ISER Working Paper Series

Heaping and leaping: Survey response behaviour and the dynamics of self-reported consumption expenditure

Stephen Pudney
Institute for Social and Economic Research
University of Essex

No. 2008-09
February 2008

INSTITUTE FOR SOCIAL
& ECONOMIC RESEARCH

E·S·R·C
ECONOMIC
& SOCIAL
RESEARCH
COUNCIL

www.iser.essex.ac.uk

Non-technical summary

Survey-based household consumption data tell us a great deal about people's standards of living and the way they make spending and saving decisions. Panel data, involving repeated interviews with the same households over a period of time, can also tell us how economic welfare and decision-making respond to changing circumstances. The drawback of consumption data is that it is difficult to collect - ideally, one would like to ask survey respondents to keep very detailed records of all their purchases over a representative period, but this is rarely feasible in practice. Instead, large-scale panel surveys like the British Household Panel Survey (BHPS) mostly use simple 'stylised' questions which ask respondents to estimate their spending on some category of consumer goods over a given recall period, without keeping detailed expenditure records. The aim of this paper is to understand the ways that survey respondents go about answering these stylised consumption questions and the consequent inaccuracies in survey data.

We concentrate on the responses to BHPS questions about spending on domestic energy (electricity, gas, etc.). We find that a large majority of respondents use some form of 'rounding' strategy, which results in 'heaped' data with large numbers of responses at particular expenditure levels. A variety of rounding strategies can be detected in the data - for example, some respondents appear to choose a round number for weekly spending and then scale that up to an annual total, while others use rounding at the monthly or annual level, or no rounding at all. Many households change their response behaviour from year to year, possibly distorting measurements of change over time in living standards and consumption behaviour. 'Standard' methods of analysing consumption data ignore survey response behaviour and are potentially unreliable. This paper develops a new statistical method of analysing consumption behaviour over time, taking account of survey respondents' tendency to round their answers and change their method of rounding idiosyncratically over time. The results obtained from this approach differ significantly from those generated by standard methods, suggesting that the measurement error problem is a serious one. The general conclusion of this study is that, when analysing consumption data from household panel surveys like the BHPS, researchers need to think in terms of two behavioural processes, not one: (1) how do people decide what to spend? (2) how do they then decide how to answer survey questions about their spending?

Heaping and leaping:

Survey response behaviour and the dynamics of self-reported consumption expenditure

Stephen Pudney

Institute for Social and Economic Research
University of Essex

June 2007

Abstract

Survey respondents often use simple strategies to answer retrospective questions about their level of consumption expenditure, resulting in the heaping of data at certain round numbers. In the panel context, wave-to-wave 'leaping' from one 'heap' to another can distort the sample pattern of consumption movements over time. Using BHPS energy expenditure data, we show that respondents use a variety of different response strategies, often switching between those strategies. We estimate a joint model of the dynamics of consumption and response behaviour and find that neglect of response behaviour can lead to serious biases in empirical studies of consumption dynamics.

Keywords: measurement error, heaping, consumption, energy, BHPS

JEL codes: C81, D12, I32, Q41

Contact: Steve Pudney, ISER, University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, UK; tel. +44(0)1206-873789; email spudney@essex.ac.uk

This work was supported by the Economic and Social Research Council through the MiSoC and ULSC Centres (award nos. RES518285001 and H562255004). I am grateful to Heather Laurie for her help and advice and to seminar participants at the Institute for Fiscal Studies, the Melbourne Institute of Applied Economic and Social Research and ISER.

1 Introduction

Consumption expenditure data from household surveys provide the basis for important research in several different fields. For example, the large literature on the dynamics of consumption behaviour rests heavily on analysis of data on food consumption from the US Panel Study of Income Dynamics (PSID) but there is a long-established concern about the possible impact of measurement error in these data (see Altonji and Siow, 1987; Colera, 1993; Attanasio and Low, 2004).

Another important area of consumption research is the measurement of inequality and poverty, where it is widely accepted that survey measures of consumption expenditure provide a more reliable basis than measured income, both in developed countries (Meyer and Sullivan, 2003; Headey *et al.*, 2005) and in low-income countries (Deaton, 1997; Pudney and Francavilla, 2006). There has been work on the impact of measurement error on consumption-based inequality analysis, for example Battistin (2003), who found rather different trends implied by simple retrospective interview-based consumption data and by the more detailed information from expenditure diaries incorporated in the same US expenditure survey. See also Attanasio *et al.* (2004), who used methods for combining these measures.

Unfortunately, few surveys are able to collect consumption data in the depth we would like, particularly in the longitudinal context necessary for an understanding of dynamics of behaviour. For many classes of commodity, the most credible consumption data come from expenditure diaries, as used in the UK Expenditure and Food Survey or US Consumer Expenditure Survey, or from electronic sources such as supermarket scanner records. However, diaries are very time-consuming and electronic records usually provide little or no household contextual information. For these reasons, large-scale household panels in developed countries mostly collect limited expenditure data, using survey instruments that rely on respondents' recall of past expenditure levels or estimation of 'normal' levels. We refer to this type of expenditure estimate as a 'stylised' consumption variable.

The consumption category examined in this paper, expenditure on domestic energy, is a major element in studies of consumer demand functions (Baker *et al.*, 1989). Household energy demand is also a focus of research in energy security, environmental economics and climate change (Barker, *et al.* 1995). Domestic energy shares with many durable goods the special feature that consumption episodes are usually very frequent and not closely related to the payment schedule, over which consumers usually have a large degree of control. Energy is often paid for in large infrequent irregular amounts or by regular more frequent payments that are adjusted occasionally. Demand is also highly seasonal. These factors combine to make data collection through expenditure diaries largely infeasible, so recall is the principal question design underlying survey evidence on domestic energy.

It has been argued that stylised consumption questions can be designed in such a way that reliable statistical analysis is possible (Browning *et al.*, 2003). However, there are worrying features of the available data, particularly the pronounced heaping of observations at particular 'round' numbers. The closely related phenomena of heaping, rounding and digit preference have been studied extensively, particularly in the demographic, epidemiological and historical literatures relating to reported self-reported ages (see Bachi, 1951, for early work and Ó'Gráda, 2006, who uses heaping as an indication of innumeracy). Heaping has also been considered as a source of bias in the context of survival analysis, where distortions are often evident in retrospectively-reported durations or event dates (Baker, 1992; Torelli and Trivellato, 1993). There has been relatively little analysis of the heaping problem for variables with a cash metric and little consideration, in any area of application, of its consequences for panel data analysis.

Our aim is to investigate the impact of response error, focusing mainly on the expenditure questions relating to annual energy consumption included in the British Household Panel Study (BHPS). Most existing work on reporting error in consumption data (see Browning *et al.*, 2003 for a review) deals with the properties of stylised data in a cross-section context, whereas much of the most important research on consumption behaviour is explicitly dynamic and requires repeated observations for a panel of households. Our main contributions to the literature are an examination of the dynamic properties of error induced in stylised consumption data by the heaping phenomenon and of its consequences for common forms of panel data analysis. We begin, in section 2, by summarising the BHPS data and identifying a set of distinct response strategies used by many survey respondents. In section 3 we propose a statistical model of this reporting behaviour and develop a full dynamic model of reported expenditure, incorporating a simple dynamic consumption model. We then evaluate the biases induced by mistakenly assuming that the data are accurate. Section 5 concludes.

2 BHPS expenditure data

The BHPS is a nationally-representative annual household panel survey that began in the UK in 1991. Although the BHPS includes personal interviews with all adult household members, we are concerned here with expenditures at the household level. Interviewers try to ensure that these questions are answered by the available household member who has the best knowledge of the household's budgeting but in practice more than one respondent may be involved. The cross-wave identifier of the principal respondent to the household questionnaire is recorded, so that changes of respondent between waves can be recorded, albeit imperfectly in some cases.

There is no unambiguous definition of a household in the intertemporal context since household membership can vary considerably over time and, consequently, the BHPS has no cross-wave household identifier to track households over time periods. In this study, we use a very simple linking strategy: all households are selected at the initial wave in 1991 and the household reference person (who can be thought of as a "head of household") is identified. In later waves, the household containing that person is deemed to be the same household (even if he or she is no longer regarded as the reference person of the household concerned). The statistical methods we use in section 3 require an uninterrupted run of observations over time, so we discard any observations that follow an episode of wave non-response. However, the summary tables appearing in this section use the full set of available observations.

2.1 Expenditure on domestic energy

Since 1997, the BHPS questionnaire has contained stylised questions on the preceding 12 month's household expenditure on four types of energy. The questions are asked sequentially about gas, electricity, heating oil and solid fuel.[1] The questions are introduced as follows:

> *In the last year, since September 1st [...], approximately how much has your household spent on domestic fuel? Starting with gas...*

2.2 Modes of rounding and heaping

Figures 1-4 show the sample distributions of responses to the BHPS questions on annual electricity, gas, heating oil and solid fuel expenditures, with all waves from 1997-2004 pooled. There are clear and very pronounced peaks in these distributions, which suggest four basic modes of responding: *annual rounding*, where an estimate of the annual total is given directly as a multiple of £50; *weekly rounding*, where the response is conceived as a weekly sum to the nearest £1 and converted to an annual figure by multiplying by 52; and *monthly rounding*, where the response is conceived as some multiple of £5 per month, multiplied by 12 to give an annual figure.[2] There are potential ambiguities in this classification, since responses which are multiples of £300 could arise from either monthly or annual rounding.[3] The summary analysis presented in this section resolves this ambiguity arbitrarily by interpreting all such points as monthly-mode estimates; this convention makes some difference to the summary

[1] Some households pay combined gas and electricity bills and report a single figure covering both. We delete the small number of such cases.

