Volume Title: Improving the Measurement of Consumer Expenditures

Volume Author/Editor: Christopher D. Carroll, Thomas F. Crossley, and John Sabelhaus, editors

Series: Studies in Income and Wealth, volume 74

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-12665-X, 978-0-226-12665-4

Volume URL: http://www.nber.org/books/carr11-1

Conference Date: December 2-3, 2011

Publication Date: May 2015

Chapter Title: The Potential Use of In-Home Scanner Technology for Budget Surveys

Chapter Author(s): Andrew Leicester

Chapter URL: http://www.nber.org/chapters/c12669

Chapter pages in book: (p. 441 – 491)

16

The Potential Use of In-Home Scanner Technology for Budget Surveys

Andrew Leicester

16.1 Introduction

In-home scanner expenditure data are collected via a barcode reader installed in the home. 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 behavior. Scanner data have long been used for marketing studies, and increasingly in the economics literature to explore questions relating to consumer, retailer, and manufacturer behavior (recent examples include Griffith et al. [2009]; Broda, Leibtag, and Weinstein [2009]; 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 (see Parker, Souleles, and Carroll [chapter 3, this volume] for a discussion of the virtues of longitudinal data for research purposes). Panel expenditure data sets are comparatively rare. National budget surveys are usually a cross section, and

Andrew Leicester is a manager in the public policy practice of Frontier Economics.
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. For acknowledgments, sources of research support, and disclosure of the author's material financial relationships, if any, please see http://www.nber.org/chapters/c12669.ack.

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 behavior in response to shocks or policy interventions. For example, Harding, Leibtag, and Lovenheim (2012) use scanner data to explore the impact of changes in cigarette taxes on prices and consumption. There is also growing interest among policymakers 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 noncommercial surveys: the US Department of Agriculture is planning to use scanners, alongside other data collection methods, to record detailed food purchase behavior as part of the Food Acquisition and Purchase Survey (FoodAPS).[1]

The main aims of this chapter are to consider what role home 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).[2] A limited role is 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 microlevel comparisons since few surveys (besides the budget survey itself) collect detailed spending information. Scanner data offer such a possibility. They also record household expenditures over long periods of time, which allows us to explore how the time-limited nature of budget surveys (the UK LCF records spending over just two weeks) affects the spending patterns that are observed.[3]

A more involved role for home scanner data in budget surveys might be as part of the data collection process itself (Mathiowetz, Olson, and Kennedy 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 that 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. Scanner data could also be used for imputation. As a way to reduce respondent burden, some commentators (e.g., Tucker 2011) have suggested asking only limited questions about aggregate category-level expenditure in budget surveys.

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

2. The discussion here represents a condensed version of Leicester (2012), which provides further detail and analysis in a number of areas.

3. This has important implications for attempts to use short-run spending information to make inferences about the distribution of living standards (for example, Attanasio, Battistin, and Leicester 2006; Brewer, Goodman, and Leicester 2006).

Information from other sources such as scanner data could then be used to break these down into detailed spending patterns using imputation methods.

This chapter focuses on how statistical agencies might make use of established home scanner surveys collected by market research companies. An alternative would be for statistical agencies to establish and maintain their own scanner data. While costly, this would offer a number of advantages. It would override concerns about outsourcing part of the data collection process to commercial organizations. It would allow controlled experimentation to explore modal effects. It would ensure that the information on demographics and the statistical properties of the data were of sufficiently high quality to be useful for national statistics and research purposes. We offer some thoughts on the scope for establishing a separate scanner survey in the conclusions.

The rest of the chapter is organized as followed. Section 16.2 describes the data sets that underlie most of our analysis: the UK LCF and home scanner data from Kantar Worldpanel. Section 16.3 briefly surveys the existing literature comparing home scanner data to other data. In section 16.4, we compare food expenditures reported in home scanner data, budget survey data, and national accounts data. 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 16.5 explores the prospects for using detailed spending patterns from home scanner data to impute budget shares for households when all we observe are total expenditures. Section 16.6 offers some overall thoughts and conclusions.

16.2 Data

16.2.1 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. The survey has been renamed twice, each time undergoing some structural changes (though the coverage and main methods of the data have remained essentially unchanged). Until 2001 it was the Family Expenditure Survey (FES). It then merged with a second, related survey recording nutritional intake at the household level and became known as the Expenditure and Food Survey (EFS). It was then renamed the LCF in 2008. We use LCF throughout to refer to this data. Sampling is carried out via stratified random sampling, with strata based on region, socioeconomic status, and car ownership. Northern Ireland is oversampled, but survey weights are provided to ensure the weighted sample is nationally representative. The response rate in 2010 was 50 percent; 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 age sixteen and over. Children age seven to fifteen receive a simplified diary.[4] Participants record all their expenditures over the period, attaching till receipts where possible 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*.[5]

16.2.2 Kantar Worldpanel

Kantar is a market research company that operates a number of surveys of consumer behavior, including Worldpanel, which is collected in a number of countries. In Britain, one use of Worldpanel data is to estimate market shares of the major supermarkets.[6] 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, since then it has risen to around 25,000. Participants are recruited from a range of address sources using quota sampling, though Northern Ireland is excluded. Household weights are derived that 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 participate for as long as they wish, and receive points redeemable for consumer goods as an incentive to do so. Participating households are issued 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 "fast-moving consumer goods"—essentially food and grocery products, including things like cleaning products and personal care items. Alcohol (purchased off-licence) is included, but tobacco is not. Leicester (2012) estimates that the set of products contained in the Worldpanel data make up something like 18 percent of all nonhousing expenditure.[7]

4. 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.

5. 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.

6. See http://www.kantarworldpanel.com/.

7. His analysis suggests that just over one-third of all spending would be amenable to in-home scanning.

Purchases from all retailers, not just supermarkets, are in principle recorded, as are online grocery purchases. The data are at the transaction level—typically up to a million separate transactions are recorded in a week of data. Detailed information on the characteristics of the products purchased is recorded, including the macronutritional composition of food items. Until 2006, all households were asked to report nonbarcoded food and grocery purchases using a booklet of generic barcodes. Details of the product characteristics for these items (such as weight, country of origin, flavor, and so on) were also entered manually via the scanner device. This increases responder burden, and so since 2006 some households were no longer required to report these items and were issued with a simpler scanner unit. Information on the price paid is obtained from till receipts that are mailed in to Kantar, who match the price to the purchase record. Where no receipts are available, prices are taken from centralized 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.

Household demographic characteristics are recorded in a baseline telephone interview, and then updated every nine months or so. 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.

16.2.3 Mapping Kantar Expenditure Data to LCF Data

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. Making comparisons requires us to match individual products from the Kantar data into equivalent LCF expenditure codes. We use detailed information on the sorts of products that make up each expenditure code supplied with the LCF documentation, and the detailed product characteristics in the Kantar data, to make this match as accurately as possible, though inevitably there is some judgement in this process. Having created this mapping, we then further aggregate expenditures into commodity groups to match those defined in the UK Consumer Prices Index (CPI). This definition matches the level at which disaggregated expenditure information is available from the national accounts, making comparisons to aggregate data more straightforward. Our analysis covers only food and drink purchases.

In principle, of course, when making comparisons across data sets, we could look at much more disaggregate commodity groups. The LCF includes seventy-three distinct food and drink codes, so this would be the most disaggregate comparison possible.[8] Finer disaggregation may be useful to under-

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

stand exactly where differences between scanner and other data sets arise and what might be driving that. However, as discussed in Leicester and Oldfield (2009), the more disaggregate the comparison the less confident we can be about the mapping between Kantar and LCF expenditures. 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, for example, plain frozen fish fillets) or "processed" (largely meaning they are preprepared or flavored in some way). Fresh and processed products have distinct LCF codes but fall into the same CPI groups, such that at the CPI level of aggregation there is more certainty that we are comparing like-with-like (spending on fish or meat, say).

16.3 Previous Research

This chapter adds to a small but growing literature exploring data quality issues for home scanner data. There is a parallel literature on "storescan" data collected from items passing through tills. There are clearly complementarities between store-level and consumer-level scanner data. Store-level scanner data might be employed for some of the uses to which budget survey data are put. In a previous NBER volume, Feenstra and Shapiro (2003) considered the possible usefulness of store-level scanner data for measuring prices and price indices, and provides a good introduction to such data.

A number of previous studies have compared home scanner data to budget survey data. Any regularities that emerge from these comparisons would be strongly suggestive of survey mode effects from the use of scanners. However, definitive statements about modal effects would require experimental evidence, which held, as far as possible, other factors constant. Such evidence does not appear to exist at the moment. A valuable contribution from statistical agencies collecting their own scanner survey data would be the ability to carry out controlled experimental analysis.

