate. One shows the proportion of wave 1 respondents who also responded at every subsequent wave,

wave on wave. This proportion has fallen markedly as the panel matured from around 88 per cent at wave 2 to 55 per cent at wave 13, but with the pace of decline decreasing over time (perhaps suggesting that there is a ‘hard core’ of compliant respondents). Although the decline in response rates over the period as a whole appears large, observe that requiring participation at every wave is unnecessary for many types of analysis.

The final column of Table 1 shows wave on wave response rates for individuals who have responded at each wave up to the previous one. After the initial fall-off – a phenomenon observed for all household panels – the wave on wave retention rate is very high, at around 95 per cent or higher. Again, this pattern and level is shared by other household panel surveys such as the German SOEP and Australian HILDA (Kroh 2009; Watson and Wooden 2006). If the calculations in Table 4.1 are repeated, but using numbers of people enumerated in wave 1 households (a group more relevant for household income distribution analysis) rather than respondents, then the trends in each column of the new table are the same, but corresponding percentages are slightly higher in the new table (Lynn 2006: Table 68).

Table 1. Interview outcomes for BHPS wave 1 respondents

Wave	% ineligible	% of eligible responding	% of eligible responding at all waves	Wave on wave response rate (interviewed all waves up to previous)
2	1.4	87.7	87.7	87.7
3	2.9	81.5	79.1	90.3
4	4.3	79.9	74.8	94.9
5	5.6	76.8	70.6	94.8
6	6.9	77.3	68.7	97.6
7	8.4	76.0	66.7	97.6
8	9.5	74.1	64.7	97.4
9	10.5	72.1	62.4	97.0
10	12.0	70.4	60.0	96.7
11	12.8	68.4	59.3	96.1
12	13.7	66.6	57.1	96.5
13	14.8	64.9	55.1	96.8

Source: Lynn (2006: Table 67).

Lynn (2006) and Uhrig (2008) document which types of sample member are most likely to drop out of the BHPS. Uhrig (2008) fits multivariate discrete time hazard regression models to data from waves 1–14, modelling the hazard rate of drop-out from the sample for wave 1 respondents, and hence the number of waves to first drop-out from the sample. Respondents who become ineligible are treated as right-censored observations. Uhrig fits

models of overall non-response, and also separate models for sample drop-out due to non-contact and due to refusal since the determinants are likely to differ. (Nicoletti and Peracchi 2005 demonstrate the importance of this distinction in the context of the ECHP.) Uhrig's estimates suggest that higher rates of non-contact are associated with physical impediments to contact (such as living in gated accommodation or apartment blocks), and characteristics associated with a respondent being more likely to be away from home or to be geographically mobile. Indicators of lack of interest in the survey and of a low motivation to participate are predictive of refusal per se. The study does not, however, examine the magnitude of the differential attrition associated with each characteristic.

Lynn (2006: Tables 67–71) compares the distribution of characteristics among wave 1 respondents, with the distribution of characteristics of those who responded at some wave t but not at every wave from 1 to t , and those who responded at every wave from 1 to t , taking t = wave 5 or wave 13. Differences between the distributions are indicative of differential attrition. Lynn's summary of his findings states that those who failed to respond on at least one occasion included disproportionate numbers of people with the following characteristics at wave 1: aged 16–24 years; never married; unemployed; no qualifications; not active in any organisations; resident of Inner London, West Midlands conurbation, or Merseyside; tenant of local authority or housing association housing; and in the poorest 40 per cent of the income distribution. He also remarks, however, that 'although under-representation of these groups is statistically significant, the actual magnitude of under-representation is generally small. Furthermore, these differences apparent at the data collection stage are largely removed by the application of the weighting' (2006: 76). I discuss the BHPS weights below.

4. Item non-response

Even if sample members are counted as being respondents at a particular wave, they may not give complete responses to every question, either because they simply don't know the answer or because they are unwilling to provide the information. (Data may also be missing due to interviewer error such as skipping a question, though the introduction of CAPI should largely eliminate this problem.) This is the situation known as item non-response. Income details are examples of relatively sensitive items likely to be subject to this problem, though its prevalence may fall as the panel matures and respondents establish trust in the survey (Lynn 2006: 42).

Lynn (2006) provides information about levels of item non-response in BHPS waves 1–13. When non-response is considered in relation to all BHPS variables, its prevalence is relatively small, fluctuating around about 2 per cent in both the Individual questionnaire (Table 50) and the Household questionnaire (Table 51). In the former case, and restricting attention to variables with more than 100 cases eligible to answer, item non-response ranges from 1.22 per cent (wave 7) to 2.46 per cent (wave 13), with no obvious trend over time or break points associated with the introduction of CAPI in wave 9. In the latter case, the range is from 1.78 per cent (wave 2) to 5.73 per cent (wave 10). The higher rates, apparent at waves 7–10, were associated with the introduction of additional follow-up questions concerning amounts spent on white goods. By wave 13, the item non-response rate was below 2 per cent again.

Of particular concern for the study of income dynamics are, not the overall rates of item non-response, but the rates associated with income and related items. Lynn (2006: Table 51) reports that these rates are markedly higher than the overall rates. For example, a core component of the calculation of total household income is ‘usual pay at last payment’ for those in employment (see below for details). The non-response rate among employees for this variable was 15.1 per cent at wave 1 (the maximum among the wave 1–13 rates), 6.91 per cent at wave 9 (the minimum) and 11.03 per cent at wave 13, with fluctuation over time. For ‘net profit’ from self-employment, the rates of non-response are substantially higher, ranging between 32.8 per cent (wave 3) to 47.16 per cent (wave 9), again with fluctuation over time. (Note that the numbers of cases is much smaller: self-employment is much less prevalent than employment.) These rates can be contrasted with the rates for marital status for which item response is near zero, or health status for which the rate is always less than 1 per cent (Lynn 2006: Table 51).

The discussion so far has been of item non-response on items provided by individual respondents on behalf of themselves (for example their pay if an employee) or on behalf of the household as a whole (for example questions related to the dwelling). But, for analysis of household income, it is non-response at the household level and on a combination of variables that is important. Data on total household income is incomplete if there is item non-response for any member of the household or, indeed, unit non-response by any one household member – even if there is complete response by all the other members.

In sum, item non-response is clearly an important issue for analysis of BHPS data on income. Researchers may simply omit cases with non-response, which may reduce sample numbers to unacceptable levels and introduce sample selection biases into estimates if not

controlled for. Alternatively, researchers might use imputations for the missing data, thereby maintaining sample sizes, but run the risk of introducing measurement errors that will contaminate estimates. This second approach is what most analysts employ. BHPS imputation procedures and the treatment of partial household non-response are discussed shortly.

5. Adjusting the data after collection: (i) BHPS weights

The BHPS data release contains a large number of weighting variables that may be used to account for non-response in estimation, and which have been derived following conventional survey methodological practice. Here I discuss only those variables applicable to the original BHPS sample, and ignore the weights constructed for use with the extension samples. There are separate sets of weights for households, respondent individuals, and enumerated individuals (all persons within sample households). And there are cross-sectional weights for analysis of each wave taken separately, and longitudinal weights for longitudinal analysis.

The foundation of all the weighting variables for all waves is the set derived for wave 1, as these account for the unequal probabilities of selection of each address (determined as part of the design of the survey). These design weights are adjusted to take account of non-response at the household level, and non-response of individuals within households. There are then some ‘post-stratification’ adjustments to make the sample more representative of Britain’s private household population, with the modifications aligning sample distributions with data on the distributions of housing tenure, household size, number of cars, age and sex, available from the 1991 national Census. Finally, the resulting weights are trimmed in order that sampling variances are not unduly affected by outlier values, and then scaled so that their sum corresponds to the relevant achieved sample size. This procedure is used to derive wave 1 weights for households, respondent individuals, and enumerated individuals.