[2] The existence of a quarterly billing option suggests another possibility: that the response is conceived as a rounded quarterly figure and multiplied by 4 to give an annual estimate However, this would be expected to lead to spikes at £40, £80, £160, *etc.* which are absent from the data.

[3] There are a few cases of reported expenditures of £1300, which could result either from annual rounding or from weekly rounding at £25 a week; we have interpreted them as the latter. Other potential points of ambiguity lie outside the sample range.

tables that follow but no qualitative difference to the conclusions we draw from them.[4] Our formal modelling approach, set out in section 3, resolves the ambiguity more satisfactorily by allowing explicitly for both modes of reporting.

Figure 1 shows the distribution of (positive) responses for each of the four categories of energy expenditure. In every case, the distribution is dominated by very large 'spikes' at particular values, with different configurations of spikes for the four commodities.

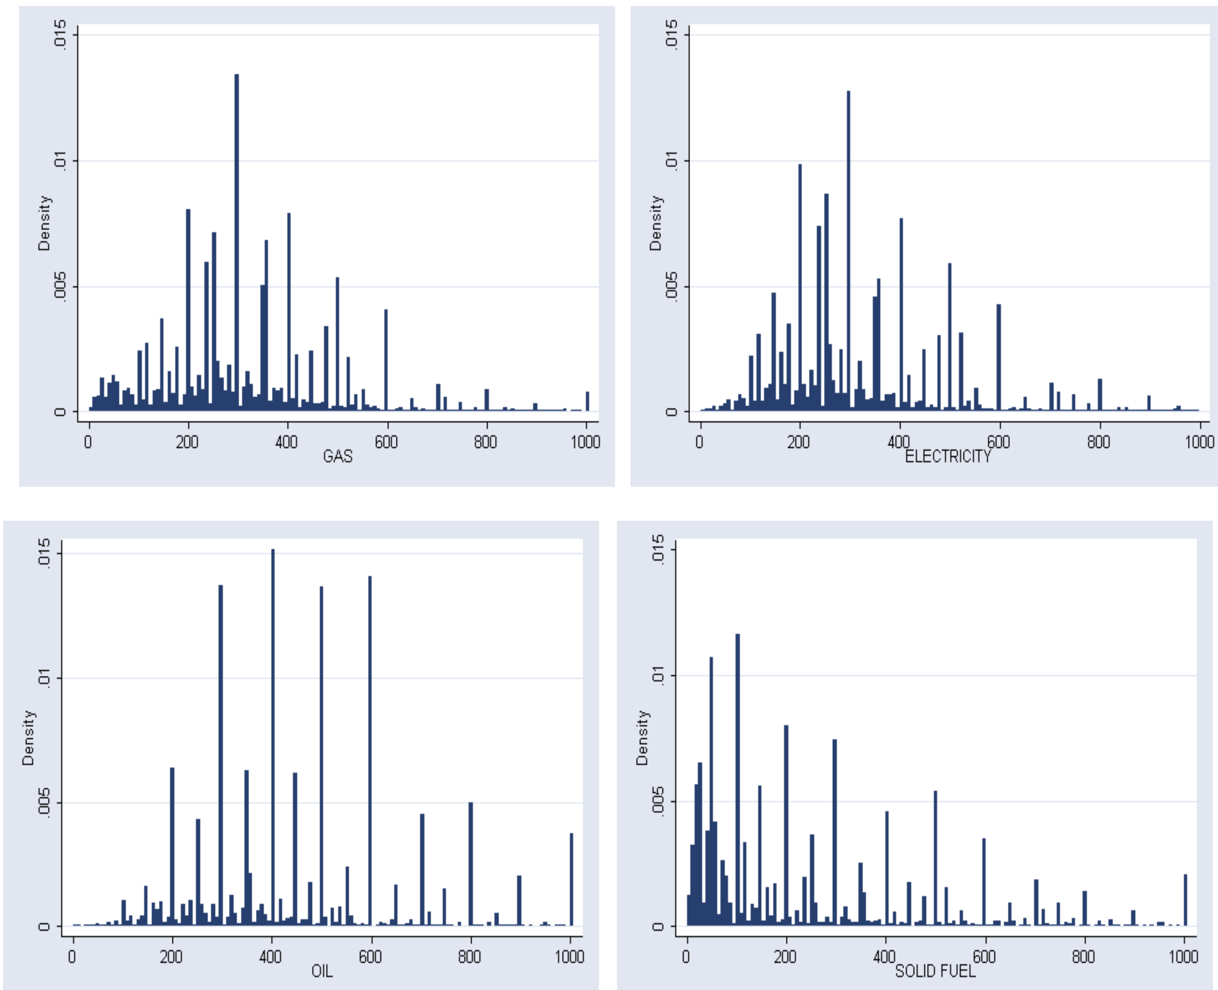

Figure 1: Distributions of reported annual energy expenditures, BHPS 1997-2004 pooled

[4]See Appendix Tables A2-A4 for variants of Tables 2-4, based on the alternative interpretation of multiples of £300 as annual rounding.

Table 1 sets out the twelve most frequent sample values for each of the expenditure categories, together with the sample mean of expenditure and the mean and standard deviation of its first difference. These are calculated only for the sample of households reporting a strictly positive value for the expenditure in question. Rounding is clearly a very important feature of the data, with the dozen most frequent sample values accounting for 54-69% of sample numbers for the four specific energy types. There are very high variances for year-to-year expenditure changes, in comparison to the mean difference (note that the median expenditure change is £0 for all expenditure categories). The observed patterns of rounding are rather different for the energy types, with annual rounding being more dominant for oil and solid fuel. Weekly rounding is uncommon and only reaches 5% of the sample for electricity.

Table 1 The dozen highest-frequency heaping points for each energy category BHPS 1997-2004 pooled

	ELECTRICITY		GAS		OIL		SOLID FUEL		TOTAL	
Rank	£p.a.	%	£p.a.	%	£p.a.	%	£p.a.	%	£p.a.	%
1	300	9.3	300	10.0	400	11.0	100	8.5	600	4.9
2	200	7.1	200	5.9	600	10.3	50	7.8	500	3.6
3	250	6.1	400	5.9	300	9.9	200	5.9	400	2.8
4	400	5.6	250	5.1	500	9.9	300	5.5	800	2.8
5	240	5.3	360	5.1	350	4.6	500	4.0	700	2.7
6	500	4.3	240	4.4	200	4.5	150	4.0	480	2.3
7	360	3.8	500	4.0	450	4.5	20	3.9	1000	2.1
8	350	3.2	350	3.6	800	3.6	30	3.6	550	2.1
9	600	3.1	600	2.9	700	3.3	400	3.4	720	2.0
10	150	3.1	480	2.5	250	3.0	60	2.9	900	1.8
11	180	2.3	150	2.4	1000	2.7	250	2.7	450	1.8
12	520	2.3	120	2.0	550	1.7	600	2.5	650	1.7
% of sample	55.4		53.6		69.0		54.4		30.6	
Mean C	345.3		327.9		481.8		264.8		672.3	
Mean (ΔC)	-0.43		0.58		6.64		-11.64		-3.50	
S. D. (ΔC)	189.6		175.6		226.6		196.7		294.8	

2.3 Transitions in the response mode

A striking feature of the sequences of responses over time is the large number of changes in the response mode used in different periods. Table 1 shows the matrices of year-to-year transition rates. Although the diagonal is mostly dominant, implying a degree of persistence, there are large off-diagonal elements, reflecting considerable temporal variation in the response strategies used by BHPS respondents.

Table 2 Transition matrix for response mode, BHPS 1997-2004 pooled

	Rounding type				
Mode of origin	Unrounded	Weekly	Monthly	Annual	**Frequency**
Electricity					
Unrounded	37.5	4.3	28.5	29.7	10,875
Weekly rounding	20.4	27.8	25.5	26.4	1,866
Monthly rounding	21.8	4.3	39.0	34.9	10,817
Annual rounding	19.1	4.0	30.1	46.8	12,953
All origins	25.5	5.4	32.0	37.1	36,511
Gas					
Unrounded	42.3	3.7	26.9	27.1	9,839
Weekly rounding	24.9	26.2	24.4	24.5	1,166
Monthly rounding	22.5	3.3	43.2	31.0	8,625
Annual rounding	21.1	3.5	30.9	44.5	9,197
All origins	28.9	4.4	33.0	33.7	28,827
Oil					
Unrounded	26.8	1.9	26.1	45.2	949
Weekly rounding	22.6	11.3	30.2	35.9	53
Monthly rounding	11.1	1.0	33.1	54.8	1,156
Annual rounding	11.8	1.1	28.8	58.3	2,041
All origins	15.1	1.4	29.4	54.1	4,199
Solid fuel					
Unrounded	47.3	2.6	14.0	36.2	1,554
Weekly rounding	22.2	21.4	18.0	38.5	117
Monthly rounding	22.2	4.3	33.3	40.3	658
Annual rounding	21.4	2.1	21.5	55.0	1,438
All origins	32.3	3.3	20.3	44.2	3,767

Response mode is related to the degree of year-to-year variability evident in the expenditure data. Table 3 shows that respondents who, on our definition, do not use rounding show much less volatility in the annual change in their reported expenditure than those who do round their responses. This is consistent with the normal expectation that heaping tends to increase the extent of measurement error. More surprising is the fact that respondents who use the same rounding mode at every wave display just as much volatility in annual reported expenditure change as respondents who switch mode. Thus there is no evidence to suggest that people adapt their rounding method as necessary in order to track the true consumption level closely (implying lower variance for 'switchers'), nor is there evidence that consistency of rounding method leads to a more stable series of expenditure reports.

