

Design of the Millennium Villages Project Sampling Plan: a simulation study for a multi-module survey

Shira Mitchell, Rebecca Ross, Susanna Makela, Elizabeth A. Stuart, Avi Feller, Alan M. Zaslavsky, Andrew Gelman

Abstract

The Millennium Villages Project (MVP) is a ten-year integrated rural development project implemented in ten sub-Saharan African sites. At its conclusion we will conduct an evaluation of its causal effect on a variety of development outcomes, measured via household surveys in treatment and comparison areas. Outcomes are measured by six survey modules, with sample sizes for each demographic group determined by budget, logistics, and the group's vulnerability. We design a sampling plan that aims to reduce effort for survey enumerators and maximize precision for all outcomes. We propose two-stage sampling designs, sampling households at the first stage, followed by a second stage sample that differs across demographic groups. Two-stage designs are usually constructed by simple random sampling (SRS) of households and proportional within-household sampling, or probability proportional to size sampling (PPS) of households with fixed sampling within each. No measure of household size is proportional for all demographic groups, putting PPS schemes at a disadvantage. The SRS schemes have the disadvantage that multiple individuals sampled per household decreases efficiency due to intra-household correlation. We conduct a simulation study (using both design- and model-based survey inference) to understand these tradeoffs and recommend a sampling plan for the Millennium Villages Project. Similar design issues arise in other studies with surveys that target different demographic groups.

1 Background

The Millennium Villages Project (MVP) is an economic development project that targets rural populations across ten countries in sub-Saharan Africa, implementing a multi-sector package of interventions at a village level (Sachs and McArthur, 2005; Sanchez et al., 2007). See Mitchell et al. (2015a) for background on the project, study site selection, outcomes of interest, and a comprehensive description of the plan to evaluate its effectiveness. Mitchell et al. (2015b) describe our plan for causal inference about the MVP's effect on a variety of development outcomes measured in different demographic groups. These outcomes will be measured via survey modules administered in both treatment and comparison villages.

A design analysis described in Mitchell et al. (2015b) was used to recommend the number of control villages and magnitude of sampling in each. Next, we must determine how to select households and individuals within households. We propose a two-stage sample: households will be sampled in stage I, followed by individuals within households in stage II (Lohr, 2010; Särndal et al., 1992). In the first stage, we must decide between simple random sampling and probability

proportional to size sampling of households. Because the project operates at the village level, a sampling plan that efficiently estimates outcome means per village is an efficient sampling plan for the overall causal evaluation. In this paper we conduct a simulation study to decide on a sampling plan for estimating finite population village means.

We aim to minimize the design effect, the ratio between the actual and effective sample sizes. One factor in determining the efficiency of a sampling design is the intraclass correlation, i.e. the correlation among individuals within a household. If more than one individual is sampled per household, the intraclass correlation increases the design effect, reducing the effective sample size relative to the actual sample size.

Another factor in the efficiency of a sampling design is the distribution of individuals' sampling probabilities. Sampling probabilities can be optimized for a specific outcome, e.g. by sampling with probability approximately proportional to the outcome (Särndal et al., 1992, p.88). However, with many outcomes of interest, such tailored optimization is difficult or impossible. Therefore, a *self-weighted* sample design is preferred, such that all individuals are sampled with equal probability (Kish, 1992; Lohr, 2010, p.287). Such samples are representative without weighting adjustments, and unbiased point estimates can be obtained from standard statistical procedures.

Given a fixed precision, we aim to minimize time and resources for the survey enumerator teams. This includes minimizing the numbers of people surveyed (i.e. the actual sample size), but also considering the number of households visited, and the effort required to prepare a sampling frame. To conduct the first stage of sampling, a scheme that samples households with equal probability only requires a list of all households with GPS coordinates identifying their locations. However, a scheme which samples households with probability proportional to size requires some measure of household size (e.g. the total number of household members). This additional piece of information requires more effort for enumerators, especially for larger villages with many households. After either method of first stage sampling, we will conduct a demographic census in the sampled households to create the sampling frame for the second stage.

In this paper we conduct a simulation study to understand the tradeoffs between simple random sampling and probability proportional to size sampling of households in the context of the MVP evaluation. Additionally, our simulations explore design-based versus model-based inference, a dichotomy which has implications for our the analysis of our outcome data.

