Combining national surveys with composite calibration to improve the precision of estimates from the UK's Living Costs and Food Survey

Takis Merkouris[*], Paul A. Smith[†] and Andy Fallows[‡]

Abstract

The UK's Living Costs and Food (LCF) Survey has a relatively small sample size, but produces estimates which are widely used, notably as a key input to the calculation of weights for consumer price indices (CPI). There has been a recent call for the use of additional data sources to improve the estimates from the LCF. Since some LCF variables are shared with the much larger Labour Force Survey (LFS), we investigate combining data from these surveys using composite calibration to improve the precision of estimates from the LCF. We undertake model selection to choose a suitable set of common variables for the composite calibration using the effect on the estimated variances for national and regional totals of important LCF variables. The variances of estimates for common variables are reduced to around 5% of their original size. Variances of national estimates are reduced (across several quarters) by around 10% for expenditure and 25% for income; these are the variables of primary interest in the LCF. Reductions in the variances of regional estimates vary more but are mostly large when using common variables at the regional level in the composite calibration. The composite calibration also makes the LCF estimates for employment status almost consistent with the outputs of the LFS, which is an important property for users of the statistics. A novel alternative method for variance estimation, using stored information produced by the composite calibration, is also presented.

Keywords: household expenditure survey, composite estimation, data integration, data harmonisation, linearised jackknife

Significance statement: We show that integrating two data sources using composite calibration gives substantial reductions in the variances of important estimates in a household expenditure survey, the Living Costs and Food Survey. The approach has several advantages: it makes LCF estimates almost consistent with LFS outputs for common variables and provides weights that can

[*]Dept. of Statistics, Athens University of Economics and Business, Athens, Greece, merkouris@aueb.ac.gr
[†]S3RI and Dept. of Social Statistics, University of Southampton, UK, p.a.smith@soton.ac.uk
[‡]Office for National Statistics, Newport, UK, andy.fallows@ons.gov.uk

be used for all variables. Variances account appropriately for the uncertainty in estimation, and are straightforwardly calculated. Variances of LCF estimates of common variables are very substantially reduced. Variance reductions for other variables depend on the correlation with the common variables in the composite calibration. We select common variables by their impact on the estimated variances of income and expenditure. Variance reductions vary between variables and across periods.

1 Introduction

The household expenditure survey in the UK is the Living Costs and Food Survey (LCF). It is an important survey source, providing estimates of household expenditure broken down by commodities, used for many purposes (ONS, 2009), but notably as one of the main inputs to the calculation of weights for consumer price indices (CPI). The weights are derived through the input-output framework of the national accounts, which integrates multiple data sources for specific commodities. Nevertheless the LCF estimates are the main source of expenditure information for many commodities. Its achieved sample size is approximately 5,000 households per year, which is relatively small for a household expenditure survey, particularly relative to the UK's population size (for example the biennial Survey of Household Spending in Canada sampled 48,570 households in 2021 (Statistics Canada, 2021) for a population about a third as large as the UK). The LCF sample size is, however, just above the median achieved sample size in the 2015 round of Household Budget Surveys in Europe (Eurostat, 2020, Table 1). The small sample size results in estimates for expenditure by commodity (particularly rarer commodities) with relatively large variances. These contribute to larger variances for the weights in the CPI (Dawber et al., 2022), particularly if there is a requirement to produce estimates for domains, for example by regions. Other estimated breakdowns also have relatively large variances, and are therefore not suitable for some detailed uses. This has been highlighted as a weakness in a recent assessment by the Office for Statistics Regulation (2021).

The quality of the estimates from household expenditure surveys is influenced by different sources of error. Measurement errors can result from the relatively complex data collection processes, which require significant effort from respondents (Jäckle et al., 2021; Eckman, 2022). There is also a tendency, across many similar household expenditure surveys, for households with very low incomes to have relatively high expenditure (Brewer et al., 2017), from a mismatch in expenditure and income. Brewer et al. (2017) suggest that income is underestimated, which has led to attempts to reconcile income (Burton et al., 2020) in the data collection process.

The Office for Statistics Regulation also suggested the use of additional data sources as a creative approach to improving the estimates from the LCF (Office for Statistics Regulation, 2021). Jäckle et al. (2021) review the available data sources for the measurement of household finances, including data obtained by linkage to administrative data sources (such as credit card and loyalty card data) and data from volunteer collections (such as receipt scanning). They also review the different types of error, setting them within a total survey error framework.

We assume that the effects of these kinds of errors on data quality are well handled, since the LCF is a well-designed national survey (see Ralph and Manclossi (2016) for a more comprehensive view). Among the effects of various errors on the quality of the survey estimates, we focus on the relatively large sampling errors resulting from the small sample size of the LCF. One consequence, identified in the LCF Quality Review (Ralph and Manclossi, 2016) is a mismatch between the employment/unemployment/inactivity estimates from the LCF and the corresponding official figures from the Labour Force Survey (LFS), which has a substantially larger sample size. Users identified this inconsistency as a challenge in using and understanding estimates from LCF and LFS, and sought greater consistency between them. We therefore follow Office for Statistics Regulation (2021)'s suggestion to utilise alternative data sources by seeking to use the information from the two surveys together. We want to incorporate the information from the LFS into the estimation for the LCF for two purposes: to reduce the variance of estimates for variables collected in the LCF, at national and regional level; and to correct for differential representation of employment statuses in the LCF to improve consistency between LCF and LFS estimates.

Information from two surveys may be combined for more efficient estimation and consistency of estimates for common variables in several ways; see Yang and Kim (2020) for a review. In this paper the proposed method is an adaptation of the regression method of Merkouris (2004, 2010) for integrating independent surveys with some common variables through composite calibration, whereby two surveys are weighted together so that they produce consistent estimates for selected common variables. Other design-based approaches that could be considered include the regression methods of Renssen and Nieuwenbroek (1997) and Hidiroglou (2001), the model-based method of Kim and Rao (2012) and the empirical likelihood methods of Wu (2004), Chen and Kim (2014) and Berger and Kabzińska (2020). The proposed method is adopted because it serves effectively all the set objectives, is also very practical and fits readily into the system of regression estimation and variance estimation of the LFS and the LCF. We also present a novel alternative method for variance estimation, using stored information produced by the composite calibration.

We use the LCF and LFS data to pursue a number of questions:

1. can we substantially reduce the variance of the LCF estimates by incorporating data from the LFS?

2. what is the best choice of common variables to include in the composite estimation? There are several variables common to LCF and LFS, and we seek common variables correlated with the outcomes of interest, but which will also be useful in a general calibration approach suitable for many variables, since only one set of weights will be used for all variables for ease of processing and consistency.

3. what is the effect of the composite estimation on the final weights, under different calibration models?

4. does composite estimation improve the consistency between LCF and LFS estimates, and therefore reduce the variability of LCF estimates across quarters to obtain a more stable time series?

The remainder of the paper is structured as follows. Section 2 contains an overview of the designs and relevant features of the LCF and LFS. Section 3 describes the composite calibration approach, and section 4 explains where the methodolology has been adapted to the situation of the LCF and LFS. Section 5 presents the results of applying composite calibration to the LCF in 2017-18. Section 6 discusses the impacts on the LCF estimates and the effectiveness of the approach.

2 Survey descriptions

2.1 Outline description of Living Costs and Food Survey

The UK's LCF is a nationally representative household survey, with an achieved sample size of around 5,000 households per year. Sampling follows a two-stage design (Ralph and Manclossi, 2016, section 3.1), with postcode sectors (geographic areas based on postal delivery information) forming the primary sampling units (PSUs), which are stratified by region. Within regions, PSUs are ordered by percentage of households without a car, percentage of households in the National Statistics Socio-economic Classification (NSSEC) groups 1-3, and percentage of pensioners, all derived from the previous population census. 683 postcode sectors are selected across Great Britain (that is, the UK excluding Northern Ireland) by systematic selection in this ordered list. Within postcode sectors, addresses are ordered geographically, and then 18 addresses in each postcode sector are selected by systematic sampling. The sample in Great Britain is therefore clustered by postcode sector. Northern Ireland is treated as a separate stratum, within which there is no clustering, just a systematic random selection of addresses.

The primary outputs from LCF are annual, but there is also interest in expenditure and income by quarter, for which the sample size does not support any detailed breakdown. The annual sample is therefore processed in four quarterly tranches, with the annual totals derived as the average (for stock variables) or sum (for flow variables) of the quarterly estimates.

All adults in sampled households are included, and (before COVID affected survey operations) collection was in two parts, a household interview administered by an interviewer using CAPI, and an expenditure diary covering a two-week period following the interview, to be completed by all household members. The main interest is in household (not personal) expenditure, so clustering within households is not important for expenditure estimates. Within-household clustering does, however, have an effect for person level variables such as labour force status. The interview involves substantial input from the whole household to the expenditure diary, and therefore the LCF response rate is relatively low (Office for Statistics Regulation, 2021); it was 43% in the financial year 2017-18, covering most of the period we investigate below (ONS, 2019).

The LCF is weighted to household population estimates (demographic projections of the numbers of people in the non-institutional population), published regularly by the ONS, with a breakdown by age, sex and region. All people in a household receive the same weight (integrated weighting, (Lemaître and Dufour, 1987)), ensuring consistency when working with household-level variables. A key function of the weighting is to adjust for differential nonresponse, and therefore we always include these variables when exploring composite calibration.

The LCF estimates of primary interest are expenditure and its breakdown by commodities, but it also includes a range of other variables, including labour force status, tenure and income. Some of these were introduced when household surveys in the UK were harmonised as part of an integration project (Smith, 2009); many are important classifiers or predictors. LCF estimates are used mainly at national level, for example in the National Accounts, but there is also interest in breakdowns for regions (Dawber et al., 2022), income deciles and other characteristics (ONS, 2020, 2022).

There has been concern that the LCF operation may not produce sufficiently reliable estimates and may not have kept pace with the latest methodology and technology. This led to a National Statistics Quality Review (Ralph and Manclossi, 2016) and an assessment of the LCF against the National Statistics Code of Practice, including an appraisal of progress in implementing the quality review's recommendations (Office for Statistics Regulation, 2021). We address this issue in this paper and explore the approach outlined in section 1 to improve the quality of estimates from the LCF.

2.2 Outline description of the Labour Force Survey

The UK's LFS is also a nationally representative household survey, with an achieved sample size of around 40,000 households per quarter (ONS, 2015, 2021). The larger sample size means that addresses are sufficiently dense that they can be interviewed efficiently with an unclustered design, and therefore addresses are selected as a single-stage sysematic selection. The LFS has a rotating panel design where sampled households are interviewed five times at quarterly intervals (a 1-2-1(5) design, Steel (1997)). Sampled households are grouped, so data collection follows a weekly pattern of defined areas; weeks are aggregated consistently into months for the production of estimates. Therefore the sample is unclustered for any consecutive three-month period, but clustered over shorter periods. The main estimates are produced monthly, covering the most recent three months of data, giving 'rolling quarterly' outputs. Because the rotation pattern involves repeat surveys of the same household, the annual sample size (that is, the number of households included in the sample at any point during a year), is only *c.*64,000, around 1.6 times the quarterly sample size.