After wave 1, there are both cross-sectional and longitudinal weights for each of these groups, except that there are no longitudinal household weights because there is no valid concept of a longitudinal household. See the discussion in Jenkins (2011: Chapter 2).

The BHPS longitudinal respondent weights for some wave t are non-zero for all individuals who gave a full interview at every wave up to and including wave t , and also for children at wave $t-1$ who became full sample members at t , but the weights are zero for TSMs. The longitudinal enumerated individual weights at t are non-zero for all those enumerated in respondent households at every wave up to and including wave t . For both sets

of weights, the longitudinal weight at some wave s is the product of the initial wave 1 weight and weights adjusting for sample drop-out between each successive pair of waves thereafter (wave 1 to wave 2, wave 2 to wave 3, and so on up to and including wave s).

To derive the weighting adjustments, sample members were allocated to a large number of classes according to characteristics perceived as predictive of non-response or of particular interest to researchers. Within each class, that is, conditional on observed characteristics, it is assumed that response status is random. The inverse of the within-class response rate is used as the weight for all the responding cases who fall within the class (which is then further adjusted using post-stratification weights as described above). Clearly, the construction of the classes is crucial and, for this, the BHPS staff use an ‘automatic interaction detection’ procedure (as implemented in the SPSS CHAID module), which facilitates derivation of a meaningful number of classes while at the same time avoiding problems of small cell sizes. The procedure is analogous to running a probit or logit regression with response status as the dependent variable and a large number of explanatory variables and their interactions, and then using the inverse of the predicted response probability as the weight.

For the longitudinal respondent weights, the classification variables include: whether moved from the previous address; age, sex, employment status, income total and composition, race, level of organisational membership, and educational qualifications, and various household characteristics such as region, housing tenure, number of cars, and ownership of consumer durables (Lynn 2006: 51). Children reaching the age of 16 are allocated a longitudinal respondent weight equal to the minimum of that of their parents. A similar procedure is used to derive longitudinal enumerated weights, with the main difference being that weighting classes were mainly based on the characteristics of the household and the household head. New-born children receive the average of the weights for their parents.

Derivation of the cross-sectional weights after wave 1 is complicated by the need to derive weights for new entrants after wave 1. A person marrying OSM does not have a wave 1 weight or a longitudinal weight. Moreover, their initial sample inclusion or response probabilities are not known and so assumptions have to be made about these. The ‘equal shares’ method that the BHPS uses (in common with other panels like SLID) in effect derives the unknown initial sample probabilities by supposing that the new entrants are like the other members of their household and uses the information about the members who were present in wave 1 to derive these probabilities. At each wave, the ‘average’ of the weights for the original members of the household, adjusted for subsequent drop-out, is shared with the

joiners at that wave. Cross-sectional respondent weights can be derived by a similar procedure, and a household weight is set equal to the cross-sectional individual weight, rescaled to correspond to the total number of households.

Although the rationale for weighting is relatively straightforward, it is clear from this discussion that the detailed derivations of the different types is complicated. Similar procedures are used across the major household panels. For example, the PSID has weights corresponding to the BHPS's longitudinal weights. The German SOEP, like the BHPS, has longitudinal and cross-sectional weights, except that the former are provided for each pair of successive waves up to and including wave t (unconditional on response prior to wave $t-1$), rather than the one set of weights for the full sequence of waves up to and including wave t . The Australian HILDA provides both types of longitudinal weights as well as cross-sectional weights.

Additional issues concerning the use of the BHPS weights are their general purpose nature and the derivation for individuals with a particular type of response pattern. Particular outcomes of research interest may be associated with particular patterns of non-response and, ideally, one should take account of this. And the BHPS longitudinal weights are non-zero only for original and permanent sample members with complete response at every wave up to and including the current one. Those with intermittent response are excluded and this is undesirable for some types of longitudinal analysis – a common example is analysis based on wave-on-wave transitions.

This discussion suggests several options. One might be to develop one's own set of specialist weights, appropriate to the research question under consideration. This is rarely done (but see for example Jenkins 2009). Aside from the complications involved, there are also conceptual problems. For example, researchers (including me) commonly examine transitions between states between two consecutive waves, pooling transitions from multiple pairs of waves. In this case, it is unclear what population of interest the pooled transitions are intended to represent and hence how either to calculate suitable longitudinal weights or to combine the weights typically supplied. This is not a decisive argument against using the weights supplied; rather the lesson is that differential non-response can lead to biased estimates, and so analysts should check that the sensitivity of their conclusions to different assumptions about non-response. An approach commonly used is to compare weighted and unweighted estimates and to claim robustness if they are similar.

Economists are sometimes resistant to using weights in estimation, especially in analysis based on multivariate regression modelling. Reasons for this view are rarely

documented but partly represent the idea that many of the variables included as explanatory variables in the regressions are the same as those that would be used to predict non-response and thence generate weights, and so there is a form of redundancy if weights are used. (On this, see Winship and Radbill 1998.) A contrary view would be that the interpretation of the impact of these variables is made more complicated (estimated coefficients reflect the impact of non-response as well as the substantive impact on the outcome) and, in any case, non-response related to survey design and which manifest themselves via non-contact rather than refusal are typically not included as explanatory variables. A second reason for economists' scepticism about weights is that the multivariate models of response used to derive them ignore the impact of unobservables.

The issue of whether to weight or not in the multivariate regression context has been helpfully clarified by Wooldridge (2002), who shows that 'the weighted estimator is consistent if we have an appropriate ignorability assumption and if we either know or can consistently estimate the sampling probabilities' (2002: 11). Ignorability refers to there being no unobservable factors associated both with the outcome of interest and the probability of response (conditional on observable variables). This is untestable without further assumptions about the nature of the association, and the standard approach is to suppose a model in which the additive 'error' terms in the outcome and response equations that characterize unobservables are independent of observables, and distributed multivariate normal. Identification of model parameters relies on there being variables that explain response that do not also affect the substantive outcome ('instruments'). In this approach, the test for ignorability is a test of the statistical significance of a correlation. For an application to poverty transitions, see Cappellari and Jenkins (2004) and to low pay transitions see Cappellari and Jenkins (2008). In both cases, attrition was found to be non-ignorable, but the magnitude of its impact is small.

6. Adjusting the data after collection: (ii) BHPS imputation procedures

Item non-response arises when a respondent is judged to have provided a full interview, but data are missing on some variables of interest. The issue, as with attrition, is whether the non-response is differential rather than random. If it is, then analysing data consisting of only non-missing cases – which is the default in most software packages – may lead to biased estimates. As an illustration of the scope for this, Frick and Grabke (2005) show, using German SOEP data, that income mobility estimates using only cases with non-missing data

markedly understate estimates derived from all cases including those with imputed income values. The differences may represent bias or the effects of measurement error introduced by imputation.

One approach to item non-response would be to develop suitable sample weights, exploiting the parallels with the case of non-response and attrition just discussed. The alternative, more commonly followed, and also adopted by the BHPS producers, is to make some specific assumptions about the item non-response process, and to use these to generate predicted values that are used to ‘fill in’ the missing values. At the same time, additional variables (‘imputation flags’) are created in order that researchers may identify cases with imputed values, and exclude them or derive alternative values if they wish.

BHPS imputation procedures focus on variables connected with income and housing costs (see the discussions of prevalence earlier). Two imputation approaches are used depending on the nature of the variable.

Hot-deck imputation is used for variables derived from questions with a limited number of valid responses – for example banded income from investments and savings, or some cash benefits. The procedure is very similar to that described earlier for the derivation of weights. Cases are placed in classes defined by combinations of variables believed to predict item non-response and then, assuming that response is random conditional on class membership, a case with a valid value for the variable of interest is randomly selected and that value imputed to a case from the same class with missing data. Classes are constructed using the same automatic interaction detection methods as described earlier.

