

Analysis of data from a national micronutrient survey with a linear mixed model: estimates, predictions and lessons for future surveys

Journal of Public Health Research
2024, Vol. 13(4), 1–13
© The Author(s) 2024
DOI: 10.1177/22799036241274962
journals.sagepub.com/home/phj



Hakunawadi Alexander Pswarayi¹ , Edward J M Joy²,
Dawd Gashu³, Fanny Sandalinas², Adamu Belay⁴ and R Murray Lark¹

Abstract

Background: Because micronutrient deficiencies affect public health, countries monitor population status by national-scale, multi-stage, micronutrient surveys (MNS). In design-based surveys, inclusion probabilities are specified for sample units and the corresponding sample weights allow design-unbiased estimates to be made of population parameters. Corrections may be possible on departures from the design; an alternative is to use linear mixed models (LMM), with an estimated covariance structure reflecting the sampling design, to obtain model-based estimates.

Design: The Ethiopia National Micronutrient Survey (2016) specified inclusion probabilities at enumeration area (EA) and household (HH) levels, and sample weights are provided. However, the design was not followed as it would have resulted in insufficient sampling from women of reproductive age.

Results: Having found no evidence that sample weights were informative for target serum micronutrient concentrations (Zn), we estimated LMM parameters, with Regions as fixed effects, and the variation of individuals nested within households, households within EA, and EA within regions, random effects. We obtained LMM standard errors, Best Linear Unbiased Estimates (BLUEs) of regional means, and empirical Best Linear Unbiased Predictions for sampled/unsampled EA and HH. The probability that each true regional mean exceeded the sufficiency threshold ($65\mu\text{g dL}^{-1}$) was evaluated. The variances of BLUEs of regional means, under alternative sampling designs, were bootstrapped from LMM variance components.

Conclusions: We demonstrate use of LMM to obtain model-unbiased estimates and predictions when surveys deviate from the original design; and the use of LMM variance components to evaluate alternative designs for further sampling, or for sampling comparable populations.

Keywords

Micronutrient, survey, linear mixed model, sample weight, estimates, prediction, BLUE, EBLUP, inclusion probability

Date received: 7 December 2023; accepted: 27 July 2024

Introduction

Micronutrient deficiencies are a serious challenge for global public health, and contribute, *inter alia*, to impaired physical and cognitive development, susceptibility of children to pneumonia and diarrhoea and maternal mortality.^{1–4} This paper is specifically about zinc (Zn) deficiency. Zinc is essential to the human body as it is involved in many metabolic processes⁵ that include cell division, cell growth and differentiation, protein and DNA synthesis, RNA metabolism, immune function and wound healing.⁶

¹School of Biosciences, University of Nottingham, Sutton Bonington Campus, Sutton Bonington, Loughborough, UK

²London School of Hygiene & Tropical Medicine Faculty of Epidemiology and Population Health, London, UK

³Addis Ababa University, Addis Ababa, Ethiopia

⁴Center for Food Science and Nutrition, Addis Ababa University, Addis Ababa, Ethiopia

Corresponding author:

Hakunawadi Alexander Pswarayi, University of Nottingham, Sutton Bonington Campus, Loughborough, LE12 5RD, UK.
Email: hakunawadi.pswarayi@nottingham.ac.uk



Creative Commons CC BY: This article is distributed under the terms of the Creative Commons Attribution 4.0 License

(<https://creativecommons.org/licenses/by/4.0/>) which permits any use, reproduction and distribution of the work without further permission provided the original work is attributed as specified on the SAGE and Open Access pages (<https://us.sagepub.com/en-us/nam/open-access-at-sage>).

In addition, Zn supports normal growth and development during pregnancy and childhood.⁷ Zinc deficiency is associated with high morbidity and mortality among mothers and new born infants, as well as increased risk of diarrhoea in children under 5 years of age.⁸ It is believed about 17% of the global population experience Zn deficiency, with a greater prevalence in developing countries in South and South-East Asia, Sub-Saharan Africa and Central America.⁸ Zinc sufficiency is indicated by a concentration of Zn in blood serum of above $65\mu\text{g dL}^{-1}$ for adult women, in a non-fasted sample collected in the morning.⁵

Because of the importance of micronutrients in public health, it is valuable to undertake national scale micronutrient surveys (MNS). Such surveys have been undertaken, for example in Malawi as part of the Demographic and Health Survey for 2015/16.^{9,10} The results of these surveys are used by many stakeholders, and notably by national government authorities to support the implementation of policies and programs to control micronutrient deficiencies. In this paper we examine data from the Ethiopian Micronutrient Survey (EMNS).¹¹ Archived blood serum samples from EMNS have been analysed for Zn concentration because of the importance of this micronutrient,¹² and we focused on data for women of reproductive age (WRA) which is the largest demographic group in the sample.

The national-scale micronutrient surveys mentioned above are design-based multistage sampling (MSS) surveys. A design-based sample¹³ is one in which the inclusion of a particular unit from the sample frame in the final sample is a random event. In advance of randomization we do not know whether a particular unit will be included, but we do know the probability of this, the *inclusion probability*. Multistage designs consists of two or more stages of random sampling that typically follow the hierarchical structure of naturally occurring clusters, with a different type of cluster randomly sampled at each stage of sampling; the clusters are nested within each other.¹⁴ An example of clusters nested within each other are households (HHs) randomly selected from within enumeration areas (EAs) which are randomly selected within regions. An Enumeration Area is a sampling unit which comprises some grouping of households. All households in the region of interest belong to exactly one EA. In a census an EA might correspond to a region which one enumerator can cover in the sampling period, but in other household-based surveys it might be smaller. In the Ethiopia MNS studied here an EA in a rural area typically comprises 150–200 households. Once an EA is selected for sampling it is feasible to enumerate households prior to sampling, and to undertake sensitization at EA scale. The use of MSS surveys with HH grouped in EA therefore increases logistical efficiency.

In MSS sampling designs inclusion probabilities are specified at different levels of the sampling. For example, each EA in a sampled district will have an inclusion probability, π_{EA_j} for the J^{th} EA. The HH within the EA may then have inclusion probabilities assigned, e.g. $\pi_{HH_{j,i}}$ for the i^{th} HH in the J^{th} EA. If a single individual is sampled in each HH, then the inclusion probability associated with the sample from the i^{th} HH in the J^{th} EA is $\pi_{EA_j}\pi_{HH_{j,i}}$. Inclusion probabilities at some level are all equal in the case of simple random sampling, but it is common practice for π_{EA_j} to be proportional to the size of the J^{th} EA, measured, for example, by its population or the number of HH which it contains. This is called sampling with probability proportional to size (PPS). More information on sample weights is given in Appendix 1.

Design-based sampling, which incorporates the inclusion probabilities, allows unbiased estimates of population parameters. This is typically done by reference to sample weights. The sample weight for the i^{th} HH in the J^{th} EA is

$$\frac{1}{\pi_{EA_j}\pi_{HH_{j,i}}}.$$

The sample weight corresponds to the number of population units which the corresponding sample can be regarded as representing. If the sample weight is multiplied by the value of the observation for the corresponding sample unit, and the product is summed over the whole sample then the result is an unbiased estimate of the population total. If the weights are normalized to sum to 1 over the sample then an unbiased estimate is obtained of the sample mean. These weights, normalized or unnormalized, reflect the structure of the sample frame and the sample design can be specified before the survey is undertaken or any data are available. The same weights would be applied for estimation from any variable measured on the sample units.

It is not unusual, particularly when a sampling exercise is being undertaken for the first time, that the original design sample weights do not apply to the complete final sample. This may be due to factors such as refusal of a sampled HH to participate, or loss of data. In such cases, post-hoc adjustments may be made to the sample weights, for example, to adjust for non-response.¹⁵ However, the post-hoc adjustment of weights is not without problems.¹⁶ In the example of the Ethiopia MNS which we report here, substantial changes were made to the sample survey in the field, and in such instances, we propose using an alternative approach to derive estimates, based on a linear mixed model (LMM).

