

An Analysis of Consumer Panel Data

IFS Working Paper W09/09

Andrew Leicester
Zoë Oldfield

An analysis of consumer panel data

Andrew Leicester[†] and Zoë Oldfield[‡]

[†]Institute for Fiscal Studies and University College London

[‡]Institute for Fiscal Studies

Abstract

In terms of collecting comprehensive panel expenditure data, there are trade-offs to be made in terms of the demands imposed on respondents and the level of detail and spending coverage collected. Existing comprehensive spending data tends to be cross-sectional whilst panel studies include only limited expenditure questions that record spending only as broad aggregates. More recently, economists have begun to use spending information collected by market research companies that records very detailed spending down to the barcode level from a panel of households, usually recorded by in-home barcode scanners, which may provide considerable advantages over existing data more commonly used in social sciences. However, there has not been a comprehensive assessment of the strengths and weaknesses of this kind of data collection method and the potential implications of survey mode on the recorded data.

This paper seeks to address this, by an in-depth examination of scanner data from one company, Taylor Nelson Sofres (TNS), on grocery purchases over a five-year period. We assess how far the ongoing demands of participation inherent in this kind of survey lead to ‘fatigue’ in respondents’ recording of their spending and compare the demographic representativeness of the data to the well-established Expenditure and Food Survey (EFS), constructing weights for the TNS that account for observed demographic differences. We also look at demographic transitions, comparing the panel aspect of the TNS to the British Household Panel Study (BHPS). We examine in detail the expenditure data in the TNS and EFS surveys and discuss the implications of this method of data collection for survey attrition. Broadly, we suggest that problems of fatigue and attrition may not be so severe as may be expected, though there are some differences in expenditure levels (and to some extent patterns of spending) that cannot be attributed to demographic or time differences in the two surveys alone and may be suggestive of survey mode effects. Demographic transitions appear to occur less frequently than we might expect which may limit the usefulness of the panel aspect of the data for some applications.

JEL Classifications: C81, C83, C33, C41

Keywords: household panel data, scanner data, expenditure, food, duration models, attrition

Acknowledgments: Financial support from the ESRC under the Survey Design and Measurement Initiative (grant number RES-175-25-0008) and the Centre for the Analysis of Public Policy at IFS (grant number RES-544-28-5001) is gratefully acknowledged. The authors would like to thank Thomas Crossley, Rachel Griffith, Laura Blow, Ian Preston and Lars Nesheim at IFS, and participants of the Survey Design and Measurement Initiative workshop held at IFS in October 2008 for comments and helpful discussions, and to TNS for providing the consumer panel data and many useful conversations about the data. EFS data 2002–2006 is collected by the Office for National Statistics and distributed by the Economic and Social Data Service. Crown copyright material is reproduced with the permission of the Controller of HMSO and the Queen’s Printer for Scotland. Errors remain the sole responsibility of the authors.

1. Introduction

Understanding how consumers make decisions is vital to a wide range of policy issues, from the analysis of indirect taxation to the assessment of competition policy. Historically, work in this area has relied on the UK Expenditure and Food Survey (EFS) and its predecessor the Family Expenditure Survey (FES), which are repeated cross-sections. The main source of social science panel data in the UK is the British Household Panel Study (BHPS); however, this only has limited spending data reported in broad categories. In addition, none of the traditional social science data sources record the prices that households paid for individual items. These data limitations mean that it has not been possible to study a number of important issues related to consumer behaviour.

In recent years, social science researchers have turned to less-traditional commercial sources of data on consumer spending. Market research companies such as Taylor Nelson Sofres (TNS) and AC Nielsen collect consumer panel data which contains very detailed spending for a large panel of households. Data are collected by means of a hand-held scanner and respondents are instructed to scan the bar code of every purchase that they bring into their home. Such data has already been used to some extent in economic research (amongst others Griffith et al (2009a) and Aguiar and Hurst (2007)) and in the non-economics literature (Bucklin and Gupta (1999), Van Heerde et al (2000), and others).

Shortcomings of data such as the EFS are well-known (see, for example, Banks and Johnson, 1998) which means that researchers can either choose appropriate techniques to deal with data limitations or can draw informed conclusions in the light of any potential biases they may throw up. A similar understanding of these new sources of data will be important as they become more widely used. This paper attempts to provide a comprehensive analysis of the expenditure and demographic information in one particular consumer panel data set from TNS, making detailed comparisons to traditional social science surveys, in an attempt to highlight the strengths and weaknesses of this new and potentially extremely exciting pool of data for social science research. In particular, we hope to shed some light on the extent to which the mode of data collection impacts on the data that is recorded, and perhaps to highlight the availability and usefulness of consumer panel data to a wider research community.

There is, to our knowledge, relatively little existing comparative evidence of consumer panel data to other data sources. Garner et al (2003) cite unpublished evidence from Keil (2003) showing that consumer scanner data records considerably less alcohol and tobacco expenditure than in the US Consumer Expenditure Survey, and Robertson et al (2003) use TNS data between 1991 and 2000 to examine household nutrition comparing it to evidence

from the National Food Survey 1997–2000. They suggest the demographic data in the TNS matches well to NFS data. All of these comparisons are of a more limited nature than our study. Einav et al (2008) look at AC Nielsen Homescan data in the US, comparing recorded purchases in the data to those from store records and loyalty card data; they find some errors in how prices are recorded (though the Nielsen approach is different to the UK TNS approach) but less error in the recording of purchased quantities, though this study does not make comparisons of recorded spending and demographics to other social science data.

Before discussing our results and in particular our comparative findings, it is worth sounding two notes of caution. First, our broad objective with this study is to assess the extent to which the use of scanner technology affects the data that is recorded. It will be typically be very difficult to isolate the effects of the technology (survey mode) from the effects of the survey design, demographics and so on, though we do the best we can to isolate these factors where possible. This means our conclusions about the impact of survey mode are likely to be suggestive, rather than definitive. Further research using properly controlled trials of scanner versus other technologies may well be desirable. It is hoped our findings will provide guidelines for practitioners who wish to use scanner technologies which may help mitigate its shortcomings.

Second, our ambition is not to state that scanner technologies are ‘better’ or ‘worse’ than other technologies (whether paper diaries or even simple recall questions) as a means of collecting spending data. The comparison of scanner data to other data sources is not carried out on the basis that, say, the EFS spending patterns are the objective ‘truth’ and that we should expect TNS patterns to match completely. Differences are certainly to be expected and our more modest aim is to try to get a sense of how far we can attribute them to survey mode. However, it is clear that scanner technologies may provide certain advantages over traditional micro-level expenditure data for social science research and so understanding the implications of using this kind of data is extremely important and, we hope, the central contribution of this analysis.

The paper is organised as follows. Section 2 describes the main datasets used in our analysis in some detail. Sections 3 to 6 then present the main results of our analysis, looking at expenditures in the consumer panel data in isolation, then in section 4 comparing the demographic composition of the TNS and EFS data, looking also at demographic transitions in the TNS compared to the British Household Panel Study. Section 5 compares TNS and EFS expenditures in depth, asking what evidence there may be of the impact of survey mode on recorded spending. Section 6 looks at attrition from the consumer panel data in depth, a vital feature of panel studies that it is important to understand in detail. Section 7 pulls together our findings and concludes.

2. Data

This section describes in detail the two datasets that underlie the main focus of our analysis, the TNS consumer panel data and the EFS.

a) Taylor Nelson Sofres (TNS) data

Taylor Nelson Sofres (TNS) are a market research company based in London who operate in a number of countries around the world, both in the collection and analysis of data. The data is part of their UK Worldpanel dataset for the Fast-Moving Consumer Goods (FMCG) sector – in effect, food and groceries. We have data covering the six-year period between November 2001 and November 2007, though for comparability with the period of EFS data that we have, our main focus is on the five full calendar years between 2002 and 2006 inclusive.

The TNS data are collected by means of a handheld scanner sent to all participating households¹. Households scan the barcodes of all the food and grocery items they bring into the home, which records precise details of the product purchased and its characteristics; details of the store of purchase are also recorded. This information is sent electronically to TNS; households also post till receipts from their shopping trips to TNS who match the prices paid to the items, giving a complete record of the trip.

In principle, purchases from all stores should be covered – not just supermarkets, but also local corner and specialist stores, internet purchases, chemists and so on. The data also covers non-barcoded items such as loose fruit and vegetables: households are issued with a booklet containing barcodes relating to various non-barcoded products which they scan and record details of the purchase, much like non-barcoded items are scanned in stores. Participating households are rewarded with loyalty points which can be exchanged for high-street vouchers that are not typically redeemable in stores that would be covered by the Worldpanel data so that their recorded spending is not directly affected by their participation.

Until 2006, the sample size was around 15,000 households at any one time. Households continue to participate in the survey for as long as they wish. From 2006, the sample size was increased to around 25,000 households. New households were issued with a slightly different scanner technology that made it easier to scan and upload information to TNS (in particular, the information could be uploaded via the home PC and did not rely on a landline connection). However, for these households no information on non-barcoded products is recorded at all; despite this, TNS suggest households with the new scanners record typically

¹ Until the early 1990s, data was collected by written paper diaries, much like the current EFS.

around 10–15% more items than those using the older technology². We examine the differences between households using the different technologies in the next section.

The panel is maintained to be broadly representative at all times. Households are recruited from a number of different address sources in Great Britain (not Northern Ireland). Relatively detailed information on household demographics are recorded and re-assessed approximately every nine months,³ though information common to social science research but less common to marketing research, such as income and education data, are not routinely recorded (income data has been collected for a limited sub-sample of households since 2006). Household weights are calculated over rolling periods; the most common weights are adjusted every 4 weeks with non-complying households (e.g. those that drop out, or who appear to be absent for some part of the period) given a zero weight for the period. The weights gross up the sample to the GB household population in the BARB Establishment Survey (BARB is the Broadcasters' Audience Research Board, primarily responsible for calculating television and radio audience figures). The main demographics controlled for in deriving the weights are household size, housewife age, household social class, region of residence and 'lifecycle' which broadly reflects the household composition. TNS deliberately oversample larger households containing multiple shoppers (giving them commensurately lower household weights) and suggest there are particular problems recruiting and retaining the very oldest and youngest households and households in London, which are all accounted for in their weighting. Our analysis in section 4 broadly confirms these claims.

Since 2006, the data has been supplemented with information on the nutritional content of the items purchased in each shopping trip, such as calories, fat, sugar, salt, carbohydrates and so on, as well as the (self-reported) BMI of the main shopper in each household. This data may be extremely valuable for social research but is not the focus of our study here. For a subset of households, 'media' data is also collected, such as newspaper and magazines read, television programmes watched, attitudes towards a number of shopping habits and social issues and so on. Again, we do not make use of this data for our analysis but it provides an important supplement for researchers more generally. An application of the nutritional information can be found in Griffith et al (2009b).

Purchase data are recorded at the level of the individual purchase for each shopping trip, giving details of the product and its characteristics, the price, whether the price was under special offer, the store of purchase, the date of purchase and so on. Not surprisingly, given the scope of the data and size of the ongoing panel, the data are very large. Over the period, around half a million separate products (barcodes) are recorded. In a typical week, there are between 600,000 and one million recorded purchases worth some £1 million or so.

² Source: private discussions with TNS data collectors.

³ In practise, the demographic data appears to be updated less frequently than this; see section 4b.

b) Expenditure and Food Survey (EFS) data

The EFS is an annual, cross-sectional survey that has been carried out every year since 1957. Around 6,500–7,000 households are surveyed each year. There are two major components to the survey – a two-week paper-based diary that records all expenditures, and an interview that collects information on household demographics, incomes and some retrospective information on regular purchases (such as rent, mortgage payments and utility bills) and irregular, expensive purchases (such as durables and holidays). Expenditures are calculated and recorded as household-level weekly averages in a number of relatively disaggregate categories – for food there are around 100 such categories.

EFS households are sampled from across the UK⁴. Households are sampled from randomly selected postcode sectors stratified according to region, car ownership and socio-economic status. For the 2006 data, the response rate was around 55%. We use data that covers the same calendar year period as the TNS, from 2002 to 2006.⁵ Since 1998, weights have been included to account for non-response and for sampling, using information from the 1991 and 2001 Censuses on region, age group and sex to construct the sampling weights. Full details of the data collection and weighting are available in the annual ONS publication *Family Spending*⁶.

c) British Household Panel Study (BHPS) data

The British Household Panel Study is a survey of around 10,000 adults in around 5,000 households, designed to be representative of the British population⁷. The same individuals are interviewed annually although as in any such study, some respondents drop out of the panel over time. Information on a wide range of topics is collected for households and individuals including detailed questions on income, employment, household composition, education and housing.

⁴ Northern Ireland households are sampled on a slightly different basis, and are oversampled relative to their population weight. As NI is not included in the TNS data, we exclude these households from any analysis we perform of the EFS data.

⁵ Because the EFS is carried out on fiscal year basis between 1993/4 and 2005/6 this means information in some calendar years is combined from two different waves of data (for example, data for calendar year 2004 comes from the EFS 2003/4 and 2004/5), though this should not affect representativeness as the data is collected to be representative within quarter as well as within year.

⁶ The publication based on the 2006 data is available at http://www.statistics.gov.uk/downloads/theme_social/Family_Spending_2006/FamilySpending2007_web.pdf - see Appendix B of that report for methodological details.

⁷ Note that since 1997, there have been an number of non-representative booster samples added to the BHPS sample. We do not use these households in our analysis.

3. TNS Expenditures

This section provides an overview of the TNS expenditure data, focusing on the breakdown of expenditure over time across spending category and across household types. Our analysis then turns to examining how expenditures change *within* households over time, to assess the extent to which there is evidence of survey ‘fatigue’ owing to the relatively onerous demands of continued participation.

Our sample covers purchases made in the full calendar years 2002–2006. In each year, we select those households that are observed purchasing some FMCG items in at least four separate weeks⁸, though we do not condition that these weeks should be consecutive. There is no particular reason to make the cut-off four weeks, though it does not restrict the sample very much as few households are observed in the data for less than this length of time (see table 1). However it eliminates households that drop out very quickly after signing up or which record expenditures very sporadically. Note that the substantial increase in sample size from 2006 in our selected sample is larger in terms of households (around 38%) than household-weeks (around 29%). On average, each household is observed for around 32–34 weeks in a given year, conditional on purchasing for at least 4.