The strongest common finding is that average expenditure levels are markedly lower in scanner data than in budget survey data. Duly et al. (2003) compare AC Nielsen Homescan data to CE diary survey data from 2000. Overall, scanner expenditures were about two-thirds of the budget survey level. Alcohol and tobacco expenditures were about half the budget survey level in the scanner data. Using British data from Kantar Worldpanel and the Living Costs and Food Survey, Leicester and Oldfield (2009) find that, in 2005, weekly total food expenditures in scanner data were about 20 percent below those in budget surveys on average. Using comparisons of nutritional intake, Griffith and O'Connell (2009) find that the number of calories reported in British scanner data were around 23—52 percent lower than in budget survey data depending on the household type studied. They find strong evidence, as do we below, that a large part of this gap is driven by weeks in which no food at all is purchased in scanner data, though it does not account for all of the difference.

Previous studies have not found common results on the extent to which this under-reporting in scanner data is consistent across product categories. Using US data from 2002–2005, Zhen et al. (2009) find scanner expenditures 50 percent below CE levels for a number of commodity groups, particularly categories where nonbarcoded items are common such as meat, fruit, vegetables, and fish. For groups where almost all items are barcoded, such as confectionery and processed fruits, they find essentially no differences in expenditure. In contrast, Leicester and Oldfield (2009) find little evidence in British data of significant differences across groups, with the exception of alcohol. One likely reason for the difference is that in the US data, only 20 percent of households were required to record nonbarcoded items, whereas all households in the British data were asked to do so. Leicester and Oldfield (2009) can explore this directly, since from 2006 only some households were asked to report nonbarcoded purchases. They find this has a substantial effect. Households reporting their nonbarcoded purchases recorded 24 percent less expenditure on fruit than found in budget surveys; those not reporting them recorded 44 percent less spending.

There is also disagreement on the extent to which demographic differences between samples help account for the lower expenditures in scanner data. Zhen et al. (2009) use a regression model to strip out observable demographic differences between data sets and argue that, in combination with the nonbarcoded items issue, they largely account for the spending differences. Leicester and Oldfield (2009) conclude that demographic differences *accentuate* the gaps between data sets. 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 data sets rises from 20 percent to 25 percent. These contrasting findings probably result in part from differences across countries in the relative sample compositions between scanner and budget survey data. For example, in the United States, scanner households appear to have fewer members on average than those in the budget survey (Huffman and Jensen 2004) while the reverse is true in Britain (Leicester and Oldfield 2009). The contrast may also reflect differences in the set of observable demographics common to scanner and budget survey data sets. 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.

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 among a sample of households headed by someone employed and age fifty or over, just 2.9 percent were observed to be unemployed a year later. This compared to 11.4 percent of a similar sample constructed from the British Household Panel Survey (BHPS), the main panel data set in the United Kingdom. There does not appear to be similar evidence from any study of other scanner panels that would shed light on whether this issue was common to scanner data in general or particular to the Kantar Worldpanel.

The main lessons from these findings seem to be that scanner data record less spending than budget surveys, but differences in methods across scanner data sets lead to different conclusions about the extent to which this is driven by observable factors rather than being a modal effect of scanner data *per se*. 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 and what that might mean for the data that are collected. There would appear to be a strong case for collaboration between statistical agencies, researchers, and data collectors to better understand these issues. Without experimental methods, the next best approach to tease out modal effects may be to try and make comparisons of scanner and budget survey data across countries that follow, as closely as possible, identical methodologies to see which findings are robust.

Aside from comparisons to budget survey data, some papers have attempted to explore reporting issues in scanner data more directly. Einav, Leibtag, and Nevo (2008) 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 percent or so of trips recorded in Homescan were not found in the retailer data, and that around half the trips that were reported in the retailer data were not observed in Homescan. On matched trips, the scanner data reported on average 10–15 percent fewer items, mostly small consumables that may be consumed before entering the home. The authors found significant problems in reporting prices in scanner data. The price reported in the Homescan record failed to match the loyalty card recorded price about half the time. However, this appears to be a particular problem with the way in which prices are imputed into Homescan data based on centralized records of chain week-level prices. This means store- or consumer-specific prices are missed in Homescan. In British Worldpanel data, prices are taken from till receipts and are rarely imputed.[9] Einav,

9. Where receipts are not sent in, imputation methods may be used. Since national supermarket chains in the United Kingdom 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.

Leibtag, and Nevo (2008) recommend that a similar approach be adopted for Homescan.

With in-home scanners, there may concern about people taking some time to adapt to the technology before they report reliable data. Leicester and Oldfield (2009) find that reported expenditures were highest in the first few weeks of participation. After about six months, households spent about 5 percent less than in their first week, on average. 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 behavioral reaction to participation. For statistical agencies thinking about scanner data for budget surveys, the interesting comparison is with the current survey approach. Ahmed, Brzozowski, and Crossley (2006) find that in the Canadian Food Expenditure Survey, spending drops by 9 percent between the first and second week, on average.

A further issue relates to quota sampling methods used in scanner data (Tucker 2011; Zhen et al. 2009; USDA 2009; Harris 2005) rather than random probability or stratified sampling used in budget 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.

A final point relates to *unobservable* differences in the characteristics of households who participate in scanner data and budget surveys. If significant, such differences could have important repercussions for spending records and researchers. Lusk and Brooks (2011) find that households in two large US scanner samples, Homescan and IRI, appear to be more price responsive than the population at large, even conditional on observable characteristics. They offer two possible explanations. First, participating in scanner data may make households more aware of their purchasing behavior and thus more price sensitive. Second, 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, which would be an interesting extension. A fascinating study would be to estimate demand models using budget survey data and scanner data aggregated to the same level to see whether they give similar results.

16.4 Comparing Scanner and Other Expenditure Data

Surveys of household expenditure are prone to error. Participants could deliberately or accidentally misreport their purchases, or change their usual shopping behavior as a result of participation. Data validation is therefore 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 survey data from Kantar Worldpanel and the LCF to aggregate data from the ONS national accounts. Second, we make microlevel comparisons between the two surveys. In each case we explore not just total spending, but also expenditure patterns. 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 or estimating demand models. Comparisons of both are therefore important.

16.4.1 Comparison to National Expenditure Aggregates

Several recent papers explore the quality of budget survey data by comparing them to aggregate expenditure data. Examples in the United States include Triplett (1997), Slesnick (2001), Attanasio, Battistin, and Leicester (2006), and in the United Kingdom include Tanner (1998), Blow, Leicester, and Smith (2004), Attanasio et al. (2006), and O'Dea and Crossley (2010). Key findings from these studies are:

- Spending reported in the US CE makes up about 70 percent of aggregate levels. In the UK LCF, the figure is around 80 percent. Both have worsened over time. The decline in UK coverage is particularly noticeable from the early 1990s.
- Similar trends occur for food at home. In the United States, coverage fell from more than 75 percent in the 1980s to around 65 percent early in the first decade of the twenty-first century. In the United Kingdom, coverage fell from more than 95 percent in the 1970s to less than 90 early in the first decade of the twenty-first century.

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.[10] We are not aware of other papers that have made similar comparisons of scanner data to national spending figures in this way. Our interest is not just in how much 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. The NA figures include expenditures by people living in nonhousing accommodation (student halls,

10. For information on data used to compile NA expenditure figures, see Office for National Statistics (2010a).

old age homes, army barracks, and so on) and spending by tourists in the United Kingdom, which are not included in the surveys. The 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.

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 fifty-two) 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., four weeks, fifty-two 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 four-week periods using the appropriate weights. These are then further aggregated into annual totals by adding up the thirteen four-week periods that generate a fifty-two-week "year," which closely (though not perfectly) covers a single calendar year.[11] We convert aggregated figures to weekly averages.

Figure 16.1 shows average total weekly food and drink expenditure in the NA and the survey data sets between 2002 and 2009, the period for which full-year comparisons can be made. Figure 16.2 reports year-on-year growth rates.

Levels of spending in the LCF are about 79 to 81 percent of those in NA data.[12] Those in the Kantar data are about 51 to 53 percent of NA values. Looking at growth rates, there is much more volatility in the LCF. Aggregate spending on food and drink grew by around 8 percent between 2005 and 2006, a spike not seen in other data. The NA and Kantar figures suggest a sharp slowdown in spending in 2009, but this is not reflected in the LCF data. There is no evidence that the Kantar data perform worse than the LCF when compared to NA data. Indeed, the larger sample size in the Kantar data mitigates the volatility in growth rates observed in the LCF. Such volatility cautions against making year-on-year inferences about changes in living standards from consumption changes in the LCF.