When monetary amounts are missing, a regression-based imputation method known as ‘predictive mean matching’ is used for a number of primary variables from which some other income-related variables are derived. Taking cases with non-missing values of the variable of interest, a regression model is fitted with this variable as the dependent variable and a large number of explanatory variables thought to be predictive of response and their interactions. Predictions of the amount are derived from the fitted model for all cases, including those with missing amounts. The closest valid value to the predicted value of a missing case is then determined, and imputed. Using the closest valid value rather than the closest predicted value of a non-missing case ensures that only possible real values are imputed, and that the imputation process does not reduce the variance of imputed values relative to valid values. The imputation regression used for a particular wave also makes use of information about the value of the variable in other waves for some key components of household income, including gross usual pay from employment. The idea is that past or

future realisations of the variable are informative about the missing current value (in addition to current characteristics). The BHPS procedures use information from up to three waves – previous, current and next (BHPS Documentation Team 2009: A5–24 and Table 27). The result of this cross-wave imputation is that ‘[t]he imputed value should ... imply a rate of change drawn from a randomly selected similar case. This approach will avoid introducing spurious change for panel analysis, which would be likely to arise if only single wave imputation was used’ (Lynn 2006: 55).

These imputation procedures lead to non-missing values (and imputation flags) for individual level income variables. For total gross household income, there is also the problem of household members who refuse to complete the questionnaire altogether. For these refusers, income totals are imputed using the methods described above. Total gross household income can then be derived for every household.

The BHPS imputation procedures are relatively conventional, but not the only possibilities. Multiple imputation methods (Rubin 1987) have not been used, for example. Other panels use different approaches. For example, the German SOEP mainly uses the row-and-column method proposed by Little and Su (1989): see Frick and Grabke (2005).

The discussion so far has focused on the data that are in the main public-release BHPS files, made available to any bone fide researcher who is registered with the United Kingdom Data Archive (<http://www.data-archive.ac.uk>). The net household income variables are also made available in the same way, but have been created separately from the main release and on an ad hoc basis. I now turn to discuss their derivation in detail.

7. Derivation of the net household income variables

The BHPS net household income definition is modelled on the one used in the UK’s principal official source of information about the personal income distribution – the so-called *Households Below Average Income* (HBAI) series prepared by the Department for Work and Pensions (formerly the Department for Social Security). The HBAI publication provides detailed information about inequality and poverty using repeated cross-sectional data from the annual *Family Resources Survey*, a large specialist income survey. Some twenty editions have been produced to date; the most recent at the time of writing is Department for Work and Pensions (2010) covering the period between 1994/5 and 2008/09.

Since the early 1990s, my colleagues and I have derived net household variables to provide a longitudinal complement to the HBAI statistics. The first edition of our data was

for BHPS waves 1 and 2 (documented in the Appendix to Jarvis and Jenkins 1995, Appendix), and the latest covers waves 1–16 (Levy and Jenkins 2008). The computer code required to derive the variables is extensive. The structure is designed to be as modular as possible in order to facilitate updating to take account of changes in the taxes and benefit system but, inevitably, it remains complicated.

In the early 2000s, the Department for Work and Pensions began to create its own net household income variables from the BHPS, and to use them in its *Low Income Dynamics* (LID) statistics, first publishing these summaries along with the main HBAI statistics but more recently publishing them separately on the internet: the latest edition is Department for Work and Pensions (2009a). The definitions of net household income used in the LID and by me are broadly similar, but differ in matters of detail. The LID definitions are not documented but the principal difference from in net income definition appears to be that the LID one does not include a deduction for local taxes.

Improvements and corrections have been made to the BHPS net income variables at every edition (and documented in the materials accompanying their release). These changes have been applied to every wave of data retrospectively (where relevant) in each new edition of the files. For brevity, the discussion that follows here refers to the definitions used in the latest edition, with little mention of the changes introduced earlier.

The BHPS net income variable has three key features:

1. The reference period for the majority of income sources is the period round about the time of the interview, that is, it is a current rather than annual definition, with income converted pro rata to be expressed in terms of pounds per week. (Some comparisons between current and annual income are made later.)
2. The unit over which incomes are aggregated is the household (as defined in the previous section).
3. The sources of income and deductions from income that are included in the definition of net income are summarized in Table 2. (Non-cash income from other sources including imputed rent from owner-occupied housing, and capital gains, are not included.) *Gross* income is the sum of sources (a) to (g).

Table 2. The income sources included in net household income

	(a)	usual gross earnings from employment
+	(b)	earnings from subsidiary employment
+	(c)	profit or loss from self-employment
+	(d)	social security benefits and tax credits
+	(e)	private and occupational pensions
+	(f)	income from investments and saving
+	(g)	private transfers and other income
-	(h)	income tax (employees and self-employed)
-	(i)	National Insurance contributions (employees and self-employed)
-	(j)	contributions to occupational pension schemes
-	(k)	local taxes
=		Total net household income

Notes: The income definition refers to net income before the deduction of housing costs, i.e. net income ‘BHC’ in HBAI terminology. Gross income is the sum of sources (a) to (g).

These three components define an income variable that corresponds to the ‘before housing costs’ (BHC) net household income variable used by the Department for Work and Pensions in the HBAI statistics. (Details differ because the BHPS does not collect as much detailed information as the *Family Resources Survey*.) Post-calculation adjustments to account for differences in household size and composition using equivalence scales, and adjustments to constant-purchasing power terms using prices are summarized after discussion of the derivation of the nominal household net income variable.

The steps involved in constructing income components (a) to (k) are as follows:

1. Derive a measure of taxable income from employment and self-employment for each individual (components a–c in Table 2);
2. Estimate the income tax and National Insurance Contribution liabilities implied by this estimated taxable income, together with estimates of contributions to occupational pension schemes (components h–j);
3. Add on the sources of non-labour income (components d–g);
4. Estimate liabilities for local taxes (currently Council Tax; formerly the Community Charge).

Estimation of tax and National Insurance Contribution (NIC) liabilities is based on labour income only, reflecting the limited information available (see below) and, moreover, all such liabilities and also the deductions for occupational pension and local taxes are estimated rather than observed in the data. The use of simulation methods to estimate income deductions is common practice and employed by all UK tax-benefit microsimulation models,

but may lead to the introduction of measurement error. However, estimated liabilities for wave 1 and 2 respondents who provided both gross and net amounts are remarkably similar to the difference between gross and net labour income (Jarvis and Jenkins 1995: Tables A–9 and A–10). This suggests that the use of simulation does not lead to major problems. In any case, there are advantages of consistency in applying the same derivation procedure to all households. I now discuss steps 1–4 in more detail.

Income from employment and self-employment

The BHPS asks employees to report their gross and net (take home) pay at last payment, the time period it covered, and whether their last payment was equal to what they are usually paid. If last and usual pay differed, respondents are then asked to give their usual pay and to explain why the last amount was unusual. A majority of respondents provides both gross and net amounts. If possible, the interviewer checks a recent payslip and sees them in around one third of the cases. A small minority of employees either refuse to give information or do not know the amount or time period of their last earnings. The BHPS data include imputed values for these cases (see above), which are used. The survey also asks about earnings from second jobs, but this information is reported only as a gross figure.

Income from self-employment is difficult to measure in household surveys because the degree of non-response and under reporting tends to be higher for the self-employed than employees, and income from self-employment varies considerably over time making it difficult for respondents to assess their incomes and for researchers to derive a measure of ‘current income’ from the data provided. Both of these problems occur in the BHPS. The survey asks the self-employed to provide details of their most recent accounts or (where this is not available) an estimate of their usual monthly gross earnings. Approximately one fifth of self-employed respondents either refuse to give information or do not know how much they earn. The BHPS contains imputed values for these cases, and these are used as an estimate of gross earnings.

The data refer to the most recent period for which the respondent has either kept profit and loss accounts or has a record of his or her gross earnings. This information may be out of date by up to four years, and therefore underestimated. To correct for this, the incomes are updated to allow for inflation using the not-seasonally-adjusted Average Earnings Index (AEI) for the whole economy (Office for National Statistics series LNMM). Where earnings from self employment have been imputed in the BHPS, the modal reporting period from the

non-imputed cases is used, which is the financial year ending in the April before the interview.

