

Leicester, Andrew

Working Paper

How might in-home scanner technology be used in budget surveys?

IFS Working Papers, No. W12/01

Provided in Cooperation with:

The Institute for Fiscal Studies (IFS), London

Suggested Citation: Leicester, Andrew (2012) : How might in-home scanner technology be used in budget surveys?, IFS Working Papers, No. W12/01, Institute for Fiscal Studies (IFS), London, <https://doi.org/10.1920/wp.ifs.2012.1201>

This Version is available at:

<https://hdl.handle.net/10419/91509>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

How might in-home scanner technology be used in budget surveys?

IFS Working Paper W12/01

Andrew Leicester

How might in-home scanner technology be used in budget surveys?

Andrew Leicester

Institute for Fiscal Studies

Abstract: This paper considers what role in-home barcode scanner data could play in collecting household expenditure information as part of national budget surveys. One role is as a source of validation. We make detailed micro-level comparisons of food and drink expenditures in two British datasets: the Living Costs and Food Survey (the main budget survey) and Kantar Worldpanel scanner data. We find that levels of spending are significantly lower in scanner data. A large part (but not all) of the gap is explained by weeks in which no spending at all is recorded in scanner data. Demographic differences between the surveys accentuate rather than close the gap. We also demonstrate that patterns of expenditure across the surveys are much more similar, as are Engel curves relating food commodity budget shares to total food expenditures. A key finding is that the period over which we observe households in the scanner data significantly alters the distribution, but not the average, of weekly food expenditures and budget shares, which has important implications for whether two-week spending diaries common to budget surveys are giving a truly accurate reflection of a household's typical spending patterns. A second, more involved use of scanner data would be to impute detailed commodity-level expenditure patterns given only information on total expenditures, as a way of reducing respondent burden in budget surveys. We find that observable demographics in the scanner data explain little of the variation in store-specific expenditure patterns, and so caution against relying too heavily on imputation.

JEL Classifications: C81, C83, D12

Keywords: Scanner data, expenditure, inflation, food, measurement

Acknowledgements: Funding from the ESRC Centre for the Microeconomic Analysis of Public Policy (RES-544-28-5001) at IFS is gratefully acknowledged. The author would like to thank Chris Carroll and participants at the December 2011 NBER-CRIW conference on *Improving the Measurement of Consumer Expenditure* in Washington, DC for useful comments; Kantar Worldpanel for supplying the data on which much of the analysis is based; Kate Davies, Gareth Clancy and Neil Price from the UK Office for National Statistics (ONS) for advice on National Accounts and Retail Sales data; and Giles Horsfield of the ONS for information on the use of scanners by the ONS. Data from the UK Living Costs and Food Survey 2001/2–2009 are 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. Views and mistakes are those of the author alone.

1. Introduction

In-home scanner expenditure data are collected via a barcode reader installed in the home.¹ Information about purchases is collected when participants scan the barcodes of any items brought home. Matched with information on prices, stores and the characteristics of participants, such data offer in principle a detailed, complete record of purchasing behaviour. Scanner data have long been used for marketing studies, and increasingly in the economics literature to explore questions relating to consumer, retailer and manufacturer behaviour (recent examples include Griffith et al, 2009; Broda et al, 2009 and Aguiar and Hurst, 2007). For researchers, the appeal of scanner data lies both in the detailed purchase information and in the fact that the data are typically longitudinal. Panel expenditure datasets are comparatively rare. National budget surveys are usually cross-section, and measures of spending in panel data (such as the UK Household Longitudinal Study) tend to be limited and highly aggregated. Panel data offer the chance to explore changes in purchasing behaviour in response to shocks or policy interventions. For example, Harding et al (2012) use scanner data to explore the impact of changes in cigarette taxes on prices and consumption. There is also growing interest amongst policy makers in what can be learned from scanner data. The UK Department for the Environment, Food and Rural Affairs carried out a study of ethical shopping decisions in conjunction with a commercial scanner data collector (DEFRA, 2011). Scanners are also starting to feature in non-commercial surveys: the US Department of Agriculture is planning to use scanners, alongside other data collection methods, to record detailed food purchase behaviour as part of the Food Acquisition and Purchase Survey (FoodAPS).²

The main aims of this paper are to consider what role scanner data could play for collecting household expenditure information as part of budget surveys such as the US Consumer Expenditure Survey (CE) or the UK Living Costs and Food Survey (LCF). At one level, the role could be one of validation. Comparisons of budget survey data to aggregate data have led to increasing concern about the quality of survey expenditure data. However, there has been little scope to make micro-level comparisons since few surveys (besides the budget survey itself) collect detailed spending information. Scanner data offer such a possibility.

A more involved role for scanner data in budget surveys might be as part of the data collection process itself (Mathiowetz et al, 2011). This could involve scanners being used in place of or alongside current survey methods such as paper diaries and recall questions. The key question for statistical agencies is to understand the modal effect of using scanners on the data which is obtained. Comparative studies between scanner and other expenditure data offer some insights here, but fully disentangling modal effects from other differences between surveys (such as demographic and sampling differences) is likely to require experimental methods.

¹ We should distinguish these in-home data from 'storescan' datasets collected at the store- or chain-level from items passing through the tills. A previous NBER volume, Feenstra and Shapiro (2003), considered the possible usefulness of store-level scanner data for measuring prices and price indices. There are clearly complementarities between store-level and consumer-level scanner data. The main advantage of the former is that it provides in principle a complete record of purchases in that store without problems of sampling variation or mis-measurement; the advantage of the latter is that it matches what is bought to who is buying.

² See http://www.ers.usda.gov/Briefing/SNAP/food_aps.htm and Cole (2011).

Scanner data could also be used for imputation. As a way to reduce respondent burden, some commentators have suggested asking only limited questions about household aggregate expenditures in the survey, and using information from other sources such as scanner data to drill down into how this breaks down in more detail.³ The value of scanner data from this perspective is that we know not just what was bought, but also the store of purchase, which would allow disaggregation by store type. This is important because spending patterns (even within the same broad aggregate commodity like food) may vary by store. Evidence on this role for scanner data is provided in Section 5.

Our focus in this paper is on how statistical agencies might make use of established in-home scanner surveys collected by commercial market research companies. However, an alternative would be for statistical agencies to establish and maintain their own scanner data. Whilst probably costly, this would offer a number of advantages. It would override concerns about effectively outsourcing part of the data collection process to commercial organisations. It would allow controlled experimentation to explore modal effects, incentive effects and so on. It would ensure that the information on, for example, demographics and demographic transitions was of sufficiently high quality to be useful for national statistics and research purposes. It would also ensure that the statistical properties of the scanner data were the same as those of the budget survey data (as we discuss below, commercial scanner data are normally collected using quota sampling). Further, scanner methods offer the possibility of collecting detailed spending information over a longer period than is common in diary-based budget surveys. The UK LCF, for example, asks households to report spending only for two weeks. We find evidence that whilst this period is sufficient to capture the average pattern of how food expenditures break down across households, it does not give an accurate record of the distribution of detailed expenditure patterns when compared against long-term spending habits.⁴ We offer some thoughts on the scope for establishing a separate scanner survey, and how it might overcome some of the limitations of existing scanner datasets, in the conclusions.

The rest of the paper is organised as followed. Section 2 describes the datasets which underlie most of our analysis: the UK Living Costs and Food Survey (the main UK budget survey) and in-home scanner data from Kantar Worldpanel. Our discussion is based on UK data but the main lessons are almost certainly transferable to the US context. Section 3 surveys the existing literature comparing scanner data to other data. In Section 4, we make direct comparisons between scanner data, budget survey data and aggregate spending information from retail sales and national accounts data to try and understand better the differences between datasets and what might be driving them. We assess the impact of using budget shares taken directly from the budget survey and scanner data as basket weights in calculating food price indices, in place of current weights derived from aggregate expenditure data. We also explore how expenditure patterns vary with the duration for which we observe household spending. Section 5 explores the prospects for using detailed spending patterns from scanner data to impute budget shares for households when all we observe are total expenditures. Section 6 offers some overall thoughts and conclusions.

³ See for example Tucker (2011), *Using Multiple Data Sources and Methods to Improve Estimates in Surveys* (http://www.bls.gov/cex/hhsrvywrkshp_tucker.pdf)

⁴ This has important implications for attempts to use short-run spending information to make inferences about the distribution of living standards (for example Attanasio et al, 2006; Brewer et al, 2006)

2. Data

2.1 Living Costs and Food Survey

The Living Costs and Food Survey (LCF) is the main UK source of household budget information. Collected by the Office for National Statistics (ONS), it is an annual cross-section of around 6,000 households and began in 1957. The survey has undergone some significant changes over time. Until 1993, it was carried out on a calendar year basis, but switched to a fiscal year (April–March) basis from 1993/4 and then reverted back to a calendar year basis from 2006. It was originally known as the Family Expenditure Survey (FES). From 2000/01, it merged with the former National Food Survey – the main source of micro-data on household nutrition in the UK – and was renamed the Expenditure and Food Survey (EFS). From 2008, the survey was renamed the Living Costs and Food Survey, and brought into the wider Integrated Household Survey (IHS). We use LCF throughout to refer to this data.

Sampling is carried out via stratified random sampling, with strata based on region, socio-economic status and car ownership. Northern Ireland is over-sampled but survey weights are provided to ensure the weighted sample is nationally representative. The response rate in 2010 was 50%; this has declined substantially in recent years.

The data are made up of two main parts. The first is a two-week diary issued to all household members aged 16 and over (children aged 7 to 15 receive a simplified diary).⁵ Participants record all their expenditures over the period (including anything purchased with credit cards). Where possible, they are asked to attach till receipts to the diary to reduce the extent to which handwritten records of spending have to be maintained. A £10 incentive (£5 for children) is paid for successful completion of the diary. Household members are also interviewed to obtain detailed demographic and income information, as well as data on large irregular purchases (such as furniture and holidays) and regular expenses like household energy and housing payments. Data from the diary and the questionnaire are coded into a large number of separate spending items for each household, all of which are reported on a per-week basis. Details of methods and the main findings are collated each year into an ONS publication *Family Spending*.⁶

2.2 Kantar Worldpanel

Kantar are a market research company, formerly called Taylor Nelson Sofres (TNS). They operate a number of surveys of consumer behaviour. One of the largest is Worldpanel, carried out in a number of countries. In the UK, Worldpanel data is used, amongst other things, to estimate market shares of the major supermarkets.⁷

We have access to Worldpanel data covering a nine-year period between November 2001 and November 2010. A large, representative sample of households is active in the data at any one time. Until 2006, the average sample size was around 15,000 households. Since then it has risen to around 25,000. Participants are recruited from a range of address sources, though geographic coverage is limited to Great Britain (Northern Ireland is not

⁵ Children were first asked to keep a diary in 1995/6. We use LCF data including spending reported by children, since in principle children's purchases should also be captured in scanner data.

⁶ The report on the 2010 LCF data is available from <http://www.ons.gov.uk/ons/rel/family-spending/family-spending/family-spending-2011-edition/family-spending-2011-pdf.pdf>.

⁷ See <http://www.kantarworldpanel.com/>.

included). Recruitment uses quota sampling methods. Household-specific weights are derived which ensure that the weighted sample (over a particular period of observation) of active households is representative based on household size, housewife age, social class and region. Households can remain members of the Worldpanel for as long as they wish. Ongoing participation is rewarded with points redeemable for consumer goods.

Participating households are issued with a barcode reader which is installed in the home, and are asked to record the purchases of all barcoded products brought home. Our data contain information on all 'fast-moving consumer goods' – essentially food and grocery products, including things like cleaning products, personal care items and so on. The data includes alcohol purchases (from off-licences and supermarkets) but not tobacco or, oddly, baby foods. Purchases from all retailers, not just supermarkets, are in principle recorded, as are online grocery purchases. The data are at the transaction level, and record detailed information on the characteristics of the products purchased.⁸ Since 2006, these characteristics have included the macro-nutritional composition of foods purchased (carbohydrates, fats, protein and so on). Until 2006, all households were asked to report non-barcoded food and grocery purchases using a booklet of generic barcodes. Details of the product characteristics for these items (such as weight, country of origin, flavour and so on) were also entered manually via the scanner device. This increases the responder burden for these items. From 2006, some households were no longer required to report these items and were issued with a simpler scanner unit. Records of purchases are transmitted to Kantar either by phone line or via a home computer.

Information on the price paid is obtained from till receipts which are mailed in to Kantar who match the price to the purchase record. Where no receipts are available, prices are taken from centralised databases of store- and product-specific prices, or otherwise imputed. The data also record any promotional deal attached to a purchase. Information on the store visited is recorded by the participants.

Demographic characteristics of the participants are recorded at a baseline telephone interview held at sign up, and then updated every nine months or so. We have annual snapshots of the demographics of all households. The set of demographic questions is typically much less comprehensive than those recorded in the LCF, and an interview is held only with the 'main shopper' in each household rather than with each household member separately. All household members should, though, report their expenditures.

2.3 To what extent is household expenditure 'scannable'?

A basic question for in-home scanner data is what types of spending could, in principle, be scanned. Most obviously, items not brought into the home would be excluded.⁹ Scanner methods are also clearly less suited to items without barcodes. They could be covered by generic barcodes supplemented with user-supplied information, but this is costly in terms of survey effort and would be hard to apply for product categories where barcodes are not used at all, such as services. At the very least, it is hard to see what the advantage of in-home scanning for these purchases would be over a written diary or asking participants to keep hold of receipts.


⁸ A typical week of data can record up to a million separate transactions.

⁹ One possibility might be to use mobile devices to record these items – a survey of different technologies can be found in Westat (2011a).

In principle, what proportion of expenditures could be captured by in-home scanners? A rough estimate can be obtained by adding up the UK Consumer Prices Index (CPI) weights of commodities which would be most amenable to scanning. Defining this set of products is a matter of judgement, but it seems reasonable to assume that services, energy products and durable goods (which are often delivered without packaging) are not scannable. The remaining CPI items are then the set of scannable goods: food and beverages at home; tobacco and off-licence alcohol; non-durable, non-energy goods (e.g. cleaning products, personal care items, newspapers and goods for home maintenance and semi-durable goods (e.g. clothing, books, tools, household ornaments, DVDs).

Figure 2.1 shows the total CPI basket between 1996 and 2011 broken down into 'scannable' categories (coloured green) and 'non-scannable' categories (shaded grey). The total scannable CPI weight fell from around 42% in 1996 to around 35–36% by 2003, and has been at a similar level since then. This was mainly caused by the growth of services in the late 1990s and early 2000s, coming mostly at the expense of food, alcohol and tobacco. More recent years have seen a fall in the share of the CPI covered by services but this has been offset by growth in the energy component (also not scannable) as oil prices have risen. Thus at best just over a third of total expenditures (by CPI weight) appear to be readily amenable to in-home scanner technology.

Figure 2.1: The extent to which the UK CPI is amenable to in-home scanning, 1996–2011




Source: Calculated from UK Office for National Statistics Data. Notes: Green colours indicate categories which may be best suited to in-home scanning. Grey colours indicate those which are not suitable. The ONS codes used to calculate these categories are Food and Beverages (CJUX+CJUY); Alcohol (CJUJ); Tobacco (CJWP); Non-durables excluding water (A9ET - CJWW); Semi-durables (A9EU); Energy (A9F3); Durables (A9ES) and Services (ICVI).

Commercially-available data sources, which might be used by agencies as an alternative to implementing a new in-home scanner dataset, are typically even more limited than this. For example, Duly et al (2003) compared AC Nielsen Homescan data in the US to CE

Diary Survey data from 2000, and found that only 45% of the CE diary items (but 83% of food items) were covered in the Homescan data.¹⁰

We approximate the proportion of consumer spending captured by the Kantar Worldpanel data in the UK as the total weight of the following CPI categories: food and non-alcoholic beverages; alcohol; non-durable household goods (e.g. matches, foil); pharmaceutical and medical products (e.g. sticking plasters, medicines); personal care products (e.g. toothbrushes, deodorants) and pet care (e.g. pet food, cat litter). This probably slightly overstates the coverage (as this definition includes things such as make up, electric razors, contact lenses and NHS prescription charges which are not in the Kantar data) but is a good approximation at the publicly-available level of CPI disaggregation. Based on this, the Worldpanel data currently covers something like 18% of all expenditure, having fallen from around 24% in 1996 (see Figure 2.2). Not surprisingly, this is driven mainly by the CPI weight of food, beverages and alcohol, which fell from 19% in 1996 to 12% in 2006.

Figure 2.2: The extent to which the UK CPI is covered by Kantar Worldpanel data, 1996–2011



Source: Calculated from UK Office for National Statistics Data. Notes: The ONS codes included in this total are Food and Beverages (CJUX+CJUY); Alcohol (CJUJ); Non-durable household goods (CJXK); Pharmaceutical products (CJYA); Other medical and therapeutic equipment (CJYH); Appliances and products for personal care (CJYO) and Pets, related products and services (CJYJ).