In LMM we treat our data as a combination of fixed effects which determine the mean of the dependent variable, and random effects which account for variation about the mean. In a MSS, for example, EA and HH means are modelled as random variables of mean zero which account

for differences between EA means and the mean expected from the fixed effects, and HH means and the EA mean. Their variances are key model parameters. By virtue of their design, MSS designs collect data from purported clusters (e.g., regions, EAs within regions, and HHs within EAs within regions). Although many individual factors determine an individual's micronutrient status, we may expect some degree of correlation between the biomarkers for individuals within an EA due to causes which show less variation within an EA than between EAs. For example, individuals within the same EA may have a more similar diet than individuals in different EAs, with staple crops grown over soils with similar micronutrient status. Access to health care, fortification programmes and more localized food sources like fish are also expected to vary between EAs, but to be relatively uniform within them. This is reflected in the LMM by the random EA mean which contributes to modelled values within any EA, and to larger variances at between-EA scale than within-EA. Similarly, individuals within the same HH are likely to be a cluster of genetically related people, with a common diet and environment. While the use of appropriate sample weights would allow us to avoid bias from two-stage surveys, the alternative model-based approach of linear mixed modelling (LMM) is to propose a covariance structure which reflects the nested structure of the sample, and then to estimate its parameters by residual maximum likelihood (REML). These parameters then provide what is in effect a weighting of the observations for the fixed effects from the data. Where this weighting differs from the sample weights referred to above is that they depend on the observed between EA and between HH within EA correlation of the observations, and so are variable-specific and not known until the data are collected and modelled. The design-based method has the advantage of being unbiased and free from model-type assumptions (i.e. that the data are a realization of a multivariate normal random model), but where the model assumptions are plausible the use of model information in the weighting can be an advantage.

We may expect a design-based analysis of a set of survey data to give an unbiased estimate of the sample mean, and, with suitable estimators, a meaningful characterization of the uncertainty of the mean in its standard error. However, an LMM with a random effects structure which is correctly specified so that it reflects the survey design also provides a basis to estimate the mean, and to characterize the uncertainty with a standard error based on the fitted model.¹⁷ To illustrate this we undertook a simulation exercise. A large population of values was simulated, comprising a total of 42 913 HH in 991 EA (these were random quantities, the example can be reproduced with R code in the supplementary material). These were then sampled by an MSS design with 50 EA selected by simple random sampling, and 100 HH per EA, also by simple random sampling. The mean value and its standard error was estimated by design-based

Table 1. Simulation results.

Mean type	Mean	Standard error	Coverage (95%)
Design	99.3	2.6	0.93–0.96
Model	98.8	2.6	0.93–0.96
Naive	98.9	0.8	0.42–0.48

estimation using the `svymean` function from the survey package for the R platform.¹⁸ The LMM mean and standard error were estimated from the same sample, using the `lme` function from the `nlme` library.¹⁹ For both estimates it was observed whether the (known) population mean was included in the 90% confidence interval. This procedure was repeated 1000 times, and the proportion of cases in which the confidence interval included the population mean was computed as an estimate of the coverage probability of the interval. The 95% confidence interval for the coverage estimates was computed by the method of (Blaker)²⁰ using the `blakerci` function from the `PropCIs` library.²¹ The coverage estimates were both close to the specified probability (90%) which was included in the confidence interval in each case (Table 1). This demonstrates that both the design and model-based approaches give reliable results for the analysis of MSS data, provided that the survey information is known for the design-based procedure. As the model does not depend on the survey weights, the LMM provides an alternative when the true survey weights are unknown due to departures from survey design in the field.

Linear mixed models (LMMs) with REML parameter estimation are also useful when there are missing values in survey data. In fact, REML was developed specifically for the task of recovering information from data sets which are unbalanced because of missing observations,²² with the missing values treated as a random process. Missing values are a common phenomenon in HH surveys due to non-response (e.g. one might be less likely to find people whose work takes them far from home), or refusals at individual level within households. Non-response is considered an increasing problem in HH surveys,²³ and may exceed 50%.²⁴ Non-participation in a survey may introduce bias, for example if households not engaged with current outreach programmes are more likely than average to refuse participation).

In an MSS hierarchical sampling design, means at the level where all units are sampled can be treated as fixed effects. For example, at the regional level in the Ethiopian National Micronutrient Survey (ENMS) examined in this paper. These mean values may be of direct interest in themselves, for example, used to prioritize Zn interventions among regions by ranking the regional mean serum Zn concentration for women of reproductive age (WRA) of Ethiopia. Estimates of these means have error variances, which can be obtained from the LMM, and used to obtain confidence intervals (CIs).

The motivation of this paper is to show how LMM can be used to make population parameter estimates from national-scale surveys that departed in practice from the original sample plan. The idea being to raise awareness among of those involved in micronutrient research, for example, data analysts, policy makers, non-governmental organizations (NGOs), etc., who encounter such data in their day-to-day work, and subsequently to increase the accuracy of estimates of micronutrient status with robust estimates of uncertainty. We describe the data that departed from the sample plan (ENMS), highlighting how it departed from the plan. And we show how to (a) specify the LMM, (b) interpret the estimated regional means, understood either as estimates or as predictions of the means of unsampled included units and (c) calculate and interpret CIs and PIs. We also explain how variance components estimated for EAs and HHs can be used to improve the efficiency of future national-scale surveys. Finally, we consider how estimates and predictions for biomarker values from sample surveys can be compared with nutritionally relevant threshold values to support intervention decisions, while quantifying the inevitable uncertainty.

Methods

Ethiopian National Micronutrient Survey (ENMS) description

We used data on the concentration of Zn in blood serum of WRA sampled in the ENMS.¹¹ The objectives of the survey were to support estimation of prevalence, at National scale, of anaemia, deficiencies of six vitamins and minerals, and the proportion of households with adequately iodized salt.¹¹ The specific biomarker data analysed here were not in the original EMNS, but were obtained from archived samples.¹² Serum Zn concentration was measured by inductively coupled plasma mass spectrometry (ICP-MS), which was considered more reliable than the method used to measure serum Zn in the original EMNS. The serum Zn concentrations were adjusted for inflammation, using the procedure from the BRINDA project.²⁵ We chose to analyse the WRA demographic group because they were the largest sample (1181) across all regions of Ethiopia (Table 2).

Sampling in the EMNS was by a simple hierarchical design, with EAs nested within regions and city administrations, and HHs nested within EAs within regions. All regions and city administrations were sampled. Therefore, the key elements of the sampling design were EAs and HHs. The sampled populations were in the nine regions and two city administrations, comprising pre-school children 6–59 months (PSC), school-age children 5–14 years (SAC), non-pregnant women of reproductive age 15–49 years (WRA) and men (Men).

The regions and city administrations, which were treated as strata for the survey, were not sampled

Table 2. Summary table of sample.

Region	EAs	HHs	WRA
Addis Ababa	33	118	119
Afar	27	165	80
Amhara	41	299	181
Benshangul-Gumuz	24	149	86
Dire Dawa	27	135	72
Gambella	22	127	83
Harari	24	117	64
Oromia	46	331	188
SNNPR	41	270	143
Somal	26	136	70
Tigray	32	196	95
Totals	343	2043	1181

proportional to size to ensure adequate representation of smaller regions. Within each stratum a two-stage sampling design was planned. The primary units (PU) were Enumeration Areas (EA). A list of these, provided by the Central Statistical Authority, was used as the sampling frame. From this list, EAs were randomly selected with probability proportional to size. And the selected EAs were then visited to compile HH lists, and to identify logistical problems for fieldwork. Based on the listing, 11 HHs were selected by simple random sampling (so the inclusion probability of each household depends on the number in the EA and the sample size). Households deemed to be inaccessible were excluded and replaced by another drawn from the remaining HHs in the EA by simple random sampling.