Total gross earnings from all sources (employment, self employment, and income from second or occasional jobs) are computed using the most recent *usual* gross payment received. There are a small number of respondents who are not employed or self employed but who report income from occasional jobs in the month previous to the interview. It is assumed that this income is untaxed and net labour income is set equal to the gross amount reported for these cases.

Income tax

The first step in estimating income tax payments is derivation of each individual's taxable income. This is defined to be equal to gross income minus certain tax allowances and tax-deductible contributions to employer pension schemes. The rules have changed over time, and the calculations take account of this. Other minor tax allowances which can be set against income are ignored as there is insufficient information collected by the BHPS. Under independent taxation (introduced in 1990–1) each taxpayer is entitled to a personal allowance, the value of which is higher for those aged 65 or over. A married man can also claim a married couple's allowance in addition to his personal allowance. If his income is insufficient to make full use of this allowance then the unused part can be transferred to his wife. It was only from the tax year 1993–4 onwards that couples could choose to allocate the whole allowance to the wife or split it equally between them.

Each individual's tax allowance is estimated using demographic information on age and marital status reported in the BHPS. Data for husbands and wives are matched in order to be able to use information on spouses' age and earnings when calculating the married couples' allowance (MCA). It is assumed that any unused MCA is transferred from the husband to the wife. (As Sutherland and Wilson (1995) point out, this transfer does not happen automatically, but depends on decisions made by the couple concerned. In cases where the husband's income level is likely to increase in the near future (for example temporary unemployment) then the couple may decide not to transfer the allowance.) The procedure for computing the MCA was modified in the 10-wave release of the net income variables to take account of the fact that the part of the MCA that is age-related has to go to the husband and cannot be transferred to the wife. However, the husband can transfer to the wife the part that he is not able to use (and it is sensible to do so). In practice, the old and the

new procedure produced similar results, but the allocation between husband and wife is slightly different (in a few cases even the total MCA the couple is entitled to). From 2000–01, the married couple allowance for people born after 5 April 1935 was withdrawn. Hence, the general MCA no longer exists and the MCA for older people will progressively disappear.

Having deducted the appropriate tax allowances and pension contributions from gross income, tax paid is calculated by applying the schedule of tax rates for the relevant year. Net labour income is equal to gross earnings minus estimated tax, NICs and occupational pension contributions.

National Insurance Contributions (NICs)

Employees are liable to pay Class 1 NICs if they are aged sixteen or over and earn more than the ‘lower earnings limit’. All of the earnings of an employee who earns at least the lower earnings limit are subject to NICs up to the upper earnings limit. The rate of contribution is calculated as a percentage of gross earnings and depends on whether the employee is a full member of the State retirement pension scheme or whether their employer has contracted out of the earnings related part of the State scheme and provides a separate occupational pension. Employees in contracted-out employment pay NICs at a rate 2 per cent lower than the non-contracted-out rate on earnings between the upper and lower limits. Prior to 1977, married and some widowed women could elect to pay NICs at a reduced rate of 3.85 per cent. This rate is the same for both contracted-out and non-contracted-out employment. Those who chose to do so (before 1977) could then continue to pay reduced rate contributions thereafter. Administrative statistics (Department of Social Security 1994) show that approximately 10 per cent of women paid reduced rate contributions in 1991/2 and this can be expected to have fallen after that due to some of these women leaving the labour market.

Since the BHPS collects no information about NICs, they are estimated for employees using data on gross earnings and membership of occupational pension schemes. It is assumed that members of an employer’s scheme pay NICs at the lower contracted out rate, and that all others make full contributions. (Since there is no information to identify the women who opted to pay reduced rate contributions, it is assumed that all are paying at the non-reduced rate.) This may overstate the number of contracted-out employees by approximately 10 per cent: see the discussion in Jarvis and Jenkins (1995).

Self-employed people are liable for two types of NICs. Class 2 contributions are paid as a flat rate weekly amount with exemption given to those whose profits fall below a

specified amount. Class 4 contributions are calculated as percentage of annual taxable profits between an upper and lower earnings limit. Half of Class 4 contributions can be offset against income tax. NICs for self-employed people are estimated using data about their most recent gross earnings or profit. There are insufficient data available to estimate lump sum tax or NIC payments or refunds, and so these factors are ignored.

Occupational pension contributions

For the respondents who report making contributions to their employer's pension scheme, pension contributions are imputed at a rate equal to 4.7 per cent of gross earnings. This figure is the average of the figure reported by respondents to the *Family Expenditure Survey* in survey years 1991, 1992, and 1993. The average rate for 2003/04 according to *Family Resources Survey* data is not much different (4.9 per cent). The earlier figure has been retained simply for consistency. Clearly these estimates are an approximation of reality, but it was thought that the benefits of attempting to derive of a more accurate individual-specific amount (for example using occupation-specific data) were not justified by the time required.

Social security benefits and tax credits

Using respondents' retrospective recall at the interview, the BHPS collects detailed information on the type of social security benefit received by each member of the household on a month by month basis for the whole of the period from September of the previous year to the date of the interview. The survey also asks about the amount of the last payment of each benefit.

To construct the net income variable, the BHPS derived variable which measures the total benefit income of the household in the month before the interview (and therefore includes imputed values) is used, with one important caveat concerning the housing benefit component. A change in the wording between waves 1 and 2 of the prompt card used by BHPS interviewers to remind respondents of their various sources of income appears to have led to a large drop in the number of people reporting housing benefit receipt. Whereas in wave 1 the card referred to 'Housing Benefit (Rent Rebate and Allowances)', in wave 2 it referred to 'Housing Benefit paid directly to you'. This appears to have led some individuals who do not receive their housing benefit directly to fail to report it in wave 2. To deal with this discontinuity, an alternative measure of housing benefit is created using information from

the household questionnaire (following Webb 1995). Households are asked to report their rent as both a gross and net amount, the latter taking into account housing benefits received. For households reporting a 100% rent rebate housing benefit is set equal to gross rent. For other households, the estimate of housing benefit is equal to gross rent minus net rent.

There have been revisions to the details of our calculations in different editions of the net income files, the most significant of which was in the most recent (16-wave) edition. There was a coding error, now corrected, which meant that, for households reporting a 100 per cent rent rebate, housing benefit was set equal to zero (rather than gross rent). The correction increases the income of low-income households and reduces measured inequality – though only for waves 2–8 is the effect particularly marked (Levy and Jenkins 2008).

The other notable modifications to the calculations have been related to the introduction of and reforms to tax credits. In October 1999, Working Families Tax Credit (WFTC) replaced the Family Credit and Disabled Person's Tax Credit (DPTC) replaced Disability Working Allowance. In April 2003, WFTC and DPTC were replaced by Child Tax Credit (CTC) and Working Tax Credit (WTC). WFTC, DPTC and WTC can be delivered in two ways: through the employer in the pay slip or as a benefit. It appears that the pay slip amounts are usually recorded also in the BHPS's income grid, but not always. Also, in a few cases there are discrepancies between the two sections. As a rule, calculations use the amount that is recorded in the income section, ignoring the discrepancies between this amount and the one reported in the employment section. When no amount is recorded in the income section (and only in this case), but some positive amount is recorded in the employment section, the household's benefit income is assumed to include the amount recorded in the employment section.

The annual amount of Tax Credit received through the employer is computed as follows. First, the weekly amount is computed using the amount and pay period variables. Second, the annual amount is computed using the reported number of weeks worked in the relevant year – therefore assuming that those who are currently receiving the Tax Credit through their employer have received the same amount in every week they were in work during the relevant year. Finally, the amount computed in this way (and summed across individuals within each household to derive a household figure) is set equal to zero for those households where somebody declared receipt of the Tax Credit as a benefit, in order to avoid double counting.