¹⁰ The paper is internal to the US Bureau of Labor Statistics, but the results are cited in Garner et al (2009).

3. Previous research

A number of previous studies have compared scanner data to budget survey data. These comparisons shed light on the possible modal effects of using scanners to collect spending information, both in terms of what is reported and who participates. A common finding is that raw average expenditure levels are markedly lower in scanner data than in budget survey data. Duly et al (2003) compare AC Nielsen Homescan data with CE Diary Survey data from 2000. They find food at home spending in Homescan was about two-thirds of the spending reported in the CE. Alcohol and tobacco spending was only half that reported in the CE. Zhen et al (2009) also compare Homescan and CE data, this time for 2002–2005. They map expenditures into 18 categories. Homescan spending was around 50% below that reported in the CE for a number of categories, including beef, poultry, pork, fresh fruit and vegetables, eggs, fish and fats. By contrast for several categories – sugar and sweets, other meats, processed fruits and vegetables, miscellaneous foods and other dairy products, they find essentially no differences between CE and Homescan spending. Notably, the categories where expenditures are substantially lower in Homescan were those with a substantial proportion of random weight items which do not have barcodes, and only around 20% of the Homescan panel in their data recorded such purchases. Thus they find substantial variation in expenditure patterns in the surveys: for example, in the CE data fresh fruit is the 6th largest item amongst the 18 categories studied, but ranks only 12th in the Homescan data.

In the UK, Leicester and Oldfield (2009) and Griffith and O’Connell (2009) compare Kantar Worldpanel scanner data to the Living Costs and Food Survey (LCF) budget data. Leicester and Oldfield (2009) compare expenditures whilst Griffith and O’Connell (2009) use the fact that both surveys also record information about the nutritional content of food purchases to compare macronutrient intakes. In common with US studies, Leicester and Oldfield (2009) find that average spending is lower in the scanner data. In 2005, average weekly household spending on food and alcohol in the Worldpanel was £42.52, compared to £53.15 in the budget survey – a difference of some 20%. However, in contrast to the Zhen et al (2009) study, they found little difference across commodities in these expenditure gaps and thus very similar patterns of spending in the two surveys. The exception was alcohol, where Worldpanel spending was just 58% of LCF levels compared to 71% to 86% across other food groups. One reason for the relative lack of variation across commodities compared to the US results is that before 2006, all Worldpanel households were asked to record non-barcoded items. From 2006, however, only a subset of households was asked to do so. The authors find this has a significant effect on reported spending. In 2006, households not reporting random weight items spent 44% less on fruit than households in the LCF; the gap for those reporting random weight items was just 24%. However, those not reporting random weight items appeared to report more spending on other goods and more spending overall. Kantar suggested this was because the scanning device given to these households was simpler to operate, though Leicester and Oldfield (2009) suggest it may also be driven by the demographic characteristics of those not asked to report random weight items.

Griffith and O’Connell (2009) find a similar gap between Worldpanel and LCF in terms of calorie intake. They conclude that the size of the gap between the scanner and diary data depends substantially on the treatment of weeks in which no purchases are recorded. Amongst households who remained in the Worldpanel data for all of 2006, average

calorie purchases among single adult households amounted to 1,850 per day compared to 2,282 in the LCF, a gap of about 19%. When they look just at the longest period of purchases without a zero spending week, however, the gap shrinks to just 1% (though remains larger for other household types). Their conclusion that "... the treatment of periods during which there are zero purchases is one of the reasons that differences arise between the [datasets] ... it is important to understand more about why such periods arise and how they should be treated" is critical, and not yet fully understood.

They also show that dietary patterns vary according to the period over which households are observed in scanner data. Observed over a full year, the proportion of calories obtained from saturated fat is higher than when households are observed for shorter durations. We explore both these issues – of zero purchase weeks and changing purchase patterns over different observation periods – in Section 4.2.

One reason for the difference between spending reported in scanner data and in household budget survey data is survey mode. Difficulties in scanning random weight items have been discussed already. It may also be the case that certain types of shopping trips or purchases are less well recorded in scanner data than others – for example, households may forget to scan small 'top up' trips. Leicester and Oldfield (2009) find some evidence for this: for example, bread and milk (traditional 'top up' items) are relatively underreported in the Worldpanel data compared to the LCF. An interesting exercise to explore this, only possible in the UK using detailed versions of the LCF data which are not made freely available to researchers, would be to compare not just levels and patterns of spending, but also the pattern of grocery shopping trips (e.g. the number of items purchased, frequency of shopping and so on).

Another explanation, however, could be the demographic composition of the samples. Interestingly, results from the US and UK reach different conclusions about the extent to which demographics account for the differences in spending. Zhen et al (2009) use a regression model to strip out observable demographic differences between datasets and argue that, in combination with the random weight issue, this largely accounts for the spending differences. Leicester and Oldfield (2009) conclude that demographic differences *accentuate* the gaps between datasets. They estimate a 'propensity weight' for each household in the Worldpanel data which reflects how similar its observed demographics are to those of LCF households. Using this weight, they find that the average gap between total spending in the two datasets rises from 20% to 25%.

These differences could result from quite different findings across the countries about the nature of the scanner samples compared to the budget survey samples. In the US, for example, Homescan households appear to have fewer members on average than those in the CE (Huffman and Jensen, 2004) whilst in the UK, Kantar households have more members on average (Leicester and Oldfield, 2009). In terms of income, there is conflicting evidence in the US: Huffman and Jensen (2004) suggest Homescan households have higher incomes than those in the CE, whereas Perloff and Denbaly (2007) find that scanner data households (in data collected by IRI) have lower incomes than Census averages. In the UK, Worldpanel households were found to be poorer on average and more likely to be unemployed, or part-time workers (Leicester and Oldfield, 2009).

Whilst there is a consistent finding across countries that expenditures are lower in scanner data than budget surveys, there is then no common consensus on what drives this. Without an experimental approach which holds, as far as possible, other differences constant, it is hard to make any definitive assessment of the pure survey mode effect.

The demographic information available in scanner data is often much less comprehensive than that found in budget surveys. Kantar Worldpanel, for example, did not routinely collect information on household incomes until 2008, and even then only a banded measure of gross total income is collected from a single question asked of the main shopper. By contrast the LCF contains detailed questions on unbanded incomes by source for each household member. Similarly, information in the Worldpanel on education and employment status are not consistently collected for each adult household member, and common variables like tenure are also not always reported. The problems may be less acute in other surveys and other countries, but in general it is important for any users of scanner data to bear in mind that existing data are collected for market research purposes rather than as a source of national statistics or for social science research. As a result, if there were a serious move to use existing scanner data in budget surveys, it would be important for the statistical agency to work with the data collectors, perhaps to obtain more detailed and useful demographic information about the respondents, and certainly to share knowledge about issues in how the scanner data are recorded. A particular problem noted by Leicester and Oldfield (2009) was poor reporting of demographic *transitions* over time in the Kantar data. For example, using data from 2002 to 2005, they find that amongst a sample of households headed by someone employed and aged 50 or over, just 2.9% were observed to be unemployed a year later. This compared to 11.4% of a similar sample constructed from the British Household Panel Survey (BHPS), the main panel data set in the UK. This could reflect transitions being more associated with attrition in scanner data than in social science survey data, but seems likely to at least partly reflect a failure to maintain the demographic information about households over time. There does not appear to be similar evidence from any study of other scanner panels which would shed light on whether this issue was common to scanner data in general or particular to the Kantar Worldpanel.

It may well be the case that there are *unobservable* differences in the characteristics of households in different surveys. A fascinating paper by Lusk and Brooks (2011) suggests that households in two large US scanner samples, Homescan and IRI, appear to be more price responsive than the population at large. They take random samples from the scanner surveys and from a random-digit-dialing survey of all households, and use web-based questionnaires to elicit a series of discrete choices over grocery products from which elasticity estimates are derived. Even controlling for observable differences, they find that those in scanner samples are more price responsive – for example, the average whole milk own-price elasticity was -1.39 in the national sample compared to -1.51 in the IRI sample and -1.54 in the Homescan sample. They offer two possible explanations. Firstly, participating in scanner data may make households more aware of their purchasing behaviour and thus more price sensitive. Secondly, those who agree to participate in scanner data may be a self-selected sample of more price conscious households. Of course, these findings do not tell us whether the participation and self-selection effects are greater in scanner data than in budget surveys. It is possible that those who agree to participate in the CE or LCF are also more price responsive, or become more so as a result of participating. We may expect the participation effect to be smaller

for the budget survey which takes place over a limited period whereas participation in the scanner survey is potentially long-term, though in the US households participate in the CE (interview survey) for up to five quarters. It would be nice therefore to carry out a similar exercise on a sample who had previously participated in the CE data.

With in-home scanners, there may concern about people taking some time to adapt to the technology before they report reliable data. Indeed, the Homescan survey monitors early records of new panel members and only includes them into the reported data after a few weeks of regularity checking (Harris, 2005). The Worldpanel data studied by Leicester and Oldfield (2009) includes all purchases reported by households irrespective of length of participation. They found that expenditures were highest in the first few weeks of participation but fell away slightly such that after about 6 months households spent, on average, about 5% less than in their first week. This might be evidence of survey 'fatigue', with households being less assiduous about reporting all their spending after the initial novelty wears off. It might also be evidence of a settling-in process in which households make small errors early on (multiple recording, say) which inflate expenditures relative to their true values. It could also be a genuine behavioural reaction to participation. For statistical agencies thinking about scanner data for budget surveys, the interesting comparison is with the current survey approach. For example, Ahmed et al (2006) find that in the Canadian Food Expenditure Survey, collected over two weeks, spending drops by 9% between the first and second week on average – larger than the decline in spending found in the Kantar data over a 6-month period.

Comparing scanner data to other expenditure records provides additional insight into data quality issues. Einav et al (2008), for example, perform a detailed matching exercise of shopping trips at a particular store, comparing purchases reported in Homescan data to what should be the same shopping trips in loyalty card records. They find that 20% or so of trips recorded in Homescan were not matched in the retailer data, suggesting inaccurate store or date information in Homescan. From their best estimates, around half the trips that were reported in the store data were not observed in Homescan. Where trips could be matched up, the Homescan record appeared to miss on average 10-15% of the items purchased in the store record, most often small consumables like soft drinks which may be consumed on the way home. However the largest problems were found with price records: on matched trips, the price reported in the Homescan record failed to match the loyalty card recorded price about half the time. This seems to have been driven mostly by the way prices are imputed into Homescan data. To avoid households having to input the price of each purchase manually, purchases are linked to a chain-level database which holds the price of each product in each week. This means any store-specific discounts or offers will be missed, as would any individual loyalty card or coupon discounts. Whilst no similar evidence is available for the UK (though in principle a similar comparison would be possible if data could be obtained from the stores) it is unlikely that the price issue is so pronounced, since participants in Worldpanel are asked to mail in till receipts from their shopping trips and prices are matched from these receipts to the individual purchase records.¹¹ Einav et al (2008) recommend that a similar approach be adopted for Homescan.

¹¹ Where receipts are not sent in, similar imputation methods may be used. Since national supermarket chains in the UK all use national pricing, this imputation should still capture chain-level deals and promotions, though individual discounts from coupons or loyalty cards would be missed.

One issue raised by a number of papers and commentators relates to quota sampling methods used in scanner data (Tucker, 2011; Zhen et al, 2009; USDA, 2009; Harris, 2005). National budget surveys use random probability or stratified sampling instead. If a statistical agency were planning to introduce a scanner component into a random sample then this may not be a problem (other than possibly non random refusal to use the technology), but if the idea were to use existing scanner data then this might be a drawback for any statistical comparisons or imputation across surveys. Westat (2011b) and Perloff and Denbaly (2007) are both critical of commercial scanner data collectors for releasing little information on sampling methods, response rates, attrition rates and so on, and suggest caution in relying on existing scanner data for these reasons.

The main lesson from the existing literature seems to be that there is no single approach to how scanner data are collected (just as there is no single approach to how different countries implement their budget surveys), and the sorts of errors and problems that there may be with a given scanner dataset depend very much on precisely how that data is collected. Thus any statistical agency (or indeed researcher) planning to use scanner data ought to be aware in detail of the methods that underlie its collection.

4. Comparing scanner and other expenditure data

For any source of household expenditure data, a natural question is whether accurate information is recorded. Participants could deliberately or accidentally mis-report their purchases, or change their usual shopping behaviour as a result of participation. Data validation is vital. Without any clear way to obtain a 'gold standard' benchmark of actual expenditures against which to compare surveys, the most promising approach to validation is to compare data sources against one another to see whether they provide different impressions of spending levels, patterns and trends.

In this section, we make two distinct sets of comparisons. First, we compare UK expenditure survey data from Kantar Worldpanel and the LCF to aggregate data from the ONS National Accounts and Retail Sales. Second, we compare expenditures across the Worldpanel and LCF surveys. We explore not just total spending but also patterns of spending across commodity. Differences in total spending will matter for issues like living standards and inequality where spending is used as a measure of well-being. But in some cases it is the pattern of spending that matters – for example, in deriving expenditure weights for price indices. Comparisons of both are therefore important.

Mapping Kantar expenditure data to LCF data

Before detailing the results of these comparison exercises, it is worth saying something about how we match expenditures in the Kantar and LCF datasets.

The Kantar data are reported at the barcode level. There are more than 568,000 individual products. The LCF records household-level expenditures in a large number of fairly disaggregated expenditure codes. The challenge is therefore to match individual products from the Kantar data into equivalent LCF expenditure codes. We make use of detailed information on the sorts of items which make up each expenditure code provided with the LCF documentation, and the detailed product-level characteristics supplied with the Kantar data, to make this match as accurately as we can. We then further aggregate expenditures into broader commodity groups. We choose to aggregate to the commodity groups which underlie the Consumer Prices Index (CPI). In part, this is because this is the level at which disaggregated expenditure information is available from the National Accounts, so making comparisons to aggregate data more straightforward. This aggregation also still generates an informative set of broad commodities (9 food groups, 2 non-alcoholic beverage groups and 3 alcohol groups) as a basis for comparison. We exclude all non food and drink purchases, as we cannot be completely confident that the set of products we observe in the Kantar data represent the full set of items which make up the relevant non-food CPI commodities.

In principle, of course, when making comparisons across datasets, we could look at much more disaggregate commodity groups. The (publicly-available) LCF includes 73 distinct food and drink codes, so this would be the most disaggregate comparison possible.¹² An alternative aggregation would be to the level of the Retail Price Index (RPI), which is broken down into 31 food and drink commodities. Finer disaggregation may be useful to understand exactly where differences between scanner and other datasets arise and what might be driving that. However, as discussed in Leicester and Oldfield (2009) who use the

¹² Statistical agencies would, of course, have access to even more disaggregate budget survey data.

RPI as the starting point for their comparison of Kantar and LCF data, the more disaggregate the comparison the less confident we can be about the mapping between Kantar and LCF expenditures.¹³ The problems are particularly acute where it is not clear in the Kantar product information whether meats, fish, fruits and vegetables are ‘fresh’ (largely meaning unadulterated, so including e.g. plain frozen fish fillets) or ‘processed’ (largely meaning they are pre-prepared or flavoured in some way), but there are different LCF and RPI groups for fresh and processed goods.¹⁴ The CPI provides more certainty that we are comparing like-with-like (spending on ‘fish’ or ‘meat’, say) at the cost of aggregation. However, the challenge of accurately mapping many thousands of different products from scanner data to the budget survey should not be underestimated.

4.1 Comparison to national expenditure aggregates

Several recent papers have explored the quality of budget survey data by comparing them to national aggregate household expenditure data. Examples in the US include Triplett (1997), Slesnick (2001), Attanasio et al (2006), and in the UK include Tanner (1999), Blow et al (2004), Attanasio et al (2006) and O’Dea and Crossley (2010). Key findings from these studies are:

- In the US, expenditures reported in the CE make up about 70% of those reported in national aggregates. The trend is worsening over time: the figure fell from around 80% in the mid-1980s to 60% in the early 2000s.
- The proportion of national expenditures recorded in the UK LCF is somewhat higher, around 80% or so on average. But again the trend is worsening over time, from almost 90% in the mid-1970s to under 70% by the mid-2000s. The decline in UK coverage is particularly noticeable from the early 1990s.
- For food at home, both surveys perform better but the trends are similar. In the US, the coverage of national accounts food spending in the CE data fell from more than 75% in the 1980s to around 65% by the early 2000s. In the UK, coverage fell from more than 95% in the 1970s to less than 90% by the early 2000s.