Ethical approval for the ENMS was obtained from the National Research Ethical Review Committee of the Ethiopian Science and Technology Ministry (Reference 3.10/433/06). Informed consent was obtained from all adult participants and assent for all child participants in the survey.

Departures from sample practice

In the original sampling design, 11 HHs, 7 WRA (1 per HH, maximum), 3 men and 6 SAC were to be sampled from each EA. And all PSCs in the selected HHs were to be sampled. The sampling was based on the random ordering of the HHs from which the individuals were selected according to fixed rules. For example, the six SAC were to be selected from the 2nd, 4th, 6th, 8th, 9th and 10th HHs of the 11.

However, the original sampling design was not followed. Instead of sampling 1 WRA per HH, maximum, the teams collected data from all WRA encountered in a sampled HH to compensate for some HHs that did not have WRA. Also, the design was complicated further for PSC in that it was intended to sample all children in the 11 HHs, which gave rise to multiple samples within many of the basic sample units of the design, leading to some degree of within-household correlation.

We therefore have somewhat different designs for the different demographic groups, not a simple balanced two-stage sample with 11 sample units within the primary units, as was planned. For these reasons, the design inclusion probabilities do not apply to any of the data collected from individuals, as designed, which can have implications on the efficient estimation of micronutrient deficiencies. As we could not use sampling weights, one option is to use linear mixed modelling (LMM) to make estimates and predictions from the survey. Linear mixed models have an added advantage of making estimates of variance components.

Specification of the linear mixed model (LMM)

We specified region means as fixed effects, and EAs and HHs, random effects. This specification paralleled the sampling design where all regions were sampled, and EAs and HHs, randomly sampled.

In the linear mixed model (LMM) a vector of n observations, y , is treated as a combination of fixed and random effects. The fixed effects may be continuous covariates or factors, in this study the fixed effect was the mean value of the biomarker for each region. The LMM takes the following form,

$$\mathbf{y} = \mathbf{X}\boldsymbol{\tau} + \mathbf{Z}\mathbf{u} + \boldsymbol{\varepsilon},$$

where \mathbf{X} is an $n \times p$ design matrix which associates each of the n observations with a value for each of p fixed effects, and $\boldsymbol{\tau}$ contains the fixed effects coefficients. In this study we used a parameterization of the model in which $X[i, j] = 0$ unless the i^{th} observation corresponds to the j^{th} region in which case $X[i, j] = 1$.

There are r random effects, with values in \mathbf{u} , and \mathbf{z} is an $n \times r$ design matrix which associates each observation with a subset of these. In our case the random effects are the r_{EA} enumeration areas (EA) and the r_{HH} households (HH), with HH nested within EA. We may therefore think of \mathbf{u} as comprising two subvectors, the $r_{\text{EA}} \times 1$ vector of EA random effects, \mathbf{u}_{EA} , and the $r_{\text{HH}} \times 1$ vector of HH random effects, \mathbf{u}_{HH} , where $r = r_{\text{EA}} + r_{\text{HH}}$, and $\mathbf{u} = \begin{bmatrix} \mathbf{u}_{\text{EA}}^T, \mathbf{u}_{\text{HH}}^T \end{bmatrix}^T$. The design matrix \mathbf{Z} associates each observation with exactly one EA and one HH within that EA. The term $\boldsymbol{\varepsilon}$ is an independent and identically distributed residual.

It is assumed that the random effects and residual terms are independent, and normally distributed. The random effects at each level of a nested structure have a common variance, so in this case the model parameters include a separate variance component for EA within regions, σ_{EA}^2 ; one for HH within EA, $\sigma_{\text{HH:EA}}^2$, and a residual variance, $\sigma_{\boldsymbol{\varepsilon}}^2$. The joint distribution of \mathbf{u} and $\boldsymbol{\varepsilon}$ is therefore modelled as

$$\begin{bmatrix} \mathbf{u} \\ \boldsymbol{\varepsilon} \end{bmatrix} \sim N \left\{ \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \begin{bmatrix} \mathbf{G} & \mathbf{0} \\ \mathbf{0} & R_v \end{bmatrix} \right\}, \quad (1)$$

where the random effects have an $r \times r$ covariance matrix \mathbf{G} . In the case of the nested design-based sample used here \mathbf{G} can be written in terms of the variance components for the random effects and identity matrices \mathbf{I} :

$$\mathbf{G} = \begin{bmatrix} \sigma_{\text{EA}}^2 \mathbf{I}_{r_{\text{EA}}}, & \mathbf{0} \\ \mathbf{0}, & \sigma_{\text{HH:EA}}^2 \mathbf{I}_{r_{\text{HH}}} \end{bmatrix},$$

and the residual term has an $n \times n$ covariance matrix $\mathbf{R}_v = \sigma_{\boldsymbol{\varepsilon}}^2 \mathbf{I}_n$, which is diagonal because of the assumption that the residuals are independent. The unknown parameters, the variance components σ_{EA}^2 , $\sigma_{\text{HH:EA}}^2$ and $\sigma_{\boldsymbol{\varepsilon}}^2$ are estimated by residual maximum likelihood which avoids the well-known bias in ordinary maximum likelihood estimation. We did this estimation using the lmer function from the lme4 library for the R platform.²⁶

BLUE, BLUP and E-BLUP

Once the random effects parameters are estimated, the Mixed Model equation of Henderson et al.²⁷ can be applied to obtain the Best Linear Unbiased Estimates of the fixed effects coefficients (BLUE, $\hat{\boldsymbol{\tau}}$) and the Best Linear Unbiased Predictions of the random effects (BLUP, $\hat{\mathbf{u}}$). The equation is as follows:

$$\begin{bmatrix} \mathbf{X}^T \mathbf{R}_v^{-1} \mathbf{X} & \mathbf{X}^T \mathbf{R}_v^{-1} \mathbf{Z} \\ \mathbf{Z}^T \mathbf{R}_v^{-1} \mathbf{X} & \mathbf{Z}^T \mathbf{R}_v^{-1} \mathbf{Z} + \mathbf{G}^{-1} \end{bmatrix} \begin{bmatrix} \hat{\boldsymbol{\tau}} \\ \hat{\mathbf{u}} \end{bmatrix} = \begin{bmatrix} \mathbf{X}^T \mathbf{R}_v^{-1} \mathbf{y} \\ \mathbf{Z}^T \mathbf{R}_v^{-1} \mathbf{y} \end{bmatrix}. \quad (2)$$

The error covariance matrix of the estimates/predictions $\begin{bmatrix} \hat{\boldsymbol{\tau}}^T & \hat{\mathbf{u}}^T \end{bmatrix}$, \mathbf{C} , is estimated by

$$\hat{\mathbf{C}} = \left(\mathbf{W}^T \mathbf{R}_v^{-1} \mathbf{W} + \mathbf{G}^* \right)^{-1}, \quad (3)$$

where $\mathbf{W} \equiv [\mathbf{X}, \mathbf{Z}]$ and $\mathbf{G}^* \equiv \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{G}^{-1} \end{bmatrix}$.

The BLUP for some random quantity, described by a LMM, is the mean for the prediction distribution of that quantity, conditional on the model and observations. When the REML estimates of the random effects parameters are used to specify the model then the BLUP is sometimes called the empirical BLUP or EBLUP. Equation (2) can be solved to find the EBLUP of the individual random effects by using the estimated variance parameters to specify \mathbf{R}_v and \mathbf{G} .

EBLUP for the units within the sample

The EBLUP for the mean value for a sampled EA in a sampled region is the sum of the BLUE of the regional mean

and the BLUP of the random effect for that EA. The EBLUP for the mean of a sampled HH within an EA is, similarly, the sum of the BLUE for the regional mean and the BLUPs for the EA and HH random effects. The error variance of these predictions can be obtained from the variance and covariances of the model BLUEs and BLUPs given in equation (3). Such predictions would be useful, for example, if EA or HH from the original survey were under consideration for participation in further research where we wanted to select units where there is good reason to expect zinc deficiency.