All members of sampled households are interviewed, but the main interest is in person-level variables, so it is less critical that all members of a household participate. The first contact and interview are face-to-face, with follow-up interviews by telephone wherever possible. The LFS is the largest official social survey in the UK, with substantial regional sample sizes. This makes it the best source for estimates which need a regional breakdown, and has led to a relatively long questionnaire as there has been pressure to include many topics in order to obtain regional estimates.

The weighting uses the same population totals as the LCF, broken down by age, sex and region, but with additional breakdowns of population sizes by Local Authority. Some specific communal establishments are included in the sample and therefore in the population totals too. The weighting is undertaken at person level (so different people in the same household in general get different weights). This weighting produces the main headline figures, though a supplementary household-level weighting is also produced for use with household variables. As with the LCF, a key function of weighting is to adjust for differential nonresponse, so we always include these calibration variables when exploring composite calibration. The LFS response rate averaged 42.7% over quarter 1 (Q1) of 2017 to quarter 2 (Q2) of 2018, the quarters we use in our application below (ONS, 2021, Fig. 5.1).

The LFS is the official source in the UK for employment, unemployment and inactivity according to the International Labour Organisation definition. However, the dataset used here has some small differences from the official LFS (see section 4.1). Note therefore that the LFS estimates presented here are different from the official estimates, and should not be treated as official.

Regional LFS estimates are improved by a supplementary survey, the annual population survey (APS), which boosts the LFS sample to ensure a minimum sample size in each local authority, with an annual rotation pattern. The combination of LFS and APS is therefore useful for regional variables, but we do not consider this additional complication in the present paper.

3 Composite estimation for the LCF

3.1 Current estimation systems for the LFS and the LCF

This section describes the current estimation systems of the LFS and the LCF, in a unified manner on account of their similarities. In particular, it explains the production of the generalized regression estimators of totals through a procedure of calibration of the survey's design weights, and the method of computing the variances of these estimators.

Let U denote the sampling frame of households, and s denote a sample of households from U, with sample size n, in which data are collected from all eligible members of each household. Let w_i be the design weight of the i-th household in the sample s, and let y be a household-level variable with value y_i for the i-th household and population total $t_y = \sum_U y_i$. The basic estimator of t_y is the Horvitz-Thompson estimator $\hat{t}_y = \sum_s w_i y_i$.

The calibration procedure for both the LFS and the LCF involves three categorical auxiliary variables defined for the survey population of persons, namely sex, age categories, and regions. Calibration groups, i.e., subpopulations of known sizes, are defined by the various categories of particular categorical variables, or by cross-classification of categories of different variables. Two sets of calibration groups, called partitions, are specified so that each responding individual is in one calibration group of each partition. One partition comprises a number of calibration groups defined by cross-classifying sex with age groups, and a second partition includes as groups geographic regions of the country.

Suppose that the various categorical auxiliary variables used in calibration define a total of p distinct calibration groups, in the two partitions. Let m_i denote the size of the i-th household in the sample s, and for each partition define the indicator variable x_{ijk} such that $x_{ijk} = 1$ if the j-th member of household i belongs to group k of that partition and $x_{ijk} = 0$ otherwise. For each member of the household this defines a p-dimensional vector, of which as many elements as the number of partitions are set equal to 1 and the rest are set equal to 0. Then $x_{ik} = \sum_{j=1}^{m_i} x_{ijk}$ is the number of members of the i-th household that belong to the calibration group k, for $k = 1, \dots, p$. Thus, for the i-th household, $\mathbf{x}_i^{'} = (x_{i1}, \dots, x_{ip})$ is the row vector of membership counts in the p calibration groups, where $(')$ denotes vector transpose. Of course, for any particular household, not all p counts are necessarily greater than 0. The associated $p \times 1$ vector of population totals, called calibration totals, is $\mathbf{t_x} = \sum_U \mathbf{x}_i$, and its Horvitz-Thompson estimator is $\hat{\mathbf{t}}_\mathbf{x} = \sum_s w_i \mathbf{x}_i$.

The calibrated weight, c_i, of the i-th household in the sample is defined by

$$c_i = w_i \left[1 + \mathbf{x}_i^{'} \left(\sum_s w_i \mathbf{x}_i \mathbf{x}_i^{'} \right)^{-1} \left(\mathbf{t_x} - \sum_s w_i \mathbf{x}_i \right) \right], \tag{1}$$

where the inverted matrix $\left(\sum_s w_i \mathbf{x}_i \mathbf{x}_i^{'} \right)$ is of dimension $p \times p$. For this matrix to be invertible, it is required that one (arbitrary) group be removed from all but one of any partitions which have the same aggregate population totals. The calibrated weights given by (1) satisfy the set of p equations $\sum_s c_i \mathbf{x}_i = \mathbf{t_x}$. Equation (1) can be written in matrix form as

$$\mathbf{c} = \mathbf{w} + \mathbf{W}\mathbf{X} \left(\mathbf{X}^{'} \mathbf{W} \mathbf{X} \right)^{-1} \left(\mathbf{t_x} - \mathbf{X}^{'} \mathbf{w} \right), \tag{2}$$

where $\mathbf{w}$ and $\mathbf{c}$ are the vectors of design and calibrated weights, respectively, and $\mathbf{W}$ is the $n \times n$ diagonal matrix with diagonal elements the design weights, and

$$\mathbf{X} = \begin{pmatrix} \mathbf{x}_1^{'} \\ \vdots \\ \mathbf{x}_n^{'} \end{pmatrix} = \begin{pmatrix} x_{11} & \cdots & x_{1p} \\ \vdots & & \vdots \\ x_{n1} & \cdots & x_{np} \end{pmatrix} \tag{3}$$

is the $n \times p$ 'design matrix'. The expression (2) facilitates the computation of the calibrated weights, and the description of the calibration setup for composite estimation.

The calibrated household weights given by (1) are used for the estimation of household-level population characteristics in the LFS and the LCF. As mentioned in section 2.2, calibration in the LFS is also undertaken at person level for use with person characteristics. However, for the combination of data from the two surveys, only calibrated household weights are needed. Then, using the calibrated weights defined by (1), the calibration estimator of the total t_y of a household-level variable y, constructed as $\sum_s c_i y_i$, takes the form of (generalized) regression estimator

$$\hat{t}_y^R = \hat{t}_y + \hat{\mathbf{B}}(\mathbf{t_x} - \hat{\mathbf{t}}_\mathbf{x}), \tag{4}$$

where $\hat{\mathbf{B}} = \sum_s w_i y_i \mathbf{x}_i' \left(\sum_s w_i \mathbf{x}_i \mathbf{x}_i'\right)^{-1}$. The alternative form of the estimator $\hat{t}_y^R$, based on (2), that will be used later is $\hat{t}_y^R = \mathbf{Y}'\mathbf{c}$, where $\mathbf{Y}$ is the $n \times 1$ vector of the sample values of y. Obviously, by calibration we have $\hat{t}_{\mathbf{x}}^R = \mathbf{X}'\mathbf{c} = \mathbf{t}_{\mathbf{x}}$. This means that the regression estimate of the population total of any calibration group is exactly equal to that total.

Calculation of the approximate (for large samples) variance of $\hat{t}_y^R$ uses the linearized form $\hat{t}_y^R \approx \hat{t}_y + \mathbf{B}(\mathbf{t}_{\mathbf{x}} - \hat{\mathbf{t}}_{\mathbf{x}})$, where $\mathbf{B} = \sum_U y_i \mathbf{x}_i' \left(\sum_U \mathbf{x}_i \mathbf{x}_i'\right)^{-1}$ and U denotes the sampling frame of households. Writing $\hat{t}_y^R \approx \mathbf{B}\mathbf{t}_{\mathbf{x}} + \sum_s w_i(y_i - \mathbf{B}\mathbf{x}_i) = \mathbf{B}\mathbf{t}_{\mathbf{x}} + \sum_s w_i e_i$, where e_i are the regression residuals, it follows that the variance $V(\hat{t}_y^R)$ is approximately the variance $V(\sum_s w_i e_i)$. Then the estimated approximate variance of $\hat{t}_y^R$ is $\hat{V}(\sum_s w_i \hat{e}_i)$, where $\hat{e}_i = y_i - \hat{\mathbf{B}}\mathbf{x}_i$ are the estimated residuals. These estimated residuals are conveniently generated in vector form as $\hat{\mathbf{e}} = \mathbf{Y} - \mathbf{X}\hat{\mathbf{B}}' = [\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{W}\mathbf{X})^{-1}\mathbf{X}'\mathbf{W}]\mathbf{Y}$. For stratified multi-stage sampling, used in both LFS and LCF, the exact formula for $\hat{V}(\sum_s w_i \hat{e}_i)$ is impracticable, and so the alternative approximate formula for the estimated variance of $\hat{t}_y^R$ that is typically used is

$$\hat{V}(\hat{t}_y^R) = \sum_h^H \frac{n_h}{n_h - 1} \sum_l^{n_h} (e_{hl} - \bar{e}_h)^2, \tag{5}$$

where $h = 1, \ldots, H$ denotes stratum, n_h is the number of primary sampling units (clusters) in stratum h, $e_{hl} = \sum_i w_{hli} \hat{e}_{hli}$ with $\hat{e}_{hli}$ and w_{hli} being the estimated residual and design weight, respectively, associated with the sample element (hli), and $\bar{e}_h = (1/n_h) \sum_l e_{hl}$. Replacing the design weights w_{hli} in e_{hl} by the calibrated weights c_{hli} gives the jackknife linearization variance estimator; see Yung and Rao (1996). This approach of jackknife linearization, involving residuals weighted by the calibrated weights, is used in variance estimation in both the LFS and the LCF; see Holmes and Skinner (2000).