Table 16.1 disaggregates total food expenditures in the surveys and NA into groups based on CPI commodity definitions. Figures are shown for

11. For example, the period labeled 2002 in the Kantar figures covers Jan. 7th 2002—Jan. 5th 2003.

12. 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 percent of those reported in the NA; adding Northern Ireland back in raises this to just 83 percent.

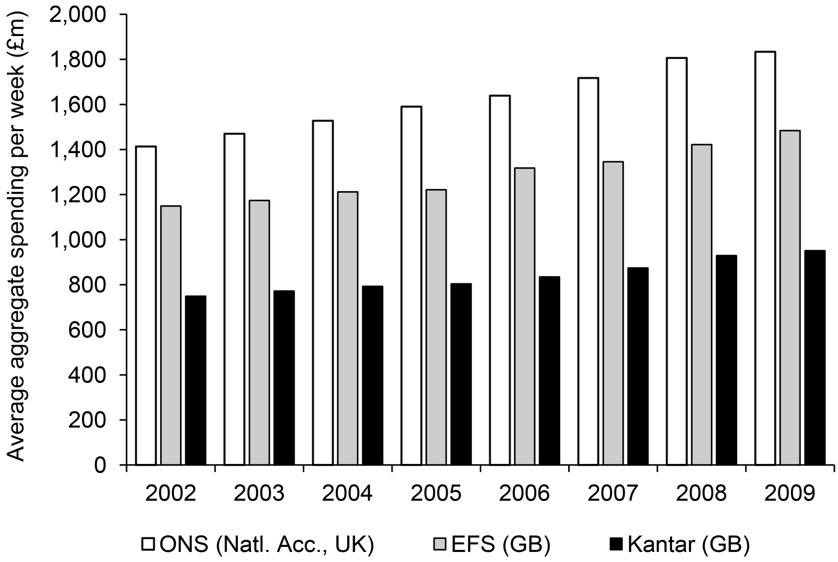

Fig. 16.1 Gross weekly food and drink expenditures, 2002–2009
Source: Calculated from UK Office for National Statistics Data, LCF data, and Kantar Worldpanel.

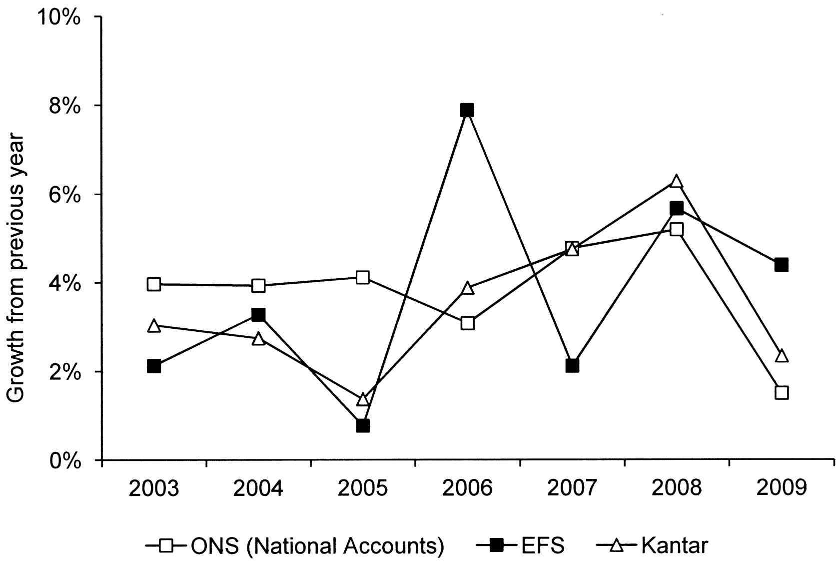

Fig. 16.2 Growth rates of aggregate expenditure, ONS and survey data
Source: Calculated from UK Office for National Statistics Data, LCF data, and Kantar Worldpanel.

Table 16.1 **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. The LCF and Kantar data are for Great Britain, while NA data are for the United Kingdom (including Northern Ireland). The 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.

2002 and 2009, the start and end of our comparison period. Food items are better reported than either alcohol or nonalcoholic beverages in both survey data sets. The relative underreporting of drinks compared to food is striking.[13] Interestingly, the LCF data appear, if anything, relatively worse at recording nonalcoholic drinks than alcoholic drinks, while the reverse is true in the Kantar data. 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) categorization method. Since we map the Kantar products onto LCF codes, these too should then translate directly into comparable NA commodities.

Comparing across the LCF and Kantar data sets, those commodities that 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 percent of NA expenditure (compared to 61 percent for beer and 75 percent for wine). In the Kantar data, it is beer that is least well recorded (29 percent 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, while in principle each household member should record items brought home, in practice it may be that main shopping trips are well reported while 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.

{insert table 16.1 about here}

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 percent 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

13. By 2009 only 54 percent of soft drink expenditure in the NA was captured in the LCF, compared to 75 percent of wine expenditure and 61 percent of beer expenditure. Much attention has been paid to the relative underrecording of alcohol in surveys, but these results show the same is also true of soft drinks.

fruit spending. Again, this is driven mostly by a single year: fruit spending grew by more than 16 percent in 2005 in the NA data, but only by around 7 percent 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.

16.4.2 Comparing Scanner and Budget Survey Expenditure Data

We now make detailed comparisons between the survey measures of spending from Kantar and the LCF. 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 16.5). If we find that expenditure *patterns* are very different in the Worldpanel and LCF, we might be less confident about using scanner data to try and predict detailed expenditures in the survey data given high-level information on total outlays.

When making cross-data set 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 (Ayers, Hossack, and Payne 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, but in general attempts to ensure good compliance are limited. 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 noncompliance. Leicester and Oldfield (2009), for example, look at average weekly expenditures in the Kantar data among 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 percent lower in the Kantar data than in the LCF data in 2005, once observable demographic differences in the samples were taken into account.[14]

14. They find a smaller gap of around 16 percent 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 earlier, which sees recorded spending drop off slightly with the length of participation.

Sample Selection

Our main comparisons between LCF and Kantar data cover calendar year 2009.[15] In the LCF, we exclude all households in Northern Ireland to ensure the geographical coverage of the two data sets 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 "nondropout" sample of 21,093 households (79.1 percent 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 six weeks.[16] 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 nonreporting might be reconciled as holiday periods, for example, but it seems somewhat unlikely that many households would legitimately purchase no food items for that long.

This "regular reporter" sample includes 15,781 households. Leicester (2012) shows that households excluded by this sample selection are significantly more likely to be in London, be headed by someone under forty or over sixty years of age, 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 nonbarcoded 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. Income and employment status have no independent effects on the likelihood of selection into this sample. The fact that our preferred regular reporter sample is clearly a nonrandom set of all Kantar households should be kept in mind when comparing raw expenditures. Later in this section, we condition on observable demographics across the data sets to see how far they can explain spending differences.[17]

Comparisons of Average Expenditure and Budget Shares

Table 16.2 shows average expenditures per week, by CPI commodity, from the LCF and Kantar data during 2009. To strip out the effects of household

15. As in the comparison to national accounts aggregates, we use a fifty-two-week Kantar period, which does not quite overlap with the calendar year, running from December 29th, 2008, until December 27th, 2009.

16. We do not consider whether or not households buy nonfood items during this period.

17. In the Kantar data, we observe household spending over a full-year period. 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.

Table 16.2 Weekly average equivalized expenditure levels by CPI commodity group, LCF, and Kantar (2009)

			Kantar							
	LCFS		Non-dropouts		Regular reporters					
	2 weeks		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
Food	*45.10*	*24.03*	*28.82*	*20.82*	*32.50*	*19.87*	*32.53*	*13.56*	*36.70*	*14.32*
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
Beverages	*3.73*	*3.31*	*2.60*	*2.95*	*2.92*	*2.99*	*2.92*	*1.92*	*3.29*	*2.19*
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
Alcohol	*6.63*	*13.05*	*3.61*	*8.40*	*4.10*	*8.88*	*4.11*	*6.52*	*4.56*	*7.44*
Total spending	**55.46**	**31.28**	**35.03**	**26.03**	**39.52**	**25.03**	**39.56**	**17.18**	**44.56**	**18.29**
No. of households	5,220		21,093.00		15,781		15,781		15,781	
Percent zero weeks	2.7		21.7		13.2		13.2		0.0	
Avg. nonzero weeks/hh	1.95		1.57		1.74		45.13		25.10	

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

composition, expenditures are equivalized using the before housing costs-modified Organisation for Economic Co-operation and Development (OECD) equivalence scale.[18] 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. This approach assumes that weeks in which households record no food spending are accurate reflections of true purchasing behavior. We then make the additional selection described above and drop households with long reporting gaps. Within this sample, we look at average weekly spending when we choose observation periods of different length.[19] Finally, for each Kantar household in this sample, we select the longest consecutive set of weeks containing any recorded expenditure at all. This, in effect, assumes that all weeks in which zero spending is recorded are inaccurate. The mean duration of observation in this sample is 25.1 weeks.