EBLUP for unsampled units within the sample frame

Often we want to make an inference from sample data about unsampled units. Consider, for example, a case where a community within a Region of Ethiopia, corresponding to an EA, not included in the EMNS survey, is to be incorporated into a study and we want, in advance of field work to evaluate the probability that zinc is in deficiency there on the basis of the EMNS data.

If we want a prediction for an unsampled EA within a region, then, given the assumptions in our model that EAs within a region are independent, our best prediction is the BLUE for the regional mean. The error variance of this prediction is given by the sum of the error variance of the BLUE, and the between EA within region variance, σ_{EA}^2 .

Similarly, if we want a prediction for an household within an unsampled EA the BLUE of the regional mean is our best estimate and, this time the error variance of the prediction is given by the sum of the error variance of the BLUE, the between EA within region variance, σ_{EA}^2 , and the between HH within EA variance, $\sigma_{HH:EA}^2$.

If we want a prediction for an unsampled household within a sampled EA, then the EBLUP for the EA mean (i.e. the sum of the BLUE and the BLUP for the EA random effect) is our best prediction. The sum of the prediction error variance for this EBLUP and the between HH within EA variance, $\sigma_{HH:EA}^2$ gives us the variance of this prediction.

EBLUP for a new observation

The EBLUP for a new observation (i.e. an individual measurement), conditional on our data and the estimated model, can be written as

$$\begin{aligned} \tilde{y}_0 &= \mathbf{x}_0^T \hat{\boldsymbol{\tau}} + \mathbf{v}_{o,p}^T \mathbf{V}^{-1} (\mathbf{y} - \mathbf{X} \hat{\boldsymbol{\tau}}) \\ &= (\mathbf{x}_0^T - \mathbf{v}_{o,p}^T \mathbf{V}^{-1} \mathbf{X}) \hat{\boldsymbol{\tau}} + \mathbf{v}_{o,p}^T \mathbf{V}^{-1} \mathbf{y}, \end{aligned} \quad (4)$$

where \mathbf{x}_0 contains the entries for a design matrix for the observation (i.e. in this case indicators for the regional mean which is the fixed effect), $\mathbf{v}_{o,p}$ is a vector of covariances between the observations and the predicted value, and

$$\mathbf{v}_0 = \mathbf{ZGz}^T$$

where \mathbf{z} is a $1 \times r$ vector which indicates any corresponding random effects from the data, that is, sample units to which the individual belongs, if any) and

$$\mathbf{V} = \mathbf{ZGZ}^T + \mathbf{R}.$$

Information for further sampling

The LMM provides us with estimates of the variance components at each level of the sampling design. All other components of the estimation or prediction error variances for BLUEs and BLUPs depend only on the sample design (i.e. on the number of samples, and the number of EA and HH over which they are distributed within the regions). It is therefore possible to make estimates of the uncertainty of inferences based on different future sample designs, either within the same population or in analogous ones, given the estimated variance components. See, for example, Webster and Oliver.²⁸

To illustrate this we envisage a sampling task. This is of more limited scope than EMNS, focused within a single region. We are interested in forming a reliable estimate of the regional mean for the biomarker across a single region. We used a parametric bootstrap approach to the problem. For some specified sample design (a given number of EA in the region, and HH per EA, with a single individual sampled per household), we wrote a covariance matrix for a sample of observations based on the estimated variance components for EA, HH within EA, and observations within HH obtained from the EMNS analysis. A set of sample data were then simulated from this matrix using the `mvrnorm` function from the MASS library for the R platform.²⁹ The LMM was then fitted to the data and the BLUE of the regional mean was estimated. This procedure from the simulation step on was repeated 5000 times. The standard deviation of the BLUE values were treated as estimates of their standard error, and a 95% confidence interval was obtained for this. This was done for various designs with total sample sizes of 50, 100, 200 and 300, and with the number of EA varying between 5 and 150.

Results

Exploratory data analysis

We identified three large values of serum Zn concentration (outliers) which exceeded the upper ‘outer fence’ of Tukey,³⁰ that is, a threshold value equal to the third quartile of the data set plus three times the H-spread or inter-quartile range. Because our focus here is on the variability of serum Zn concentrations and its implications for prediction uncertainty and future sample design, we elected to remove these values before fitting the LMM. That is not to say that they should be discarded for other interpretations of the data set. A histogram of the data on the sample data is given in Figure A1 in the Appendix.

Table 3. Residual summary statistics.

Level	Mean	Median	Skewness	Octile skewness	Number of outliers
Population level	0.11	-0.28	0.41	0.01	0
EA level	0.01	-0.47	0.43	0.05	1
HH level	0.0	-0.40	0.44	0.06	0

An initial fit of the LMM to the data was undertaken and the residuals at population-, EA- and HH-levels were extracted. Their summary statistics are shown in Table 3, and histograms in Figure A2 in the Appendix. On the basis of these exploratory plots and statistics, the assumption of normally-distributed random effects in the model was deemed to be plausible.

EA sample weights and EA random effect BLUPs

We have noted that the EMNS survey was originally designed with inclusion probabilities specified at EA level (proportional to size at EA level, and equal across the EA at HH level). However, for WRA the selection of individuals for sampling was not based on the original sample plan, and so the inclusion probabilities, from which sample weights would be calculated, do not hold. Our analysis is therefore not based on the sample weights. However, we should consider the risk that the inclusion probabilities at EA level are *informative*. That is to say, EAs with larger inclusion probabilities might be more likely to have Zn deficiency (or sufficiency) than average. There are formal methods to check for informativeness of sample weights, for example by Pfeffermann.³¹ However, in this study we do not know the actual sample weights for individual observations (this the motivation for our use of LMM) so these tests cannot be applied. We therefore used an informal method to evaluate any evidence that the EA inclusion probabilities might be informative. We extracted the BLUPs for the EA random effects, and plotted them in order of EA inclusion probability. This is shown in Figure 1. Where the BLUPs are shown as points distributed about their mean (zero). The red line shows the moving average of the random effect BLUPs within a window. There is no evidence from this plot of any trend in the Zn variation at EA level with EA inclusion probability, and so we proceed on the assumption that the original design weights are uninformative.

BLUEs of region means as weighted means

If one examines the mixed model equation, equation (2), then it is apparent that the BLUEs for the regional means, and the BLUPs for the random effects for the sample units, are linear combinations of the data:

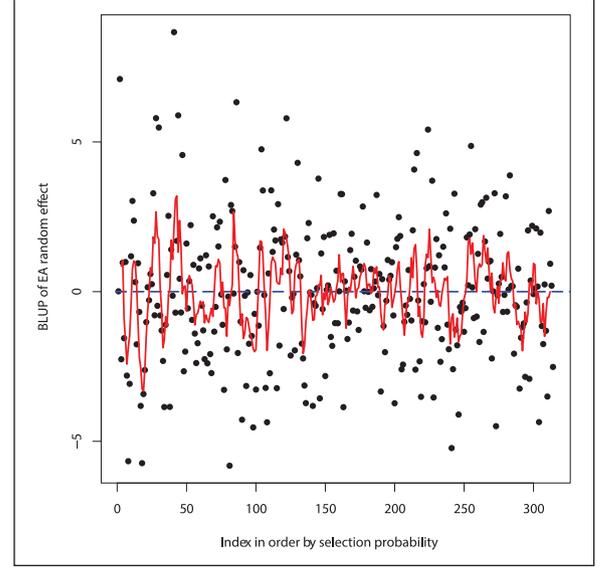


Figure 1. BLUP for the EA random effect of each EA in EMNS (serum Zn concentration for WRA) plotted in order of increasing selection probability. The red line shows the moving average of the BLUP in five successive ordered EA.