Income from investments and savings

Obtaining reliable information about income from investments and savings is notoriously difficult in household surveys. In the first few BHPS waves of the BHPS, the questions asked were not very detailed. At waves 1 and 2, only banded responses were sought, using four categories. At waves 3–8, the top band was split into three. From wave 9 onwards, respondents have been asked for an exact amount and banded amounts sought from those who do not provide one. For those with non-zero amounts or non-exact amounts, BHPS staff impute a value from the banded responses (common to all reporting the same band), and those imputations are also used in the net income calculations. The use of banding is likely to result in an understatement of income from this source, particularly for those with very high incomes. Unlike as for other income sources, the BHPS asks respondents about amounts received in the past 12 months, and this annual figure is then converted pro-rata to a monthly or weekly amount as required.

Transfers and other income

This category includes educational grants, maintenance and alimony payments, foster allowances, payments from family members not living in the household and any other payments received by household members. The derived BHPS variable, which gives the household total for these income sources, is used for the net income calculation.


Local taxes

Local tax payments are estimated for all households using external data on average Council Tax levels by local authority. For waves 1 and 2 when the community charge ('poll tax') was in operation rather than Council Tax, data on the average community charge payment in each local authority district is used. For waves 3–6, council tax payments are imputed using data on the average council tax payment in each local authority area. (See Redmond 1997 for details.) From wave 7 onwards, the BHPS has collected data on the council tax band of households, and these are used to estimate more precise council tax liabilities in conjunction with information identifying the local authority in which the household resides.

Partial unit non-response: non-response by some household members

As mentioned earlier, the measure of net household income cannot be derived if one or more of the adult members of the household are not interviewed – the problem of partial unit non-response. Figure 2 summarizes its prevalence. The proportion of all households at each wave with all adults providing an interview fluctuated between 80 per cent and 90 per cent at the beginning of the panel but, at the most recent waves, the fraction is around 80 per cent. The increase over time is mostly accounted for by a growing fraction of individuals who refuse to provide an interview and, at more recent waves, by a small but growing proportion who provide only telephone interviews.

Figure 2. Full and partial response within BHPS households, rates (%) by wave



Note: The chart shows breakdowns of the BHPS variable wIVFHO (not available at wave 1).
Source: Author’s calculations from BHPS Documentation Team (2009).

There are a number of ways of accounting for the impact of partial unit non-response. They range from doing nothing (the approach adopted in the construction of the BHPS net income variables) through to use of methods employing imputation of the missing data or re-weighting of the sample. Observe that partial unit non-response cannot occur within one-person households and the risk of occurrence is greater the larger the household.

The ECHP used a form of imputation in which there was a ‘flat correction’ method that re-scaled household incomes for households with partial unit non-response (Eurostat 2000). A range of methods are applied to German SOEP data and compared by Frick, Grabka, and Groh-Samberg (2009). Their results show that the importance of the issue depends on the context. For example, cross-sectional poverty rates estimated using the

methods are very similar in any given year and follow similar trends over time (Figure 7). Income mobility, including the probability of moving into or out of poverty, is greater for households with partial unit non-response but arguably the differences in estimates from applying different methods is not large (Figures 9–11). This may be because the prevalence of partial unit non-response, while non-trivial, is not large. In the following section, I show indirectly that partial unit non-response appears not to be a critical problem in BHPS net income data since estimated distributions match counterparts in the benchmark HBAI distributions remarkably well.

8. Equivalence scales and price indices

In order to compare incomes for households of different size and composition, and to make comparisons of incomes in different years, each net household income value should be adjusted by an equivalence scale factor and by a monthly price index. These adjustments are standard ones that are commonly-used.

The BHPS net household income data releases includes the two equivalence scales that are used in Britain's official income statistics, *Households Below Average Income* (HBAI). These are the 'McClements' scale and the 'modified-OECD' scale (the 'before housing costs' versions in each case). In the late 1990s, Eurostat began employing the modified-OECD scale for its cross-national comparisons of income distributions and the UK has now adopted it for its headline statistics as well, switching over from the McClements scale. The modified-OECD scale distinguishes between individuals aged between 0 and 14 years, and those aged 15+ ('adult'). The scale equals one for the first adult and adds a weight of 0.5 for each additional 'adult' and a weight of 0.3 for each child. In addition, the Department for Work and Pensions normalizes scale values so that the (normalized) scale rate for a childless couple household is equal to 1.0 (rather than 1.5). This has the convenience of aligning the scale with the McClements scale, which is normalized to equal 1.0 for a childless married couple household. The weights used to construct the McClements scale take account of differences in household size and composition in finer detail (for example with different weights for children of different ages or additional adults beyond the first two). However, the relativities for different household types are quite similar for the majority of households according to both the McClements and modified-OECD scales. For further details of the scales, see for example Department for Work and Pensions (2009*b*: Appendix 2). Coulter, Cowell, and Jenkins (1992*a*, 1992*b*) and Jenkins and Cowell (1994)

compare the McClements scale with other scales using parametric approximations, and their analysis also suggests that the choice between these two scales is not a major issue.

The price index used to convert household net income values from different time periods to constant price terms are the same as used in the official income statistics. The index is the ‘all-items Retail Price Index excluding Council tax’, created by the Department for Work and Pensions by the Office for National Statistics (ONS). Values of the index are reproduced in the appendices to Levy and Jenkins (2008). The index is a monthly one (not seasonally adjusted), and matched to respondents using data about the interview month in each survey year. As in the official statistics, no account is taken of potentially different inflation rates between different groups (such as low-income versus high-income households, or between young and old, or between different regions of Britain). On these issues, see for example Crawford and Smith (2002).

9. Current versus annual measures of income

The definition of net household income is essentially a measure of current income because it is mostly derived from respondents’ reports about income received round about the time of the survey interview – as virtually all UK survey measures of income are. It is not a definition of annual income, as used by surveys for most other countries. As discussed by Jenkins (2011: Chapter 2), use of a current income definition might be expected to produce estimates of inequality, poverty, and mobility that are larger than those derived using an annual income measure, other things being equal. In addition, the differences in types of measure may compromise comparisons between patterns for Britain and those of other countries.

Böheim and Jenkins (2006) show, however, that estimates of cross-sectional and longitudinal income distribution summary statistics derived from BHPS measures of current and annual income are remarkably similar. Almost all differences between corresponding estimates for the two measures are in the expected direction, but the magnitude of the differences is small both in aggregate and also when looking at breakdowns by family type and employment status. Although our published paper is based on comparisons of gross income measures rather than net income ones, our unpublished work shows that similar results apply in this case too (2006: n. 8). The results suggest that, for practical purposes, the distinction between current and annual income measures is a minor one.

The reasons underlying this result mainly hinge on the fact that in Britain survey measures of income are rarely purely current or purely annual in the sense of every

constituent component having a current or annual reference period. Böheim and Jenkins (2006) emphasise several specific factors related to this in the BHPS context. First, the measure of employment earnings included in household income refers to usual pay, and not the amount most recently received. Second, some other income sources use a reference period that may often be as long as a year. Self-employed workers who keep accounts report income (net profit or loss) over a year, and the BHPS question about income from investments and savings specifically refers to receipts over the previous twelve months. At the same time, third, the BHPS annual income definition is not derived from reports of annual receipts for every source. Instead, it is a measure constructed using information about incomes received at the current interview and at the previous interview, combined with information from retrospective monthly histories of employment and benefit receipt, and information from external sources such as administrative statistics, in order to build up a picture of incomes received between interviews. For each source, this information yields a series of monthly income estimates that are summed to produce an annual aggregate (see Böheim and Jenkins 2006 for more details). Total income is derived by summing the annual receipts from each income source.