We focus on food at home and off-licence alcohol expenditures in the UK National Accounts (NA) and compare them to spending reported in the LCF and, for the first time, Kantar Worldpanel data. We are not aware of other papers which have made similar comparisons of scanner data to national spending figures in this way. Our interest is not just in how much of total aggregate expenditure is reported in the surveys, but also in whether trends over time are similar. There are a number of reasons why we would expect food spending to be higher in NA data than survey data. NA figures include expenditures by people living in non-housing accommodation (student halls, old age homes, army barracks and so on) and spending by tourists in the UK which are not included in the surveys. NA figures are also based on UK-wide expenditures (including Northern Ireland). Since the Kantar data covers only Great Britain (excluding Northern Ireland), for consistency across surveys we also look at LCF data for Great Britain.

We would also expect a closer correlation between LCF data and the NA than between Kantar data and the NA simply because the LCF is one of the main inputs used to derive

¹³ The bulk of their analysis is carried out on 13 commodity groups aggregated from the 31 RPI groups.

¹⁴ There is also often ambiguity as to which LCF code a particular Kantar product should be allocated – to give one example, LCF code 11151 (‘other breads and cereals’) seemingly includes ‘filo pastry’ and ‘puff pastry’ whilst code 11931 (‘yeast, dessert preparations and soups’) includes ‘ready-made pastry’.

NA figures.¹⁵ We therefore make an additional comparison to a measure of expenditures as reported in Retail Sales (RS) data, which are based on a survey of retailer turnover.¹⁶ In principle, we would expect aggregate food expenditures in the LCF and Kantar data to line up quite closely with total food retail sales. Again, though, the RS figure will include purchases made by foreigners in the UK and those living in institutional accommodation. Unlike the NA or survey data, the RS figure will also include purchases made for non-final consumption (e.g. purchases by restaurants or caterers from food retailers). In addition, RS data are not available at a disaggregated commodity level. We observe only an aggregate measure of total expenditure on food, drinks (alcoholic and non-alcoholic) and tobacco. As the Kantar data exclude tobacco, we also remove it from the LCF figures to provide a cleaner comparison of the survey data. However, unlike the NA data, the RS figures are reported for Great Britain rather than the UK which means the geographical coverage matches those of the survey data.

To make the comparisons, we need to aggregate the LCF and Kantar survey data to national totals. The LCF reports weekly household-level expenditure by commodity group and provides sampling weights for each household which gross up the data to national figures. Thus we convert weekly expenditure figures to annual figures (multiplying by 52) and use the weights to generate national annual expenditures. In the Kantar data, sampling weights are provided for each household covering different periods of time (e.g. 4 weeks, 52 weeks). Households who fail to report expenditures consistently over that period are assigned a zero weight, with the weights of other households adjusted such that the figures gross up to national totals. We calculate total aggregated expenditures by commodity over a series of 4-week periods using the appropriate weights. These are then further aggregated into annual totals by adding up the thirteen 4-week periods that generate a 52-week 'year' which closely (though not perfectly) covers a single calendar year.¹⁷ Since the RS figures are reported as average sales per week, we convert all aggregated figures from the NA and survey data to per-week averages.

Figure 4.1 shows average total weekly expenditure in the two ONS datasets (RS and NA) and the survey datasets (LCF and Kantar) between 2002 and 2009, the period for which full-year comparisons can be made. We make detailed comparisons of the survey data and NA figures below so for the moment focus on the RS comparison. Levels of spending in the LCF are about 67 to 69% of those in RS data (though recall that tobacco is excluded from the LCF data). Those in the Kantar data are about 44 to 46% of those in the RS data. The fact that spending in the Kantar data is lower than the LCF is consistent with our direct comparison of these datasets in Section 4.2. There is no obvious deterioration in the proportion of RS figures captured in either dataset over this period.

Figure 4.2 shows year-on-year growth rates across the different datasets. Strikingly, there is much more volatility in the LCF figure: for example aggregate spending on food and drink grew by around 8% between 2005 and 2006, a spike not seen in other data. A sharp slowdown in spending in 2009 evident in the NA figure was also seen in the Kantar data but, interestingly, not in the RS figure or the LCF. Indeed the divergence in 2009 between


¹⁵ For information on data used to compile NA expenditure figures, see Office for National Statistics (2010a). The LCF is not the only input into NA data; for example, adjustments are made based on retail sales information, and to try and take account of illicit consumption of items like tobacco and alcohol.

¹⁶ See <http://www.ons.gov.uk/ons/rel/rsi/retail-sales/august-2011/index.html> for data. The figures give an index value each year and a cash value of average retail sales per week in the base year (2006) from which we derive the cash value of sales in other years.

¹⁷ For example, the period labelled 2002 in the Kantar figures covers Jan 7th 2002 – Jan 5th 2003.


the RS and NA figure is very noticeable; until then the two ONS sources had tracked one another quite closely. Certainly, there is no evidence that the Kantar data perform worse when compared to national expenditure figures than the LCF survey data; indeed, the larger sample size in the Kantar data helps mitigate against the sort of volatility observed in the LCF. Such volatility might caution us against making year-on-year inferences about changes in living standards from consumption changes in the LCF.

Figure 4.1 Gross weekly expenditures, ONS and survey data, 2002–2009



Source: Calculated from UK Office for National Statistics Data, LCF data and Kantar Worldpanel.

Figure 4.2: Growth rates of aggregate expenditure, ONS and survey data



Source: Calculated from UK Office for National Statistics Data, LCF data and Kantar Worldpanel.

We now go on to look at more disaggregated comparisons between the LCF and Kantar surveys and the National Accounts measure of spending. Table 4.1 shows per-week CPI commodity-level gross expenditures across datasets in 2002 and 2009, the start and end of our comparison period, and the fraction of NA expenditure reported in each survey. As seen above, expenditures in the LCF and Kantar surveys are lower than those in the NA data.¹⁸ Comparing across commodities, it is clear that food is better reported than either alcohol or non-alcoholic beverages in both survey datasets. Within food, spending on bread and fish appears to be particularly well captured in survey data. There appears to be some issue in how sugar and confectionery products are reported as well as 'other food'. Given what appear to be high relative reports of survey expenditure on other food and low reports of sugar and confectionery, it could be some coding issue where items that are included in the NA definition of confectionery are included in the LCF definition of other food. It is not clear what drives this, and we cannot drill down into the NA figures in more detail. In principle, the LCF expenditure codes should map directly on to the NA commodity codes since both use the COICOP (Classification of Individual Consumption by Purpose) categorisation method. Since we map the Kantar products onto LCF codes, these too should then translate directly into comparable NA commodities.

The relative under-reporting of drinks compared to food is striking. Interestingly, the LCF data appear if anything relatively worse at recording non-alcoholic drinks than alcoholic drinks whilst the reverse is true in the Kantar data. Also striking is that whilst over time there has been on average no noticeable deterioration in the proportion of NA food expenditure reported in either survey, the proportion of beverage and alcohol spending reported in the LCF has fallen whilst in the Kantar data there is no clear trend for drinks. It is easy to imagine that the relatively poor reporting of soft drinks might be driven by the high frequency with which they are purchased by children or as part of smaller shopping trips. The fact, though, that by 2009 only 54% of soft drink expenditure in the NA was captured in the LCF compared to 75% of wine expenditure and 61% of beer expenditure is quite surprising – much attention has been paid to the relative under-reporting of alcohol in surveys but these results show the same is also true of soft drinks.

Comparing across the LCF and Kantar datasets, those commodities which are relatively well reported in one also tend to be relatively well reported in the other. The most notable differences come in alcohol where reported Kantar spending is relatively lower than reported LCF spending. Within alcohol there are differences too: for example, the LCF does comparatively badly at recording spirits purchases, capturing just 44% of NA expenditure (compared to 61% for beer and 75% for wine). In the Kantar data, it is beer that is least well recorded (29% of NA spending). This could be due to the way the data is collected. In the LCF, each household member has an individual diary to fill out. In the Kantar data, whilst in principle each household member should record items brought home, in practice it may be that main shopping trips are well-reported whilst those carried out by secondary shoppers are less well captured. If main shoppers are mostly female and men buy more beer, this might explain this finding.

¹⁸ This is not simply because of geography – the absence of Northern Ireland from the survey data is not nearly enough to account for the lower spending. In 2009, for example, we find LCF expenditures are 81% of those reported in the NA; adding Northern Ireland back in raises this to just 83%.

Table 4.1: Gross expenditures by commodity in National Accounts, LCF and Kantar data, 2002 and 2009

	£ million / week						Ratios					
	2002			2009			LCF / NA		Kantar / NA		Kantar / LCF	
	NA	LCF	Kantar	NA	LCF	Kantar	2002	2009	2002	2009	2002	2009
TOTAL	1,413.6	1,149.3	749.0	1,833.5	1,484.0	951.2	81%	81%	53%	52%	65%	64%
FOOD	1,057.5	922.4	620.9	1,377.3	1,209.3	787.5	87%	88%	59%	57%	67%	65%
Bread and cereals	176.9	177.9	120.5	230.8	234.0	156.4	101%	101%	68%	68%	68%	67%
Meat	248.3	227.3	155.2	315.7	284.1	182.4	92%	90%	62%	58%	68%	64%
Fish	46.1	44.4	27.8	56.6	58.3	35.7	96%	103%	60%	63%	63%	61%
Milk, cheese and eggs	145.6	125.6	84.1	210.0	180.7	120.2	86%	86%	58%	57%	67%	67%
Oils and fats	24.2	21.1	15.4	28.7	28.5	20.7	87%	99%	64%	72%	73%	73%
Fruit	87.2	74.1	47.8	125.6	99.4	59.5	85%	79%	55%	47%	65%	60%
Vegetables	166.5	142.3	97.6	208.5	182.8	118.9	85%	88%	59%	57%	69%	65%
Sugar and confectionery	134.8	66.3	43.4	160.9	80.4	56.4	49%	50%	32%	35%	66%	70%
Other food	27.9	43.3	29.0	40.6	61.3	37.3	155%	151%	104%	92%	67%	61%
BEVERAGES	141.7	86.3	58.8	183.9	103.1	71.4	61%	56%	41%	39%	68%	69%
Coffee, tea and cocoa	35.1	25.7	18.3	47.3	29.5	21.4	73%	62%	52%	45%	71%	73%
Fruit juices and soft drinks	106.6	60.6	40.4	136.6	73.7	50.0	57%	54%	38%	37%	67%	68%
ALCOHOL	214.4	140.5	69.3	272.3	171.6	92.3	66%	63%	32%	34%	49%	54%
Spirits	55.8	28.1	18.9	74.9	32.7	24.4	50%	44%	34%	33%	67%	75%
Wines, cider and perry	103.2	75.8	34.8	131.1	98.3	49.0	74%	75%	34%	37%	46%	50%
Beer	55.4	36.6	15.7	66.4	40.6	19.0	66%	61%	28%	29%	43%	47%

Source: Author's calculations from ONS, LCF and Kantar Worldpanel data. Notes: NA = National Accounts, LCF = Living Costs and Food Survey. LCF and Kantar data are for Great Britain, whilst NA data are for the UK (including Northern Ireland). LCF and Kantar data are converted to gross national annual totals using supplied household sampling weights; all data are then expressed as weekly average expenditures.

Comparing results for 2002 and 2009 tells us where survey expenditure measures have grown more or less quickly than those in the NA. A noticeable shift occurs for oils and fats, where spending growth was much faster in the both surveys than the NA. This is driven by a sharp fall of around 15% in spending between 2008 and 2009 in the NA, with much smaller declines in the surveys. There is also a relative decline in both surveys for fruit spending. Again, this is driven mostly by a single year: fruit spending grew by more than 16% in 2005 in the NA data, but only by around 7% in the two surveys. Detailed figures for spending ratios and growth rates in each year for each commodity are available on request, but broadly the conclusion from earlier that (a) spending growth is more volatile in the LCF than either the NA or Kantar data sources and that (b) there is no clear 'winner' between the LCF and Kantar as to which tracks growth rates observed in the NA holds across commodities as well.

4.2 Comparing scanner and budget survey expenditure data

Section 3 discussed a number of papers which made comparisons between UK and US scanner data and budget survey data. Here we extend a previous comparison of the Kantar Worldpanel and LCF (Leicester and Oldfield, 2009) to more recent years of data, and explore new issues around how the period of observation and household selection affects the comparisons. This analysis is useful not only as a way to compare, contrast and validate the different expenditure surveys but also to inform us about how useful existing scanner data might be for imputing detailed expenditures into budget survey data (see Section 5). If we find that expenditure *patterns* are very different in the Worldpanel and LCF data, we might be less confident about using the scanner data to try and predict detailed expenditures in the survey data given high-level information on total outlays.

When making cross-dataset comparisons, it is important to bear in mind that they are collected in very different ways. The LCF is recorded over two weeks based on diaries kept by each household member. Respondents are contacted at least once during the two week period to check for any problems filling in the diary, and a thorough check is made of the diaries at the end of the period to ensure they have been properly completed (Ayres et al, 2010). In the Kantar data, each household can participate for as long as they wish. They are contacted every nine months or so to check that demographic information is up to date. Households who do not appear to be participating may be contacted to check if there are any problems, but in general attempts to ensure good compliance are more limited than in the LCF. Thus if we want to compare average spending levels and patterns in the LCF to those in the Kantar data, the crucial issue is what sample of households we select from the Kantar data, over what period of time we choose to observe them and how we deal with seeming periods of non-compliance. Leicester and Oldfield (2009), for example, look at average weekly expenditures in the Kantar data amongst households who report spending in at least four separate weeks (not necessarily consecutive) in a given year. They include only those weeks in which some spending is observed. On this basis, they find that average total food and drink expenditures were about 25% lower in the Kantar data than in the LCF data in 2005, once observable demographic differences in the samples were taken into account.¹⁹

¹⁹ They find a smaller gap of around 16% when comparing the first two full weeks of expenditure for households newly signed up to the Kantar data (again excluding cases where either week includes zero expenditure). This could reflect the 'fatigue' issue mentioned above which sees recorded spending drop off slightly with the length of participation.

Here, we follow the approach of Griffith and O’Connell (2009) and report how spending levels and patterns vary with the period over which we observe Kantar households and how we deal with zero spending weeks. The main conclusion is that, at least at the mean (though not for the distribution of total spending or commodity budget shares), the period of observation does little to change observed scanner data expenditures but that the treatment of zero spending weeks makes a significant difference.

Sample selection

Our main comparisons between LCF and Kantar data cover calendar year 2009.²⁰ In the LCF, we exclude all households in Northern Ireland to ensure the geographical coverage of the two datasets is comparable. This gives a sample size of 5,220 households.

In the Kantar data, a total of 26,655 separate households are observed making at least one purchase in 2009. We look first at households who, according to the dates at which they signed up to and dropped out of the survey, were active during the whole period. This gives a ‘non-dropout’ sample of 21,093 households (79.1% of the full sample). As an additional selection, we also condition on households who have no reporting gap (period during which no food and drink expenditures at all are recorded) exceeding 6 weeks. While this threshold is somewhat arbitrary, the intention is to try and exclude households who do not appear to be fully compliant over the whole year. Shorter periods of non-reporting might be reconciled as holiday periods, for example, but it seems somewhat unlikely that many households would legitimately purchase no food items at all for a 6-week period.²¹ Whilst there are households with longer reporting gaps who then appear to re-enter the data, it could be that the gap was associated with some shock to household composition which might also make us prefer to exclude it from our analysis. This sample, which we deem the ‘regular reporter’ sample, includes 15,781 households (59.2% of the full sample). Appendix A gives the results of a probit regression comparing the demographic characteristics of those in the ‘regular reporter’ sample to the characteristics of the full sample. Households excluded by this sample selection are significantly more likely to be in London, be headed by someone aged under 40 or over 60, be headed by a female, to have larger numbers of adults or children, and to have missing demographic information either on income, employment status or the number of cars. Households who are not required to report non-barcoded items are significantly *more* likely to be part of the ‘regular reporter’ sample, which suggests that the reduced respondent burden may encourage households to report spending consistently. Interestingly, we find no impact of income or employment status on whether households are selected into the regular reporter sample or not: we might have expected richer households or those who are full-time employed to have higher valuations of time which makes them less likely to be compliant given the time costs of participation in the data.

The fact that our preferred ‘regular reporter’ sample is clearly a non-random set of all Kantar households should be borne in mind when comparing raw expenditures. Later in this section, we condition on observable demographics across the datasets to see how far they can explain spending differences.²²

²⁰ As in the comparison to National Accounts aggregates above, we use a 52-week Kantar period which does not quite overlap with the calendar year, running from 29th December 2008 until 27th December 2009.

²¹ We do not consider whether or not households buy non-food items during this period.