3.2 Composite estimation incorporating LFS data in LCF

Let s_1 and s_2 denote quarterly samples of the LFS and the LCF, respectively, with n_1 and n_2 denoting their household sample sizes, and let p_1 and p_2 be the numbers of calibration groups (with specified partitions) in the two surveys. Also, let $\mathbf{w}_1$ and $\mathbf{w}_2$ denote the vectors of household design weights for the two samples. The proposed composite calibration procedure requires simultaneous calibration for the LFS and LCF household samples. For this calibration the appropriate design matrix of household membership counts, having dimension $(n_1 + n_2) \times (p_1 + p_2)$, is

$$\mathbf{X} = \begin{pmatrix} \mathbf{X}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{X}_2 \end{pmatrix}, \tag{6}$$

where $\mathbf{X}_1$ and $\mathbf{X}_2$ are the design matrices for the LFS and the LCF, respectively, and can be built using the combined LFS and LCF data files sorted by sample, i.e., treating the samples s_1 and s_2 as superstrata. With this matrix $\mathbf{X}$, calibration of the household weights of both samples is carried out by applying the generic formula (2), using $\mathbf{w} = (\mathbf{w}_1^{'}, \mathbf{w}_2^{'})^{'}$, $\mathbf{W} = \text{diag}(\mathbf{W}_1, \mathbf{W}_2)$ with obvious definition of $\mathbf{W}_1$ and $\mathbf{W}_2$, and the vector of calibration totals $\mathbf{t_x} = (\mathbf{t}_{\mathbf{x}_1}^{'}, \mathbf{t}_{\mathbf{x}_2}^{'})^{'}$ for the LFS and the LCF. In the vector of the resulting calibrated weights, the first n_1 and the remaining n_2 of them are the calibrated weights that would be obtained from separate calibrations of the LFS and LCF vectors $\mathbf{w}_1$ and $\mathbf{w}_2$, respectively. In particular, these calibrated weights give rise to two independent regression estimates for any variable that is common to the LFS and the LCF.

Next, to combine data from the LFS and the LCF, we choose categorical variables (with unknown totals) that are common to the two surveys, and (ideally) strongly correlated with main LCF variables. These variables define a total of q distinct categories to which each household member belongs. These categories form the additional calibration groups in an extended calibration procedure. Now, for any household, say i, of the LFS or the LCF let $\mathbf{z}_i^{'} = (z_{i1}, \ldots, z_{iq})$ be the row vector of membership counts in the q calibration groups involved in composite estimation. The associated design submatrices $\mathbf{Z}_1$ and $\mathbf{Z}_2$, of the form (3) and dimension $n_1 \times q$ and $n_2 \times q$, for the LFS and the LCF, respectively, can be built using the combined LFS and LCF data file. Composite estimation is accomplished through an extended calibration procedure, whereby, in addition to the simultaneous calibration of the LFS and the LCF described above, the estimated totals for all q common variables from the two surveys are calibrated to each other, i.e., they are aligned. This involves the augmented design matrix

$$\boldsymbol{\mathcal{X}} = \begin{pmatrix} \mathbf{X}_1 & \mathbf{0} & -\mathbf{Z}_1 \\ \mathbf{0} & \mathbf{X}_2 & \mathbf{Z}_2 \end{pmatrix}, \tag{7}$$

whose dimension is $(n_1 + n_2) \times (p_1 + p_2 + q)$, and the associated vector of calibration totals $\mathbf{t}_{\boldsymbol{\mathcal{X}}} = (\mathbf{t}_{\mathbf{x}_1}^{'}, \mathbf{t}_{\mathbf{x}_2}^{'}, \mathbf{0}^{'})^{'}$, where $\mathbf{0}$ is a vector of q zeros. In the matrix $\boldsymbol{\mathcal{X}}$, the row corresponding to the i-th household of the LFS is of the form $\boldsymbol{\mathcal{x}}_i = (x_{i1}, \cdots, x_{ip_1}, 0, \cdots, 0, -z_{i1}, \cdots, -z_{iq})$, where there are p_2 zero elements. The row corresponding to the i-th household of the LCF is of the form $\boldsymbol{\mathcal{x}}_i =$

$(0, \cdots, 0, x_{i1}, \cdots, x_{ip_2}, z_{i1}, \cdots, z_{iq})$, where there are p_1 zero elements. The specification of the additional q calibration groups should satisfy the condition that $(\mathbf{X}_i, \mathbf{Z}_i)\,, i = 1, 2$, is of full rank $p_i + q$.

Using the design matrix $\boldsymbol{\mathcal{X}}$, an extension of the simultaneous calibration for LFS and LCF described above is carried out by applying (2), using $\mathbf{w} = (\mathbf{w}_1^{'}, \mathbf{w}_2^{'})^{'}$ and the vector $\mathbf{t}_{\boldsymbol{\mathcal{X}}}$. In this case, the weighting matrix $\mathbf{W}$ in (2) is the diagonal matrix (of dimension $(n_1 + n_2) \times (n_1 + n_2)$) whose first n_1 diagonal elements are those of $\mathbf{w}_1$ multiplied by $1 - \phi$, and the remaining n_2 are those of $\mathbf{w}_2$ multiplied by ϕ, where $\phi = n_1/(n_1+n_2)$. This adjustment of the weights of the two samples accounts for the differential in the sample sizes n_1 and n_2, thus leading to more efficient estimation; see Merkouris (2004). The described calibration procedure generates a $(n_1 + n_2) \times 1$ vector $\mathbf{c} = (\mathbf{c}_1^{'}, \mathbf{c}_2^{'})^{'}$ of calibrated household weights for both surveys given by

$$\mathbf{c} = \mathbf{w} + \mathbf{W}\boldsymbol{\mathcal{X}}(\boldsymbol{\mathcal{X}}^{'}\mathbf{W}\boldsymbol{\mathcal{X}})^{-1}(\mathbf{t}_{\boldsymbol{\mathcal{X}}} - \boldsymbol{\mathcal{X}}^{'}\mathbf{w}). \tag{8}$$

It follows immediately that $\boldsymbol{\mathcal{X}}^{'}\mathbf{c} = \mathbf{t}_{\mathbf{x}}$, consolidating the calibration constraints $\mathbf{X}_1^{\prime}\mathbf{c}_1 = \mathbf{t}_{\mathbf{x}_1}$, $\mathbf{X}_2^{\prime}\mathbf{c}_2 = \mathbf{t}_{\mathbf{x}_2}$ and $\mathbf{Z}_1^{\prime}\mathbf{c}_1 = \mathbf{Z}_2^{\prime}\mathbf{c}_2$. The $n_2 \times 1$ sub-vector $\mathbf{c}_2$ is the vector of composite household weights for the LCF, and is given by

$$\mathbf{c}_2 = \mathbf{w}_2 + \mathbf{W}_2\boldsymbol{\mathcal{X}}_2(\boldsymbol{\mathcal{X}}^{'}\mathbf{W}\boldsymbol{\mathcal{X}})^{-1}(\mathbf{t}_{\boldsymbol{\mathcal{X}}} - \boldsymbol{\mathcal{X}}^{'}\mathbf{w}), \tag{9}$$

where $\mathbf{W}_2$ is the lower block diagonal submatrix of $\mathbf{W}$, and $\boldsymbol{\mathcal{X}}_2 = (\mathbf{0}\ \ \mathbf{X}_2\ \ \mathbf{Z}_2)$ is the lower submatrix of $\boldsymbol{\mathcal{X}}$ in (7). These composite weights, which incorporate information from the LFS sample, can be used to estimate any household characteristic of interest to LCF; there is no further need of LFS data for parameter estimation in the LCF. Thus for any LCF variable y, the composite regression estimate of its population total t_y is obtained as $\hat{t}_y^{CR} = \sum_{s_2} c_i y_i$, and more conveniently as $\hat{t}_y^{CR} = \mathbf{Y}_2^{'}\mathbf{c}_2$, where $\mathbf{Y}_2$ is the $n_2 \times 1$ vector of LCF sample values of y. In particular, for the vector $\mathbf{z}$ of common variables the composite regression estimate of the population total $\mathbf{t}_{\mathbf{z}}$ using LCF data is $\hat{\mathbf{t}}_{\mathbf{z}}^{CR} = \mathbf{Z}_2^{'}\mathbf{c}_2$. In composite regression form, $\hat{t}_y^{CR}$ is given by

$$\hat{t}_y^{CR} = \hat{t}_y + \hat{\mathbf{B}}\left(\mathbf{t}_{\boldsymbol{\mathcal{X}}} - \hat{\mathbf{t}}_{\boldsymbol{\mathcal{X}}}\right), \tag{10}$$

where $\hat{\mathbf{B}} = \mathbf{Y}_2^{\prime}\mathbf{W}_2\boldsymbol{\mathcal{X}}_2(\boldsymbol{\mathcal{X}}^{'}\mathbf{W}\boldsymbol{\mathcal{X}})^{-1}$ is the regression coefficient, incorporating information from both surveys, and $\hat{\mathbf{t}}_{\boldsymbol{\mathcal{X}}} = (\hat{\mathbf{t}}_{\mathbf{x}_1}^{'}, \hat{\mathbf{t}}_{\mathbf{x}_2}^{'}, \hat{\mathbf{t}}_{\mathbf{z}_2}^{'} - \hat{\mathbf{t}}_{\mathbf{z}_1}^{'})^{\prime}$. The components of $\hat{\mathbf{t}}_{\boldsymbol{\mathcal{X}}}$ are the Horvitz-Thompson estimators of $\mathbf{t}_{\mathbf{x}_1}$, $\mathbf{t}_{\mathbf{x}_2}$ and $\mathbf{t}_{\mathbf{z}}$ based on the indicated samples. Clearly, by the calibration property the composite regression estimators of these totals are $\hat{\mathbf{t}}_{\mathbf{x_1}}^{CR} = \mathbf{X}_1^{\prime}\mathbf{c}_1 = \mathbf{t}_{\mathbf{x}_1}$ and $\hat{\mathbf{t}}_{\mathbf{x_2}}^{CR} = \mathbf{X}_2^{\prime}\mathbf{c}_2 = \mathbf{t}_{\mathbf{x}_2}$. Also, $\hat{\mathbf{t}}_{\mathbf{z}}^{CR} = \mathbf{Z}_1^{\prime}\mathbf{c}_1 = \mathbf{Z}_2^{\prime}\mathbf{c}_2$.