Table 3 Volatility of change in reported expenditure by response mode, BHPS 1997-2004 pooled

		Rounding method constant across waves:				**Method varies across waves**
		No rounding	Weekly	Monthly	Annual	
Electricity	Mean	-1.54	-11.92	0.97	-1.11	-0.37
	Std.dev.	86.64	124.77	180.54	211.65	190.63
	n	*718*	*96*	*741*	*1,401*	*31,753*
Gas	Mean	0.23	3.53	9.05	3.45	0.28
	Std.dev.	86.60	82.99	134.21	166.82	179.33
	n	*901*	*59*	*577*	*856*	*24,421*
Oil	Mean	-3.58	-	57.60	10.45	3.94
	Std.dev.	127.13	-	296.44	217.58	226.15
	n	*48*	*1*	*125*	*598*	*3,024*
Solid fuel	Mean	3.09	62.40	-35.12	-21.17	-11.32
	Std.dev.	87.38	411.75	160.98	179.73	205.96
	n	*219*	*5*	*41*	*359*	*2,438*

2.4 Combinations of response mode

Most households use electricity and at least one other form of domestic energy. One might expect respondents to use the same method of answering questions about each form of energy. Table 4 cross-tabulates the response mode for electricity against those for gas, oil and solid fuel, and shows a more complicated pattern of behaviour. The diagonals are not consistently dominant in these tables, which implies a high rate of dissonance between the response modes for electricity and other energy types. Response behaviour is to some degree tailored to the nature of the commodity in question and the way it is purchased.

Table 4 Joint distributions of response mode for electricity and other energy types, BHPS 1997-2004 pooled

Electricity mode	Other energy type: Unrounded	Weekly	Monthly	Annual	**Frequency**
			Gas		
Unrounded	54.5	3.7	22.8	19.1	10,279
Weekly rounding	28.9	38.2	18.9	14.0	2,107
Monthly rounding	21.3	2.1	51.6	25.0	12,610
Annual rounding	17.0	1.8	24.5	56.7	14,268
All origins	28.8	4.3	32.4	34.4	39,264
			Oil		
Unrounded	24.9	1.9	28.1	45.2	1,404
Weekly rounding	21.8	2.5	25.9	54.5	243
Monthly rounding	11.5	1.6	33.2	53.8	1,936
Annual rounding	10.1	0.9	28.9	60.2	2,672
All origins	14.3	1.4	29.9	54.5	6,255
			Solid fuel		
Unrounded	39.7	3.4	19.8	37.1	1,203
Weekly rounding	29.5	14.2	20.4	36.0	339
Monthly rounding	30.3	2.5	22.0	45.3	1,781
Annual rounding	27.1	2.4	19.3	51.1	2,257
All origins	31.0	3.4	20.3	45.3	5,580
			Total Energy		
Unrounded	52.1	3.8	21.2	22.9	15,142
Weekly rounding	37.7	15.2	23.8	23.4	1,402
Monthly rounding	35.4	3.9	31.6	29.2	9,119
Annual rounding	31.4	3.2	24.8	40.6	10,751
All origins	41.3	4.1	24.9	29.7	3,767

2.5 Aggregation of expenditure categories

What happens when expenditure categories are combined into an aggregate expenditure? A simple heuristic based on the central limit theorem might suggest that, as a large number of misreported component expenditures are aggregated, the result will be a total expenditure figure deviating from the true total consumption figure by an approximately normal, zero-mean measurement error. However, this rests on some strong assumptions: that there is no

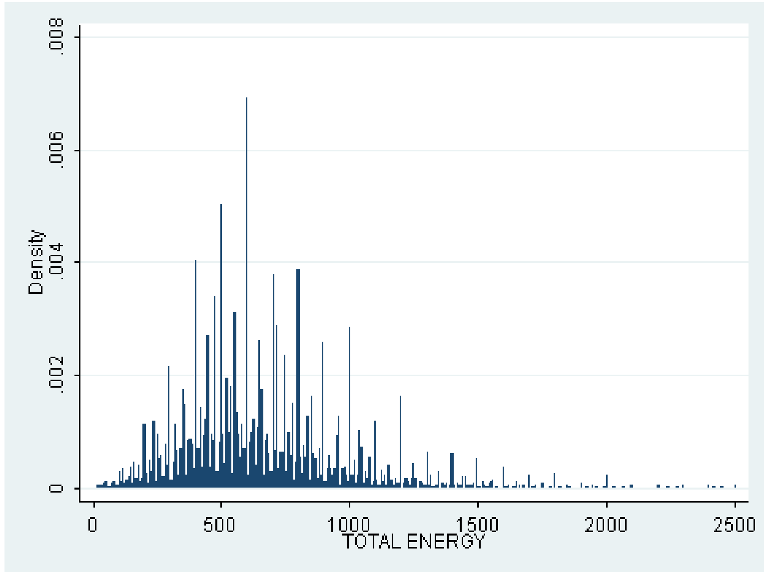

Figure 2: Distribution of annual energy expenditures aggregated over energy types, BHPS 1997-2004 pooled

systematic error across component categories; that the number of categories is large; and that the dependence between them is sufficiently weak to allow the aggregation process to reduce variance effectively. Figure 2 shows that aggregation over the four categories does not eliminate the heaping problem but it does introduce a degree of smoothing. Table 1 (above) shows that the top dozen heaping points account for only 31% of the sample, compared with 54-69% for the individual categories, and the proportion of totals classified as unrounded under our conventions increases to 41% for the aggregate compared with a maximum of 31% for any of its sub-categories (Table 3, above). However, the variance of year-to-year changes in reported total energy expenditure remains extremely large (Table 1), so it cannot be claimed that aggregation over these categories achieves much reduction in measurement error variance. It should be borne in mind here that, although we are aggregating over four energy categories, 80% of households consume only two energy types and only 8% consume more than two. It may be possible to achieve a much higher degree of variance reduction by aggregation over a larger number of expenditure categories.

2.6 Non-energy consumption expenditure

Aggregation can also be done implicitly by asking a recall question about a single broad expenditure category, rather than building up a total from components. The BHPS has such a recall question relating to food and groceries. In 1991, this was worded as follows:

> *Thinking about your weekly food bills approximately how much does your household usually spend in total on food and groceries? Write in to nearest £*

From 1992 onwards, respondents were asked instead to select one of twelve pre-defined expenditure ranges [5] from a showcard. The question wording was changed to:

> *Please look at this card and tell me approximately how much your household spends each week on food and groceries. Include all food, bread, milk, soft drinks, etc; exclude pet food, alcohol, cigarettes and meals out*

Consequently, 1992 data are subject to interval censoring by design. Figure 3 compares the distribution of responses from the 1991 and 1992 waves in density form. Note that the official rates of price inflation for food and for all consumer goods between quarter 3 of 1991 and 1992 were only 1.3% and 2.7% respectively, so the comparison is largely unaffected by price variation.

[5] Six £10 intervals over £0-59, five £20 intervals over £60-159 and an open interval of £160 or over.

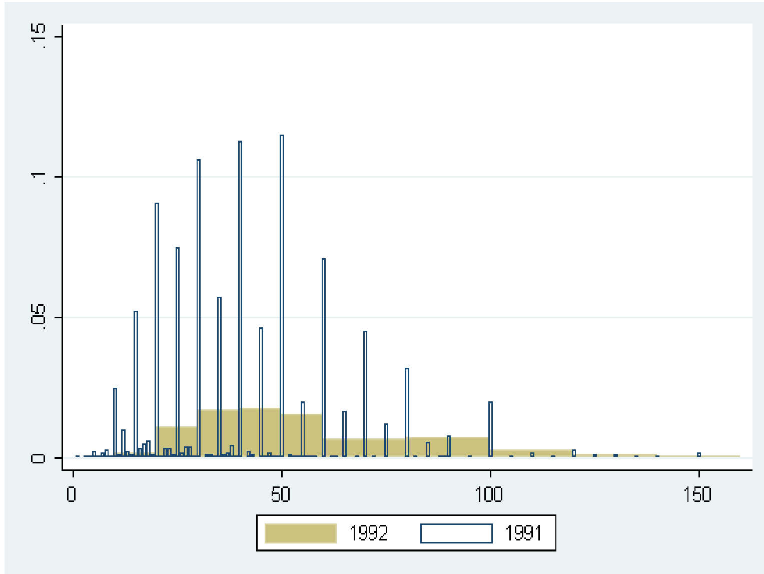

Figure 3: Distributions of weekly expenditures on food and groceries, BHPS 1991 and 1992

There are clearly difficulties with both question designs and the comparison should not be seen as a simple contrast of a continuous variable with a discrete one. The numerical responses in 1991 are dominated by 20 or so heaping points (the ten principal heaping points make up 77% of the sample and the top twenty account for 93%), so the 'continuous' expenditure variable from 1991 is essentially discretised by respondents. The 1992 variable sacrifices some detail by using only 12 response intervals, but it seems likely that there would be little difference between their information content if the question were re-designed by increasing the number of response categories. There seems no convincing case for preferring the 'continuous' variable on grounds of analytical convenience, since both question designs produce essentially discrete data.