Average household weekly equivalized 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 percent below the LCF figure. Excluding those with long gaps in reporting, average spending rises to £39.52 (29 percent below the LCF figure) when households are observed over a random two weeks or £39.56 observed over the full fifty-two 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, once we ignore zero spending weeks altogether, average spending rises to £44.56, a gap of just under 20 percent compared to LCF levels.

These figures make clear that the treatment of weeks in which zero expenditure is reported is hugely important. Around half of the gap between Kantar and LCF expenditures is eliminated once we strip these weeks out. 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 behavior that 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.

Table 16.3 shows Kantar expenditures relative to LCF expenditures for various sample definitions, and table 16.4 shows the expenditures in terms of budget shares. Even if average expenditures are lower in the scanner data, if the extent of underreporting is quite consistent such that the patterns of expenditure are similar, this acts as a useful validation (of both data sources)

18. See appendix A of Jin et al. (2011) for details of equivalence scales.

19. Results over two and fifty-two weeks are shown; figures for periods of four, twelve, and twenty-six weeks essentially show the same results and are available on request.

Table 16.3 **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
Food	*63.9*	*72.1*	*72.1*	*81.4*
Coffee, tea, and cocoa	71.3	81.7	81.7	91.3
Mineral water, soft drinks	69.0	76.7	76.7	86.8
Beverages	*69.7*	*78.3*	*78.3*	*88.2*
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
Alcohol	*54.4*	*61.8*	*62.0*	*68.8*
Total spending	**63.2**	**71.3**	**71.3**	**80.3**

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

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 16.3, 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. Once again, observing households for two weeks or fifty-two weeks makes little difference to the average proportion. Using the "longest uninterrupted" measure of Kantar spending (right-most column of table 16.3), relative to LCF spending, Kantar expenditure levels match up most closely for nonalcoholic 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 percent as much spending on average for fish, fruit, and other foods than the LCF, but about 90 percent of the expenditure on oils and fats, and coffee and tea. Differences across alcohol types are particularly clear.

Expressed as shares of the total food and drink budget, table 16.4 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

Table 16.4 **Food and drink budget shares, by survey (2009)**

		Kantar			
	LCFS	All	Regular reporters		
	2 weeks (%)	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
Food	*81.3*	*82.3*	*82.2*	*82.2*	*82.4*
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
Beverages	*6.7*	*7.4*	*7.4*	*7.4*	*7.4*
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
Alcohol	*12.0*	*10.3*	*10.4*	*10.4*	*10.2*

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

in budget share is for wine, which makes up on average 5.4 percent of food and drink spending in the Kantar data, but 6.9 percent of spending in the LCF. In the opposite direction, bread and cereals make up 16.1 percent of total Kantar food and drink spending compared to 15.3 percent of LCF spending.

These results are for 2009, but figures from earlier years are not very different. Figure 16.3 shows average equivalized weekly expenditures in the Kantar data (based on the uninterrupted sample definition) as a proportion of LCF values each year between 2002 and 2009. There is a small increase over time. In 2002, Kantar households reported 78.5 percent as much spending as LCF households on average. This rose to 80.3 percent by 2009.

More noticeable is what appears to be a step increase between 2006 and 2007 for beverages, from about 84 percent to 90 percent, and a longer-term upward trend for alcohol beginning in 2006 (though stalling somewhat in 2009). This may be related to the introduction of a new scanner technology for some Kantar households in 2006, who were no longer required to scan nonbarcoded items. If this made compliance costs lower it may have increased reporting of barcoded items—likely to cover almost all spending on alcohol and nonalcoholic drinks—at the cost of reducing reporting of some other categories of

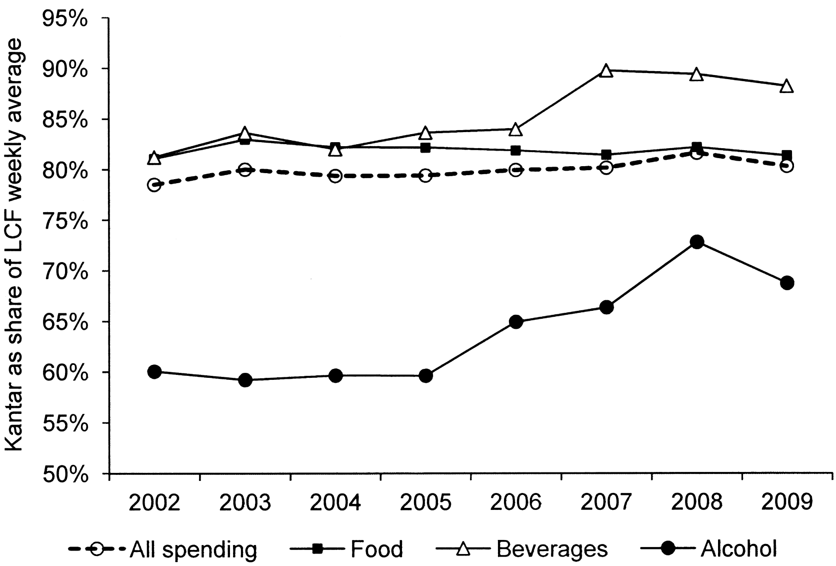

Fig. 16.3 Average Kantar expenditures as a proportion of LCF by year and broad commodity aggregate, 2002–2009

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

Note: Kantar figures relate to "longest uninterrupted" period of Kantar reporting among households that are active across the full calendar year and have no single reporting gap in excess of six weeks.

spending.[20] 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 percent to 80.2 percent of the LCF average, for dairy products (including cheese) from 85.0 percent to 83.1 percent, for fruit from 77.1 percent to 74.3 percent, and for vegetables from 83.7 percent to 81.6 percent.

Table 16.5 uses the 2009 data and makes a direct comparison between households who report nonbarcoded products and those who do not. We focus on the longest uninterrupted spending period. Households who do not record nonbarcoded items spend on average £1.65 per week more than those who do, a statistically significant difference. The largest effects are for beer (28 percent higher spending), soft drinks (22 percent), other food (14 percent), and wine (13 percent). Spending is lower in only two categories: fruit (22 percent lower) and vegetables (3 percent). This leads to quite different expenditure patterns across the groups. 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.

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

Table 16.5 **Weekly expenditure comparisons, LCF and Kantar "uninterrupted" sample 2009, by reporting of nonbarcoded items**

	Average weekly spending				Budget shares						
		Kantar (longest continuous)				Kantar (longest continuous)			Kantar/LCFS		
	LCFS	All	Yes RW	No RW	LCFS (%)	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
Food	*45.10*	*36.70*	*36.50*	*37.01*	*81.3*	*82.4*	*83.1*	*81.3*	*81.4*	*80.9*	*82.1*
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
Beverages	*3.73*	*3.29*	*3.07*	*3.63*	*6.7*	*7.4*	*7.0*	*8.0*	*88.2*	*82.3*	*97.3*
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
Alcohol	*6.63*	*4.56*	*4.33*	*4.90*	*12.0*	*10.2*	*9.9*	*10.8*	*68.8*	*65.3*	*73.9*
Total spending	**55.46**	**44.56**	**43.90**	**45.55**					80.3	79.2	82.1
No. of households	5,220	15,781	9,508	6,273							
Avg. nonzero weeks/hh	1.95	25.10	25.06	25.14							

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

Note: Expenditures are equivalized using the before housing costs-modified OECD scale. "Yes RW" are households who report random-weight (nonbarcoded) products; "No RW" are households who do not.