2 Outcomes and survey modules

The Millennium Villages Project (MVP) defines 51 outcomes of interest, including measures of poverty alleviation, agriculture, education, gender equality, health, environmental sustainability, and infrastructure (Mitchell et al., 2015a). These outcomes are measured in six different survey modules, whose content is discussed in Section ?? of Mitchell et al. (2015a). These modules include:

- a **household survey**, administered to all household heads (or other knowledgeable household

members) within the sampled households;

- a sex-specific **adult survey**, administered to men and women of reproductive age (15 to 49 years) within the sampled households;
- within the adult-female survey, a **birth history** section, administered to women of reproductive age (15 to 49 years) both in the sampled households and in additional sampled households to reach sample size sufficient for estimating child mortality;
- a **nutrition survey**, administered to men and women age 15 to 49 years in sampled households;
- **blood (malaria and anemia)** testing, administered to four age-sex groups in sampled households: children age 6 to 59 months, school-aged children (5 to 14 years old), men age 15 to 49 years, women age 15 to 49 years; and
- **anthropometry** measurements, administered among children age 6 to 59 months in sampled households.

For each module and age-sex group combination, the project has budgeted a target sample size based on a combination of budget, logistics, and relative importance of different vulnerable populations and intervention beneficiaries.

3 Sampling plans considered

For the purpose of our simulation study, we consider all survey modules except for birth history and the nutrition survey. Our sampling will be performed in two phases. First, we will sample households using either simple random sampling (SRS, without replacement) or probability proportional to size sampling (PPS, with replacement), with household size defined as $N_{h,\text{total}}$, the number of household members under 50 years old in household h . Let s_I be the set of (unique) sampled households. In the PPS scheme, we use r_I to denote the set of sampled households with repeats. Let $n_I = |s_I|$ be the number of households sampled without replacement in the SRS scheme and let $m_I = |r_I|$ be the number of households sampled with replacement in the PPS scheme. We let $n_I = m_I = 300$ based on the project's previous survey rounds and budget for the final survey round.

To describe the within-household sampling plans for each survey module, we use the following notation. Let N_h be the total number of people in household h that are in the target age-sex group for a particular module. Let $n_h \leq N_h$ be the number of people in household h that we sample and survey. For example, if considering the anthropometry module, then N_h is the number of children under five years of age in household h and n_h is the number of those sampled for the anthropometry module. Let $N = \sum_h N_h$ be the total number of people in the sampling frame (an MV1 or a control village) that are in the module's target age-sex group.

We now outline the within-household sampling plans considered in our simulation study.

Adult, anthropometry, and blood modules

For each module and age-sex group combination, the project has budgeted a target sample size, n_{target} :

- **adult survey** - $n_{\text{target}} = 400$ men and $n_{\text{target}} = 400$ women of reproductive age (15 to 49 years);
- **blood (malaria and anemia)** - $n_{\text{target}} = 300$ children age 6 to 59 months, $n_{\text{target}} = 100$ school-aged children (5 to 14 years old), $n_{\text{target}} = 100$ men age 15 to 49 years, $n_{\text{target}} = 100$ women age 15 to 49 years; and
- **anthropometry** - $n_{\text{target}} = 300$ children age 6 to 59 months.

The second stage sampling schemes we consider in this simulation study are, for a given module and age-sex group:

- For **SRS** sampling of households - combine all $N_{s_1} = \sum_{h \in s_1} N_h$ people in the sampled households in the target age-sex group. If $N_{s_1} \leq n_{\text{target}}$, then survey all. Otherwise, we consider two options:

- **stratify** by sampling n_h individuals from each household, where n_h is proportional (up to rounding) to N_h , and the constant of proportionality is determined by the total in the sampled households,

$$n_h = \text{round} \left(N_h * \frac{n_{\text{target}}}{N_{s_1}} \right); \text{ or}$$