However, the choice between the two question designs may not rest only on a possible sacrifice of information. If someone is offered a list of specific possibilities, rather than asked for a completely unguided response, he or she is being asked to follow a different cognition-decision process. It is quite conceivable that these different processes may generate conflicting answers (see Blair *et al.* 1977, Schwartz *et al.* 1985, Schaeffer and Charng 1991 for examples in various contexts where the use and design features of banded questions generates significantly different response distributions than open questions). Figure 4 examines this issue by comparing the 1991 and 1992 distributions in terms of the empirical distribution function, evaluated at the interval limits used in the 1992 questionnaire. The two curves are remarkable close, differing mainly in a slightly larger proportion of extremely low expenditures in 1992 and a net shift of cases from upper-middle expenditure levels in 1991 to high levels in 1992. Although this is not conclusive evidence, it does suggest that there is no great behavioural difference between the response processes for the two question types.

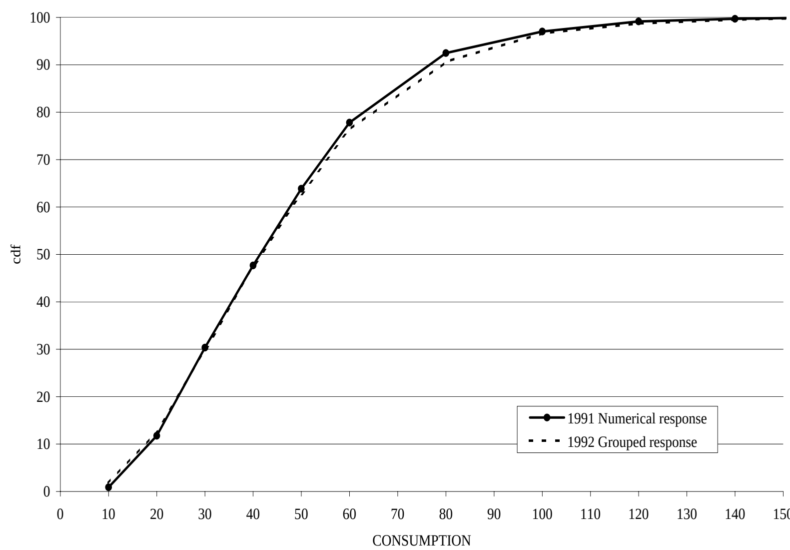

Figure 4: Distribution functions of weekly expenditures on food and groceries, BHPS 1991 and 1992

3 A joint model of consumption and response mode

We see survey response behaviour as a two-stage process. Before producing an estimate of expenditure, the respondent has to make a decision about how to form that estimate

- for example, he or she may decide to search for documentary evidence to give a precise figure or use some simple method of approximation, such as scaling up an approximate monthly amount. The important point here is that the response mode is logically prior to the response itself, just as a statistician's choice of estimator is logically prior to the calculation of an estimate. At the second stage, we assume that the respondent gives the most accurate response possible, conditional on the chosen response strategy.[6] This view of the heaping process thus interprets it as endogenously-generated interval censoring since, for example, a reported value such as £100 in the annual rounding response mode implies only that the true expenditure figure lies in some interval containing the value £100. The censoring process is endogenous in the sense that the mode of response is the outcome of a choice made by the respondent.

3.1 Response mode

Let there be M modes of response and assume initially that we can infer the response mode unambiguously by inspection of the raw expenditure data. Let the observed sequence of modes used by household i be $\{m_{it}, t = 0...T_i\}$. In our application there are M=4 modes: exact unrounded ($m_{it} = 0$); scaled-up weekly ($m_{it} = 1$), scaled-up monthly ($m_{it} = 2$); and annual rounded to the nearest £50 ($m_{it} = 3$). The first stage of the response decision entails a discrete choice between these M modes. Since there is likely to be some persistence in modes of response, we allow for (first-order) autoregressive dynamics and unobservable individual effects embedded in a multinomial logit probability structure with separate components for the transition model and the initial mode:

$$P\left(m_{it} = k|\boldsymbol{Z}_i, m_{it-1} = j, \xi_i\right) = \frac{exp\left(\gamma_{jk} + \boldsymbol{z}_{it}\boldsymbol{\beta}_k + \kappa_k\xi_i\right)}{\sum_{m=0}^{M} exp\left(\gamma_{jm} + \boldsymbol{z}_{it}\boldsymbol{\beta}_m + \kappa_m\xi_i\right)}\ , \quad t = 1...T_i \tag{1}$$

$$P\left(m_{i0} = j|\boldsymbol{z}_i^0, \xi_i\right) = \frac{exp\left(\boldsymbol{z}_i^0\boldsymbol{\alpha}_j + \kappa_{j0}\xi_i\right)}{\sum_{m=0}^{M} exp\left(\boldsymbol{z}_i^0\boldsymbol{\alpha}_m + \kappa_m^0\xi_i\right)} \tag{2}$$

[6]It is possible that there is error in the respondent's underlying perception, so that unrounded responses contain measurement error and rounded values are based on erroneous perceptions. This raises a fundamental identification problem that cannot be avoided without external validation data or further strong assumptions.

where: $\boldsymbol{Z}_i = \{\boldsymbol{z}_i^0, \boldsymbol{z}_{i1}...\boldsymbol{z}_{iT_i}\}$ is a sequence of (row) vectors of strictly exogenous observed covariates describing the characteristics of the respondent and the interview conditions; ξ_i is an unobservable N(0,1) random effect; κ_j and κ_j^0 are scale parameters and $\{\boldsymbol{\alpha}_m, \boldsymbol{\beta}_m\}$ is the vector of coefficients specific to mode m. The parameters γ_{jk} represent the autoregressive effect of the previous period's response mode $m_{it-1} = j$ on the current response $m_{it} = k$. The mode choice parameters are subject to the following normalisation restrictions:

$$\boldsymbol{\alpha}_0 = \mathbf{0}; \;\; \boldsymbol{\beta}_0 = \mathbf{0}; \;\; \gamma_{00} = ... = \gamma_{M0} = 0; \;\; \kappa_0 = \kappa_0^0 = 0 \tag{3}$$

Under our assumptions, a likelihood function for the model of mode choice alone can be constructed as the product of an initial term (2) and a sequence of terms (1).

3.2 Ambiguous response modes

It is not always possible to distinguish unambiguously between alternative possibilities for the response mode in use by a given individual. There are two issues: (*i*) the unrounded mode might by chance generate some observations situated at heaping points; (*ii*) two different rounding modes might be capable of generating the same value. The first of these is dealt with by assuming that, in unrounded mode, expenditure is observed as a continuous variable with a smooth density, implying that the heaping points are a set of probability measure zero and the rounded observations can thus be identified almost surely *a priori*. This is, of course, a simplifying approximation since the question asks for expenditures only to the nearest £.

The second case is more complicated. For individuals with a number $T_i^* \in \{0...T_i\}$ of observations which are multiples of £300 (and therefore consistent with both monthly and annual rounding) there are $2^{T_i^*}$ possible mode sequences consistent with the time series of reported expenditures. In such cases, we compute the likelihood as the sum of the separate likelihoods for each of these possible sequences. This method is only feasible for short panels. In longer panels, for cases where $2^{T_i^*}$ is a large number, it is possible to calculate an unbiased estimate of the likelihood by repeatedly drawing at random (with replacement) one of the $2^{T_i^*}$ sequences, calculating the sub-likelihood for each and then averaging over the sampled sequences.

3.3 Consumption

Let the true (log) consumption process be the following autoregressive random-effects structure:

$$c_{it} = \boldsymbol{x}_{it}\boldsymbol{\delta} + \rho c_{it-1} + \psi\xi_i + \sigma_\upsilon \upsilon_i + \sigma_\epsilon \epsilon_{it} \tag{4}$$

where: c_{it} is the logarithm of expenditure; υ_i is a persistent individual effect specific to consumption; ψ, σ_υ and σ_ϵ are scale parameters; and $\{\epsilon_{it}\}$ is a sequence of independent N(0,1) variates.

Our data come from panel with waves indexed by $t = -\tau...0...T_i$, where period 0 is the wave at which observation of consumption begins. The sequence of covariates $\boldsymbol{X}_i = \{\boldsymbol{x}_{i,-\tau}...\boldsymbol{x}_{iT_i}\}$ is fully observed. We use the following approximation to the initial observation at wave 0:

$$c_{i0} = \boldsymbol{x}_i^0\boldsymbol{\delta} + \frac{\psi}{1-\rho}\xi_i + \frac{\sigma_\upsilon}{1-\rho}\upsilon_i + \frac{\sigma_\epsilon}{\sqrt{1-\rho^2}}\epsilon_{i0} \tag{5}$$

where:

$$\boldsymbol{x}_i^0 = \sum_{s=0}^{\tau-1} \rho^s \boldsymbol{x}_{i,-s} + \frac{\rho^\tau}{1-\rho}\boldsymbol{x}_{i,-\tau} \tag{6}$$

This approximation rests on the assumption that $|\rho| < 1$, the consumption process is long-established and that $\boldsymbol{x}_{i,-\tau}$ is a good proxy for the sequence of preceding $\boldsymbol{x}$-vectors. Given that ρ^τ is likely to be small ($\tau = 5$ in our application), this seems an innocuous approximation.

Under our assumptions, c_{it} is observed exactly if the unrounded response mode $m_{it} = 0$ is used at time t, but is interval censored if any other response mode is used. Methods of dealing with the heaping problem often make the dual assumptions of a single universal set of heaping points and, for the affected observations, of rounding to the closest such value. A much-cited extension due to Heitjan and Rubin (1990) allows for an ordered set of rounding methods, but remains within a static cross-section framework. Our approach generalises previous work in three ways. First, we allow for the coexistence of different (unordered) methods of rounding by respondents; second, we model behavioural dependence between the mode of rounding used and the underlying true consumption behaviour; and third, we relax the assumption that, within any rounding mode, all heaping points have the same degree of attractiveness as 'round' values.