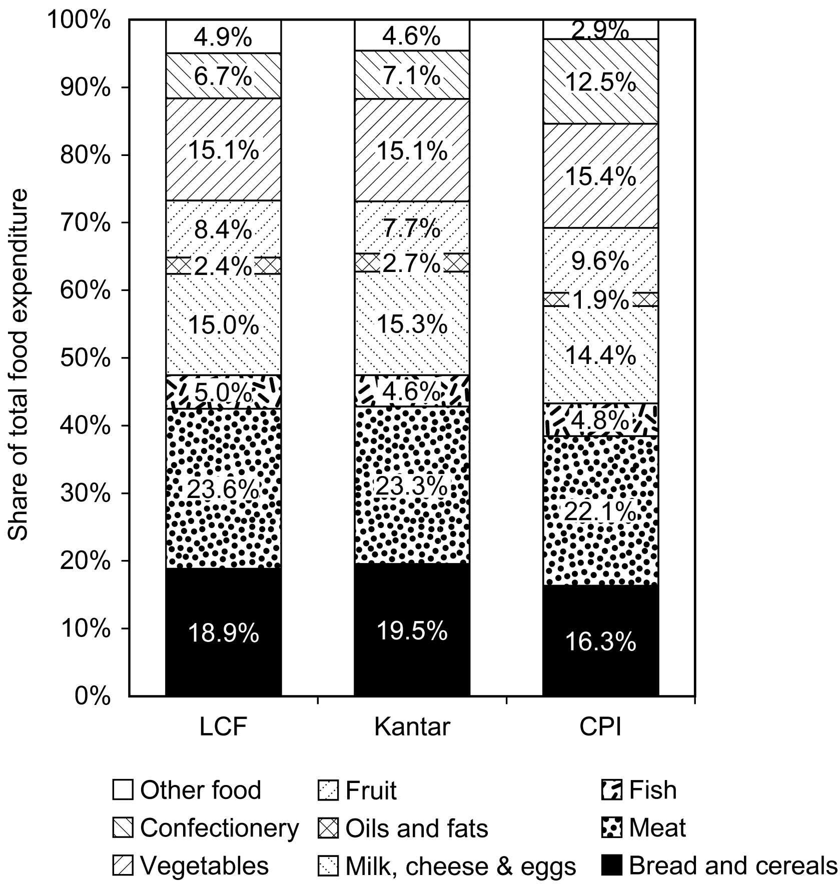

Fig. 16.4 Food commodity weights, LCF, Kantar, and CPI, 2009
Source: Author's calculations based on 2009 Kantar Worldpanel, LCF, and ONS data.

One way to illustrate the economic significance—or otherwise—of differences in spending patterns across data sources 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 16.4 shows the 2009 food budget shares based on Kantar data (using the longest uninterrupted sample for all households), LCF data, and from the CPI expenditure weights.[21] Note that

21. Note that CPI weights are based on national accounts expenditure aggregates, which are in turn based on slightly outdated 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).

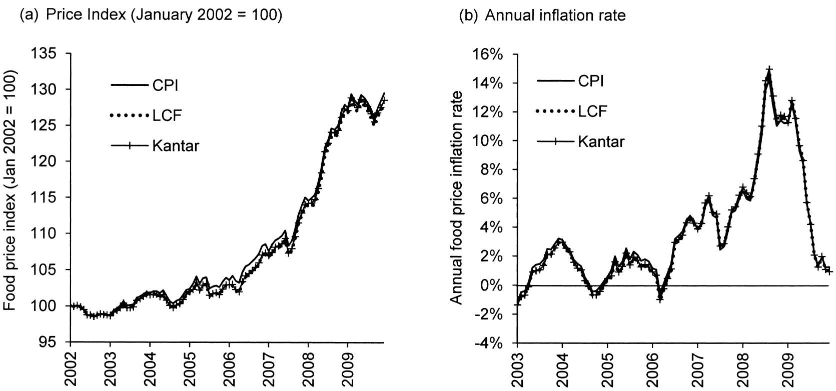

Fig. 16.5 CPI food price indices and inflation rates based on CPI, LCF, and Kantar expenditure weights, 2002–2009

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

we look only at food (a higher-level CPI aggregate) and exclude drinks. The differences between LCF and Kantar data are small: the Kantar basket more heavily weights bread, dairy, fats, and confectionery while 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, fairy, fats, and other food are lower in the CPI than either of the survey baskets, while 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 16.4.1.

Figure 16.5 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.[22] Overall, the effect of reweighting the food CPI using LCF and Kantar-specific commodity weights is small. Over the whole period, the food CPI rose by 29.6 percent, whereas an index based on LCF weights rose by 28.1 percent, and one based on Kantar weights rose by 28.6 percent. The average annual food inflation rate between 2003 and 2009 was 3.8 percent based on CPI weights, and 3.7 percent based on weights from both surveys. The largest absolute gap for any single month between the CPI-weighted inflation rate and the LCF-weighted rate was 0.5 percent. The largest gap

22. Figures were calculated by calculating within-year price indices for each food subgroup 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.

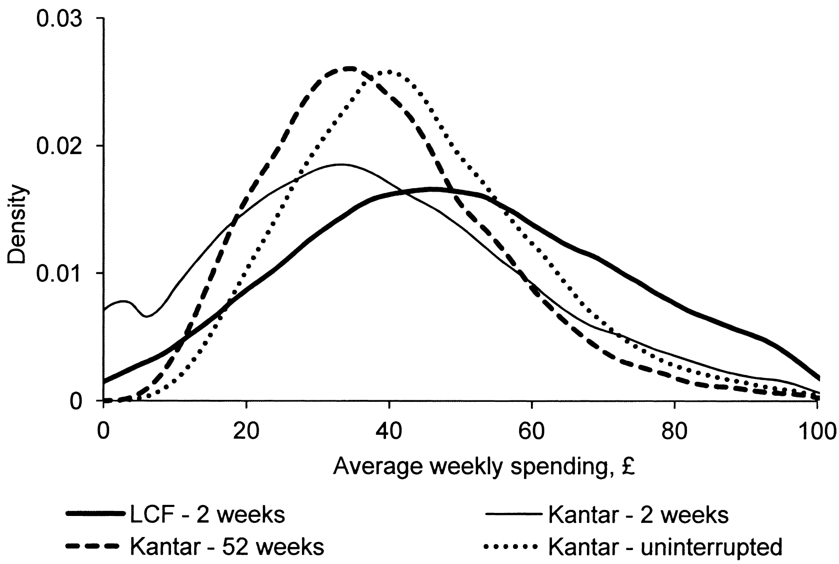

Fig. 16.6 Distribution of equivalized weekly food and drink expenditures, 2009
Source: Author's calculations based on 2009 Kantar Worldpanel and LCF data.

comparing CPI-weighted and Kantar-weighted inflation rates was 0.7 percent. The largest gap comparing LCF-weighted and Kantar-weighted inflation rates was just 0.3 percent.

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. It helps to understand what might be driving differences in the averages. Further, 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 that matters. Figure 16.6 shows a density plot of the distribution of average equivalized weekly food and drink spending in 2009 for the LCF, and regular reporter Kantar households observed over two weeks, fifty-two weeks, and for the longest uninterrupted number of weeks.

The distribution for LCF households is relatively smooth over the range of spending shown. However, 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 toward lower expenditures than the LCF figures. There are notably far fewer high-spending households in the Kantar data than the LCF.

Even more interesting is the impact of the 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 stylized example, imagine that all households shop once a week and consume one can of beer per week. Beer is only sold in four-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.

Leicester (2012) takes the regular reporter Kantar sample and shows that, observed over a full year, almost all households are observed to purchase from all CPI commodity groups whereas observed over two weeks, large numbers report zero expenditure. The exception is alcohol—over a year, around 39 percent of households do not buy spirits, 32 percent do not buy beer, and 16 percent do not buy wine. However, over just two weeks the proportions not observed to buy are 89 percent, 84 percent, and 71 percent, respectively. 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.

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 16.7 illustrates this for four different commodities (clockwise from top left for bread and cereals, meat, vegetables, and fish; figures for other goods available on request). This is crucial for some applications. For example, past research in the United Kingdom has used data from the two-week LCF to estimate household-specific expenditure patterns from which household-level inflation rates have been estimated (Leicester, O'Dea, and Oldfield 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.

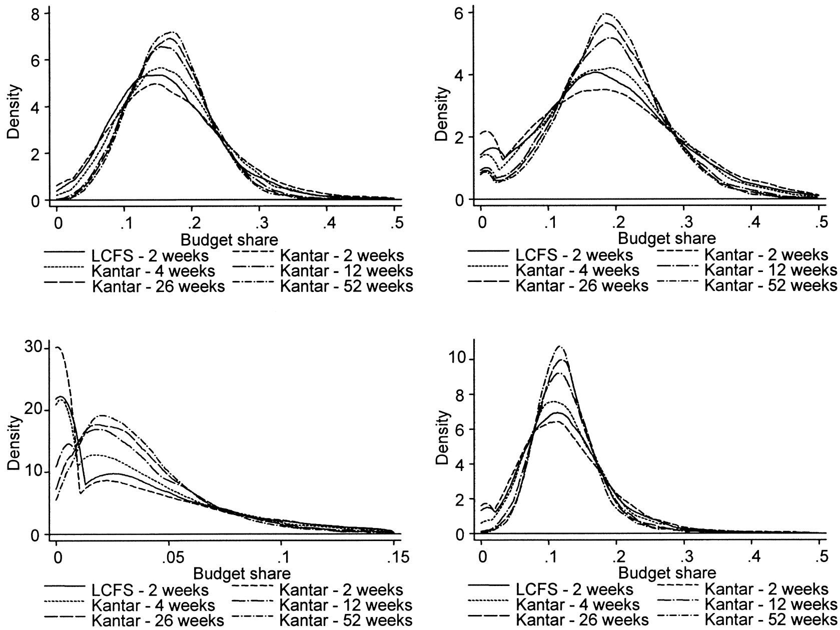

Fig. 16.7 Distribution of budget shares for particular commodity groups based on period of observation, 2009

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 nonrandom 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 that would typically make them low spenders on food at home (for example, being poorer), then this could account for the spending gap between the data sets.