TABLE 1
Sample sizes, TNS

<i>Calendar year</i>	<i>No selection</i>	<i>Purchases in min. 4 weeks</i>	<i>Households excluded (%)</i>	<i>Household-week observations</i>
2002	16,972	16,050	5.4%	548,113
2003	17,272	16,415	5.0%	552,903
2004	18,083	16,957	6.2%	574,801
2005	18,327	17,289	5.7%	599,656
2006	26,180	23,776	9.2%	772,538

Note: “No selection” counts the number of households that are recorded making at least one purchase during each calendar year.

The TNS data is recorded at the level of the individual purchase. To make our analysis more tractable, we need to aggregate the purchases into expenditure groups. There are clearly many ways in which this can be done. For both this analysis and the comparison between TNS and EFS expenditures in section 5, we aggregate purchases into Retail Price Index (RPI) expenditure categories for food, beverages and alcohol purchases⁹. There are 31 of these expenditure categories, though for analysis purposes we typically focus on 13 even more aggregated groups. For non-food grocery items such as cleaning materials and medicines, the TNS data were not always complete enough to be confident that we were fully covering the

⁸ We look only at full weeks; households that join the data in the middle of a week are not counted for that first part-week.

⁹ Appendix A gives a full list of the categories used. Annex A of ONS (2007) gives examples of the sorts of products that are contained within these categories. http://www.statistics.gov.uk/elmr/04_07/downloads/ELMR_April07_Wingfield.pdf

range of products that are included in the RPI grouping, and so for this paper we ignore non-food groceries (though include beverages and alcohol).

a) TNS expenditure across households

Our analysis of expenditure levels in the TNS data in this section takes weekly-level expenditures on each RPI category and simply assumes that each weekly observation is an independent observation when constructing average spending data. Note, too, that we look only at weeks in which households record some food, beverage or alcohol purchases; households observed spending nothing in a particular week are excluded from these averages in that week even though in some cases it may be entirely legitimate that they spent nothing at all on food and drink¹⁰. In this section, we focus on unweighted data; we return to a discussion of weighting in our comparison between TNS and EFS demographic data in section 4.

Between 2002 and 2006, average weekly household expenditures on food, beverages and alcohol in the TNS for households that reported at least four weeks of positive spending in a given calendar year rose from £39.41 to £44.56, or by just over 13% in nominal terms. Households spent the largest part of their budget on meat, which accounted for 19.8% of spending in 2006. Bread, cereals and biscuits accounted for 14.6%, dairy products 11.8% and vegetables (excluding potatoes) 8.5%. Alcohol made up 9.9% of spending in 2006. Table 2 details expenditure levels in each year, the change between 2002 and 2006 and the budget shares for 2002 and 2006. A more detailed breakdown of these expenditures is given in Appendix A.

Over the reasonably short time period considered, expenditures rose in nominal terms for all of the broad categories, with increases ranging from 4.1% for hot beverages (from £0.97 per week to £1.01) to 23.5% on alcohol (from £3.58 to £4.42). Budget shares remained relatively stable over the period: meat spending fell from 21.0% of the total in 2002 to 19.8% in 2006 whilst the largest increase was for alcohol, from 9.1% to 9.9%. Amongst food items, the largest increase came in dairy products, from 11.3% to 11.8%.

¹⁰Although zero spending in any week for a given household may be legitimate, there is a much higher percentage of households in the TNS recording zero expenditure than in the EFS (see section 4a). It is unlikely that this many household genuinely buy no food or drink at all and so assuming that all 'missing weeks' should be treated as missing rather than zero is probably more reasonable than including them.

TABLE 2
TNS expenditure patterns, 2002–2006

Category	Expenditures (£/week)					Change (2002- 2006)	Share of total (2002)	Share of total (2006)
	2002	2003	2004	2005	2006			
Bread, cereals & biscuits	5.91	6.00	6.16	6.21	6.51	+10.2%	15.0%	14.6%
Meat	8.28	8.56	8.76	8.65	8.82	+6.5%	21.0%	19.8%
Fish	1.43	1.48	1.52	1.56	1.65	+15.4%	3.6%	3.7%
Oils and fats	0.81	0.81	0.82	0.83	0.89	+9.9%	2.1%	2.0%
Dairy	4.45	4.64	4.88	5.10	5.24	+17.8%	11.3%	11.8%
Hot beverages	0.97	0.95	0.94	0.94	1.01	+4.1%	2.5%	2.3%
Soft drinks	2.16	2.30	2.28	2.29	2.58	+19.4%	5.5%	5.8%
Sugar & chocolate	1.85	1.93	2.02	2.02	2.15	+16.2%	4.7%	4.8%
Potatoes	1.64	1.56	1.61	1.66	1.75	+6.7%	4.2%	3.9%
Other vegetables	3.26	3.40	3.39	3.76	3.80	+16.6%	8.3%	8.5%
Fruit	2.42	2.53	2.62	2.83	2.86	+18.2%	6.1%	6.4%
Other food	2.64	2.69	2.71	2.72	2.86	+8.3%	6.7%	6.4%
Alcohol	3.58	3.73	3.91	3.93	4.42	+23.5%	9.1%	9.9%
Total	39.41	40.61	41.63	42.52	44.56	+13.1%	—	—

Note: Figures are unadjusted for price changes over the period and are unweighted. Figures are average weekly expenditures for households that purchase in at least 4 separate weeks which commence within the calendar year. Each weekly observation for each household is treated as an independent observation.

Figure 1 shows how expenditure aggregates varied in 2006 according to the composition of the household. Larger households with more adults or children spent more. Households made up of two non-pensioner adults¹¹ spent around 65% more per week on food, on average, than single adult households suggesting some economies of scale in food purchases. Interestingly, two adult households spent around 64% more than single adults on beverages (hot beverages and soft drinks) but 93% more on alcohol, suggesting much smaller scale economies. Pensioner households tend to spend less than non-pensioner households but again with differences over spending categories: single pensioners spent on average around 19% less per week on beverages than single non-pensioners and 12% less on alcohol, but only 2% less on food. Though overall food spending was similar for pensioner and non-pensioner households, the composition of foods bought was very different. Single pensioners spent 17% more on fish, 27% more on oils and fats and 25% more on fruit than single non-pensioners, but 15% less on potatoes, 12% less on other vegetables and 5% less on both cereals and meats.

Adding children to a household increases expenditure on food (households made up of two adults with children spend 27% more than two adults without) and beverages (29% more) but reduces spending on alcohol (around 28% less)¹². Again, significant variation across


¹¹ It is not always easy to discern familial relationships from the information about individual household members in the TNS. We have therefore chosen a household structure breakdown that does not make any particular assumptions about the nature of the family relationship and instead is based on numbers of adults and children, and the ages of the adults.

¹² In all cases, spending refers to items brought into the home for domestic consumption or preparation, excluding things like takeaway foods. We do not observe expenditure on food outside the home (including takeaways and restaurant meals, as well as alcohol in pubs and bars) which is clearly a close substitute for many

expenditure categories emerges – two adult households with children spent 50% more than two adult childless households on soft drinks but 15% less on hot beverages, 52% more on bread and cereals, 41% more on dairy, 50% more on sweets and chocolates, 50% more on potatoes but only 6% more on other vegetables, and only 1% more on both fish and fats.

FIGURE 1

Total expenditures on food, beverages and alcohol in TNS, 2006, by household composition



Notes: Figures are average weekly expenditures for households that purchase in at least 4 separate weeks which commence within the calendar year. Each weekly observation for each household is treated as an independent observation.

b) TNS expenditure changes within household

We now examine how spending changes over time for households that participate in the survey. One of the key advantages of this data is that as a panel, we can use it to follow household expenditures over time and assess how households respond to shocks and changes to their circumstances. However, one fear with this kind of data is that households who participate may eventually tire of doing so. This may manifest itself in several ways – scanning only the main shopping trips and not top-up trips, being less diligent in scanning non-barcoded items, scanning purchases for some but not all household members and so on. We call this general phenomenon ‘survey fatigue’.

Eventually, once households become completely fatigued, they will drop out or attrit from the survey; we examine this in section 6. Note, though, that fatigue and attrition are not of course unique to consumer panel data. Even in one-off spending surveys where participation is limited to two weeks, there is evidence that spending drops off in the second week¹³.

households. One example is in alcohol spending – grouping households according to the age of the household head, the group spending the *least* on alcohol in the TNS are the young aged under 30; this may be because this group spends more on alcohol outside the home.

¹³ See Ahmed at al (2006) who look at this issue in the Canadian Food expenditure survey and Central Statistical Office (1985) who look at the issue in the Family Expenditure Survey.

However, the nature and extent of attrition and fatigue in the TNS is likely to be different given the effort that is required of on-going participation

In general, all else unchanged, we would expect that households would record higher amounts of spending over time owing to grocery price inflation. If expenditures fall, this might be attributable to fatigue. To examine this, we take the sample of households that we observe in the TNS that *first* begin to scan their grocery purchases some time between November 2001 and the end of 2005 (this allows us to ensure all participating households have the opportunity to take part for at least one full year, and that we only use households that we actually observe beginning their purchases during the data period we have). For these households, we record their total spending in each week (conditional on spending something) and the number of weeks since they began active participation in the survey. Using this, plus information on household characteristics, we can examine the relationship between household spending and duration of participation using a simple regression analysis.

Our model takes the following form. Define x_t^h as the total expenditure (on food, beverages and alcohol) of household h in week t , and w_t^h as the number of full weeks since household h began to participate in the survey (thus $w_t^h=1$ will be the first full week of active participation; households that begin scanning purchases in the middle of a week will not be included until the next week of data). Then we perform a household fixed-effects regression of the form:

$$\ln x_t^h = \alpha^h + \beta w_t^h + \gamma X_t^h + \varepsilon_t^h$$

Where X_t^h is a vector of household characteristics at time t (that can vary over time as constant household characteristics will be eliminated by the fixed effect) and ε_t^h is an unexplained residual. The coefficients of interest are the vector β . As expenditure is expressed in logs, these vectors represent (approximate) percentage deviations from the expenditure in the base week for each week in which we observe households (in this case, the first full week of participation). We restrict attention to 100 weeks of participation (setting $w_t^h=100$ for any weeks greater than 100).

Note that we use a fixed-effects regression rather than assuming a common constant for all households. It may be that certain households that are pre-disposed to have higher expenditures are also more likely to participate in the survey for a long time and that observable household characteristics cannot adequately capture this (think of some unobserved household characteristic like ‘diligence’ or ‘commitment to participation’). This would tend to bias upwards the coefficients on w_t^h . Estimates of the model that exclude household fixed effects suggest this bias may be quite substantial (see Appendix B).

Figure 2 shows the results of our regression where we plot the β coefficients on ‘weeks since starting to scan’ from the fixed effects model. The model also controls for the year and month of purchase and a dummy variable indicating whether the household is within 4 weeks of the date it drops out of the survey¹⁴. The thin lines simply plot the estimated coefficients; there is some volatility week to week though the magnitude of this variation is quite small. The dotted lines are 95% confidence intervals around the point estimates and the thicker lines are a cubic fitted to the individual coefficients to smooth out the variation and make trends more visible¹⁵.

Our results suggest that around 6 months after beginning to participate, households spend on average some 5% less than they recorded in their first week of participation; this 5% limit appears quite consistent though there is some evidence of a further decline towards the end of the second year of participation. Generally, these declines are statistically significant.


There appear to be some differences across different types of household in the extent to which this result is observed. We run separate fixed-effects regressions according to the composition of the household at the time they are first observed. Figure 3 shows the results comparing, in the top panel, households made up of a couple with children compared to a couple without children; and in the bottom panel households consisting of a single pensioner and a single non-pensioner. For visual clarity we omit the confidence intervals from these figures and show only the coefficients and the fitted cubic; note, though, that the smaller sample sizes of individual household types means that the coefficients are often insignificantly different from zero.

¹⁴ We do not include any other observed time-varying demographics. Comparisons of the TNS data to other panel studies show there are very few households observed making demographic transitions over time; see section 4b.

¹⁵ Full results for all models discussed in this section are available from the authors.

FIGURE 2

Survey 'fatigue' results: random- and fixed-effects models



Notes: Thin lines plot the coefficient from a regression of log expenditure on weeks since beginning participation; dotted lines are 95% confidence intervals and thick lines fit a cubic to the coefficients. Results include households that begin participating between November 2001 and December 2005.

It is clear there are some interesting differences in the patterns across household types: households with children see their spending fall more quickly but there appears to be no sustained evidence of 'fatigue' – spending levels are quickly and consistently around 5% lower than in the first week. For households with children, there is a more sustained and marked decline in spending up to around one year of participation. On average, it takes around 20 weeks for spending to fall below its week one level, but the decline is more sustained: after around a year, spending for childless households is even further below week one values than is the case for households with children. Comparing by age shows an even more marked difference – there is no evidence at all of fatigue for pensioner households, where spending is rarely below week one levels on average even after up to two years of participation. For non-pensioner households, spending levels tend to be below week one levels almost immediately, though again there is no real evidence of continued 'fatigue' as the length of participation increases.

There are interesting differences not only across households but also across goods. Figure 4 repeats our fixed effects regression but replaces the left-hand variable with spending on different items rather than total expenditure (and thus conditions on having purchased the item in the first full week of participation). We might expect to see the greatest 'fatigue' for items which are liable to feature in top-up expenditures (bread, milk, sweets, say) and perhaps less fatigue for products that feature regularly in standard weekly shops. The results below (here we plot for the greatest clarity only the fitted cubics over the regression coefficients)


suggest this may be true – for items like fish and fruit, say, which are unlikely to be purchased as top up items, the fatigue level is low, with spending typically only around 2% lower than recorded in week one on average. For breads and cereals, the results look quite similar to the average effects found in figure 2. For sweets and chocolates, the drop in spending is much quicker and more severe; spending is around 8% lower after 6 months or so. Intriguingly, a very similar pattern is found for alcohol spending. There could be many possible explanations for these patterns and a more detailed analysis of fatigue that explores these possibilities would be extremely interesting. For example, the results here are consistent not only with a hypothesis that there is spending fatigue, but also that households taking part are suddenly faced with a record of their expenditures and seek consciously or unconsciously to change (reduce) their spending as a result of having to carefully monitor what they buy. We cannot disentangle these effects using observed data but experimental evidence may be able to do so.