Given the respondent's choice of response mode, the true consumption variable, c_{it}, is converted into reported consumption, c_{it}^{+}, as follows. In mode 0, $c_{it} = c_{it}^{+}$ with probability 1. For each mode $m > 0$, there is a countable set of (log) heaping points, $\Gamma_j^m, j = 1, 2, ...$ Not all of these points are equally attractive as rounded values: for example, in annual rounding mode, even multiples of £50 are more heavily used by respondents than odd multiples. Define a set of (inverse) attraction levels A_j^m associated with the points Γ_j^m, where the A_j^m are normalised by $A_j^m = 1$ for some j within each m. Then assume that c_{it} is rounded as Γ_j^m if and only if mode m is chosen and $c_{it} \in \left[L_j^m, U_j^m\right)$, where:

$$L_j^m = \begin{cases} -\infty, & j = 1 \\ \frac{A_{j-1}^m \Gamma_{j-1}^m + A_j^m \Gamma_j^m}{A_{j-1}^m + A_j^m}, & j > 1 \end{cases} \tag{7}$$

$$U_j^m = \frac{A_j^m \Gamma_j^m + A_{j+1}^m \Gamma_{j+1}^m}{A_j^m + A_{j+1}^m} \tag{8}$$

3.4 Estimation

To estimate the combined consumption-response model (2), (1) and (5)-(8), we use the Geweke-Hajivassiliou-Keane maximum simulated likelihood (MSL) method, extended to allow for unobserved heterogeneity in addition to the interval censoring induced by rounding behaviour. Write the vector of true consumption values as $\mathbf{c}_i = (c_{i0}...c_{iT_i})$. Under a normality assumption, (5) and (4) imply $\mathbf{c}_i|\xi_i, \boldsymbol{X}_i \sim N(\boldsymbol{\mu}_i(\xi_i), \boldsymbol{\Omega})$ where the typical elements of the mean vector and covariance matrix are:

$$\mu_{it}(\xi_i) = \left(\rho^t \boldsymbol{x}_i^0 + \sum_{s=0}^{t-1} \rho^s \boldsymbol{x}_{it-s}\right) \boldsymbol{\delta} + \frac{\psi}{1-\rho}\,\xi_i \tag{9}$$

$$\omega_{st} = \sigma_\epsilon^2 \frac{\rho^{t-s}}{1-\rho^2} + \sigma_\upsilon^2 \frac{1}{(1-\rho)^2} \quad , \quad s \leq t \tag{10}$$

where $V_t(\rho) = \frac{1-\rho^t}{1-\rho}$ if $\rho \neq 1$ and $V_t(\rho) = t$ if $\rho = 1$. In (9) for $t = 0$, the summation from $s = 0$ to -1 is interpreted as zero. Now partition $\mathbf{c}_i$ into a subvector $\mathbf{c}_{i1}$ of elements assumed to be reported accurately (*i.e.* in annual unrounded form) and the remaining elements $\mathbf{c}_{i2}$ reported in one of the rounded modes, for which we only observe the interval within which consumption falls. Partition μ_i and $\boldsymbol{\Omega}$ conformably and decompose the joint distribution into the following marginal and conditional components:

$$\mathbf{c}_{i1} | \boldsymbol{X}_i, \xi_i \sim N\left(\mu_{i1}, \boldsymbol{\Omega}_{11}\right) \tag{11}$$

$$\mathbf{c}_{i2} | \mathbf{c}_{i1}, \boldsymbol{X}_i, \xi_i \sim N\left(\mu_{i2} + \boldsymbol{\Omega}_{21}\boldsymbol{\Omega}_{11}^{-1}\left(\mathbf{c}_{i1} - \mu_{i1}\right), \boldsymbol{\Omega}_{22} - \boldsymbol{\Omega}_{21}\boldsymbol{\Omega}_{11}^{-1}\boldsymbol{\Omega}_{12}\right) \tag{12}$$

Conditional on the individual effects ξ_{iq} and the response mode determined at the first stage of the response process, the likelihood of the observation $\mathbf{c}_i$ is the product of the normal marginal density (11) multiplied by the probability for the conditional distribution (12) over the hyper-rectangle defined by the vector of intervals $(\mathbf{a}_i, \mathbf{b}_i)$. The latter rectangle probability can be approximated using the GHK simulator, based on the following representation:

$$P\left(\mathbf{c}_{i2} \in (\mathbf{a}_i, \mathbf{b}_i)\,|\mathbf{c}_{i1}, \boldsymbol{X}_i, \xi_i\right) = E h_i\left(\boldsymbol{\zeta}_i, \xi_i\right) \tag{13}$$

where: $h_i\left(\boldsymbol{\zeta}_i, \xi_i\right)$ is a function defined as a product of sequentially-conditioned normal densities. For example, if $\mathbf{c}_{i1}$ contains all T_i consumption variables:

$$h_i\left(\boldsymbol{\zeta}_i, \xi_i\right) = g_{i0}(\boldsymbol{\zeta}_i, \xi_i) \prod_{t=1}^{T_i} g_{it}(\boldsymbol{\zeta}_i, \xi_i) \tag{14}$$

where g_{i0} and g_{it} are the truncated normal densities $f\left(c_{i0} | c_{i0} \in (a_{i0}, b_{i0}), \boldsymbol{x}_i^0, \xi_i\right)$ and $f\left(c_{it} | c_{i0} ... c_{it-1}, c_{it} \in (a_{it}, b_{it}), \boldsymbol{X}_i, \xi_i\right)$ evaluated at $\mathbf{c}_i = \boldsymbol{\zeta}_i$ (see Hajivassiliou and Ruud, 1994 for details); and $\boldsymbol{\zeta}_i$ is a vector of variates drawn independently from these truncated densities.

The likelihood for household i can then be written:

$$L_i = E_{\xi\zeta}\{P(\boldsymbol{m}_i|\boldsymbol{Z}_i, \xi)h_i(\boldsymbol{\zeta}, \xi)\phi(\mathbf{c}_{i0}; \mu_{i0}(\xi), \boldsymbol{\Omega}_{00})\} \tag{15}$$

where $\phi(.; \mu, \boldsymbol{\Omega})$ is the multivariate normal density and $\boldsymbol{m}_i = (m_{i0}...m_{iT_i})$. The simulated likelihood approximation is:

$$\tilde{L}_i = \frac{1}{R}\sum_{r=1}^{R} P\left(\boldsymbol{m}_i|\boldsymbol{Z}_i, \xi_i^{(r)}\right) h_i\left(\boldsymbol{\zeta}_i^{(r)}, \xi_i^{(r)}\right) \phi\left(\mathbf{c}_{i0}; \mu_{i0}(\xi_i^{(r)}), \boldsymbol{\Omega}_{00}\right) \tag{16}$$

where r indexes the R replications. The simulated log likelihood for the whole sample is the sum of logs of terms like (16) for each household. The log-likelihood function is maximised numerically with respect to the parameters of the model, with the R sets of pseudo-random variates $\{\boldsymbol{\zeta}_i^{(r)}, \xi_i^{(r)}, i = 1...n\}$ held fixed through the optimisation process.

In the application reported in the following section, we implement this estimator in a two-stage process. First, initial approximate MSL estimates are produced, using crude Monte Carlo simulation with $R = 50$ replications. The second stage starts from this point in the parameter space, increasing the number of replications to $R = 150$, with the use of antithetic acceleration for each case i where antithetic acceleration gives a reduced estimated simulation variance for the log-likelihood, relative to simple Monte Carlo.

4 Results

The combined consumption-rounding model has been estimated for two definitions of consumption expenditure: electricity and total domestic energy. In both cases, the issue of zero expenditures is of negligible importance: almost all households use electricity (97.6% of the original pooled sample[7]) and a still larger proportion (98.7%) use some form of purchased energy. We also compare this model with results from conventional estimators which ignore the heaping problem. The explanatory covariates used in the main consumption

[7]This figure under-estimates electricity usage: some of the apparent non-buyers of electricity receive it included in rent, others may pay a combined gas-electricity charge.

model (4) include household size, housing tenure, household per capita income, the age of the main household-questionnaire respondent (entered as a quadratic), the change in income and household size since last year, and the fuel used for central heating (if any). The model of response mode (1) has the same covariates, except that the central heating dummies proved insignificant and have been omitted and additional variables are included to reflect the gender and any change in the identity of the principal respondent from the previous year. The latter variable is interacted with the lagged mode dummies to allow the possibility that a change of respondent might disrupt the reporting dynamics.

We allow for different degrees of attractiveness of rounding points within the weekly, monthly and annual rounding modes by specifying the parameters A_j^m appearing in (7) and (8) as follows. For the weekly rounding mode ($m = 1$), $A_j^1 = 1$ if point Γ_j^1 is a multiple of £260 and equal to a parameter a^1 otherwise; for monthly rounding, $A_j^2 = 1$ when Γ_j^2 is a multiple of £120 and $A_j^2 = a^2$ otherwise; and for annual rounding, $A_j^3 = 1$ when Γ_j^3 is a multiple of £100 and $A_j^3 = a^3$ otherwise. The parameters $a^1 ... a^3$ are estimated and expected to be greater than 1.

For the initial state of the response mode choice, equation (2), we use the vector $\boldsymbol{z}_i^0 = (\boldsymbol{z}_{i0} \quad \overline{\boldsymbol{z}}_i^*)$ as covariates, where $\overline{\boldsymbol{z}}_i^*$ is the average, over all waves $t = -\tau \; ... \; T_i$, of the time-varying elements of $\boldsymbol{z}_{it}$.