Appendix A compares the observable characteristics of the LCF and Kantar (regular reporter) samples in 2009.[23] There are relatively few demographics common to both data sets. To the extent that we can strip out the effects of these common demographics, there may well be a large set of

23. 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 equivalized average weekly spending on food and drink (over two weeks in the LCF sample, and over the longest uninterrupted reporting period in the Kantar sample) of less than £5.

unobserved demographics or variables observed in one data set but not the other, which affect expenditures but on which we cannot condition. To summarize the key differences, we find that:

- Kantar households tend to have lower income. Among those with nonmissing incomes, only 12.5 percent (3.7 percent) of Kantar households have gross annual household income in excess of £50,000 (£70,000) compared to 20.8 percent (9.9 percent) of LCF households.
- Kantar households are much more likely to own a home computer: 89.5 percent of the Kantar sample do so compared to 75.6 percent of the LCF sample. This probably reflects the fact that many Kantar households also participate in online surveys run by Kantar, and many use personal computers (PCs) to upload their expenditure records to Kantar from the scanner units.
- Only 12.9 percent of Kantar households do not own a car, compared to 21.6 percent in the LCF.
- The regional composition of the two data sets is similar. In the Kantar sample the southeast 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 percent of the Kantar sample is headed by someone age eighty or over compared to 7.1 percent of the LCF sample. Less than 1 percent of Kantar households are headed by someone age twenty-four or under compared to 3.0 percent of the LCF sample.
- Fifty-six percent of Kantar households are headed by a female, compared to 25 percent of the LCF sample.
- Kantar households are much more likely to be headed by someone who is not working (11.3 percent compared to 8.6 percent in the LCF) or part-time employed (15.8 percent working less than thirty hours, compared to just 6.1 percent in the LCF).
- Kantar households are much more likely to contain three or more adult members. This accords with information received from discussions with Kantar that they oversample multiple adult households because of difficulties in obtaining purchase information from secondary shoppers.
- Kantar households have about the same number of young (preschool age) children as LCF households on average, but slightly older children.

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 life cycle 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.[24] 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 observable demographic effects and see what happens to the coefficient on the Kantar dummy. Table 16.6 reports the key results, 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 percent—this is markedly less than the 20 percent seen earlier in this section, but our analysis here is based on a different (unequvalized) measure of spending and drops households with very low average food spending of less than £5 per week. Those who are asked to report random weight purchases spend 12 percent less than LCF households, while those who are not asked to do so have a raw spending difference of less than 5 percent. As above, this suggests that reducing respondent burden increases expenditures in scanner data.

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 percent. This is consistent with Leicester and Oldfield (2009), who found the gap between Kantar and LCF spending in 2005 rose from 20 percent to 25 percent 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 among the groups asked or not asked to report random weight items. This suggests that much of the seeming improvement in reported spending is attributable to demographics.

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.[25] Table 16.7 shows the main significant interaction terms. The base group is households in the southeast of England, with incomes between

24. As we now control for demographics, including household composition, here we use unequivalized expenditure figures.

25. This follows the approach of Zhen et al. (2009), equation (1), and Leicester and Oldfield (2009), table 9. Zhen and colleagues 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" (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 twenty-five with income under $5,000 of "other race" living in the northeast 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.

Table 16.6 **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***	–0.123***	–0.047***
	(0.011)	(0.011)	(0.012)
R^2	0.005	0.010	0.002
Controlling for common observed demographics	–0.182***	–0.179***	–0.179***
	(0.009)	(0.010)	(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 unequivalized. Households with missing information on number of cars or employment status are dropped, as are those spending less than £5 (equivalized) 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 birth to four years, five to ten years, eleven to seventeen years, and numbers of people age sixty-five or over. Standard errors are robust. Full results are available on request.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

Table 16.7 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
		Head 60–64	–0.087**
0 adult males	–0.095*	Head 65–69	–0.093
2 adult males	+0.019	Head 70–74	–0.070
3+ adult males	–0.205***	Head 75–79	–0.146*
		Head 80+	–0.121
0 adult females	+0.100***		
2 adult females	–0.052	Female head	+0.106**
3+ adult females	+0.022		
1 child age 0–4	–0.127***	1 person age 65+	+0.080
2 children age 0–4	–0.033	2 people age 65+	+0.137**
3+ children age 0–4	–0.330***	3+ people age 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 (unequivalized). Other variables controlled for are region, numbers of children age five to ten and eleven to seventeen, head of household employment status, number of cars and presence of a PC in the household. Households with missing information on number of cars or employment status are dropped, as are those spending less than £5 (equivalized) per week on average in either survey.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

£10,000 and £20,000 per year, with one car and a home computer, where the head is a male age forty-five to forty-nine and full-time employed, where there is one adult male and female but no children or anyone over sixty-five years of age. For this group, average spending is 9 percent lower per week in the Kantar data than the LCF. The coefficients in table 16.7 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.[26]

26. Region is not shown. There is one significant coefficient of –0.087 in the East Midlands. Full results are available on request.

A striking finding is that lower-income households report relatively higher spending in the Kantar data, while 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 underreporting. 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.[27]

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.[28]

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 age sixty-five 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.

27. 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.

28. 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 that was not observed in their estimates.

A brief comparison to the results of a similar exercise in Zhen et al. (2009) is worthwhile. They compare AC Nielsen and CE commodity-level expenditures, regressing spending on demographic variables interacted with a Nielsen dummy. They find a number of demographic effects on relative expenditures, which are different from 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 16.3: findings from one particular comparison of scanner and other data sets 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 data set 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 that 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 underrecording their spending on each broad commodity group at roughly the same rate, on average. Alternatively, the Kantar 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, one approach is to estimate Engel curves for each food group in the two data sets. Engel curves relate the share of total food expenditure devoted to each food commodity to the log of total food spending. For the purposes of how scanner data may be useful for statistical agencies, as discussed earlier (and detailed in section 16.5), 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 making this kind of imputation. If not, it may shed more light on where particular problems in measuring expenditures arise.

This exercise is carried out in Leicester (2012) (see, in particular, appendix C). He imposes no particular restriction on the shape of the Engel curves, but finds them to be approximately linear for most food groups. Most importantly, he concludes that the relationship between food expenditures and commodity-level budget shares is very similar across data sets. 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, and other food and beverages, the Engel curves are broadly flat. Perhaps the biggest divergence in the Engel relationship between data sets 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 data set suggests a commodity to be a luxury and the other data set a necessity. These findings are quite reassuring that the scanner and budget survey data sets tell similar stories about how spending patterns change with total food outlays.

16.5 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.

We now assess the potential for using scanner data to impute detailed expenditures given only knowledge of total spending. A redesigned CE could decide to reduce respondent burden by asking people only about their total expenditures (or total category-level expenditures such as food spending), and then use detailed expenditure records to estimate how this breaks down. These records could come from a subset of households who agree to provide a more detailed account of their expenditures (whether using scanners, till receipts, diaries, or whatever), or from external data sets like commercial scanner data. This section aims to provide evidence on how successful such an approach might be. We use scanner data to predict household-level budget shares for each commodity as a function of total expenditure and observable demographics. How well these predictions compare to actual budget shares 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. This is useful as spending patterns appear to vary by store, so performing this exercise separately for different store types may improve the results.

We take an agnostic view on how the information on total expenditures is obtained. One approach would be a series of questions that first ask about how much households typically spend in total over a given period and then how this breaks down into broad commodity aggregates like food, clothing, leisure, and so on. Further questions could be asked about the proportion of food spending in different store types, such as supermarkets, corner shops, and specialist food stores. Another approach would be to use individual bank records and credit card statements to get total store-specific expenditures. This becomes more feasible as cash spending, which cannot be attributed to a particular store, declines in importance.[29]

We base our analysis on the regular reporter sample in the 2009 Kantar data (see section 16.4.2), 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 nonalcoholic 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 eleven 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 percent of expenditure in our data)—and we look at each of these stores separately.[30] Other supermarket chains are grouped according to quality segment: discount supermarkets (Aldi, Lidl, Netto, Cash and Carry stores), quality supermarkets (Waitrose, Marks & Spencer), and other supermarkets (Co-op, Somerfield, Iceland, 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. These account for a low share of spending, suggesting there is little to be gained from disaggregating them still further. Table 16.8 shows the eight store types considered and their market shares by expenditure. Figure 16.8 shows how total food and drink expenditure breaks down for each store across the eleven commodities we consider.[31]

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, while dairy and meat products account for smaller shares. In quality supermarkets, fruit, vegetables, meat,

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

30. 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.