- take an equal-probability **systematic sample** of n_{target} people. We order the households randomly, and people (in the module’s target age-sex group) within households randomly, so that the people within a household are listed consecutively. We then take a sample using the fractional interval method described in Särndal et al. (1992, p.77) and Appendix A. This procedure enables us to control sample sizes and spread the sample across households such that the sample size in a household is always either the ceiling or the floor of the expected sample size in that household under simple random sampling (see Appendix A). Conceptually, this is similar to stratifying on household, except that there is dependence of the samples between strata (i.e. households).
- For **PPS** sampling of households - if $n_{\text{target}} \geq m_1$, sample a fixed number of people, $n_h = 1$, per household (regardless of household size) if available.^a If $n_{\text{target}} < m_1$, take a simple random sample of n_{target} households from r_1 to obtain a smaller PPS sample of households. Then sample $n_h = 1$ per household if available.

^aIt is possible that a household is sampled without any members of the target age-sex group. Therefore, if

Household survey

For both the SRS and PPS schemes, the household survey module is administered to the head of household in each sampled household.

3.1 Simulated data

For each survey module, we simulate one outcome measured by that module. For the household survey we use the total household consumption; for the adult male survey we use the number of days after illness began when the man first sought advice or treatment; for the adult female survey we use the number of times a woman received antenatal care during her most recent pregnancy; for malaria and anemia testing we use hemoglobin blood concentration; for the anthropometry module we use the weight for age z-score. In generating simulated data, we make the simplifying assumption that all individuals in a target age-sex group have non-missing outcomes. For example, we generate antenatal care outcomes for all women of reproductive age.

To generate data, we use the multilevel model

$$\begin{aligned} \mathbf{y}_i &\sim \text{Normal}(\boldsymbol{\alpha}_{h[i]}, \boldsymbol{\sigma}_y) \text{ for individuals } i \\ \boldsymbol{\alpha}_h &\sim \text{Normal}(\boldsymbol{\mu} + \beta_1 \mathbf{N}_{h,\text{total}} + \beta_2 \mathbf{N}_{h,\text{total}}^2, \boldsymbol{\sigma}_\alpha) \text{ for households } h, \end{aligned} \quad (1)$$

For total household consumption we use a model analogous to model 1.^b

$$\mathbf{t}_h \sim \text{Normal}(\boldsymbol{\mu} + \beta_1 \mathbf{N}_{h,\text{total}} + \beta_2 \mathbf{N}_{h,\text{total}}^2, \mathbf{N}_{h,\text{total}} \boldsymbol{\sigma}_t) \text{ for households } h. \quad (2)$$

We also use a model for the log total consumption (which in our data is more Normally distributed than the total consumption),

$$\log(\mathbf{t}_h) \sim \text{Normal}(\boldsymbol{\mu} + \beta_1 \mathbf{N}_{h,\text{total}} + \beta_2 \mathbf{N}_{h,\text{total}}^2, \boldsymbol{\sigma}_t) \text{ for households } h. \quad (3)$$

We use the demographic information from the census and the multilevel model with estimated parameter values (from the survey data) to generate simulated populations. If when models 1, 2 or 3 are fit to past survey data, the 50% posterior interval of β_1 or β_2 contains 0, we set the parameter to 0 when simulating populations. This prevents us from using very noisy estimates of coefficients. Within each simulated population, we randomly sample according to the sampling plans described above, and estimate the finite population mean using either model-based or design-based inference.

$n_h = 1$, then the PPS scheme will result in a smaller sample size than the SRS scheme. Additionally, for the adult module (where $n_{\text{target}} = 400$), if $n_h = 1$ then the PPS scheme will at most sample only 300 adults, one per sampled household.

^bModel 2 can be motivated by assuming that model 1 holds for individual-level consumption (this would assume that within a household consumption is identically distributed, not taking into account age-sex differences). This model implies that $\mathbf{t}_h = \sum_{i|h[i]=h} \mathbf{y}_i \sim \text{Normal}(\mathbf{N}_{h,\text{total}} \boldsymbol{\alpha}_h, \mathbf{N}_{h,\text{total}} \boldsymbol{\sigma}_y)$ and that the marginal variance of \mathbf{t}_h is $\mathbf{N}_{h,\text{total}}^2 (\boldsymbol{\sigma}_\alpha^2 + \boldsymbol{\sigma}_y^2)$.

