Office for National Statistics

{77}

Survey Methodology Bulletin

Summer 2017

Contents

Office for National Statistics
Swyddfa Ystadegau Gwladol

The Survey Methodology Bulletin is primarily produced to inform staff in the Office for National Statistics (ONS) and the wider Government Statistical Service (GSS) about ONS survey methodology work. It is produced by ONS, and ONS staff are encouraged to write short articles about methodological projects or issues of general interest. Articles in the bulletin are not professionally refereed, as this would considerably increase the time and effort to produce the bulletin; they are working papers and should be viewed as such.

The bulletin is published twice a year and is available as a download only from the ONS website.

The mission of ONS is to improve understanding of life in the United Kingdom and enable informed decisions through trusted, relevant, and independent statistics and analysis. On 1 April 2008, under the legislative requirements of the 2007 Statistics and Registration Service Act, ONS became the executive office of the UK Statistics Authority. The Authority's objective is to promote and safeguard the production and publication of official statistics that serve the public good and, in doing so, will promote and safeguard (1) the quality of official statistics, (2) good practice in relation to official statistics, and (3) the comprehensiveness of official statistics. The National Statistician is the principal advisor on these matters.

www.ons.gov.uk

Edited by: Philip Lowthian

methodology@ons.gov.uk

The Effect of Variance in the Weights on the CPI and RPI

Jim O'Donoghue[1]

1. Introduction

The consumer prices index (CPI) is one of the most important economic statistics produced by the Office for National Statistics. It is the target inflation measure for the Bank of England's Monetary Policy Committee; it is used for deflation of National Accounts aggregates and for price adjustment for a range of benefits and thresholds[2].

Despite its importance, there are no formal estimates of the precision of the CPI. This is due in large part to the complexity of the sample design. It is generally assumed that the 12-month inflation rate is accurate to +/-0.1 percentage points. Research has been undertaken in the past on this subject, but produced results that were much higher than this and were judged to be implausible. Nor has it proved possible to derive estimates for the precision of the Retail Prices Index: the RPI is no longer a National Statistic but it is still used for indexation of many government bonds, for price adjustment in commercial contracts and for up-rating of many pensions[3].

This article reports on research to produce estimates of the effect of variance in the weights used to compile the CPI and RPI and for selected sub-groups of the population. It originated in the need to produce evidence to feed into the National Statistics Quality Review of the Living Costs and Food Survey (LCF)[4]. The LCF is a household survey run continuously by the ONS; one of its key uses is the derivation of weights for the RPI and, indirectly, the CPI. The initial research was centred on providing estimates of the effect of random error arising from the LCF on the precision of the RPI, and investigating the impact on precision of cuts in the LCF sample size.

This work was later extended to investigate the impact of variability of the weights on the precision of subsets of the population. This follows on from the Johnson Review of Consumer Price Statistics[5] which recommended that ONS should reduce the range of indices it publishes (by ceasing publication of RPI-J and reducing the use of RPI) but suggested publishing more detailed breakdowns of the existing ones to provide more insight around inflation as experienced by different household groups. The household

[1] Jim.O'Donoghue@ons.gov.uk

[2] Consumer Price Indices: A brief guide, ONS, 2013, https://www.ons.gov.uk/ons/guide-method/user-guidance/prices/cpi-and-rpi/consumer-price-indices--a-brief-guide.pdf

[3] Users and uses of consumer price inflation statistics, ONS, 2013, https://www.ons.gov.uk/ons/guide-method/user-guidance/prices/cpi-and-rpi/users-and-uses-of-the-consumer-price-inflation-statistics.pdf

[4] National Statistics Quality Review Series 2 – Living Costs and Food Survey, ONS, 2016, https://www.ons.gov.uk/file?uri=/peoplepopulationandcommunity/personalandhouseholdfinances/incomeandwealth/methodologies/nsqrseries2reportnumber3livingcostsandfoodsurvey/lcfnsqrreport.pdf

[5] UK Consumer Price Statistics – A Review, UKSA, 2015, https://www.statisticsauthority.gov.uk/reports-and-correspondence/reviews/uk-consumer-price-statistics-a-review/

groups used for this analysis were those published in the article by Flower & Wales[6]: expenditure decile, income decile, pensioner/non-pensioner household, and households with and without children.

2. Background

2.1 Sample design

2.1.1 Representative items

The CPI covers the whole of household expenditure on goods and services and price indices are calculated and published for each category in the international classification of consumption – COICOP, the Classification of Individual Consumption by Purpose[7]. In practice, it is not feasible to price every single different product in each COICOP class, so expenditure is divided into relatively homogeneous groups of products (this can be viewed as a form of stratification) and representative item(s) is/are selected purposively to represent the range of characteristics for products in each group. In total there are nearly 700 items.

One or more items may be selected to represent a particular product group. In some instances, the representative item includes all expenditure covered by the product category. Examples are electricity consumption, air fares, national daily newspapers, car insurance. These are typically "central items" where the price data are collected and compiled centrally by ONS staff. Generally, they are weighted, either fully or partially, with component sub-indices representing different categories of consumption within that item. For instance, electricity consumption is weighted by supplier, area of the country and type of tariff.

For other product groupings, items are chosen to represent the different types of products covered by the group. For instance, the COICOP class fish covers frozen and fresh products; processed and non-processed products; and white and non-white fish. The items are chosen purposively (i.e. they are not randomly selected) and are generally those with the greatest expenditure. So, this might be viewed as a form of cut-off sample.

01.1.3 Fish	
Fresh white fish fillets	Fresh salmon fillets
Canned tuna	Fish fingers
Frozen prawns	Frozen breaded/battered white fish

For some product groupings, some locally collected items may represent themselves. An example is fruit, where bananas, oranges and apples carry the weight associated with expenditure on those items. However, fruit also includes stone fruit (for example,

[6] Variation in the Inflation Experience of UK Households 2003 – 2014, ONS, 2014, http://www.ons.gov.uk/peoplepopulationandcommunity/personalandhouseholdfinances/expenditure/compendium/variationintheinflationexperienceofukhouseholds/2014-12-15

[7] COICOP, Detailed Classification and Notes, United Nations Statistics Division, http://unstats.un.org/unsd/cr/registry/regcst.asp?Cl=5

peaches, plums, nectarines), other citrus fruit (for example, lemons, limes, grapefruit) and other fruit (for example, pineapples, kiwi fruit). Representative items are chosen to represent each of these sub-categories; each will carry the weight for the whole sub-category. If more than one item is chosen to represent a particular sub-category, the weight is frequently divided equally between them. An example is men's clothing.

2.1.2 Outlet selection

Outlets (shops) are selected randomly. In the first stage, the UK is stratified by region with the number of locations in a region being proportional to regional expenditure share[8]. Locations are selected in each region with probability proportional to an expenditure proxy. Within each location, shops are enumerated, grouped by item category (this is a set of items that are typically found together in a particular type of shop; this is not necessarily all the items in a specific COICOP class) and usually one shop chosen to represent each item category, with probability proportional to floor space frequently used as an expenditure proxy. A shop may be selected for more than one item category.

The largest chains of shops (for example, Tesco, Marks & Spencer) are sampled with certainty, with the number of desired price quotes for a particular shop proportional to its overall market share (for the particular category of product where available). In practise, it is unusual to collect from more than one shop per region for a particular chain, so the prices for regional centrals are given shop weights in line with their market share. These are called *regional centrals*. In addition, some major retailers with central pricing such as Argos, Ikea, Next, Great Universal Stores and on-line supermarkets are sampled with certainty. These are called *central shops.*

In each selected shop, the specific products whose prices are tracked are selected purposively by the price collector. Usually, one product is selected in each shop for each representative item. The selected products tend to be those that are most sold within the shop, as determined by shelf space or as advised by the retailer.

The selection processes for outlets and items are illustrated in Annex A.

2.2 Method of index calculation

2.2.1 Index calculation

The CPI and RPI are annually chain-linked Laspeyres indices calculated as the arithmetically (expenditure) weighted average of item indices[9]. In calculating the CPI and RPI, items are stratified by region, shop type (multiple or independent), region and shop type, or not at all. There are no hard and fast rules as to which form of stratification is used, although items where the relative of averages (Dutot) formula is used for the RPI tend to be stratified by region or region multiplied by shop type. The type of stratification used for a particular item rarely changes between years.

[8] The method described for the selection of locations is the one that was introduced in 2000. A new method is being piloted in 2016 and 2017, for possible roll-out in future years.

[9] A Practical Introduction to Index Numbers, Ralph J, O'Neill R and Winton J, 2015, Wiley, http://eu.wiley.com/WileyCDA/WileyTitle/productCd-1118977815.html

Item indices are generally built up from stratum indices. This is achieved by combining together all price quotes that correspond to a particular stratum into a single elementary aggregate index using the appropriate formula (geometric mean for the CPI and arithmetic mean for the RPI). For instance, clothing items are generally stratified by shop type (multiple and independents). The multiples stratum index for a man's suit is calculated from prices collected locally in each location and also from central shops and regional centrals. Regional central shops are given weights in line with their overall market share and the geographical spread of its shops. Prices from central shops are given shop weights in line with their overall market share. These prices with weights are combined with other prices with unit weights to give the elementary aggregate (stratum) indices. These in turn are weighted together to give the item indices.

So for the CPI, the index $I_t^{i,s}$ in month t for stratum s for item i is calculated as follows:

$$I_t^{i,s} = exp\left(\frac{\sum_l ln\left(\frac{p_{t,l}^s}{p_{o,l}^s}\right) + \sum_c w_c^s ln\left(\frac{p_{t,c}^s}{p_{o,c}^s}\right)}{n_l^s + \sum_c w_c^s}\right)$$

Where *c* represents central and regional central shops; *l* represents other (local) outlets; n_l^s is the number of locally collected prices; and w_c^s the weight of central shop *c* in stratum *s*.

The item index is then:

$$I_t^i = \frac{\sum_s w_s \ I_t^{i,s}}{\sum_s w_s}$$

The representative items used in the calculation of the CPI and RPI are reviewed annually and changed, as necessary, to keep pace with changes in household spending habits. For the CPI, the item indices are aggregated together to form indices classified according to the international standard classification of individual consumption by purpose (COICOP), while for the RPI they are combined together to form RPI section indices.

2.2.2 RPI weights

Weights for the 84 RPI sections are calculated directly from the LCF[10]. Wealthy households (those in the top 4% of households by income) are excluded, as are poor pensioner households (those dependent on the state for at least three-quarters of their income). There are a few expenditure categories where LCF data are modified before being used as the section weights. Principal among these are: alcohol and tobacco

[10] Consumer Price Inflation – 2016 Weights, ONS, 2016, http://www.ons.gov.uk/file?uri=/economy/inflationandpriceindices/articles/consumerpricesindexandretailpricesindexupdatingweights/2016/cpi2016weightsarticle.pdf

whose weights are adjusted to allow for under-reporting of expenditure, and housing depreciation which is an imputed expenditure intended to represent major maintenance and repair of the property by owner-occupiers. The weights for year y are based on expenditure data collected in the LCF for the 12-month period running from July y-2 to June y-1, and section level expenditure is up-rated by movements in the section index from January y-1 to January y, before the weights are calculated.

The weights for RPI item indices are mainly based on LCF data, although a variety of other sources are also used. There is an element of subjectivity used in the determination of the RPI item weights. Many of the detailed LCF expenditure codes do not map directly to a specific item, so the weight associated with that expenditure may be split between items equally or pro rata to another variable.

2.2.3 CPI weights

COICOP weights for the CPI are based on household final monetary consumption expenditure (HFMCE) estimates from the National Accounts. All monetary expenditure by private individuals in the UK is included (imputed expenditures are not included, although imputed rents will be included in CPIH). This covers expenditure by private households, institutional residents (such as people living in retirement homes or university halls of residence) and by overseas visitors to the UK. The weights for year y are based on consumers' expenditure for the calendar year y-2 and are up-rated by movements in the relevant COICOP class index from year y-2 to December y-1.

The National Accounts estimates of consumers' expenditure are derived from a variety of sources, including the LCF. The LCF estimates are primarily used for expenditure on services that cannot be obtained from other sources. For instance, expenditure on education tuition fees is sourced from universities, while travel by air comes from the International Passenger Survey. The National Accounts estimates are subject to balancing adjustments to ensure that the supply of products in the economy matches the demand for the same products.

CPI item weights are based on the RPI item weights. They are derived by mapping the RPI items to the most detailed COICOP-consistent category for which National Accounts consumers' expenditure exists (this may be more detailed than COICOP class) and then scaling the RPI item weights to match the CPI COICOP weight for that category.

2.3 Living costs and food survey

The LCF is a sample survey covering the United Kingdom; it is run continuously[11]. In Great Britain, a two-stage sampling approach is used, with primary sampling units (PSUs) selected within 26 strata defined by region and type of area (metropolitan, non-metropolitan), and households randomly sampled within PSUs. In Northern Ireland, households are the PSU. Annual results are based on a sample of an achieved sample of around 4,500 households, which represents a response rate of just under 50 per cent.

[11] Living Costs and Food Survey, ONS, 2016, http://www.ons.gov.uk/peoplepopulationandcommunity/personalandhouseholdfinances/incomeandwealth/methodologies/livingcostsandfoodsurvey

Expenditure recorded by the LCF is coded based on a provisional version of the COICOP classification system, which differs in some respects from the final version. In particular, there are some noticeable differences in Division 09, recreation and culture. There are also codes for expenditure on categories not covered by COICOP but that form part of the RPI – most notably, expenditure on mortgage interest payments and council tax.

The LCF breakdown of expenditure can be seen in Table A1[12] which is updated each year when the latest LCF results are published on the ONS website. There are up to four levels of classification. There is a complete breakdown of expenditure to the third (class level), with some broken down further to a fourth (sub-class) level. The most detailed published category is called COICOP-LCF elsewhere in this article; there are 164 such categories.

In the analysis of the impact of variation in the LCF on the accuracy of the price indices, the LCF data used covered the period 2013Q3 to 2014Q2.

The published LCF results are weighted to take account of the design of the survey and non-response bias. The latter involves recalibration to align with the 2011 census population profile as updated in the successive mid-year estimates.

2.4 Bootstrapping

The estimates of variance in this article have been derived using the technique of bootstrapping[13]. This is the practice of estimating properties of an estimator (such as its variance) by measuring those properties when sampling from an approximating distribution. One standard choice for an approximating distribution is the empirical distribution function of the observed data. In the case where a set of observations can be assumed to be from an independent and identically distributed population, this can be implemented by constructing a number of resamples with replacement, of the observed dataset (and of equal size to the observed dataset).

Bootstrapping is of particular use when the theoretical distribution of a statistic of interest is complicated or unknown. Since the bootstrapping procedure is distribution-independent it provides an indirect method to assess the properties of the distribution underlying the sample and the parameters of interest that are derived from this distribution.

[12] http://webarchive.nationalarchives.gov.uk/20160105160709/http://www.ons.gov.uk/ons/rel/family-spending/family-spending/2015-edition/rft-a1.xls

[13] https://en.wikipedia.org/wiki/Bootstrapping_(statistics)

3. Impact of variance in the weights on the RPI

3.1 Method

The use of bootstrapping to calculate the impact on the RPI of variance in the weights involved drawing multiple re-samples from the LCF for 2013q3 to 2014q2. The LCF uses a stratified random sample. If there are n_k primary sampling units (PSUs) in a stratum then for the bootstrapping, samples of size n_k-1 were drawn in each stratum, with replacement. The household weights were adjusted to take account of the changed sampling probabilities by multiplying them by $n_k/(n_k-1)$ but, in order to simplify the calculations, they were not recalibrated to align with the 2011 census population profile. This may have a small effect on the results, but is unlikely to alter the overall conclusions. This can be seen in Chart 3.1 which compares the standard deviation of the mean RPI obtained from 100 resamples, with and without weights. The biggest difference is 0.012 index points. The actual effect of not re-calibrating the weights is likely to be much smaller than this.

Chart 3.1: Standard deviation of RPI (index points)

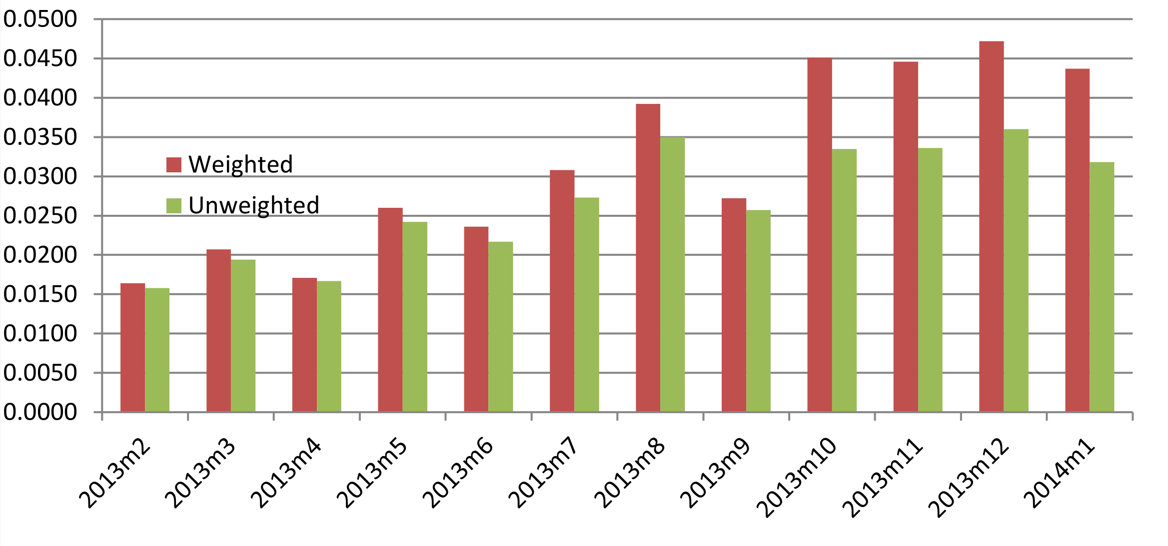

Estimates of the effect of variance in the LCF expenditures on the RPI were obtained by running 200 bootstrap iterations. This yielded 200 sets of COICOP-LCF weights. These were applied to the COICOP-LCF in-year indices; these are indices based on the previous January=100, calculated by weighting together the item indices that form part of each COICOP-LCF category, using the published RPI item weights. This was repeated for 25 months running from January 2013 to January 2015. For each month, 200 calculations of the RPI all items index were produced, and the variance of these estimates calculated.

3.2 Results

The calculation of the impact of variance in the LCF on the RPI involves a number of approximations and simplifications. For instance, the LCF expenditure data used in this analysis, which covers the period 2013q3 to 2014q2, was used in the actual RPI for indices covering the period February 2015 to January 2016. However, in this analysis the resulting weights have been applied to indices covering earlier periods. Simplifications included the use of actual LCF data without adjustments for under-reporting of alcohol and tobacco expenditure and the absence of price-updating of the weights.

Chart 3.2 shows how the average index of the 200 simulations calculated for the RPI compares with the published RPI. Generally the results are close to the actual RPI. Differences are most pronounced in the latest periods, peaking at 0.33 index points in Jan 2015. The biggest single factor contributing to this was air fares which had a higher weight in the LCF weights. Owner-occupier housing costs, which have a lower weight in the LCF, were also a large contributor, as was tobacco.