31. These estimates are based on the full 2009 Kantar sample, rather than the selected "good reporter" sample.

Table 16.8 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.

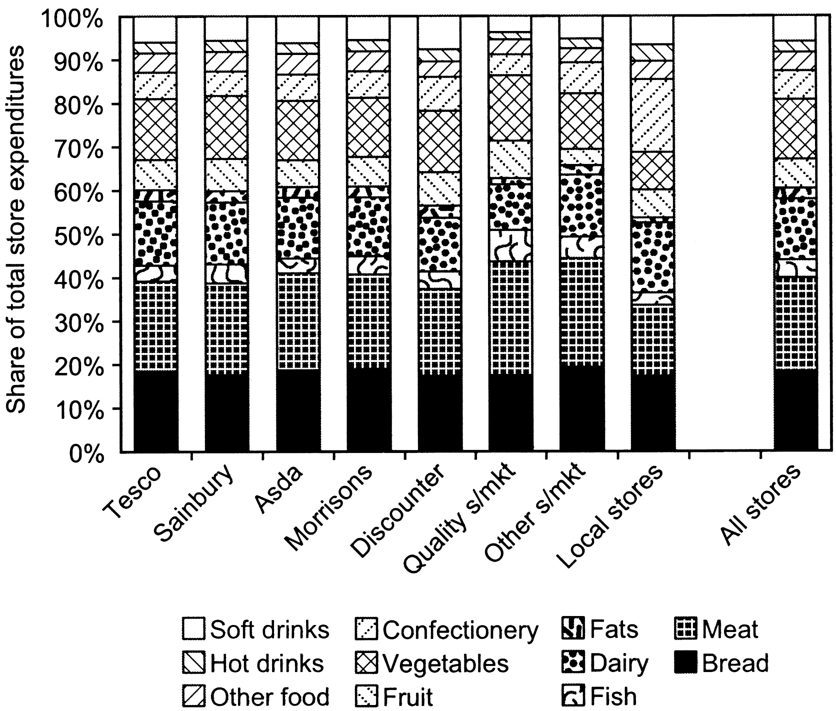

Fig. 16.8 CPI commodity budget shares by store type, 2009
Source: Author's calculations based on 2009 Kantar Worldpanel data.

and fish 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 percent of expenditure, but in local stores the share is 16.7 percent. This probably reflects the different sorts of shopping done in local stores that are more associated with top-up and impulse purchases.

If households systematically underreport different types of expenditure in scanner data, imputed budget shares would be biased. A useful check would be to compare expenditure patterns at the store level obtained from store data to those seen in the in-home scanner data. The former would presumably be complete, accurate records of store-specific sales patterns. We do not have access to such data and are not aware of any comparisons having been made of this nature. It would, though, be a useful line of inquiry for statistical agencies thinking of using in-home scanner data.

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 eleven 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 ordinary least squares (OLS) regression methods to estimate a system of equations for store-specific budget shares for each commodity:

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

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 $\boldsymbol{Z}_i$ is a vector of observable demographic household characteristics.[32] 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 ten of the eleven 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 16.9 reports

32. The model includes gross household income, region, household composition (numbers of adult males, females, children in different age groups, and people age sixty-five and older), 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 nonlinear, 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 among 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 that 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.

Table 16.9 **Explanatory power (R^2) of store-specific commodity budget share model**

	Tesco	Sainsbury	Asda	Morrisons	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.

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 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.

Appendix B shows the distribution of actual and predicted household-level budget shares resulting from this modeling exercise. Broadly, this approach captures the average shares quite well, but not the distribution: the modeled 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 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. We could also use more sophisticated imputation and econometric methods to predict the budget shares. For example, a particular problem with the OLS approach is that households do not buy from each commodity group at each store type they visit. This is particularly problematic once we condition on store type: 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 percent of shoppers never buy from the hot beverage category, more than 25 percent buy no fish, and more than 20 percent 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 percent of households are predicted to have at least one negative budget share (these households are dropped from the results in appendix B). 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.

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 toward that goal. While 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 modeling 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.

16.6 Summary and Conclusions

In-home scanner data offer a hugely exciting opportunity for researchers to explore detailed questions about household expenditure behavior and firm-pricing decisions. Scanner data sets 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. Nevertheless, in terms of the potential for using existing commercial scanner data sets, organizations 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 nondurable 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. The evidence in this chapter 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 nonbarcoded products, and that there are differences in the relative underreporting across household groups (in particular relating to income and numbers of children, which plausibly reflect time constraints). These differences point to modal effects. However, in some cases, differences between scanner and budget survey expenditures, which might in isolation be attributed to survey mode may, in fact, be specific to the particular data sets 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 cooperate effectively, and so on. In some countries, notably Nordic ones, detailed population-level data linked to identification records is maintained that 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. Evidence from British data suggests that 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 reweighting the food CPI using survey-specific budget shares has very little impact on inflation rates. The relationship between total expenditures and commodity-level budget shares is 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, while 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 data sets are limited in coverage to food at home and a small number of nonfood 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 data set collected by a statistical agency would also enable experimental analysis of the impact of survey mode, different scanner devices, different reporting requirements (e.g., nonbarcoded items), different incentives for participation, and so on. By providing "gold standard" evidence of these issues using randomized 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 chapter and previous studies, on other ways in which new scanner data could improve on existing commercial data sets. First, one of the major limitations of the Kantar Worldpanel relative to the LCF (and other data sets 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. While 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 data sets in general, but we are not aware of research that 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 analyzing how detailed food-spending behavior 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.

Second, a key unresolved problem with scanner data is 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 data set. 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 nongenuine ones where trips were made but not reported for whatever reason. 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.

Third, 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 nonbarcoded purchases. In the Kantar data, this was justified by better reporting compliance for nonbarcoded items. However, our analysis suggests that this result is at least partly driven by demographic differences between those who do and do not report nonbarcoded purchases. It may also be differences in the type of scanner device, rather than the lack of random weight reporting that 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 minimize 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 centralized store-level databases given the growth of personalized 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.

Appendix A
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 section 16.4.2). Respective sample sizes are 5,150 and 15,752 households.

Table 16A.1 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

Table 16A.2 **Computer in the household?**

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

Table 16A.3 **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

Table 16A.4 **Region**

	Kantar (%)	LCF (%)	K ÷ LCF
Northeast	4.8	4.5	1.07
Northwest	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
Southeast	15.1	13.4	1.12
Southwest	9.5	10.0	0.95
Wales	5.1	5.2	0.97
Scotland	8.2	10.4	0.79
Northeast	4.8	4.5	1.07

Table 16A.5 **Gender of household head**

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

Table 16A.6 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

Table 16A.7 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

Table 16A.8 Number of adults (age eighteen and older)