4 Bayesian model-based inference

To generalize from the data to the population, both design-based and model-based inference must take into account how the data are collected. Let $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_N)$ denote data for the population of interest, and $\mathbf{I} = (I_1, \dots, I_N)$ indicators of the observation of \mathbf{y} , where $I_i = 1$ if \mathbf{y}_i is sampled,^c and $I_i = 0$ if \mathbf{y}_i is missing. Let ‘obs’ = $\{i : I_i = 1\}$ and ‘mis’ = $\{i : I_i = 0\}$. Thus, the information available is $\mathbf{y}_{\text{obs}}, \mathbf{I}$, and the likelihood is $\mathbf{p}(\mathbf{y}_{\text{obs}}, \mathbf{I} | \mathbf{x}, \boldsymbol{\theta}, \boldsymbol{\phi}) = \int \mathbf{p}(\mathbf{y} | \mathbf{x}, \boldsymbol{\theta}) \mathbf{p}(\mathbf{I} | \mathbf{x}, \mathbf{y}, \boldsymbol{\phi}) d\mathbf{y}_{\text{mis}}$, where \mathbf{x} are observed covariates. Bayesian inference computes the posterior distribution $\mathbf{p}(\boldsymbol{\theta}, \boldsymbol{\phi} | \mathbf{x}, \mathbf{y}_{\text{obs}}, \mathbf{I})$ (superpopulation inference) and $\mathbf{p}(\mathbf{y}_{\text{mis}} | \mathbf{x}, \mathbf{y}_{\text{obs}}, \mathbf{I}, \boldsymbol{\theta}, \boldsymbol{\phi})$ (finite population inference). Under the *ignorability* condition, these inferences can be simplified to $\mathbf{p}(\boldsymbol{\theta} | \mathbf{x}, \mathbf{y}_{\text{obs}})$ and $\mathbf{p}(\mathbf{y}_{\text{mis}} | \mathbf{x}, \mathbf{y}_{\text{obs}}, \boldsymbol{\theta})$. Ignorability is satisfied if both the *missing at random* and *distinct parameters* conditions are satisfied (Gelman et al., 2014, p.202, 206-211). Missing at random requires that the missingness be independent of the missing values conditional on observed variables and a parameter $\boldsymbol{\phi}$: $\mathbf{p}(\mathbf{I} | \mathbf{x}, \mathbf{y}, \boldsymbol{\phi}) = \mathbf{p}(\mathbf{I} | \mathbf{x}, \mathbf{y}_{\text{obs}}, \boldsymbol{\phi})$. The distinct parameters condition requires that the parameters of the missingness mechanism ($\boldsymbol{\phi}$) be independent of the parameters of the data generating process ($\boldsymbol{\theta}$), conditional on covariates: $\mathbf{p}(\boldsymbol{\phi} | \mathbf{x}, \boldsymbol{\theta}) = \mathbf{p}(\boldsymbol{\phi} | \mathbf{x})$.

We include design variables such that the data collection mechanism is ignorable with respect to this model. For example, in our SRS-stratified sampling plan, the data collection mechanism is:

$$\mathbf{p}(\mathbf{I} | \mathbf{x}, \mathbf{y}, \boldsymbol{\phi}) = 1 / \left[\sum_{\substack{\mathbf{s}_I \subseteq \{1, \dots, N_I\} \\ |\mathbf{s}_I| = n_I}} \prod_{\mathbf{h} \in \mathbf{s}_I} \binom{N_{\mathbf{h}}}{n_{\mathbf{h}}} \right] \text{ where } n_{\mathbf{h}} = \text{round} \left(N_{\mathbf{h}} \frac{n}{N_{\mathbf{s}_I}} \right)$$

if $\exists \mathbf{s}_I \subseteq \{1, \dots, N_I\}$ s.t. $|\mathbf{s}_I| = n_I$ and $\sum_{i: \mathbf{h}[i]=\mathbf{h}} I_i = \text{round} \left(N_{\mathbf{h}} \frac{n}{N_{\mathbf{s}_I}} \right)$ for all $\mathbf{h} \in \mathbf{s}_I$. Otherwise, the probability of missingness pattern \mathbf{I} is zero.