Interestingly, there appears to be little relationship between the extent of survey fatigue and baseline employment status of the household head; we may expect, for example, that households that are employed have less time and may see faster and more persistent fatigue levels than those households out of work (unemployed or not seeking employment, though excluding retired households). There appears, however, to be little difference between them (results are available on request); if anything, it is the non-employed households that exhibit a more rapid decline in their spending. Around two years after beginning participation, unemployed households record around 10% less spending than in week one, compared to 6% less for employed households. Of course this is unconditional on other particular household characteristics such as age, family composition or income.

One key result from these models of survey fatigue is that the number of weeks since beginning participation, time controls and household fixed effects explain very little of the variation in expenditures: R^2 values in all cases are very low, typically less than 1% (though slightly higher in the product by product specification). Thus whilst we can find evidence that time in the survey alone can significantly affect recorded spending, other factors are clearly much more important.


FIGURE 3

Survey 'fatigue' results: by household type



Notes: Thin lines plot the coefficient from a regression of log expenditure on weeks since beginning participation, and thick lines fit a cubic to the coefficients. Results include households that begin participating between November 2001 and December 2005.

FIGURE 4
Survey 'fatigue' results: by product type



Notes: Lines represent a cubic fitted to the weekly-level coefficients. Results include households that begin participating between November 2001 and December 2005.

4. A comparison of TNS household demographics to other surveys

In this section, we examine the extent to which household demographics differ across the TNS and EFS datasets. In addition to comparing demographic composition in cross-section, we also compare demographic transitions to the patterns found in the British Household Panel Study.

a) Cross-sectional demographic comparisons

We begin with a comparison of the demographic make up of the TNS sample to that of the EFS. As discussed in our data description, the EFS is a two-week diary-based record of all expenditures whereas the TNS is an ongoing survey of grocery expenditures. The EFS sample contains one observation for each household that records their spending and demographic information; in the TNS we have multiple observations for each household for each week they are observed.¹⁶

¹⁶ It is not clear what the appropriate unit of observation is in the TNS data for the purposes of comparison with the EFS. If we were to take one observation per household as in the EFS, the resulting sample would be one observation for any household that was ever recruited to the TNS. However, because households are continuously recruited to the TNS in order to keep the demographics broadly representative, households who are more likely to drop out are also more frequently recruited, resulting in a survey population less representative than if all household-week observations are pooled.


To compare the demographics, we use the same sample of household-weeks that were used in the previous section when we looked at TNS expenditures. We make use of the household-level weights provided by TNS and compare unweighted and weighted demographics to weighted data from the EFS. As discussed in section 2a, these TNS weights are constructed on a 4-week basis so that when the weights are applied to the entire sample of households in each 4 week period, the resulting sample is broadly representative of the general GB population with respect to particular demographics. The EFS weights perform a similar task for each annual sample. In order to ensure the broad sampling frame of the two surveys is similar, we exclude households from Northern Ireland from the EFS sample in each year.

In the EFS, we exclude households that do not record any food, beverage or alcohol expenditure in one or both of the two weeks for which they participate. In the TNS, as in section 3 we exclude households only in weeks in which they observed to spend nothing. There is a relatively large difference between the surveys is in how often households are observed spending nothing in a given week. In the EFS, around 4% of households record zero spending on food and drink in one or both weeks of their participation; in the TNS, by contrast, if we look only at the first two full weeks in which households participate, around 15% spend nothing in the second week. It is not obvious why this difference should emerge. The EFS asks households who are away from home to record spending as they would do at home whereas in the TNS households away from home would not be able to participate. This may explain some of the difference; whether it accounts for the entire difference is questionable. It may be that there are more households in the TNS who simply fail to scan any food item that was brought home in a given week. This is less likely to happen in the EFS since households which have nil spending are subject to checks to ensure that such behaviour is genuine and not simply a failure to keep a diary. More analysis would be needed to understand the contribution of these two factors in explaining the difference in the proportion of households with zero spending on food.

There are a number of variables along which significant differences between the EFS and TNS samples emerge which are largely accounted for by weighting the TNS data. For example, figure 5 shows the number of adults in the household in the two surveys for 2006; the hashed bar is the weighted TNS figure. Relative to EFS, the TNS data contains far more households with multiple adults and far fewer single-adult households (22.5% of the 2006 sample in the unweighted TNS data are single adult compared to 32.5% of the EFS); weighting substantially eliminates the differences. This confirms TNS's assertion that they oversample multiple adult-shopper households where it is difficult to get adults other than the main shopper to comply but that this is accounted for by the weighting scheme. A similar

analysis for the number of children finds smaller differences; around 70% of observations in both surveys are from childless households.


FIGURE 5
Number of adults, TNS and EFS (2006)



Notes: TNS data are household-week observations for households recording spending for at least 4 weeks of calendar year 2006; EFS figures are household-level observations from EFS 2006, conditional on households spending something in both weeks for which they are observed and excluding Northern Ireland. TNS weights are 4 week weights as described in the data section.

If we look at household composition more generally (see figure 6), we find that some but not all of the main differences across datasets are accounted for by weighting. Household composition takes account of the numbers and ages of adults and the presence of children (household members under age 18) but not, for example, marital status which is recorded differently in the two surveys. The aim is to take a relatively agnostic view about the relationship of household members to one another since this cannot be inferred in the TNS (unlike the EFS which codes explicitly how members are related). Unsurprisingly, given the findings on adult numbers, the TNS records fewer single-adult households (both pensioners and non-pensioners) than the EFS but the differences are largely eliminated by weighting. The TNS records far more households comprising two adults with children, but again this difference disappears with weighting. Where weighting is less successful is for single adults with children, who remain slightly underrepresented in the TNS relative to EFS, and in households made up of multiple adults with children. On average, however, the weighted TNS sample much more closely resembles the EFS sample than the unweighted sample along this dimension.

FIGURE 6
Household composition, TNS and EFS (2006)




Notes: TNS data are household-week observations for households recording spending for at least 4 weeks of calendar year 2006; EFS figures are household-level observations from EFS 2006, conditional on households spending something in both weeks for which they are observed and excluding Northern Ireland. TNS weights are 4 week weights as described in the data section.

Another characteristic for which weighting is only partially successful at explaining the differences across datasets is age; here, we focus on the age of the *eldest* person in the household¹⁷. Unweighted, the 2006 TNS sample appears to contain far fewer observations where the eldest household member is under 25 (only 0.5% of observations compared to 2.7% in the EFS) and over 80 (3.8% compared to 8.1%). By contrast, there appear to be many more observations where the eldest household member is 50 – 74 in the TNS (around 48.6% of observations) than EFS (40.7%). Whilst in some cases, weighting reduces the differences (for example, ages 50-64, 25-29 and 35-39) in other cases the differences are reduced only marginally (the very youngest and oldest households) and in others accentuated (ages 40-49 and 65-79). Broadly, the weighted data looks comparable to EFS averages for households of prime child-rearing age and beyond (ages 25-64), with a noticeable dip in the 40s, seems to contain more households of “younger old age” (65-79) and far fewer very young and very elderly households.

¹⁷ Social science surveys more commonly focus on the “head of household” or “household reference person”. Though TNS data includes a flag for the head of household, it is not clear what the variable means in this case and how comparable it is to the EFS definition. TNS themselves tend to focus on the “main shopper” as their key contact in a household (the main shopper is much more likely to be female than the head of household). We use eldest person as it is clearly unambiguously comparably defined across surveys. TNS report their weights are based on the age of the main shopper whilst the EFS weights are based on the age of the head, so it is not necessarily surprising that weighting does not remove the differences more precisely across datasets.

FIGURE 7

Age of eldest household member, TNS and EFS (2006)




Notes: TNS data are household-week observations for households recording spending for at least 4 weeks of calendar year 2006; EFS figures are household-level observations from EFS 2006, conditional on households spending something in both weeks for which they are observed and excluding Northern Ireland. TNS weights are 4 week weights as described in the data section.

There are a number of other comparisons we could make; rather than presenting all as charts we make some general observations. Geographically, the spread of households in the two surveys is quite similar with one particular exception – there appear to be fewer households in London in the TNS than EFS. After weighting, London households make up 12.1% of the EFS sample in 2006 compared to 10.9% of the TNS sample (and only 9% in the unweighted sample). In terms of the employment status of the household head (subject to the caveat of comparably defining who the head is across the surveys, see footnote 15), the datasets are comparable; in 2005 (the variable is not present for all TNS households in 2006), there are slightly fewer retired TNS household observations, consistent with the findings on the very elderly in figure 7, though weighting reduces the differences. There are also more part-time workers (8-29 hours) in the TNS than the EFS (8.8% of the weighted TNS sample compared to 4.6% of the EFS sample).

The final comparison we make in detail concerns household incomes. In 2006, a subset of the TNS sample were asked about their gross household income (before taxes), with responses banded into one of 8 bands of £10,000 width. Calculating a similar figure from the detailed income information in the 2006 EFS, and making comparisons only for households that report a non-missing income observation in 2006, the distribution of incomes is shown in figure 8 below. Strikingly, incomes appear substantially lower in the TNS than in the EFS. In the weighted data, 69% of TNS observations have gross household incomes below £30,000,

compared to only 56% of EFS households; only 5.3% of TNS observations record income above £60,000 compared to 13.2% of EFS households.

FIGURE 8
Gross household incomes, TNS and EFS (2006)



Notes: TNS data are household-week observations for households recording spending for at least 4 weeks of calendar year 2006; EFS figures are household-level observations from EFS 2006, conditional on households spending something in both weeks for which they are observed and excluding Northern Ireland. TNS weights are 4 week weights as described in the data section.

If we look at other variables that may proxy income, it seems hard to replicate this result – for example, TNS households are far less likely not to own a car (19.3% TNS versus 24.1% EFS, which may also be related to the lower share of London households) and are more likely to be owner-occupiers rather than social or private renters (77.2% of TNS household observations are owner-occupiers, compared to 70.0% of EFS households), and, as discussed, the employment status variables tend to line up reasonably well. One possibility is that the sample of households for whom income was recorded (around 72% of household-weeks observed in 2006) was drawn more heavily from the lower-income sample and that the overall sample may have a more similar distribution of income; another is that the TNS tends to sample lower-income people who are in work, and who are also homeowners, say. The other possibility is that income is recorded with less accuracy in the TNS or at least on a basis that is not comparable with the EFS. Our measure of income in the EFS is derived from a set of detailed questions about different components of income for each family member. In contrast, the measure of income recorded in the TNS is simply a one-shot self-reported question that asks one person to report the entire household gross income.

Clearly there are a number of dimensions along which demographics in the two datasets are different, even using household weights as provided in both surveys. However our analysis so far has been to look at the demographic characteristics one by one. Clearly, some

of the differences will be correlated – for example, the relative undersampling of very elderly and very young households in the TNS compared to the EFS will help explain why there are fewer single-person households. In the next section, we will compare expenditure in the TNS and in the EFS. In order to eliminate observable demographic differences as a explanation for any differences in expenditure, we will use propensity score weights (based on Rosenbaum and Rubin, 1983).

To calculate propensity weights, we first pool the households from both surveys into a single dataset, defining demographics in each on a common basis. We then estimate a logit model in which the dependent variable is a binary variable taking the value 1 if the observation comes from the EFS and 0 if it comes from the TNS data and the independent variables are common demographics to both datasets (we use region, age and gender of the oldest household member, number of adults, number of children, head of household employment status, housing tenure, household composition, number of cars, whether there is a computer in the household, month of observation, and, for 2006 only, income; we also interact age and gender). The model is estimated separately for each year of data. From the model, we predict $e(z_i)$, the conditional probability that that observation i with characteristics z_i would be observed in the EFS. The propensity score weight P_i is then defined as:

$$P_i = \frac{e(z_i)}{1 - e(z_i)} \times \frac{\Pr(TNS)}{\Pr(EFS)}$$

Where $Pr(TNS)$ is the share of observations in the pooled sample which originate from the TNS (and similar for $Pr(EFS)$). The propensity score weight essentially applies a larger weight to those observations which “look like” observations from the EFS and a smaller weight to those observations which have a higher probability of being observed in the TNS. These weights will be used in our expenditure comparisons below – as far as possible, they ensure the weighted TNS sample is as demographically comparable to the EFS sample in each year as possible and so using these weights allows us to assess how far expenditure differences between the two surveys may be accounted for by their demographic makeup.

As well as estimating the models separately by year, we also estimate the models separately for sub-groups of the TNS sample which have missing characteristics due to the TNS survey design. So for example, in 2005, 15% of the observed household-weeks in the TNS sample have no recorded housing tenure (but none of the EFS sample do); we estimate the model separately for those with and without tenure, excluding tenure as an explanatory variable where it is missing. We carry out a similar process in 2006 where approximately 28% the sample have missing income (note that income is not present in any other year) and

7% have missing employment status. In addition, in 2006, we estimate the model separately for households with the new scanner technology (see section 2a). Since by design, all households using the new kit type own a computer, we exclude this variable from the model for this group. The results of our model for 2005 are contained in Appendix C¹⁸.

b) Demographic transitions

One of the strengths and appealing properties of the TNS data is that there are repeated observations on the same household. Panel data on detailed expenditure is not currently available elsewhere in the UK. The value of the longitudinal data on expenditure would be enhanced if we could associate changes in spending with changes in demographics. For example, one area which has attracted much research is the question about what happens to spending around retirement – the so-called retirement saving puzzle (see for example Banks, Blundell and Tanner (1998)). Such research has traditionally relied on repeated cross-sectional data to make pseudo-panels. Other examples of areas of research interest might include what happens to expenditure in the event of job loss, or when people marry or have children.

TNS suggest that demographics are updated approximately every nine months although the effort that they put in to get these updates varies across different demographics depending on their own interests and those of their clients. For this reason, it is important that any researcher using the demographic data to look at transitions is aware of whether demographic changes are captured adequately. This section examines two variables of particular interest – work status of the head of household and number of children present in the household – and compares recorded transitions in the TNS data to those recorded in the BHPS.

The TNS is a panel of households where the household is defined around the main shopper. So the main shopper is the person who is followed and other people in the household may move in or out of the household. The BHPS is a panel of individuals and all individuals are followed over time. To try to replicate the household definition in the TNS, we take all heads of households¹⁹ in 2002 (the first full year of TNS data) and follow them over time until 2006. The household characteristics will depend on who the head is living with in each year. For the TNS, observations on demographics are recorded once for every household in each year that they are observed to record spending. However, since we are looking at annual transitions, households have to be observed in at least two separate years to enter our transition tables.