²² In the Kantar data, we observe household spending over a full year period. The demographic information is essentially baseline information – we have demographic data once a year for each household that are updated roughly each November. For the 2009 sample, then, the demographics refer to November 2008 values. As

Comparisons of average expenditure and budget shares


Table 4.2 shows average expenditures per week, by CPI commodity, from the LCF and Kantar data during 2009. To strip out the effects of household composition, expenditures are equalised using the before housing costs modified OECD equivalence scale.²³

We look at a number of different Kantar samples. First, we take the ‘no dropout’ sample and pick a random consecutive two week period (matching the LCF diary period) over which to observe expenditures. Here we treat weeks with zero spending as genuine.

We then make the additional selection described above and drop households with long reporting gaps. For this ‘regular reporter’ sample, we look at average weekly spending when we choose observation periods of different length. Table 4.2 shows results over two weeks and the full 52-week period (figures for periods of 4, 12 and 26 weeks essentially show the same results).

Finally, we strip out the effect of weeks in which no purchases of food and drink at all are recorded. For each Kantar household in the regular reporter sample, we take the longest consecutive set of weeks containing any recorded expenditure. Figure 4.3 shows the distribution of observation periods. Around one in eight households report some food or drink spending in every week. The mean duration is 25.1 weeks.

Figure 4.3: Distribution of longest uninterrupted continuous spell of food and drink reporting (weeks) for 2009 Kantar ‘regular reporter’ sample



Source: Author’s calculations from 2009 Kantar Worldpanel. *Notes:* Sample includes 15,781 households who are ‘live’ during the entire 52-week period and have no single reporting gap exceeding 6 weeks.

noted in Section 3, demographic transitions in the Kantar data do not seem to be well observed other than the numbers and ages of household members.

²³ See Appendix A of Jin et al (2011) for details of equivalence scales.

Table 4.2. Weekly average equivalised expenditure levels by CPI commodity group, LCF and Kantar, 2009

	LCFS		Kantar							
	2 weeks		Non-dropouts		Regular reporters					
			2 weeks		2 weeks		52 weeks		Longest continuous	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Bread and cereals	8.51	5.20	5.65	4.63	6.34	4.54	6.36	2.98	7.17	3.24
Meat	10.65	8.73	6.78	6.86	7.64	6.94	7.60	4.64	8.54	5.17
Fish	2.25	3.43	1.33	2.20	1.50	2.28	1.51	1.40	1.70	1.65
Milk, cheese and eggs	6.75	4.48	4.38	3.71	4.96	3.67	4.97	2.55	5.61	2.77
Oils and fats	1.08	1.23	0.78	1.07	0.88	1.11	0.88	0.61	0.98	0.71
Fruit	3.81	3.94	2.19	2.91	2.48	3.04	2.48	2.25	2.83	2.57
Vegetables	6.81	4.87	4.32	3.95	4.86	3.94	4.88	2.70	5.56	3.08
Sugars, confectionery	3.01	3.32	2.06	2.90	2.34	3.03	2.34	1.76	2.61	2.04
Other food	2.23	3.46	1.35	1.65	1.50	1.67	1.51	0.91	1.69	1.11
<i>FOOD</i>	<i>45.10</i>	<i>24.03</i>	<i>28.82</i>	<i>20.82</i>	<i>32.50</i>	<i>19.87</i>	<i>32.53</i>	<i>13.56</i>	<i>36.70</i>	<i>14.32</i>
Coffee, tea and cocoa	1.15	1.68	0.82	1.51	0.94	1.58	0.94	0.85	1.05	0.99
Mineral water, soft drinks	2.58	2.77	1.78	2.35	1.98	2.40	1.98	1.67	2.24	1.92
<i>BEVERAGES</i>	<i>3.73</i>	<i>3.31</i>	<i>2.60</i>	<i>2.95</i>	<i>2.92</i>	<i>2.99</i>	<i>2.92</i>	<i>1.92</i>	<i>3.29</i>	<i>2.19</i>
Spirits	1.36	4.69	0.94	4.05	1.06	4.32	1.09	3.07	1.17	3.48
Wine	3.81	10.44	1.90	5.32	2.17	5.67	2.17	4.17	2.42	4.82
Beer	1.45	3.89	0.77	2.97	0.86	3.19	0.85	2.04	0.96	2.43
<i>ALCOHOL</i>	<i>6.63</i>	<i>13.05</i>	<i>3.61</i>	<i>8.40</i>	<i>4.10</i>	<i>8.88</i>	<i>4.11</i>	<i>6.52</i>	<i>4.56</i>	<i>7.44</i>
TOTAL SPENDING	55.46	31.28	35.03	26.03	39.52	25.03	39.56	17.18	44.56	18.29
# households	5,220		21,093		15,781		15,781		15,781	
% zero weeks	2.7%		21.7%		13.2%		13.2%		0.0%	
Avg non-zero weeks/hh	1.95		1.57		1.74		45.13		25.10	

Source: Author's calculations from 2009 Kantar Worldpanel and LCF 2009. Expenditures are equivalised using the before housing costs modified OECD scale.

Average household weekly equivalised food and drink expenditures in 2009 were £55.46 in the LCF. Expenditures in the Kantar data were lower for all samples. For the 'no dropout' sample observed over two weeks, expenditures were £35.03, almost 37% below the LCF figure. Excluding those with long gaps in reporting, average spending rises to £39.52 per week (29% below the LCF figure) when households are observed over a random two weeks or £39.56 observed over the full 52 weeks. It is striking how little difference there is in average expenditures when households are observed for a full year rather than a single two week period, though the standard deviation of expenditures falls markedly. We return to this issue shortly. Finally, amongst the 'regular reporter' sample observed for the longest period without a zero spending week, average spending rises to £44.56 per week (a gap of just under 20% compared to LCF levels).

What is clear from these figures is that the treatment of weeks in which zero expenditure is reported is hugely important for the level of spending we observe in the Kantar data. Around half of the gap between Kantar and LCF expenditures seen in the 'no dropout' sample is eliminated once we strip out zero expenditure weeks altogether. Indeed, the gap of 20% in this latter sample is consistent with the Leicester and Oldfield (2009) result, who found a gap of 20% in 2005 when comparing raw Kantar and LCF data. They too excluded weeks in which no spending at all was reported.

The greater propensity for zero spending weeks in the Kantar data than the LCF is striking and should be a priority for further analysis. It could reflect households who have effectively attrited but not formally dropped out. However, many households have reporting behaviour which is not consistent with this – for example, they report nothing for a few weeks then start scanning again. Understanding what drives this in scanner data would be useful – is it to do with prompting from Kantar, that only certain types of large, infrequent trips are reported, or something else?

Table 4.3 shows Kantar expenditures relative to LCF expenditures. Table 4.4 shows the budget share of each commodity in each sample. As discussed earlier, even if average expenditures are lower in the scanner data, if the extent of 'under-reporting' is quite consistent such that the patterns of expenditure are similar, this acts as a useful validation (of both data sources) and gives us more confidence in trying to use scanner data as a means to impute detailed budget shares from aggregate expenditure data.

From Table 4.3, several features emerge. The average proportion of LCF expenditures reported in the Kantar data rises for all commodities as we remove the impact of zero spending weeks, but again there is very little impact at all of the period over which we observe Kantar households. Using the 'longest uninterrupted' measure of Kantar spending (rightmost column of Table 4.3), relative to LCF spending, Kantar expenditure levels match up most closely for non-alcoholic beverages and least closely for alcohol. Food spending is somewhere between. There are differences across disaggregate commodities: for example, the Kantar data picks up about 75% as much spending on average for fish, fruit and other foods than the LCF, but about 90% of the expenditure on oils and fats, and coffee and tea. Differences across alcohol types are particularly noticeable, with average Kantar expenditures less than two-thirds the LCF level for wine and beer but 86% as high for spirits.

Table 4.3. Kantar as a proportion of LCF expenditure, 2009

	All	Regular reporters		
	2 weeks	2 weeks	52 weeks	Longest
Bread and cereals	66.4%	74.5%	74.7%	84.3%
Meat	63.7%	71.7%	71.4%	80.2%
Fish	59.1%	66.7%	67.1%	75.6%
Milk, cheese and eggs	64.9%	73.5%	73.6%	83.1%
Oils and fats	72.2%	81.5%	81.5%	90.7%
Fruit	57.5%	65.1%	65.1%	74.3%
Vegetables	63.4%	71.4%	71.7%	81.6%
Sugars, confectionery	68.4%	77.7%	77.7%	86.7%
Other food	60.5%	67.3%	67.7%	75.8%
<i>FOOD</i>	<i>63.9%</i>	<i>72.1%</i>	<i>72.1%</i>	<i>81.4%</i>
Coffee, tea and cocoa	71.3%	81.7%	81.7%	91.3%
Mineral water, soft drinks	69.0%	76.7%	76.7%	86.8%
<i>BEVERAGES</i>	<i>69.7%</i>	<i>78.3%</i>	<i>78.3%</i>	<i>88.2%</i>
Spirits	69.1%	77.9%	80.1%	86.0%
Wine	49.9%	57.0%	57.0%	63.5%
Beer	53.1%	59.3%	58.6%	66.2%
<i>ALCOHOL</i>	<i>54.4%</i>	<i>61.8%</i>	<i>62.0%</i>	<i>68.8%</i>
TOTAL SPENDING	63.2%	71.3%	71.3%	80.3%

Source: Author's calculations from 2009 Kantar Worldpanel and LCF 2009. Expenditures are equivalised using the before housing costs modified OECD scale.

Table 4.4. Food and drink budget shares, by survey, 2009


	LCFS	Kantar			
	2 weeks	All	Regular reporters		
		2 weeks	2 weeks	52 weeks	Longest
Bread and cereals	15.3%	16.1%	16.0%	16.1%	16.1%
Meat	19.2%	19.4%	19.3%	19.2%	19.2%
Fish	4.1%	3.8%	3.8%	3.8%	3.8%
Milk, cheese and eggs	12.2%	12.5%	12.6%	12.6%	12.6%
Oils and fats	1.9%	2.2%	2.2%	2.2%	2.2%
Fruit	6.9%	6.3%	6.3%	6.3%	6.4%
Vegetables	12.3%	12.3%	12.3%	12.3%	12.5%
Sugars, confectionery	5.4%	5.9%	5.9%	5.9%	5.9%
Other food	4.0%	3.9%	3.8%	3.8%	3.8%
<i>FOOD</i>	<i>81.3%</i>	<i>82.3%</i>	<i>82.2%</i>	<i>82.2%</i>	<i>82.4%</i>
Coffee, tea and cocoa	2.1%	2.3%	2.4%	2.4%	2.4%
Mineral water, soft drinks	4.7%	5.1%	5.0%	5.0%	5.0%
<i>BEVERAGES</i>	<i>6.7%</i>	<i>7.4%</i>	<i>7.4%</i>	<i>7.4%</i>	<i>7.4%</i>
Spirits	2.5%	2.7%	2.7%	2.8%	2.6%
Wine	6.9%	5.4%	5.5%	5.5%	5.4%
Beer	2.6%	2.2%	2.2%	2.1%	2.2%
<i>ALCOHOL</i>	<i>12.0%</i>	<i>10.3%</i>	<i>10.4%</i>	<i>10.4%</i>	<i>10.2%</i>

Source: Author's calculations from 2009 Kantar Worldpanel and LCF 2009.

Expressed as shares of the total food and drink budget, Table 4.4 makes it clear that the particular sample selected from the Kantar data makes very little difference to the pattern of expenditure observed. Comparing LCF budget shares to those from the uninterrupted Kantar sample also reveals relatively small differences. For any single commodity, the largest difference in budget share is for wine, which makes up on average 5.4% of food and drink spending in the Kantar data but 6.9% of spending in the LCF. In the opposite direction, bread and cereals make up 16.1% of total Kantar food and drink spending compared to 15.3% of LCF spending.

These results are for 2009, but figures from earlier years are not very different. Figure 4.4 shows the average equivalised weekly expenditure in the Kantar data for the uninterrupted sample as a proportion of the LCF figure between 2002 and 2009, for broad commodity aggregates. There is a small increase over time. In 2002, Kantar households reported 78.5% as much spending as LCF households on average. This rose to 80.3% by 2009. More noticeable is what appears to be a step increase between 2006 and 2007 for beverages, from about 84% to 90%, and a longer term upward trend for alcohol beginning in 2006 (though stalling somewhat in 2009).

Figure 4.4. Average Kantar expenditures as a proportion of LCF, by year and broad commodity aggregate, 2002–2009



Source: Author’s calculations from Kantar Worldpanel and LCF data. Note: Kantar figures relate to “longest uninterrupted” period of Kantar reporting amongst households that are active across the full calendar year and have no single reporting gap in excess of 6 weeks.

One possible reason for the recent improvement in the proportion of beverage and alcohol expenditures reported in the Kantar data relative to LCF might be the introduction of different scanner technology. As the sample size was enlarged between 2005 and 2006, there was a move towards a new scanning device which was more portable, and which no longer required households to scan non-barcoded products. If this made compliance costs lower it may have increased reporting of barcoded items – likely to cover almost all spending on alcohol and non-alcoholic drinks – at the cost of reducing

reporting of some other categories of spending.²⁴ More detailed analysis lends some support to this: between 2005 and 2009, for example, weekly average spending on meat in the Kantar data fell from 82.5% to 80.2% of the LCF average; for dairy products (including cheese) from 85.0% to 83.1%; for fruit from 77.1% to 74.3% and vegetables from 83.7% to 81.6%.

Table 4.5 uses the 2009 data and makes a direct comparison between households who report non-barcoded products and those who do not. We focus on the longest uninterrupted spending period. In the Kantar data, households who do not report non-barcoded items spend significantly more than those who do, £1.65 (3.8%) per week. On average, they report 82.1% of LCF spending levels compared to 79.2% for those who report non-barcoded products. The largest relative differences are in beer (28% higher spending for the sample not recording non-barcoded items), soft drinks (22%), other food (14%) and wine (13%). Spending is lower in only two groups: fruit (22% lower) and vegetables (3%) – those for which non-barcoded items are particularly important. This leads to some quite different expenditure patterns across the groups. Amongst those not reporting loose items, fruit makes up just 5.3% of the food and drink budget, compared to 7.1% for those who do (and 6.9% in the LCF). By contrast the budget share for non-alcoholic beverages is 8% for those not reporting loose items and 7% for those who do. We explore below the extent to which these differences might also be attributed to demographic differences between the groups as well as the technology they use.

The findings so far have illustrated that expenditure levels in the Kantar Worldpanel are, on average, about 20% lower per week than those in the LCF once we strip out weeks in which no spending at all is reported, and almost 40% lower if we do include them. Expenditure patterns are not much affected by the particular Kantar sample that we draw, though being required to report random weight goods or not does have a substantial impact. Though there are some differences in spending patterns across the Kantar and LCF surveys, they are typically quite small.

One way to illustrate the economic significance – or otherwise – of differences in spending patterns is to ask what food inflation rates would have looked like had CPI weights for different food groups been drawn directly from the LCF or Kantar surveys, rather than based (as now) on the National Accounts aggregates. Figure 4.5 shows the 2009 food budget shares based on Kantar data, LCF data and from the CPI expenditure weights.²⁵ Note that we look here only at food, not alcohol and beverages. As noted earlier, the differences between LCF and Kantar data are not particularly large: the Kantar basket more heavily weights bread, dairy, fats and confectionery whilst the LCF basket more heavily weights fruit, fish and other food. There are larger differences between the weights based on survey data and those from the CPI basket. Weights for bread, meat, dairy, fats and other food are lower in the CPI than either of the survey baskets, whilst weights for confectionery and fruit are higher. The much lower spending on other food in the CPI accords with the much higher expenditure on other foods observed in the surveys than the National Accounts aggregates in Section 4.1 above.

²⁴ This hypothesis does not really explain what appears to be quite a sustained improvement in alcohol reporting, at least between 2005 and 2008, however.


²⁵ Note that CPI weights are based on National Accounts expenditure aggregates which are in turn based on slightly out-dated expenditure data from the LCF. For example, 2010 weights in the CPI are heavily influenced by LCF data from 2008 and 2009. Our estimates based on LCF and Kantar data use contemporaneous data, e.g. the 2009 weights are based on 2009 data.