An additional reason for the minor differences between current and annual income estimates is also investigated by Böheim and Jenkins (2006), namely that the numbers of people moving into or out of jobs, or experiencing changes in the demographic composition of their household, are relatively small and hence consequential within-year income variability is relatively small. To examine this hypothesis further, Böheim and Jenkins (2006) analyze whether differences between statistics based on current and annual income measures are larger for households which experience changes in labour market attachment or changes in household composition – but find that there was no conclusive evidence one way or the other. Differences are relatively small for most subgroups considered.

I conclude from this research that the distinction between the BHPS measures of current and annual income is unimportant relative to other issues.

10. Comparisons of BHPS and HBAI net income distributions from a cross-sectional perspective

In this section, I compare BHPS net income distributions with their HBAI counterparts from a cross-sectional perspective, drawing on the more detailed comparisons provided by Jenkins (2010). Ideally, one would like to have a longitudinal benchmark data set as well but none is

currently available for Britain. The cross-sectional comparisons are an important validation exercise, none the less, since getting the cross-sectional estimates for each year right is an essential part of getting right the estimates of the joint distribution for any pair of years. The HBAI distributions are taken as the reference point because they are derived from the Family Resources Survey, a specialist income survey that has a sample size almost six times larger than the BHPS samples used here (almost 30,000 households per year compared to about 5,000). The HBAI net income data are used to generate Britain's official income distribution statistics and are regarded as being of high quality.

The comparisons reported in this section are based on calculations using the 16-wave release of the BHPS net income files (Levy and Jenkins 2008) and the 2009 release of the HBAI files (Department for Work and Pensions 2009*b*). The latter cover financial years 1994/95 to 2007/08 and are the latest available at the time of writing. In both sources, the variable of primary interest is net household income (before the deduction of housing costs), as described earlier, equivalized using the modified-OECD scale (with the HBAI variable re-normalized so that the scale rate equals one for a single-person household, as in the BHPS case), and expressed in January 2008 prices in pounds per week using the same monthly price index (see above). The BHPS calculations are based on the Original sample only, and households from Northern Ireland have been dropped from the HBAI files, so that comparisons refer to Britain in both sources. All calculations for each series used the relevant sample weights.

An important difference between the HBAI and BHPS net income distributions is that the former include an 'SPI adjustment' in order to better measure incomes at the very top of the distribution. For each FRS year, the Department for Work and Pensions identifies a small number of rich households defined as being households containing a rich individual, with pensioner and non-pensioner households considered separately. The threshold defining 'rich' is set at a level above which it is considered that incomes are not measured reliably in the FRS because the sample size is too small.


'Year' refers to survey year in the BHPS (the modal interview month is October), and to financial year in the HBAI (interviews spread from April to following March). Because of the secular growth in incomes on average over the period, the financial year coverage of the HBAI may lead to lower incomes in the HBAI than the BHPS, *ceteris paribus*. However, the impact of this is likely to be relatively small.

Estimates of selected quantiles of the income distribution are shown for each source, by year, in Figure 3. Panel (a) summarizes differences in the top half of the distribution,


showing the median ($p50$) and the 75th, 95th and 99th percentiles ($p75$, $p95$, $p99$). Panel (b) refers to the bottom half of the distribution, showing the median ($p50$) and the 25th, 10th and 1st percentiles ($p25$, $p10$, $p1$). There is a remarkably close correspondence between corresponding estimates from each source, with the notable exception of the very top of the distribution ($p99$), and also at the bottom of the distribution ($p1$).

Figure 3. Selected quantiles of net household income, BHPS and HBAI, by year

(a) Median and higher



(b) Median and lower




Source: Jenkins (2010).

The difference at the very top is readily explained by the use of the SPI adjustment in the HBAI data. (Estimates from unadjusted FRS data would be closer to the BHPS ones.) The SPI adjustment is also likely to explain why the $p99$ series for the two sources were relatively close in the mid-1990s but diverged thereafter. Atkinson, Piketty, and Saez (2010) report trends in the share of total income held by the richest 1% in Britain for almost a century through to 2005, with their estimates derived from administrative record tax data on incomes, including SPI data for the most recent years. Their estimates (2010: Figure 7A) show that although the share was rising throughout the period 1990–2005, there was a step change upward round about 1995. Atkinson (2005) reports similar trends for income shares within the top 1%. Although the definition of income and the income recipient in the tax data are not exactly the same as those employed here, the trends at the very top of the distribution are likely to explain what is shown for $p99$ in Figure 3(a).

Accurate measurement of very low incomes using household surveys is also a problem, and is reflected in the estimates from both sources of $p1$. For further evidence for Britain about this issue, see Brewer *et al.* (2009b). There is greater year-on-year fluctuation in the series for $p1$ compared to other percentiles, and more so for the BHPS (with the smaller sample size). Overall, the estimates presented in Figure 3 suggest that BHPS estimates of the net income distribution are relatively good, except at the very top and very bottom of the distribution.

What about summary statistics such as poverty rates and inequality indices? The similarities in estimates of quantiles throughout most of the income range means that estimates of the proportion of persons with an income below 60 per cent of the median (Britain's headline poverty rate) are close for the two sources. This is shown in Figure 4. Even when the BHPS and HBAI series differ most (during the 1990s), the difference is at most about one percentage point.


Figure 4. BHPS and HBAI estimates of the percentage of individuals with a net household income less than 60% of the median, by year



Source: Jenkins (2010).

When one looks at inequality using a portfolio of commonly-used indices, differences between the sources are more apparent: see Figure 5. Inequality is higher according to the HBAI series, particularly reflecting differences at the top of the distribution and the SPI adjustment, with the divergence beginning in the second half of the 1990s – as discussed earlier. Consistent with this, the differences between the series are greatest for the GE(2) inequality index which, of the indices considered, is the most sensitive to income differences at the top of the distribution. The differences between series are smallest (and trends are most similar) for the p90/p10 percentile ratio measure, which is not affected at all by incomes above the 90th percentile or below the 10th percentile.

Figure 5. BHPS and HBAI estimates of inequality, by index and year

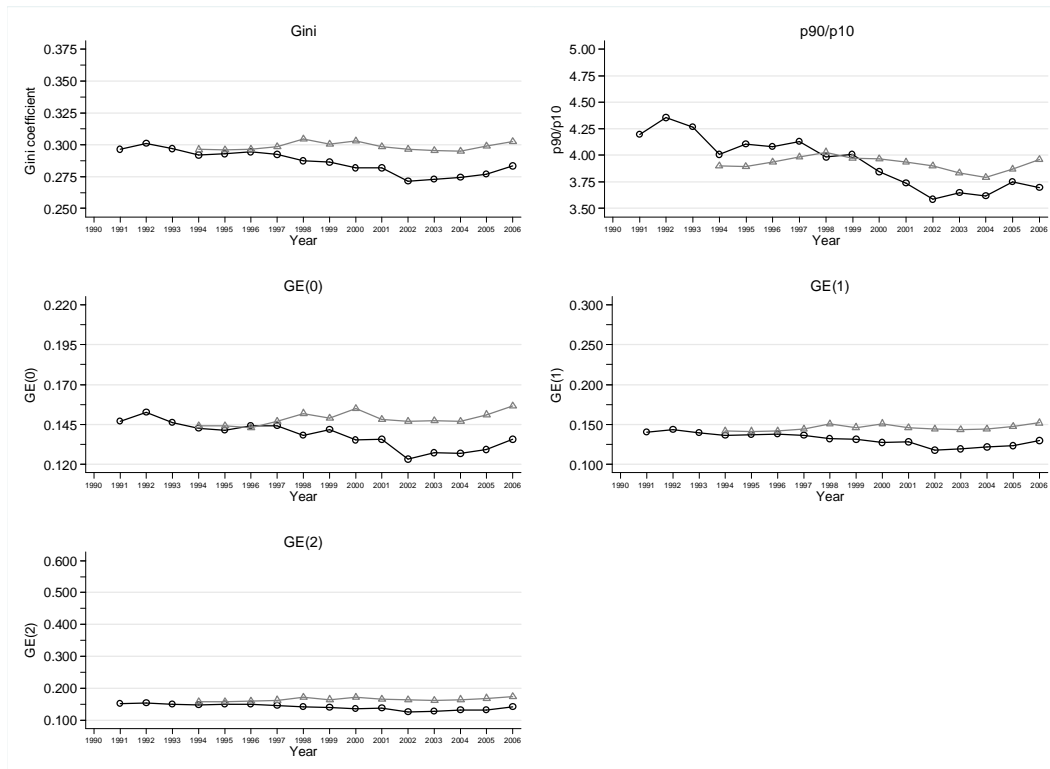


Notes: Within each chart, BHPS estimates are shown using round markers (and black lines) and the HBAI estimates using triangle markers (and grey lines). The inequality indices are the Gini coefficient (Gini), the ratio of the 90th percentile to the 10th percentile (p90/p10), mean logarithmic deviation (GE(0)), Theil index (GE(1)), and half the coefficient of variation squared (GE(2)).