Using a decomposition of $\mathbf{c}$ based on partitioning the matrix $\boldsymbol{\mathcal{X}}$ by the two column submatrices involving the auxiliary variables $\mathbf{x}$ and $\mathbf{z}$, the estimator $\hat{t}_y^{CR}$ can be expressed (Merkouris 2004) in terms of regression estimators involving only $\mathbf{x}$ as

$$\hat{t}_y^{CR} = \hat{t}_y^R + \hat{\boldsymbol{\mathcal{B}}}_y \left(\hat{\mathbf{t}}_{\mathbf{z}_1}^R - \hat{\mathbf{t}}_{\mathbf{z}_2}^R\right), \tag{11}$$

where $\hat{\boldsymbol{\mathcal{B}}}_y = \mathbf{Y}_2'\mathbf{L}_2\mathbf{Z}_2[\mathbf{Z}_1'\mathbf{L}_1\mathbf{Z}_1 + \mathbf{Z}_2'\mathbf{L}_2\mathbf{Z}_2]^{-1}$, with $\mathbf{L}_i = \mathbf{W}_i\left[\mathbf{I} - \mathbf{X}_i(\mathbf{X}_i'\mathbf{W}_i\mathbf{X}_i)^{-1}\mathbf{X}_i'\mathbf{W}_i\right]$, is the partial regression coefficient accounting for the additional information on $\mathbf{z}$ from the two surveys. The expression of the composite regression estimator $\hat{t}_y^{CR}$ in (11) allows a direct comparison with its single-sample counterpart $\hat{t}_y^R$. Regarding the effect of $\hat{\boldsymbol{\mathcal{B}}}_y$ on the efficiency of $\hat{t}_y^{CR}$, it is worth noting that in certain sampling designs, e.g., simple random sampling without replacement, an appropriately specified weight multiplier ϕ in $\mathbf{W}$ turns this coefficient into $\hat{\boldsymbol{\mathcal{B}}}_y = \widehat{AC}(\hat{t}_{y2}^R, \hat{\mathbf{t}}_{\mathbf{z}_2}^R)[\widehat{AV}(\hat{\mathbf{t}}_{\mathbf{z}_1}^R) + \widehat{AV}(\hat{\mathbf{t}}_{\mathbf{z}_2}^R)]^{-1}$, where $\widehat{AC}$ and $\widehat{AV}$ denote estimated approximate (large-sample) covariance and variance. In such cases $\hat{\boldsymbol{\mathcal{B}}}_y$ minimizes the estimated large-sample variance of $\hat{t}_y^{CR}$; see Merkouris (2004). For general sampling designs, the coefficient $\hat{\boldsymbol{\mathcal{B}}}_y$ serves as an approximation to the optimal one, much as in the case of single-sample regression estimation.

Similarly, for the q-dimensional common variable $\mathbf{z}$ we obtain the composite regression estimator of $\mathbf{t}_{\mathbf{z}}$, using LCF data, as the weighted average of the estimators $\hat{\mathbf{t}}_{\mathbf{z}_1}^R$ and $\hat{\mathbf{t}}_{\mathbf{z}_2}^R$

$$\hat{\mathbf{t}}_{\mathbf{z}}^{CR} = \hat{\mathbf{t}}_{\mathbf{z}_2}^R + \hat{\boldsymbol{\mathcal{B}}}_{\mathbf{z}}\left(\hat{\mathbf{t}}_{\mathbf{z}_1}^R - \hat{\mathbf{t}}_{\mathbf{z}_2}^R\right) = \hat{\boldsymbol{\mathcal{B}}}_{\mathbf{z}}\hat{\mathbf{t}}_{\mathbf{z}_1}^R + \left(\mathbf{I} - \hat{\boldsymbol{\mathcal{B}}}_{\mathbf{z}}\right)\hat{\mathbf{t}}_{\mathbf{z}_2}^R,$$

where $\hat{\boldsymbol{\mathcal{B}}}_{\mathbf{z}} = \mathbf{Z}_2'\mathbf{L}_2\mathbf{Z}_2[\mathbf{Z}_1'\mathbf{L}_1\mathbf{Z}_1 + \mathbf{Z}_2'\mathbf{L}_2\mathbf{Z}_2]^{-1}$. For the mentioned designs and with the same multiplier ϕ the coefficient $\hat{\boldsymbol{\mathcal{B}}}_{\mathbf{z}}$ turns into $\hat{\boldsymbol{\mathcal{B}}}_{\mathbf{z}} = \widehat{AV}(\hat{\mathbf{t}}_{\mathbf{z}_2}^R)[\widehat{AV}(\hat{\mathbf{t}}_{\mathbf{z}_1}^R) + \widehat{AV}(\hat{\mathbf{t}}_{\mathbf{z}_2}^R)]^{-1}$, giving $\hat{\mathbf{t}}_{\mathbf{z}}^{CR}$ as the weighted combination of $\hat{\mathbf{t}}_{\mathbf{z}_1}^R$ and $\hat{\mathbf{t}}_{\mathbf{z}_2}^R$ of minimum variance.

It is instructive to observe that (11) can be written as

$$\hat{t}_y^{CR} = \hat{t}_y^R + \mathbf{Y}_2'\mathbf{L}_2\mathbf{Z}_2\left(\mathbf{Z}_2'\mathbf{L}_2\mathbf{Z}_2\right)^{-1}\left(\hat{\mathbf{t}}_{\mathbf{z}}^{CR} - \hat{\mathbf{t}}_{\mathbf{z}_2}^R\right). \tag{12}$$

This is the regression estimator that would be obtained through an extended calibration of the LCF sample only, in which in addition to the calibration total $\mathbf{t}_{\mathbf{x}_2}$ the composite estimator $\hat{\mathbf{t}}_{\mathbf{z}}^{CR}$ would be used as calibration total. However, the described procedure of simultaneously calibrating the two samples is particularly practical in producing composite estimates and their variance estimates for all variables of LCF at national and regional level.

The described composite calibration procedure is devised to produce composite regression estimates for LCF national-level population totals that are more precise than the regression estimates that are based only on LCF data, more so for the common variables because of the direct correlation of their values from the two samples. To improve the precision of key LCF estimates at regional level, the composite calibration should incorporate the additional constraints that estimates of totals of common variables are aligned at regional level; see Merkouris (2010). Then for selected categories of the common categorical variables, the corresponding columns of the matrices $\mathbf{Z}_i$, $i = 1, 2$, are replaced by matrices with R columns, one for each of the R regions, with entries the household counts in the particular category and region, and zero entries for the rest of the regions. The associated vector of calibration totals contains as many additional zeros as the additional calibration constraints. Since regions constitute already one of the partitions of the calibration scheme, one of the R columns should be dropped to avoid collinearities in the design matrix. Then, for any of those key variables, say y, the composite regression estimate of its regional total t_{y_r}, $r = 1, \cdots, R$, based on the sample s_{2r} of region r, will be $\hat{t}_{y_r}^{CR} = \mathbf{Y}_{2r}'\mathbf{c}_2 (= \sum_{s_{2r}} c_i y_i)$, where $\mathbf{Y}_{2r}$ is the vector $\mathbf{Y}_2$ with zero entries for households not in region r.

3.3 Variance estimation for composite estimators of the LCF

An adaptation of the method of jackknife linearization for variance estimation, outlined in Section 3.1, is applied to the composite estimators of LCF totals described in Section 3.2. Thus, writing $\hat{t}_y^{CR} = \mathbf{Y}_2'\mathbf{c}_2$ equivalently as

$$\hat{t}_y^{CR} = \mathbf{Y}'\mathbf{c} = \mathbf{Y}'\mathbf{w} + \mathbf{Y}'\mathbf{W}\boldsymbol{\mathcal{X}}(\boldsymbol{\mathcal{X}}'\mathbf{W}\boldsymbol{\mathcal{X}})^{-1}(\mathbf{t_x} - \boldsymbol{\mathcal{X}}'\mathbf{w}),$$

where $\mathbf{Y}' = (\mathbf{0}, \mathbf{Y}_2')$, and following the general derivation of regression residuals in Section 3.1, it is easily verified that the vector of residuals in the present context is $\hat{\mathbf{e}} = \mathbf{Y} - \boldsymbol{\mathcal{X}}\hat{\mathbf{B}}'$, where $\hat{\mathbf{B}} = \mathbf{Y}'\mathbf{W}\boldsymbol{\mathcal{X}}(\boldsymbol{\mathcal{X}}'\mathbf{W}\boldsymbol{\mathcal{X}})^{-1} = \mathbf{Y}_2'\mathbf{W}_2\boldsymbol{\mathcal{X}}_2(\boldsymbol{\mathcal{X}}'\mathbf{W}\boldsymbol{\mathcal{X}})^{-1}$. The vector $\hat{\mathbf{e}}$ is partitioned by sample (LFS, LCF) as

$$\begin{pmatrix} \hat{\mathbf{e}}_1 \\ \hat{\mathbf{e}}_2 \end{pmatrix} = \begin{pmatrix} -\boldsymbol{\mathcal{X}}_1\hat{\mathbf{B}}' \\ \mathbf{Y}_2 - \boldsymbol{\mathcal{X}}_2\hat{\mathbf{B}}' \end{pmatrix}, \tag{13}$$

where $\boldsymbol{\mathcal{X}}_1$ is the submatrix $(\mathbf{X}_1 \;\; \mathbf{0} \;\; -\mathbf{Z}_1)$ of $\boldsymbol{\mathcal{X}}$, and the residuals at the level of sample unit (household) are given by

$$\hat{e}_i = \begin{cases} -\hat{\mathbf{B}}\boldsymbol{x}_i, & i \in s_1 \\ y_i - \hat{\mathbf{B}}\boldsymbol{x}_i, & i \in s_2 \end{cases}, \tag{14}$$

where $\boldsymbol{x}_i$ is the row vector element of $\boldsymbol{\mathcal{X}}$ corresponding to the i-th household (the form of which is described in the previous section following (7)). Then the estimated approximate (for large samples) variance of the composite estimator $\hat{t}_y^{CR}$ is the variance of the weighted sum of residuals, i.e., $\hat{V}(\sum_{s_1} w_i\hat{e}_i) + \hat{V}(\sum_{s_2} w_i\hat{e}_i)$, the additive variance over the two samples being due to the sampling independence of the LFS and the LCF. These two variance components can be calculated approximately through jackknife linearization, as outlined in Section 3.1 and following the current variance estimation practice for the LFS and the LCF, i.e., using formula (5) with the stratification and clustering of the particular survey. The residuals given by (14) should be weighted by the calibrated weights assembled in vectors $\mathbf{c}_1$ and $\mathbf{c}_2$ for s_1 and s_2, respectively. These residuals can also be used to estimate the variance of other estimated parameters of the LCF.

Note that unlike the estimation of the LCF parameters which requires only the vector of composite weights $\mathbf{c}_2$ and data on the LCF variables, and could also be carried out after the simultaneous composite calibration of the two samples, variance estimation for LCF involves the variable-dependent residuals for both LFS and LCF data. It is possible to compute variance estimates associated with any variable of the LCF later in a separate procedure if necessary, but this requires storing a large amount of information necessary to form the residuals in (14). A computationally convenient option for variance estimation of any LCF parameter using stored information produced by the composite calibration is explained in the Appendix. This could also be a useful approach for a statistical office when the calibration microdata cannot be made publicly available.

4 Methodological choices in the LCF/LFS application

4.1 Defining a common approach

In order to combine the data from the LCF and the LFS we have to make some choices. Our focus here is on reducing the variances of the LCF estimates, so the main outputs of interest are the household-level ones. In addition, although the LCF does in principle provide some additional information which could support the LFS, the small size of the LCF sample and its slower (quarterly) processing means that it is of marginal benefit and not practical to include it in the LFS. We may therefore produce a combined dataset with differences from the published LFS, since it will not be used for LFS outputs, and this allows us to choose a household-level weighting in line with the LCF. Further, we can use the regional breakdown of calibration totals from the LCF calibration, without needing the extra constraints related to local authorities in the LFS part (see section 2.2). We additionally remove cases belonging to nurses' accommodation and students, which are within the scope of the LFS but not in the scope of LCF. And finally, we use a common set of quarterly age-sex and region calibration totals for both surveys, also without nurses or students.