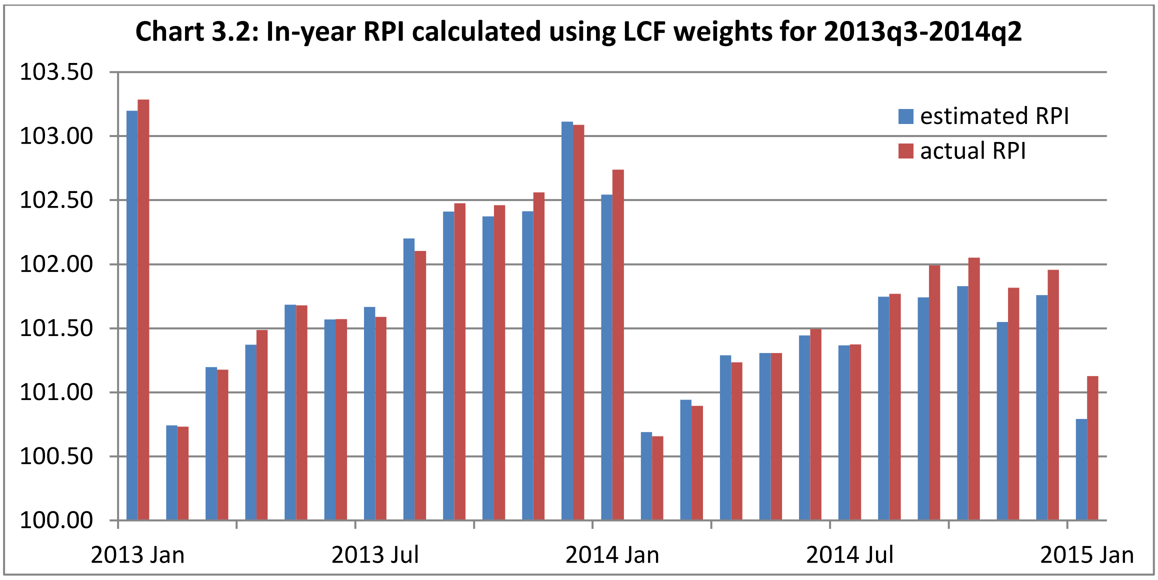

The headline results of the bootstrapping, given in Table 1 in the first column (and shown in Chart 3.3), shows the standard deviation rising during the course of the year, reaching around 0.05 index points towards the end of the year. The result for Jan 2013 is over twice as high, in large part due to the start of the phasing in of the £9,000 per year university tuition fees in England. This is large, lumpy expenditure paid by a minority of households. There is also a small step up in August of the two years shown; this is largely attributable to variation in the weight associated with air fares, which tend to increase sharply in the summer months. It should be noted that the increase in tuition fees was phased in over a three year period from October 2012 to October 2014, so the increased variance attributable to it is a temporary effect.

These results suggest that the 95 per cent confidence intervals for the all items RPI 12-month rate arising from variance in the weights are up to +/-0.1 percentage points. The results for sub-divisions of the RPI are shown in Chart 3.4. The education effect is clearly visible, with average standard deviations around four times higher than for the next highest groups.

Table 1: Standard deviation of mean RPI (in-year index points)

		PSU subsample			case subsample		
	100%	95%	90%	75%	95%	90%	75%
2013 Jan	0.114	0.125	0.118	0.131	0.117	0.125	0.138
2013 Feb	0.012	0.012	0.014	0.014	0.013	0.013	0.014
2013 Mar	0.016	0.018	0.019	0.021	0.018	0.019	0.019
2013 Apr	0.014	0.016	0.017	0.018	0.015	0.015	0.017
2013 May	0.022	0.025	0.026	0.029	0.026	0.027	0.027
2013 Jun	0.021	0.023	0.024	0.027	0.024	0.025	0.025
2013 Jul	0.029	0.032	0.033	0.037	0.034	0.034	0.036
2013 Aug	0.036	0.040	0.042	0.046	0.042	0.044	0.045
2013 Sep	0.026	0.028	0.029	0.032	0.029	0.030	0.030
2013 Oct	0.047	0.049	0.049	0.054	0.047	0.050	0.055
2013 Nov	0.045	0.050	0.048	0.053	0.046	0.050	0.055
2013 Dec	0.048	0.050	0.051	0.055	0.049	0.051	0.057
2014 Jan	0.043	0.047	0.045	0.051	0.045	0.048	0.053
2014 Feb	0.015	0.016	0.016	0.017	0.016	0.017	0.018
2014 Mar	0.017	0.018	0.019	0.020	0.019	0.019	0.019
2014 Apr	0.020	0.022	0.024	0.025	0.023	0.022	0.024
2014 May	0.019	0.022	0.022	0.024	0.023	0.021	0.023
2014 Jun	0.025	0.027	0.028	0.030	0.028	0.027	0.029
2014 Jul	0.030	0.032	0.035	0.038	0.035	0.034	0.037
2014 Aug	0.041	0.044	0.047	0.051	0.047	0.047	0.051
2014 Sep	0.027	0.030	0.030	0.033	0.030	0.030	0.032
2014 Oct	0.044	0.048	0.047	0.051	0.045	0.048	0.052
2014 Nov	0.044	0.049	0.047	0.052	0.046	0.050	0.054
2014 Dec	0.050	0.053	0.052	0.057	0.051	0.053	0.058
2015 Jan	0.046	0.052	0.050	0.056	0.050	0.054	0.058

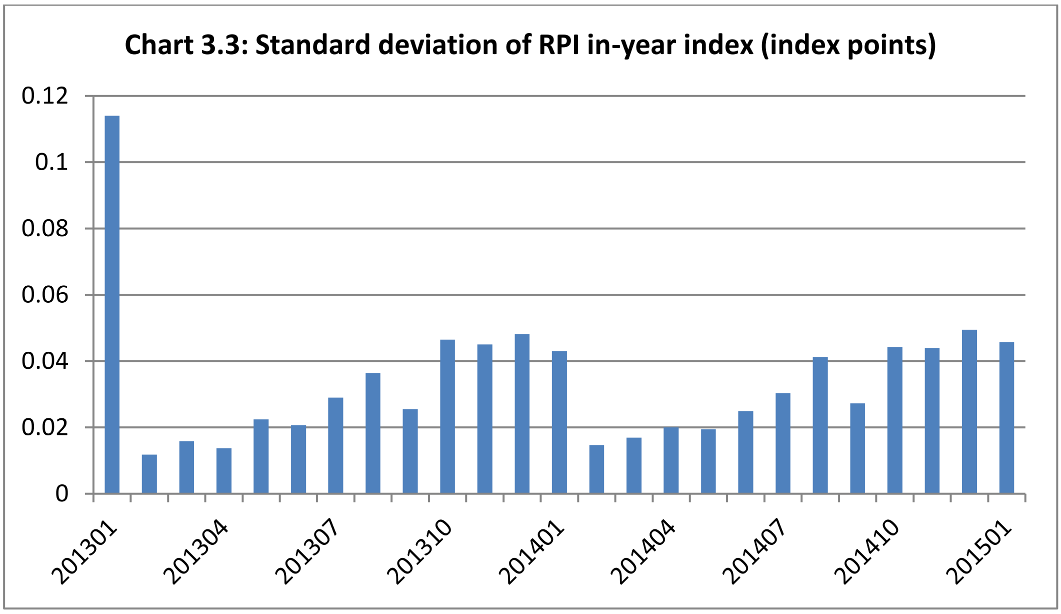
Chart 3.3: Standard deviation of RPI in-year index (index points)
0.12
0.1
0.08
0.06
0.04
0.02
0
201301
201304
201307
201310
201401
201404
201407
201410
201501

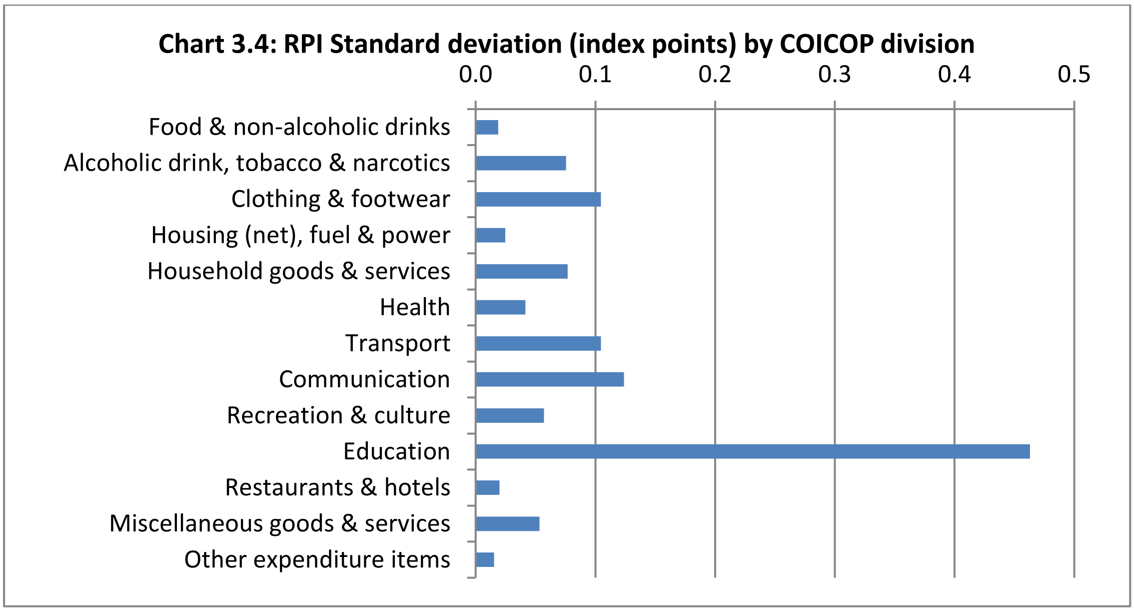
Chart 3.4: RPI Standard deviation (index points) by COICOP division
0.0
0.1
0.2
0.3
0.4
0.5
Food & non-alcoholic drinks
Alcoholic drink, tobacco & narcotics
Clothing & footwear
Housing (net), fuel & power
Household goods & services
Health
Transport
Communication
Recreation & culture
Education
Restaurants & hotels
Miscellaneous goods & services
Other expenditure items

4. Impact of reducing the sample size of the LCF

4.1 Method

The impact on the RPI of reducing the sample size of the LCF by 5 per cent, 10 per cent and 25 per cent has been assessed. There are different ways in which a reduction in sample size can be achieved:

1. the number of PSUs can be reduced (this is the most economical way of achieving a reduction);
2. the number of households can be reduced (in principle, this minimises the reduction in precision);
3. a combination of (1) and (2).

The first two options can be considered to be the two extremes. Ten samples were drawn for each of options (1) and (2), for each of the three reductions in sample size. The reductions in sample size were achieved by randomly deleting PSUs or households from the LCF data set (before bootstrapping) so that the desired number of PSUs or households was obtained for each stratum.

Fifty bootstrap iterations were then performed on each sample, yielding 500 sets of weights for each reduction in sample size. These weights were then applied to the COICOP-LCF indices for January 2013 to January 2015, and the variance of the resulting all items RPI determined.

For some of the options, a further ten samples were drawn to check the sensitivity of the results to the number of samples on which the bootstrapping was performed. The results from these repeat analyses were very similar to those obtained from the first set of randomly drawn samples.

4.2 Results

Table 1 above shows that the method of reducing the sample size (PSU or case/ household sampling) has minimal effect on precision, and that the standard deviation increases with the size of the cut, although in absolute terms the impact on the headline RPI is minimal with the 75% sample having an standard deviation that is a maximum of 0.010 index points higher than the 100% (ignoring Jan 2013).

These results are not altogether surprising. The weights are constrained to sum to 100% and indices tend to move in a similar fashion, so the potential for sampling error in the weights to impact substantially on the overall index is likely to be limited. This contrasts with sampling errors in the indices which are unconstrained.

It should be noted that these results are indicative only. The weights are calculated by randomly varying the original weights and then dividing by the sum of the varied weights. So sampling error in the weights arises from the numerator (the randomly varied weights) and the denominator (the sum of the randomly varied weights). However, only the former is explicitly being allowed for.

5. CPI inflation rates for different types of household

5.1 Background

The Johnson Review of consumer price indices recommended regular (probably annual) analyses of the inflation experience of different types of households. In response to this, Flower & Wales produced estimates of the inflation rates experienced by the following different types of households:

1. Household expenditure decile
2. Household income decile
3. Retired and non-retired households
4. Households with and without children

Their calculations were done using CPI class indices, coupled with weights largely based on the LCF. The analysis presented in this article can be used to assess the relative accuracy of these results, arising from variation in the weights.

5.2 Methods

The published CPI indices are based on COICOP but differ at the most detailed level from the LCF classification, which provides a finer breakdown for some categories, for example, food and clothing.

For the calculation of the sampling error due to the expenditure weights, CPI item indices were aggregated to COICOP-LCF class level using CPI item weights for the period January 2013 to January 2015 inclusive. This exercise was repeated using a finer classification of items to COICOP sub-categories, more consistent with the final version of COICOP.

Estimates of the standard deviation of the mean indices were obtained by bootstrapping. Expenditure weights were obtained by aggregating LCF expenditure codes to COICOP-LCF class level. No adjustments were made to these weights to try to align them with the actual CPI weights, which are based on estimates of annual household expenditure taken from the National Accounts. Chart 5.1 compares the actual in-year CPI index against that obtained by calculating weights directly from the LCF. The broad trends can be seen to be consistent, although there are some noticeable differences, particularly between September 2014 and January 2015. However, it is likely that there is little difference in their variances.

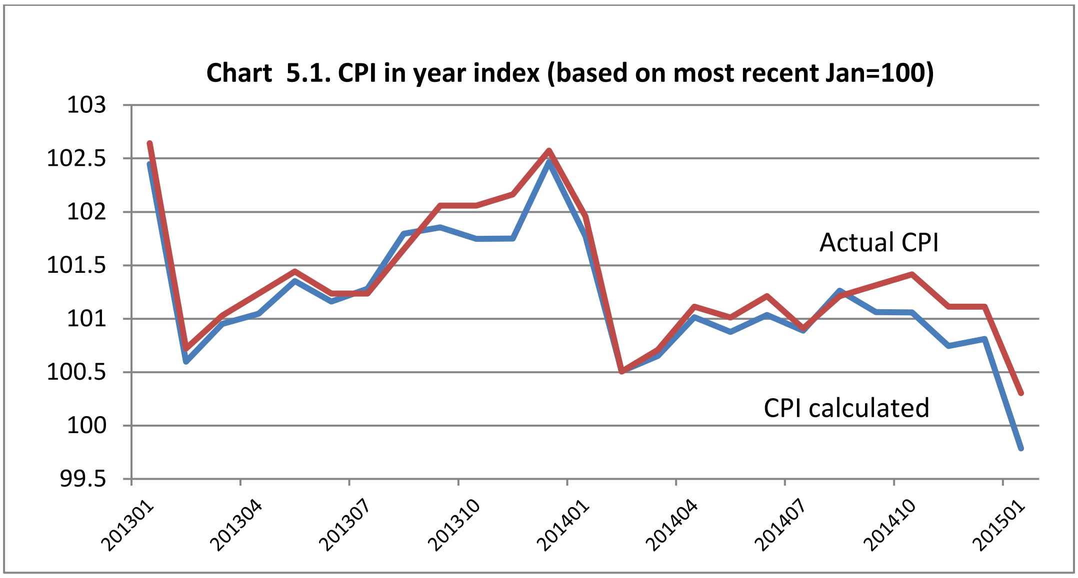

Separate sets of expenditure weights were produced for each category of household. For households categorised by decile, equivalised income and expenditure totals were used. Equivalisation is a process which adjusts income or expenditure by the size of the household[14]. The modified OECD scale was used. In this, the first person in the household was given a weight of 1.0; subsequent persons (aged 16+) were each given a weight of 0.5; and children (under the age of 16) were each given a weight of 0.3.[15]

Retired households were defined as those households where all females are aged 60 and over, and all males are aged 65 and over. This is a slightly different definition to that used by Flower & Wales which defined a retired person as anyone who describes themselves (in the LCF) as 'retired'; or anyone over minimum National Insurance pension age describing themselves as 'unoccupied' or 'sick or injured but not intending to seek work'. A retired household is then defined as one where the combined income of retired members amounts to at least half the total gross income of the household.

Households with children were defined as those households containing at least one child aged under 16.

The unadjusted expenditure weights for each type of household were combined with the COICOP-LCF indices to give estimates of COICOP class, group and division level indices, as well as the all items index, for January 2013 to January 2015.

[14] Family Spending, ONS, 2015, Chapter 3, http://www.ons.gov.uk/peoplepopulationandcommunity/personalandhouseholdfinances/incomeandwealth/compendium/familyspending/2015/chapter3equivalisedincome

[15] This is a slight approximation. The actual modified OECD scale defines 14 year-olds as adults, but the data used in this analysis contains banded ages, rather than the specific ages of individual household members.

5.3 Results for the overall CPI

The results of the bootstrapping are summarised in Charts 5.2 and 5.3 and the tables below, with the detailed results tabulated in Annex B. The results shown are the standard deviation of the in-year indices – i.e. the index calculated with reference to the previous January.

The effect on the overall CPI of variation in the weights if calculated from the LCF are very similar to those found for the RPI, with the standard deviation rising during the course of the year, reaching around 0.05 index points towards the end of the year. In practice, this is likely to be an over-estimate as the CPI weights are based on National Accounts estimates of household expenditure after they have been put through the balancing process. This process will tend to smooth out random movements in expenditure.

The main focus of interest on the CPI (and other inflation measures) is the 12-month rate of change of the index. In the analyses presented in this article, this is only shown for the January results because each 12-month period is treated separately with the first month fixed. Because of this, some of the results presented for income and expenditure deciles only show the January results for 2013, 2014 and 2015. The standard deviations for January 2013 tend to be higher than for other periods. This is in large part due to the phasing in of the £9,000 per year university tuition fees in England which is paid by a minority of households. As noted above, in the discussion of the results for the RPI, this is a temporary effect.

Chart 5.2 also shows peaks in the standard deviation in the summer months which is primarily due to air fares. Like education, air fares tend to be lumpy expenditure paid by a minority of households. This effect is more pronounced than for the RPI, probably reflecting the exclusion of the higher income households (who are more likely to travel by air) from the RPI.