¹⁸ Full results for other years are available on request. A separate model is run for the TNS household sample made up of households in their first two full weeks of participation; we also use these estimates in our expenditure comparisons below.

¹⁹ As we do not have a main shopper definition in the BHPS, we define the household around the head of household.

Table 3 reports annual transitions in the number of children in the BHPS and the TNS. The first three rows take all childless households in year t and reports the percentage of those who are still observed to be childless at time $t+1$ and the percentage who are observed to have had at least one child. The numbers show that the percentage of childless households who transition into a state with children is around twice as high in the BHPS as in the TNS (2.1% compared to 0.9%). A similar result is found for the probability that households who have children at time t are not observed to have children at time $t+1$: 6.7% in the BHPS and only 2.0% in the TNS.

As TNS households can only be observed transiting at all if they participate for at least one year, our TNS sample is highly selected and not necessarily representative. To get around this problem, the third set of numbers in Table 3 takes a more homogenous group of households: childless couples where the head is aged less than 35. These household are also in a group where we might expect to see a higher proportion of transitions as couples form families. However, in the BHPS, the percentage of childless couples aged under 35 at time t who have children at $t+1$ is 12.1%, compared to only 6.2% in the TNS. A similar story unfolds if we take older households (those where the head is aged 50 or over) who had children at time t . Of those households, 18.6% were observed to be childless at time $t+1$ in the BHPS compared to only 11.3% in the TNS.

Table 4 shows similar numbers as in Table 3 but for transitions into and out of work. The first set of numbers takes all households where the head is observed to be working at time t . Of those, 7.0% of heads are not working at time $t+1$ compared to a small percentage in the TNS (1.2%). The second set of figures compares households where the head is not working at time t : of these, in the BHPS, 9.4% are working at time $t+1$ but only 1.4% are observed to be working in the TNS. The third set of numbers shows work transitions for a more homogenous group of households in both surveys: those aged 50 or over who are working at time t . Of this group, 11.4% are observed to leave work in the BHPS compared to 2.9% in the TNS.

TABLE 3
Annual transitions in the number of children in the BHPS and TNS

	<i>BHPS</i>	<i>TNS</i>
All Childless households at time t		
No children at time $t+1$	97.9%	99.1%
Has children at time $t+1$	2.1%	0.9%
All Households with children at time t		
No children at time $t+1$	6.7%	2.0%
Has children at time $t+1$	93.3%	98.0%
Childless couples, head aged <35 at time t		
No children at time $t+1$	87.9%	93.8%
Has children at time $t+1$	12.1%	6.2%
Households with children, head aged ≥ 50 at time t		
No children at time $t+1$	18.6%	11.3%
Has children at time $t+1$	81.5%	88.8%

Notes: A couple is defined as a household with just two adults who may or may not be married or co-habiting

TABLE 4
Annual transitions in the work status in the BHPS and TNS

	<i>BHPS</i>	<i>TNS</i>
Head is working at time t		
Head not working at time $t+1$	7.0%	1.2%
Head working at time $t+1$	93.0%	98.8%
Head is not working at time t		
Head not working at time $t+1$	90.6%	98.6%
Head working at time $t+1$	9.4%	1.4%
Head aged ≥ 50 or over and is working at time t		
Not working at time $t+1$	11.4%	2.9%
Working at time $t+1$	88.6%	97.1%

At least in these dimensions, it appears that changes in demographics are not well recorded in the TNS, with employment status transitions particularly poorly accounted for. There could be a number of reasons for this. One reason might be that those households who are most likely to go through transitions are least likely to be recruited to the TNS. Another reason might be that those households who are most likely to go through a transition are those that are most likely to drop out of the sample within one year (or perhaps even drop out as a result of the transition itself). The final, most straightforward, reason is that the demographic data is not accurately updated. Untangling all these effects would be difficult and the full explanation is likely to be a combination of all three. However, at least for work status, the very small number of transitions observed suggests that a large part of the difference is due to the demographic data not being updated consistently.

5. Expenditure comparisons between TNS and EFS

This section contains a more detailed focus on expenditures, comparing spending patterns observed in the TNS data to those we see in the EFS and trying to assess how far the different modes of survey design affect the spending patterns observed. Section 3a briefly examined the TNS spending information in its own right; here we make much more detailed comparisons of TNS and EFS data across goods and households.

a) The distribution of total expenditures in the TNS and EFS

We begin by broadly comparing expenditure patterns in the two surveys, focusing as in section 3a on food, beverage and alcohol expenditures for 13 aggregate expenditure groups based on RPI definitions. We use the same household sample in the TNS as in sections 3 and 4, households that are observed purchasing in at least 4 separate weeks of a given calendar year. Our EFS sample replicates that used in section 4a: households in Great Britain that purchase some food or alcohol in both of the weeks for which they participate. All EFS expenditure data are weighted by sampling weights included in the survey.

Table 5 compares average (mean) weekly expenditures for 2005 in the two surveys²⁰, contrasting the results for the TNS both before and after we apply our propensity score weights to account for demographic differences (note that EFS data is always weighted by sampling weights provided as part of the data); table 6 presents the same data for each category as a share of total spending rather than as an absolute spending level. Two things are immediate from these comparisons: spending levels in TNS are lower than those in the EFS for all spending categories, and this difference is not driven by observable demographic differences between the surveys. TNS spending levels are around 80% as high as EFS levels in total if we look at unweighted TNS figures; once we correct for demographics, this falls to 75%. Thus, relative to EFS, demographic corrections actually reduce the level of TNS spending. This should not be a surprise, however, since we know the TNS oversamples households with children and households with more than one adult, and undersamples single-person households. Thus the TNS has a larger sample of larger households who, all else equal, will spend more. Correcting for this reduces TNS expenditure levels.

Looking across goods, there are some considerable differences in spending levels. Spending on all goods is lower than that recorded in the EFS and in all cases, demographic weighting accentuates the differences. The gap (based on weighted figures) ranges from 14% for oils and fats to 30% for fruit and 42% for alcohol. This is despite the fact that alcohol is

²⁰ The comparisons are not much different in other years. We choose 2005 rather than 2006 since there are no issues with households using different technologies in the earlier year – see section 5d for more on the impact of the particular scanner technology employed.

known to be substantially under-recorded (relative to, for example, the levels of alcohol spending recorded in the UK National Accounts) in the EFS²¹.

TABLE 5
EFS/TNS expenditure comparisons (mean spending levels), 2005

<i>Category</i>	<i>EFS</i> (£/week)	<i>TNS</i> (£/week, <i>unweighted</i>)	<i>TNS</i> (£/week, <i>weighted</i>)	<i>TNS/EFS</i> (%, <i>unweighted</i>)	<i>TNS/EFS</i> (%, <i>weighted</i>)
Bread, cereals and biscuits	7.39	6.21	5.74	84.0%	77.7%
Meat	10.43	8.65	8.11	82.9%	77.8%
Fish	2.13	1.56	1.51	73.2%	70.9%
Butter, oils and fats	0.92	0.83	0.79	90.2%	85.9%
Eggs, milk, cheese	5.96	5.10	4.72	85.6%	79.2%
Hot beverages	1.06	0.94	0.89	88.7%	84.0%
Soft drinks	2.77	2.29	2.11	82.7%	76.2%
Sugar, sweets, chocolates	2.48	2.02	1.92	81.5%	77.4%
Potatoes	1.86	1.66	1.54	89.2%	82.8%
Other vegetables	4.71	3.76	3.51	79.8%	74.5%
Fruit	3.77	2.83	2.63	75.1%	69.8%
Other food	3.04	2.72	2.52	89.5%	82.9%
Alcohol	6.62	3.93	3.81	59.4%	57.6%
Total expenditure	53.15	42.52	39.79	80.0%	74.9%

TABLE 6
EFS/TNS expenditure comparisons (mean budget shares), 2005

<i>Category</i>	<i>EFS</i>	<i>TNS</i> (<i>unweighted</i>)	<i>TNS</i> (<i>weighted</i>)
Bread, cereals and biscuits	13.9%	14.6%	14.4%
Meat	19.6%	20.3%	20.4%
Fish	4.0%	3.7%	3.8%
Butter, oils and fats	1.7%	2.0%	2.0%
Eggs, milk, cheese	11.2%	12.0%	11.9%
Hot beverages	2.0%	2.2%	2.2%
Soft drinks	5.2%	5.4%	5.3%
Sugar, sweets, chocolates	4.7%	4.8%	4.8%
Potatoes	3.5%	3.9%	3.9%
Other vegetables	8.9%	8.8%	8.8%
Fruit	7.1%	6.7%	6.6%
Other food	5.7%	6.4%	6.3%
Alcohol	12.5%	9.2%	9.6%
Total expenditure	100.0%	100.0%	100.0%

Notes: TNS data are household-week observations for households recording spending for at least 4 weeks of calendar year 2005; EFS figures are household-level observations from EFS 2005, conditional on households spending something in both weeks for which they are observed and excluding Northern Ireland. Weighted TNS data use the propensity score weights derived in the previous section. Figures in table 4 may not sum exactly due to rounding.

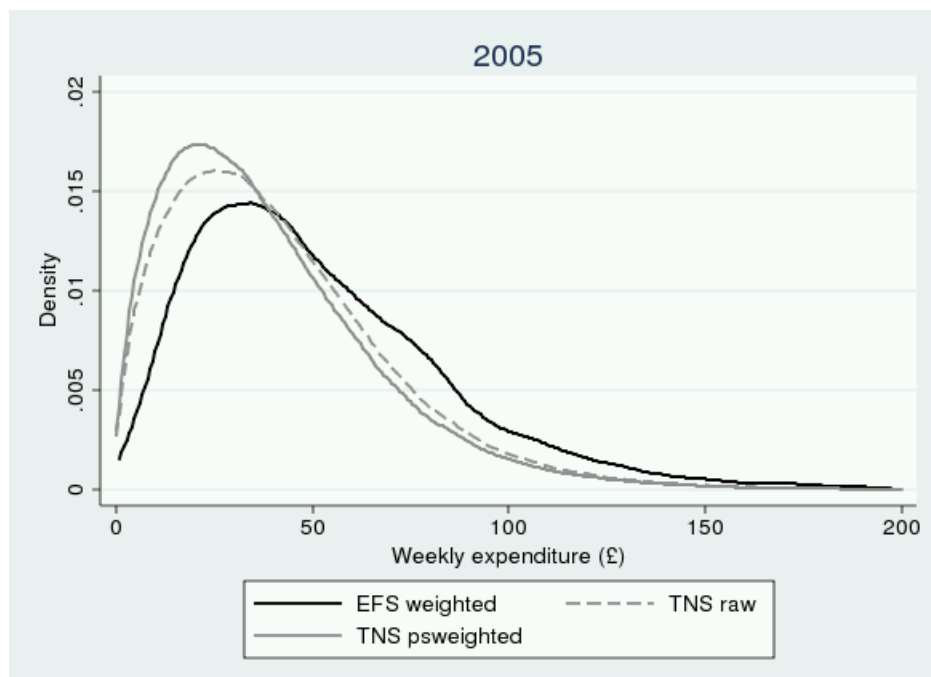
Crucially, however, although the absolute level of expenditure recorded in the TNS is around 25% below that recorded in the EFS (and this figure is remarkably stable across the five years of data), differences in recorded budget shares across the two surveys are small, as shown in table 4. Other than alcohol, there is never more than a one percentage point difference between the budget share recorded in the TNS and the budget share recorded in the

²¹ See, for example, Tanner (1998) and Attanasio et al (2006).

EFS. This similarity is reassuring for researchers wishing to use the TNS scanner data – in broad terms, expenditure patterns accord closely with those observed in much more familiar data. We consider in section 5b whether a more narrowly detailed analysis of spending across finer categories of goods shows the same result.

Though the average level of spending is much lower in the TNS than the EFS, it is also interesting to compare the distribution of spending levels across observations in the two datasets rather than just looking at single point comparisons. Figure 9 shows a kernel density estimate of the spending distribution in the two surveys for 2005 (results for other years are similar). What is clear from the figure is that modal expenditures in the two surveys are very similar at around £25-£30 per week. However, the TNS data appears to contain fewer high-spending households and more low-spending households than the EFS, a result accentuated once we demographically-weight the TNS sample which shifts the entire distribution slightly left. This is confirmed if we look at table 7 which contains summary statistics comparing the distribution of total spending (on food, beverages and alcohol) in the two surveys in each year.

FIGURE 9
Distribution of expenditures, TNS and EFS, 2005



Notes: TNS data are household-week observations for households recording spending for at least 4 weeks of calendar year 2005; EFS figures are household-level observations from EFS 2005, conditional on households spending something in both weeks for which they are observed and excluding Northern Ireland. Weighted TNS data are weighted by propensity score weights.

TABLE 7
Distribution of total expenditures, EFS and TNS, 2002–2006 (£/week)

2002	EFS	TNS (unwt.)	TNS (wt.)	2003	EFS	TNS (unwt.)	TNS (wt.)
Mean	49.13	39.41	36.82	Mean	50.42	40.61	38.14
10 th percentile	16.10	10.18	9.10	10 th percentile	16.13	10.47	9.48
25 th percentile	26.07	19.53	17.56	25 th percentile	27.00	20.07	18.26
Median	43.41	34.29	31.44	Median	44.40	35.19	32.58
75 th percentile	65.11	53.10	49.87	75 th percentile	67.25	54.59	51.52
90 th percentile	89.81	74.22	70.55	90 th percentile	91.85	76.65	73.05
2004	EFS	TNS (unwt.)	TNS (wt.)	2005	EFS	TNS (unwt.)	TNS (wt.)
Mean	52.24	41.63	39.23	Mean	53.15	42.52	39.79
10 th percentile	16.90	10.64	9.69	10 th percentile	17.07	10.91	9.83
25 th percentile	27.95	20.54	18.78	25 th percentile	28.10	21.07	19.02
Median	45.96	36.17	33.47	Median	45.65	36.97	33.90
75 th percentile	68.89	56.14	52.94	75 th percentile	70.50	57.41	53.80
90 th percentile	94.35	78.69	75.07	90 th percentile	97.00	80.21	76.36
2006	EFS	TNS (unwt.)	TNS (wt.)				
Mean	54.99	44.56	42.61				
10 th percentile	17.89	11.07	9.90				
25 th percentile	29.76	21.79	19.78				
Median	48.06	38.53	35.97				
75 th percentile	71.69	59.98	57.66				
90 th percentile	99.78	84.43	78.86				


Notes: TNS data are household-week observations for households recording spending for at least 4 weeks of each calendar year; EFS figures are household-level observations, conditional on households spending something in both weeks for which they are observed and excluding Northern Ireland. Weighted TNS data use the propensity score weights derived in the previous section. EFS data are weighted by sampling weights provided in the data.