$$\begin{bmatrix} \hat{\tau} \\ \hat{\mathbf{u}} \end{bmatrix} = \begin{bmatrix} \mathbf{X}^T \mathbf{R}_v^{-1} \mathbf{X} & \mathbf{X}^T \mathbf{R}_v^{-1} \mathbf{Z} \\ \mathbf{Z}^T \mathbf{R}_v^{-1} \mathbf{X} & \mathbf{Z}^T \mathbf{R}_v^{-1} \mathbf{Z} + \mathbf{G}^{-1} \end{bmatrix}^{-1} \times \begin{bmatrix} \mathbf{X}^T \mathbf{R}_v^{-1} \\ \mathbf{Z}^T \mathbf{R}_v^{-1} \end{bmatrix} \mathbf{y}. \quad (5)$$

In this sense, the BLUEs for the regional means are weighted means of the observations, as would be estimates based on survey weights, but with weights which reflect the modelled degree of within HH and within EA correlation. To give an intuitive insight into this we extract in Table 4 the effective weights for the estimation of the mean serum Zn concentration for Amhara region for WRA in different households in 5 EA from the region. From these we can see the following.

1. *Individual weights are smaller within EA with more observations.* To see this compare the EAs ENMS037, ENMS036 and ENMS312 with one, three and four observations respectively (each from a single household), and corresponding weights for each observation of 0.009, 0.007 and 0.006. This reflects the fact

Table 4. Some example LMM weights from enumeration areas (EAs) and households (HHs) in Amhara region.

Region	EA	HH	Weight
Amhara	ENMS036	C036-04	0.007070486
Amhara	ENMS036	C036-07	0.007070486
Amhara	ENMS036	C036-02	0.007070486
Amhara	ENMS312	C312-05	0.006372480
Amhara	ENMS312	C312-06	0.006372480
Amhara	ENMS312	C312-10	0.006372480
Amhara	ENMS312	C312-04	0.006372480
Amhara	ENMS039	C039-04	0.005832112
Amhara	ENMS039	C039-02	0.005832112
Amhara	ENMS039	C039-04	0.005832112
Amhara	ENMS039	C039-02	0.005832112
Amhara	ENMS037	C037-02	0.009053915
Amhara	ENMS040	C040-09	0.004807782
Amhara	ENMS040	C040-07	0.004807782
Amhara	ENMS040	C040-04	0.004807782
Amhara	ENMS040	C040-02	0.004807782
Amhara	ENMS040	C040-01	0.004807782
Amhara	ENMS040	C040-05	0.003680640
Amhara	ENMS040	C040-05	0.003680640
Amhara	ENMS040	C040-05	0.003680640

that repeated observations within an EA are to some extent correlated.

- Individual weights within an EA from different households are larger than for multiple observations within households. Compare EAs ENMS312 and ENMS039; both include four observations, but in the former these are each from a different household, whereas in the latter the four observations comprise two from each of two households. The weights in the latter case are smaller, reflecting the correlation between observations within the same household.
- Within an EA, single observations within households have larger weights than repeat observations within one. Consider EA ENMS040, with three observations within a single household, each with a smaller weight than the remaining five observations in that EA from five different households.

BLUEs and EBLUPs

Regional means: BLUEs and their confidence intervals. The BLUEs for regional mean serum Zn concentration for women of reproductive age (WRA) in the Ethiopian regions ranged from $57\mu\text{g dL}^{-1}$ (Oromia) to $69.2\mu\text{g dL}^{-1}$ (Addis Ababa) (Table 5). Three regions had point estimates below $60\mu\text{g dL}^{-1}$, and the rest, (8), were above $60\mu\text{g dL}^{-1}$ (Table 5). However, Addis Ababa was the only region with a mean serum Zn concentration that was above the threshold cut-off for sufficiency of $65\mu\text{g dL}^{-1}$

Table 5. Estimates of regional means, and standard errors (SE) of the estimates.

Regions	Estimated mean	SE of estimated mean
Addis Ababa	69.18	1.18
Afar	60.57	1.46
Amhara	62.35	1.01
Benishangul-Gumuz	59.90	1.38
Dire Dawa	60.11	1.47
Gambella	62.29	1.46
Harari	61.08	1.60
Oromia	57.78	0.96
SNNPR	57.99	1.07
Somali	61.13	1.50
Tigray	60.03	1.35

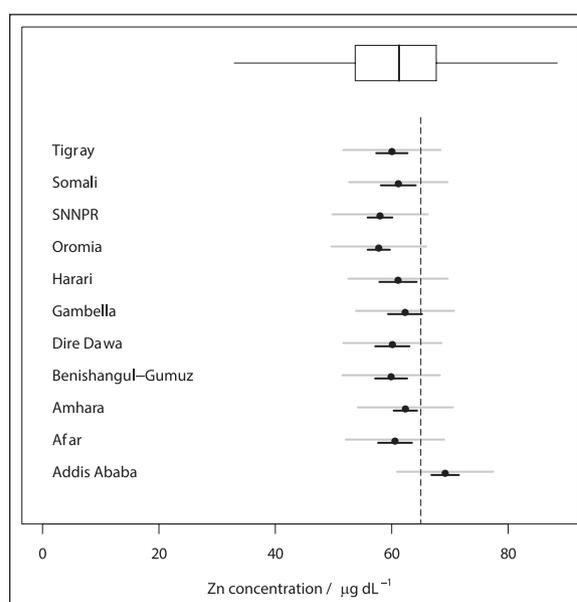


Figure 2. The regional mean concentration of Zn in blood serum of WRA, estimated from EMNS data by weighted least squares after REML estimation of variance parameters of a LMM. The solid black line shows the confidence interval of the estimate. The solid grey line shows the prediction interval for the mean, treated as the EBLUP of the mean of an unsampled EA within each region. The boxplot at the top shows the distribution of the sample data. The vertical dashed line is a threshold concentration for Zn sufficiency.

demarcated by the vertical dotted line in Figure 2. The means are plotted in Figure 2 (black dots), along with 95% confidence intervals (solid horizontal black lines straddling the means) which are based on the standard errors in Table 5. The standard errors vary between 0.96 and 1.5, corresponding to error variances of 0.92 and 2.25 respectively. These estimates could be used directly to plan Zn intervention for regions in Ethiopia. For example, Oromia, the region with the smallest estimated mean ($57.7\mu\text{g dL}^{-1}$) could be prioritized for Zn intervention, whereas Addis

Table 6. EBLUPs of individuals in sampled and unsampled enumeration areas (EAs) and households (HHs), and sampled and unsampled (EAs) and HHs, all within Oromia Region.

Scenario	Enumeration area (EA)	Household (HH)	EBLUP type	EBLUP	Variance	Probability < 65µg dL ⁻¹	Interpretation
1.	Sampled	–	EA	56.6	8.9	>0.99	Virtually certain
2.	Unsampled	–	EA	57.8	16.6	0.96	Very likely
3.	Sampled	Sampled	HH	56.5	18.8	0.98	Very likely
4.	Unsampled	Unsampled	HH	57.8	31.3	0.90	Very likely
5.	Sampled	Sampled	Individual	56.5	100.2	0.80	Likely
6.	Sampled	Unsampled	Individual	56.6	104.9	0.79	Likely
7.	Unsampled	Unsampled	Individual	57.8	112.6	0.75	Likely

Ababa, with the largest estimated mean, may require no interventions, unless these are targeted to particular vulnerable subgroups. The regional mean, however, masks substantial variation at the EA, HH and individual levels, so while Oromia may be the priority region, there could be substantial regions of severe deficiency elsewhere, including in Addis Ababa.