Thus, we include as design variables the household identifiers and the $N_{\mathbf{h}}$ (e.g. the number of women per household, if the survey module targets women). Similar computations show that under the SRS-systematic sampling scheme these variables are also sufficient to satisfy missing at random. For the PPS scheme, we also will need the measure of household size used to select the households (e.g. the total number of household members under 50) (Gelman et al., 2014, p.211). For simplicity, we fit the same ignorable model for both the SRS and PPS schemes. For the anthropometry, blood, and adult survey modules we fit

$$\begin{aligned} \mathbf{y}_i &\sim \text{Normal}(\boldsymbol{\alpha}_{\mathbf{h}[i]}, \boldsymbol{\sigma}_y) \text{ for individuals } i \\ \boldsymbol{\alpha}_{\mathbf{h}} &\sim \text{Normal}(\boldsymbol{\mu} + \beta_1 N_{\mathbf{h}, \text{total}} + \beta_2 N_{\mathbf{h}, \text{total}}^2 + \beta_3 N_{\mathbf{h}} + \beta_4 N_{\mathbf{h}}^2, \boldsymbol{\sigma}_{\alpha}) \text{ for households } \mathbf{h}, \end{aligned} \tag{4}$$

For the household survey, we fit models 2 and 3.

Our parameter of interest is the finite population mean $\bar{Y} = \frac{1}{N} \sum_{\mathbf{h}=1}^{N_I} N_{\mathbf{h}} \bar{y}_{\mathbf{h}}$, where $\bar{y}_{\mathbf{h}} =$

^cWe assume that all units that are sampled are observed.

$\frac{n_h}{N_h} \bar{y}_{h,obs} + \frac{N_h - n_h}{N_h} \bar{y}_{h,mis}$ (Gelman et al., 2014, p.205). We obtain posterior simulations of \bar{Y} as follows: if household h is sampled, we use a simulation of α_h to generate $N_h - n_h$ simulated y_i 's. If household h is not sampled, we use simulations of μ and σ_α to simulate a new α_h , then generate N_h simulated y_i 's.

5 Frequentist design-based inference

We use the `survey` package to compute design-based estimates and variances (Lumley, 2004). Though we perform our SRS schemes without replacement, we compute all variances without finite population corrections, using the Horvitz Thompson (Hajek) ratio estimator and its with-replacement variance (Lohr, 2010, p.247).

Our SRS schemes are two-phase rather than two-stage designs, since the sampling within a household depends on which households were sampled in the first stage (Särndal et al., 1992, p.134-135). This dependence is reflected in the design weights we compute, see below. For the SRS-systematic sampling scheme, the independence assumption of two-stage sampling is also violated, with the sampling in each household dependent on the sampling in other households. Our design-based analysis approximates these two-phase designs with a two-stage analysis. In contrast, in model-based inference the details of the design do not matter in the analysis once we include design variables in our model (Gelman et al., 2014, p.202, 206-211).

5.1 Design weights

For the SRS-systematic design, the inclusion probabilities are:

$$\begin{aligned} \pi_{hi} &\equiv \text{P[person } i \text{ in household } h \text{ is sampled]} \\ &= \text{P}(h \in s_I) \text{P}(i \in s_h | h \in s_I) \\ &= \frac{n_I}{N_I} \sum_{s_I | h \in s_I} \text{P}(i \in s_h | h \in s_I, s_I) * \text{P}(s_I | h \in s_I) \\ &= \frac{n_I}{N_I} \sum_{s_I | h \in s_I} \underbrace{\min\left(\frac{n}{N_{s_I}}, 1\right)}_{(*)} * \frac{1}{\binom{N_I-1}{n_I-1}} \end{aligned}$$

For SRS-stratified, we replace $(*)$ with $\min\left(\frac{\text{round}\left(\frac{n}{N_{s_I}} N_h\right)}{N_h}, 1\right)$. In the simulations, instead of computing this precisely, we estimate it by randomly sampling s_I such that $h \in s_I$. This avoids the computationally intensive loop over all $\binom{N_I-1}{n_I-1}$ such sets. Although these weights are not equal for all individuals, because the distributions of household sizes (from the MVP demographic data) have no extreme outliers, in our simulations the weights are nearly equal.