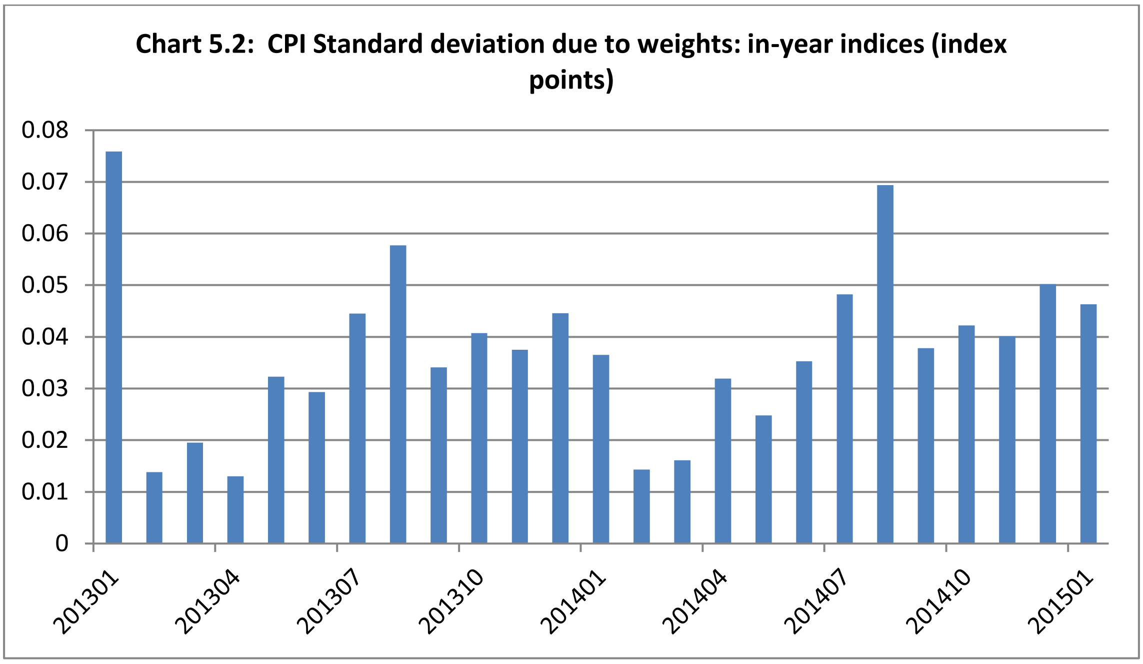

Chart 5.3 shows how the average standard deviation varies by CPI division, and that education is the highest. Next highest is Communications, largely due to mobile phone handsets which tend to show large falls in price during the course of a year. However, its impact on the overall CPI is less than transport costs, because the weight of the latter is considerably higher – 14.5% compared with 1.7% in 2014. Compared with the RPI, both transport and communication have higher variances.

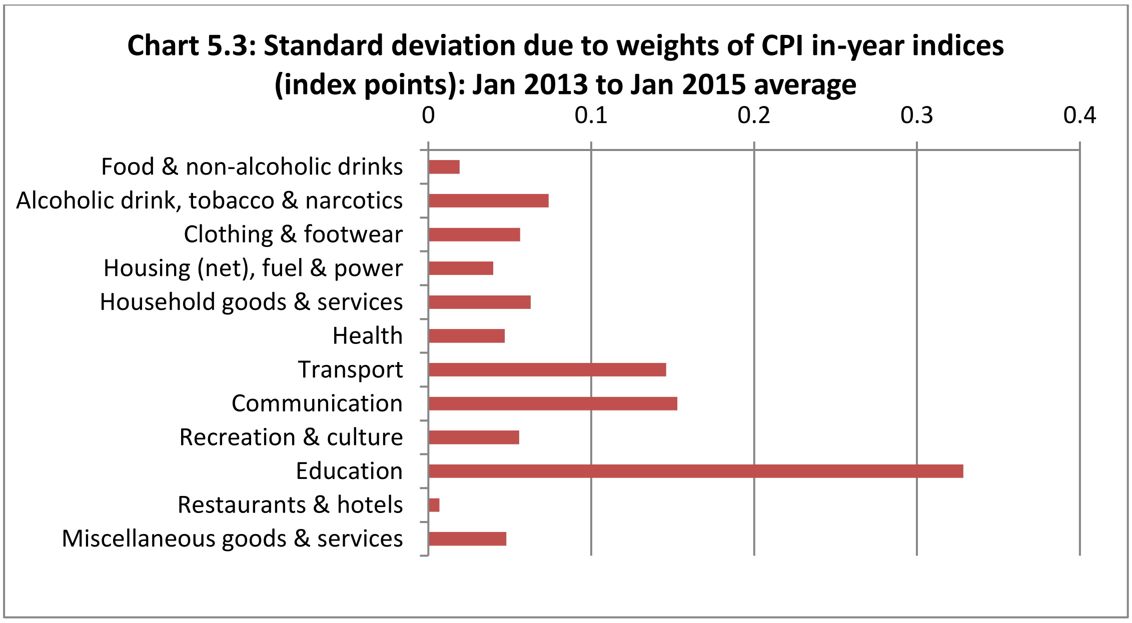

5.4 Results for expenditure deciles

Chart 5.4 shows the average standard deviation over the 25 months by expenditure decile. It can be seen that:

- The standard deviation for the expenditure deciles increases as expenditure increases. This implies greater homogeneity of expenditure patterns for the lowest spenders, which is unsurprising – Flower & Wales found that what might be considered the essentials (food, clothing and housing costs) represent half of the lowest decile's expenditure compared with one-fifth for the highest spending households.
- The standard deviation for the lowest spending decile (0.034) is one-sixth that of the highest spending decile (0.189 index points), and below the overall average (0.037).

Chart 5.4: Average standard deviation of CPI in-year indices by expenditure decile (index points)

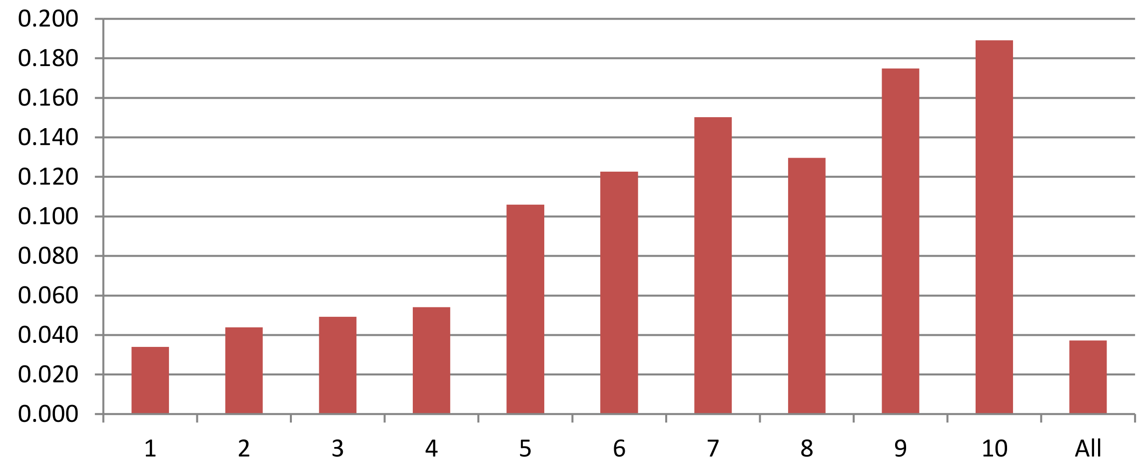

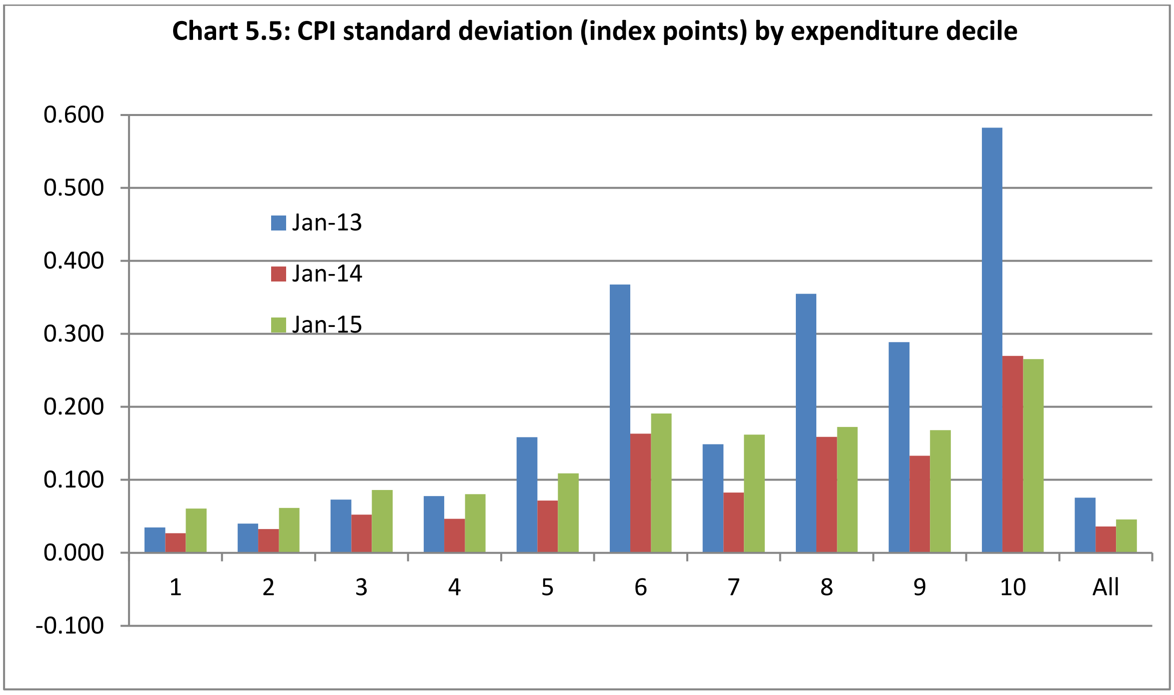

The January results for expenditure deciles are shown in Chart 5.5. Here the pattern is similar to the overall average standard deviations for expenditure deciles shown in Chart 5.4, with the lowest deciles being much less variable than the highest ones. In fact, the standard deviations for January 2013 and 2014 for the highest decile are around ten times that for the lowest decile; with the standard deviation for January 2015 being five times higher.

5.5 Results for income deciles

Chart 5.6 shows the average standard deviation over the 25 months by income decile. It can be seen that the standard deviations for the highest earning 30 per cent of households are noticeably higher than for lower earning households, with the 8th decile exceeding 0.2 index points and the 10th decile exceeding 0.3 index points, with 95% confidence intervals that are roughly twice this. Lower income households have standard deviations of around 0.1 index points.

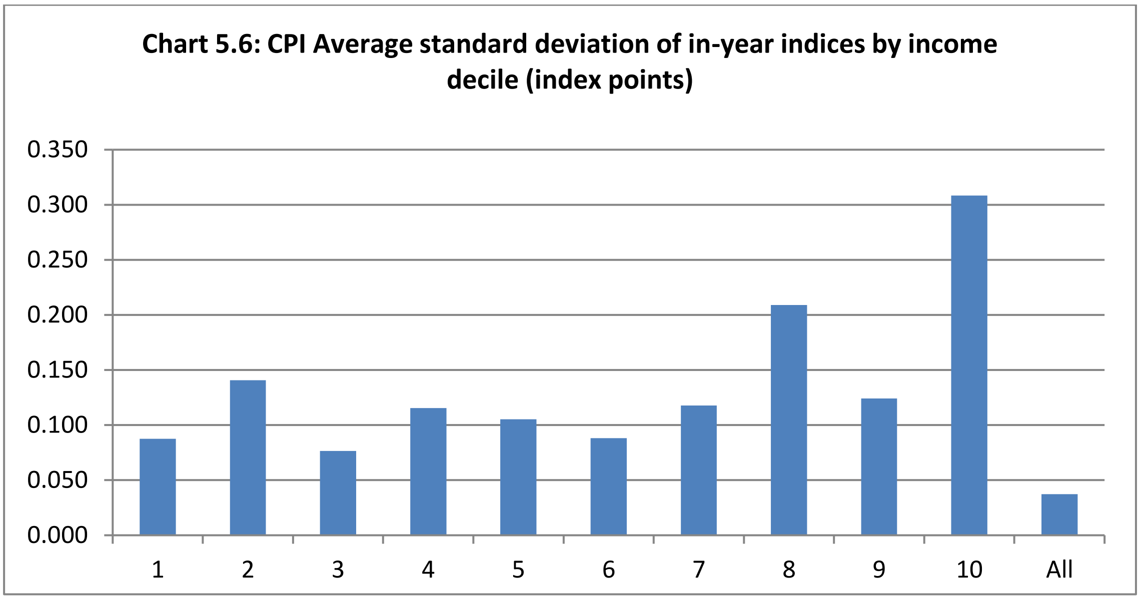

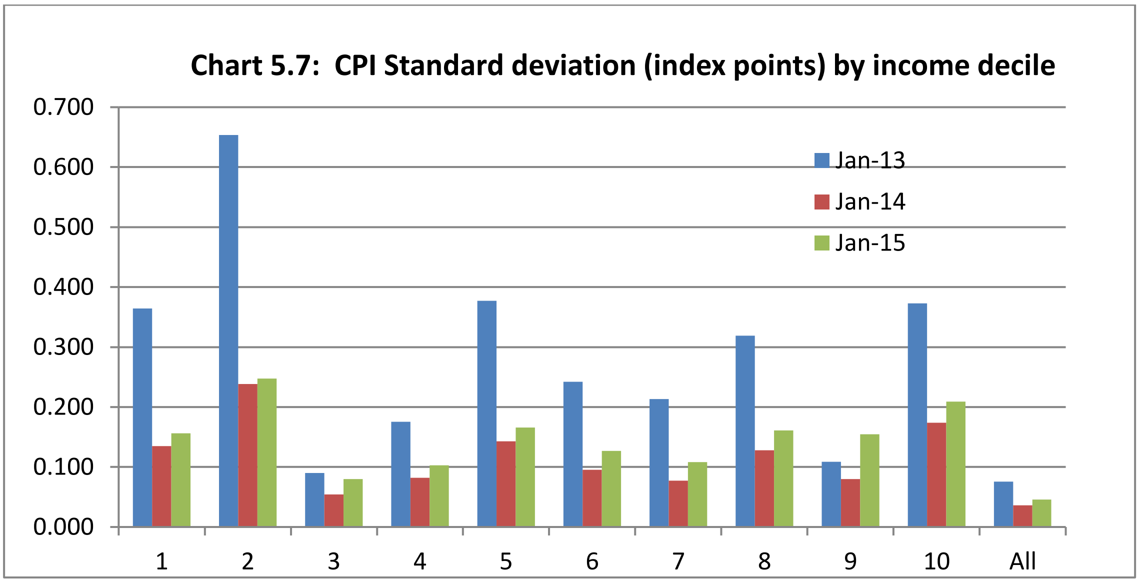

The January 12-month rates for income deciles are summarised in Chart 5.7. The picture differs somewhat from Chart 5.5, with the second lowest income decile having the highest standard deviations, with little to choose between the 1st, 5th, 8th and 10th deciles. The 3rd income decile is the lowest. The explanation for the high standard deviation for the two lowest deciles lies, in large part, in the variability of the weight attached to the education component. The index for education increased sharply in the preceding October in each of the years shown, reflecting the phasing in of the £9,000 annual university tuition fees.

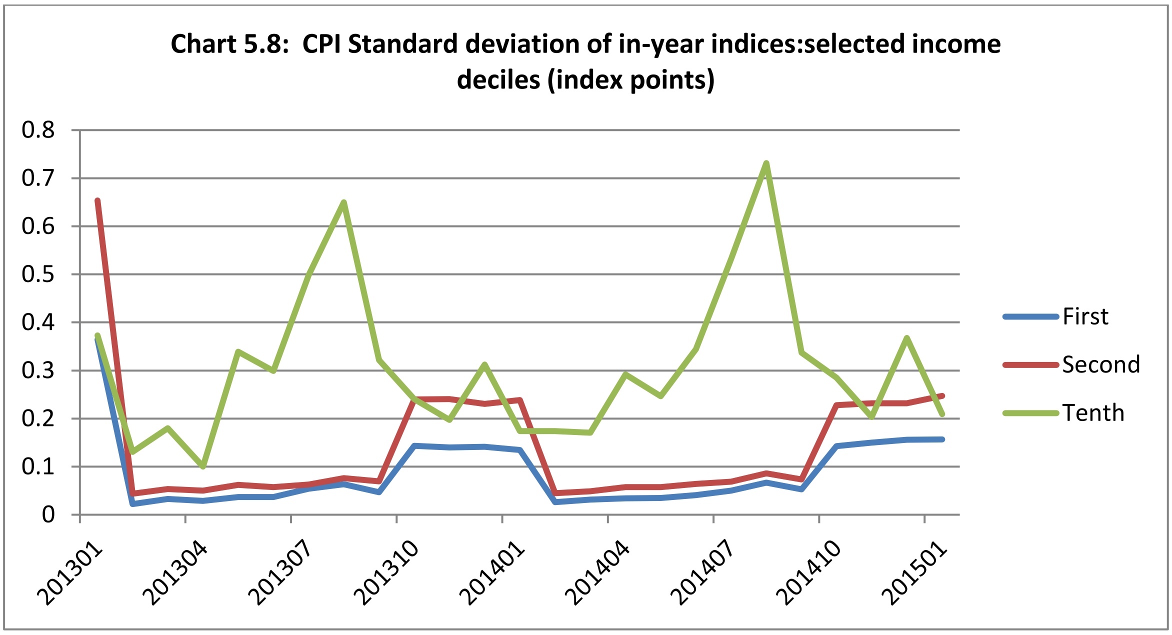

The lowest two deciles have relatively few households paying tuition fees, but for those that do, their household inflation rate will be considerably higher than other households, which serves to increase their standard deviation. This can be seen in Chart 5.8 where the standard deviations for the two lowest income deciles increase in October. For the tenth decile, a different effect is at work. The peak standard deviation occurs in the summer. This is when air fares increase sharply. Like education, air fares can be large, lumpy expenditure paid by a minority of households, which as the chart shows are more likely to be paid by the highest earning households.

5.6 Results for retired and non-retired households

The standard deviations for retired and non-retired households are shown in Chart 5.9. It can be seen that the standard deviation for retired households tends to be higher than for non-retired households, which in turn tends to be slightly higher than for all households. This may simply be due to there being fewer retired households in the sample. Or it may be that retired households have more heterogeneous spending patterns than non-retired households, perhaps reflecting a greater range of incomes, from those mainly dependent on state benefits to those on index-linked occupational pensions. It may also reflect differences in the health and mobility of this group of the population.

The vast majority of standard deviations for retired households are in the range 0.04 to 0.06 index points which imply 95% confidence intervals for the in-year indices of between +/-0.08 to 0.12 index points. The standard deviations for non-retired households tend to be lower, particularly in the first part of the year, where 95% confidence intervals are below +/-0.08 index points.

The summer spike, associated with air fares, is apparent for both household types but the education effect is less pronounced, which is not surprising as most expenditure on undergraduate tuition fees will be concentrated in the non-retired households, resulting in relatively more homogeneous expenditure patterns than for all households, other things being equal.