EBLUPs and Prediction intervals (PIs). Table 6 presents EBLUPs of serum Zn for some example cases in Oromia Region. These are at different levels of organization (individual, EA mean or HH mean) and based on different sampling situations (sampled or unsampled HH and EA). In each case the sampled EA and sampled HH considered are the same. The BLUPs and their prediction error variances are as described in the Methods section. Note that all these BLUP error variances are at least an order of magnitude larger than the estimation error variance of the BLUE. The cases in rows 2, 4 and 7 are all for cases where the prediction is outside the sampling frame at all levels (in an unsampled EA). In all these cases the BLUP is equal to the BLUE for the regional mean. The variance is largest for BLUPs at the individual level.

For EBLUPs of EA means (rows 1 and 2) or HH means (rows 3 and 4), the variances are larger in the latter. The prediction error variance for the mean of a sampled EA is smaller than that of an unsampled EA, but is not negligible because of the contribution of variances at within-EA levels. The same is seen for HH mean predictions.

Each row of Table 6 represents a possible situation in which a prediction is required to support a decision. For example, consider a case where we were planning a further study in the Oromia region which was to be held in a sampled EA, and we wanted a case where we were confident that the mean serum Zn concentration in WRA is below the 65µg dL⁻¹ threshold. For the sampled EA in row 1, the EBLUP of the mean is well below this threshold. Assuming normal variation (justified by our exploratory analysis of the LMM), the variance can be used to compute the probability that the EA mean is < 65µg dL⁻¹. Table 6 shows that this probability is > 0.99. Table 6 also

presents the interpretation of this probability in terms of the Intergovernmental Panel on Climate Change's calibrated phrases³²: < 0.01 (*exceptionally unlikely*), < 0.33 (*unlikely*), ≥ 0.66 (*likely*), ≥ 0.9 (*very likely*), ≥ 0.99 (*virtually certain*). Note that it is *virtually certain* that the mean for this sampled EA is below the threshold.

Now if we wished to expand the EAs in the second study, including one which had not been sampled, then the prediction 'shrinks' to the regional mean BLUE (on Table 5). The effect is small, but the prediction error variance nearly doubles. The probability that the unsampled EA mean is less than the threshold is smaller than for the sampled case, but is still large at 0.96, which is interpreted as *very likely*.

Row 3 of the Table shows that, for a particular sampled household, the uncertainty is close to that for an unsampled EA mean. Because the sampled HH EBLUP is smaller than the regional mean BLUE (in Table 5), the probability that the HH mean is below the threshold is slightly larger than for the unsampled EA mean (0.98). Again, for an unsampled HH in an unsampled EA (row 4) the EBLUP shrinks to the regional mean BLUE, and the prediction error variance is notably larger than for the sampled EA, but the probability that the HH mean is below the threshold is 0.90, still interpreted as *very likely*.

Rows 5–7 are all predictions for individuals. The prediction error variances are all an order of magnitude larger than those for HH or EA mean BLUPS, and the probability that the individual Zn serum concentration is < 65µg dL⁻¹ ranges from 0.75 to 0.80. Note that in all cases in Table 6, the EBLUP is below the threshold, the probability that the true value is below the threshold decreases from EA to HH to individual level because the uncertainty of the prediction also increases, and so the probability that the true value is actually above the threshold.

Table 7 presents the probability that the mean of an unsampled EA is < 65µg dL⁻¹ for each sampled region. In each case the BLUP shrinks to the BLUE of the regional mean, as presented in Table 5. The prediction intervals are shown in Figure 2. They are the solid horizontal grey lines straddling the means, which are clearly larger than the

Table 7. Probability enumeration area (EA) mean is < 65µg dL⁻¹.

Region	Probability EA is < 65µg dL ⁻¹	Interpretation
Addis Ababa	0.16	Unlikely
Afar	0.85	Likely
Amhara	0.74	Likely
Benishangul-Gumuz	0.89	Likely
Dire Dawa	0.88	Likely
Gambella	0.74	Likely
Harari	0.82	Likely
Oromia	0.96	Very likely
SNNPR	0.96	Very likely
Somali	0.82	Likely
Tigray	0.88	Likely

confidence intervals (the solid horizontal black lines also straddling the means). H_ireturn;

Lessons for future sampling. The bootstrapped standard error for the BLUE of the regional mean is reduced below 1.0 by a sample of 300 individuals from 25 EAs – so 12 HH per EA, (Figure 3). If the same total sample size were distributed over just 10 EAs (30 HH per EA), then the uncertainty of the BLUE would be larger. Note that if we sample 50 EA with 100 samples the standard error exceeds 1, it is brought below 1 if the EAs are sampled with 200 individuals. The reduction in uncertainty from increasing the sample size over 50 EA from 200 to 300 is small, despite the 50% increase in the number of biomarker analyses, and the intrusive sampling of individuals. Such results would clearly help the rational choice of a sample design.

If the marginal costs of an additional sample EA, C_{EA}, in the survey, and an additional HH within a sampled EA, C_{HH}, or at least their ratio R_{EA,HH} = C_{EA} / C_{HH} can be approximated, then the optimum number of HH to sample per EA, \bar{m} , can be estimated on the assumption of a uniform total number of HH in each EA, \bar{M} (REF):

$$\bar{m} = \sqrt{R_{EA,HH} \frac{\sigma_{HH:EA}^2 + \sigma_{\epsilon}^2}{\sigma_{EA}^2 - \frac{\sigma_{HH:EA}^2 + \sigma_{\epsilon}^2}{\bar{M}}}} \quad (6)$$

If we assume that in our target region $\bar{M} = 200$ (the mean over the MNS is 193), and that R_{EA,HH} = 30, then the optimum number of HH per EA, given the variance components estimated for serum zinc in the EMNS, is 14

Conclusions

Design-based surveys, such as those used in national-scale micronutrient surveys (MNS) may depart in practice from the sample design due to one or more problems encountered

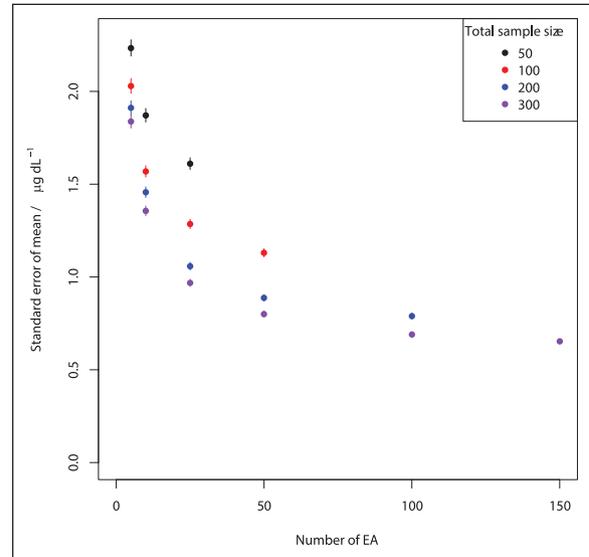


Figure 3. Parametric bootstrapped standard error for the BLUE of a regional mean based on sampling within the region based on different total sample sizes, and numbers of EA within the sample. The vertical lines show the 95% confidence interval of the bootstrap estimates (which in many cases is narrower than the plotting symbol).

during the sampling exercise. In such cases, the specified inclusion probabilities, from which sample weights are calculated no longer hold. If the original sample weights are used, then we can no longer be sure that population parameters are estimated without bias. Alternatively, linear mixed models (LMM), with covariance structures specified to reflect the nested structure of the sample and the estimated variance parameters at each level, can be used to make model-based parameter estimates.