To investigate the questions set out in section 1 there are a number of options for which common variables to use, and we explore combinations of the following, which we expect *a priori* will be related to expenditure, the key output from the LCF:

- labour force status (employed, unemployed, inactive)
- tenure, recoded into three groups – 'owned outright', 'owned with a mortgage' and all remaining tenure categories. This last category is principally made up of people who rent their home, but also includes some other categories of tenure with small numbers of cases
- household size, coded to 1, 2 and 3 or more people

For each of these categorical variables we may include a single category in the calibration, or all categories simultaneously. In addition, all of these composite constraints can be applied at the national level, or separately within regions.

4.2 Adjustments for practical application of composite estimation

Although the theory presented in section 3 is general, some adjustments are needed for its implementation in the LCF application. The vector $\mathbf{c}$ of calibration weights derived in section 3.2 contains negative weights in some of the quarters (and most often for units of the LFS). These are not intrinsically a problem for calculation of estimates (except in small domains where they may cause negative estimates), but they do not fit with the approach to variance estimation with the linearised jackknife in section 3.3. There are several possible approaches to constrain calibration to produce positive weights, including changing the distance function in the calibration, or bounding the weights (Deville and Särndal, 1992). In our application we choose to handle this by repeating the calibration with the initial weights $\mathbf{w}$ set equal to the final weights $\mathbf{c}$ except where those are negative, in which case they are set equal to the original design weights, and then repeating the calibration. This generates a new set of calibration weights which meet the calibration constraints, are minimally different from the original calibration weights and reduce the influence of observations whose auxiliary variables result in negative weights. In some cases the procedure must be repeated before a solution with all positive weights is obtained. A procedure close to this has been used as a practical approach to calibration with positive weights by Statistics Canada (2017, section 6.3.3); in that case the iterative procedure takes place at most twice, with negative weights at the last iteration being replaced by 1. In response to some comments from a reviewer we also used the truncated weights (case 7 in Deville and Särndal, 1992) in one quarter, and found negligible differences in the estimates and variances.

5 Results

We have implemented composite calibration as described in section 3.2 using data from the LCF and LFS from 2017 Q1 to 2018 Q2, six quarters in total. These data have been collected with a promise of confidentiality, and so are not publicly available, but equivalent microdata (though without the full detail of variables needed for variance estimation) can be found in ONS and DEFRA (2020).

Table 1 shows the achieved household sample sizes for the LCF and LFS in the period under consideration.

Table 1: Quarterly achieved household sample sizes of the LCF and LFS for the periods used in this study. The LFS quarterly sample is a much larger proportion of the LFS annual sample size than the quarterly LCF is of the annual LCF sample because of the LFS rotation pattern.

year	quarter	sample size	
		LCF	LFS
2017	1	1,204	33,238
	2	1,275	33,223
	3	1,252	33,310
	4	1,402	33,693
2018	1	1,466	33,983
	2	1,240	33,898

5.1 Composite calibration for national estimates

First we explore the effect of composite calibration with a single category of the labour force status variable as the common variable. Table 2 shows the estimates and corresponding standard errors for labour force status derived from the LCF data for Q1 2017 as an example. The first row is produced with the standard calibration estimation in the LCF with age, sex and region constraints, and the following rows with the same set of constraints with the addition of composite calibration using the LFS cases. The composite estimation in turn uses only employment, then only unemployment and finally only inactivity. In each case the estimate of the variable used in the composition shows a very substantially reduced standard error, while the other variables show little (unemployment) or a smaller but still substantial (employment/inactivity) reduction in standard error. This shows the effectiveness of the composite calibration for the estimate of the variable used, but also indicates some correlation between employment and inactivity leading to smaller reductions in the corresponding standard errors.

Composite calibration with two of the labour force statuses (which effectively includes all three statuses in the calibration, because the total population size is included through the age-sex categories and the final category is implicit by differencing) achieves these reductions in the standard error on all three of the labour force status variables simultaneously. So including all levels of labour force status (and by deduction the same for any categorical variable) is a good strategy.

Table 2: Estimates of labour force status and their standard errors (in 000s) when weighting with composite estimation with a single category of the labour force status variable in addition to the standard LCF calibration.

	employment		unemployment		inactivity	
	est	s.e.	est	s.e.	est	s.e.
Base weighting:						
age, sex and region	32,128	562	1,470	243	18,744	489
Composite weighting additionally using:						
employment	31,796	117	1,517	221	19,029	238
unemployment	32,009	487	1,502	42	18,741	485
inactivity	31,834	268	1,465	243	19,044	111

We therefore undertake model selection on the variables to use as composite variables (always in addition to the standard age-sex and region calibration as used in the LCF and LFS) according to their effect on the variance estimates from the LCF, particularly the effect on expenditure, which is the key variable, although we also examine income. The best model using national level variables in the composite calibration is one with labour force status (all levels), household size (1, 2, 3+) and tenure (owned outright, owned with a mortgage coded together, the residual category (all other tenure types) obtained implicitly because the total number of households is provided by the household size variable).

Table 3 shows the reduction in variance for the national estimates from the LCF resulting from the composite calibration with this model (negative figures in the tables show a reduction in variance and are good outcomes). The full effect of the much larger quarterly sample size of the LFS is seen in the very large reductions in the variances of estimates of labour force status and tenure, both to around 5% of their original size (with some variation from quarter to quarter). This is because these variables are used explicitly in the composite calibration model. The effects on the variances of estimates of income and expenditure, however, are determined by the correlations between these variables and the variables used in the composite calibration (on top of the correlations with the base age, sex and region variables). The reduction in the variance of income is around 20% in most quarters (lower in Q2 2018), and the reduction in variance for expenditure, the LCF variable of key interest, is around 10-15% except in Q2 2017. These are all substantial reductions in the variances, which make an important contribution to the quality of the LCF outputs.

Table 3: Relative change in estimates and change in variance for national estimates from the LCF using composite calibration with labour force status, tenure and household size in addition to the base weighting by age, sex and region. Changes are relative to estimates from the base calibration of the LCF by age, sex and region.

year	quarter	Emp't	Unemp't	Inactivity	Owned outright	Owned with mortgage	Income	Expenditure
change in estimates (%)								
2017	1	-1.03	2.13	1.60	-1.69	-4.06	-0.93	-0.29
	2	-3.15	-14.76	7.16	0.34	1.20	-0.61	-0.15
	3	-3.22	-2.02	6.11	-0.79	1.04	-1.49	-0.48
	4	-0.26	-12.36	1.52	-4.70	7.52	0.11	-0.02
2018	1	0.80	-14.86	-0.06	-4.43	9.24	0.98	0.49
	2	-1.76	-18.38	4.76	-9.13	-0.35	-1.01	-1.32
change in variances (%)								
2017	1	-95.69	-96.97	-94.87	-96.61	-95.76	-26.05	-17.01
	2	-94.99	-95.42	-94.27	-96.17	-93.80	-19.25	-5.65
	3	-93.99	-95.54	-93.77	-94.77	-94.12	-21.18	-10.98
	4	-93.83	-96.18	-93.30	-93.55	-94.87	-24.18	-9.80
2018	1	-94.66	-94.73	-94.93	-93.95	-93.29	-24.34	-13.46
	2	-93.66	-96.57	-93.14	-95.10	-96.34	-10.98	-14.76

We can also examine the effects on the short time series of quarterly estimates. The largest proportional differences are seen in the smallest category, unemployment, and in Fig. 1b the time series of quarterly estimates from the LCF and the composite estimates are both quite volatile, but the trends are quite different. The trend of the composite estimates is much closer to the estimates from the LFS, making the two series approximately consistent. For employment in Fig. 1a the same consistency is introduced but the time series of quarterly changes in the composite calibration estimates are also much smoother because of the influence of the LFS data.

Fig. 2 shows the same kind of plot for expenditure (now with no LFS estimate since this variable is

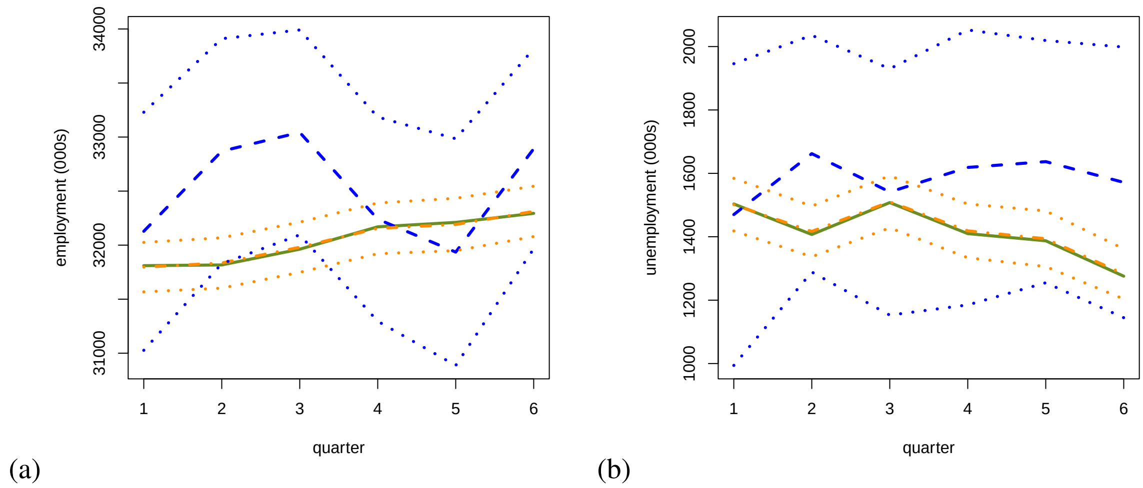

Figure 1: Time series of (a) employment and (b) unemployment estimates from LCF (dashed), LFS (solid) and from composite calibration (dash-and-dot). Normal approximation confidence intervals for the LCF and composite estimates are shown with dotted lines.

not collected on the LFS), which shows that the composite calibration has a relatively small effect on these estimates. Although the variances of the quarterly estimates are reduced, there is no noticeable impact on the changes between quarters, and the main effect of composite calibration is to improve inferences about quarterly changes rather than to generate a different pattern of expenditure in this case study.