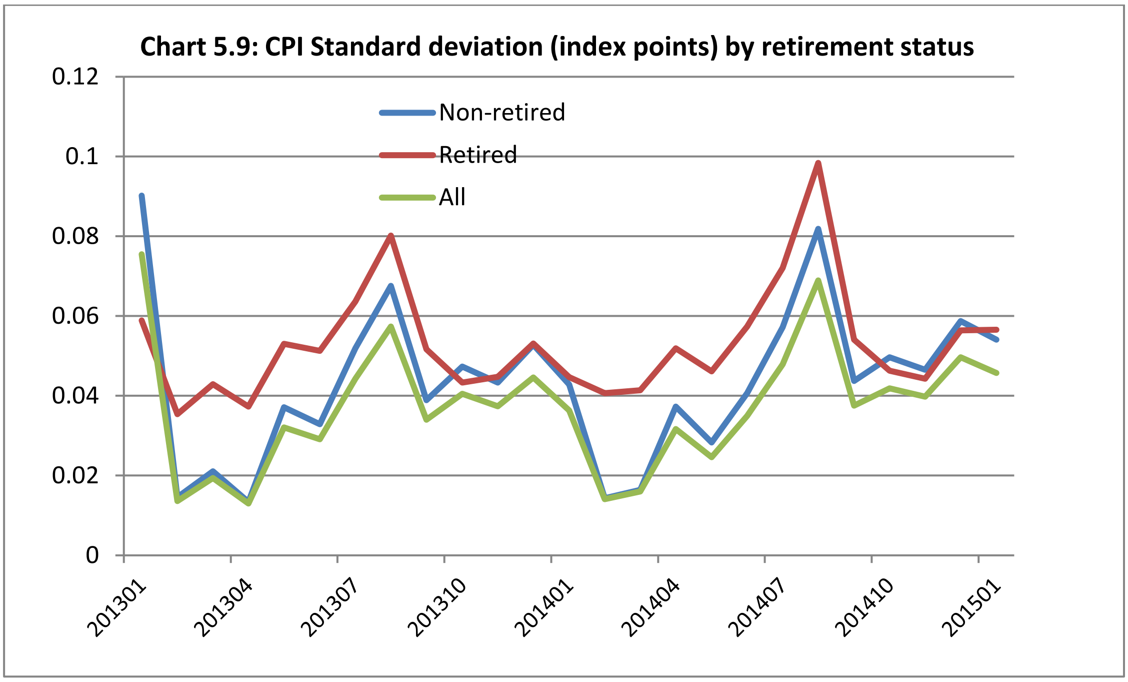

5.7 Results for households with and without children

Chart 5.10 shows the standard deviations for households with and without children. The picture is similar to that for non-retired households. Standard deviations for households with children tend to be higher than for those without children, particularly in the second half of the year. This may be the tuition fees effect coming through, as households with university students are more likely to contain younger children than those without. As a result, sampling errors for households with children tend to be higher, with most having 95% confidence intervals of up to +/-0.14 index points, compared with +/-0.10 index points for those without children.

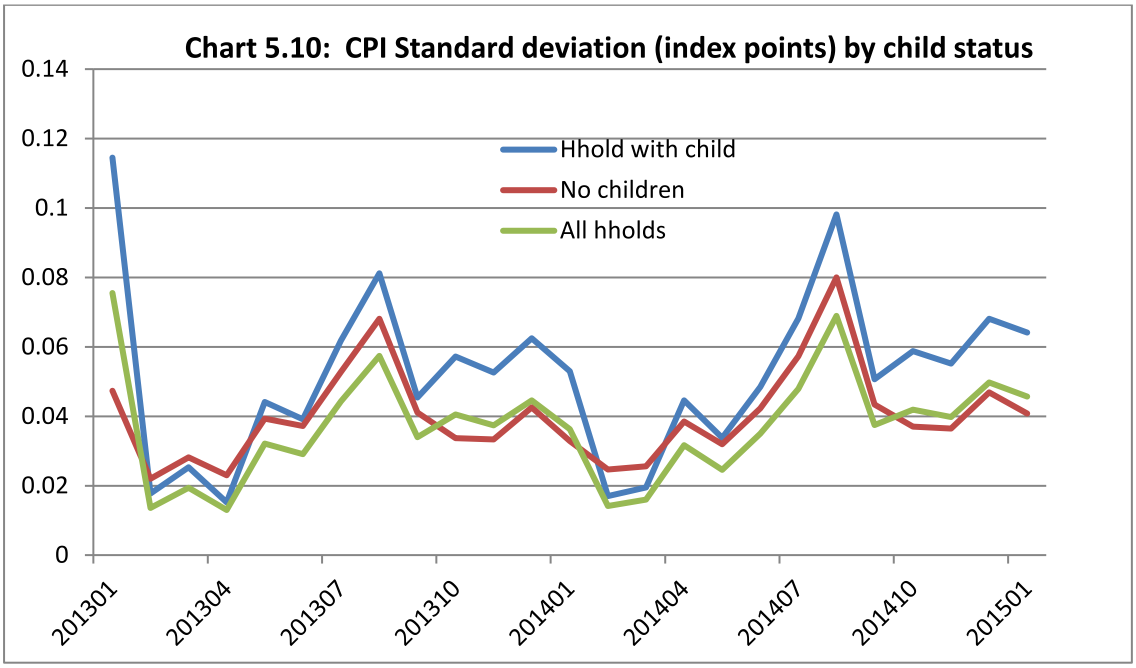

5.8 Conclusions

The variability of the results for the in-year indices are heavily influenced by expenditure by some, but not all, households on products with sharp price movements. Two specific examples of this are shown above: university tuition fees and air fares. The tuition fees particularly influence the periods covered by this analysis; they are likely to be less of an issue in other periods. The summer spike in the cost of air travel is persistent through time and will tend to increase sampling errors for all categories of households that incur these costs.

This may go some way towards explaining why the average standard deviation of the in-year indices increases as total household expenditure rises. The 95% confidence interval for the lowest expenditure decile is +/-0.07 index points, compared with nearly +/-0.4 index points for the highest decile. The impact on the highest earning households is even more pronounced with the 10th decile having 95% confidence intervals of +/-0.6 index points.

Another factor that might be contributing to these results may be that lower expenditure households have more homogeneous expenditure patterns, concentrated on the essentials such as food and housing costs. This contrasts with higher expenditure and higher income households who are more likely to spend varying amounts on a range of discretionary products.

These results indicate that the results presented by Flower & Wales are more robust for the lower expenditure deciles, but less so for the highest income and expenditure deciles.

The sampling errors for retired/non-retired households and for households with and without children are less variable than for the deciles analysis. All exhibit a peak in the summer, most likely due to air fares. At other times of the year retired households generally have 95% confidence intervals of around +/-0.08 to 0.12 index points. The other household types exhibit similar patterns with smaller confidence intervals in the first part of the year. However, households with children tend to have greater confidence intervals in the last few months of the year, which most likely reflects the impact of tuition fees.

Finally, it should be noted that these results are sensitive to the number of COICOP sub-categories to which the items are categorised. The results presented above are based on 160 categories. A slightly more detailed classification of 180 categories was also tested. This more detailed classification raises the estimates of the standard deviation of the all items CPI by an average of 0.027 index points, with a maximum difference of 0.053 index points (see Chart 5.11). The increase is primarily due to air fares, and to a lesser extent recreation and culture.

In the original classification air fares had been classified to "other travel costs" alongside sea fares, taxi and minicab fares, and bicycle purchases. The inclusion of these other categories tends to smooth out fluctuations in expenditure on air fares; in the revised classification, it has its own category.

In the case of recreation & culture, the number of categories has increased from 23 to 31. This finer categorisation leads to greater variability in the weights of the different categories, some of which (particularly the technology goods, such as computers) have pronounced index movements.

Chart 5.11: Standard deviation of CPI in-year indices due to weights (index points)

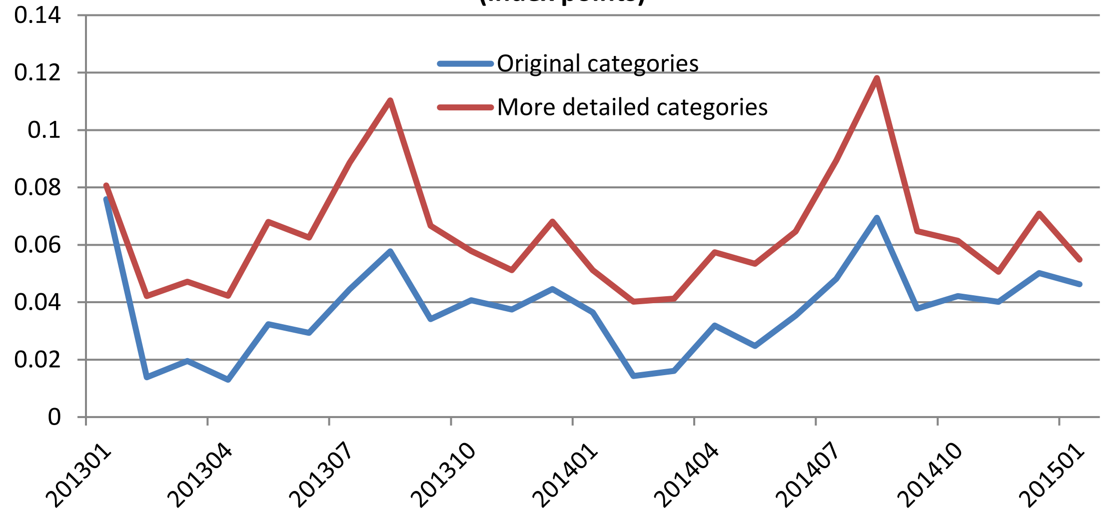

The impact of education tuition fees and air fares on the CPI's variance can be seen in Chart 5.12. Almost all of the fluctuations in the variance of the CPI is attributable to these two categories in the period shown.

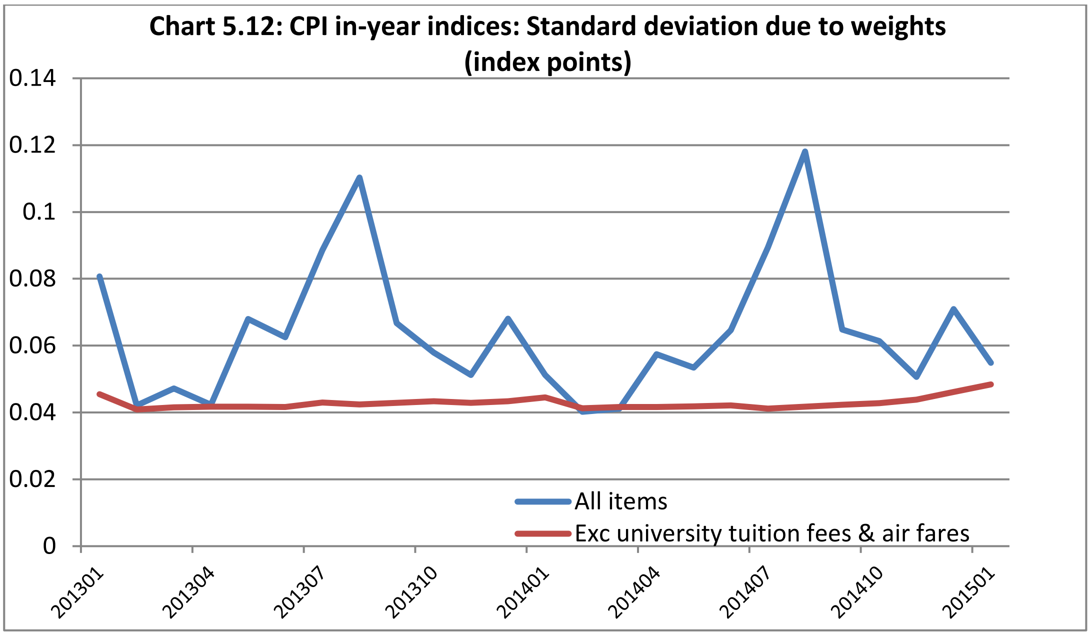
Chart 5.12: CPI in-year indices: Standard deviation due to weights (index points)
0.14
0.12
0.1
0.08
0.06
0.04
0.02
0
201301
201304
201307
201310
201401
201404
201407
201410
201501
All items
Exc university tuition fees & air fares

Annex A: Illustration of the sample selection

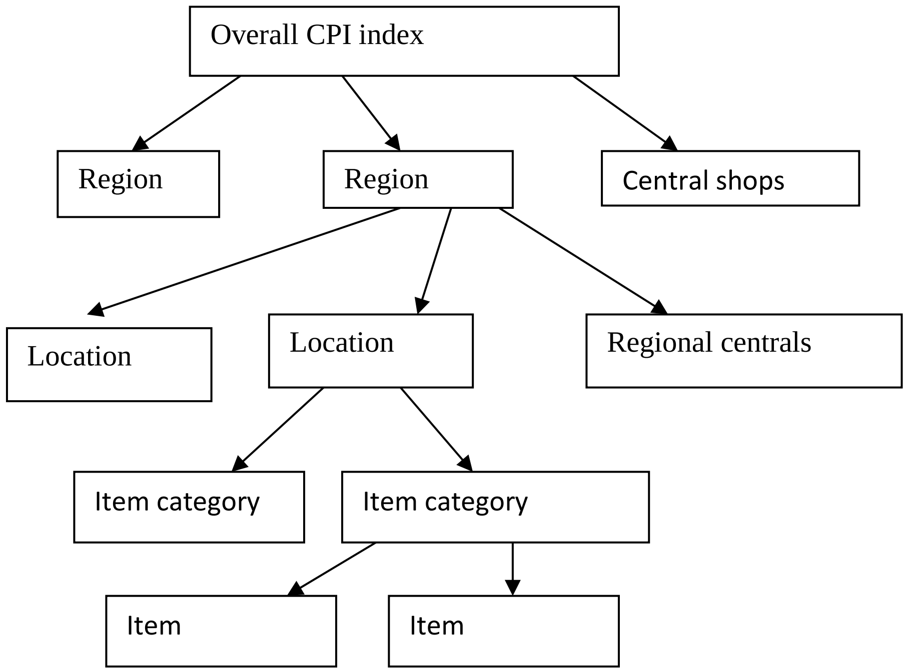

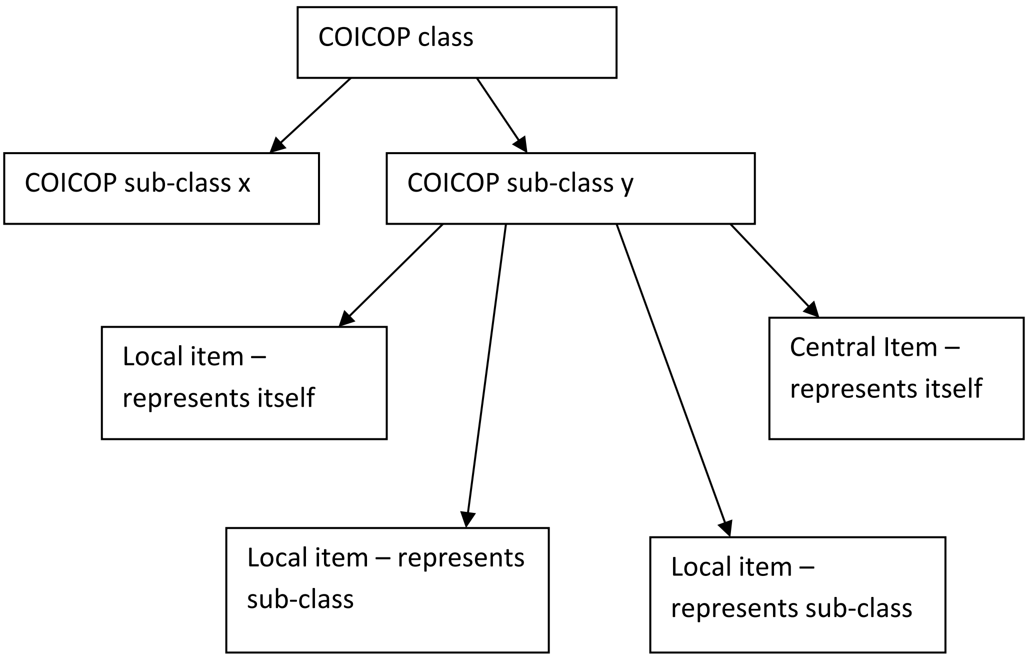

Annex B1. Standard Deviation of in-year indices (index points) for CPI by income decile

		Equivalised household income decile									
	All items	1	2	3	4	5	6	7	8	9	10
201301	0.076	0.364	0.654	0.090	0.176	0.377	0.242	0.213	0.319	0.109	0.373
201302	0.014	0.022	0.044	0.039	0.048	0.033	0.040	0.062	0.102	0.059	0.131
201303	0.019	0.033	0.053	0.056	0.068	0.049	0.052	0.087	0.143	0.094	0.180
201304	0.013	0.029	0.050	0.048	0.060	0.044	0.042	0.078	0.090	0.068	0.100
201305	0.032	0.037	0.062	0.076	0.123	0.069	0.065	0.121	0.223	0.129	0.339
201306	0.029	0.037	0.058	0.067	0.112	0.064	0.061	0.106	0.197	0.111	0.299
201307	0.044	0.054	0.063	0.088	0.169	0.081	0.074	0.142	0.274	0.171	0.498
201308	0.057	0.063	0.076	0.115	0.217	0.105	0.094	0.185	0.4	0.223	0.650
201309	0.034	0.047	0.070	0.083	0.118	0.078	0.075	0.140	0.229	0.133	0.322
201310	0.041	0.143	0.240	0.079	0.102	0.149	0.104	0.114	0.177	0.109	0.241
201311	0.037	0.140	0.240	0.079	0.098	0.146	0.106	0.109	0.152	0.100	0.197
201312	0.045	0.141	0.230	0.087	0.123	0.144	0.112	0.127	0.233	0.131	0.313
201401	0.036	0.135	0.239	0.054	0.082	0.143	0.096	0.077	0.128	0.080	0.174
201402	0.014	0.026	0.045	0.043	0.051	0.037	0.054	0.064	0.092	0.059	0.174
201403	0.016	0.031	0.049	0.053	0.055	0.048	0.059	0.075	0.115	0.079	0.170
201404	0.032	0.034	0.057	0.067	0.102	0.057	0.059	0.108	0.208	0.111	0.292
201405	0.025	0.035	0.057	0.066	0.084	0.058	0.061	0.097	0.180	0.099	0.247
201406	0.035	0.041	0.064	0.080	0.115	0.069	0.075	0.121	0.241	0.127	0.344
201407	0.048	0.050	0.069	0.094	0.182	0.089	0.079	0.146	0.297	0.187	0.532
201408	0.069	0.067	0.086	0.127	0.247	0.119	0.106	0.199	0.420	0.252	0.731
201409	0.038	0.053	0.073	0.082	0.119	0.074	0.079	0.120	0.245	0.136	0.337
201410	0.042	0.143	0.228	0.080	0.103	0.139	0.106	0.111	0.206	0.119	0.285
201411	0.040	0.150	0.232	0.083	0.097	0.145	0.113	0.104	0.162	0.109	0.203
201412	0.050	0.156	0.232	0.098	0.132	0.153	0.126	0.131	0.255	0.155	0.368
201501	0.046	0.156	0.247	0.080	0.103	0.166	0.127	0.108	0.161	0.155	0.209
Average	0.037	0.087	0.141	0.077	0.115	0.105	0.088	0.118	0.209	0.124	0.308
Jan Average	0.053	0.218	0.380	0.075	0.120	0.229	0.155	0.133	0.203	0.114	0.252
Jan av 2014 & 15	0.041	0.146	0.243	0.067	0.092	0.154	0.111	0.093	0.145	0.117	0.192

Annex B2. Standard Deviation of in-year indices (index points) for CPI by expenditure decile