Table 4.5. Weekly expenditure comparisons, LCF and Kantar ‘uninterrupted’ sample 2009, by reporting of non-barcoded items

	Average weekly spending				Budget shares				Kantar / LCFS		
	LCFS	Kantar (longest continuous)			LCFS	Kantar (longest continuous)			All	Yes RW	No RW
		All	Yes RW	No RW		All	Yes RW	No RW			
Bread and cereals	8.51	7.17	6.99	7.44	15.3%	16.1%	15.9%	16.3%	84.3%	82.1%	87.4%
Meat	10.65	8.54	8.42	8.73	19.2%	19.2%	19.2%	19.2%	80.2%	79.1%	82.0%
Fish	2.25	1.70	1.69	1.71	4.1%	3.8%	3.8%	3.8%	75.6%	75.1%	76.0%
Milk, cheese and eggs	6.75	5.61	5.55	5.71	12.2%	12.6%	12.6%	12.5%	83.1%	82.2%	84.6%
Oils and fats	1.08	0.98	0.98	0.99	1.9%	2.2%	2.2%	2.2%	90.7%	90.7%	91.7%
Fruit	3.81	2.83	3.10	2.41	6.9%	6.4%	7.1%	5.3%	74.3%	81.4%	63.3%
Vegetables	6.81	5.56	5.63	5.45	12.3%	12.5%	12.8%	12.0%	81.6%	82.7%	80.0%
Sugars, confectionery	3.01	2.61	2.53	2.73	5.4%	5.9%	5.8%	6.0%	86.7%	84.1%	90.7%
Other food	2.23	1.69	1.60	1.83	4.0%	3.8%	3.6%	4.0%	75.8%	71.7%	82.1%
<i>FOOD</i>	<i>45.10</i>	<i>36.70</i>	<i>36.50</i>	<i>37.01</i>	<i>81.3%</i>	<i>82.4%</i>	<i>83.1%</i>	<i>81.3%</i>	<i>81.4%</i>	<i>80.9%</i>	<i>82.1%</i>
Coffee, tea and cocoa	1.15	1.05	1.01	1.11	2.1%	2.4%	2.3%	2.4%	91.3%	87.8%	96.5%
Mineral water, soft drinks	2.58	2.24	2.06	2.52	4.7%	5.0%	4.7%	5.5%	86.8%	79.8%	97.7%
<i>BEVERAGES</i>	<i>3.73</i>	<i>3.29</i>	<i>3.07</i>	<i>3.63</i>	<i>6.7%</i>	<i>7.4%</i>	<i>7.0%</i>	<i>8.0%</i>	<i>88.2%</i>	<i>82.3%</i>	<i>97.3%</i>
Spirits	1.36	1.17	1.17	1.18	2.5%	2.6%	2.7%	2.6%	86.0%	86.0%	86.8%
Wine	3.81	2.42	2.30	2.61	6.9%	5.4%	5.2%	5.7%	63.5%	60.4%	68.5%
Beer	1.45	0.96	0.87	1.11	2.6%	2.2%	2.0%	2.4%	66.2%	60.0%	76.6%
<i>ALCOHOL</i>	<i>6.63</i>	<i>4.56</i>	<i>4.33</i>	<i>4.90</i>	<i>12.0%</i>	<i>10.2%</i>	<i>9.9%</i>	<i>10.8%</i>	<i>68.8%</i>	<i>65.3%</i>	<i>73.9%</i>
TOTAL SPENDING	55.46	44.56	43.90	45.55					80.3%	79.2%	82.1%
# households	5,220	15,781	9,508	6,273							
Avg non-zero weeks/hh	1.95	25.10	25.06	25.14							

Source: Author’s calculations from 2009 Kantar Worldpanel and LCF 2009. Expenditures are equalised using the before housing costs modified OECD scale. “Yes RW” are households who report random-weight (non-barcoded) products; “No RW” are households who do not.

Figure 4.5. Food commodity weights, LCF, Kantar and CPI, 2009



Source: Author's calculations based on 2009 Kantar Worldpanel, LCF and ONS data.

Figure 4.6 shows the different inflation rates for food that result from applying survey- and year-specific commodity weights. The left-hand panel gives the results as index numbers between January 2002 and December 2009; the right-hand panel as annual inflation rates starting in January 2003.²⁶ Overall, the effect of re-weighting the food CPI using LCF and Kantar-specific commodity weights is small. Over the whole period, the food CPI rose by 29.6% whereas an index based on LCF weights rose by 28.1% and one based on Kantar weights rose by 28.6%. The average annual food inflation rate between 2003 and 2009 was 3.8% based on CPI weights, and 3.7% based on weights from the survey data. The largest absolute gap for any single month between the CPI-weighted inflation rate and the LCF-weighted rate was 0.5%. The largest gap comparing CPI-weighted and Kantar-weighted inflation rates was 0.7%. The largest gap comparing LCF-weighted and Kantar-weighted inflation rates was just 0.3%.


Distributions of spending and budget shares

The comparisons so far have focused on average spending levels and budget shares. However, looking at their distribution across households is informative. First, it helps to understand what might be driving differences in the averages. Second, even if (as we saw above) changing the period over which we observe Kantar households makes little difference to the average budget share or spending level, it may still affect the distribution. For issues like poverty and inequality it is the distribution which matters.

Figure 4.7 shows a density plot of the distribution of average equivalised weekly food and drink spending in 2009 for the LCF, and 'regular reporter' Kantar households observed over 2 weeks, 52 weeks and for the longest uninterrupted number of weeks. The distribution for LCF households is relatively smooth over the range of spending shown. When we observe Kantar households for just two weeks, there is a bulge in the distribution at zero, reflecting the high prevalence of zero expenditure weeks. Observed over longer periods, the distribution of expenditures in the Kantar data become more smoothly distributed though clearly skewed somewhat more towards lower expenditures than the LCF figures. There are notably far fewer high-spending households in the Kantar data than the LCF.


²⁶ Figures were calculated by calculating within-year price indices for each food sub-group based at 100 in January and using the different weights to calculate a within-year food index. These indices are then 'chained' to give a series over the whole time period. See Section 2.5 of Office for National Statistics (2010b) for more on chaining.

Figure 4.6: CPI food price indices and inflation rates, based on CPI, LCF and Kantar expenditure weights, 2002–2009



Source: Author’s calculations based on 2002–2009 Kantar Worldpanel, LCF and ONS data.

Figure 4.7. Distribution of equivalised weekly food & drink expenditures, 2009



Source: Author’s calculations based on 2009 Kantar Worldpanel and LCF data.

Even more interesting is the impact of observation period on commodity-level budget shares. One problem with observing expenditure over a short horizon like two weeks is that households may purchase and consume some goods relatively infrequently. To take a stylised example: imagine that all households shop once a week and consume one can of beer per week. Beer is only sold in 4-packs which sell for £5. Households therefore spend £5 on beer once every four weeks, and average weekly beer consumption is £1.25. If we took a random two week period, we would observe half of households buying beer

(consuming £2.50 per week) and half of households buying no beer (consuming £0 per week). The average value of consumption across all households would be right, but the distribution would be wrong. Given the wide availability of freezer and refrigerator space and the ability to store some food and drink items (like canned goods) for a long time, there is also scope for households to engage in stockpiling: buying goods when they are cheap (perhaps on a temporary special offer) for consumption over a long period. The longer the horizon over which we can observe spending patterns, therefore, the more accurate a record of true consumption that data is likely to represent.

Table 4.6 reports some evidence on this. It compares, for different periods of observation and datasets, the fraction of households reporting zero spending on each commodity group. Kantar results are based on the regular reporter sample defined earlier (15,781 households). It is notable that, aside from alcohol, once observed over a full year almost all households purchase from each food commodity type. Observed over just two weeks, a large fraction of zero expenditures are reported. For example, over a fortnight, more than 42% of the Kantar sample reports no expenditure on fish, but over 52 weeks only 1.7% report zero. Also striking is that the probability of zero purchases over two weeks is markedly higher in the Kantar data than the LCF, though as noted above this reflects the much higher prevalence of weeks in which no spending at all is reported. For researchers interested in estimating price responsiveness or demand models, the ability to observe spending over an extended period is a key advantage of scanner data, since it drastically reduces the problem of how to deal with zero expenditure values.

Table 4.6. Proportion of households reporting zero commodity expenditures, by dataset and period of observation

	LCFS	Kantar, regular reporter sample				
	2 weeks	2 weeks	4 weeks	12 weeks	26 weeks	52 weeks
Bread and cereals	1.3%	6.1%	1.5%	0.0%	0.0%	0.0%
Meat	5.7%	12.2%	5.1%	1.9%	1.2%	0.9%
Fish	33.5%	42.2%	23.9%	7.3%	3.2%	1.7%
Milk, cheese and eggs	2.3%	7.4%	1.9%	0.1%	0.0%	0.0%
Oils and fats	28.7%	35.9%	16.4%	2.6%	0.7%	0.3%
Fruit	12.8%	20.8%	9.7%	2.3%	0.9%	0.4%
Vegetables	3.4%	8.0%	2.0%	0.1%	0.0%	0.0%
Sugars, confectionery	13.9%	21.3%	8.1%	1.1%	0.3%	0.2%
Other food	17.8%	23.8%	9.0%	1.1%	0.3%	0.1%
Coffee, tea and cocoa	43.7%	52.0%	31.0%	9.4%	3.4%	1.2%
Mineral water, soft drinks	17.9%	25.6%	12.6%	3.2%	1.2%	0.5%
Spirits	85.8%	89.2%	84.1%	72.4%	61.5%	39.1%
Wine	59.9%	70.5%	59.6%	41.4%	29.2%	16.2%
Beer	76.2%	84.4%	76.6%	60.1%	46.4%	32.0%


Source: Author's calculations based on 2009 Kantar Worldpanel and LCF data.

For alcohol, about 71% of households in the Kantar data report no spending on wine during a two week period, 84% no spending on beer and 89% no spending on spirits. Over a full year, these figures are 16%, 32% and 39% respectively. Alcohol products are the only commodities in our set for which more than 2% of households report zero spending over a year.

Aside from the impact on the likelihood of observing zero expenditure, increasing the duration over which spending is recorded substantially reduces the variance in the distribution of household-level budget shares of each commodity. Figure 4.8 illustrates this for four different commodities (clockwise from top left, for bread, meat, vegetables and fish; figures for other goods available on request). This is crucial for

some applications. For example, past research in the UK has used data from the two-week LCF to estimate household-specific expenditure patterns from which household-level inflation rates have been estimated (Leicester et al, 2008; Levell and Oldfield, 2011). If at least some of the variation in household budget shares is driven by the short period of observation, then these studies would overstate the variation in household-specific inflation rates across different types of household groups. Estimation of demand models based on cross-sectional expenditure data also often use household-specific budget weights to estimate price and income elasticities of demand. To re-iterate, even if the observation period has little material effect on average spending patterns (and thus would not lead us to change, for example, CPI basket weights), its impact on the distribution of expenditures is nevertheless important.

Figure 4.8. Distribution of budget shares for particular commodity groups based on period of observation, 2009



Source: Author's calculations based on 2009 LCF and Kantar Worldpanel data. Notes: Commodities shown, clockwise from top left, are bread and cereals, meat, vegetables and fish.

The impact of demographics

Are differences in expenditures across scanner and budget survey data driven by demographic effects? Recall that we select a particular non-random sample of households from the scanner data – those who do not drop out of the sample in a given year and have no long gap in their reported expenditures. If the households in this sample have characteristics which would typically make them low spenders on food at home (for example, being poorer) then this could account for the spending gap between the datasets.²⁷

²⁷ The comparisons earlier in this section were based on equivalised spending, which should at least partly address differences in household composition across the two surveys, but there may clearly be other observable differences.

Appendix B compares the observable characteristics of the Kantar and LCF 2009 samples along a number of dimensions.²⁸ At this point it is worth re-stating that there are relatively few common observable demographic characteristics in the LCF and Kantar surveys, reflecting the very different purposes for which they are collected. To the extent that we can strip out the effects of these common demographics, there may well be a large set of unobserved demographics, or variables observed in one dataset but not the other, which affect expenditures but for which we cannot control in this analysis.

To summarise the key differences:

- Kantar households tend to have lower income. Amongst those with non-missing incomes, only 12.5% (3.7%) of Kantar households have gross annual household income in excess of £50,000 (£70,000) compared to 20.8% (9.9%) of LCF households.
- Kantar households are much more likely to own a home computer: 89.5% of the Kantar sample do so compared to 75.6% of the LCF sample. This probably reflects the fact that many Kantar households also participate in online surveys run by Kantar, and many use PCs to upload their expenditure records to Kantar from the scanner units.
- Kantar households are more likely to own a car: only 12.9% do not compared to 21.6% in the LCF.
- The regional composition of the two datasets is similar. In the Kantar sample the South East and East of England are more heavily represented. Scotland appears to be slightly less represented.
- Kantar households are noticeably more middle aged: only 2.7% of the Kantar sample is headed by someone aged over 80 compared to 7.1% of the LCF sample; similarly 0.6% is headed by someone under 25 compared to 3.0% of the LCF sample.
- Kantar households are much more likely to be female-headed: 56% of heads of household are female compared to 25% of the LCF sample. This is almost certainly correlated with the fact that Kantar households are much more likely to be headed by someone who is not working (11.3% compared to 8.6% in the LCF) or part-time employed (15.8% working less than 30 hours, compared to just 6.1% in the LCF) and less likely to be headed by someone working full-time (41.8% compared to 51.5%).
- Strikingly, though, fewer Kantar households contain no adult males at all and there are more multiple adult households. This accords with information received from discussions with Kantar that they over-sample multiple adult households because of difficulties in obtaining purchase information from secondary shoppers.
- Kantar households have about the same number of pre-school aged children as LCF households, but usually slightly more older children. This reflects the age profile of the two samples.

Ex ante it is not clear what these demographic differences imply for expenditures. Kantar households are poorer, on average, and less likely to be headed by a full-time employee which might mean we would expect them to spend less than LCF households. But along other dimensions they are better off: for example, being more likely to have cars and computers, and more likely to be middle aged (where lifecycle expenditures peak). Thus an empirical study is needed.

We first pool observations from the LCF and Kantar samples in 2009. We take average total food and drink weekly spending figures over two weeks from the LCF and from the longest uninterrupted spending period from the Kantar sample.²⁹ We regress the log of expenditure on a dummy variable which takes the value 1 for households from the Kantar sample, to give the raw average proportional difference between the surveys. We then add a vector of common demographic controls from the surveys to strip out

²⁸ In these results and for the rest of this section, we exclude a small number of households from the Kantar sample who report missing information on the number of cars they own or the employment status of the household head, and households in either survey who report equalised average weekly spending on food and drink (over 2 weeks in the LCF sample, and over the longest uninterrupted reporting period in the Kantar sample) of less than £5.

²⁹ As we now control for demographics, including household composition, here we use unequalised expenditure figures.

observable demographic effects and see what happens to the coefficient on the Kantar dummy. Table 4.7 reports the results (full regression results available on request), including separate results for households who do and do not report random weight items in the Kantar survey.

Across all households, the raw gap between the LCF and Kantar surveys is just over 9% - this is markedly less than the 20% seen earlier in this section, but our analysis here is based on a different (unequivalised) measure of spending and drops households with very low food spending of less than £5/week on average. Those who are asked to report random weight purchases spend 12% less than LCF households whilst those who are not asked to do so have a raw spending difference of less than 5%. The raw figures, as above, therefore suggest that reducing respondent burden increases expenditures, but all these unconditional differences are statistically significant.

Table 4.7. Coefficients on Kantar dummy from regression of log average weekly food and drink expenditures, pooled Kantar and LCF sample, 2009

	All households	Records random weight	No random weight
No demographic controls	-0.093*** (0.011)	-0.123*** (0.011)	-0.047*** (0.012)
R ²	0.005	0.010	0.002
Controlling for common observed demographics	-0.182*** (0.009)	-0.179*** (0.010)	-0.179*** (0.013)
R ²	0.406	0.412	0.414
N	20,875	14,643	11,382

Source: Author’s calculations based on 2009 LCF and Kantar Worldpanel data. *Notes:* Expenditure figures are unequivalised. Households with missing information on number of cars or employment status are dropped, as are those spending less than £5 (equivalised) per week on average in either survey. Demographic controls are household gross annual income group, number of cars, region, age group of household head, sex of household head, employment status of household head, number of adult males, number of adult females, numbers of children in age groups 0-4, 5-10 and 11-17 and numbers of people aged 65 or over. *** = p<0.01; ** = p<0.05; * = p<0.1. Standard errors are robust.

Adjusting for demographic differences between the surveys, however, gives very different results. The coefficient on the Kantar dummy across all households almost doubles, suggesting a conditional expenditure gap of more than 18%. This is consistent with Leicester and Oldfield (2009) who found the gap between Kantar and LCF spending in 2005 rose from 20% to 25% once observable demographics were taken into account. They also concluded that controlling for demographics made little difference to overall spending patterns, suggesting that the effects are similar for each commodity group. From our results, it is notable that once we adjust for demographics, there is no difference at all in the Kantar dummy amongst the groups asked or not asked to report random weight items.