In this paper, we demonstrated the use of LMM to make estimates (region means), and mean predictions (EA, HH and individual means) of Zn in blood serum of WRA using the Ethiopia National Micronutrient Survey (EMNS) in which the actual sampling practice departed from the design for which sample weights were available. We first checked for evidence that the sample weights were informative for the target variable. There was no evidence that this was the case. The BLUEs for regional mean serum Zn concentration for women of reproductive age (WRA) in the Ethiopian regions ranged from 57µg dL⁻¹ (Oromia) to 69.2µg dL⁻¹ (Addis Ababa). The Addis Ababa estimate (69.2µg dL⁻¹) was the only region with a mean serum Zn concentration that was above the threshold cut-off for sufficiency of 65µg dL⁻¹. These estimates, although generally slightly higher (by 0.02 to 4.75), were generally close to the design-based estimates on Table 2 of Belay et al.¹² However, except for Addis Ababa and Oromia regions, the ranking of the regional means were changed. The largest changes in rank were Dire Dawa (dropped 5 positions), and Harari and Amhara regions (both gained five positions). The second largest

changes were Benishangul-Gumuz, which lost three positions, and SNNPR which lost two positions. Two of the remaining four regions lost a position each while the other two gained a position each. The differences in ranking could be due to differences in weighting, with the LMM weighting according to actual number of individuals within units (e.g. HH), correlated individuals. The standard errors for the BLUES varied between 0.96 and 1.5. The BLUP error variances were at least an order of magnitude larger than the estimation error variance of the BLUES. This is because the BLUPs were for smaller units within the regions (EA, HH or individuals). The variance of the BLUP is also larger for an unsampled unit than for one which contains some of the sample data.

The regional mean serum Zn concentration was *likely* to *very likely* to be smaller than the threshold concentration for deficiency in all regions, except Addis Ababa. These calibrated expressions are used to communicate the underlying probability of the specified state, computed from the fitted LMM.

The variance components from the ENMS could be used to explore options for future micronutrient surveys (MNS) in Ethiopia. Some general conclusions could be drawn. For example, if a total sample size of 200 or more is used, then there are limited marginal gains in precision of estimates of the regional mean from sampling more than 50 EA in total. It is also possible to draw more specific conclusions about optimal sampling schemes if the marginal costs of an additional sample EA in the survey, and an additional HH within a sampled EA, or at least their ratio, could be approximated.

To conclude, these sampling issues also apply to estimation of other micronutrient deficiencies. Thus, many micronutrient surveys, including those capturing multiple MNs, and other surveys involving MSS would benefit from LMM when sampling weights are no longer valid due to changes in sampling practice.

Acknowledgments

The authors acknowledged the Ethiopian Public Health Research Institute for sharing archived national micronutrient survey serum samples for mineral analysis.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported, in whole or in part, by the Bill & Melinda Gates Foundation [INV-002855]. Under the grant conditions of the Foundation, a "Creative Commons Attribution 4.0 Generic License has already been assigned to the Author Accepted Manuscript version that might arise from this submission."

ORCID iD

Hakunawadi Alexander Pswarayi  <https://orcid.org/0009-0000-8009-5852>

Supplemental material

Supplemental material for this article is available online.

References

- Gashu D and Stoecker BJ. Selenium and cognition: mechanism and evidence. In: Preedy V and Patel VB (eds) *Handbook of famine, starvation, and nutrient deprivation: from biology to policy*. Cham: Springer International Publishing, 2018, pp.1–17.
- Shenkin A. The key role of micronutrients. *Clin Nutr* 2006; 25(1): 1–13.
- Allen LH. Multiple micronutrients in pregnancy and lactation: an overview. *Am J Clin Nutr* 2005; 81(5): 1206S–1212S.
- World Health Organization. *Worldwide prevalence of anaemia 1993-2005: WHO global database on anaemia*. World Health Organization, 2008. Available at: <https://apps.who.int/iris/handle/10665/43894> (accessed 6 September 2024).
- King JC, Brown KH, Gibson RS, et al. Biomarkers of Nutrition for Development (BOND)—Zinc Review 12345. *J Nutr* 2016; 146(4): 858S–885S.
- Lowe NM, Fekete K and Decsi T. Methods of assessment of zinc status in humans: a systematic review. *Am J Clin Nutr* 2009; 89(6): 2040S–2051S.
- Terhune MW and Sandstead HH. Decreased RNA polymerase activity in mammalian zinc deficiency. *Science* 1972; 177(4043): 68–69.
- Wessells KR and Brown KH. Estimating the global prevalence of zinc deficiency: results based on zinc availability in national food supplies and the prevalence of stunting. *PLoS One* 2012; 7(11): e50568.
- Phiri FP, Ander EL, Bailey EH, et al. The risk of selenium deficiency in Malawi is large and varies over multiple spatial scales. *Sci Rep* 2019; 9(1): 6566.
- National Statistical Office. *Malawi micronutrient survey key indicators report 2015–16*. Atlanta, GA: NSO, CHSU, CDC and Emory University, 2017.
- Hailu S, Wubshet M, Woldie H, et al. Iodine deficiency and associated factors among school children: a cross-sectional study in Ethiopia. *Arch Public Health* 2016; 74: 1–7.
- Belay A, Gashu D, Joy EJ, et al. Zinc deficiency is highly prevalent and spatially dependent over short distances in Ethiopia. *Sci Rep* 2021; 11(1): 1–13.
- De Gruijter J, Brus DJ, Bierkens MF, et al. *Sampling for natural resource monitoring*. vol. 665. Berlin, Heidelberg: Springer, 2006.
- Sedgwick P. Multistage sampling. *BMJ* 2015; 351: h4155.
- Oh HL. Weighting adjustment for unit nonresponse. In: Nisselson H, Olkin I and Madow WG (eds) *Incomplete data in sample survey*. vol. 2. London: Academic Press, 1983, pp.143–184.
- Henry K and Valliant R. Comparing alternative weight adjustment methods. In: *Proceedings of the Joint Statistical Meetings, section on survey research methods*, 2012, pp.4696–4710. Alexandria, VA: American Statistical Association. Available at: http://www.asasrms.org?Proceedings/y2012/Files/306157_76012.pdf

17. Verbeke G and Molenberghs G. *Linear mixed models for longitudinal data*. New York, NY: Springer, 2000.
18. R Core Team. *R: a language and environment for statistical computing*. Vienna, Austria, 2023. Available at: <https://www.R-project.org/> (accessed 6 september 2024).
19. Pinheiro JC and Bates DM. *Mixed-effects models in S and SPLUS*. New York, NY: Springer, 2000.
20. Blaker H. Confidence curves and improved exact confidence intervals for discrete distributions. *Can J Stat* 2000; 28: 783–798.
21. Scherer R. PropCIs: Various confidence interval methods for proportions. R package version 0.30, 2018. Available at: <https://CRAN.R-project.org/package=PropCIs>
22. Patterson HD and Thompson R. Recovery of inter-block information when block sizes are unequal. *Biometrika* 1971; 58(3): 545–554.
23. Skinner C and Wakefield J. Introduction to the design and analysis of complex survey data. *Stat Sci* 2017; 32(2): 165–175.
24. Yan T and Curtin R. The relation between unit nonresponse and item nonresponse: a response continuum perspective. *Int J Public Opin Res* 2010; 22(4): 535–551.
25. Namaste SM, Rohner F, Huang J, et al. Adjusting ferritin concentrations for inflammation: Biomarkers Reflecting Inflammation and Nutritional Determinants of Anemia (BRINDA) project. *Am J Clin Nutr* 2017; 106(suppl 1): 359S–371S.
26. Bates D, Mächler M, Bolker B, et al. Fitting linear mixed-effects models using lme4. *J Stat Softw* 2015; 67(1): 1–48.
27. Henderson CR, Kempthorne O, Searle SR, et al. The estimation of environmental and genetic trends from records subject to culling. *Biometrics* 1959; 15(2): 192–218.
28. Webster R and Oliver MA. *Statistical methods in soil and land resource survey*. Oxford: Oxford University Press, 1990.
29. Venables W, Venables WN and Ripley BD. *S programming*. New York, NY: Springer Science & Business Media, 2000.
30. Tukey JW. *Exploratory data analysis*. vol. 2. Reading, MA: Addison-Wesley Publishing Company, 1977.
31. Pfeffermann D. The role of sampling weights when modeling survey data. *Int Stat Rev* 1993; 61(2): 317.
32. Mastrandrea MD, Mach KJ, Plattner GK, et al. The IPCC AR5 guidance note on consistent treatment of uncertainties: a common approach across the working groups. *Clim Change* 2011; 108(4): 675–691.
33. Korn EL and Graubard BI. Examples of differing weighted and unweighted estimates from a sample survey. *Am Stat* 1995; 49(3): 291–295.
34. United Nations. Construction and use of sample weights. In: United Nations (ed.) *Designing household survey samples*. New York, NY: UN; 2008, pp.109–127.
35. Horvitz DG and Thompson DJ. A Generalization of sampling without replacement from a finite universe. *J Am Stat Assoc* 1952; 47(260): 663–685.