So far, we have illustrated that TNS expenditure levels are below those in the EFS and that this is not explained by demographics. One other factor that may account in part for the difference is the fatigue result highlighted in section 3b. Since household spending tends to decline with length of participation, we would expect recorded TNS spending to be lower as part of our sample contains households that have been in the survey for some time, whereas all EFS households are ‘fresh’ (and are asked to record spending only for two weeks). To assess the impact this has on the spending comparisons, we compare in figure 10 very broad spending on food, beverages and alcohol for 2005 in the EFS, and two versions of the weighted TNS sample. The first is the same sample used so far, of household-week observations for all households that purchase for at least 4 weeks in 2005. The second contains *only* households that are observed participating for their first two full weeks in 2005 (and who purchase in both those weeks), for whom fatigue is presumably much less of an issue²².

²² Propensity score weights are estimated separately for this sample. Since the sample of new households is quite different from the sample of established households (new households will be over-represented amongst those types that frequently attrit) making demographically adjusted comparisons is very important.

FIGURE 10

Broad expenditure breakdown, TNS (full sample and new starters) and EFS, 2005



Notes: TNS ‘full sample’ data are household-week observations for households recording spending for at least 4 weeks of calendar year 2005; ‘new starters’ are those recording their first full two weeks during 2005 and purchasing in both weeks; EFS figures are household-level observations from EFS 2005, conditional on households spending something in both weeks for which they are observed and excluding Northern Ireland. TNS data are weighted using the propensity score weights derived in the previous section.

Clearly, fatigue may play some role in the expenditure differences, but is not enough to explain them altogether. The spending gap of 25% in the full sample falls to around 16% for the sample of new starters (£44.51 total expenditure in the new starter TNS sample compared to £53.15 in the EFS sample). Spending levels are around 12% higher in the new starter sample than in the full sample, and are around 21% higher for alcohol which accords with the earlier findings of greater alcohol ‘fatigue’ and relatively lower alcohol spending. Note, though, that alcohol spending amongst new TNS starters is still 30% below that of the EFS.

b) Detailed expenditure comparisons – does TNS underrecording vary by product?

Tables 5 and 6 suggested that, other than alcohol, most expenditure categories exhibited a similar degree of ‘underrecording’ in the TNS relative to EFS. Once propensity score weights are applied to TNS data, TNS spending levels on non-alcohol categories range from 70.4% of EFS levels for fruit to 86.4% for oils and fats.

In part, this similarity may be driven by the relatively large amount of product aggregation. Looking at more specific categories of spending may tell us something more about the effects of survey mode on recorded spending. There might be types of goods which are recorded relatively less well than others that are not picked up by a relatively broad aggregation of products.

There is of course a trade-off to be made between tractability and what we can discern about expenditure patterns from very narrowly-defined groupings. We could look at the 31

RPI groups as described in Appendix A rather than the 13 more aggregate categories used so far. In principle, one could compare expenditures across the datasets at an even more disaggregate level such as individual EFS expenditure codes of which there are around 100 relating to food, beverages and groceries. However there is considerable ambiguity about how half a million observed TNS products should be allocated to some individual EFS codes that is mitigated by this RPI-level analysis. Even here, however, the problem is not eliminated – it is not totally clear how TNS codes should be allocated between particular RPI groups, notably fresh and processed meats, fish and vegetables which is why we have preferred to look at even more aggregate spending groups than these 31 for most of our analysis.

With this in mind, Table 8 repeats the analysis for the more narrowly defined expenditure categories of the 31 RPI groups (Appendix A also looks at expenditures in 2002 and 2006 at this level, the results are similar). These figures suggest more variation in spending by good and support the idea that scanner technology impacts on recorded spending. In 2005, for example, comparing demographically-weighted TNS spending to sample-weighted EFS spending, the (non-alcohol) RPI categories for which the TNS data was furthest below EFS levels were lamb (-49%), bread (-37%), beef (-36%), poultry (-34%), fresh fruit (-31%), processed fish (-31%), fresh vegetables (-30%) and fresh milk (-28%). Other than processed fish, all the other groups are either categories where non-barcoded products are particularly important (fresh meats, fruit and vegetables) or those where we may think top-up and secondary shopping are particularly important (milk, bread and so on). The recording of non-barcoded products is a more lengthy process²³ and for top-up items the concern may be that households simply neglect to record these purchases and focus instead on their full weekly shopping trips²⁴. These figures provide suggestive evidence for both. A more detailed analysis of these issues may be possible and extremely interesting. For example, using the data, it is possible to discern which products are more typically purchased as part of very small shopping trips that may represent top-up trips, or to look at which products tend to be purchased at local corner stores rather than at supermarkets; similarly the product descriptions allow researchers to determine which categories contain larger fractions of non-barcoded products. We may also be able to look at even finer disaggregation of spending in the two surveys where we can be more confident that misallocation issues are not very important.

²³ Respondents are required to find the generic bar code in a book, scan that and then enter details including weight of the product and (for fruit and vegetables at least), country of origin.

²⁴ Einav et al (2008) find some evidence for this failure to scan top-up items using US Nielsen data; comparing scanner data to store records, they find that consumable items like snacks and soft drinks are often not reported, perhaps because they were eaten on the way home and the packaging with barcode discarded before the item was brought into the house.

TABLE 8

Detailed expenditure patterns, TNS and EFS, 2005, levels (£/week) and % of total

	TNS		EFS		Difference (%)
	Level	%	Level	%	
Bread	1.36	3	2.17	4	-37
Cereals	1.85	5	2.26	4	-18
Biscuits	2.53	6	2.95	6	-14
Beef	0.98	2	1.53	3	-36
Lamb	0.34	1	0.67	1	-49
Pork	0.46	1	0.58	1	-21
Bacon	0.69	2	0.88	2	-22
Poultry	1.15	3	1.74	3	-34
Other meat	4.49	11	5.04	9	-11
Fresh fish	0.71	2	0.97	2	-27
Processed fish	0.8	2	1.16	2	-31
Butter	0.24	1	0.28	1	-14
Oils & fats	0.55	1	0.64	1	-14
Cheese	1.37	3	1.56	3	-12
Eggs	0.34	1	0.46	1	-26
Fresh milk	1.55	4	2.15	4	-28
Milk products	1.46	4	1.78	3	-18
Tea	0.32	1	0.44	1	-27
Coffee & hot drinks	0.56	1	0.62	1	-10
Soft drinks	2.11	5	2.77	5	-24
Sugar & preserves	0.46	1	0.57	1	-19
Sweets & chocolate	1.46	4	1.91	4	-24
Fresh potatoes	0.55	1	0.69	1	-20
Processed potatoes	0.99	2	1.17	2	-15
Fresh vegetables	2.49	6	3.58	7	-30
Processed vegetables	1.02	3	1.13	2	-10
Fresh fruit	2.17	5	3.15	6	-31
Processed fruit	0.46	1	0.63	1	-27
Other foods	2.52	6	3.04	6	-17
Beer	0.99	2	1.79	3	-45
Wine & spirits	2.82	7	4.83	9	-42

Note: TNS data includes households that purchase in at least 4 separate weeks during 2005; EFS data exclude households that purchase no food, beverages or alcohol in one or both weeks of their spending diary. Difference is the percentage difference between TNS and EFS spending. TNS data are weighted using the propensity score weights derived in the previous section; EFS data is sample-weighted.


c) Does TNS underrecording vary by household type?

Our broad analysis in section 3a found that recorded expenditure is around 25% lower in the TNS than in the EFS. We now ask whether this finding is uniform across different demographic groups or whether this relative difference is bigger (or smaller) for certain groups. If there are differences in the relative spending between datasets across demographic groups, this might partly be driven by survey mode. For example, households that find the technology more difficult to use or who have less leisure time might under-record their spending to a greater extent than other households who are more at ease with the technology or who have more leisure time. Of course, this recording behaviour may be true in any survey (including the EFS) which require effort from respondents. However, the extent to which we find differences between the EFS and TNS *within* the demographic groups might be

informative about the survey mode effect. Again, we focus on demographically-weighted TNS data and sample-weighted EFS data from 2005.

Figure 11 shows spending by broad household composition. The bars make comparisons of spending levels across datasets: hatched bars represent spending for the EFS and solid bars spending for the same demographic group in the TNS. Spending for all demographic groups is lower in the TNS than in the EFS. However, there are differences in the relative spending of different household types – the figures above each TNS bar represent the level of TNS spending for that group as a percentage of EFS spending. Single adults in the TNS spend 93% of what single adults in the EFS spend; by contrast multiple adult households with children spend only two-thirds of what the same household group records in the EFS (recall that on average across the whole data, TNS figures for 2005 were around 25% below EFS figures). Our findings suggest single adult households have spending levels more closely aligned to EFS levels than multiple-person households. Adding children appears to increase the gap between EFS and TNS spending (for example, a couple without children record 76% as much as EFS couples without children; couples with children 70% as much) and the lowest relative spending comes from households with three or more adults. We know these households are oversampled in the (unweighted) TNS because of difficulties in ensuring adult shoppers other than the main shoppers comply with recording spending; these results appear to back up that claim.


FIGURE 11
Expenditures by household composition, TNS and EFS, 2005



Notes: TNS data are household-week observations for households recording spending for at least 4 weeks of calendar year 2005; EFS figures are household-level observations from EFS 2005, conditional on households spending something in both weeks for which they are observed and excluding Northern Ireland. TNS data are weighted using the propensity score weights derived in the previous section. Hatched bars are EFS figures and solid bars are TNS figures. Numbers represent the relative spending in the TNS to the EFS for the same household group.

Figure 12 repeats the analysis according to the age of the oldest person in the household. Spending by age follows a similar but not quite identical pattern in both datasets; rising towards middle-age, then declining into older age. This of course is related to family formation patterns over the lifecycle which increase expenditure demands. Note, though, that the peak is different across datasets, coming at ages 40–44 in the TNS and 45–49 in the EFS. Further, households in which the eldest person is age 80+ spend much less than those aged 75–79 in the EFS (£34.92 against £43.10) but this is not the case in the TNS (£30.83 against £31.90). Indeed, relative to the spending levels recorded in the EFS, it is this eldest age group that records the highest spending in the TNS, at 88%. Over all the age groups, the ages at which TNS households record the least relative to EFS are those middle-aged groups when spending is highest; further, the pattern of spending across age is much flatter in the TNS data than in the EFS data.

FIGURE 12
Expenditures by age of eldest household member, TNS and EFS, 2005



Notes: TNS data are household-week observations for households recording spending for at least 4 weeks of calendar year 2005; EFS figures are household-level observations from EFS 2005, conditional on households spending something in both weeks for which they are observed and excluding Northern Ireland. TNS data are weighted using the propensity score weights derived in the previous section. Hatched bars are EFS figures and solid bars are TNS figures. Numbers represent the relative spending in the TNS to the EFS for the same household group.

Do these results tell us anything about the effect of survey mode? The fact that the very oldest households record higher relative spending may concur with suggestions from TNS that older households are more assiduous in recording their spending, though this is only clear really for the very oldest aged 80+ and not for ‘younger elderly’ households. Young households recording relatively high levels of spending may well be consistent with the results for single adult households we saw in figure 11. The key message appears to be that relative spending is particularly low for households in middle-age who are more likely to have children.

Another interesting finding from figure 12 comes if we look only at alcohol expenditures. We have seen that alcohol is particularly poorly recorded in TNS relative to EFS (and in turn is badly recorded in EFS data as well). Across age groups, the group with the lowest record of alcohol spending in the TNS relative to EFS are the very youngest aged under 25, who record only 30% as much alcohol spending in the TNS as in the EFS (in both surveys, this relates to alcohol brought into the home). By contrast, older households record much higher levels of alcohol spending and in the case of the oldest 80+ households, levels of alcohol spending are *higher* in the TNS than in the EFS data. It may be possible to speculate why alcohol spending is so variably recorded (embarrassment, changing habits, forgetfulness) but it is clear that alcohol is particularly problematic in this kind of survey and especially so for younger households, even relative to known problems in the EFS.

These univariate comparisons are interesting but it is not clear how the various demographic characteristics interact with one another to drive the findings. To assess this, we carry out a simple regression of log total expenditure on dummy variables of household characteristics interacted with the dataset in which households are observed, using the 2005 data. These interactions tell us along which demographic dimensions TNS expenditures are relatively higher or lower than EFS expenditures controlling for other common observed characteristics and whether the differences in relative spending levels across datasets are significant²⁵. Table 9 shows the coefficients on selected interaction terms and their statistical significance: negative numbers suggest relative spending for that group is, conditional on other observables, lower in the TNS than the EFS.

²⁵ Given the large sample size in the TNS data owing to our use of household-week observations, we cluster standard errors at the household level. This has no implication for EFS households since only one observation per household is observed in any case.

TABLE 9

Coefficients on interaction of household demographics and TNS dummy, 2005

<i>Variable</i>	<i>Coeff.</i>		<i>Variable</i>	<i>Coeff.</i>	
PC in household	-0.076	***	Region E Midlands	-0.085	*
Zero cars	-0.047	**	Region W Midlands	-0.069	
Two cars	-0.009		Region Eastern	-0.104	**
Eldest aged <25	-0.049		Region London	-0.125	***
Eldest aged 25-29	+0.086	*	Region South East	-0.094	**
Eldest aged 30-34	+0.092	**	Region South West	-0.077	
Eldest aged 35-39	+0.099	***	Region Wales	-0.094	*
Eldest aged 40-44	+0.074	**	Region Scotland	-0.077	
Eldest aged 50-54	+0.015		Tenure Council Rent	+0.074	***
Eldest aged 55-59	-0.029		Tenure Private Rent	+0.048	
Eldest aged 60-64	-0.019		Tenure Other	+0.005	
Eldest aged 65-69	+0.004		Head works 8-29 hrs	+0.038	
Eldest aged 70-74	+0.003		Head works <8 hrs	+0.102	
Eldest aged 75-79	-0.026		Head unemployed	+0.131	**
Eldest aged 80+	+0.073		Head retired	+0.039	
Eldest person female	-0.100	***	Head in FT education	+0.087	
Region North West	-0.076	*	Head inactive	+0.098	***
Region Yorkshire	-0.045				

Note: included in the regression but omitted from the table are number of adults, number of children and household composition; none of these exhibited significant differences. The R^2 of the regression is 0.151. *** = significant at 1% level; ** = significant at 5% level; * = significant at 10% level. Omitted variables from each category are no PC, one car, eldest aged 45-49, eldest male, region North East, tenure owner-occupier, head works 30+ hours per week.