5.2 Effect of composite calibration on the weights

We are also interested in the effect of the composite calibration on the LCF weights. The base age–sex–region calibration in the LCF involves 32 marginal constraints (some additional constraints have been used to compensate for changes in interviewing during COVID, but they do not affect the example period used here). We then add seven more constraints (two for labour force status, three for household size, and two for tenure; the third labour force status and third tenure category are implied by differencing from the population totals of people and households respectively) across the three composite variables in our preferred model. This is not a specially large number of calibration constraints considering the size of the combined LCF/LFS data, or even for the size of the LCF only. Fig 3 shows the adjustments made to the final weights. There are a very small number of weights with larger changes, but almost all changes are between 0.5 and 2 with mean very close to 1, the theoretically desirable value. Comparison of the weights without and with composition (not shown) demonstrates that there is a small increase in the variability of the weights under the composite calibration, but this is clearly not large enough to

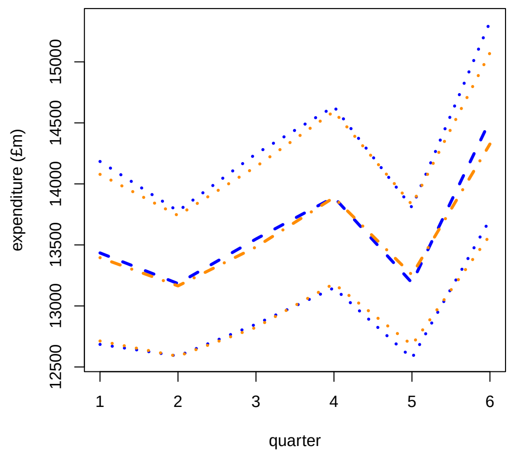

Figure 2: Time series of expenditure estimates (expressed as total weekly household expenditure in £m) from LCF (dashed) and from composite calibration (dash-and-dot). Normal approximation confidence intervals for the estimates are shown with dotted lines.

increase the variance of the estimates.

It is interesting that the extent of the variance reduction seen in table 3 is variable between quarters for income and expenditure. The larger reduction in variance for the expenditure estimate in Q1 2017 reflects that the composite calibration does not have to perturb the weights much to meet the additional composite constraints. This is shown in the leftmost boxplot in Fig. 3, which has a narrower range than those for other quarters. Two households in quarter 1 have very small but positive g-weights with the standard calibration, giving small final weights. In the first iteration of composite calibration the g-weights are negative, so the calibration is rerun starting from the final weights, which are replaced by the design weights for these (and other negative weight) cases only. The process completes at the next iteration and, as expected, the final weight is near to the design weight. The difference between the final weights from composite calibration and from the standard LCF calibration is large (≈ 13), because of the small weight of these cases under the standard calibration. These two observations are omitted from the leftmost boxplot in Figure 3.

5.3 Regional composite calibration and regional estimates

The larger sample size of the the LFS gives subnational estimates with much lower variances than the LCF, so we can also borrow strength from the LFS cases to improve the precision of the regional estimates from the LCF, by setting up the composite calibration to use the common variables at regional level, following the approach of Merkouris (2010). In the same way as before we undertake a model

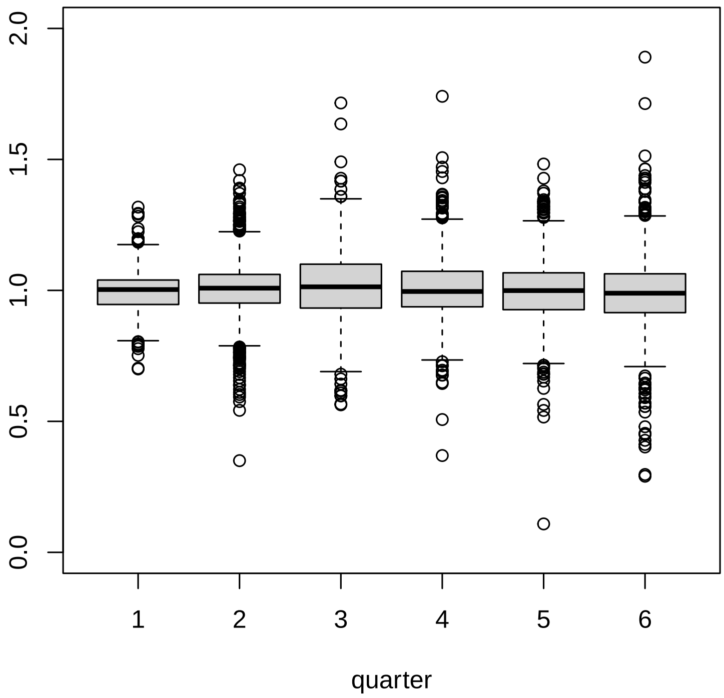

Figure 3: Boxplots of the ratio of the final weight of LCF cases after composite calibration to the final weight from the base LCF weighting for the six quarters investigated. Two extreme outlier observations of ≈ 13 have been omitted from the leftmost boxplot – see the text for an explanation of these weights.

selection process by comparing the estimated variances from the corresponding composite calibration using different combinations of regional and national constraints. This leads us to a model where the common variables used in the composite calibration are employment $\times$ region, unemployment, tenure, and household size (1, 2, 3+) $\times$ region.

Table 4 shows the effect on the relative change in variance for national estimates of income and expenditure. The gains in precision are slightly smaller in some cases and slightly larger in others when using the composite calibration with regional constraints compared with the composite calibration with national constraints. The variance under the regional model is larger than the variance before composite calibration in one quarter for each of income and expenditure.

Table 5 shows the impact on the variance of regional estimates of expenditure, which are of interest since there is a programme to assess the feasibility of developing regional consumer price indices. These indices rely on regional expenditure estimates for weighting information (Dawber et al., 2022).

In general, adding regional constraints in the composite calibration reduces the estimated variance of the regional estimates of expenditure from the LCF, and in many cases the reduction is substantial.

Table 4: Relative change (%) in variance for national estimates from national composite calibration (using labour force status, tenure and household size, taken from Table 3) and from regional composite calibration (using employment × region, unemployment, tenure, household size (1, 2, 3+) × region).

year	quarter	composite calibration			
		national		regional	
		income	expenditure	income	expenditure
2017	1	-26.05	-17.01	-19.47	-12.93
	2	-19.25	-5.65	-36.60	-8.83
	3	-21.18	-10.98	-24.46	-11.55
	4	-24.18	-9.80	-22.87	8.81
2018	1	-24.34	-13.46	-27.82	-14.05
	2	-10.98	-14.76	6.67	-9.22
2017 annual		-22.47	-11.42	-26.10	-5.70

Table 5: Relative change in variance for regional estimates of expenditure from national composite calibration ("n", using labour force status, tenure and household size) and from regional composite calibration ("r", using employment × region, unemployment, tenure, household size (1, 2, 3+) × region).

year	quarter	region							
		1		2		3		4	
		n	r	n	r	n	r	n	r
2017	1	26.59	-52.19	-2.55	-36.25	1.24	-20.17	-0.14	-5.92
	2	7.53	-23.86	0.14	-40.26	1.86	-61.77	1.63	-41.67
	3	-11.04	-29.76	-0.87	-7.21	5.46	-16.62	-1.36	-28.71
	4	-22.28	-64.92	1.65	26.35	-2.78	1.46	-7.82	-31.53
2018	1	-1.02	-5.99	-16.05	-23.28	-9.64	43.25	7.18	-50.23
	2	11.64	71.27	-7.89	-48.50	5.20	-14.45	-20.63	33.56
year	quarter	region							
		5		6		7		8	
		n	r	n	r	n	r	n	r
2017	1	-7.02	-24.31	2.37	-36.31	-4.00	6.08	1.23	-18.21
	2	0.48	-2.98	-4.80	-39.02	3.97	15.40	14.18	-7.23
	3	-0.29	-41.70	-4.14	-6.36	-7.98	-22.14	10.66	-2.36
	4	-5.73	-42.77	-10.28	3.52	-3.15	-1.44	-0.19	41.89
2018	1	-0.83	-40.18	-8.06	-24.13	-10.15	-29.61	7.31	8.29
	2	-3.38	-33.44	3.91	-16.92	-20.22	-20.47	-2.33	-1.51
year	quarter	region							
		9		10		11		12	
		n	r	n	r	n	r	n	r
2017	1	-2.14	-14.77	-1.23	-8.90	-10.42	-24.39	7.59	-31.67
	2	17.96	79.99	-0.71	-6.60	1.85	10.35	-1.44	-2.92
	3	8.00	-3.51	3.26	-18.25	-1.77	-1.21	0.25	19.77
	4	-11.81	18.26	-11.77	34.11	4.33	-28.47	-0.69	-29.86
2018	1	2.03	7.83	-9.60	-10.31	7.21	-15.70	-3.53	8.91
	2	-9.86	33.36	0.71	-64.90	-2.06	-35.21	5.39	-7.16

But there are regions and quarters where there is not much effect, or where the variance is larger under this more complex model. This is particularly so in 2017 Q4, where five of the twelve regional estimates have larger variances, leading to the larger variance of the national estimate of expenditure seen in Table 4.

The change in the weights from introducing composite calibration with common variables at regional level is larger than at national level, as greater adjustments are needed to meet the regional constraints.

5.4 LCF annual estimates

We can construct annual estimates by averaging (for stock variables) or summing (for flow variables) the quarterly estimates. The reductions in the variances of estimates of annual total income and expenditure are shown in the last row of Table 4. Although the annual variance estimates are smaller than the quarterly variances, the reduction is in line with the reductions seen for the individual quarters. The same pattern of variance reductions and consistency with LFS estimates that has already been described is also seen for the common variables included in the composite calibration, and also in the regional results (results not shown).

The LCF data have been processed a quarter at a time, using the cases from the LFS in the same quarter. This means that some LFS households have been used in more than one quarter, because they were interviewed in different waves in the different quarters. If we had actually processed a complete annual dataset, we would include almost all the available information by including only the wave 1 and wave 5 cases, and this is the method used in production of annual estimates from the LFS. The estimated annual variances may therefore be slightly too large because they do not account for the rotating panel design of the LFS, but we leave this detail for further investigation.

6 Discussion

There is a need to utilise the information from the much larger sample of the LFS to improve the estimates originating from the LCF; the described composite calibration methodology (Merkouris, 2004) is an effective approach to integrate these two sources. Using the composite calibration weights gives estimates for variables from the LCF with reduced variances. Estimates for the common variables in the composite calibration, calculated using the LCF cases and the composite calibration weights, show very substantial reductions in variance compared with estimates using the LCF-only weights, because they take direct advantage of the much larger sample size of the LFS. Furthermore, estimates for these variables derived from the LCF with composite estimation are very close to those from the LFS, and this consistency and stability over time is already valuable for users. However, the real benefits we seek are for variables which are collected only in the LCF, and in particular the expenditure variables which are widely used (ONS, 2009). Variance reductions for these variables rely on their correlation with the variables included in the calibration, and for expenditure we see smaller, but still important reductions in the variances of estimates derived from the LCF.