	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

Table 16A.9 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

Table 16A.10 Number of people age sixty-five and older

	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 B

Actual and Predicted Household-Level Budget Shares by Commodity, 2009

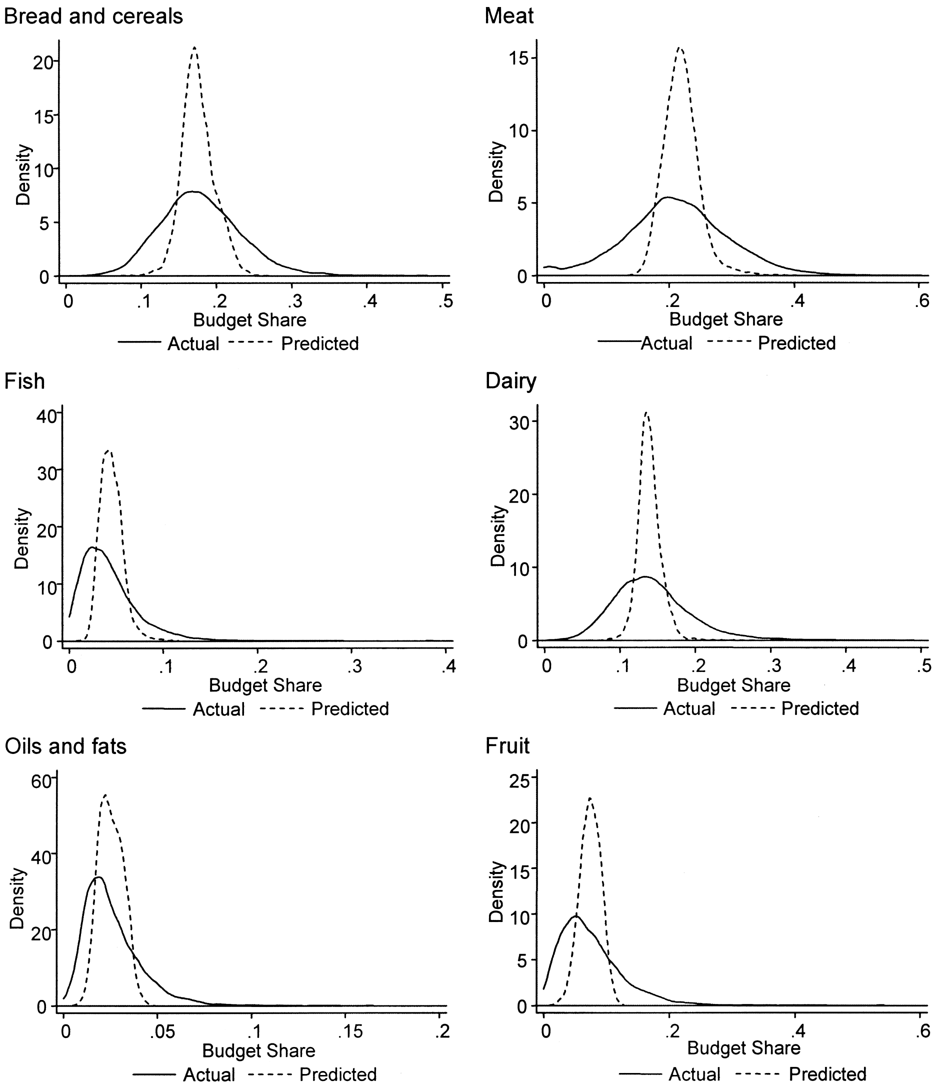

Fig. 16B.1 Actual and predicted household-level budget shares by commodity, 2009

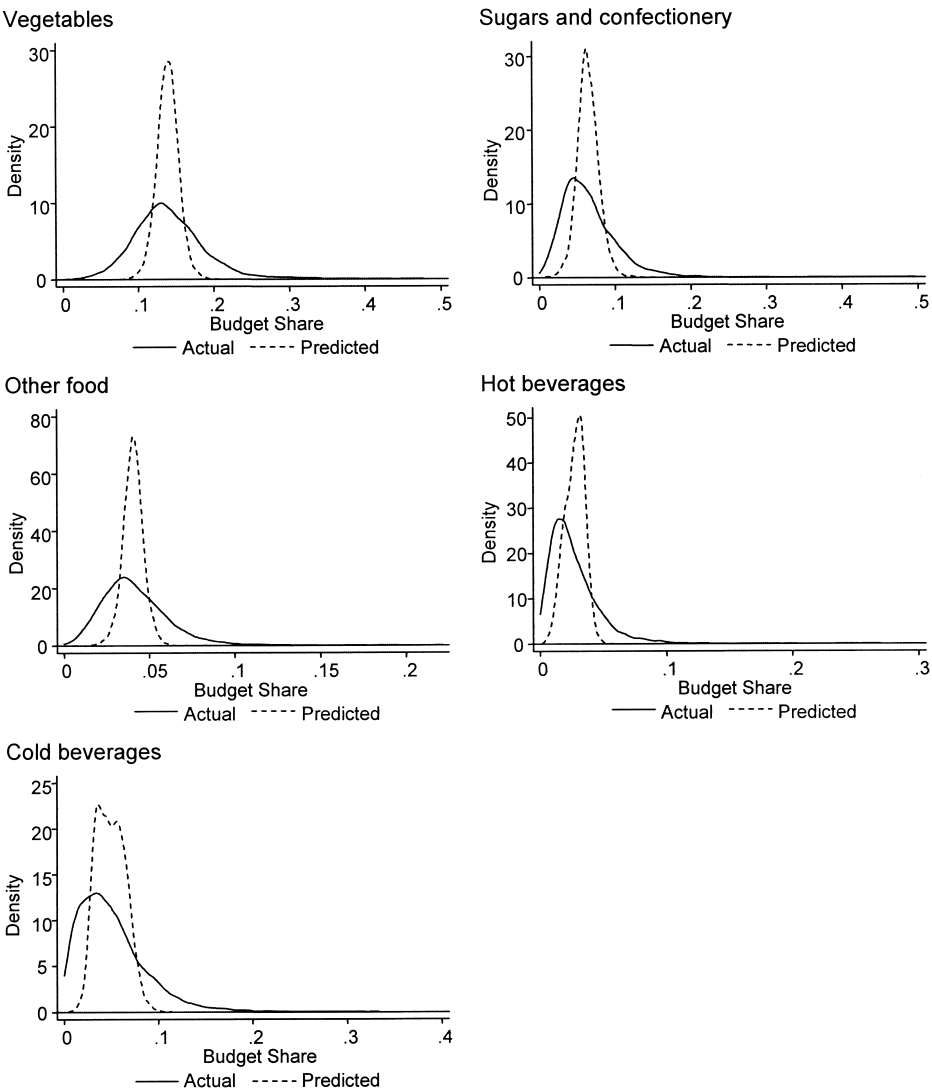

Fig. 16B.1 (cont.)

References

Aguiar, M., and E. Hurst. 2007. "Life-Cycle Prices and Production." *American Economic Review* 97:1533–59.

Ahmed, N., M. Brzozowski, and T. Crossley. 2006. "Measurement Errors in Recall Food Consumption Data." IFS Working Paper no. 06/21, Institute for Fiscal Studies. http://www.ifs.org.uk/publications/3752.

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." Conference Paper, National Poverty Center, University of Michigan. 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*. Newport, South Wales: Office for National Statistics.

Blow, L., A. Leicester, and Z. Smith. 2004. "Consumption Trends in the UK, 1975–99." IFS Report no. 65, Institute for Fiscal Studies. 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, England: 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.

Department for Environment, Food & Rural Affairs (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 manuscript, Bureau of Labor Statistics.

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 no. 08–07, Stanford University. http://www-siepr.stanford.edu/repec/sip/08-007.pdf.

Feenstra, R. C., and M. D. Shapiro, eds. 2003. *Scanner Data and Price Indexes*. Chicago: University of Chicago Press.

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): 339–65.

Harding, M., E. Leibtag, and M. Lovenheim. 2012. "The Heterogeneous Geographic and Socioeconomic Incidence of Cigarette Taxes: Evidence from Nielsen Homescan Data." *American Economic Journal: Economic Policy* 4 (4): 169–98.

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/sp04/hu05.pdf.

Jin, W., R. Joyce, D. Philips, and L. Sibieta. 2011. "Poverty and Inequality in the UK: 2011." IFS Commentary no. 118, Institute for Fiscal Studies. http://www.ifs.org.uk/comms/comm118.pdf.

Leicester, A. 2012. "How Might In-Home Scanner Technology Be Used in Budget Surveys?" IFS Working Paper no. 12/01, Institute for Fiscal Studies. http://www.ifs.org.uk/wps/wp1201.pdf.

Leicester, A., C. O'Dea, and Z. Oldfield. 2008. "The Inflation Experience of Older Households." IFS Commentary no. 106, Institute for Fiscal Studies. 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, Institute for Fiscal Studies. http://www.ifs.org.uk/publications/4468.

Levell, P., and Z. Oldfield. 2011. "The Spending Patterns and Inflation Experience of Low-Income Households over the Past Decade." IFS Commentary no. 119, Institute for Fiscal Studies. 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): 226–40.

Mathiowetz, N., K. Olson, and C. Kennedy. 2011. "Redesign Options for the Consumer Expenditure Survey." Unpublished manuscript, National Academy of Sciences. 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 no. 71, Institute for Fiscal Studies. http://www.ifs.org.uk/comms/r71.pdf.

Office for National Statistics. 2010a. *Consumer Trends, Q2 2010*. Newport, South Wales: ONS. http://www.ons.gov.uk/ons/rel/consumer-trends/consumer-trends/q2–2010/consumer-trends-no—57.pdf.

———. 2010b. *Consumer Price Indices Technical Manual: 2010 Edition*. Newport, South Wales: 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): 1282–87.

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 *How Reliable is the Family Expenditure Survey? Trends in Incomes and Expenditures over Time*, edited by J. Banks and P. Johnson. London: Institute for Fiscal Studies. http://www.ifs.org.uk/comms/r57.pdf.

Triplett, J. 1997. "Measuring Consumption: The Post-1973 Slowdown and the Research Issues." Federal Reserve Bank of St. Louis Review (May/June):9–42. 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 Bureau of Labor Statistics 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.

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

Westat. 2011. *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): 470–92.