Broadly, our results support what the univariate analyses above suggest: relative spending is higher for younger households aged between 25 and 44 (for the very oldest households, we estimate spending to be relatively higher in the TNS but not significantly so, reflecting the small sample size in this age group). Interestingly, none of the variables relating to numbers of adults, children or household composition are significant, though the signs go as we would expect (progressively more negative as we add adults and children), suggesting these factors are strongly related to other observable characteristics that influence expenditures²⁶. In terms of variables that could proxy income (since in 2005 we do not have any direct measures of income), the results tend to suggest that poorer households record relatively higher spending in the TNS compared to EFS levels: those with PCs have lower relative spending, households headed by someone who is unemployed or out of the labour market (except retirees) have higher relative spending, as do council tenants. The only case where this is reversed is number of cars: households with zero cars tend to report lower relative spending in the TNS. These results tend to suggest that households with lower valuations of time may be relatively more assiduous recorders of expenditure in the TNS; further investigation could reveal whether, for example, those households without cars who report relatively lower spending do so because they tend to shop more frequently purchasing top-up items which may be less well covered,

²⁶ Einav et al (2008) compare US scanner data from AC Nielsen to store records to identify whether households record all of their trips and items from those trips; they suggest that smaller households, particularly single person households, are less prone to misrecording shopping trips though their results compare characteristics one by one and are based on a relatively small household sample for which matching trips could be recorded.

or because they tend to live in London, say, where spending is particularly poorly recorded relative to EFS levels.

d) The effect of survey mode within survey: scanner technology changes in the TNS

In the description of the TNS data in Section 2, we discussed how in 2006, a new type of scanner technology was introduced for the booster sample of households which meant they no longer had to scan non-barcoded items as part of their participation, but which TNS suggest have increased overall compliance. We assess how the recorded spending of households using the different scanner technologies compares to EFS and how far differences across the scanner types can be explained by the demographic makeup of households using them. On at least one front, we know there must be a demographic difference between households with the different technologies since those with the new kit upload their data to TNS via their home computer rather than a fixed landline connection, so all those with the new kit have a computer in the home. This also suggests they may be richer and younger on average. Table 10 shows the breakdown of spending for 2006 by the scanner type – old and new – and shows both demographically-weighted and unweighted averages for comparison. If we just look at the raw unweighted figures, then it appears that those with the new kit type do indeed record significantly more spending – £48.03 per week compared to £43.51 per week for those with the old technology – and have spending levels much closer to those recorded in the EFS (£54.99 per week). But note that these total spending differences across scanner types are almost *entirely* explained by demographics: once we adjust using demographic weights (where the weights are defined separately for the sample using each kit type), households with the new technology spend £43.11 per week compared to £42.08 per week for those with the old technology. The gap is now just 2.4% and both types record spending levels more than 20% below EFS levels.

TABLE 10
Expenditure breakdown, TNS (by scanner type) and EFS, 2006

<i>Category</i>	<i>EFS</i>	<i>TNS old kit (unweighted)</i>	<i>TNS new kit (unweighted)</i>	<i>TNS old kit (weighted)</i>	<i>TNS new kit (weighted)</i>
Bread, cereals & biscuits	7.66	6.36	6.99	5.99	6.23
Meat	10.62	8.63	9.44	8.24	8.44
Fish	2.31	1.63	1.71	1.65	1.61
Butter, oils and fats	0.97	0.86	0.99	0.82	0.89
Eggs, milk, cheese	6.11	5.18	5.44	4.97	4.92
Hot beverages	1.07	0.95	1.21	0.91	1.07
Soft drinks	3.11	2.45	2.99	2.34	2.74
Sugar, sweets, chocolates	2.42	2.05	2.49	1.96	2.12
Potatoes	1.91	1.72	1.84	1.6	1.62
Other vegetables	4.85	3.84	3.67	3.81	3.38
Fruit	4.02	3.04	2.26	3.02	2.18
Other food	3.19	2.74	3.27	2.63	2.98
Alcohol	6.74	4.03	5.73	4.14	4.93
Total expenditure	54.99	43.51	48.03	42.08	43.11

Notes: TNS data are household-week observations for households recording spending for at least 4 weeks of calendar year 2006; EFS figures are household-level observations from EFS 2006, conditional on households spending something in both weeks for which they are observed and excluding Northern Ireland. Weighted TNS data use propensity score weights derived separately by kit type.

However, there are clear differences in the expenditure *composition* across scanner types even after adjusting for demographic differences. Households with the new type, on average, record more expenditure on drinks – 19.1% more on alcohol, 17.6% more on hot beverages and 17.1% more on soft drinks. They also tend to record a little more on most food categories. The categories where households with the new technology spend substantially less are vegetables (-11.3%) and fruit (-27.8%). It is almost certainly loose fruit and vegetables without barcodes that drives this expenditure difference, though it is not immediately clear why households with the new technology would be so much better at recording purchases of alcoholic and non-alcoholic drinks.²⁷ It is not clear to us that the trade-off that TNS have made in sacrificing the need to record non-barcode items in return for greater compliance in other areas has been entirely successful since much of the difference in total spending can be attributed to demographic differences across households. Recording non-barcode items meant that the TNS grocery purchase data was at least in principle a complete record; other scanner datasets have excluded them or included them for only a subsample of households. However, whilst there is a significant sample still using the old technology this data is at least still available. More investigation into how and why spending differs across the technology types may be informative, as would experiments on whether there may be better ways to encourage households to record non-barcode items which are an important component of

²⁷ One likely explanation is that there are differences in demographics which we are unable to account for since only a limited set of characteristics are common to both datasets.

spending in some product groups²⁸. In the future, if the new technology kit becomes permanent and if existing respondents who are currently using the old kit are changed over to the new technology, it will be interesting to see more directly what the effects are as we will be able to control for within household effects.

6. Attrition from TNS data

Our final section of analysis we focuses on attrition from the TNS and examines factors that are associated with households leaving the data and the length of time they participate. Attrition in panel studies is a well known phenomenon. Following individuals or households over time inevitably means that subjects are lost. The time involved in taking part in any survey means that some respondents drop out because they no longer want to or are no longer able to dedicate that time. Other things being equal, one would expect that the less time it takes to take part in a survey, the lower the chance that a respondent will attrit. Other reasons for attrition may include contact failure following a geographical move or death of the respondent.

The TNS data is unlike most traditional social science surveys in the way that respondents are recruited and in the way that cooperation is on-going. Typically, social science panel surveys are conducted on an annual or bi-annual basis (the British Household Panel Study or the English Longitudinal Study of Ageing for example). These surveys invest substantial resources into “panel maintenance” in order to minimise respondent drop-out. The TNS data relies on respondents taking part on a continuous basis for as long as they wish to. The burden on the respondent is therefore higher in the TNS than in an annual survey. In addition, the data are designed from a marketing perspective and it might not be efficient for respondents to be encouraged to stay for indefinite periods of time. At any one time, the panel is designed to be broadly representative, so households who tend to drop out more often are also recruited more often; attrition might be less of an issue for the market researcher. However, for the social science researcher, where the panel nature of the data might be very valuable, attrition is likely to be more important.

The nature of causes of attrition from social science panel surveys is a well-studied topic (see for example Uhrig, 2008 who looks at attrition in the BHPS). Attrition is a complex issue and it would certainly be possible to write an entire paper on this topic with respect to the TNS data. However, our intention here is simply to provide a very descriptive analysis of attrition rather than to carry out any complex modelling. We start with a descriptive analysis

²⁸ In particular as we have seen fruit and vegetables, where there may be a lot of interesting work that uses the nutritional data in the TNS sample to investigate intake of vitamins and fruit and vegetable ‘portions’ across demographic groups, say.

of overall attrition rates and hazard functions and then use a duration model to see how attrition rates vary with demographics.

a) Overall attrition rates in the TNS

Before we begin to look at pattern of attrition, it is necessary to describe the nature of our data in a little more detail. The data we are considering, records purchases that span the period November 2001 to the end of 2006. However, respondents who record their expenditure during that time were recruited from as far back in time as 1989. As such, we do not observe the first spell in the data for any respondent that signed up to the survey prior to November 2001. These data are commonly known as “left truncated” and make up 36% of our sample. Because by definition, the left truncated observations are those which have completed a spell length sufficient to be included in our data period, these observations upwardly bias any estimates of duration (length of participation). On the other side of the coin, we have 47% of our sample who are “right censored” – in other words, as these observations are still in the sample at the end of period of data, we do not observe their final spell in the data. We observe completed spells for only 31% of our sample (this implies that 14% of our sample are both right censored and left truncated).

The first question is how long, on average, do respondents stay in the TNS?²⁹ Of those for whom we observe a completed spell, the mean number of weeks spent in the data is 48. If we include the right censored households (but exclude the left truncated data) and count the number of weeks they have been in the data so far, the mean number of weeks spent in the data is 62.

Table 11 shows the distribution of spell length for all respondents and for those whom we observed from their first spell (final column labelled “not truncated”). Of those observed from their first spell, there is a relatively high level of attrition in the first 4 weeks (7.2% drop out during this time). Drop out rates then tail off somewhat during the first year. Nearly 8% drop out between years 1 and 2.

Can we compare attrition rates in the TNS to attrition in other panel studies such as the BHPS? The very different nature of the BHPS and the TNS means that we would not expect drop-out rates to be comparable but nevertheless such a comparison might provide a useful benchmark against which to judge the amount of attrition in the TNS. If we take all non-truncated households who signed up in 2005 or before and therefore have a chance to be observed for at least a year, we find that the probability of observing a household for at least a year is 63%. Taking the sample of BHPS households who gave a full interview in 1991 (the

²⁹ Note that we define a household’s last week to be the last week that they are observed to scan items. TNS record an official drop-out week but this can be a number of weeks after the household is last observed to scan anything.

first wave of BHPS), 86% gave a full interview the following year. So although the probability of observing a household one year later is clearly lower in the TNS, given that those household have co-operated for a full year we consider a follow up response rate of 63% to be quite high, despite the relatively onerous burden of continued participation.

One point worth noting is that even though our definition of participation requires a household to be scanning items, they may not be fully complying. As they approach the time that they drop out, households might scan less of their shopping as their commitment level diminish. Our analysis of survey fatigue (see section 3b and Appendix B) suggests some evidence for this: within four weeks of dropping out, households record around 10% less than in their first week on average, on top of the estimated level of fatigue for the duration of their participation at that point.

Figure 13 shows the Kaplan-Meier survival function for the observations which are not left truncated. The probability of a household surviving falls fairly rapidly during the first year. The survival function then becomes less steep. The probability of a household surviving for 2 years is just under 50%. After 5 years, the probability of survival is 18%.


TABLE 11

The distribution of spell lengths in the TNS

<i>Probability of attriting between:</i>	All respondents	Not truncated
weeks 1-4	4.6	7.2
weeks 5-8	2.5	3.9
weeks 9-12	1.9	2.9
weeks 13-16	1.9	2.9
weeks 17-20	1.9	2.7
weeks 21-24	1.4	2.1
weeks 25-28	1.4	2.0
weeks 29-32	1.3	1.8
weeks 33-36	1.2	1.8
weeks 37-40	1.3	1.7
weeks 40-44	1.1	1.5
weeks 45-48	1.3	1.8
weeks 49-52	0.9	1.1
1-2 years	8.2	7.9
2-3 years	5.5	4.0
3-4 years	3.7	2.0
4-5 years	2.6	0.9
5+ years	9.9	0.0
<i>right censored</i>	<i>47.5</i>	<i>51.8</i>

FIGURE 13

Kaplan- Meier survivor function for non-truncated observations



b) Attrition rates and demographics

This section examines the association between attrition and household demographics to identify factors which increase or decrease the risk of attriting. Our intention here is provide a preliminary descriptive analysis into the correlation of attrition and demographics rather than to construct a model of the attrition process which is beyond the scope of this paper.

It is well-known that certain demographic groups are harder to attract to surveys and harder to retain in panel surveys. For example, Uhrig (2008) found that older respondents were more likely to refuse to participate in subsequent years in the BHPS whereas respondents with children were more likely to remain in the study. Fitzgerald et al (2000) found that attrition is concentrated amongst lower socioeconomic status individuals in the Panel Study of Income Dynamics which is fielded in the United States. A priori, one might expect those individuals who have less leisure time to be more likely to drop out of the TNS survey; our analysis of relative spending found unemployed households spent relatively more compared with EFS and so it is not unreasonable to assume a similar correlation in terms of attrition. We know that elderly households (particularly the very elderly) are less likely to sign up given our demographic comparisons, but conditional on having signed up, are they more or less likely to remain in the panel? In order to try to answer some of these questions, we use a Cox proportional hazard model to estimate the effect of demographics on the

duration of time spent in the survey. We estimate a Cox proportional hazard model rather than a fully parametric duration model to avoid the need to specify the shape of the hazard function.

Our results are presented in Table 12. We estimate the duration model for all observations (left two columns of results) and separately for those observations which are not left truncated (rightmost two columns of results). We include region, age, number of adults, number of children, employment status of head, household composition, number of cars and sex of head as explanatory variables.³⁰ In addition, we also include a dummy variable which indicates whether the household is using the new scanner technology which is designed to make reporting easier and hence might encourage sample retention. However because the new kit was only introduced to those signing up in 2006 we also include dummy variables to indicate which year the household signed up (these are not reported in the table). Around 20% or so of households who joined the survey in 2006 do not have a recorded employment status so, in addition to the employment status indicator variables, we also include a dummy for missing employment status. The coefficients reported are hazard ratios. A hazard ratio of one means that characteristic is not associated with duration, a hazard ratio of less (greater) than one means that characteristic is associated with a lower (higher) risk of attriting or equivalently, a longer survival time.

For both models (all observations and the non-truncated observations), age is significantly related to duration in the study. Relative to 30-34 year olds, younger age groups have a higher hazard ratio and older age groups have lower hazard ratio meaning that being in a younger age group, all else being equal, increases the risk of attriting from the survey. This contrasts with our finding that younger households had relatively higher expenditures compared to EFS spending levels, suggesting that these households are relatively compliant whilst they participate but simply do not continue to do so for long.