		Equivalised household expenditure decile									
	All items	1	2	3	4	5	6	7	8	9	10
201301	0.076	0.035	0.040	0.073	0.078	0.159	0.368	0.149	0.355	0.289	0.583
201302	0.014	0.016	0.022	0.021	0.024	0.043	0.036	0.059	0.063	0.071	0.059
201303	0.019	0.023	0.030	0.031	0.039	0.061	0.054	0.085	0.071	0.106	0.093
201304	0.013	0.013	0.023	0.024	0.025	0.039	0.032	0.046	0.070	0.066	0.055
201305	0.032	0.031	0.041	0.036	0.048	0.109	0.079	0.162	0.099	0.171	0.120
201306	0.029	0.027	0.040	0.039	0.044	0.099	0.072	0.145	0.099	0.148	0.109
201307	0.044	0.044	0.058	0.054	0.062	0.165	0.115	0.241	0.124	0.226	0.161
201308	0.057	0.056	0.072	0.059	0.079	0.213	0.155	0.316	0.153	0.309	0.215
201309	0.034	0.030	0.043	0.043	0.048	0.104	0.082	0.158	0.103	0.176	0.126
201310	0.041	0.024	0.036	0.045	0.058	0.082	0.158	0.111	0.162	0.164	0.261
201311	0.037	0.022	0.037	0.049	0.054	0.076	0.157	0.089	0.162	0.144	0.250
201312	0.045	0.033	0.042	0.055	0.062	0.100	0.159	0.144	0.172	0.202	0.260
201401	0.036	0.027	0.033	0.052	0.047	0.072	0.163	0.083	0.159	0.133	0.270
201402	0.014	0.012	0.022	0.027	0.022	0.029	0.039	0.040	0.066	0.061	0.102
201403	0.016	0.014	0.024	0.027	0.029	0.034	0.044	0.050	0.066	0.075	0.106
201404	0.032	0.033	0.042	0.037	0.047	0.112	0.086	0.166	0.090	0.173	0.127
201405	0.025	0.026	0.036	0.041	0.038	0.080	0.065	0.118	0.076	0.125	0.113
201406	0.035	0.034	0.047	0.052	0.052	0.114	0.090	0.171	0.099	0.182	0.148
201407	0.048	0.052	0.065	0.059	0.071	0.184	0.136	0.269	0.131	0.247	0.176
201408	0.069	0.071	0.087	0.073	0.096	0.260	0.191	0.383	0.173	0.367	0.247
201409	0.038	0.040	0.049	0.057	0.053	0.112	0.090	0.176	0.092	0.198	0.143
201410	0.042	0.038	0.046	0.057	0.063	0.097	0.163	0.141	0.158	0.184	0.235
201411	0.040	0.038	0.046	0.058	0.062	0.083	0.165	0.115	0.158	0.172	0.242
201412	0.050	0.051	0.056	0.076	0.075	0.115	0.178	0.177	0.170	0.220	0.264
201501	0.046	0.061	0.062	0.086	0.080	0.109	0.191	0.162	0.172	0.168	0.266
Average	0.037	0.034	0.044	0.049	0.054	0.106	0.123	0.150	0.130	0.175	0.189
Jan Average	0.053	0.041	0.045	0.070	0.068	0.113	0.241	0.131	0.229	0.197	0.373
Jan av 2014 & 15	0.041	0.044	0.047	0.069	0.063	0.090	0.177	0.122	0.166	0.151	0.268

Annex B3. Standard Deviation of in-year indices (index points) for CPI for selected household types

	All items	Hholds with children?		Pensioner Hhold?	
		No	Yes	No	Yes
201301	0.076	0.047	0.115	0.090	0.059
201302	0.014	0.022	0.018	0.015	0.035
201303	0.019	0.028	0.025	0.021	0.043
201304	0.013	0.023	0.015	0.013	0.037
201305	0.032	0.039	0.044	0.037	0.053
201306	0.029	0.037	0.039	0.033	0.051
201307	0.044	0.053	0.062	0.052	0.064
201308	0.057	0.068	0.081	0.068	0.080
201309	0.034	0.041	0.045	0.039	0.052
201310	0.041	0.034	0.057	0.047	0.043
201311	0.037	0.033	0.053	0.043	0.045
201312	0.045	0.043	0.063	0.053	0.053
201401	0.036	0.033	0.053	0.043	0.045
201402	0.014	0.025	0.017	0.014	0.041
201403	0.016	0.026	0.020	0.016	0.041
201404	0.032	0.039	0.045	0.037	0.052
201405	0.025	0.032	0.034	0.028	0.046
201406	0.035	0.042	0.048	0.041	0.057
201407	0.048	0.057	0.068	0.057	0.072
201408	0.069	0.080	0.098	0.082	0.098
201409	0.038	0.043	0.051	0.044	0.054
201410	0.042	0.037	0.059	0.050	0.046
201411	0.040	0.037	0.055	0.047	0.044
201412	0.050	0.047	0.068	0.059	0.056
201501	0.046	0.041	0.064	0.054	0.057
Average	0.037	0.040	0.052	0.043	0.053
Jan Average	0.053	0.040	0.077	0.062	0.053
Jan av 2014 & 15	0.041	0.037	0.059	0.048	0.051

An integrated sample design for the ONS household finance surveys

Folasade Ariyibi, Salah Merad [1], Steven Dunstan

Summary

The Office for National Statistics (ONS) currently runs several household surveys that collect data on household finances and living conditions: the Survey on Living Conditions (SLC), the Living Costs and Food (LCF) survey and the Wealth and Assets Survey (WAS). In order to improve coherence, reduce survey costs and produce more precise income and living conditions estimates, ONS is working towards the integration of these surveys. The first stage of an integrated sample design began in 2017, with the integration of the LCF and SLC samples, along with the harmonisation of key components of the survey questionnaire.

In this paper, we describe the integrated sample design of LCF and SLC, and present the method for determining the required sample size and its geographic allocation in view of improving the precision of regional estimates and satisfying the EU regulations on precision of poverty indicators.

1. Introduction

The Living Costs and Food (LCF) survey is a cross-sectional survey, collecting data on household income and detailed expenditure data on food and other items through the use of a diary. The data from this survey are used to produce household income and inequality statistics, including statistics on the effect of taxes and benefits on household income.

The Survey on Living Conditions (SLC) is an annual longitudinal survey with a rotational panel design, collecting detailed income data. Respondents were interviewed for four consecutive years, or waves, before being rotated out, and a new panel rotated in, but from 2017 the SLC is transitioning to a six-wave design. SLC data are used to produce statistics on income and living conditions (SILC); SLC adopts harmonised standards applied across the EU to meet the EU regulation for comparable EU statistics on income and living conditions (EU-SILC) (see Eurostat (2015)).

By exploiting existing similarities in the design and content of the SLC and LCF, it will be possible to combine the survey datasets to produce coherent and precise income statistics. This has been achieved by integrating the sample designs for the surveys,

[1] Salah.merad@ons.gov.uk

whilst harmonising key components of the survey questionnaires. The need to improve coherence and precision of income statistics has been highlighted in the LCF National Statistics Quality Review (NSQR) (see NSQR Series (2) Report Number 3 - Living Costs and Food Survey (2016)), the UK Statistics Authority's income monitoring review (see UK Statistics Authority Monitoring Review (2015)) and in the context of European statistics for monitoring poverty and social exclusion at a regional level.

In Section 2 we present the integrated sample design, including an assessment of some of the changes in stratification on LCF and SLC statistics. In Section 3 we present the method we used to determine the SLC sample size and the regional allocation of the combined LCF and SLC sample. In Section 4 we discuss the challenges of data collection in the field and outline a proposal to retain half of the primary sampling units across consecutive years to improve the efficiency of data collection and the precision of estimates of change. In Section 5 we discuss the next steps and future work.

2. The integrated sample design

2.1 Overview of the integrated design

A key requirement for the new 2017 integrated design is to maintain consistency with the 2016 LCF. This is achieved by adopting s the same two-stage probability sample design, as used on the 2016 LCF and SLC surveys, with postcode sectors being the primary sampling units (PSU). The PSUs are selected using probability-proportional-to-size (PPS) systematic sampling from a sorted list of postcode sectors. Up to 2016, the PSUs were stratified broadly by the NUTS1 (Nomenclature of Units for Territorial Statistics level 1 international classification) geography, consisting of the 9 English regions, Scotland and Wales that are further broken down by metropolitan/non-metropolitan for certain regions, with London split into 4 quadrants, resulting in a total of 26 major strata. Within each major stratum, the PSUs are sorted by Census factors, hence providing an implicit stratification. The Census factors used by LCF are: the percentage of households in the National Statistics Socio-economic Classification (NS-SEC) classes 1-3) and the percentage of households without a car. The Census factors used in FRS, from which the SLC panels originate, are: the percentage of households in NS-SEC classes 1-3, the percentage of economically active adults and the percentage of male unemployed adults.

As one of the objectives of an integrated design is to improve the quality of regional estimates below the NUTS1 geography, and to meet EU requirements on precision at the NUTS2 level geography, we decided to use NUTS2 to define the major strata. There are between 1 and 5 NUTS2 regions in each of the 9 English regions, Wales, Scotland, adding up to a total of 39 regions.

We also decided to use the Census factors used in LCF in the integrated design. We will discuss the impact of the change in regional stratification and Census factors below.

Under the new design, the LCF component continues to have 638 PSUs, with 18 addresses in each (11,484 in total), and be self-weighted (as in the 2016 sample). For SLC, we continue to select 15 addresses per PSU, as in 2016, but the number of PSUs is determined with respect to precision requirements by NUTS2; the method is presented in Section 3.

The combined set of PSUs is selected jointly within each major stratum using PPS systematic random sampling, as described above. The selected PSUs are then allocated systematically at random within each regional stratum between LCF and SLC, according to the required stratum sample sizes for each component survey. The allocation is also random over the months of the year, which results in a uniform allocation throughout the year. Hence, the joint selection and random allocation provides a representative sample by region and month for the combined sample, as well as for the component survey samples.

2.2 Assessing the change in stratification

To assess the change in stratification, as given by the geography that defines the major strata and the Census factors, we fitted regression models for key statistical outputs for the SLC and LCF. For SLC, the "At Risk Of Poverty or Social Exclusion (AROPE)" rate per PSU was modelled, where the model covariates are some specified stratification variables.

To obtain a sorted list of PSUs that is robust, we decided to use only two Census factors. Preliminary models indicated that NS-SEC1-3 is by far the best predictor of AROPE, so we wanted to identify which other factor to add; the major stratifier was set to NUTS2. The other Census factors that we considered are:

- HNOCAR - Percentage of households with no car (car ownership)
- ADECACT - Percentage of economically active adults
- MUEMPRT- Percentage of male unemployed adults
- P60PLUS - Percentage of people aged 60+
- HRENTSOC - Percentage of social rented households
- POPDEN - Population density

Table 1 shows the model R^2 for each model using SLC data from 2012, 2013 and 2014. We can see that overall MUEMPRT performs marginally better than HNOCAR and the other factors. However, as the difference is negligible, for continuity with the LCF stratification, we decided to use HNOCAR.

Table 1. R^2 values of regression models of AROPE using SLC data

Model	2012	2013	2014
2016 geography, NS-SEC1-3, HNOCAR	0.29	0.28	0.28
2016 geography, NS-SEC1-3, ADECACT	0.28	0.25	0.27

2016 geography, NS-SEC1-3, MUEMPRT	0.31	0.28	0.28
2016 geography, NS-SEC1-3, P60 PLUS	0.28	0.24	0.26
2016 geography, NS-SEC1-3, HRENTSOC	0.29	0.27	0.26
2016 geography, NS-SEC1-3, POPDEN	0.27	0.26	0.27

To assess the change in the geography that defines the major strata on SLC and LCF estimates, we compared models that include either the 2016 geography or the NUTS2 geography in addition to the two Census factors NSSEC and HNOCAR.

For LCF, the models were fitted to average income and expenditure per PSU in LCF data from 2012, 2013 and 2014, after variable transformation to reduce the skewness of the distributions. Because the NUTS2 geography has more regions than the 2016 geography, we compared the model adjusted R^2 values.

As can be seen from Table 2, NUTS2 performs marginally better for SLC in 2012 but is similar to the 2016 geography in 2013 and 2014.

Table 2. Modelling AROPE using SLC data - Adjusted R^2 values

Model	2012	2013	2014
2016 geography, NS-SEC1-3, Car Ownership	0.23	0.24	0.24
NUTS2, NS-SEC1-3, Car Ownership	0.25	0.24	0.24

Table 3 shows that NUTS2 performs similarly to, or marginally better in, LCF than the 2016 geography, which could lead to an improvement in the precision of LCF estimates, albeit a small one.

Table 3. Modelling average expenditure using LCF data - Adjusted R^2 values

Model	2012	2013	2014
2016 geography, NS-SEC1-3, Car Ownership	0.34	0.34	0.34
NUTS2, NS-SEC1-3, Car Ownership	0.35	0.35	0.34

The integrated sample design where the PSUs are stratified by NUTS2 and sorted by the Census factors NS-SEC and car ownership has a stratification that is of similar efficiency to the 2016 designs for both LCF and SLC; moreover, it allows us to specify the required sample sizes at the NUTS2 level to have control over the achieved precision at this level. We next describe the method used to determine the sample sizes and how the sample is allocated between the two surveys.

3. The method for sample size calculation

EU-SILC precision requirements are specified by Eurostat for the income and living conditions domain (at the national and regional level). The requirements, which are specifically for the AROPE variable, are expressed through the inequality

$$s.e.(\hat{p}) \leq \sqrt{\frac{\hat{p}(1-\hat{p})}{a\sqrt{N}+b}} \quad (1)$$

where *N* is the population to which the survey refers and *a* and *b* are specified parameters and $\hat{p}$ is the estimator of the proportion of interest p and $s.e.(\hat{p})$ denotes its standard error estimator. The parameter values for *N*, *a* and *b* are detailed in Table 4.

Table 4. Parameter estimates for the income and living conditions domain

Requirement	*p*	*N*	*a*	*b*
1	Ratio of AROPE count to population	Number of private households in the country in millions and rounded to 3 decimal digits	900	2600
2	Ratio of AROPE count to population in each NUTS2 region	Number of private households in the NUTS 2 region in millions and rounded to 3 decimal digits	600	0

Now, ignoring the finite population correction factor, the standard error of the estimator $\hat{p}$ under a complex design, $se_{complex}$, is given by

$$se_{complex}(\hat{p}) = \sqrt{deff}\sqrt{\frac{\hat{p}(1-\hat{p})}{n-1}} \quad (2)$$

where $deff$ is the design effect and n is the achieved sample. Hence, if we denote a NUTS2 region by *h*, the right-hand sides of inequality (1) and equation (2) may be compared to show that the effective sample size, n_{eff} , in stratum h is

$$n_{eff,h} = \frac{n_h}{deff_h} \cong a\sqrt{N_h} + b. \quad (3)$$

Therefore, the required achieved sample in stratum h is given by

$$n_h = deff_h\big(a\sqrt{N_h} + b\big). \tag{4}$$

The effective sample sizes were calculated for the years 2012 to 2014 but varied very little; we used the 2013 values in the calculations.

The design effect is a function of the achieved samples of the stratum PSUs and the level of homogeneity in the PSUs as given by the intracluster correlation coefficient. An approximation of the design effect is given by

$$deff_h \approx 1 + (\overline{m}_h - 1)\rho_h \tag{5}$$

where $\overline{m}_h$ is the average number of observations in a primary sampling unit and ρ_h is the intracluster correlation coefficient (see Kish (1965)).

The values of the intracluster correlation coefficients were calculated using SLC data from the years 2012 to 2014. This required first the calculation of the NUTS2 standard errors for AROPE estimates.

The AROPE variable is a complex, non-linear population parameter, being a function of three other variables:

- At Risk of Poverty Rate (ARPR60) – proportion of people in poverty (defined as 60% of the median income)
- Severe Deprivation (SD) – proportion of people lacking four or more items from a list of nine
- Low Work Intensity (LWI) – proportion of people who live in households where working age members (jointly) work less than 20% of the months of the year

The AROPE rate is the proportion of people who fall into any of the three other variables.

For the ARPR60 and AROPE variables, which are functions of quantiles, the linearisation approximation method for variance estimation proposed in Deville (1999) was used. In comparison to the Taylor linearisation method one might use to calculate standard errors (Särndal et al, 1992), the Deville method correctly accounts for the fact that the median used in the calculation of AROPE is an estimate from survey data. This method requires the derivation of an influence function; see Berger and Skinner (2003) and Osier (2009) for more detail on the derivations.

The NUTS2 standard error estimates from the Deville linearisation method were compared with those computed using the Taylor linearisation approach and were found to be similar.

The design effects were derived and Equation (5) was then used to calculate the intracluster correlation coefficients for each of the years 2012, 2013 and 2014. They were found to vary a lot between NUTS2 areas and within the same NUTS2 area over the three years.

To determine the required achieved sample size for each NUTS2 area, we needed to estimate the design effect for AROPE under the new integrated design; its structure is illustrated in Figure 1, where the core questions include the harmonised household finance and living conditions questions. These questions are also included in SLC waves 2 to 6. The achieved sample sizes per PSU were computed for LCF and SLC wave 1 assuming response rates of 2016. Even though LCF has a larger sample of addresses per PSU than SLC (18 addresses against 15), the achieved sample sizes are similar because of the better response rate in SLC. Also, the achieved sample sizes in SLC decrease over the waves because of attrition. To compute the design effect under the new design, we decided to be conservative; hence, we used the achieved sample size in wave 1 of SLC, which represents the highest expected achieved sample per PSU, and the maximum value of the intracluster correlation coefficient over the three years in each NUTS2 region. On the other hand, we allowed the standard error to exceed the EU threshold by up to 50%. This ensures that, under the worst case scenario, the precision of the NUTS2 estimates will not be too far from the specified EU precision requirement. Let n_h^* denote the resulting total achieved sample size in stratum h.

Figure 1. The new integrated LCF-SLC survey design

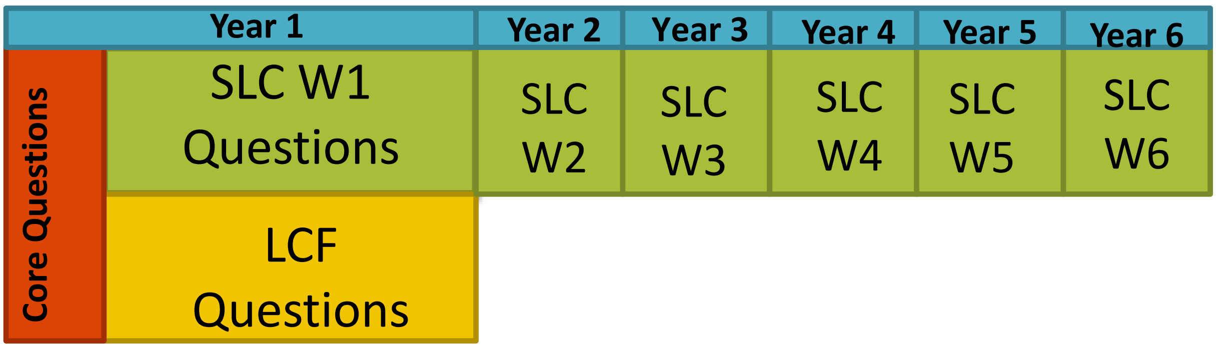

Sample Allocation

The data used to estimate the AROPE rate will come from both LCF and SLC waves 1 to 6. Once the required achieved sample sizes are determined for each NUTS2 region, we need to allocate the sample across the primary sampling units, for the SLC and LCF modules.