Source: Jenkins (2010).

Reflecting the problems of securing reliable measurement of incomes at the very top and the very bottom of the income distribution, and the potential lack of robustness of summary measures to outlier values at the top and the bottom of the distribution, it is often suggested that income data should be trimmed prior to analysis. Figure 6 shows what happens to the inequality estimates if this suggestion is implemented. Specifically, the bottom 1% and top 1% of the distribution for each year and source are dropped prior to calculations of each index.

Figure 6. BHPS and HBAI estimates of inequality (trimmed distributions), by index and year



Notes: As for Figure 4.5, except that each inequality estimate is derived using a distribution from which the richest 1% and the poorest 1% of observations have been dropped.

The result is that each BHPS index series is now much closer to its HBAI counterpart. There is a suggestion of a slight decrease in inequality up to around 2002 according to the BHPS series but not the HBAI ones. But both series suggest that inequality increased slightly after 2002 according to all indices. The trimming removes SPI-adjusted observations from the HBAI distributions, and note the impact on estimates of top-sensitive GE(2) in particular. The differences remaining between the series arise from the combination of relatively small differences throughout the income range, above and below the median (see Figure 3).

Of course, comparisons of the two sources need to consider potential differences at the level of population subgroups and specific income sources, not only for differences in the distribution of total net income among the population as a whole in a given year. In Jenkins (2010), I show that BHPS estimates of the proportions of individuals in different family types or in different groups defined by the economic status of the family to which they belong are remarkably similar to those derived from HBAI data, using definitions of the subgroups that are the same as those employed by the Department for Work and Pensions (2009b).

Some differences between series appear when the focus is more detailed, however. For example, I also consider the subgroup composition of the poorest fifth and of the richest

fifth of the distribution in each year according to the two sources (Jenkins 2010). The family type breakdowns suggest that BHPS produces an over-estimate of the proportion of single pensioners, especially women, in the poorest fifth of the distribution throughout the period as a whole. (For example, in 2006, the proportion of the poorest fifth who are female single pensioners is 14 per cent compared with 9 per cent in the HBAI.) Consistent with this, the economic status breakdowns indicate that the BHPS over-estimates the proportion of individuals in workless families containing a head or spouse aged 60+ in the poorest fifth. Estimates of the subgroup composition of the richest fifth are relatively close in both sources.

My comparisons of the composition of net household income distinguish between six sources for which comparable definitions in both sources are possible: income from employment, income from self-employment, benefits and tax credits (including the state retirement pension), income from investments and savings, other income (for example transfers from private individuals outside the household), plus payments of income and local taxes and national insurance contributions (deductions which are 'negative' income). For this analysis, income is not equivalized.

My breakdowns (Jenkins 2010) suggest that the BHPS under-records labour income relative to the HBAI. For example, for 2006, average household income from employment is £487 per week according to the BHPS, but £526 per week according to the HBAI. Deductions are also under-estimated, which most likely reflects the fewer details available in the BHPS to estimate them relative to the FRS/HBAI. For example, for 2006, deductions are £202 in the HBAI, but £166 in the BHPS. These two features offset each other, so that net household income totals are quite similar across sources (£551 in the BHPS and £571 in the HBAI in 2006, a difference of 4 per cent). Differences between sources do not appear as large if one summarizes income composition in terms of shares of the total. For instance, for 2006, the share of employment income in total net household income is 88 per cent according to the BHPS and 92 per cent according to the HBAI, and the shares of deductions are -30 per cent and -35 per cent, respectively. For other income sources, differences in corresponding shares of the total across sources are always less than two percentage points. When the focus is on income packaging for the poorest fifth or for richest fifths of the distribution, there are similar off-setting patterns. And for both the distribution as a whole, and for the richest and poorest fifths, there is no apparent change in these patterns over the period 1994–2006.

Earlier sections raised the possibility that changes in the BHPS survey design may have affected income distribution estimates. Reassuringly, my analysis in Jenkins (2010) reveals no obvious discontinuities in the BHPS series associated with either the introduction

of CAPI (wave 9, survey year 1999) or of dependent interviewing (wave 16, survey year 2006).

Overall, use of the BHPS net income distributions data as a longitudinal complement to the HBAI appears to be valid, especially if the focus is not on the very poorest or the very richest incomes.

References

- Atkinson, A. B. (2005). 'Top Incomes in the UK over the 20th Century', *Journal of the Royal Statistical Society, Series A*, 168: 325–43.
- Atkinson, A. B., Piketty, T., and Saez, E. (2010). 'Top Incomes in the Long Run of History'. NBER Working Paper 15408. Cambridge MA: National Bureau of Economic Research. *Journal of Economic Literature*, forthcoming.
- BHPS Documentation Team (2009). BHPS Documentation (On-Line Materials). Colchester: Institute for Social and Economic Research, University of Essex. <http://www.iser.essex.ac.uk/survey/bhps/documentation>
- Böheim, R, and Jenkins, S. P. (2006). 'A Comparison of Current and Annual Measures of Income in the British Household Panel Survey', *Journal of Official Statistics*, 22: 733–58.
- Brewer, M., O'Dea, C., Paull, G., and Sibieta, L. (2009). *The Living Standards of Families with Children Reporting Low Incomes*. Research Report 577. London, Department for Work and Pensions. <http://research.dwp.gov.uk/asd/asd5/rports2009-2010/rrep577.pdf>
- Cappellari, L. and Jenkins, S.P. (2004). 'Modelling Low Income Transitions', *Journal of Applied Econometrics*, 19: 593–610.
- Cappellari, L. and Jenkins, S. P. (2008). 'Estimating Low Pay Transition Probabilities Accounting for Endogenous Selection Mechanisms', *Journal of the Royal Statistical Society, Series C (Applied Statistics)*, 57: 165–86.
- Coulter, F. A.C., Cowell, F. A., and Jenkins, S. P. (1992a). 'Differences in Needs and Assessment of Income Distributions', *Bulletin of Economic Research*, 44: 77–124.
- Coulter, F. A.C., Cowell, F. A., and Jenkins, S. P. (1992b). 'Equivalence Scale Relativities and the Extent of Inequality and Poverty', *Economic Journal*, 102: 1067–82.
- Crawford, I. and Smith, Z. (2002). *Distributional Aspects of Inflation*. Commentary 90. London: Institute for Fiscal Studies. <http://www.ifs.org.uk/comms/comm90.pdf>
- Department of Social Security (1994). *Social Security Statistics*. London: HMSO.