For the PPS scheme, the inclusion probabilities are:

$$\begin{aligned}
\pi_{hi} &= E_k [P(\text{person } i \text{ in household } h \text{ is sampled} \mid \text{household } h \text{ is chosen } k \text{ times})] \\
&= E_k [1 - (1 - n_h/N_h)^k] \text{ since we independently subsample a household as many times as it is drawn.} \\
\text{Since } k &\sim \text{Bin}(m_I, p_h), \text{ by its probability generating function, we obtain} \\
&= 1 - (p_h(1 - n_h/N_h) + (1 - p_h))^{m_I} \\
&= 1 - \left(1 - p_h \frac{n_h}{N_h}\right)^{m_I} \\
\text{if } p_h \frac{n_h}{N_h} &\text{ is small, we can approximate this as:} \\
&= m_I p_h \frac{n_h}{N_h}
\end{aligned}$$

In PPS sampling, $p_h \propto x_h$, where x_h is a measure of household size (Särndal et al., 1992, p.97). So the PPS weights are:

$$w_{hi} = \frac{\sum_{h \in U_I} x_h N_h}{m_I x_h n_h}.$$

If $x_h \propto N_h$, and $n_h \propto 1$, then the design is self-weighted. We take $n_h = c$, a constant, but we cannot choose x_h such that $x_h \propto N_h$ for all modules, since the target age-sex groups differ from module to module. We chose $x_h = N_{h,\text{total}}$, the number of household members under 50 years of age, because it represented a compromise between the different target age-sex groups. Thus, our weights are $w_{hi} \propto \frac{N_h}{N_{h,\text{total}}}$.

6 Comparisons between sampling schemes: variances and design effects

We want to compare the PPS and SRS designs (in either the Bayesian model-based or the design-based paradigms). In general, the two schemes will have slightly different sample sizes, making direct comparisons of variances less relevant. For the household survey module, we fix the sample sizes to be equal, and for the adult, anthropometry, and blood modules, we adjust for the differing sample sizes by computing a design effect, defined below.

The household survey module is administered to the heads of households only, not individual members. Therefore, the time cost of the household module is mostly determined by the number of households surveyed. We set up our simulations such that the number of household heads to be interviewed (i.e. sample size) is the same for the SRS and PPS sampling schemes. We first perform a PPS sampling of households. Then, we use the number of unique sampled households to obtain the number of households to sample for the SRS scheme. We then directly compare the variances in estimating \bar{Y} , the finite population mean consumption per person.

For the remaining modules, we compute design effects. To define the design effect (often abbreviated as “deff”), we first introduce the following notation. Let $\widehat{\theta}_\pi = \widehat{\theta}_\pi(\mathbf{I}, \mathbf{y}_{\text{obs}})$ be the estimator of θ (in our case, $\theta = \bar{Y}$) where π is the sampling distribution assumed to have been used in drawing sample \mathbf{S} . Let $V_{\pi_1}(\widehat{\theta}_{\pi_2}; \mathbf{y})$ be the sampling variance of an estimator of θ that assumes sampling distribution π_2 , and π_1 is the distribution with respect to which we want the variance. Let $\widehat{V}_{\pi_1}(\widehat{\theta}_{\pi_2}; \pi_3; \mathbf{I}, \mathbf{y}_{\text{obs}})$ be an estimator where π_3 is the sampling distribution assumed to have been used in drawing sample \mathbf{S} . The population design effect is defined as $= V_{\mathbf{p}}(\widehat{\theta}_{\mathbf{p}}; \mathbf{y}) / V_{\text{SRS}}(\widehat{\theta}_{\text{SRS}}; \mathbf{y})$. The estimated design effect is defined as $= \widehat{V}_{\mathbf{p}}(\widehat{\theta}_{\mathbf{p}}; \mathbf{p}; \mathbf{y}_{\text{obs}}) / \widehat{V}_{\text{SRS}}(\widehat{\theta}_{\text{SRS}}; \mathbf{p}; \mathbf{y}_{\text{obs}})$.

In the design-based setting, we compute design effects assuming sampling with-replacement in both numerator and denominator variances. This is done in the `survey` package by specifying `deff = ‘replace’`.