To maintain the quality of LCF outputs, the LCF sample is apportioned between the NUTS2 regions proportionally to their size; let $n_{I,LCF,h}$ be the number of LCF PSUs in stratum h. The expected achieved LCF sample size is then given by

$$n_{LCF,h} = n_{I,LCF,h} \times \bar{m}_{r,LCF,h},$$

where $\bar{m}_{r,LCF,h}$ is the expected number of responding households to LCF in a primary sampling unit in stratum *h.*

The required achieved sample size for the six waves of SLC is then given by

$$n_{SLC,h} = n_h^* - n_{LCF,h}$$

The required number of PSUs for the SLC component is then given by

$$n_{I,SLC,h}^* = \frac{n_{SLC,h}}{\bar{m}_{r,v1-v6,h}},$$

where $\bar{m}_{r,v1-v6,h}$ is the expected total number of responding households across the six waves. We used past data to estimate the response rate at wave 1 and attrition rates between consecutive waves, between wave 1 and wave 4. Attrition rates after wave 4 we assumed to be equal to the wave 3 to wave 4 rate.

Applying the method we have described, we found that the sum of the required numbers of SLC PSUs across all NUTS2 regions is 378, which is just over half the number of SLC PSUs under the 2016 design.

These calculations are based on the assumption that we have six SLC waves selected under the new design, which will hold in six years' time. However, we would like to achieve these levels of precision starting from 2018. In this year, although SLC will have six waves, only the first two will be based on the new design. We revised the allocation by performing the following two steps:

(a) We calculated the expected achieved sample sizes is each NUTS2 in 2018; the large NUTS2 regions were found to have an achieved sample size above the required minimum, but in the small NUTS2 areas, the expected achieved sample sizes were below the required limits.

(b) The allocation of each NUTS2 region where the expected 2018 achieved sample size was found to be below the required limit was increased so that the revised expected achieved sample would be above the minimum.

The total number of SLC PSUs increased from 378 to 640, which is similar to that in the 2016 SLC design, but the distribution of the sample over the NUTS2 regions is different: the sampling fractions across the NUTS2 regions vary under the new design.

4. Practicalities in data collection

As SLC is longitudinal and the number of waves will increase from four to six, the number of households to be interviewed per PSU will decrease across the waves because of attrition as very few PSUs will be selected in consecutive years if the samples are selected independently from year to year. Assuming similar attrition rates to 2016, it is expected that the sample size per PSU will be very small after wave 3. Figure 4 shows the expected number of households, per postcode sector, over time. The number of issued cases if we do not retain clusters is given by the individual year rows; it is below 6 from wave 4, down from 15 in wave 1. Sending interviewers to any particular PSU to interview only a small number of households would not be efficient for the field force. To address this, we consider retaining half of the PSUs selected in one year for the following year (this method is already adopted on the Department for Work and Pensions' Family Resources Survey). SLC PSUs will then have larger sample sizes than without PSU retention. Figure 2 shows that only a sixth of the LCF PSUs will have fewer than seven addresses (only those cases on Wave 6 interview) whereas without PSU retention the number of PSUs containing fewer than seven households would be over half (those on interviews in waves 3, 4, 5 and 6).

Figure 2. Expected number of issued households in any given postcode sector for SLC

	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7
Year 1 of selection	Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6	
Number of issued cases	15	8	6	5	4	3	
Year 2 of selection		Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6
Number of issued cases		15	8	6	5	4	3
Total number of issued cases per cluster	15	23	14	11	9	7	3

5. Future Work

As data collection is currently in progress, the weighting strategies for this new integrated survey will need to be developed as well as quality assurance procedures. The integrated sample design with a six-wave SLC will lead to improved region estimates. However, the six panels will be selected from populations spanning a wide time period, which results in differential coverage bias between the panels. In the future, we will examine how this affects the ways the panels are combined to produce the pseudo cross-sectional weights.

The retention of half the PSUs from one year to the next should make data collection more efficient but the overall impact on precision will need to be assessed. The retention of PSUs should lead to positive covariates between the estimates at two consecutive time points, which should lead to an improvement in the precision of estimates of change. On the other hand, the level estimates will have slightly larger variances because of the covariance between the waves. The trade-off between the impact of PSU retention on estimates of change and level estimates needs further investigation.

In the next steps of the integration, we will be seeking to add another financial survey (WAS). We will also consider the use of administrative data to further improve the design of the sample, by oversampling high income households, for example, in weighting and to replace some questions, and hence reduce burden.

References

Berger, Y. and Skinner, C. (2003). Variance estimation for a Low-Income Proportion, *Social Statistics Research Centre*, pp 1-20.

Deville, J. C. (1999). Variance estimation for complex statistics and estimators: linearization and residual techniques. Survey Methodology, 25, 193-203.

Eurostat (2015). Draft framework regulation on IESS (Integrated European Social Statistics), pp 20-21.

Kish, L. (1965) *Survey Sampling*. New York, Wiley.

NSQR Series (2) Report Number 3 - Living Costs and Food Survey (2016), available at https://www.ons.gov.uk/peoplepopulationandcommunity/personalandhouseholdfinances/incomeandwealth/methodologies/nsqrseries2reportnumber3livingcostsandfoodsurvey.

Osier, G. (2009). Variance estimation for complex indicators of poverty and inequality using linearization techniques, *Survey Research Methods*, 3(3), pp 167-195.

Särndal, C. E., Swensson, B., & Wretman, J. (1992). Model Assisted Survey Sampling. New York: Springer Series in Statistics.

UK Statistics Authority Monitoring Review (2015) The Coherence and Accessibility of Official Statistics on Income and Earnings, available at https://www.statisticsauthority.gov.uk/archive/assessment/monitoring/monitoring-reviews/monitoring-review-1-2015---the-coherence-and-accessibility-of-official-statistics-on-income-and-earnings.pdf.

The Impact of Moving Holidays on Official Statistics Time Series

Bethan Russ[1], Tariq Aziz[2]

Abstract

A major challenge faced when seasonally adjusting time series is accounting for annual events that move dates from one calendar year to the next, for example, Easter. If these events are not accounted for appropriately it will impact on the estimation of seasonal factors, and leave systematic calendar related effects in the seasonally adjusted series.

Currently the Time Series Analysis Branch (TSAB) tests for Easter effects and, if identified, estimates and removes them as part of seasonal adjustment. This method assumes that daily activity changes by a fixed amount or proportion for a given number of days before Easter Sunday and remains at this level until Easter Saturday.

There are other moving holidays celebrated in the UK, which may have an impact on time series despite not being public holidays. These are Chinese New Year, Ramadan, Eid al-Fitr and Eid al-Adha. Currently these holidays are not adjusted for in any seasonally adjusted time series published by the Office for National Statistics (ONS).

TSAB has undertaken research to test alternative windows for Easter effects and whether other moving holidays have identifiable effects on ONS time series.

This paper will present findings from this research on a range of ONS time series.

1. Background

A moving holiday is defined as a calendar event that moves between periods, where a period can be a week, month or quarter, from one year to the next. (*In this research a year is defined as a solar year, based on dates of the Gregorian calendar.*) Well-known examples of these events are Easter, Chinese New Year, Ramadan, Eid al-Fitr and Eid al-Adha - all of which are widely celebrated in the UK.

As with many annual events, a moving holiday can cause seasonality in a time series. However, since these events do not occur in the same period each year, without an appropriate adjustment the estimation of the seasonal component may become distorted and the resulting seasonally adjusted series may contain systematic variation due to the arrangement of the calendar.

At present TSAB only tests and adjusts, where appropriate, ONS time series for an Easter effect, where activity throughout the Easter period is assumed to be of a fixed amount or proportion and only pre-Easter Sunday windows are considered (*see Section 2.1*). Any effect resulting from other moving holidays, alternative windows or amounts which are not constant are not currently accounted for.

[1] Office for National Statistics; bethan.russ@ons.gov.uk

[2] Office for National Statistics; tariq.aziz@ons.gov.uk

Moving holiday effects are estimated through the inclusion of appropriate regressors in a regARIMA model (*for more details see Section 1.3*). This research project aims to develop, and test, alternative regressors. These regressors aim to account for the following:

- UK moving holidays:
  - Chinese New Year
  - Eid al-Adha
  - Eid al-Fitr
  - Ramadan
- Different lengths and positions of windows around the moving holiday
- Alternative shapes of regressor about the moving holiday (*ie an effect which is not assumed to be a fixed amount or proportion*).

1.1 Moving Holiday Dates

The holidays analysed in this research are Easter, Chinese New Year, Ramadan, Eid al-Fitr and Eid al-Adha. Subsections 1.1.1, 1.1.2 and 1.1.3 describe these moving holidays in more detail. Subsection 1.1.4 contains a table, showing the changing dates of each event since 2001.

Information on the dates of the moving holidays, summarised below, was taken from the *Time and Date* website.

1.1.1 Easter

Each calendar year the date of Easter Sunday changes from the previous year. The date of Easter Sunday can fall anywhere between March 22nd and April 25th since the date of Easter is defined as;

"*Easter Sunday falls on the first Sunday after the first full moon following the northern spring equinox.*"

Due to this pattern, the Easter public holiday period may fall wholly in March, wholly in April or start in late March and end in early April. Similarly, the Easter holiday period may fall wholly in Quarter 1, wholly in Quarter 2 or start in late Quarter 1 and end in early Quarter 2. The movement of the Easter period can have a direct impact on aggregated time series data and is one of the few well-known calendar effects accounted for in official statistics.

1.1.2 Chinese New Year

Chinese New Year, which in modern Chinese translates as the `*Spring Festival*', is celebrated at the turn of the traditional lunisolar Chinese calendar. Chinese New Year festivities traditionally run until the 15th day of the first Chinese calendar month. The date for the first day of Chinese New Year can fall anywhere between 21st January and 20th February, on the Gregorian calendar, since it is defined as;

"*The first day of the New Year falls on the second new moon after the winter solstice.*"

Due to this pattern, the Chinese New Year period may fall wholly in February or start in late January and end in early February. This moving holiday is not currently accounted for in official statistics, but will be considered as part of this research project.

1.1.3 Islamic Calendar Events

The three most well-known Islamic events are Ramadan, Eid al-Fitr and Eid al-Adha. Ramadan is the ninth month in the Islamic calendar, not an event in itself, and lasts for an entire lunar cycle. Eid al-Fitr marks the end of Ramadan and starts on the first day of Shawwal, the tenth month in the Islamic calendar, and lasts for 3 days. Eid al-Adha marks the end of the Hajj pilgrimage, starting on the tenth day of Dhu'l-Hijjah, the twelfth month in the Islamic calendar, and lasts for 3 days. Each of these three holidays are associated with specific dates in the Islamic calendar, and hence will move from year-to-year in relation to the Gregorian calendar.

Unlike the solar Gregorian calendar, which ONS time series are measured upon, the Islamic calendar is a 12-month lunar calendar which is based upon cycles of the moon's phases. The lunar Islamic calendar is 10 to 12 days shorter than a solar year, hence any Islamic events migrate throughout the seasons and occur 10 to 12 days earlier, in the Gregorian calendar, each year.

1.1.4 Table of Dates

Table 1 shows the dates of each of the above moving holidays between 2007 and 2016. This table has been included to illustrate how much the holidays can move over a small period of time.

Table 1: Changing dates of moving holidays[1], 2007-2016.

Year	Easter	Chinese New Year	Ramadan	Eid-al-Fitr	Eid-al-Adha
2007	April 8th	February 18th	September 13th	October 13th	December 20th
2008	March 23rd	February 7th	September 2nd	October 2nd	December 9th
2009	April 12th	January 26th	August 22nd	September 21st	November 28th
2010	April 4th	February 14th	August 11th	September 10th	November 17th
2011	April 24th	February 3rd	August 1st	August 31st	November 7th
2012	April 8th	January 23rd	July 20th	August 19th	October 26th
2013	March 31st	February 10th	July 9th	August 8th	October 15th
2014	April 20th	January 31st	June 29th	July 29th	October 4th
2015	April 5th	February 19th	June 18th	July 18th	September 24th
2016	March 27th	February 8th	June 7th	July 6th	September 13th

[1] For moving holidays lasting more than one day, this date represents the first day of the event.

1.2 Moving Holiday Windows

A common issue in the analysis of moving holidays is finding the optimal window over which to analyse the impact of the event. The term *window* refers to the period associated with a moving holiday over which an effect may be present. A window can be of any length; for example a day, a week or a month, and can fall before, during or after an event.

Choosing the optimum length and placement of a window will involve some prior knowledge of the series and of the moving holiday. Knowledge of the time series will aid the analysis, allowing the user to predict whether or not a moving holiday will have an impact on the measured estimate and whether this impact is likely to be in the run up to the event, during the event or afterwards.

This research analyses a range of ONS time series and as such is impractical to use prior knowledge of each of the series to determine an appropriate window. Several windows will be tested during the analysis of moving holidays. The majority of windows will fall for some number of days before an event, although many will also include the event itself during the window or include a period after the event.

1.3 Moving Holiday Regressors

A standard way to account for moving holidays in seasonal adjustment is to include regressors in the regARIMA model used for seasonal adjustment. A regARIMA model is a regression model where the errors follow an ARIMA process (*more information can be found in the X-13ARIMA-SEATS manual [Bureau, 2016]*). The user must build an appropriate regressor describing the shape and window of the moving holiday. In simple terms, a regressor is a variable that contains information about the possible influence of an event on a time series.

There are three main shapes of regressors considered in this type of research:

- Constant - *known generally as Shape 0*
- Linear - *known generally as Shape 1*
- Quadratic - *known generally as Shape 2*.

Each type of regressor assumes that the effect caused by a moving holiday follows some given pattern across a given period, and that this effect is of stable size and direction from one holiday to the next. In simple terms, it assumes that the effect is not random and exists in every cycle.

A constant regressor assumes that for all days, within the specified window for an event, activity has changed by a fixed amount or proportion in comparison to days that lie outside the window. In mathematical terms, this assumes that every k^{th} day in the window (of n days) has an effect equal to xk^0 (or x), where x is some arbitrary amount or proportion. An example of such a regressor can be seen in Figure 1. Figure 1 shows that for all days in the window, from 1 to n, there is a constant effect but that outside the window there is no effect. It should be noted that in this scenario the moving holiday can fall on any of the n days, allowing for a pre-event effect, post-event effect or an effect either side of the event.

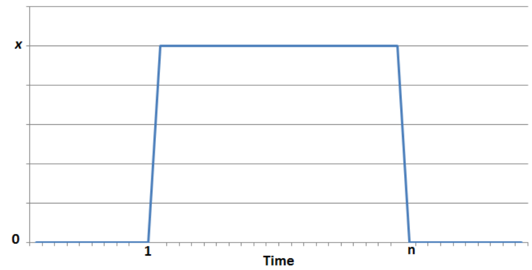

Figure 1: Example of a constant regressor.

A linear regressor assumes that, within the specified window for an event, activity increases linearly over time (ie that activity increases by the same amount between days). In mathematical terms, this assumes that every k^{th} day in the window (of n days) has an effect equal to xk^1 (or xk), where x is some arbitrary amount or proportion. Two examples of such a regressor can be seen in Figure 2. Figure 2 (a) and (b) show that for all days in the window, from 1 to n, there is a linear effect but that for days that fall outside the window there is no effect. It should be noted that in scenario (a) the moving holiday can fall on any of the n days, allowing for a pre-event effect, post-event effect or an effect either side of the event. In scenario (b) the moving holiday will generally fall at the midpoint of the n days, showing a symmetric, linear effect pre and post the event. A non-symmetric linear effect is also possible, both in the case of having unequal gradients for the slope and having the effect fall not at the midpoint.

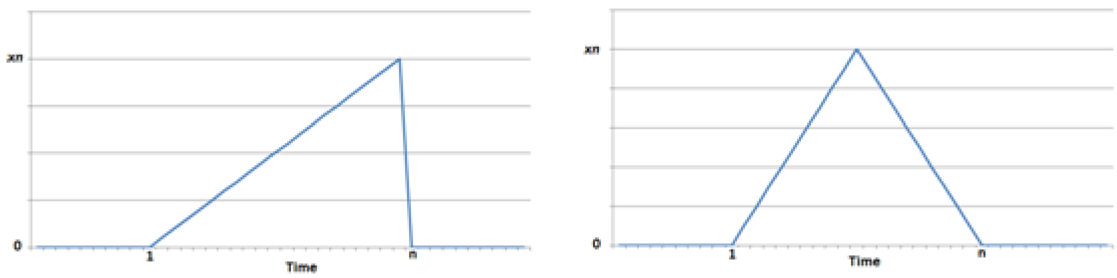

(a) (b)

Figure 2: Examples of a linear regressor.

A quadratic regressor assumes that, within the specified window for an event, activity increases quadratically over time (ie that activity increases by a non-linear amount between days). In mathematical terms, this assumes that every k^{th} day in the window (of n days) has an effect equal to xk^2, where x is some arbitrary amount or proportion. Two examples of such a regressor can be seen in Figure 3. Figure 3 (a) and (b) show that for all days in the window, from 1 to n, there is a quadratic effect but that outside the window there is no effect. It should be noted that in scenario (a) the moving holiday can fall on any of the n days, allowing for a pre-event effect, post-event effect or an effect either side of the event. In scenario (b) the moving holiday will generally fall at the midpoint of the n days, showing a symmetric effect pre and post the event. A non-symmetric linear effect is also possible, both in the case of having unequal gradients for the slope and having the effect fall not at the midpoint.

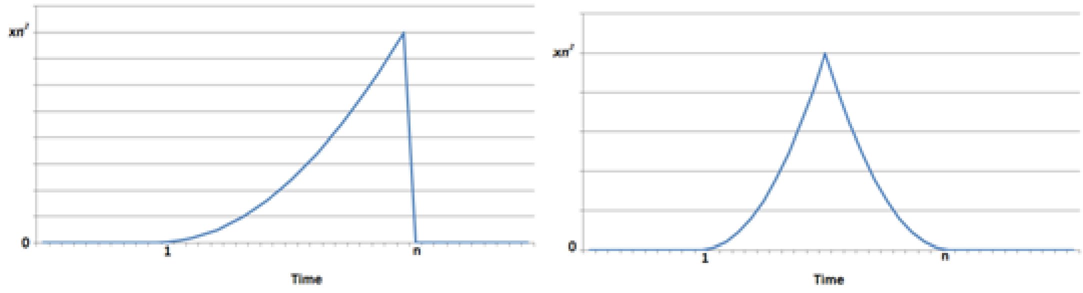

(a) (b)

Figure 3: Examples of a quadratic regressor.

1.4 Series to be Analysed

Every year TSAB reviews the seasonal adjustment specifications of thousands of ONS time series. A subset of these was chosen for this analysis as it was impractical to run each regressor on all time series.

A selection of time series from a range of ONS outputs were used for this analysis. These included Construction, Index of Production, Index of Services, Trade in Goods and Service and Leisure and Tourism. For the Easter regressors a selection of 599 time series were used. For the other holidays, 290 were used.

These series were selected as many are already adjusted for Easter effects. It was thought that many of these series could be affected in different ways by moving holidays, for example:

- A decrease in output or trade in certain goods and services due to workers taking leave over the holiday
- Increased output and trade of goods and services associated with each holiday
- More visits abroad around holidays, particularly Easter which falls during the school holidays.