To investigate the impact of demographics further, we repeat the above analysis, but now interact common observable demographics across the surveys with a Kantar dummy. The coefficients on these interactions tell us which household groups report relatively higher spending in the Kantar data than the LCF.³⁰ Table 4.8 shows the main significant interaction terms (full results available on request). The base

³⁰ This follows the approach of Zhen et al (2009), equation 1, and Leicester and Oldfield (2009), Table 9. Zhen et al note that in their results, “[the] coefficient on the Homescan indicator variable (H) provides a measure of the average difference in reported expenditures between Homescan and CES for the reference group. Interestingly, this coefficient is not statistically significant for any of the five [commodities]. These results suggest that much of the differences in mean expenditures [between Homescan and CES] are correlated with the observed household characteristics.” (p.479). However this interpretation is not quite right – as they say, the insignificant Homescan dummy tells us that there is no difference between CES and Homescan *for the reference group* (in their study, households under 25 with income under \$5,000 of ‘other race’ living in the North East in 2002 and so on). To compare Homescan and CES for other groups requires a test of the joint significance of the Homescan coefficient and the interaction between Homescan and the other group dummy variable.

group is households in the South East of England, with incomes between £10,000 and £20,000 per year, with one car and a home computer, where the head is a male aged 45-49 and full-time employed, where there is one adult male and female but no children or anyone aged over 65. For this group, average spending is 9% lower per week in the Kantar data than the LCF. The coefficients in Table 4.8 show the additional difference in the Kantar/LCF gap for other demographic groups; significantly positive figures suggest that households who differ from the base group only in terms of that characteristic report relatively higher expenditures in the Kantar data. Demographic variables which had no significant effects on the interaction terms are not shown.³¹

Table 4.8. Interaction terms between Kantar dummy and demographic groups

Variable	Coefficient	Variable	Coefficient
Income £0-£10k	+0.101***	Head <25	-0.019
Income £20-£30k	-0.032	Head 25-29	+0.058
Income £30-£40k	-0.049	Head 30-34	+0.003
Income £40-£50k	-0.087**	Head 35-39	+0.039
Income £50-£60k	-0.041	Head 40-44	-0.006
Income £60-£70k	-0.106**	Head 50-54	-0.036
Income £70k+	-0.220***	Head 55-59	-0.019
0 adult males	-0.095*	Head 60-64	-0.087**
2 adult males	+0.019	Head 65-69	-0.093
3+ adult males	-0.205***	Head 70-74	-0.070
0 adult females	+0.100***	Head 75-79	-0.146*
2 adult females	-0.052	Head 80+	-0.121
3+ adult females	+0.022	Female Head	+0.106**
1 child aged 0-4	-0.127***	1 person aged 65+	+0.080
2 children aged 0-4	-0.033	2 people aged 65+	+0.137**
3+ children aged 0-4	-0.330***	3+ people aged 65+	-0.065
N	20,875	R ²	0.413

Source: Author's calculations based on 2009 LCF and Kantar Worldpanel data. Notes: Left-hand side variable is the log of total weekly household average food and drink expenditure (unequalised). Other variables controlled for are region, numbers of children aged 5-10 and 11-17, head of household employment status, numbers of cars and presence of a PC in the household. *** = p<0.01; ** = p<0.05; * = p<0.1. Households with missing information on number of cars or employment status are dropped, as are those spending less than £5 (equalised) per week on average in either survey.

A striking finding is that lower income households report relatively higher spending in the Kantar data whilst higher income households report relatively lower spending. If we take LCF spending levels as 'true' (though in general we should be wary of doing so) this result is consistent with poorer households fully reporting their spending in the scanner data, and richer households being more prone to under-reporting. This might reflect the higher opportunity costs of time faced by high income participants. It could also simply reflect poorer households buying less overall and thus requiring less time and effort to scan their purchases. Notably, there is no effect of employment status on relative Kantar expenditures, which we might also expect to be related to the opportunity costs of time.³²

³¹ Region is also not shown. There is one significant interaction coefficient of -0.087 in the East Midlands.

³² Recall too that these estimates are based on the subset of Kantar households who report for the whole year without large spending gaps. If time costs are important then richer or full-time employed households might be less likely to be part of this sample – though as we saw earlier (and detailed in Appendix A) there is no significant effect of income or employment status on being part of this selected sample.

Household composition also matters. Households with young children report relatively lower expenditures than those without – again, this seems plausibly related to time constraints. The coefficients on the interaction terms for older children are negative but insignificant. The number of adults also affects relative Kantar expenditures, though in different ways according to gender. For example, having three or more male adults in the household is associated with much lower relative reported spending but there is no similar effect for three or more female adults. In general, we would tend to expect that households with multiple shoppers would report relatively low spending in the scanner data compared to the LCF to the extent that spending is better reported by the main shopper than other adults; however the evidence for this based on these figures is quite weak.³³

Age effects are interesting: there is some (albeit not particularly significant) evidence that relative spending is higher for households headed by younger people and lower for those headed by older people. This might be interpreted as a modal effect – for example, older households may find using the scanner technology more difficult than younger households. However, it is worth noting two things that go against this conclusion. First, these results are based on the sample of Kantar households who report spending consistently – presumably households who cannot use the technology will not be included. Second, the regression also includes a variable for the number of people aged 65 or over in the household (which is the UK retirement age for males). Discussion with Kantar suggests that they believe older households to be relatively more diligent recorders of their spending, perhaps because they have more time. We find that households with two people aged 65 or over report significantly higher relative spending, consistent with this. We also found (see Appendix A) that older households were more likely to be included in the regular reporter sample than younger households, again consistent with them reporting more assiduously and being less likely to drop out of the scanner sample.

A brief comparison to the results of a similar exercise in Zhen et al (2009) is worthwhile. They compare AC Nielsen and CE expenditures, and regress commodity-level expenditures on demographic variables interacted with a Nielsen dummy. They find a number of demographic effects on relative expenditures which are different to the findings here: for instance, they find no significant effect of the number of children, a negative impact of households with female heads, significant regional effects and generally positive effects for older ages. Other findings are more comparable: for example, lower relative spending for high income households and no clear employment status effects. What this suggests is a point raised in Section 3: findings from one particular comparison of scanner and other datasets do not necessarily translate across countries or surveys. If there were clear ‘modal’ effects of scanner technology we might expect the results to be quite similar across countries; instead, it seems that the particular features of each dataset might be most crucial in driving findings in different countries. Of course, the comparison we make here is not identical to that made in Zhen et al (2009) in terms of the covariates for which we can control or the selection of households, for example. One area for future work might be to explore cross country comparisons of scanner and budget survey data using, as far as possible, identical methods. Regularities which emerge from this exercise might be more credibly assigned to modal effects.

The relationship between total food expenditures and expenditure patterns

The results so far in this section have indicated that, on average, food spending levels in scanner data are lower than those in budget survey data, but that spending patterns are very similar, certainly once we strip out alcohol purchases. There may be two explanations for this. First, the Kantar and LCF surveys could be sampling similar types of households but those in the scanner data are under-recording their spending on each broad commodity group at roughly the same rate on average. Alternatively, the Kantar

³³ Leicester and Oldfield (2009) perform a similar exercise using 2005 data from Kantar and the LCF; they find no evidence of household composition effects on relative expenditures. Unlike our estimates, they find that unemployed households report relatively more spending, but this may proxy for income which was not observed in their estimates.

data could be sampling lower-spending households on average, but food spending patterns vary relatively little with total food expenditures. To explore this latter possibility, we look at how the share of the total food budget devoted to different commodities varies with total food spending. This relationship, the ‘Engel Curve’, is of particular interest in economic research, since it allows us to estimate how changes in spending power will translate into consumption of different. For the purposes of how scanner data may be useful for statistical agencies, as discussed earlier (and detailed in Section 5), one option may be to use detailed expenditures from scanner data to impute spending patterns into budget survey data if all we knew were total spending – which is precisely this Engel relationship. Therefore a key issue is whether or not the Kantar and LCF surveys give a similar impression as to how food expenditures break down as the total food budget increases. If so, we might be more confident in make this kind of imputation. If not, it may shed more light on where particular problems in measuring expenditures arise.

For this analysis we drop alcohol expenditures and focus on food and non-alcoholic drinks. We take data from 2009 and compare Engel curves for the 11 food and drink commodities in the LCF and Kantar datasets. We use the ‘regular reporter’ Kantar sample and restrict attention to those households who are asked to record random weight items. So as not to impose any particular form on the shape of the Engel curve, we plot a non-parametric relationship (using local-linear kernel regression) between the log of total equivalised food expenditures and the share of total spending devoted to the particular commodity.³⁴ We exclude the highest and lowest 1% of households by total spending from each dataset; despite this trimming, the confidence bands around our curves are still very wide at the extremes of the expenditure distribution, though the focus is on the broad slope of the curves.

Results for each food group are shown in Appendix C and summarised here. In general, the Engel curves are close to linear. The main exception is vegetables (at least in the LCF), where the budget share appears to first decline with total food spending then rises again. The results show that the relationship between food expenditures and commodity-level budget shares are very similar across datasets. The slopes of the Engel curves – where they have any slope at all – are the same. Necessities within the food budget (goods with downward sloping Engel curves) are bread and dairy products. Luxuries (with upward sloping curves) are meat, fish and fruit. For fats, vegetables, confectionery, other food and beverages, the Engel curves are broadly flat. Perhaps the biggest divergence in the Engel relationship between datasets is for other food, which shows some evidence of being a luxury in the LCF but where the budget share is essentially flat in total spending in the Kantar data. However there are certainly no clear cases where one dataset suggests a commodity to be a luxury and the other dataset a necessity. Again, of course, we are looking at relatively broad food group aggregates; more disaggregate comparisons could yield different findings. In general, though, these findings are quite reassuring that the scanner and budget survey datasets tell similar stories about how spending patterns change with total food outlays.

Summary

The main results from the comparison of budget survey and scanner data are that the within-period patterns of spending and expenditure trends over time are very similar, though spending levels are lower in scanner data. This is not simply attributable to observable demographic differences. There appears to be evidence based on how this ‘under-reporting’ varies across households and commodities, and from previous studies, to suggest that it is at least partly driven by modal effects and the particular technology used in the scanner data (such as differences between those who do and do not report random weight purchases). Comparing our findings with those based on US data highlights some differences which suggest that differences across scanner datasets are at least as important as differences between scanner data and diary-based budget survey data.

³⁴ We do not condition on other observable covariates.

5. Using detailed scanner data to predict budget shares from aggregate spending

Our analysis so far has focused on how scanner data compare to spending information from aggregate data and other expenditure surveys. These sorts of comparisons are useful as a source of validation for budget survey data and to inform statistical agencies of possible issues in using scanner methods as part of the data collection process.

As discussed in Section 1, scanner data could also be used to impute detailed commodity expenditures given only knowledge of total spending. As a way of reducing respondent burden in a redesigned Consumer Expenditure Survey, one approach would be to ask people only about their total expenditures or their total category-level expenditures (e.g. typical spending on food at home per week or per month), and then use detailed expenditure records to estimate how this breaks down. These records could come from a subset of households within the data who agree to provide a more detailed account of their expenditures (whether using scanners, till receipts, diaries or whatever), or from external datasets like commercial in-home scanner data. In this Section, we provide some evidence on how successful such an approach might be. We use scanner data to predict the budget share of each commodity as a function of total expenditures and observable demographics, and see how well these predicted budget shares compare to the observed data. This gives a sense of how well observable demographics predict expenditure patterns, and thus how successful such an imputation approach might be. One advantage of using scanner data for this exercise is that we have detailed information on *where* people shop, not just what they buy. This means we can look at store-specific expenditures, which is useful as there is evidence (see Figure 5.1) that spending patterns differ by store.

We take an agnostic view on how the information on total expenditures is obtained. One approach would be a series of questions which first ask about how much households spend in total over a week, month or year, say, then how this breaks down into spending on food, clothing, leisure and so on, and then how food expenditures break down across different store types – supermarkets, corner shops, specialist stores and so on. Another would be to use individual bank records and credit card statements to get total store-specific expenditures, an approach which becomes more feasible as cash spending, which cannot be attributed to a particular store, declines in importance.³⁵ Our intention here is to assess how far we might be able to impute detailed spending records given total store-level spending and observable household characteristics; the question of how these totals are obtained is left to future evaluation.³⁶

We base our analysis on the ‘regular reporter’ sample in the 2009 Kantar data (see Section 4.2), restricting attention to households who record random weight purchases. We exclude fourteen households from this sample who have missing demographic information on cars or employment status to give a sample size of 9,494 households. We focus attention on food and non-alcoholic beverages, and use the CPI commodity categories which have formed the basis of our analysis so far to define the patterns of expenditure we wish to impute. This gives 11 categories in total.

We look at expenditures in eight different types of store. UK food retailing is dominated by the ‘big four’ supermarkets – Tesco, Asda, Sainsbury and Morrisons (together accounting for almost 80% of

³⁵ Figures from the UK Payments Council (2010, 2011), for example, suggest that about 68% of retail purchases by value were on credit and debit cards in the second quarter of 2011, compared to about 61% in 2008.

³⁶ One particular issue concerns non-food spending: if we know from bank records or survey questions that a household spends £2,000 per year in Tesco, we would not know how this breaks down between food and non-food items sold there like toys, books, clothes, electricals and so on, which are not included in the Kantar data. Thus we would need some way to try and distinguish the food and non-food budgets if the intention were to use existing commercial scanner data for this kind of exercise.


expenditure in our data) – and we look at each of these stores separately.³⁷ Other supermarket chains are grouped according to quality segment: discount supermarkets (Aldi, Lidl, Netto and Cash and Carry stores), quality supermarkets (Waitrose and Marks & Spencer) and other supermarkets (Co-op, Somerfield, Iceland and any other chains). Remaining stores are then grouped together into a single category for local and high street expenditures – largely shopping in specialist food retailers (butcher shops, delicatessens and so on) and corner shops. More detailed analysis could try to separate these out, but given their relatively small combined share of total food expenditure doing so would probably have very little impact on our overall findings. Table 5.1 shows the eight store types considered and their market shares by expenditure, and Figure 5.1 shows how total food and drink expenditure breaks down for each store across the 11 commodities we consider.³⁸

Table 5.1. Store types in the analysis and market shares, 2009

Store	Market share (% by expenditure)
Tesco	31.5
Asda	18.8
Sainsbury	14.8
Morrisons	13.6
Other supermarkets	7.2
Discount stores	6.0
Quality stores	4.5
Local / High Street stores	3.7

Source: Author’s calculations based on 2009 Kantar Worldpanel data.

Figure 5.1. CPI commodity budget shares, by store type, 2009



Source: Author’s calculations based on 2009 Kantar Worldpanel data.

There are some differences in expenditure patterns across stores. Notably, in discount supermarkets, soft drinks and confectionery account for a larger share of spending than across all stores, whilst dairy and meat products account for smaller shares. In quality supermarkets, fruit, vegetables, meat and fish

³⁷ Spending in the big four supermarkets includes spending in all formats of stores (e.g. Tesco Extra, Tesco Express). Online expenditures are allocated to the appropriate store.

³⁸ These estimates are based on the full 2009 Kantar sample, rather than the selected ‘good reporter’ sample.

account for a larger share and dairy, confectionery and soft drinks account for lower shares. The most striking differences are in local stores – for example, across all stores confectionery accounts for about 6.5% of expenditure but in local stores the share is 16.7%. For meat the shares are 21.6% and 16.3% respectively, and for vegetables 13.7% and 8.7%. These differences probably reflect the different sorts of shopping done in local stores – more top-up items and impulse purchases, for example.

If households systematically under-report different types of expenditure in scanner data, imputed budget shares would be biased. We discussed evidence for this in Section 4. For example (see Table 4.5) amongst households in the Kantar data reporting random weight items, fish and other food appear to be less well reported relative to LCF levels, and fats and confectionery better reported (of course LCF figures themselves will also be subject to reporting problems). Store-level scanner data, or detailed chain-level sale information, would presumably be a complete record of purchases and so act as a useful check on the in-home scanner data. We are not aware of any direct comparisons having been made of this nature, and it would be a useful line of inquiry for statistical agencies thinking of using in-home scanner data. Section 4 also noted that different types of households seem to report their spending in scanner data relatively more successfully (compared to LCF figures) than others – notably that richer households and those with young children reported larger total expenditure gaps. If this relative under-reporting is driven by particular commodity groups, we may be less confident that an imputation approach will give unbiased estimates of the true expenditure patterns of such groups.

Using the Kantar sample, we calculate total annual household-level spending during 2009 across each of the store types defined above, broken down by the 11 CPI food and drink commodities. This gives us household-store specific budget shares. Weighting these by the share of each household's spending in each store type gives the household's overall observed budget share for each commodity. We then use seemingly-unrelated OLS regression methods to estimate a system of equations for store-specific budget shares for each commodity:

$$w_{ij}^k = \alpha + \beta \ln(X_{ij}) + \mathbf{Z}_i' \boldsymbol{\gamma} + \varepsilon_{ij}^k$$

where w_{ij}^k is the budget share of household i in store j for commodity k , X_{ij} is the total spend of household i in store j , and \mathbf{Z}_i is a vector of observable demographic household characteristics.³⁹ The system is run separately for each store type, allowing for store-specific Engel curves and for demographic effects to vary across store types. From the equation we predict the expected budget share \hat{w}_{ij}^k . To ensure the predicted budget shares add up to 1, we estimate the results for 10 of the 11 commodities, excluding 'other food', for which the predicted budget share is estimated as a residual. These predicted shares are then weighted into an overall predicted household-specific budget share for each commodity using the household-store specific expenditure weights.