For the model-based simulations, we estimate the numerator of the deff with the posterior variance for \bar{Y} from fitting a model that includes enough design variables such that the data collection mechanism is ignorable with respect to this model. This posterior variance includes an implicit finite population correction, so we compute a denominator variance that also includes such a correction:


$$\begin{aligned} V_{\text{SRS}}(\widehat{\theta}_{\text{SRS}}; \mathbf{y}) &= V_{\text{SRS}}(\bar{y}; \mathbf{y}) \\ &= \left(1 - \frac{n}{N}\right) \frac{S^2}{n} \end{aligned} \tag{5}$$

where $S^2 = \frac{1}{N-1} \sum_{i=1}^N (\mathbf{y}_i - \bar{Y})^2$.

To assess our estimated deff in the model-based setting, we compare the posterior variance $V_{\mathbf{p}\text{-ignorable}}(\theta | \mathbf{y}_{\text{obs}})$ from fitting an ignorable model with respect to a sampling distribution \mathbf{p} to the design-based sampling variance of the posterior means, $E_{\mathbf{p}\text{-ignorable}}(\theta | \mathbf{y}_{\text{obs}})$. The latter can be computed by simulation: we sample repeatedly from the full population using distribution \mathbf{p} , fit the \mathbf{p} -ignorable model, obtain a posterior mean of θ , and compute the variance of these across the samples from \mathbf{p} . Fixing one finite population, in Figure 1a we create a histogram of posterior variances from fitting the \mathbf{p} -ignorable model to each sample, and indicate with a vertical line the design-based variance of the posterior means, which is computed by simulation. We make the same comparison for $\mathbf{p} =$ a simple random sample (and its ignorable model with flat priors and no design variables), and include the closed-form design-based estimate (5) as a vertical line, in addition to the simulation-computed design-based estimate. See Figure 1b. We see that the posterior variances appear unbiased for the design-based variances.

7 Simulation results

Our results are displayed in Appendix B, where we see that neither the SRS nor PPS sampling of households is more efficient (i.e. has a lower design effect) in general.



(a) Sampling distribution \mathbf{p} is SRS sampling of households followed by an equal-probability systematic sample within households. (b) Sampling distribution \mathbf{p} is SRS sampling of people.

Figure 1: Fixing one finite population, we show a histogram of posterior variances from fitting a \mathbf{p} -ignorable model to each sample using sampling distribution \mathbf{p} , and indicate in a vertical line the design-based variance of the posterior means, which is computed by simulation. When \mathbf{p} is simple random sample of people, we also include the closed-form design-based estimate (5).

We see that for modules with higher target sample sizes, SRS tends to be less efficient. For example, in the under-5 blood ($n_{\text{target}} = 300$) and adult ($n_{\text{target}} = 400$) modules the SRS scheme is less efficient. One explanation for this observation is the different numbers of people sampled per household in the SRS versus PPS schemes, which has efficiency implications due to the intra-house correlation. In the PPS scheme, the households sampled in the first stage are larger and therefore more likely to include people in the target demographics. In contrast, in the SRS scheme, the sample is often drawn from fewer households, with more people sampled per household. Moreover, the PPS scheme only samples one person per household draw (though this can result in more than one person being sampled per household due to the with-replacement sampling at the first stage).

For modules where the target sample size is low, there are fewer people sampled per household in the SRS sampling scheme, and the intra-house correlation does not substantially impact the design efficiency. Therefore, because SRS has near-equal individual-level probability of sampling (see the design-weights computed above), its design effect in the absence of household clustering should be close to one. In contrast, the PPS scheme does not have near-equal individual-level sampling probabilities because the measure of household size is not proportional to the target demographic (see the design-weights computed above).

The relative efficiency of SRS versus PPS is similar between design-based and model-based simulations. In the few cases where they differ, design-based results show that SRS has higher design effects than PPS, relative to model-based results. In general, our model-based simulations show more variability across simulations than the design-based simulations. Comparing systematic to stratified sampling at the second stage of the SRS schemes, we see few differences except that stratified sampling tends to have higher variance across simulations.