While this paper presents results on whether ONS time series favour the inclusion of different holidays and different regressors, it does not present results on the direction of the effect of the moving holiday for different time series.

1.5 Literature Review

The analysis of moving holidays in time series is not a new concept and methods for detection of calendar effects, with respect to the Gregorian calendar, are widely used. Methods have been developed in official statistics, academia and economics. This section provides a summary of a literature review of current methods that was undertaken as part of this research.

Pfeffermann and Fisher proposed a new method for festival and working day adjustment for economic time series. In their research they used regression modelling to assess the relationship between the irregular components of time series data and the moving date of the holiday, in respect to the Gregorian calendar. They constructed regressors for the lunar calendar events in the Jewish calendar and for Chinese New Year. In this study the irregular component was calculated through an appropriate seasonal adjustment program and extreme values were excluded before calendar adjustments were made. The results from this study were not highly significant. The explanation of variability, through the regression model with calendar regressors, was low but the inclusion of these regressors did reduce the irregularity of time series without affecting the seasonal pattern and the trend [Pfeffermann and Fisher, 1982].

Riazuddin and Khan used a regARIMA modelling approach for the detection and forecasting of Islamic calendar effects, in financial time series in Pakistan. They applied the concept of fractional indicator variables and used dummy variables to account for the effects on the Gregorian calendar months. The forecast evaluation of the ARIMA model, with and without including the significant regressors for Islamic calendar months, was carried out. It was found that the inclusion of regressors significantly contributed to the improvement of forecast performance for money circulation data [Riazuddin et al.. 2002].

Shuja, Lazim and Wah proposed using regARIMA modelling to detect the lunar calendar effect of Chinese New Year, Eid al-Fitr and Deepavali, on Malaysian economic time series. They constructed the regressors, built with respect to the Gregorian calendar, investigating proximity effects of windows of different lengths. The results of this study showed that the investigated festivals had both a stimulating and reducing effect for the periods of the holidays. The methods proposed to account for the holidays were able to significantly eliminate any moving holiday effects [Shuja et al., 2007].

Lin and Liu used an ARIMA modelling approach to model lunar calendar effects in Taiwan. In this study ten economic time series were analysed and results showed that the inclusion of lunar calendar event's regressors (Mid-Autumn and Dragon Boat Festival) can effectively control the impact of moving holiday effects on the seasonal adjustment process. Significant moving holiday regressors for unemployment series showed that the seasonal factors cannot be consistently estimated unless the effects are controlled for [Lin et al., 2002].

The Australian Bureau of Statistics (ABS) investigated proximity effects using a regARIMA approach for the estimation of effects of Chinese New Year and Ramadan. Their research focused on the Overseas Arrivals and Departure time series [AEI]. ABS also did an investigation into alternative Easter regressors, other than those built into the then common software X-12-ARIMA [Findley et al., 1998]. They looked at implementing alternative windows, accounting for the Australian Easter holiday period, in Australian Total Retail Turnover series. Their results showed that when an Easter proximity effect was present their regressors gave a substantial improvement in the seasonally adjusted estimates, over not adjusting for Easter. The new regressors provided additional gains over default holiday regressors in X-12-ARIMA, capturing the unique characteristics of the Australian Easter holiday period [Leung et al., 1999].

Using this previous research as a benchmark, TSAB has built a number of regressors to test in this study. The proposed study will use different window lengths for the construction of the regressors, based on knowledge of the nature of each of the moving holidays. The proposed study will also use two shapes of regressor, Shape 0 and Shape

1, something that has not been investigated in previous work. (*See Sections 2.2 and 2.3 for more information on the methodology proposals.*)

2. Methodology

2.1 Current Methodology

To account for the moving date of Easter Sunday TSAB currently tests and applies, if appropriate, one of three different Easter regressors. These regressors have been built into X-13ARIMA-SEATS [Bureau, 2016] to account for the fact that Easter can fall in one of two months (March or April) in a monthly series or in one of two quarters (quarter 1 or quarter 2) in a quarterly series. (*Note: X-13ARIMA-SEATS is the recommended software for time series analysis across the Government Statistical Service (GSS)*.)

The three Easter regressors currently used in production are Easter[1], Easter[8] and Easter[15], the standard built-in regressors which account for the North American Easter holiday period. The Easter[1] regressor is used to account for Easter Saturday only. The Easter[8] regressor accounts for Easter Saturday and the 7-day period that precedes this. The Easter[15] regressor accounts for Easter Saturday and the 14-day period that precedes this. Each of the three regressors is a Shape 0 and assumes that activity changes by a fixed amount or proportion across the period it covers.

2.2 New Methodology - Easter

Whilst the built-in regressors are adequate for accounting for Easter, and the days preceding this, they do not adequately account for the Easter period in the UK. In the UK there are bank holidays either side of Easter Sunday, on Good Friday and Easter Monday, and school holidays which generally fall both sides of Easter Sunday, dependent on the date of Easter Sunday. The bank holiday and school holiday periods could have an effect on movements in a time series. Two UK specific Easter regressors were constructed and tested on a number of high profile series.

As part of the exploratory analysis, to identify whether or not an Easter effect was present, Easter proximity charts were constructed for each series analysed. These charts were constructed by mapping the irregular component from the time series which has been decomposed without accounting for an Easter effect, for both March and April, against the date of Easter Sunday in a given year.

An example of one of these series, which shows a possible Easter effect, can be seen in Figure 4 below. This chart is for Index of Production: Manufacture of other non-metallic mineral products (other). The most notable feature of Figure 4 is that when Easter Sunday falls in March, the value of the March irregular is negative and the April irregular is positive. The magnitude of the irregular decreases as the date of Easter Sunday moves closer to April. When Easter Sunday is in April there is no obvious pattern to the irregular component.

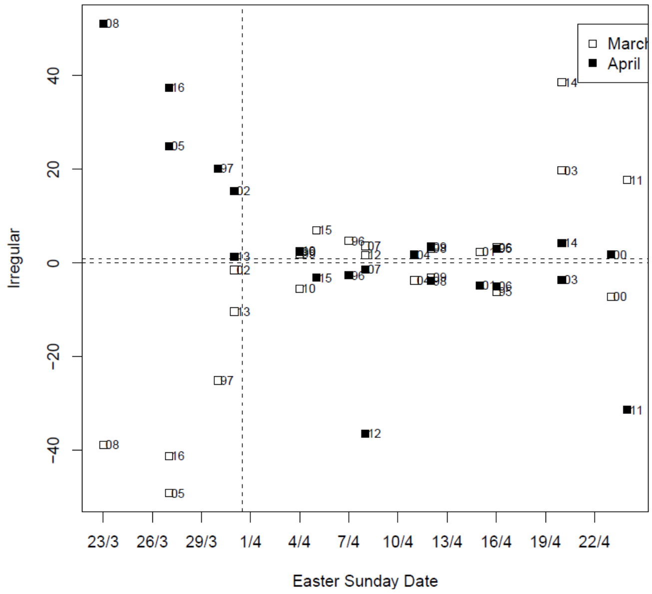

Figure 4: Easter proximity chart.

Following the exploratory analysis, two UK Easter regressors were constructed and tested, UKE[4] and UKE[12]. The UKE[4] regressor covers the period from Good Friday to Easter Monday, accounting for the entire UK bank holiday period. The UKE[12] regressor runs from the Monday before Easter Sunday until the Friday following it, a period which is similar to UK school holidays. Both the new UK Easter regressors assume that activity changes by a fixed amount or proportion over these periods and have been built as a Shape 0 to account for this.

The performance of the new UK Easter regressors was assessed on 599 current ONS time series. Each series was modelled independently with no Easter effect, the standard regressors in X-13ARIMA-SEATS; Easter[1], Easter[8], Easter[15], and the UK specific regressors; UKE[4] and UKE[12].

The Akaike information criterion-corrected (AICC) value, produced by X-13ARIMA-SEATS, was recorded for each series and each regressor. The regressor resulting in the lowest AICC value was recorded as the most appropriate Easter regressor for that series. The best Easter regressor was then compared against the Easter regressor currently used in the seasonal adjustment of that time series.

It should be noted that the AICC value is a crude estimate of the performance of a regressor. In practice the AICC value would be used in conjunction with further knowledge, further tests and continuity analysis before such regressors are put in the production of official statistics (see Section 4.1 for information on further work).

Extracts of the UKE[4] and UKE[12] regressors can be found in the Appendix. The information on the built-in regressors can be found in the X-13ARIMA-SEATS manual [Bureau, 2016].

2.3 New Methodology – Chinese New Year and Islamic Calendar Events

There are currently no methods in place to account for the moving holidays Chinese New Year, Ramadan, Eid al-Fitr or Eid al-Adha. With these events becoming more prominent in the UK it was important for TSAB to construct regressors and analyse whether or not these events are having an impact on official statistics.

For the analysis of Chinese New Year, Eid al-Fitr and Eid al-Adha 8 regressors were constructed for each moving holiday - where the regressor was centred around day 1 of each moving holiday. The regressors were constructed as two groups of four - with the first four regressors taking the form of Shape 0 and the remaining four Shape 1. Within each group there were the following four windows:

- (7,0) - the 7 days leading up to the event and the event itself
- (14,0) - the 14 days leading up to the event and the event itself
- (7,7) - the 7 days leading up to the event, the event itself and the 7 days following the event
- (14,14) - the 14 days leading up to the event, the event itself and the 14 days following the event

For the analysis of Ramadan three regressors were constructed. The first regressor, Regressor A, was constructed in the form of Shape 0 with the window covering the entire month of Ramadan itself. The second regressor, Regressor B, was also constructed as a Shape 0 and considered the fourteen days leading up to the start of Ramadan, and day 1 of Ramadan only. The third regressor, Regressor C, considered the same window as Regressor B but was constructed as Shape 1. Regressors with a window including dates after the end of Ramadan were not considered, as Eid al-Fitr is the first day after Ramadan and has been considered separately.

For each of the new moving holidays, dates spanning 100 years have been used to construct each of the regressors, centring them to align with the Gregorian calendar. Since the Chinese New Year windows considered can only fall (partially or wholly) in January, February or March, its regressors were centred by subtracting the means of January, February and March from the January, February and March values each year. The regressor is 0 for all other months.

Similarly, since the Islamic calendar events can fall in any month of the Gregorian calendar they were centred by subtracting the total mean of the build-up period from each month of the year. From here regARIMA modelling has been used to model the time series, which included the new regressors to assess the impact of these events. As with the Easter methodology, the AICC was recorded for each model to assess the most appropriate regressor for each time series that was analysed.

Extracts of a selection of the Chinese New Year and Islamic Calendar regressors can be found in the Appendix.

3. Results

3.1 Easter

3.1.1 Regressor Analysis

In total there were 599 time series which were tested for Easter effects. Of these 167 time series, or 28% of them, deemed one of the two UK Easter regressors most appropriate, over the alternatives (which included no Easter effect). Table 2 below gives the high level summary of the counts of which Easter regressor was deemed most appropriate against each of the current regressors.

Table 2: Number of time series preferring each Easter regressor.

	Easter Regressor with the Lowest AICC						
	No Effect	**Easter[1]**	**Easter[8]**	**Easter[15]**	**UKE[4]**	**UKE[12]**	**Total**
Current Regressor							
No Effect	262	26	16	19	20*	23*	376
Easter[1]	25	15	3	4	39	54	140
Easter[8]	7	2	8	0	0	10	27
Easter[15]	15	1	5	23	4	8	56
Total	309	44	32	46	63*	95*	599

*There are 10 time series, currently specifying no effect, that deem UKE[4] and UKE[12] equally most appropriate.

From Table 2 it can be seen that across all time series, those currently specifying the Easter[1] regressor saw one of the two UK Easter regressors outperforming the current regressor most often. There are currently 140 time series that contain an Easter[1] regressor and of these 93 deemed a UK Easter regressor most appropriate. This is approximately two thirds of the Easter[1] series tested. In contrast to this, series currently with no Easter effect included only deemed a UK Easter regressor most appropriate in 14% of cases.

Of the two UK Easter regressors the UKE[12] regressor was deemed the most appropriate most often. Across the 599 time series analysed the UKE[12] regressor was most appropriate in 95 instances and the UKE[4] regressor in 63 instances, in 16% and 11% of time series respectively. There were 10 instances where both the UKE[4] and UKE[12] regressors were both deemed equally most appropriate.

The assessment of which regressor is most appropriate has been made by considering the lowest AICC value only. No analysis has been done to look at the magnitude between the lowest AICC value and the AICC values for the other regressors. This analysis, along with other diagnostics, would be important in determining exactly how well each Easter regressor is performing.

To illustrate these results in a different way a chart has been created to show the difference between the seasonally adjusted series with no Easter regressor, the current Easter[1] regressor and the preferred UKE[12] regressor. Figure 5, to be consistent with Figure 4, has been plotted for Index of Production: Manufacture of other non-metallic mineral products (other). The results displayed in Figure 5 show a minimal difference in the seasonally adjusted values for the two Easter regressors, however a difference can be seen between these and no effect. It is possible that neither of these Easter regressors is fully capturing the proximity effect seen in Figure 4, and so other shaped regressors should be considered in any further analysis.

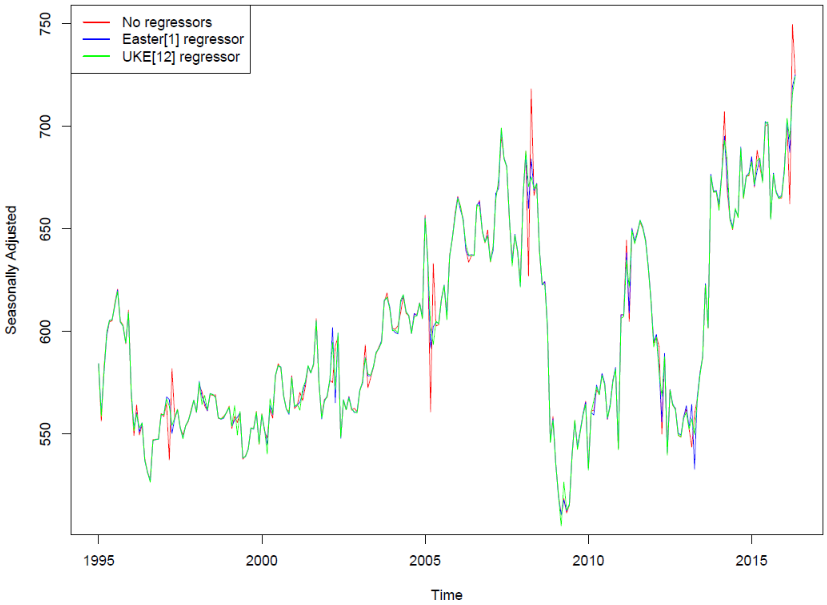

Figure 5: Seasonally adjusted time series, for three types of Easter regressors.

3.1.2 Regressor Analysis

It is difficult to see much visual difference between the seasonal adjustment with the Easter[1] and UKE[12] regressors in Figure 5. To illustrate the differences caused by the two regressors, month-on-month growth rates in the seasonally adjusted time series were calculated for this same series. The aim of this analysis was to compare the growth rate between February and March with that between March and April when different Easter regressors were included. Table 3 shows the growth rate between months, and the seasonally adjusted value, alongside the date of Easter in each year.

Table 2: Month-on-month growth rate (percent) in seasonally adjusted time series, for three types of Easter regressors.

Date of Easter	No Easter		Easter[1]		UKE[12]	
	Feb-Mar	Mar-Apr	Feb-Mar	Mar-Apr	Feb-Mar	Mar-Apr
April 24th 2011	6.0 (644)	-6.2 (604)	5.0 (638)	-4.6 (609)	4.5 (634)	-2.1 (621)
April 8th 2012	-0.9 (59)	-7.2 (550)	-2.8 (581)	-4.4 (555)	-3.0 (579)	-1.9 (568)
March 31st 2013	-1.6 (543)	3.0 (559)	2.0 (564)	-5.6 (532)	0.9 (557)	-1.4 (549)
April 20th 2014	4.6 (707)	-5.4 (668)	2.7 (695)	-3.1 (673)	2.4 (693)	-1.3 (683)
April 5th 2015	2.7 (688)	-1.2 (680)	1.1 (678)	0.8 *684)	1.7 (683)	0.2 (684)
March 27th 2016	-5.7 (662)	13.2 (749)	-2.2 (687)	4.7 (719)	-1.5 (693)	3.2 (715)

The growth rate analysis illustrated in Table 3 is consistent with the results in Section 3.1.1, identifying that an Easter effect is present in this particular series. Generally the growth rate is negative going into the month when Easter falls and positive in the other month. When either Easter[1] or UKE[12] is included the growth rates no longer exhibit a systematic effect in sign based on the date of Easter. The growth rates also become less extreme when including one of the two Easter regressors, for example, when Easter was very early in 2016 the March to April growth rate is 13.2% if Easter is not accounted for. This reduces to 4.7% and 3.2% respectively for the Easter[1] and UKE[12] regressors.

The results in Table 3 show that the growth rates when using the UKE[12] regressor are generally smaller than when using the Easter[1] regressor. This could be showing that the UKE[12] regressor is better adjusting for an Easter effect and thus reducing the irregular component, however further work is needed to understand if this is the case.

3.1.3 Span Analysis

This analysis has used the AICC to determine the most appropriate regressor. In practice there would also be other considerations when deciding which regressors to include. One such consideration would be stability, namely, is this regressor consistently the best when more data points become available? To aid the analysis further span analysis was run on the Index of Production: Manufacture of other non-metallic mineral products (other) series highlighted in Figure 5. The analysis involved removing the final 12 data points and re-running the AICC analysis. This was repeated 4 times, to get 6 years worth of analysis.

The analysis for the 5 new spans of data were consistent with the original analysis, with the AICC for the UKE[12] regressor being lower in each span.

This is a positive indication that the UKE[12] regressor is behaving consistently over time, although the regressor would need to be tested over larger spans and more datasets to confirm this.

3.2 Chinese New Year and Islamic Calendar Events

3.2.1 Regressor Analysis

In total there were 290 time series analysed for each of the further moving holidays. Of these 65 time series (22%) preferred a Chinese New Year regressor, 94 time series (32%) preferred an Eid al-Fitr regressor, 80 time series (28%) preferred an Eid al-Adha regressor and 91 time series (31%) preferred a Ramadan regressor, to no regressor for that event. Each moving holiday was assessed separately and the different windows and shapes were compared to no effect. Table 4 below gives a high level summary of the counts of which regressors were deemed most appropriate, for each moving holiday.

Table 4: Number of time series preferring each moving holiday regressor.

	No Effect	Shape 0				Shape 1			
		(7,0)	(7,7)	(14,0)	(14,14)	(7,0)	(7,7)	(14,0)	(14,14)
Chinese New Year	225	2	9	12	9	13	0	6	14
Eid-al-Fitr	196	5	7	21	12	9	6	4	30
Eid-al-Adha	210	8	2	17	9	11	8	4	21
	No Effect	**Regressor A**		**Regressor B**		**Regressor C**			
Ramadan	199	47		23		21			

To show the impact of moving holiday regressors on time series seasonal adjustment process the seasonally adjusted time series have been plotted, see Figure 6, with and without Chinese New Year and Islamic Calendar regressors. This figure is for International Passenger Survey: Overseas visits to the UK, expenditure. This figure illustrates the results of the comparison of the seasonally adjusted series with no regressor included and moving holiday regressors for Chinese New Year and Islamic Calendar Events. This series has been chosen as it was one of the few that found an effect for all four moving holidays.