Appendix I

Sampling weights in design-based sampling and analysis. Mathematically, sampling weights are the reciprocal of the inclusion probabilities for the basic sample units.³³ For example, in a survey where we sample one individual per HH, if the inclusion probability for the i th HH is (π_i) , then the sample weight for a measurement from that individual is $(\pi_i)^{-1}$. This weight (w_i), is the number of population units represented by the i th sampled unit.²³ The sampling weight of a lower stage that is nested in an upper stage is a product of the weight of that lower stage with the weight of the upper stage in which it is nested.³⁴ For example, if $(\pi_{EA,i})$ is the inclusion probability over all EAs of the i th EA and $(\pi_{HH,i})$ is the inclusion probability of the j th HH in the i th EA, then the overall inclusion probability for the sampled individual is $(\pi_{EA,i})(\pi_{HH,ij})$.

To calculate design-unbiased population estimates, for example, $E(\bar{\mu}) = \mu$, using sampling weights, the inclusion probability (π_i) multiplied by the sampling weight (w_i) should equal one (1): that is, $\pi_i \times w_i = 1$.³⁵ This holds when the sampling weight (w_i) is the inverse of the

inclusion probability (π_i^{-1}) : $\pi_i \times \pi_i^{-1} = 1$.³⁵ Hence, if an observation in a design based sample has inclusion probability π_i , and, therefore, sample weight π_i^{-1} , then the unbiased design-based estimate for the mean of variable z for observations in region j would be

$$\frac{\sum_{i=1}^n \pi_i^{-1} I_{i,j} z_i}{\sum_{i=1}^N \pi_i^{-1} I_{i,j}}, \quad (7)$$

where $I_{i,j}$ is an indicator variable which takes value 1 if observation i belongs to region j and 0 otherwise. In effect

$$\frac{\pi_i^{-1} I_{i,j}}{\sum_{i=1}^N \pi_i^{-1} I_{i,j}},$$

is a weight applied to observation z_i to obtain its contribution to the design-based estimate of the regional mean. As mentioned above, these weights sum to 1 over all the observations in the region.

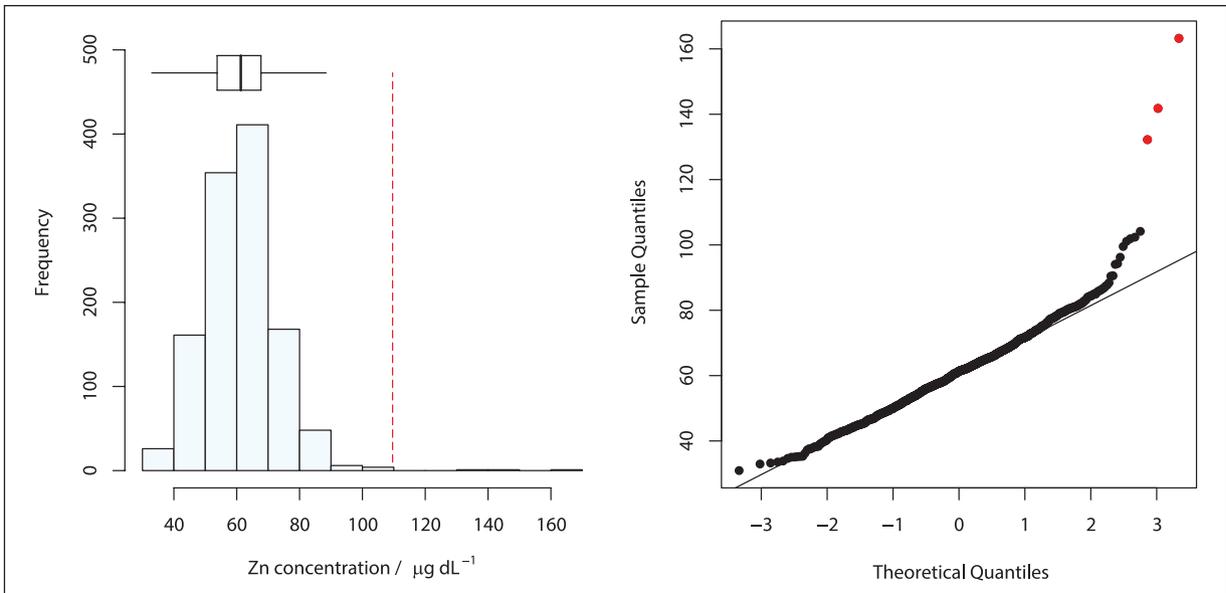


Figure A1. Histogram with boxplot (left) and QQ-plot (right) for sample data. The vertical red broken line on the histogram shows the upper outer fence for the observations,³⁰ and probable outliers are shown by red symbols on the QQ plot.

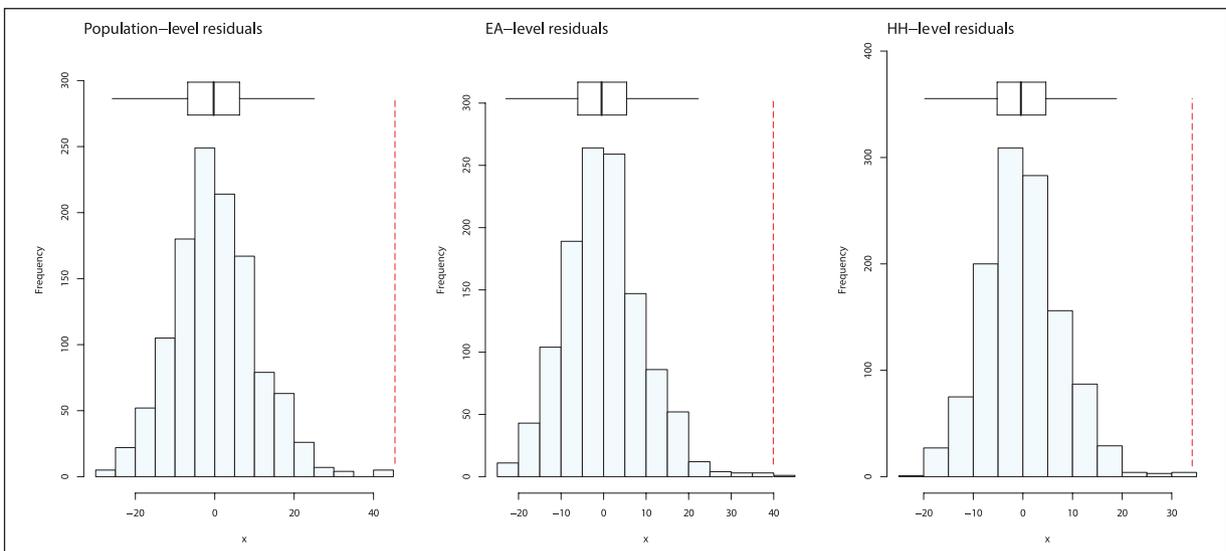


Figure A2. Histograms with boxplots for random effects from exploratory fit of model.