There is little difference in the hazard ratios across regions although, relative to the omitted category (South East), living in the North West and London increases the risk of attriting by around 6%. Both regions were associated with lower relative spending levels in our demographic comparison of expenditures.

Number of children is strongly related to duration: relative to having no children in the household, having one child (for the whole sample) increases the hazard ratio by 14% and for the non-truncated sample by 34%. Having successively more children, also increases the hazard ratio relative to having no children. This is not surprising if we expect that these households have greater demands on their time, and concurs with our finding that more

³⁰ Note that we do not include income as an explanatory variable because so few observations have income recorded. For presentational purposes, not all coefficients in all categories are reported in table 12 but full results are available.

children added to a household reduces recorded TNS expenditures relative to EFS. Having just one adult in the household decreases the risk of attrition, conditional on age. This may reflect the need in multi-adult households for everyone to co-operate, and again accords with our finding on expenditures. Single pensioner households have a higher hazard ratio relative to non-pensioner couple households, even after controlling for age. We may speculate that pensioner households are less comfortable with scanner technologies, though our analysis of relative spending levels by age group did not suggest older households had lower relative spending conditional on other demographics.

Interestingly, we find almost no significant effects of employment status on attrition, conditional on other variables. Households with a missing employment status are much more likely to attrit, though since it is unclear why employment status is missing this may be picking up non-cooperative households. Relative to heads employed full-time, households headed by part-time or non-working people are no less likely to attrit.

Looking at the impact of having the new scanner technology, we find that households using the new kit have a lower hazard ratio than those with the old kit type, conditional on when they began to participate. Our conclusion that the new kit type did not necessarily increase recorded spending once demographics are controlled for is therefore tempered somewhat by the finding that the new technology is associated with lower dropout rates. Again, this finding indicates that even within a given survey, small differences in design such as a new type of scanner can have an effect on survey outcomes in many ways.

TABLE 12

Results of duration model of length of participation in TNS survey

	<i>Hazard ratio</i>		<i>Robust SE</i>		<i>Hazard ratio</i>		<i>Robust SE</i>	
Region (base =South East)	All observations				Non-truncated			
North West	1.057	*	0.028		1.048		0.036	
London	1.060	*	0.031		1.051		0.039	
Age (base = Aged 30-34)								
Aged <25	1.703	*	0.083		1.725	*	0.091	
Aged 25-29	1.078	*	0.031		1.111	*	0.036	
Aged 35-39	0.882	*	0.023		0.848	*	0.026	
Aged 40-44	0.767	*	0.022		0.712	*	0.026	
Aged 45-49	0.736	*	0.024		0.737	*	0.030	
Aged 50-54	0.618	*	0.021		0.600	*	0.026	
Aged 55-59	0.581	*	0.021		0.584	*	0.028	
Aged 60-64	0.505	*	0.023		0.503	*	0.032	
Aged 65-69	0.513	*	0.029		0.537	*	0.043	
Aged 70-74	0.474	*	0.028		0.455	*	0.039	
Aged 75-79	0.608	*	0.038		0.582	*	0.053	
Aged 80+	0.755	*	0.049		0.600	*	0.064	
Missing age	1.243	*	0.072		0.773		0.142	
Number of adults (base = 2 adults)								
1 adult	0.809	*	0.064		0.751	*	0.074	
3 adults	0.972		0.046		1.047		0.067	
Number of children (base = 0 children)								
1 child	1.141	*	0.068		1.341	*	0.120	
2 children	1.071		0.066		1.280	*	0.117	
3 children	1.179	*	0.077		1.422	*	0.135	
4+ children	1.189	*	0.091		1.519	*	0.160	
Employment of head (base = Works >30hrs)								
Works 8-29 hrs	1.011		0.025		0.974		0.031	
Works <8 hrs	1.075		0.081		1.018		0.102	
Unemployed	1.067		0.048		1.069		0.067	
Retired	1.051		0.035		0.974		0.048	
Full time education	1.052		0.078		1.158		0.105	
Unoccupied	1.003		0.024		0.963		0.029	
Emp status missing x joined in 2006	2.209	*	0.098		2.195	*	0.100	
HH composition (base = 2 adults, non-pens)								
One adult, non-pensioner	1.118		0.094		1.206		0.127	
One adult, pensioner	1.417	*	0.134		1.409	*	0.177	
Two adults, pensioner	0.991		0.048		0.923		0.066	
Two adults with children	1.024		0.065		0.915		0.084	
One adult with children	1.333	*	0.113		1.280	*	0.146	
Three+ adults, no children	1.181	*	0.069		1.157		0.092	
Three+ adults, with children	1.185		0.107		0.997		0.132	
Number of cars (base = 1 car)								
No cars	1.098	*	0.023		1.118	*	0.030	
2+ cars	0.977		0.016		1.017		0.022	
New kit x signed up in 2006	0.454	*	0.019		0.456	*	0.019	
Male	1.042	*	0.017		1.029		0.021	

Controls also included (but not reported) for year that the household signed up.

Only significant coefficients are reported for region.

*indicates significance at the 5% level

7. Conclusions and Recommendations

This paper has analysed data collected by a market research company using in-home barcode scanners to record grocery purchases, in an attempt to assess how useful such data may be for social science researchers. We have compared the data to existing surveys and examined expenditure and demographic information in detail. Our key finding is that this sort of data may have considerable advantages over traditional social science survey data for academic research; while it has been used considerably in marketing literature, it has been less exploited by the social science research community, at least in the UK. Given its many possible advantages for researchers – including detailed product information, exact prices paid, precise store of purchase, rich information on nutrition and household attitudes and so on – the data represent an exciting source of information that could be used in a wide range of applications.

Our study of the data has thrown up several interesting results that we feel are important for researchers wishing to use this data to bear in mind. In general, our ‘guidelines’ for researchers and main findings are:

- The distribution of household demographics, in cross section, is broadly comparable to existing surveys, though with some important exceptions such as a deliberate oversampling of larger households. It may well be necessary for researchers to construct demographic weights for their own purposes, or at least to ensure they understand along which dimensions any sample they select for analysis are likely to be poorly representative.
- In particular, given that there are such differences in the propensity of households to drop out and be resampled, a sample comprising people that begin to participate in any given period will have a very different demographic balance to a sample of all active households over the same period.
- Market research data is collected primarily with the interests of clients in mind, and not for social science research. A number of variables that researchers would traditionally use are missing or incomplete (such as income and housing tenure) or not consistent over time. Crucially, it appears from our study that demographic *transitions* over time are particularly poorly recorded; again, this is because the panel element per se of the data is not necessarily particularly important for the data collectors. Researchers using this or similar data who want to exploit the panel element should look carefully at this issue for their particular purposes.
- Despite the relatively high time costs of continued participation, attrition rates from the TNS survey are surprisingly low, though higher than in more familiar panel data

- sets. Nevertheless, some types of household are more likely to attrit (the young, household with children, multiple adults) than others.
- Households with high levels of attrition are also more likely to have relatively low recorded spending compared to existing survey data. The presence of multiple adults or children in a household is associated with particularly low recorded spending which is consistent with relatively poorer compliance for these types of household.
 - One exception is younger households, who appear to record their spending relatively well whilst participating but who are also more likely to attrit. Older households are less likely to attrit, but it is not clear that their relative spending is higher, conditional on other characteristics. Households with a low marginal value of time may well be more assiduous at recording their spending whilst participating though it is not clear they are any less likely to attrit from the data.
 - Broad patterns of spending match very closely with those in existing survey data, though absolute levels of spending are typically lower in scanner data. There appears to be some correlation between relative spending levels on different items and, for example, whether they are typically ‘top up’ items or not and whether non-barcoded products are a large proportion of spending within a product group. Whilst no longer requiring participants to report non-barcoded items may be one way to reduce attrition, it does have a significant impact on expenditure patterns and may be crucial for particular areas of research.
 - Researchers need to be careful to consider the treatment of weeks in which households are not observed recording any expenditures. Although there will be circumstances where zero spending is genuine (holidays for example), the fact that zero spending weeks are a much more common occurrence in the TNS than in the EFS suggests that there are many cases where this is an extreme case of under-reporting.
 - Data collectors such as TNS are typically well aware of issues around dropout and underreporting. Though we have not made use of them in this study, product-level weights may be provided or derivable that can aggregate up reported expenditures to known national sales volumes and which will be higher for products where there are known problems such as top-up items. Researchers using this kind of data should ensure they are familiar with these kinds of issues and any information that may be provided by the data collectors on how to resolve them.
 - There do appear to be some clear effects of survey mode on demographics and expenditures. Even correcting for observed demographic differences, some households seem to report relatively low expenditures and dropout rates differ strongly by observed characteristics.

References

- Aguiar, E. and M. Hurst (2007), "Lifecycle prices and production", *American Economic Review*, Vol. 97, No. 5 (December), pp. 1533-59.
- Ahmed, N., M. Brzozowski and T. F. Crossley (2006), "Measurement errors in recall food consumption data", IFS Working Paper W06/21 (<http://www.ifs.org.uk/wps/wp0621.pdf>).
- Attanasio, O., E. Battistin and A. Leicester (2006), "From Micro to Macro, from Poor to Rich: Consumption and Income in the UK and the US", *mimeo*, http://www.npc.umich.edu/news/events/consumption06_agenda/Attanasio-Battistin-Leicester.pdf
- Banks, J., R. Blundell and S. Tanner (1998), "Is there a retirement-savings puzzle", *American Economic Review*, Vol. 88, No. 4 (September), pp. 769–788.
- Banks, J. and P. Johnson (eds) (1998), "How reliable is the Family Expenditure Survey? Trends in income and expenditures over time", IFS Report Series 57, London: IFS (<http://www.ifs.org.uk/comms/r57.pdf>).
- Bucklin, R.E. and S. Gupta (1999), "Commercial use of UPC scanner data: industry and academic perspectives", *Marketing Science*, Vol. 18, No. 3, pp. 247-273.
- Central Statistical Office (1985), *UK National Accounts: Sources and Methods: third edition*, London: HMSO.
- Fitzgerald, J.M., Gottschalk, P., Moffitt, R.A. (1998) "An Analysis of Sample Attrition in Panel Data: The Michigan Panel Study of Income Dynamics" NBER Working Paper No T0220, February 1998.
- Garner, T., G. Janini, W. Passero, L. Paszkiewicz and M. Vendemia (2003), "The Consumer Expenditure Survey in comparison: focus on personal consumption expenditures", Bureau of Labor Statistics, US Department of Labour: Washington, DC, <http://www.bls.gov/bls/fesacp1032103.pdf>.
- Griffith, R., E. Leibtag, A. Leicester and A. Nevo (2009a), "Consumer shopping behavior: how much do consumers save?", forthcoming in *Journal of Economic Perspectives*
- Griffith, R., L. Nesheim and M. O'Connell (2009b), "Empirical estimates of the impact of a fat tax", *mimeo*
- Keil, E. (2003), Unpublished aggregate expenditures based on AC Nielsen household scanner data for alcoholic beverages and tobacco and tobacco accessories, Bureau of Labor Statistics, US Department of Labour: Washington, DC.
- Einav, L., E. Liebtag and A. Nevo (2008), "On the accuracy of Nielsen Homescan data", US Department of Agriculture Research Report 69, Washington, DC: USDA.
- Robertson, C., N. Best, J. Diamond and P. Elliot (2003), "Tracing ingestion of 'novel' foods in UK diets for possible health surveillance – a feasibility study", *Public Health Nutrition*, Vol. 7, No. 2, pp. 345–352.
- Rosenbaum, Paul R., and Donald R. Rubin (1983) 'The Central Role of the Propensity Score in Observational Studies for Causal Effects', *Biometrika*, 70, 1, 41-55.
- Tanner, S. (1998), "How much do consumers spend? Comparing the FES and National Accounts", in Banks, J. and P. Johnson (1998), *op. cit.*
- Uhrig, S.C. Noah, (2008), "The Nature and Causes of Attrition in the British Household Panel Study" ISER Working Paper No 2008-05
- Van Heerde, H. J., P. S. H. Leeflang and D. R. Wittink (2000), "The estimation of pre- and postpromotion dips with store-level scanner data", *Journal of Marketing Research*, Vol. 37, Issue 3 (August), pp. 383-395.

Appendix A. Details of RPI expenditure categorisation

This appendix details the definition of the RPI expenditure categories that are used to analyse expenditures in the EFS and TNS data. There are 31 RPI expenditure categories for food, beverages and alcohol. We map expenditures in the TNS and the EFS into these categories when making comparisons between the datasets, and aggregate these into slightly larger groups to look at how important a share of the total budget they are. Below we detail the sorts of items that are included in each of the RPI expenditure groups and how each RPI group is aggregated into the broader categories that underlie, for example, figure 1, before table A2 reports details of average expenditures in the TNS and EFS in 2002 and 2006. Guidance as to the products included is based where possible on the ONS (2007).