The regressors used in Figure 6 are Chinese New Year (7,0) Shape 1, Eid-al-Fitr (14,0) Shape 0, Eid-al-Adha (14,0) Shape 0, and Ramadan regressor A (Shape 0, full month). The results displayed in Figure 6 show a minimal difference in the seasonally adjusted values for no effect versus the moving holiday regressors. This could suggest that the effect is genuinely small, or that the analysed regressors are not fully capturing the shape of the effect, despite being preferred in a number of time series.

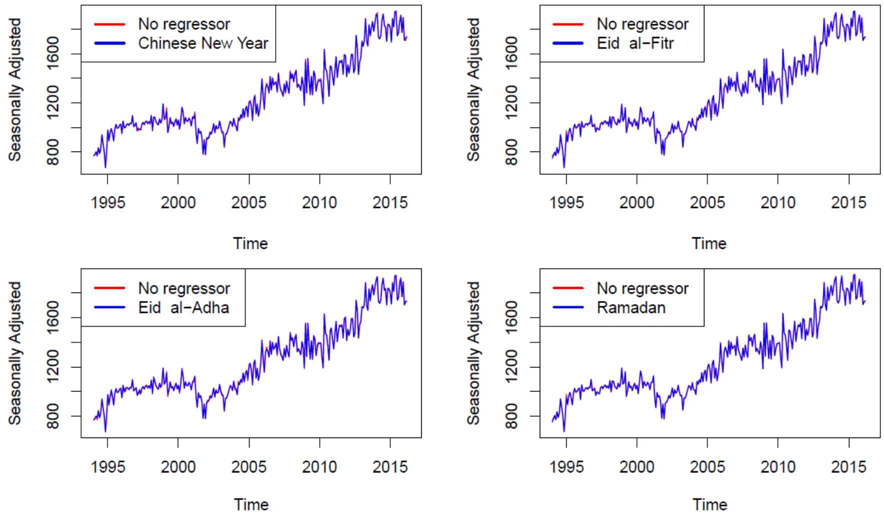

Figure 6: Seasonally adjusted time series, no regressor versus moving holiday regressor.

3.2.2 Growth Rate Analysis

Month-on-month growth rate in the seasonally adjusted time series were calculated for the International Passenger Survey: Overseas visits to the UK, expenditure, to help see any differences in Figure 6. The aim of this analysis was, for Chinese New Year, to compare the growth rate between December and January with that between January and February. For the Islamic Holidays the aim was to compare growth rates for the month in which the event took place, and the months preceding and following it. Table 5 shows the growth rate between months, alongside the date of each event in each year.

Generally the growth rates are very similar with and without the different holiday regressors. There are a few examples where the inclusion of a regressor changes the direction of growth, namely when Eid al-Fitr is included.

Table 5: Month-on-month growth rate (percent) in seasonally adjusted time series.

Date of Event	No Regressor		Event Regressor	
	Dec-Jan	Jan-Feb	Dec-Jan	Jan-Feb
Chinese New Year				
Feb 3rd 2011	-5.1	9.4	-4.9	8.8
Jan 23rd 2012	6.4	-3.3	6.9	-4.5
Feb 10th 2013	7.7	-0.7	8.7	-2.7
Jan 31st 2014	3.2	1.5	3.9	-0.6
Feb 19th 2015	-1.6	-5.1	-1.0	-6.9
Feb 8th 2016	-9.8	1.3	-9.5	-0.4

	Two Months Before - Month Before	Month Before - Month of Event	Month of Event - Month After	Two Months Before - Month Before	Month Before - Month of Event	Month of Event - Month After
Eid-al-Fitr						
Aug 31st 2011	-0.7	0.9	-7.2	-0.6	0.5	-7.3
Aug 19th 2012	2.1	18.3	-5.2	2.0	17.9	-5.0
Aug 8th 2013	10.6	-5.2	3.9	10.3	-5.4	4.3
Jul 29th 2014	7.4	2.7	-4.8	7.4	2.4	-5.0
Jul 18th 2015	0.1	-10.1	1.1	0.1	-10.5	1.0
Eid-al-Adha						
Nov 7th 2011	11.1	-0.7	-10.4	10.8	0.8	-11.9
Oct 26th 2012	-5.2	-13.5	7.1	-4.3	-13.7	7.6
Oct 15th 2013	3.9	2.8	-4.2	5.4	2.2	-4.3
Oct 4th 2014	-0.9	-6.3	8.2	0.9	-6.9	7.5
Sep 24th 2015	1.1	5.3	2.8	0.9	7.5	1.9
Ramadan						
Aug 1st 2011	-0.7	0.9	-7.2	-0.9	2.1	-8.5
Jul 20th 2012	-9.4	2.1	18.3	-9.4	2.4	18.4
Jul 9th 2013	-2.5	10.6	-5.2	-2.6	11.5	-6.0
Jun 29th 2014	1.3	7.4	2.7	1.4	7.3	3.8
Jun 18th 2015	6.8	0.1	-10.1	6.9	0.6	-10.3

* For the Islamic events as they move around the year, growth rates are shown for months relative to the month of the event. For example for Eid al-Fitr, in 2011 the event is in August. The growth rates presented are for June-July (two months before - month before), July-August (month before - month of event) and August-September (month of event - month after). In 2015 the event is in July. The growth rates presented are May-June (two months before - month before), June-July (month before - month of event) and July-August (month of event - month after).

3.2.3 Span Analysis

A span analysis of the International Passenger Survey: Overseas visits to the UK, expenditure series was conducted to see whether the most favourable holiday regressors on the whole span of data were consistently favoured over no regressor on shorter spans. Five spans were considered. Only the Chinese New Year regressor had an AICC value consistently lower than that of no regressor, however the AICC values were very close with all differences below 2. This suggests that further work is needed to determine whether the regressors are sufficiently significant and stable enough over time to be included as part of seasonal adjustment.

4. Conclusion

Although the results from this research show that the UK Easter regressors, Chinese New Year regressors and Islamic Calendar regressors do have an impact on some official time series, the results are not conclusive enough to implement these methods in production. In the cases of Easter regressors, there are instances where the current regressors are contained within the new UK Easter regressors. It would be interesting to look into what the additional effect of the extended period is, not necessarily the full four-day or two-week periods.

In comparison with any current methods for the respective holiday the Eid al-Fitr regressors were preferred the most, being preferred in 32% of time series. This was closely followed by the Ramadan regressors (31%) and the Easter & Eid al-Adha regressors (both 28%). The least preferred regressors were those for Chinese New Year, only being preferred in 22% of time series.

4.1 Further Work

There are a number of directions that this research can go from here. The main things to consider will be;

- Applying the new regressors to alternative series (*for example, regional series or unconsidered outputs*)
- Constructing and testing different shape regressors (*for example, Shape 2 regressors*)
- Constructing and testing regressors of alternative windows
- Using alternative diagnostics to assess suitability
- Considering the magnitude in the difference between the lowest AICC value and the other AICC values
- Analysing the interaction between moving holidays (*for example - between Ramadan and the day after Ramadan, Eid al-Fitr*).

Whilst attempting all these things would be ideal in expanding the research, they are not adequate in determining whether or not these regressors are appropriate for use in the production of official statistics. Every year TSAB undertakes a seasonal adjustment review of all seasonally adjusted time series produced by ONS. Within these reviews the team considers whether or not changes are required to the seasonal adjustment, making their decision based on a trade off between the most appropriate adjustment and the size of revisions to the series. As a result, should TSAB choose to roll out any of these new regressors into production, an investigation would be undertaken to look at the stability of these regressors (ie whether year-on-year these regressors would be preferred over any other possible alternatives). This would be an expansion of the span analysis in Sections 3.1.3 and 3.2.3, analysing further years and more series.

To conclude, it is likely that this research will be continued but in the meantime any new regressor discussed in this paper will not be used in production.

References

[1] Unknown. Estimating and Removing the Effects of Chinese New Year and Ramadan to Improve the Seasonal Adjustment Process. Austrailian Economic Indicators, 2005.

[2] U.S. Census Bureau. X-13ARIMA-SEATS Reference Manual. U.S. Census Bureau, Washington, DC 20233, version 1.1 edition, March 2016. URL: https://www.census.gov/srd/www/x13as/.

[3] David F Findley, Brian C Monsell, William R Bell, Mark C Otto, and Bor-Chung Chen. New capabilities and methods of the X-12-ARIMA seasonal-adjustment program. Journal of Business & Economic Statistics, 16(2):127_152, 1998.

[4] Caleb Leung, Craig H McLaren, and Mark Zhang. Adjusting for an Easter Proximity Effect. Australian Bureau of Statistics, 1999.

[5] Jin-Lung Lin, Tian-Syh Liu, et al. Modeling lunar calendar holiday effects in Taiwan. Taiwan Economic Forecast and Policy, 33(1):1_37, 2002.

[6] Danny Pfeffermann and Je_rey M Fisher. Festival and working days prior adjustments in economic time series. International Statistical Review/Revue Internationale de Statistique, pages 113_124, 1982.

[7] Riaz Riazuddin et al. Detection and forecasting of Islamic calendar effects in time series data. Technical report, State Bank of Pakistan, Research Department, 2002.

[8] Norhayati Shuja, Mohd Alias Lazim, and Yap Bee Wah. Moving holiday effects adjustment for Malaysian economic time series. 2007.

Appendix

Table 6 provides examples of the values for the regressors for the UKE[4], UKE[12], Chinese New Year (7,0) Shape 1, Eid-al-Fitr (14,0) Shape 0, Eid-al-Adha (14,0) Shape 0, and Ramadan regressor A (Shape 0, full month).

Date	Easter		Chinese New Year	Eid-al-Fitr	Eid-al-Adha	Ramadan
	UKE[4]	UKE[12]				
2011.01	0.00	0.00	0.44	-0.09	-0.09	-2.57
2011.02	0.00	0.00	-0.43	-0.09	-0.09	-2.57
2011.03	-0.25	-0.25	-0.02	-0.09	-0.09	-2.57
2011.04	0.25	0.25	0.00	-0.09	-0.09	-2.57
2011.05	0.00	0.00	0.00	-0.09	-0.09	-2.57
2011.06	0.00	0.00	0.00	-0.09	-0.09	-2.57
2011.07	0.00	0.00	0.00	-0.09	-0.09	-2.57
2011.08	0.00	0.00	0.00	-0.09	0.91	27.43
2011.09	0.00	0.00	0.00	-0.09	-0.09	-2.57
2011.10	0.00	0.00	0.00	0.45	-0.09	-2.57
2011.11	0.00	0.00	0.00	0.38	-0.09	-2.57
2011.12	0.00	0.00	0.00	-0.09	-0.09	-2.57
2012.01	0.00	0.00	0.64	-0.09	-0.09	-2.57
2012.02	0.00	0.00	-0.63	-0.09	-0.09	-2.57
2012.03	-0.25	-0.25	-0.02	-0.09	-0.09	-2.57
2012.04	0.25	0.25	0.00	-0.09	-0.09	-2.57
2012.05	0.25	0.00	0.00	-0.09	-0.09	-2.57
2012.06	0.00	0.00	0.00	-0.09	-0.09	-2.57
2012.07	0.00	0.00	0.00	-0.09	-0.09	9.43
2012.08	0.00	0.00	0.00	-0.09	0.91	15.43
2012.09	0.00	0.00	0.00	-0.09	-0.09	-2.57
2012.10	0.00	0.00	0.00	0.91	-0.09	-2.57
2012.11	0.00	0.00	0.00	-0.09	-0.09	-2.57
2012.12	0.00	0.00	0.00	-0.09	-0.09	-2.57
2013.01	0.00	0.00	-0.02	-0.09	-0.09	-2.57
2013.02	0.00	0.00	0.04	-0.09	-0.09	-2.57
2013.03	0.50	0.34	-0.02	-0.09	-0.09	-2.57
2013.04	-0.50	-0.34	0.00	-0.09	-0.09	-2.57
2013.05	0.00	0.00	0.00	-0.09	-0.09	-2.57
2013.06	0.00	0.00	0.00	-0.09	-0.09	-2.57
2013.07	0.00	0.00	0.00	-0.09	0.38	20.43
2013.08	0.00	0.00	0.00	-0.09	0.45	4.43
2013.09	0.00	0.00	0.00	-0.09	-0.09	-2.57
2013.10	0.00	0.00	0.00	0.91	-0.09	-2.57
2013.11	0.00	0.00	0.00	-0.09	-0.09	-2.57
2013.12	0.00	0.00	0.00	-0.09	-0.09	-2.57
2014.01	0.00	0.00	0.64	-0.09	-0.09	-2.57

2014.02	0.00	0.00	-0.63	-0.09	-0.09	-2.57
2014.03	-0.25	-0.25	-0.02	-0.09	-0.09	-2.57
2014.04	0.25	0.25	0.00	-0.09	-0.09	-2.57
2014.05	0.00	0.00	0.00	-0.09	-0.09	-2.57
2014.06	0.00	0.00	0.00	-0.09	-0.09	-0.57
2014.07	0.00	0.00	0.00	-0.09	0.91	25.43
2014.08	0.00	0.00	0.00	-0.09	-0.09	-2.57
2014.09	0.00	0.00	0.00	0.65	-0.09	-2.57
2014.10	0.00	0.00	0.00	0.18	-0.09	-2.57
2014.11	0.00	0.00	0.00	-0.09	-0.09	-2.57
2014.12	0.00	0.00	0.00	-0.09	-0.09	-2.57
2015.01	0.00	0.00	-0.36	-0.09	-0.09	-2.57
2015.02	0.00	0.00	0.37	-0.09	-0.09	-2.57
2015.03	-0.25	-0.08	-0.02	-0.09	-0.09	-2.57
2015.04	0.25	0.08	0.00	-0.09	-0.09	-2.57
2015.05	0.00	0.00	0.00	-0.09	-0.09	-2.57
2015.06	0.00	0.00	0.00	-0.09	-0.09	10.43
2015.07	0.00	0.00	0.00	-0.09	0.91	14.43
2015.08	0.00	0.00	0.00	-0.09	-0.09	-2.57
2015.09	0.00	0.00	0.00	0.91	-0.09	-2.57
2015.10	0.00	0.00	0.00	-0.09	-0.09	-2.57
2015.11	0.00	0.00	0.00	-0.09	-0.09	-2.57
2015.12	0.00	0.00	0.00	-0.09	-0.09	-2.57
2016.01	0.00	0.00	0.11	-0.09	-0.09	-2.57
2016.02	0.00	0.00	-0.09	-0.09	-0.09	-2.57
2016.03	0.75	0.67	-0.02	-0.09	-0.09	-2.57
2016.04	-0.75	-0.67	0.00	-0.09	-0.09	-2.57
2016.05	0.00	0.00	0.00	-0.09	-0.09	-2.57
2016.06	0.00	0.00	0.00	-0.09	0.51	21.43
2016.07	0.00	0.00	0.00	-0.09	0.31	2.43
2016.08	0.00	0.00	0.00	0.05	-0.09	-2.57
2016.09	0.00	0.00	0.00	0.78	-0.09	-2.57
2016.10	0.00	0.00	0.00	-0.09	-0.09	-2.57
2016.11	0.00	0.00	0.00	-0.09	-0.09	-2.57
2016.12	0.00	0.00	0.00	-0.09	-0.09	-2.57

Methodology Advisory Service (MAS)

The Methodology Advisory Service is a service of the Office for National Statistics (ONS); it aims to spread best practice and improve quality across official statistics through methodological work and training activity. The ONS has about one hundred methodologists - highly qualified statisticians and researchers; their primary role is to provide expert support, advice and methodological leadership to the ONS in producing and analysing National Statistics.

Methodology staff are arranged into Centres of Expertise, each comprising a team of specialists who keep abreast of research and developments in their area of expertise through contacts with academia, other national statistical institutes and the wider research community. Many of these Centres have international reputations and present research and applied work at conferences and at other meetings of experts in their fields. Examples of these centres are Sample Design and Estimation and Time Series Analysis.

The Methodology Advisory Service has a remit to extend the services of ONS methodologists beyond ONS into other public sector organisations. Every year, MAS carries out projects with customers addressing a wide range of statistical requirements. As well as calling on methodology staff, MAS can also draw on the wider expertise of statisticians, researchers and subject area specialists across the ONS. Further expertise is available though links with Universities.

Contact ***MAS@ons.gov.uk***

GSS Methodology Series

Latest reports in the GSS Methodology Series:

38. 100 Years of the Census of Production in the UK, Paul Smith

39. Quality of the 2010 Electoral Register in England & Wales, Neil Hopper

40. Modelling sample data from smart-type electricity meters to assess potential within Official Statistics, Susan Williams and Karen Gask

41. Using geolocated Twitter traces to infer residence and mobility, Nigel Swier, Bence Komarniczky and Ben Clapperton

42. Assessing the Generalised Structure Preserving Estimator (GSPREE) for Local Authority Population Estimates by Ethnic Group in England, Solange Correa-Onel, Alison Whitworth and Kirsten Piller

Reports are available from:

https://www.ons.gov.uk/methodology/methodologicalpublications/generalmethodology/currentmethodologyarticles

Forthcoming Courses

GSS Statistical Training Programme

A series of government specific short courses (between 0.5 and 2 days in length) delivered by methodological experts in the field. These courses are delivered at ONS sites in London, Newport and Titchfield.

For further information on the available courses see the Statistical Training Service prospectus:

https://gss.civilservice.gov.uk/learning-and-development/training-events/training-co-ordinated-by-the-statistical-training-service/

or contact gss.capability@ons.gov.uk

The timetable for 2017/18 will be available for downloaded through the GSS Learning Curriculum in the above link shortly.

Details of additional opportunities for learning can also be found in in the training events page. In summary these are:

MSc in Data Analytics for Government

This is available at the following universities: University College London, Oxford Brookes University and Southampton University. More details can be found via this link.

https://gss.civilservice.gov.uk/blog/2017/07/msc-data-analytics-government/

MSc in Official Statistics

Available at Southampton University. This has been replaced by the MSc. in Data Analytics but will continue for those already enrolled.

The degree in Official Statistics is part of the network of Master programmes provided at European level. Further details can being via this link.

https://ec.europa.eu/eurostat/cros/content/emos_en

European Statistical Training Programme 2017

The purpose of the European Statistical Training Programme (ESTP) is to provide statisticians the opportunity to participate in international training courses, workshops and seminars at postgraduate level. It comprises courses in Official Statistics, IT applications, Research and Development and Statistical Management. More information on the courses can be found on the Eurostat website

http://ec.europa.eu/eurostat/c/portal/layout?p_l_id=747782&p_v_l_s_g_id=0

Enquiries

The Survey Methodology Bulletin is usually published twice a year, in Spring and Autumn. Copies of many previous editions are available electronically at:

http://webarchive.nationalarchives.gov.uk/20160105160709/http://www.ons.gov.uk/ons/guide-method/method-quality/survey-methodology-bulletin/index.html

If you would like to be added to the distribution list please email ONS Methodology at:

methodology@ons.gsi.gov.uk

Or write to us at:

Philip Lowthian
Survey Methodology Bulletin
2nd Floor
Office for National Statistics
Drummond Gate
London
SW1V 2QQ

ons.gov.uk
visual.ons.gov.uk