Most of the information on the level of the expenditure estimates, and therefore on the evolution of their time series, is derived from the data collected in the LCF. This means that the composite calibration has a limited effect on the smoothness of the time series, as seen in Fig. 2, which makes sense because the LFS contributes relatively little information about expenditure, only through the correlation between the composite calibration variables and expenditure. The main effect therefore is on the estimated variances themselves, and the main benefit is to inferences about the estimates of income and expenditure. A further possible approach to combining these series would be at an aggregate level through a multivariate state space model (Durbin and Koopman, 2012, section 3.3), but this would be practical only for one or a very small number of variables, and does not give the multipurpose solution available with a single set of weights. We note that the consistency between analyses involving multiple variables is very important to users, so this multipurpose solution is beneficial.

Introducing the composite calibration variables makes only small adjustments to the calibrated weights compared with the original survey calibration in this example. No large changes are needed in order to meet the additional constraints, and adding them has only a small effect on variance through the variability of the weights.

The reductions in the variances of the expenditure and income estimates are significant, and worth the additional complexity of processing (for which the biggest challenge is just to bring the data from the two surveys together). However, the improvements are quite variable from quarter to quarter as the LCF dataset changes, which suggests that the correlations between the expenditure and income variables and the variables used in composite calibration fluctuate too, possibly as part of the randomness due to sampling, but also possibly in a way related to the pattern of expenditure and income themselves. For example, differences between spending patterns may be different in Q4, which is influenced by Christmas spending. The improvements in the estimates of income are somewhat larger, and generally more stable, than those for expenditure.

The differences from quarter to quarter in the reduction in variances is also related to the differences in the adjustments from the base weights to the composite weights (Fig. 3). The largest reduction in variance is associated with the smallest range of the changes in weights, in the first of the quarters in our example.

The differences between the pre- and post-composite calibration variances for estimates of expenditure at the regional level seem to be much greater, and correspondingly more variable. In several cases there are large reductions in the variances for regional estimates, but these are not observed consistently across quarters, and there are examples where the variance after composite calibration is larger. Further research is needed to understand how the characteristics of the survey responses influence the variance gains from using composite calibration.

The effect on the key annual estimates is essentially the same as observed for the quarterly estimates, with reductions in variance of similar magnitude, at both national and regional level. Overall the regional estimates are substantially improved by the inclusion of a regional breakdown of the common variables in the composite calibration. There are, however, some instances where the estimated variance of the regional estimates is increased somewhat with composite calibration, and this effect is also seen in the annual estimates.

We have shown that integrating two data sources using composite calibration can result in a substantial reduction in the size of the variances of important estimates in a household expenditure survey. This approach has a number of advantages, including that it improves the consistency of outputs of the two surveys on the common variables used in the composite calibration (in principle they could be made completely consistent, but timing and other issues make improved consistency a more practical objective), it accounts straightforwardly and appropriately for the uncertainty in estimation and variance estimation, and it provides a general solution with weights that can be used for all variables. The extent of the reduction in the variances of estimates depends on the correlation between the variable of interest and the common variables in the composite calibration. The results here illustrate that this effect varies between variables and across periods.

In this study we have demonstrated that data integration involving large-scale national surveys, through composite calibration, is a valuable tool to use existing survey information as efficiently as possible. It improves the quality of the outputs through a reduction of the variance estimates. It has wider potential for use in other similar situations, and we encourage other applications of this approach to integrate survey sources with common variables.

Appendix

A computationally convenient option for variance estimation of any LCF parameter using stored information produced by the composite calibration is based on the representation of the composite estimator $\hat{t}_y^{CR} = \hat{t}_y^R + \hat{\boldsymbol{\mathcal{B}}}_y \left(\hat{\mathbf{t}}_{\mathbf{z}_1}^R - \hat{\mathbf{t}}_{\mathbf{z}_2}^R\right)$ in (11), where $\hat{t}_y^R = \hat{t}_y + \hat{\mathbf{B}}_{y/\mathbf{x}}(\mathbf{t}_{\mathbf{x}_2} - \hat{\mathbf{t}}_{\mathbf{x}_2})$ is the LCF regression (calibration) estimator of t_y, using LCF data and the vector of the LCF auxiliary variables $\mathbf{x}_2$, and $\hat{\mathbf{B}}_{y/\mathbf{x}} = \sum_{s_2} w_i y_i \mathbf{x}_i^{'}(\sum_{s_2} w_i \mathbf{x}_i \mathbf{x}_i^{'})^{-1}$. Also, $\hat{\mathbf{t}}_{\mathbf{z}_2}^R = \hat{\mathbf{t}}_{\mathbf{z}_2} + \hat{\mathbf{B}}_{\mathbf{z}/\mathbf{x}}^{(2)}(\mathbf{t}_{\mathbf{x}_2} - \hat{\mathbf{t}}_{\mathbf{x}_2})$ is the LCF regression estimator of the total $\mathbf{t}_{\mathbf{z}}$, where $\hat{\mathbf{t}}_{\mathbf{z}_2}$ and $\hat{\mathbf{t}}_{\mathbf{x}_2}$ are the Horvitz-Thompson estimators defined before, and $\hat{\mathbf{B}}_{\mathbf{z}/\mathbf{x}}^{(2)} = \sum_{s_2} w_i \mathbf{z}_i \mathbf{x}_i^{'}(\sum_{s_2} w_i \mathbf{x}_i \mathbf{x}_i^{'})^{-1}$, and $\hat{\mathbf{t}}_{\mathbf{z}_1}^R = \hat{\mathbf{t}}_{\mathbf{z}_1} + \hat{\mathbf{B}}_{\mathbf{z}/\mathbf{x}}^{(1)}(\mathbf{t}_{\mathbf{x}_1} - \hat{\mathbf{t}}_{\mathbf{x}_1})$ is the LFS counterpart of $\hat{\mathbf{t}}_{\mathbf{z}}^R$, with $\hat{\mathbf{B}}_{\mathbf{z}/\mathbf{x}}^{(1)} = \sum_{s_1} w_i \mathbf{z}_i \mathbf{x}_i^{'}(\sum_{s_1} w_i \mathbf{x}_i \mathbf{x}_i^{'})^{-1}$. The expression of the vector $\hat{\boldsymbol{\mathcal{B}}}_y$ is given following (11).

Now, collecting terms of $\hat{t}_y^{CR}$ from the two independent samples s_1 and s_2, and using their expanded form given above, we obtain for the estimated variance of $\hat{t}_y^{CR}$

$$\begin{aligned}
\hat{V}(\hat{t}_y^{CR}) &= \hat{V}\left(\hat{t}_y^R - \hat{\boldsymbol{\mathcal{B}}}_y \hat{\mathbf{t}}_{\mathbf{z}_2}^R\right) + \hat{V}\left(\hat{\boldsymbol{\mathcal{B}}}_y \hat{\mathbf{t}}_{\mathbf{z}_1}^R\right) \\
&= \hat{V}\left(\hat{t}_y + \hat{\mathbf{B}}_{y/\mathbf{x}}(\mathbf{t}_{\mathbf{x}_2} - \hat{\mathbf{t}}_{\mathbf{x}_2}) - \hat{\boldsymbol{\mathcal{B}}}_y\left(\hat{\mathbf{t}}_{\mathbf{z}_2} + \hat{\mathbf{B}}_{\mathbf{z}/\mathbf{x}}^{(2)}(\mathbf{t}_{\mathbf{x}_2} - \hat{\mathbf{t}}_{\mathbf{x}_2})\right)\right) \\
&+ \hat{V}\left(\hat{\boldsymbol{\mathcal{B}}}_y\left(\hat{\mathbf{t}}_{\mathbf{z}_1} + \hat{\mathbf{B}}_{\mathbf{z}/\mathbf{x}}^{(1)}(\mathbf{t}_{\mathbf{x}_1} - \hat{\mathbf{t}}_{\mathbf{x}_1})\right)\right).
\end{aligned}$$

Next, retaining the necessary terms for linearized variance estimation we obtain the approximate estimated variance

$$\begin{aligned}
\hat{V}(\hat{t}_y^{c}) &\approx \hat{V}\left(\hat{t}_y - \hat{\mathbf{B}}_{y/\mathbf{x}}\hat{\mathbf{t}}_{\mathbf{x}_2} - \hat{\boldsymbol{\mathcal{B}}}_y\left(\hat{\mathbf{t}}_{\mathbf{z}_2} - \hat{\mathbf{B}}_{\mathbf{z}/\mathbf{x}}^{(2)}\hat{\mathbf{t}}_{\mathbf{x}_2}\right)\right) + \hat{V}\left(\hat{\boldsymbol{\mathcal{B}}}_y\left(\hat{\mathbf{t}}_{\mathbf{z}_1} - \hat{\mathbf{B}}_{\mathbf{z}/\mathbf{x}}^{(1)}\hat{\mathbf{t}}_{\mathbf{x}_1}\right)\right) \\
&= \hat{V}\left(\sum_{s_2} w_i\left(y_i - \hat{\mathbf{B}}_{y/\mathbf{x}}\mathbf{x}_i - \hat{\boldsymbol{\mathcal{B}}}_y\left(\mathbf{z}_i - \hat{\mathbf{B}}_{\mathbf{z}/\mathbf{x}}^{(2)}\mathbf{x}_i\right)\right)\right) \\
&+ \hat{V}\left(\sum_{s_1} w_i \hat{\boldsymbol{\mathcal{B}}}_y\left(\mathbf{z}_i - \hat{\mathbf{B}}_{\mathbf{z}/\mathbf{x}}^{(1)}\mathbf{x}_i\right)\right) \\
&= \hat{V}\left(\sum_{s_2} w_i \hat{e}_i\right) + \hat{V}\left(\sum_{s_1} w_i \hat{e}_i\right),
\end{aligned}$$

where the "composite" residuals $\hat{e}_i$ are the linear transformations of the elementary residuals $\hat{e}_{y_i}$ and $\hat{e}_{\mathbf{z}_i}$

$$
\begin{aligned}
\hat{e}_i &= \begin{cases} \hat{\boldsymbol{\mathcal{B}}}_y \hat{e}_{\mathbf{z}_i}, & i \in s_1 \\ \hat{e}_{y_i} - \hat{\boldsymbol{\mathcal{B}}}_y \hat{e}_{\mathbf{z}_i}, & i \in s_2 \end{cases} \\
&= \begin{cases} \hat{\boldsymbol{\mathcal{B}}}_y \left(\mathbf{z}_i - \hat{\mathbf{B}}^{(1)}_{\mathbf{z}/\mathbf{x}} \mathbf{x}_i \right), & i \in s_1 \\ y_i - \hat{\mathbf{B}}_{y/\mathbf{x}} \mathbf{x}_i - \hat{\boldsymbol{\mathcal{B}}}_y \left(\mathbf{z}_i - \hat{\mathbf{B}}^{(2)}_{\mathbf{z}/\mathbf{x}} \mathbf{x}_i \right), & i \in s_2 \end{cases} .
\end{aligned}
\tag{15}
$$