Attempts to implement the full model (1), (2), (4) (5), (7) and (8) led to a corner solution of the likelihood maximisation problem, with $\hat{\sigma}_v = 0$. Consequently, we work with a single-factor model with σ_v restricted to be zero. Full parameter estimates for the electricity and total energy models are given in appendix tables A2 and A3.

4.1 Electricity consumption

Table 5 summarises the influence of household characteristics on the mode of rounding by comparing the predicted probability of mode persistence across a number of hypothetical household types. Mode persistence is defined as recurrence of the response mode used at the previous year's interview and there are consequently four conditional persistence probabilities, one for each of the four possible lagged modes. These predicted probabilities are first calculated for a baseline household type, specified as containing two members, who own their own home, have an income of £10,000 per head; and where the primary respondent to the expenditure questions is a 50-year old woman. For this household, the probability of retention of the initial response mode ranges from 0.192 for weekly rounding to 0.525 for non-rounding. Annual rounding and non-rounding are considerably more persistent than weekly or monthly rounding.

A number of separate changes to this baseline are made to explore the impact of particular household characteristics, specifically: adding a household member; moving from owner-occupation to social rented housing or to private rental; a 10% income increase; a male rather than female respondent; and variation in the respondent's age to 40 and to 60. Few of these variations lead to large changes in the origin-specific probabilities of mode persistence. The one striking result is the strong influence of housing tenure on the persistence of weekly rounding. Social tenants and, to a lesser degree, private tenants are much more likely than owner-occupiers to persist in a weekly basis for expenditure reporting. It is likely that the housing tenure effect is a combination of factors, including the more frequent use of short-term bill-payment arrangements in many rented properties and also a more general tendency towards short-term budgeting among the 'lower' social classes, where home-ownership is less common. It is striking that income itself is not a significant factor.

Table 5 Influence of household characteristics on response mode: electricity

Change	**Initial mode** Unrounded	Weekly	Monthly	Annual
Pr(mode persistence)	0.525	0.192	0.346	0.517
	(0.015)	(0.033)	(0.017)	(0.015)
	Δ Pr(mode persistence)			
+1 household member	-0.036	0.018	0.008	0.023
	(0.014)	(0.027)	(0.017)	(0.016)
Social tenant	-0.070	0.248	-0.003	-0.022
	(0.021)	(0.044)	(0.023)	(0.022)
Private tenant	0.003	0.123	-0.079	0.032
	(0.033)	(0.065)	(0.033)	(0.031)
10 % income increase	0.000	-0.001	0.001	-0.000
	(0.016)	(0.029)	(0.019)	(0.017)
Male respondent	-0.007	-0.020	0.007	0.006
	(0.015)	(0.026)	(0.018)	(0.017)
Age 40	-0.035	0.011	0.027	0.005
	(0.018)	(0.031)	(0.019)	(0.018)
Age 60	0.026	-0.015	-0.024	0.001
	(0.015)	(0.026)	(0.018)	(0.016)

Baseline: 2-person household; homeowner; annual income £10,000 p.c.; female respondent aged 50; standard errors in parentheses

Table 6 compares the coefficients of the consumption part of the model with results from a dynamic random effects regression (estimated using simulated maximum likelihood) and the Blundell-Bond (1998) GMM estimator, both of which ignore the heaping problem. All three approaches use essentially the same set of assumptions about the initial value of the consumption series. The GMM estimator avoids the assumption of independence between the explanatory covariates and the individual effect ξ_i, which underlies both the random effects model and our extension which allows for heaping.

There are substantial differences between the three sets of estimates. The heaping model differs from the other two primarily in terms of the estimated impact of: housing tenure (where home ownership has a significant positive, rather than negative or insignificant, effect); income (where the estimated relationship is U-shaped rather than increasing and mildly concave); and form of central heating (where the impact of oil-fired central heating is to increase, rather than decrease, electricity consumption. Table 7 quantifies these differences in terms of the predicted long run comparative statics effects of changes in household characteristics. These impacts are calculated as $100 \times [\exp\{\Delta \boldsymbol{x}\boldsymbol{\delta}/(1-\rho)\} - 1]$, where $\Delta \boldsymbol{x}$ is the change in the baseline characteristics.

All three models suggest that the autoregressive parameter ρ is of moderate size, but the heaping model suggests a negative, rather than positive sign. The heaping model and random-effects regression also differ substantially in terms of the importance of the individual effect ξ_i, as shown by the intra-class correlation, $\psi^2/(\psi^2 + \sigma_\epsilon^2)$, which is only 34% for the random effects regression but 76% for the heaping model. Thus, allowing explicitly for response behaviour has changed the estimated autocorrelation structure from a mixed persistent effect-autoregressive structure to one where the persistent effect dominates.

Table 6 Coefficient estimates of the consumption process: electricity

	Heaping model		RE regression		Blundell-Bond	
	coefficient	std. error	coefficient	std. error	coefficient	std. error
Intercept	5.236	0.048	3.927	0.076	5.207	0.236
H/hold size-2	0.170	0.004	0.162	0.006	0.096	0.017
(H/hold size-2)^2	-0.018	0.001	-0.018	0.002	-0.011	0.005
Homeowner	0.076	0.009	0.019	0.017	-0.097	0.053
Social tenant	0.001	0.010	-0.037	0.018	0.013	0.056
Income p.c.	-0.086	0.052	0.387	0.058	0.066	0.119
Income p.c.^{2}	0.068	0.013	-0.021	0.005	-0.007	0.008
Age/10	0.204	0.013	0.176	0.021	0.004	0.080
(Age/10)^2	-0.204	0.011	-0.162	0.018	0.004	0.070
Electric c.h.	0.211	0.008	0.309	0.014	0.113	0.043
Gas c.h.	-0.114	0.006	-0.162	0.011	-0.164	0.033
Oil c.h.	0.172	0.008	-0.029	0.020	-0.081	0.064
c_{it-1}	-0.049	0.004	0.228	0.008	0.104	0.012
σ_ϵ	0.274		0.336		-	
Intra-class corr.	0.764		0.311		-	

Table 7 Long-run percentage impact of household characteristics on consumption: electricity

Change	Heaping model	RE model
+1 household member	15.81	18.64
	(0.44)	(0.64)
Social tenant	-6.93	-6.90
	(0.76)	(1.42)
Private tenant	-7.00	-2.38
	(0.77)	(2.12)
10 % income increase	-0.07	0.50
	(0.05)	(0.07)
Age 40	-1.97	-3.91
	(0.41)	(0.74)
Age 60	-1.88	-0.20
	(0.29)	(0.48)
Electric c.h.	36.34	84.03
	(0.82)	(2.82)
Oil c.h.	31.32	18.83
	(1.08)	(2.91)
Solid fuel c.h.	11.47	23.40
	(0.60)	(1.71)

Baseline: 2-person household; homeowner; annual income £10,000 p.c.; female respondent aged 50; gas central heating standard errors in parentheses

4.2 Total energy consumption

Tables 8, 9 and 10 give the same model summaries for total energy expenditure. The general conclusions are remarkably similar to those for electricity expenditure. A greater proportion of sample observations appear unrounded, as a consequence of the combination of different rounding methods used for the constituent categories. An implication of this is that the estimated degree of persistence in unrounded mode is higher for total energy than for electricity alone, with lower persistence for all other expenditure categories. The influence of household characteristics on response mode appears weaker than for the single energy category, with housing tenure again being the only strong influence.

As we found for electricity expenditure, response error causes substantial bias in estimates of the consumption process, particularly in terms of its dynamic properties. When allowance is made for the rounding process, the positive autoregressive coefficient becomes small and negative, with the unobserved household effect carrying virtually all of the persistence of consumption behaviour across waves. There are also substantial biases in the effects of some explanatory covariates in terms of both short-run and long-run coefficients.

Table 8 Influence of household characteristics on response mode: total energy

	Initial mode			
Change	Unrounded	Weekly	Monthly	Annual
Pr(mode persistence)	0.620	0.114	0.249	0.441
	(0.012)	(0.021)	(0.016)	(0.014)
	Δ Pr(mode persistence)			
+1 household member	-0.006	-0.004	-0.009	0.020
	(0.012)	(0.014)	(0.014)	(0.015)
Social tenant	-0.030	0.125	0.011	-0.041
	(0.017)	(0.024)	(0.020)	(0.020)
Private tenant	-0.036	0.007	-0.003	0.046
	(0.028)	(0.034)	(0.030)	(0.031)
10 % income increase	-0.002	-0.002	0.003	0.000
	(0.012)	(0.014)	(0.014)	(0.014)
Male respondent	-0.008	0.002	0.003	0.006
	(0.012)	(0.014)	(0.015)	(0.016)
Age 40	-0.016	0.014	0.017	-0.004
	(0.015)	(0.017)	(0.017)	(0.017)
Age 60	0.014	-0.013	-0.018	0.007
	(0.011)	(0.013)	(0.013)	(0.014)

Baseline: 2-person household; homeowner; annual income £10,000 p.c.; female respondent aged 50; standard errors in parentheses