Detailed regression results for each store type are available on request. As a summary of the predictive power of these regressions, Table 5.2 reports the R^2 estimates by store and commodity (the number of observations varies because the equations are estimated only for households with positive expenditures in each store during the year). The key conclusion is that observable characteristics, including total store expenditures, have little predictive power for store-specific commodity budget shares. This implies that

³⁹ The model includes gross household income, region, household composition (numbers of adult males, females, children in different age groups and people aged 65+), the age band and sex of the household head, the employment status of the household head, the number of cars and whether or not there is a home computer. Because we use annual expenditures we do not include any seasonal controls. We experimented with various specifications, including adding a squared term on total expenditure to allow store-specific Engel curves to be non-linear, but found this not to be important. We also experimented with including the share of total food spending by household i in store j as a right-hand side covariate, the idea being to capture variations in budget shares amongst households who rely on a particular store type for most of their shopping against those who use the store more as a top up or secondary store. We found this had little additional explanatory power for the model but led to a substantial rise in the number of predicted budget shares which were negative. Similarly, conditioning the regression estimates on households who spent more than some minimum amount in each store type made the predicted number of negative budget shares much larger, presumably because the model did not perform well out of sample for those with low total store expenditures.

there is a large amount of unobserved heterogeneity in within-store spending patterns. Of course, our estimates here are somewhat constrained by the limited set of demographic variables available in the Kantar data, though our covariates include most of the usual explanatory variables such as age, household composition and income that would feature in demand models.

Table 5.2. Explanatory power (R^2) of store-specific commodity budget share model

	Tesco	Sains	Asda	Morris	Discount	Quality	Other	Local
N	7,988	6,026	6,558	5,970	6,054	4,736	7,109	7,512
Bread	0.055	0.055	0.039	0.045	0.021	0.078	0.029	0.028
Meat	0.038	0.044	0.038	0.050	0.046	0.064	0.050	0.085
Fish	0.034	0.032	0.025	0.036	0.028	0.047	0.018	0.025
Dairy	0.013	0.017	0.019	0.023	0.023	0.032	0.019	0.090
Fats	0.037	0.021	0.021	0.026	0.018	0.018	0.020	0.012
Fruit	0.045	0.029	0.026	0.041	0.028	0.023	0.030	0.063
Vegetables	0.028	0.028	0.021	0.025	0.017	0.036	0.012	0.015
Sugars	0.029	0.034	0.024	0.020	0.030	0.036	0.017	0.109
Hot beverages	0.026	0.019	0.018	0.019	0.019	0.016	0.014	0.020
Soft drinks	0.037	0.035	0.033	0.045	0.038	0.024	0.08	0.070


Source: Author's calculations based on 2009 Kantar Worldpanel data.

Figure 5.2 shows the distribution of actual and predicted household-level budget shares resulting from this modelling exercise. Broadly, this approach captures the average shares quite well but not the distribution: the modelled shares lie over a much narrower range than is observed in the data.


One way in which we might be able to capture some of this unobserved heterogeneity is to use the implied variance of the error terms from the model to predict a (mean-zero) vector of random 'noise' for each store-specific budget share which is then added to the predicted share, and perhaps to use more sophisticated imputation methods such as multiple imputation in predicting the budget shares. In addition, a particular problem with the OLS approach used here is that households do not buy from each commodity group at each store type they visit. From the distribution of actual budget shares above (and our earlier analysis in Table 4.6), it is clear that for some commodities, even over a full year, nothing at all is spent by a small but important fraction of households – notably for fish, meat and beverages. At a store-specific level the problem is even more acute: Table 5.3 shows the proportion of households (conditional on purchasing something from a given store type) who buy nothing from each commodity group. This figure can be very large – for example, more than half of households who ever use local stores buy no meat, fish, fats or hot beverages at all, even over a year. Even within the big four supermarkets, more than 30% of shoppers never buy from the hot beverage category, more than 25% buy no fish and more than 20% no fats. This problem might suggest running a system of Tobit equations for each store to help better model the zero shares. This would also mean predicted shares could not be negative: in our estimates, around 16% of households are predicted to have at least one negative budget share (these households are dropped from the results in Figure 5.2). However, the main intention of this exercise is to assess the extent to which observable covariates are able to explain variation in within-store spending patterns, for which these simple OLS estimates provide initial evidence. If this approach were taken further by statistical agencies then more attention should be paid to the precise econometric methods used.

Figure 5.2. Actual and predicted household-level budget shares, by commodity, 2009


Bread and cereals




Meat




Fish




Dairy




Oils and fats




Fruit




Vegetables




Sugars and confectionery



Other food



Hot beverages



Cold beverages




Table 5.3. Proportion of zero store-specific budget shares, by commodity

	Tesco	Sains	Asda	Morris	Discount	Quality	Other	Local
Bread	4.4%	6.8%	6.4%	5.6%	10.9%	14.9%	8.8%	24.5%
Meat	11.1%	16.2%	12.7%	13.7%	21.6%	26.4%	16.6%	50.3%
Fish	25.4%	38.0%	32.1%	31.1%	45.5%	52.7%	37.5%	72.9%
Dairy	7.5%	13.1%	10.4%	11.7%	22.3%	37.8%	17.4%	43.6%
Fats	20.5%	33.4%	28.1%	29.2%	45.6%	76.4%	50.9%	79.9%
Fruit	13.2%	21.4%	17.7%	15.6%	20.5%	48.4%	43.9%	47.8%
Veg	7.3%	12.4%	9.2%	9.1%	13.2%	29.4%	14.7%	36.6%
Sugars	13.3%	21.5%	17.9%	18.8%	20.1%	43.6%	25.5%	20.6%
Other food	15.7%	23.0%	18.2%	19.2%	31.7%	56.9%	38.6%	52.7%
Hot beverages	30.4%	43.3%	38.9%	40.2%	53.4%	79.8%	57.8%	60.8%
Soft drinks	17.8%	29.0%	23.7%	25.4%	30.7%	65.4%	41.6%	49.7%

Source: Author's calculations based on 2009 Kantar Worldpanel data.

The findings in this section suggest that attempts to use detailed store-level expenditure patterns to impute household-specific budget shares if we observe only total spending may not be particularly successful. There is a large amount of unobserved heterogeneity in store-specific expenditure patterns not captured by the usual demographic covariates commonly featured in models of household spending. We may be able to do a reasonable job of predicting average budget shares but would be unlikely to replicate the distribution of actual budget shares, though of course more sophisticated econometrics may get us further towards that goal. Whilst from the perspective of estimating CPI budget shares getting the average right is the key objective, from a research perspective having accurate household-specific budget shares is vital (for example, in demand modelling and accurately estimating household responsiveness to price shocks, or in estimating the distributional consequences of policies, or how inflation rates vary across households). A consumer budget survey that relied heavily on imputed expenditure patterns might therefore be very undesirable. At the very least it would seem important to have a large enough subset of respondents provide detailed expenditure data without imputation, even if imputation were used to 'fill in' the spending patterns for households unable or unwilling to provide more than broad aggregates.

6. Summary and conclusions

In-home scanner data offer a hugely exciting opportunity for researchers to explore detailed questions about household expenditure behaviour and firm pricing decisions. Scanner datasets and scanners as a method of data collection are also potentially of interest for statistical agencies in terms of how they might inform, or be integrated into, budget survey data. Existing scanner data typically only cover a relatively narrow subset of total expenditures – our best guess is that the Kantar Worldpanel covers around 18% of spending, for example. Nevertheless, in terms of the potential for using existing commercial scanner datasets, organisations like the ONS or BLS may see three key opportunities:

1. Comparison and validation – do expenditures reported in scanner data tell a different story to those reported in budget survey data, and what can we learn from a more micro analysis of differences in spending across households?
2. Detail – having full knowledge of precisely what is bought helps inform the choice of representative items for inflation measurement, the weights that should be given to these items at the lowest level of disaggregation in inflation calculations, the importance of product turnover (particularly in dynamic sectors like grocery retail) and so on.
3. Duration – scanner data report spending over a long time period whereas budget surveys collect detailed spending only for a short duration. The scanner data can therefore give insight into whether short observation periods are appropriate even for non-durable commodities like food and whether the distribution of spending patterns across households is well measured.

However, crucial for any use of scanner data is an understanding of data quality. Researchers and data collectors have, over a long period, developed a good knowledge of the strengths and weaknesses of budget survey data. Scanner data are more novel and less-studied. The results in this paper and the previous literature point to some consistent differences between scanner data and budget survey data collected via diaries – in particular, that expenditures are lower in scanner data, that expenditures appear to be particularly low for certain commodity groups like alcohol and soft drinks, that it matters whether or not participants have to record non-barcoded products, and that there are differences in the relative under-reporting across household groups (in particular relating to income and numbers of children, which plausibly reflect time constraints). These differences point to modal effects. However, another finding from this study is that some differences between scanner and budget survey expenditures, which might in isolation be attributed to survey mode, may in fact be specific to the particular datasets studied. This implies that researchers and data collectors would need to be aware of the features and methods of the particular scanner data they are using in order to assess its likely benefits. Detailed cross-country comparisons of scanner and budget survey data which, as far as possible, apply common methods would be a useful way to tease out modal effects. Collaboration between national statistical agencies and researchers in different countries to share knowledge and carry out such research would be desirable. For example, statistical agencies have access to much more disaggregate information about households and shopping trips from the budget survey than are made available to external researchers. This information could give more detailed insight into the differences between diary-based and scanner expenditures, and so should be made available for this kind of analysis. Another possibility would be for statistical agencies and commercial data collectors to collaborate directly. For example, to address issues around sample representativeness in scanner surveys, agencies could supply a household sample drawn at random from the population and ask the data collectors to try and incorporate them into the sample, seeing who refuses, who drops out quickly, who appears to co-operate effectively and so on. In some countries, notably Nordic ones, detailed population-level data linked to identification records is maintained which would allow for detailed analysis of these sorts of issues to be carried out.

Our comparisons between scanner data and budget survey data point to a number of key conclusions. Most notably, scanner data are prone to large numbers of weeks in which no purchases at all are recorded, even by ostensibly active households. We do not have a clear understanding of what drives this, but these zeroes explain a large part of the raw gap between scanner and budget survey expenditures. However, even when we restrict ourselves to a sample of households who are faithful reporters, and eliminate zero spending weeks altogether, a large expenditure gap remains. From the UK data at least, this gap is not accounted for by observable demographic differences between samples. More reassuringly, patterns of spending across surveys are similar even though the levels are different, and re-weighting the food CPI using survey-specific budget shares has very little impact. Further, the relationship between total expenditures and commodity-level budget shares are also very similar across surveys.

We find compelling evidence that asking households to record detailed expenditures over a short horizon leads to an inaccurate picture of the distribution of commodity-level budget shares, but not average spending patterns. The two week duration of budget diaries both in the LCF and the CE is probably not long enough to get a good measure of true food consumption patterns, generating too much variability in the distribution of budget shares and too many zero expenditures for particular commodities.

We also consider the possibility that scanner data may be useful to ‘drill down’ into aggregate store-level expenditures and impute commodity-specific expenditures. We find substantial differences in expenditure patterns across store types which supports breaking down spending by store rather than imputing based on total spending in all stores. However, we find very little relationship between expenditure patterns and observable covariates. Thus whilst an imputation approach might get the averages broadly right, the distribution of imputed budget shares is much less dispersed than the distribution of observed budget shares. More sophisticated imputation or econometric methods may help, and we could also be hamstrung in this analysis by the relatively limited set of demographic information in the scanner data. In general, though, it would not appear to be sensible to rely too heavily on imputation methods to obtain detailed measures of spending.

As mentioned in the introduction, rather than relying on existing commercial scanner data, agencies like the BLS or ONS might be interested in the idea of establishing their own panel of households using scanner methods to record their spending on an ongoing basis. If this were integrated with the budget survey, there would be clear benefits in terms of immediately stripping out demographic differences as a source of variation between scanner and other expenditure data, and in ensuring that the scanner sample were collected using proper random sampling methods. It would also allow, for the first time, detailed scanner data for particular commodities to be linked to more general expenditure patterns. Existing scanner datasets are limited in coverage to food at home and a small number of non-food purchases. Knowing how much households spent on food out as well as other expenditure categories would be extremely beneficial, as would any other links that could be made between the detailed information in scanner data and wider household characteristics relating to health, dwelling characteristics, durable ownership and so on. A new dataset collected by a statistical agency would also enable experimental analysis of the impact of survey mode, different scanner devices, different reporting requirements (e.g. non-barcoded items), the incentives for participation, and so on. By providing ‘gold standard’ evidence of these issues using randomised trials, this would be extremely valuable not just for the agencies themselves but also for researchers into survey design and data users.

We end by offering some thoughts, based on this paper and previous studies, on other ways in which new scanner data could improve on existing commercial datasets. Firstly, one of the major limitations of the Kantar Worldpanel relative to the LCF (and other datasets collected by statistical agencies and used frequently by researchers) is the relatively poor demographic information. Scanner data are collected for

commercial marketing and market research purposes. The main clients are retailers and grocery manufacturers who may need only relatively basic demographic profiles of the households in the sample. Whilst the data follow the same households over time, one of the key findings from previous analysis of the Kantar data is that important demographic transitions such as changes in employment status are not well recorded. This may not be true of scanner datasets in general, but we are not aware of research which has looked at this question for other data. Improperly recorded transitions hugely restrict the usefulness of the panel nature of the data (for example, in analysing how detailed food spending behaviour changes around retirement or in response to unexpected income shocks). New scanner data collected by statistical agencies could presumably obtain much more detailed demographic information about household members to match the kinds of data familiar in budget surveys, and ideally would also take more care to ensure changes in demographics were captured if the same households were followed for an extended period. This should extend not only to knowing whether a transition occurred but also when.

Secondly, a key unresolved problem with scanner data is to be the high prevalence of weeks in which no expenditures are reported. This does not match up to budget survey data. Understanding this better would be a key contribution of a new scanner dataset. For example, if no expenditures are recorded over a week or a fortnight, households could be prompted with some form of contact from the statistical agency to see what has driven this and then code this into the data. This would allow the agency and researchers to distinguish genuine zeroes – holidays, weeks in which households simply used stocks of food or ate out and so on – from non-genuine ones where trips were made but not reported for one reason or another. Indeed it might be possible in this latter case to get retrospective information (from till receipts or recall questions) or to impute missing trips altogether.

Thirdly, it seems important that any scanner data aims to be, as far as possible, a complete record of the shopping trip. Most existing panels now require only a subset of respondents to report non-barcoded purchases. In the Kantar data, this was justified by better reporting compliance for non-barcoded items. However, our analysis suggests that this result is at least partly driven by demographic differences between those who do and do not report non-barcoded purchases. It may also be differences in the type of scanner device, rather than the lack of random weight reporting which gives this result. Our preference would be that households were asked to record random weight purchases as well since not doing so risks substantial bias in estimating consumption of certain commodities, particularly fruit and vegetables. The aim should be to minimise the additional burden imposed by this requirement. One problem for commercial scanner data is the need to include very detailed product characteristics which needed to be manually input by respondents for random weight items (e.g. the country of origin of different fruits, whether organic, the weight purchased). To the extent that such details are not needed by statistical agencies they could be dispensed with altogether, or taken where available from till receipts, or imputed. If recording random weight items simply required scanning a barcode from a generic booklet or tapping an icon on the barcode reader itself then it may be little more onerous than recording barcoded products.

Finally, we would certainly advocate that any new scanner data made use of till receipts as well as the in-home scanner. Studies of Nielsen data highlighted the problems of relying on imputed prices obtained from centralised store-level databases given the growth of personalised pricing through vouchers and loyalty cards, as well as store-specific special offers. Participants should be encouraged to send in receipts, and it may be that technology could enable this to be done digitally (scanned receipts or optical character recognition devices attached to computers) to integrate it with the data collection process more closely.