The vector of these composite residuals, partitioned by sample, is expressed as

$$
\begin{pmatrix} \hat{\mathbf{e}}_1 \\ \hat{\mathbf{e}}_2 \end{pmatrix} = \begin{pmatrix} \left(\mathbf{Z}_1 - \mathbf{X}_1 \hat{\mathbf{B}}^{(1)'}_{\mathbf{z}/\mathbf{x}} \right) \hat{\boldsymbol{\mathcal{B}}}_y^{'} \\ \mathbf{Y}_2 - \mathbf{X}_2 \hat{\mathbf{B}}^{'}_{y/\mathbf{x}} - \left(\mathbf{Z}_2 - \mathbf{X}_2 \hat{\mathbf{B}}^{(2)'}_{\mathbf{z}/\mathbf{x}} \right) \hat{\boldsymbol{\mathcal{B}}}_y^{'} \end{pmatrix}, \tag{16}
$$

using the vector expressions $\hat{\mathbf{B}}_{y/\mathbf{x}} = \mathbf{Y}_2^{'} \mathbf{W}_2 \mathbf{X}_2 (\mathbf{X}_2^{'} \mathbf{W}_2 \mathbf{X}_2)^{-1}$, $\hat{\mathbf{B}}^{(1)}_{\mathbf{z}/\mathbf{x}} = \mathbf{Z}_1^{'} \mathbf{W}_1 \mathbf{X}_1 (\mathbf{X}_1^{'} \mathbf{W}_1 \mathbf{X}_1)^{-1}$ and $\hat{\mathbf{B}}^{(2)}_{\mathbf{z}/\mathbf{x}} = \mathbf{Z}_2^{'} \mathbf{W}_2 \mathbf{X}_2 (\mathbf{X}_2^{'} \mathbf{W}_2 \mathbf{X}_2)^{-1}$. It may be computationally easier to generate the residuals in the vector form of (16).

Note that the LCF variable of interest y enters the calculation of the residuals through the vector of coefficients $\hat{\boldsymbol{\mathcal{B}}}_y$ and $\hat{\mathbf{B}}_{y/\mathbf{x}}$. Then variance estimation for the LCF as a procedure following the composite calibration can be carried out as follows. First, the $n_1 \times q$ matrix $\left(\mathbf{Z}_1 - \mathbf{X}_1 \hat{\mathbf{B}}^{(1)'}_{\mathbf{z}/\mathbf{x}} \right)$ of the residuals $\hat{e}_{\mathbf{z}_i}$, produced in a step of the composite calibration procedure, is multiplied by the $q \times q$ matrix $\mathbf{A} = \left[\mathbf{Z}_1^{'} \mathbf{L}_1 \mathbf{Z}_1 + \mathbf{Z}_2^{'} \mathbf{L}_2 \mathbf{Z}_2 \right]^{-1}$, also produced during the composite calibration. The resulting $n_1 \times q$ matrix, say $\mathbf{C}$, associated with the LFS sample s_1, is saved for later use. Next, for any LCF variable y the $q \times 1$ vector $\mathbf{Z}_2^{'} \mathbf{L}_2 \mathbf{Y}_2$, generated by LCF data only, is multiplied by the matrix $\mathbf{C}$ to give the $n_1 \times 1$ vector $\hat{\mathbf{e}}_1$ in (16). The $n_2 \times 1$ vector $\left(\mathbf{Z}_2 - \mathbf{X}_2 \hat{\mathbf{B}}^{(2)'}_{\mathbf{z}/\mathbf{x}} \right) \hat{\boldsymbol{\mathcal{B}}}_y^{'}$ is produced in a similar manner. Finally the $n_2 \times 1$ vector $\mathbf{Y}_2 - \mathbf{X}_2 \hat{\mathbf{B}}^{'}_{y/\mathbf{x}}$ of the residuals $\hat{e}_{y_i}$ is computed easily using the LCF data, and then the vector $\hat{\mathbf{e}}_2$ in (16) is generated in an obvious way.

The estimated variances produced by the residuals given by (16) are the same variances produced by the residuals given by (14).

Acknowledgements

We are grateful to Jo Bulman and Steve Martin-Drury who were involved in an earlier phase of this project. The work of Takis Merkouris and Paul Smith was supported by ONS under contract PU-16-0031-6.009. The views expressed in this article are those of the authors, and do not necessarily reflect

those of the ONS or their respective organisations.

References

Berger, Y. G. and Kabzińska, E. (2020). Empirical likelihood approach for aligning information from multiple surveys. *International Statistical Review*, 88(1):54–74.

Brewer, M., Etheridge, B., and O'Dea, C. (2017). Why are households that report the lowest incomes so well-off? *The Economic Journal*, 127:F24–F49.

Burton, J., Crossley, T. F., Fisher, P., Gaia, A., and Jäckle, A. (2020). Understanding and reducing income reporting error in household surveys. Understanding society working paper series no. 2020-01, ISER, University of Essex.

Chen, S. and Kim, J. K. (2014). Population empirical likelihood for nonparametric inference in survey sampling. *Statistica Sinica*, 24(1):335–355.

Dawber, J., Würz, N., Smith, P. A., Flower, T., Thomas, H., Schmid, T., and Tzavidis, N. (2022). Experimental UK regional consumer price inflation with model-based expenditure weights. *Journal of Official Statistics*, 38:213–237.

Deville, J.-C. and Särndal, C.-E. (1992). Calibration estimators in survey sampling. *Journal of the American Statistical Association*, 87(418):376–382.

Durbin, J. and Koopman, S. J. (2012). *Time series analysis by state space methods*. Oxford University Press, Oxford, 2nd edition.

Eckman, S. (2022). Underreporting of purchases in the US Consumer Expenditure Survey. *Journal of Survey Statistics and Methodology*, 10(5):1148–1171.

Eurostat (2020). Household Budget Survey, 2015 Wave. EU Quality Report, European Commission, Luxembourg. `https://ec.europa.eu/eurostat/documents/54431/1966394/HBS_EU_QualityReport_2015.pdf`.

Hidiroglou, M. (2001). Double sampling. *Survey Methodology*, 27(2):143–154.

Holmes, D. J. and Skinner, C. J. (2000). Variance estimation for labour force survey estimates of level and change. GSS Methodology Series no. 21, Government Statistical Service. `https:`

//webarchive.nationalarchives.gov.uk/ukgwa/20110824221331/http://www.statistics.gov.uk/StatBase/Product.asp?vlnk=9220.

Jäckle, A., Couper, M. P., Gaia, A., and Lessof, C. (2021). Improving survey measurement of household finances: A review of new data sources and technologies. In Lynn, P., editor, *Advances in Longitudinal Survey Methodology*, pages 337–367. Wiley, Hoboken.

Kim, J. K. and Rao, J. N. K. (2012). Combining data from two independent surveys: a model-assisted approach. *Biometrika*, 99(1):85–100.

Lemaître, G. and Dufour, J. (1987). An integrated method for weighting persons and families. *Survey Methodology*, 13(2):199–207.

Merkouris, T. (2004). Combining independent regression estimators from multiple surveys. *Journal of the American Statistical Association*, 99(468):1131–1139.

Merkouris, T. (2010). Combining information from multiple surveys by using regression for efficient small domain estimation. *Journal of the Royal Statistical Society: Series B*, 72(1):27–48.

Office for Statistics Regulation (2021). Assessment of compliance with the code of practice for statistics: The living costs and food survey. Assessment report 358, Office for Statistics Regulation, Newport. https://osr.statisticsauthority.gov.uk/publication/assessment-report-the-living-costs-and-food-survey/.

ONS (2009). Summary quality report for the Living Costs and Food Survey (LCF). , Office for National Statistics, Newport. https://webarchive.nationalarchives.gov.uk/ukgwa/20160105160709/http://www.ons.gov.uk/about-statistics/methodology-and-quality/quality/qual-info-economic-social-and-bus-stats/quality-reports-for-social-statistics/sum-rpt-exp-food-surv.pdf.

ONS (2015). Labour Force Survey (LFS) QMI. , Office for National Statistics, Newport. https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/methodologies/labourforcesurveylfsqmi.

ONS (2019). Family spending in the UK: April 2017 to March 2018. , Office for National Statistics, Newport. https://www.ons.gov.uk/peoplepopulationandcommunity/

personalandhouseholdfinances/expenditure/bulletins/familyspendingintheuk/financialyearending2018.

ONS (2020). CPIH-consistent inflation rate estimates for UK household groups: October to December 2019. , Office for National Statistics, Newport. https://www.ons.gov.uk/economy/inflationandpriceindices/articles/cpihconsistentinflationrateestimatesforukhouseholdgroups20052017/octobertodecember2019.

ONS (2021). Labour Force Survey user guide vol. 1: LFS background and methodology. , Office for National Statistics, Newport. https://www.ons.gov.uk/file?uri=/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/methodologies/labourforcesurveyuserguidance/lfsuserguidevolume1oct2020.pdf.

ONS (2022). CPI-consistent inflation rate estimates for UK household groups: 2005 to 2021. , Office for National Statistics, Newport. https://www.ons.gov.uk/economy/inflationandpriceindices/bulletins/cpiconsistentinflationrateestimatesforukhouseholdgroups/2005to2021.

ONS and DEFRA (2020). Living costs and food survey, 2017-2018 [data collection], 3rd edition. UK Data Service SN: 8459. doi: 10.5255/UKDA-SN-8459-3.

Ralph, J. and Manclossi, S. (2016). Living Costs and Food Survey. National Statistics Quality Review: Series 2., Office for National Statistics, Newport.

Renssen, R. H. and Nieuwenbroek, N. J. (1997). Aligning estimates for common variables in two or more sample surveys. *Journal of the American Statistical Association*, 92(437):368–374.

Smith, P. (2009). Survey harmonisation in official household surveys in the United Kingdom. *Proceedings of the ISI World Statistical Congress 2009*.

Statistics Canada (2017). Methodology of the Canadian Labour Force Survey. , Statistics Canada, Ottawa. https://www150.statcan.gc.ca/n1/en/pub/71-526-x/71-526-x2017001-eng.pdf.

Statistics Canada (2021). Survey of household spending (SHS). , Statistics Canada, Ottawa. https://www23.statcan.gc.ca/imdb/p2SV.pl?Function=getSurvey&SDDS=3508.

Steel, D. (1997). Producing monthly estimates of unemployment and employment according to the International Labour Office definition (with discussion). *Journal of the Royal Statistical Society: Series A*, 160(1):5–46.

Wu, C. (2004). Combining information from multiple surveys through the empirical likelihood method. *Canadian Journal of Statistics*, 32(1):15–26.

Yang, S. and Kim, J. K. (2020). Statistical data integration in survey sampling: A review. *Japanese Journal of Statistics and Data Science*, 3(2):625–650.

Yung, W. and Rao, J. N. K. (1996). Jackknife linearization variance estimators under stratified multi-stage sampling. *Survey Methodology*, 22:23–32.