Table 9 Coefficient estimates of the consumption process: total energy

	Heaping model		RE regression		Blundell-Bond	
	coefficient	std. error	coefficient	std. error	coefficient	std. error
Intercept	6.006	0.030	4.376	0.078	5.294	0.222
H/hold size-2	0.120	0.002	0.145	0.005	0.081	0.015
$(\text{H/hold size-2})^2$	-0.012	0.001	-0.015	0.002	-0.009	0.004
Homeowner	0.171	0.007	0.066	0.014	-0.049	0.048
Social tenant	0.124	0.006	-0.033	0.014	-0.015	0.052
Income p.c.	0.003	0.025	0.253	0.051	-0.147	0.107
Income p.c.^2	-0.003	0.004	-0.013	0.004	0.003	0.007
Age/10	0.224	0.008	0.170	0.021	0.126	0.072
$(\text{Age}/10)^2$	-0.224	0.007	-0.152	0.017	-0.085	0.063
Electric c.h.	-0.176	0.005	-0.041	0.012	-0.085	0.039
Gas c.h.	-0.020	0.004	0.001	0.009	-0.071	0.030
Oil c.h.	0.081	0.008	0.149	0.016	-0.008	0.060
c_{it-1}	-0.068	0.003	0.226	0.008	0.116	0.012
σ_ϵ	0.246		0.296		-	
Intra-class corr.	0.773		0.385		-	

Table 10 Long-run percentage impact of household characteristics on consumption: total energy

Change	**Heaping model**	**RE model**
+1 household member	10.64	17.03
	(0.16)	(0.57)
Social tenant	-4.32	-12.00
	(0.56)	(1.17)
Private tenant	-14.77	-8.14
	(0.57)	(1.62)
10 % income increase	0.00	0.32
	(0.02)	(0.07)
Age 40	-2.06	-4.21
	(0.23)	(0.71)
Age 60	-2.09	0.36
	(0.17)	(0.44)
Electric c.h.	-13.57	-5.22
	(0.35)	(1.26)
Oil c.h.	9.91	21.14
	(0.72)	(2.42)
Solid fuel c.h.	1.86	-0.07
	(0.37)	(1.20)

Baseline: 2-person household; homeowner; annual income £10,000 p.c.; female respondent aged 50; gas central heating standard errors in parentheses

5 Conclusions

This study gives rise to a number of tentative conclusions on the design of consumption questions and the impact of response error on dynamic statistical models of consumption.

5.1 Implications for survey design

A first conclusion is that different respondents answer questions in different ways and there is a need for more information on the factors underlying this diversity. In the context of energy consumption, it would be helpful to preface questions on expenditure by questions on the method and frequency of payment. For example, it is likely that respondents will follow a different cognitive path if they pay by monthly direct debit than if they use pre-payment meters. A potential difficulty here is that technological change in payment methods makes it difficult to maintain a body of relevant survey questions and to use the responses effectively.

A second definite finding is that recall expenditure data of this kind are inevitably discrete, however the question is designed. The BHPS questionnaire asks for a response to the nearest £1 but the responses are not continuously distributed: the top dozen heaping points account for well over a half of the sample and the top twenty account for over 90%. The choice between an open question design yielding a ‘continuous’ response and a banded design yielding an ordinal response is largely illusory. There are two other issues which should govern the choice between open and banded designs: the number of bands and the properties of the (explicit or implicit) discretisation process. For a banded design, BHPS evidence suggests that around 20 bands would be required to avoid significant loss of information relative to the open design. There can be difficulties in using so many bands, particularly with showcard methods. The second issue may be more important. Open questions leave the respondent to decide how to round his or her answer and the resulting implicit discretisation process may be endogenous to the consumption behaviour we are trying to observe.

In contrast, banded questions involve exogenous discretisation specified by the question designer, which avoids the statistical complexities entailed by endogenous rounding. This is a strong argument in favour of banded questions, provided a sufficient number of bands can be used, but there are counter-arguments. Change in expenditure levels, particularly through strong price inflation, renders any fixed set of bands inappropriate after a period of time and the updating process may introduce comparability problems. Aggregation over expenditure categories is also problematic with banded data, although, arguably, the difficulties are just as great for unbanded data, which are also largely discrete.

A third important conclusion is the very large variance of year-to-year changes in reported expenditure on energy, much of which is likely to result from measurement error. This suggests the need for error-reduction devices such as dependent interviewing or external validation data. Dependent interviewing involves follow-up questions in cases where there is a large change from the previous wave; this would tend to reduce the scale of error but it also risks imposing artificial stability, causing a different, not necessarily less serious, type of distortion in the data. Another possibility is the use of validation samples, such as the billing records of energy utilities, to estimate the true volatility of expenditure movements, from which we could drive an estimate of the overall variance of measurement error in recall data. At present, in the UK at least, it is almost certainly infeasible to link utility company billing records to survey households at the micro level.

A fourth finding is that the different rounding methods used by respondents may lead to different degrees of measurement error, which suggests the possibility of substantial efficiency gains from use of a weighting strategy which takes some account of the (inferred) rounding method. However, the potential endogeneity of the rounding method is a serious complicating factor here.

5.2 Implications for consumption modelling

There is a very large and sophisticated econometric literature which uses stylised consumption data to estimate dynamic models of consumer expenditure and, from this, draw inferences about intertemporal decision-making and welfare change. The recent literature has considered the issue of measurement error in consumption data and its implications for bias in estimated dynamic models but has generally assumed classical, zero-mean independent measurement errors. This is despite the fact that the obvious heaping of responses is a very strong signal that measurement error is far from classical.

This study has demonstrated that the rounding processes responsible for sample heaping are complex in nature, with their own dynamic pattern superimposed on that of the underlying true consumption process. Extension of a simple autoregressive energy demand model to encompass this complex pattern of misreporting has led to results that differ quite substantially from those generated by standard econometric modelling techniques which ignore the heaping problem. The important conclusion to be drawn from this finding is that we should change our modelling outlook. Observed panel survey data on expenditures constitute a single observation on the joint outcome of two processes: the true consumption process and a quite separate interview-response process. A full understanding of the data would require joint modelling of both processes, using a model of response that is considerably more realistic than the classical measurement error assumptions.

References

[1] Altonji, J. and Siow, A. (1987). Testing the response of consumption to income changes with noisy panel data, *Quarterly Journal of Economics* **102**, 293-328.

[2] Attanasio, O. and Low, H. (2004). Estimating Euler equations, *Review of Economic Dynamics* **7**, 405-435.

[3] Attanasio, O., Battistin, E. and Ichimura, H. (2004). What really happened to consumption inequality in the US? NBER Working Paper 10338.

[4] Bachi, R. (1951). The tendency to round off age returns: measurement and correction. *Bulletin of the International Statistical Institute* **33**, 195-221.

[5] Baker, M. (1992). Digit preference in CPS unemployment data, *Economics Letters* **39**, 117-121.

[6] Baker, P., Blundell, R. and Micklewright, J. (1989). Modelling household energy expenditures using micro data, *Economic Journal* **99**, 720-738.

[7] Barker, T., Ekins, P. and Johnstone, N. (eds.) (1995). *Global Warming and Energy Demand*, London: Routledge.

[8] Battistin, E.(2003). Errors in survey reports of consumption expenditures, Institute for Fiscal Studies Working Paper 03/07.

[9] Blair, E., Sudman, S., Bradburn, N. M. and Stocking, C. (1977). How to ask questions about drinking and sex: response effects in measuring consumer behavior, *Journal of Market Research* **14**, 316-321.

[10] Blundell, R. and Bond, S. (1998). Initial conditions and moment restrictions in dynamic panel data models, *Journal of Econometrics* **87**, 115-143.

[11] Browning, M., Crossley, T. F. and Weber, G. (2003). Asking consumption questions in general purpose surveys, *Economic Journal* **113** , F540-F567.

[12] Colera, E. V. (1993). A note on measurement error and Euler equations, *Economics Letters* **45**, 305-308.

[13] Deaton, A. (1997). *The Analysis of Household Surveys. Microeconometric Analysis for Development Policy*. Baltimore: Johns Hopkins University Press (for the World Bank).

[14] Hajivassiliou, V. and Ruud, P. (1994). Classical estimation methods for LDV models using simulation, in Engle, R. F. and McFadden, D. L. (eds.) *Handbook of Econometrics vol. 4*, 2383-2441. Amsterdam: North-Holland.

[15] Headey, B., Muffels, R. and Wooden, M. (2005). Money and happiness: a reconsideration based on the combined effects of wealth, income and consumption, *Journal of Applied Social Science Studies (Schmollers Jahrbuch)* **125**, 131-144.

[16] Heitjan, D. F. and Rubin, D. B. (1990). Inference from coarse data via multiple imputation with application to age heaping, *Journal of the American Statistical Association* **85**, 304-314.

[17] Lillard, D. and Wang, H. (2007). A heap of trouble? Accounting for mismatch bias in retrospective data, paper presented at 2007 BHPS conference, University of Essex.

[18] Meyer, B. D. and Sullivan, J. X. (2003). Measuring the well-being of the poor using income and consumption, *Journal of Human Resources* **38** , 1180-1220.

[19] Ó'Gráda, C. (2006). Dublin Jewish demography a century ago. UCD Centre for Economic Research Working Paper Series WP 06/01.

[20] Pudney, S. E. and Francavilla, F. (2006). Income Mis-Measurement and the Estimation of Poverty Rates. An Analysis of Income Poverty in Albania, University of Essex: ISER Working Paper no. 2006-35.

[21] Schaeffer, N. C. and Charng, H-W (1991). Two experiments in simplifying response categories: intensity ratings and behavioral frequencies, *Sociological Perspectives* **34**, 165-182.

[22] Torelli, N. and Trivellato, U. (1993). Modelling inaccuracies in job-search duration data, *Journal of Econometrics* **59**, 185-211.