References

- Ahmed, N., M. Brzozowski and T. Crossley (2006), *Measurement Errors in Recall Food Consumption Data*, IFS Working Paper 06/21 (<http://www.ifs.org.uk/publications/3752>)
- Aguiar, M. and E. Hurst (2007), 'Life-cycle prices and production', *American Economic Review*, 97, 1533–1559
- 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*, National Poverty Center conference paper (http://www.npc.umich.edu/news/events/consumption06_agenda/Attanasio-Battistin-Leicester.pdf)
- Ayers, R., P. Hossack and C. Payne (2010), *Living Costs and Food Survey Technical Report for the Survey Year January – December 2009*, Office for National Statistics
- Blow, L., A. Leicester and Z. Smith (2004), *Consumption Trends in the UK, 1975 – 99*, IFS Report 65, London: IFS (<http://www.ifs.org.uk/comms/r65.pdf>)
- Brewer, M., A. Goodman and A. Leicester (2006), *Household Spending in Britain: What Can it Teach Us About Poverty?*, Bristol: The Policy Press (<http://www.jrf.org.uk/sites/files/jrf/9781861348555.pdf>)
- Broda, C., E. Leibtag and D. E. Weinstein (2009), 'The role of prices in measuring the poor's living standards', *Journal of Economic Perspectives*, 23(2), 77–97
- Cole, N. (2011), *Design of the National Household Food Acquisition and Purchase Survey (FoodAPS)*, Presentation to Committee on National Statistics Household Survey Producers Workshop (http://www.bls.gov/cex/hhsrvywrkshp_cole.pdf)
- DEFRA (2011), *Attitudes and Behaviours Around Sustainable Food Purchasing*, London: DEFRA (<http://www.defra.gov.uk/statistics/files/defra-stats-foodfarm-food-attitudes-report-110406-mainreport.pdf>)
- Duly, A., T. Garner, E. Keil, S. Reyes-Morales and C. Wirth (2003), *The Consumer Expenditure Survey and AC Nielsen Survey: a Data Comparison Study*, unpublished internal BLS document.
- Einav, L., E. Leibtag and A. Nevo (2008), *Not-so-classical measurement errors: a validation study of Homescan*, Stanford Institute for Economic Policy Research Discussion Paper 08-07 (<http://www-siepr.stanford.edu/repec/sip/08-007.pdf>)
- Feenstra, R. C. and M. D. Shapiro (2003), *Scanner Data and Price Indexes*, Chicago: University of Chicago Press (<http://www.nber.org/books/feen03-1/>)
- Garner, T., R. McClelland and W. Passero (2009), *Strengths and Weaknesses of the Consumer Expenditure Survey from a BLS Perspective* (http://www.bls.gov/cex/pce_compare_199207.pdf)
- Griffith, R., E. Leibtag, A. Leicester and A. Nevo (2009), 'Consumer shopping behavior: how much do consumers save?', *Journal of Economic Perspectives*, 23(2), 99–120

- Griffith, R. and M. O'Connell (2009), 'The use of scanner data for research into nutrition', *Fiscal Studies*, 30 (3-4), 339-365
- Harding, M., E. Leibtag and M. Lovenheim (2012, forthcoming), 'The heterogeneous geographic and socioeconomic incidence of cigarette taxes: evidence from Nielsen Homescan data', *American Economic Journal: Economic Policy*
- Harris, J. M. (2005), 'Using Homescan data and complex survey design techniques to estimate convenience food expenditures', American Agricultural Economics Association Conference Paper (<http://ageconsearch.umn.edu/bitstream/19344/1/sp05ha07.pdf>)
- Huffman, S. and H. Jensen (2004), *Demand for Enhanced Foods and the Value of Nutritional Enhancements of Food: the Case of Margarines*, American Agricultural Economics Association Conference Paper (<http://ageconsearch.umn.edu/bitstream/20205/1/sp04hu05.pdf>)
- Jin, W., R. Joyce, D. Philips and L. Sibieta (2011), *Poverty and Inequality in the UK: 2011*, IFS Commentary 118, London: IFS (<http://www.ifs.org.uk/comms/comm118.pdf>)
- Leicester, A., C. O'Dea and Z. Oldfield (2008), *The Inflation Experience of Older Households*, IFS Commentary 106, London: IFS (<http://www.ifs.org.uk/comms/comm106.pdf>)
- Leicester, A. and Z. Oldfield (2009), *An analysis of consumer panel data*, IFS Working Paper W09/09 (<http://www.ifs.org.uk/publications/4468>); see also Leicester, A. and Z. Oldfield (2009), 'Using scanner technology to collect expenditure data', *Fiscal Studies*, 30 (3-4), 309-337
- Levell, P. and Z. Oldfield (2011), *The Spending Patterns and Inflation Experience of Low-Income Households Over the Past Decade*, IFS Commentary 119, London: IFS (<http://www.ifs.org.uk/comms/comm119.pdf>)
- Lusk, J. L. and K. Brooks (2011), 'Who participates in household scanning panels?', *American Journal of Agricultural Economics*, 93(1), 226-240
- Mathiowetz, N., K. Olson and C. Kennedy (2011), *Redesign Options for the Consumer Expenditure Survey* (http://www.bls.gov/cex/redwrkshp_pap_abtsrbirecommend.pdf)
- O'Dea, C. and T. Crossley (2010), *The Wealth and Saving of UK Families on the Eve of the Crisis*, IFS Report 71, London: IFS (<http://www.ifs.org.uk/comms/r71.pdf>)
- Office for National Statistics (2010a), *Consumer Trends, Q2 2010*, Newport: ONS (<http://www.ons.gov.uk/ons/rel/consumer-trends/consumer-trends/q2-2010/consumer-trends-no--57.pdf>)
- Office for National Statistics (2010b), *Consumer Price Indices Technical Manual: 2010 Edition*, Newport: ONS (<http://www.ons.gov.uk/ons/guide-method/user-guidance/prices/cpi-and-rpi/consumer-price-indices-technical-manual---2010.pdf>)
- Perloff, J. M. and M. Denbaly (2007), 'Data needs for consumer and retail firm studies', *American Journal of Agricultural Economics*, 89 (5), 1282-1287
- Slesnick, D. (2001), *Consumption and Social Welfare*, New York: Cambridge University Press

Tanner, S. (1998), 'How much do consumers spend? Comparing the FES and National Accounts', in Banks, J. and P. Johnson (eds.), *How Reliable is the Family Expenditure Survey? Trends in Incomes and Expenditures Over Time*, IFS Report 57, London: IFS (<http://www.ifs.org.uk/comms/r57.pdf>)

Triplett, J. (1997), *Measuring Consumption: the Post-1973 Slowdown and the Research Issues* (http://www.brookings.edu/~media/Files/rc/articles/1997/0702useconomics_triplett/19970702.pdf)

Tucker, C. (2011), *Using Multiple Data Sources and Methods to Improve Estimates in Surveys*, presented at BLS Household Survey Producers Workshop, June 2011 (http://www.bls.gov/cex/hhsrvywrkshp_tucker.pdf)

UK Payments Council (2010), *The Way We Pay 2010: The UK's Payment Revolution* (http://www.paymentscouncil.org.uk/files/payments_council/the_way_we_pay_2010_final.pdf)

UK Payments Council (2011), *Quarterly Statistical Report Q2 2011* (http://www.paymentscouncil.org.uk/files/payments_council/minutes_of_payments_council_board/statistical_release_2011_q2.pdf)

USDA (2009), *The Consumer Data and Information Program: Sowing the Seeds of Research* (<http://www.ers.usda.gov/publications/ap/ap041/ap041.pdf>)

Westat (2011a), *Data Capture Technologies and Financial Software for Collecting Consumer Expenditure Data* (<http://www.bls.gov/cex/ceother2011westat.pdf>)

Westat (2011b), *Redesign Options for the Consumer Expenditure Survey* (http://www.bls.gov/cex/redwrkshp_pap_westatrecommend.pdf)

Zhen, C., J. L. Taylor, M. K. Muth and E. Leibtag (2009), 'Understanding differences in self-reported expenditures between household scanner data and diary survey data: a comparison of Homescan and Consumer Expenditure Survey', *Review of Agricultural Economics*, 31(3), 470-492

Appendix A: Kantar sample selection

Table A.1 reports the results of a probit regression, where the left-hand side takes a value of 1 for Kantar households observed in 2009 who form part of the ‘regular reporter’ sample which underlies much of the analysis in Section 3.2, and 0 for households who are selected out of this sample either because they are not active for the full 52-week period or because they have at least one continuous period of more than 6 weeks during which they report no food and drink expenditures at all. The right hand side is a vector of controls for different observable demographic characteristics.

Table A.1. Probit regression for the effect of demographic characteristics on being part of the ‘regular reporter’ Kantar sample, 2009

	Coefficient	Std. Error
<i>Number of cars (base = 0)</i>		
One car	0.0475*	0.0268
Two or more cars	0.0530*	0.0310
Unknown	-2.119***	0.147
<i>PC in household (base = no)</i>		
Has PC	-0.0488	0.0317
<i>Region (base = South East England)</i>		
NE England	-0.0151	0.0429
NW England	-0.0450	0.0318
Yorkshire & Humber	-0.0159	0.0346
E Midlands	0.0652*	0.0356
W Midlands	0.0326	0.0345
East England	0.0247	0.0327
London	-0.102***	0.0345
SW England	-0.0158	0.0340
Wales	0.0296	0.0427
Scotland	-0.0668*	0.0353
<i>Age of household head (base = 45-49)</i>		
<25	-1.324***	0.0700
25-29	-0.780***	0.0423
30-34	-0.412***	0.0388
35-39	-0.271***	0.0366
40-44	-0.127***	0.0352
50-54	0.0386	0.0377
55-59	0.133***	0.0392
60-64	0.147***	0.0434
65-69	0.0341	0.0665
70-74	0.289***	0.0709
75-79	0.409***	0.0769
80+	0.271***	0.0858
<i>Sex of household head (base = male)</i>		
Female head	-0.220***	0.0222
<i>Employment status of household head (base = works 30+ hours/week)</i>		
Works 8-29 hours	-0.0257	0.0262
Works <8 hours	0.0286	0.0583
Unemployed	-0.0609	0.0570
Retired	-0.0139	0.0360
In education	0.0000	0.0939
Not working	-0.0365	0.0287
Unknown	-0.431***	0.130
<i>Household gross annual income group (base = £10 – 20k)</i>		
£0 – 10k	0.0602*	0.0336

£20 – 30k	0.0127	0.0275
£30 – 40k	0.0201	0.0309
£40 – 50k	0.0475	0.0359
£50 – 60k	-0.0129	0.0433
£60 – 70k	-0.0505	0.0570
£70k +	-0.0523	0.0520
Unknown	-0.223***	0.0260
<i>Number of adult males (base = 1)</i>		
0 males	-0.220***	0.0268
2 males	0.0303	0.0301
3+ males	-0.0922	0.0624
<i>Number of adult females (base = 1)</i>		
0 females	-0.275***	0.0350
2 females	-0.0564*	0.0290
3+ females	-0.227***	0.0666
<i>Number of children aged 0-4 (base = 0)</i>		
1 child	-0.221***	0.0284
2 children	-0.261***	0.0446
3+ children	-0.544***	0.159
<i>Number of children aged 5-10 (base = 0)</i>		
1 child	-0.0255	0.0268
2 children	-0.128***	0.0405
3+ children	-0.219**	0.108
<i>Number of children aged 11-17 (base = 0)</i>		
1 child	-0.0354	0.0272
2 children	-0.117***	0.0390
3+ children	-0.196***	0.0929
<i>Number of people aged 65+ (base = 0)</i>		
1 person	0.0720	0.0464
2 people	-0.0702	0.0602
3+ people	-0.790	0.482
<i>Does household record non-barcode items (base = Yes)</i>		
No	0.148***	0.0184
# households	26,655	
Pseudo-R ²	0.107	

*** = p<0.01; ** = p<0.05; * = p<0.1

Appendix B: Observable demographic comparisons

The tables below compare observable demographic characteristics of the LCF 2009 sample (excluding Northern Ireland) and the Kantar 2009 'regular reporter' sample (see Appendix A). Respective sample sizes are 5,150 and 15,752 households.

Gross annual household income

	All		Excluding unknown		
	Kantar %	LCF %	Kantar %	LCF %	K ÷ LCF
Unknown	18.1	0.0			
< £10,000	9.9	13.4	12.1	13.4	0.90
£10,000–£19,999	22.4	23.9	27.3	23.9	1.15
£20,000–£29,999	17.7	17.9	21.6	17.9	1.20
£30,000–£39,999	13.1	13.6	15.9	13.6	1.17
£40,000–£49,999	8.6	10.4	10.5	10.4	1.01
£50,000–£59,999	4.9	6.6	5.9	6.6	0.90
£60,000–£69,999	2.3	4.3	2.9	4.3	0.67
≥ £70,000	3.1	9.9	3.7	9.9	0.38

Computer in the household?

	Kantar %	LCF %	K ÷ LCF
No	10.5	24.2	0.43
Yes	89.5	75.8	1.18

Number of cars

	Kantar %	LCF %	K ÷ LCF
Zero	12.9	21.6	0.59
One	50.7	45.0	1.13
Two or more	36.5	33.4	1.09

Region

	Kantar %	LCF %	K ÷ LCF
North East	4.8	4.5	1.07
North West	11.6	11.2	1.04
Yorks & Humber	8.8	9.2	0.95
East Midlands	8.5	7.5	1.13
West Midlands	9.2	10.1	0.91
East of England	10.8	9.6	1.13
London	8.5	8.8	0.96
South East	15.1	13.4	1.12
South West	9.5	10.0	0.95
Wales	5.1	5.2	0.97
Scotland	8.2	10.4	0.79
North East	4.8	4.5	1.07

Gender of household head

	Kantar %	LCF %	K ÷ LCF
Male	44.0	74.7	0.59
Female	56.0	25.3	2.21

Age of household head

	Kantar %	LCF %	K ÷ LCF
<25	0.6	3.0	0.20
25–29	4.0	5.3	0.75
30–34	7.5	7.1	1.06
35–39	10.1	8.7	1.15
40–44	11.6	9.9	1.16
45–49	10.6	10.1	1.05
50–54	10.3	9.0	1.14
55–59	10.6	9.7	1.09
60–64	9.9	9.6	1.03
65–69	8.3	7.7	1.08
70–74	8.6	6.8	1.26
75–79	5.4	6.0	0.90
≥80	2.7	7.1	0.38

Employment status of household head

	Kantar %	LCF %	K ÷ LCF
Works 30+ hours	41.8	51.5	0.81
Works 8–29 hours	13.7	5.7	2.42
Works <8 hours	2.1	0.4	5.68
Unemployed	1.9	3.9	0.50
Retired	28.6	29.1	0.98
Full-time education	0.5	0.9	0.60
Not working	11.3	8.6	1.32

Number of adults (aged 18+)

	Males			Females		
	Kantar %	LCF %	K ÷ LCF	Kantar %	LCF %	K ÷ LCF
Zero	16.5	22.7	0.72	9.4	13.5	0.69
One	71.6	69.2	1.03	79.7	78.7	1.01
Two	10.1	7.1	1.43	9.6	6.8	1.40
Three+	1.8	1.0	1.89	1.4	1.0	1.45

Number of children


	Ages 0–4			Ages 5–10			Ages 11–17		
	Kant %	LCF %	K ÷ L	Kant %	LCF %	K ÷ L	Kant %	LCF %	K ÷ L
Zero	87.9	87.7	1.00	84.4	87.5	0.96	82.9	85.5	0.97
One	9.4	9.6	0.97	11.4	9.0	1.27	11.7	9.6	1.22
Two	2.6	2.4	1.10	3.8	3.3	1.16	4.7	4.3	1.10
Three+	0.1	0.3	0.52	0.4	0.3	1.59	0.7	0.6	1.13

Number of people age 65+


	Kantar %	LCF %	K ÷ LCF
Zero	71.0	70.9	1.00
One	16.5	18.6	0.89
Two	12.5	10.5	1.19
Three+	0.0	0.1	0.33

Appendix C: Food commodity Engel Curves


Bread and cereals




Meat



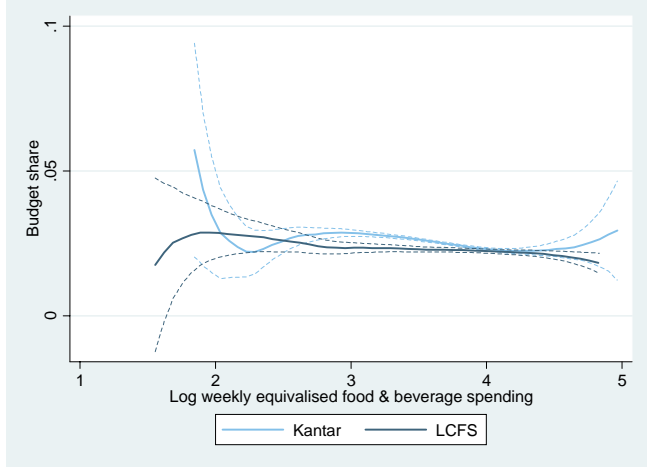
Fish




Dairy




Oils and fats




Fruit




Vegetables




Sugars and confectionery




Other food



Hot beverages



Cold beverages



Source: Author's calculations based on 2009 LCF and Kantar Worldpanel data.