- Department for Work and Pensions (2008). *Households Below Average Income. An Analysis of the Income Distribution 1994/95 – 2006/07*. London: Department for Work and Pensions. http://statistics.dwp.gov.uk/asd/index.php?page=hbai_arc#hbai
- Department for Work and Pensions (2009a). *Low Income Dynamics. 1991–2007 (Great Britain)*. London: Department for Work and Pensions. http://statistics.dwp.gov.uk/asd/index.php?page=hbai_arc#low_income
- Department for Work and Pensions (2009b). *Households Below Average Income. An Analysis of the Income Distribution 1994/95 – 2007/08*. London: Department for Work and Pensions. http://statistics.dwp.gov.uk/asd/index.php?page=hbai_arc#hbai
- Department for Work and Pensions (2010). *Households Below Average Income. An Analysis of the Income Distribution 1994/5–2008/09*. London: Department for Work and Pensions. http://statistics.dwp.gov.uk/asd/index.php?page=hbai_arc#hbai
- Eurostat (2000). ‘Imputation of Income in the ECHP. DocPan 164/00. Luxembourg: Eurostat.
- Frick, J. R. and Grabka, M. M. (2005). ‘Item nonresponse on income questions in panel surveys: Incidence, imputation and the impact on inequality and mobility’, *Allgemeines Statistisches Archiv*, 89: 49–61.
- Frick, J. R., Jenkins, S. P., Lillard, D. R., Lipps, O., and Wooden, M. (2007). ‘The Cross-National Equivalent File (CNEF) and its Member Country Household Panel Studies’, *Schmollers Jahrbuch – Journal of Applied Social Sciences Studies*, 127: 627–54.
- Frick, J. R., Grabka, M. M., and Groh-Samberg, O. (2009). ‘Dealing with Incomplete Household Panel Data in Inequality Research’. Unpublished Paper. Berlin: DIW Berlin.
- Gottschalk, P. and Danziger, S. (2005). ‘Inequality of Wage Rates, Earnings and Family Income in the United States, 1975–2002’, *Review of Income and Wealth*, 51: 231–54.
- Gottschalk, G. and Huynh, M. (2010). ‘Are Earnings Inequality and Mobility Overstated? The Impact of Non-Classical Measurement Error’, *Review of Economics and Statistics*, 92: 302–15.
- Gottschalk, P. and Moffitt, R. (1994a). ‘Welfare Dependence: Concepts, Measures, and Trends’, *American Economic Review (Papers and Proceedings)*, 84: 38–42.
- Gottschalk, P. and Moffitt, R.A. (1994b). ‘The Growth of Earnings Instability in the U.S. Labor Market’, *Brookings Papers on Economic Activity*, 2-1994: 217–72.
- Gottschalk, P. and Moffitt, R.A. (2007). ‘Trends in Earnings Volatility in the U.S: 1970–2002’. Paper presented to the Annual Meetings of the American Economic Association, Chicago, January 2007.

- Gottschalk, P. and Moffitt, R.A. (2009). 'The Rising Instability of U.S. Earnings', *Journal of Economic Perspectives*, 23: 3–24.
- Gottschalk, P. and Smeeding, T.M. (2000). 'Empirical Evidence on Income Inequality in Industrialized Countries', in A. B. Atkinson and F. Bourguignon (eds), *Handbook of Income Distribution*. Amsterdam: North-Holland, 261–307.
- Hill, M. S. (1992). *The Panel Study of Income Dynamics: a User's Guide*. Newbury Park CA: Sage Publications.
- Jarvis, S. and Jenkins, S. P. (1995). *Do The Poor Stay Poor? New Evidence about Income Dynamics from the British Household Panel Survey*. Occasional Paper No. 95–2. Colchester: Institute for Social and Economic Research, University of Essex. http://www.iser.essex.ac.uk/files/occasional_papers/pdf/op95-2_text.pdf, http://www.iser.essex.ac.uk/files/occasional_papers/pdf/op95-2_appen.pdf
- Jenkins, S.P. (2009). 'Marital splits and income changes over the longer term', in M. Brynin and J. F. Ermisch (eds), *Changing Relationships*. London: Routledge, 217–36.
- Jenkins, S. P. (2010). 'Comparisons of BHPS and HBAI distributions of net household income'. Unpublished Paper. Colchester: Institute for Social and Economic Research, University of Essex. <https://filestore.iser.essex.ac.uk/d/3a16c2>
- Jenkins, S. P. (2010). 'How Can We Improve our Knowledge of Poverty Dynamics in Europe?'. Paper presented at the EU Presidency 'Science Against Poverty' Conference, Segovia, 8–9 April 2010. <http://www.iser.essex.ac.uk/assets/327>
- Jenkins, S. P. (2011). *Changing Fortunes: Income Mobility and Poverty Dynamics in Britain*. Oxford, Oxford University Press, forthcoming.
- Kroh, M. (2009). 'Documentation of Sample Sizes and Panel Attrition in the German Socio Economic Panel (SOEP) (1984 until 2008)'. SOEP Data Documentation Paper 47. Berlin: SOEP Group, DIW Berlin. http://www.diw.de/documents/publikationen/73/diw_01.c.341747.de/diw_datadoc_2009-047.pdf
- Levy, H. and Jenkins S. P. (2008). 'Documentation for Derived Current and Annual Net Household Income Variables, BHPS Waves 1–16'. Unpublished Paper. Colchester: Institute for Social and Economic Research, University of Essex. <http://www.data-archive.ac.uk/doc/3909/mrdoc/pdf/3909userguide.pdf>
- Little, R. J. A. and Su, H. -L. (1989). 'Item Nonresponse in Panel Surveys', in D. Kasprzyk, G. Duncan, G. Kalton, and M. P. Singh (eds) *Panel Surveys*. New York: John Wiley and Sons, 400–25.

- Lynn, P. (ed.) (2006). 'Quality Profile: British Household Panel Survey Version 2.0: Waves 1 to 13: 1991–2003'. Unpublished Paper. Colchester: Institute for Social and Economic Research, University of Essex. <http://www.iser.essex.ac.uk/files/bhps/quality-profiles/BHPS-QP-01-03-06-v2.pdf>
- National Equality Panel (J. Hills, Chair). (2010). *An Anatomy of Economic Inequality in the UK. Report of the National Equality Panel*. London: Government Inequalities Office. <http://www.equalities.gov.uk/pdf/NEP%20Report%20bookmarkedfinal.pdf>
- Nicoletti, C. and Perrachi, F. (2005). 'Survey Response and Survey Characteristics: Microlevel Evidence from the European Community Household Panel'. *Journal of the Royal Statistical Society, Series A*, 168: 763–81.
- Redmond, G. (1997), 'Imputing Council Tax Bands for Households in the British Household Panel Study'. Working Paper 97–10. Colchester: Institute for Social and Economic Research, University of Essex.
- Rubin, D. B. (1987). *Multiple Imputation for Non-Response in Surveys*. New York: John Wiley and Sons.
- Sutherland, H. and Wilson, M. (1995). 'Using the FES to Simulate Income Tax in POLIMOD'. Microsimulation Unit Research Note MU/RN/15. Cambridge: Department of Applied Economics, University of Cambridge.
- Uhlig, S. C. N. (2008). 'The Nature and Causes of Attrition in the British Household Panel Study'. ISER Working Paper 2008–05. Colchester: Institute for Social and Economic Research, University of Essex. <http://www.iser.essex.ac.uk/publications/working-papers/iser/2008-05.pdf>
- Watson, N. and Wooden, M. (2006). 'Modelling Longitudinal Survey Response: the Experience of the HILDA Survey'. HILDA Discussion Paper series No. 2/06. Melbourne: Melbourne Institute, University of Melbourne. <http://www.melbourneinstitute.com/hilda/hdps/HILDA%20DP%202.06.pdf>
- Webb, S. (1995). *Poverty Dynamics in Great Britain: Preliminary Analysis from the British Household Panel Survey*. IFS Commentary No. 48. London: Institute for Fiscal Studies. <http://www.ifs.org.uk/comms/comm48.pdf>
- Winship, C. and Radbill, L. (1994). 'Sampling Weights and Regression Analysis'. *Sociological Methods and Research*, 23: 230–57.
- Wooldridge, J. M. (2002). 'Inverse Probability Weighted M-Estimators for Sample Selection, Attrition, and Stratification', *Portugese Economic Journal*, 1: 117–39.