Appendix: Additional tables

Table A1 Parameter estimates for electricity model
Multinomial logit for response mode

	Response mode (baseline = unrounded)		
Parameter	**Weekly**	**Monthly**	**Annual rounded**
Autoregressive coefficients γ_{jk}: no change of respondent			
Weekly	2.193	0.348	0.482
	(0.188)	(0.193)	(0.175)
Monthly	0.571	1.027	0.754
	(0.184)	(0.079)	(0.085)
Annual	0.680	0.679	1.499
	(0.174)	(0.083)	(0.062)
Autoregressive coefficients γ_{jk}: new respondent			
Weekly	1.236	0.216	0.243
	(0.501)	(0.431)	(0.382)
Monthly	0.195	0.786	0.670
	(0.437)	(0.204)	(0.197)
Annual	0.441	0.658	1.255
	(0.378)	(0.204)	(0.166)
Respondent characteristics β_j			
Intercept	-1.421	-0.065	-0.041
	(0.964)	(0.129)	(0.062)
H/hold size-2	-0.066	0.662	-0.035
	(0.499)	(0.295)	(0.053)
(H/hold size-2)^2	0.302	0.458	-0.134
	(0.431)	(0.176)	(0.308)
Homeowner	0.220	0.045	-0.314
	(0.097)	(0.140)	(0.164)
Social tenant	0.137	-0.241	-0.301
	(0.046)	(2.626)	(0.139)
Income p.c.	0.177	0.174	-0.025
	(0.041)	(1.225)	(0.265)
Income p.c.^{2}	-0.031	-0.165	0.153
	(0.031)	(0.997)	(0.139)
Female respondent	-0.011	-0.602	0.214
	(0.015)	(0.678)	(0.114)
Age/10	-0.026	0.094	0.326
	(0.013)	(0.314)	(0.282)
(Age/10)^2	-0.615	0.121	0.188
	(0.292)	(0.248)	(0.152)
Changed respondent	0.306	0.111	0.201
	(0.166)	(0.127)	(0.138)
Random effect: κ_k	-0.343	0.006	0.112
	(0.089)	(0.044)	(0.038)

Table A1 continued Parameter estimates for electricity model
Multinomial logit for initial response mode

	Response mode (baseline = unrounded)		
Parameter	**Weekly**	**Monthly**	**Annual rounded**
	Initial characteristics α_j		
Intercept	-0.670	-0.817	0.449
	(1.359)	(8.161)	(0.228)
H/hold size-2	0.362	-0.980	0.119
	(0.879)	(3.251)	(0.151)
H/hold size-2)^2	0.319	-0.048	0.081
	(0.724)	(2.214)	(0.120)
Homeowner	-0.380	-0.755	0.547
	(0.266)	(3.710)	(0.728)
Social tenant	-0.131	0.088	-0.463
	(0.164)	(0.997)	(0.563)
Income p.c.	0.170	0.408	-0.120
	(0.137)	(0.580)	(0.522)
Income p.c.^{2}	0.003	0.534	-0.138
	(0.045)	(0.426)	(2.706)
Female respondent	0.039	-0.105	-0.073
	(0.032)	(0.208)	(0.743)
Age/10	-0.026	0.085	-0.035
	(0.028)	(0.162)	(0.555)
(Age/10)^2	-0.702	-0.209	-0.051
	(0.503)	(0.500)	(0.524)
	Sample mean characteristics α_j		
Mean h/hold size-2	-0.031	-0.091	0.055
	(0.320)	(0.310)	(0.265)
Mean homeowner	-0.113	0.162	-0.115
	(0.250)	(0.255)	(0.211)
Mean income p.c.	0.590	-0.074	-0.256
	(0.816)	(0.449)	(0.258)
Mean age/10	0.457	-0.029	-0.221
	(0.602)	(0.276)	(0.156)
Mean electric c.h.	-0.326	-0.178	-0.401
	(0.535)	(0.227)	(0.123)
Random effect: κ_{k0}	-0.752	0.083	0.010
	(0.184)	(0.108)	(0.090)

Table A1 continued Parameter estimates for electricity model
Autoregression for consumption

Intercept	5.236	(0.048)
H/hold size-2	0.170	(0.004)
(H/hold size-2)^2	-0.018	(0.001)
Homeowner	0.076	(0.009)
Social tenant	0.001	(0.010)
Income p.c.	-0.086	(0.052)
Income p.c.^{2}	0.067	(0.013)
Age/10	0.204	(0.013)
(Age/10)^2)	-0.204	(0.011)
Electric c.h.	0.211	(0.008)
Gas c.h.	-0.114	(0.006)
Oil c.h.	0.172	(0.008)
c_{it-1}	-0.049	(0.004)
σ_ϵ	0.274	(0.000)
ψ	-0.493	(0.005)
Rounding: inverse attraction parameters A_j^m		
Weekly: non-multiples of £260	1.420	(1.871)
Monthly: non-multiples of £120	1.253	(0.046)
Annual: non-multiples of £100	1.044	(0.021)

Table A2 Parameter estimates for total energy model
Multinomial logit for response mode

	Response mode (baseline = unrounded)		
Parameter	**Weekly**	**Monthly**	**Annual rounded**
Autoregressive coefficients γ_{jk}: no change of respondent			
Weekly	1.344	0.128	-0.059
	(0.198)	(0.194)	(0.189)
Monthly	0.370	0.805	0.658
	(0.201)	(0.087)	(0.090)
Annual	0.354	0.425	1.309
	(0.179)	(0.094)	(0.061)
Autoregressive coefficients γ_{jk}: new respondent			
Weekly	1.487	0.935	1.237
	(0.567)	(0.530)	(0.411)
Monthly	-0.053	1.029	0.635
	(0.586)	(0.205)	(0.200)
Annual	0.172	0.506	1.137
	(0.406)	(0.220)	(0.150)
Respondent characteristics β_j			
Intercept	-1.647	-0.211	-0.029
	(0.942)	(0.122)	(0.067)
H/hold size-2	-1.021	0.745	-0.034
	(0.534)	(0.334)	(0.055)
(H/hold size-2)^2	-0.666	0.029	-0.191
	(0.429)	(0.174)	(0.311)
Homeowner	-0.099	-0.299	-0.100
	(0.085)	(0.137)	(0.172)
Social tenant	0.016	-1.168	-0.094
	(0.051)	(1.611)	(0.141)
Income p.c.	0.102	1.694	0.035
	(0.042)	(1.028)	(0.266)
Income p.c.^{2}	0.031	0.458	-0.007
	(0.029)	(0.539)	(0.146)
Female respondent	0.015	0.046	0.078
	(0.020)	(0.290)	(0.117)
Age/10	-0.015	-0.268	-0.583
	(0.016)	(0.209)	(0.320)
Age^2/10	-0.141	-0.026	0.174
	(0.333)	(0.053)	(0.130)
Changed respondent	-0.062	-0.036	0.216
	(0.159)	(0.126)	(0.113)
Random effect: κ_k	-0.212	0.094	0.120
	(0.081)	(0.043)	(0.035)

Table A2 continued Parameter estimates for total energy model
Multinomial logit for initial response mode

	Response mode (baseline = unrounded)		
Parameter	**Weekly**	**Monthly**	**Annual rounded**
	Initial characteristics α_j		
Intercept	-0.719	-0.521	0.362
	(1.600)	(8.694)	(0.248)
H/hold size-2	-0.739	-0.603	0.016
	(0.940)	(4.389)	(0.162)
(H/hold size-2)^2	-0.228	-0.224	-0.007
	(0.710)	(2.075)	(0.120)
Homeowner	-0.307	-0.104	1.515
	(0.329)	(3.927)	(1.120)
Social tenant	0.034	-0.866	-1.468
	(0.176)	(1.696)	(0.582)
Income p.c.	0.176	0.356	-0.348
	(0.135)	(0.516)	(0.479)
Income p.c.^{2}	0.029	0.866	-0.809
	(0.049)	(0.351)	(3.463)
Female respondent	0.010	0.068	0.849
	(0.035)	(0.217)	(0.840)
Age/10	-0.033	-0.127	0.219
	(0.027)	(0.164)	(0.553)
(Age/10)^2	0.206	-0.656	-0.830
	(0.622)	(0.575)	(0.482)
	Sample mean characteristics α_j		
Mean h/hold size-2	0.171	-0.112	-0.127
	(0.352)	(0.320)	(0.279)
Mean homeowner	-0.050	0.130	-0.016
	(0.228)	(0.257)	(0.213)
Mean income p.c.	-0.197	0.366	-0.415
	(1.146)	(0.509)	(0.282)
Mean age/10	1.707	0.035	-0.014
	(0.643)	(0.288)	(0.168)
Mean electric c.h.	-0.076	-0.172	-0.189
	(0.481)	(0.230)	(0.120)
Random effect: κ_{k0}	-0.812	-0.023	0.314
	(0.232)	(0.107)	(0.082)

Table A2 continued Parameter estimates for total energy model
Autoregression for consumption

Intercept	6.006	(0.030)
H/hold size-2	0.120	(0.002)
(H/hold size-2)^2	-0.012	(0.001)
Homeowner	0.171	(0.007)
Social tenant	0.124	(0.006)
Income p.c.	0.003	(0.025)
Income p.c.^{2}	-0.003	(0.004)
Age/10	0.224	(0.008)
(Age/10)^2	-0.224	(0.007)
Electric c.h.	-0.176	(0.005)
Gas c.h.	-0.020	(0.004)
Oil c.h.	0.081	(0.008)
c_{it-1}	-0.086	(0.004)
σ_ϵ	0.246	(0.000)
ψ	-0.454	(0.003)
Rounding: inverse attraction parameters A_j^m		
Weekly: non-multiples of £260	1.501	(4.702)
Monthly: non-multiples of £120	0.881	(0.074)
Annual: non-multiples of £100	1.205	(0.082)