TABLE A1

RPI Expenditure Categories

<i>RPI group</i>	<i>Typical items</i>	<i>Broad group</i>
Bread	Fresh wrapped and unwrapped bread, including buns, baguettes, bagels, etc; fruit and flavoured breads; prepared sandwiches brought home.	Bread, cereals and biscuits
Cereals	Rice and rice-based products (such as savoury rice); fresh, dried and tinned pasta and pasta-based products (such as pasta salads); breakfast cereals; non-potato based snack products; flour and grains.	Bread, cereals and biscuits
Biscuits	Packaged and loose biscuits; cereal and cake bars; scones, teacakes, etc.; cakes and fresh/chilled/frozen desserts (excluding milk-based desserts).	Bread, cereals and biscuits
Beef	Fresh, chilled or frozen cuts of beef, beef mince, etc. (not including beef-based products or cooked meats).	Meat
Lamb	Fresh, chilled or frozen cuts of lamb, lamb mince, etc. (not including lamb-based products or cooked meats).	Meat
Pork	Fresh, chilled or frozen cuts of lamb, pork mince, etc. (not including pork-based products or cooked meats).	Meat
Bacon	Bacon rashers and joints, gammon, ham.	Meat
Poultry	Fresh, chilled or frozen cuts of poultry, poultry mince, etc. (not including poultry-based products or cooked meats).	Meat
Other meat	Other fresh, chilled or frozen cuts of meat (rabbit, venison, etc., not including meat-based products or cooked meats); sausages; offal; meat-based pates and spreads; processed and cooked meats; meat-based ready meals and pies.	Meat
Fresh fish	Fresh, chilled or frozen fish and seafood (not including fish-based products or cooked fish).	Fish
Processed fish	Dried and smoked fish and seafood; canned or bottled fish; fish-based pates and spreads; fish-based ready meals and pies.	Fish
Butter	Branded and unbranded butters; salted and unsalted butters.	Butter, oils and fats
Oils & fats	Margarine and vegetable fats; imitation and soya creams, margarines, etc.; olive oil; other edible oils (e.g. sunflower, vegetable) and animal fats (lard, dripping, etc.)	Butter, oils and fats
Cheese	Fresh cheeses and curds; processed cheese.	Eggs, milk, cheese
Eggs	Chicken, duck, quail, etc. eggs.	Eggs, milk, cheese
Fresh milk	Whole, semi-skimmed and skimmed milk; UHT and sterilised milk.	Eggs, milk, cheese
Milk products	Baby milk and milk powders; condensed milk; flavoured milk and milkshakes; yoghurt, fromage frais and milk-based desserts (not including ice-cream).	Eggs, milk, cheese
Tea	Tea bags, loose tea, fruit teas, herbal teas, instant teas.	Tea, coffee, hot drinks
Coffee & hot drinks	Ground coffee, coffee beans, instant coffees; cocoa powder, drinking chocolate, malted hot drinks, etc.	Tea, coffee, hot drinks
Soft drinks	Fruit and vegetable juices and smoothies, squashes and concentrated drinks; carbonated beverages in cans and bottles; mineral and spring waters.	Soft drinks
Sugar & preserves	Sugar, artificial sweetener, icing sugar and ready icings; jams, marmalades, lemon curd, peanut butter, jellies; sugar products (syrup, sugar sauces, edible cake decorations, etc.)	Sugar, sweets and chocolates
Sweets & chocolate	Chocolate bars, box of chocolates; sweets and confectionery products, chewing gum.	Sugar, sweets and chocolates
Fresh potatoes	Old and new loose and packaged potatoes.	Potatoes
Processed potatoes	Potato crisps, fries, oven chips, hash browns, tinned and bottled potatoes.	Potatoes

Fresh vegetables	Unprocessed (fresh, chilled and frozen) leaf and stem vegetables (asparagus, lettuce, corn on the cob, etc.), fresh herbs, cabbages, vegetables grown for their fruit (courgettes, peppers, peas, etc.), root crops and bulbs (carrots, onions, swedes, etc.), mushrooms.	Other vegetables
Processed vegetables	Dried vegetables, lentils, pulses; canned and bottled vegetables, meat substitute products, vegetable-based ready meals and snacks.	Other vegetables
Fresh fruit	Fresh citrus fruits, bananas, apples, pears, stone fruits (cherries, dates, etc.), berries, other fresh, chilled or frozen fruits and fruit salads.	Fruit
Processed fruit	Dried fruit, crystallised fruits, nuts (whole and ground), seeds; tinned fruit and fruit products.	Fruit
Other foods	Ice cream and edible ices; table sauces and condiments, prepared cooking sauces; salt, spices and dried herbs; soups; yeast; dessert preparations (e.g. cake mixes, ready pastry, instant custard powder); food hampers and kits; savoury pastry (quiches, flans, frozen pizzas).	Other food
Beer	Beers, bitter, lager, ales, cider, perry.	Alcohol
Wine & spirits	Spirits, liqueurs, wine (from grape and other fruits), fortified wines and sherries, alcopops (spirits-based mixed drinks), champagne, sparkling wines.	Alcohol

TABLE A2

Detailed expenditure patterns, TNS and EFS, 2002–2006, levels (£/week) and share of total (%)

<i>RPI category</i>	<i>TNS (unweighted)</i>				<i>TNS (propensity score weighted)</i>				<i>EFS (weighted)</i>			
	2002	%	2006	%	2002	%	2006	%	2002	%	2006	%
Bread	1.33	3	1.57	4	1.25	3	1.48	3	1.97	4	2.32	4
Cereals	1.96	5	2.12	5	1.81	5	2.01	5	2.24	5	2.38	4
Biscuits	2.62	7	2.82	6	2.44	7	2.63	6	2.77	6	2.96	5
Beef	1.00	3	1.08	2	0.93	3	0.99	2	1.34	3	1.62	3
Lamb	0.34	1	0.37	1	0.32	1	0.34	1	0.62	1	0.69	1
Pork	0.49	1	0.49	1	0.46	1	0.45	1	0.58	1	0.59	1
Bacon	0.77	2	0.75	2	0.71	2	0.69	2	0.86	2	0.88	2
Poultry	1.24	3	1.27	3	1.13	3	1.18	3	1.50	3	1.76	3
Other meat	4.44	11	4.87	11	4.20	11	4.69	11	4.83	10	5.07	9
Fresh fish	0.64	2	0.77	2	0.61	2	0.75	2	1.00	2	1.09	2
Processed fish	0.78	2	0.88	2	0.74	2	0.88	2	0.91	2	1.22	2
Butter	0.22	1	0.28	1	0.21	1	0.28	1	0.27	1	0.30	1
Oils & fats	0.59	2	0.61	1	0.55	2	0.57	1	0.61	1	0.67	1
Cheese	1.30	3	1.53	3	1.21	3	1.47	3	1.41	3	1.58	3
Eggs	0.31	1	0.36	1	0.29	1	0.34	1	0.41	1	0.46	1
Fresh milk	1.55	4	1.65	4	1.46	4	1.50	4	2.00	4	2.20	4
Milk products	1.29	3	1.70	4	1.18	3	1.64	4	1.56	3	1.87	3
Tea	0.38	1	0.36	1	0.36	1	0.34	1	0.49	1	0.44	1
Coffee & hot drinks	0.58	2	0.65	1	0.54	2	0.66	2	0.61	1	0.63	1
Soft drinks	2.16	6	2.58	6	1.99	5	2.55	6	2.59	5	3.11	6
Sugar & preserves	0.45	1	0.49	1	0.42	1	0.46	1	0.50	1	0.54	1
Sweets & chocolate	1.40	4	1.66	4	1.32	4	1.58	4	1.85	4	1.88	3
Fresh potatoes	0.57	1	0.62	1	0.53	1	0.55	1	0.76	2	0.71	1
Processed potatoes	1.07	3	1.13	3	0.99	3	1.07	3	1.16	2	1.20	2
Fresh vegetables	2.34	6	2.65	6	2.17	6	2.44	6	3.14	6	3.71	7
Processed vegs.	0.92	2	1.15	3	0.87	2	1.15	3	1.04	2	1.14	2
Fresh fruit	2.04	5	2.33	5	1.88	5	2.07	5	2.68	6	3.34	6
Processed fruit	0.38	1	0.54	1	0.35	1	0.51	1	0.49	1	0.68	1
Other foods	2.64	7	2.86	6	2.45	7	2.81	7	2.90	6	3.19	6
Beer	0.99	3	1.26	3	0.95	3	1.24	3	1.74	4	1.88	3
Wine & spirits	2.59	7	3.16	7	2.48	7	3.31	8	4.29	9	4.86	9

Note: TNS data includes households that purchase in at least 4 separate weeks during a given calendar year, not including the first week of their participation. EFS expenditures exclude households that purchase no food, beverages or alcohol in one or both weeks of their spending diary.

Appendix B: Expenditure ‘fatigue’ regression results

	<i>Baseline</i>	<i>Time controls</i>	<i>Fixed effects</i>	<i>Fixed effects: alcohol</i>	<i>Fixed effects: fish</i>
Week 2	-0.0062 (0.0113)	-0.0071 (0.0113)	-0.0177* (0.0092)	-0.0292* (0.0175)	0.0103 (0.0156)
Week 3	-0.0056 (0.0114)	-0.0131 (0.0113)	-0.0250*** (0.0093)	-0.0274 (0.0176)	-0.0118 (0.0158)
Week 4	0.0335*** (0.0115)	0.0190* (0.0114)	0.0022 (0.0093)	-0.0362** (0.0176)	0.0057 (0.0158)
Week 5	0.0282** (0.0115)	0.0121 (0.0115)	-0.0059 (0.0094)	-0.0138 (0.0176)	-0.0045 (0.0159)
Week 10	0.0207* (0.0118)	0.0106 (0.0118)	-0.0169* (0.0096)	-0.0428** (0.0183)	0.00435 (0.0162)
Week 20	0.0173 (0.0123)	-0.0003 (0.0123)	-0.0352*** (0.0100)	-0.0643*** (0.0192)	-0.0320* (0.0168)
Week 30	0.0183 (0.0126)	0.0002 (0.0126)	-0.0361*** (0.0103)	-0.0418** (0.0196)	-0.0197 (0.0172)
Week 40	0.0367*** (0.0129)	0.0079 (0.0129)	-0.0354*** (0.0106)	-0.0957*** (0.0201)	-0.0138 (0.0177)
Week 50	0.0337*** (0.0131)	0.0067 (0.0132)	-0.0368*** (0.0109)	-0.0507** (0.0203)	-0.0336* (0.0181)
Week 60	0.0127 (0.0135)	-0.0153 (0.0137)	-0.0571*** (0.0113)	-0.0686*** (0.0216)	-0.0162 (0.0188)
Week 70	0.0511*** (0.0139)	0.0196 (0.0142)	-0.0290** (0.0117)	-0.0678*** (0.0217)	0.0178 (0.0194)
Week 80	0.0245* (0.0144)	-0.0040 (0.0148)	-0.0558*** (0.0121)	-0.0955*** (0.0230)	-0.0013 (0.0202)
Week 90	0.0599*** (0.0148)	0.0214 (0.0153)	-0.0388*** (0.0125)	-0.0763*** (0.0233)	-0.00086 (0.0206)
Week 100	0.0266*** (0.0080)	-0.0042 (0.0103)	-0.0615*** (0.0086)	-0.0979*** (0.0163)	-0.0036 (0.0144)
About to drop out	—	-0.0620*** (0.0164)	-0.0917*** (0.0142)	-0.0059 (0.0281)	-0.0093 (0.0241)
Year/month dummies	No	Yes	Yes	Yes	Yes
Year/month began dummies	No	Yes	—	—	—
Fixed effects	No	No	Yes	Yes	Yes
Observations	950,178	950,178	950,178	292,694	398,574
Households	—	—	12,240	11,007	11,905
R ²	<0.001	0.009	0.004	0.033	0.004

Notes: Fixed effects are household fixed effects. Figures in brackets are standard errors. * = significant at 10% level; ** = significant at 5% level; *** = significant at 1% level. “Began” dummies represent the year and month in which the household begins participation. “Drop out” dummy refers to households less than 5 weeks from dropping out of the survey. Results for intermediate weeks, other products and some household groups as described in section III.1 are available on request.

Appendix C: Regression results for logit model for demographic comparisons, 2005

Variable	All observations		1st 2-week observations	
	<i>co-efficient</i>	<i>std error</i>	<i>co-efficient</i>	<i>std error</i>
Male	-0.28	0.109	-0.46	0.179
age <25	1.30	0.181	-0.86	0.304
age 25-29	0.14	0.130	-0.29	0.224
age 35-39	0.14	0.112	0.82	0.214
age 40-44	0.36	0.113	1.17	0.239
age 45-49	0.06	0.122	0.82	0.235
age 50-54	-0.13	0.129	0.79	0.251
age 55-59	-0.28	0.126	0.91	0.264
age 60-64	-0.20	0.131	2.51	0.554
age 65-69	-0.39	0.158	0.63	0.366
age 70-74	-0.49	0.160	1.50	0.436
age 75-79	-0.35	0.159	2.56	0.578
age 80+	0.43	0.150	4.75	1.080
Male*age<25	-0.29	0.246	0.30	0.375
Male*age 25-29	-0.08	0.170	-0.06	0.270
Male*age 35-39	-0.09	0.144	-0.31	0.252
Male*age 40-44	0.02	0.143	0.29	0.285
Male*age 45-49	0.43	0.148	0.87	0.290
Male*age 50-54	0.12	0.154	0.22	0.295
Male*age 55-59	0.18	0.149	0.54	0.310
Male*age 60-64	0.20	0.158	-0.50	0.564
Male*age 65-69	-0.10	0.157	0.10	0.297
Male*age 70-74	-0.59	0.158	-0.74	0.367
Male*age 75-79	-0.41	0.160	-0.24	0.564
Male*age 80+	-0.91	0.151	-1.72	1.130
1 adult	1.19	0.548	0.96	0.475
3 adults	-1.14	0.243	-0.17	0.252
4+ adults	-1.63	0.266	0.02	0.350
1 child	0.97	0.322	0.27	0.349
2 children	1.02	0.325	0.55	0.361
3 children	0.85	0.334	0.23	0.373
4+ children	1.06	0.346	0.22	0.414
Employed 8-29 hrs	-0.79	0.070	-1.08	0.137
Employed <8 hrs	-0.82	0.224	-1.21	0.446
Unemployed	-0.43	0.102	0.16	0.246
Retired	0.48	0.055	-0.16	0.180
Full time education	0.49	0.139	1.97	0.435
Not working	-0.33	0.054	-1.00	0.121
One adult, non pens	-1.16	0.550	-0.69	0.489
One adult, pens	-1.33	0.554	-1.58	0.573
Two adults, pens	0.03	0.090	0.33	0.248
Two adults, w/ children	-1.35	0.328	-0.96	0.361
One adult, w/children	-2.05	0.590	-0.57	0.479
3+ adults, no children	1.33	0.250	0.65	0.309
3+ adults, w/children	0.28	0.402	0.09	0.513
Owner Occupied	-0.25	0.053	1.03	0.212
Council rented	0.03	0.059	1.25	0.226
Other	-0.12	0.128	0.97	0.228
No cars	0.43	0.040	0.38	0.103
2+ cars	0.19	0.039	0.00	0.079
Owns computer	-0.39	0.033	-0.67	0.092
Constant	-2.85	0.142	2.75	0.353

Notes: Dependent variable is an indicator variable which equal one if the observation is from the EFS survey. Since tenure is missing for all household observed in November and December 2005, this data only covers the period up to the end of October. The model is estimated separately in 2005 for those households with missing tenure. Controls also included (but not reported) for Government Office Region and month of observation. Bold text indicates significance at the 5% level.