8 Final sampling plan

As described above, the PPS scheme requires a sampling frame that includes household sizes, whereas the SRS scheme only requires a list of households. Given our results, we cannot justify the additional resources required to collect the more detailed household list for the PPS scheme. Therefore, our sampling scheme will begin with an SRS sample of households. For the second stage of sampling for the adult, anthropometry, and blood modules, we prefer the control over sample size achieved by systematic sampling (as opposed to stratified sampling).

The household and nutrition modules follow a different sampling scheme. As mentioned above, the household module is administered to all household heads (or other knowledgeable household members) within the sampled households. The nutrition module consists of a food frequency questionnaire, which takes longer to administer than other modules. We suspect that the within-household correlation is very high for data on food frequency, because household members are likely to eat similar foods. (This intra-house correlation cannot be measured from project data, because the project has always limited this module to one member per household.) For these reasons, we limit the nutrition module to one adult (age 15 to 49 years) per household.

9 Software

For fitting multilevel models we use Stan in R, (Stan Development Team, 2013; R Development Core Team, 2014).

Acknowledgements

The authors would like to thank the following people for very valuable feedback and ideas: Jeffrey D. Sachs, Joseph K. Blitzstein, Qixuan Chen, Jennifer Hill, Macartan Humphreys, Michael Clemens, Alberto Abadie, Marc Levy, Linda Pistoiesi, Keli Liu, Peng Ding, Natalie Exner, Abhishek Chakraborty, Rachael Meager, and Natalie Bau.

All mistakes are our own.

A Properties of systematic sampling

Definition 1 (The fractional interval method of systematic sampling). *Consider a population of size N^* consisting of people grouped into n_1 households indexed by h , with N_h people within-household h . Let $n_{\text{target}} < N^*$ be the desired sample size. Set $\alpha = \frac{N^*}{n_{\text{target}}}$. Order the households randomly, and randomly order the people in the target group within the households. Let $k = 1, \dots, N^*$ label the people in this order:*

$$\underbrace{1 \dots N_1}_{\text{household 1}} \underbrace{(N_1 + 1) \dots (N_1 + N_2)}_{\text{household 2}} \dots \underbrace{\dots N^*}_{\text{household } n_1}$$

Draw a random real number ξ , uniformly between 0 and α , $\xi \sim \mathcal{U}(0, \alpha)$, and sample all people with k such that

$$k - 1 < \xi + (j - 1)\alpha \leq k \text{ for } j = 1, \dots, n_{\text{target}}.$$

(Särndal et al., 1992, p.77)

Claim 1. *When performing the sampling scheme in Definition 1, the sample size will be n_{target} .*

Proof of Claim 1. Since $\alpha \equiv \frac{N^*}{n_{\text{target}}}$ and $N^* > n_{\text{target}}$, $\alpha > 1$. Since $k - 1 < x \leq k \Leftrightarrow \lceil x \rceil = k$, we can write $\lceil \xi + (j - 1)\alpha \rceil = k$. The ceiling function is monotone increasing and $\lceil x + 1 \rceil = \lceil x \rceil + 1$, so each time j increases by 1, we get a different value of k . Now we must show that the k 's stay in the set $\{1, \dots, N^*\}$, i.e. those from which we are sampling. The first k is such that $k - 1 < \xi \leq k$, where $\xi \in (0, \alpha)$. Since $\xi > 0$, we must have $k \geq 1$. The last k is such that $k - 1 < \xi + (n_{\text{target}} - 1)\alpha \leq k$, and we know $\xi + (n_{\text{target}} - 1)\alpha \leq \xi + N^* - \alpha < N^*$ because $\xi < \alpha$. Then $k \leq N^*$. Thus, since each j maps to a unique k , we've proven we get a sample size of exactly n_{target} . \square

Claim 2. *When performing the sampling scheme in Definition 1, the sample size within each household h is always the ceiling or the floor of the expected sample size in household h under simple random sampling: $\frac{N_h}{\alpha}$.*

We first prove the following lemma which is used to prove the above claim:

Lemma 1. *Consider the set $A(x) \equiv \{j \in \mathbb{Z}^+ \mid \xi + (j - 1)\alpha \leq x\}$. The maximum of $A(x)$ is $\frac{x}{\alpha}$ if $\frac{x}{\alpha} \in \mathbb{Z}$, $\lfloor \frac{x}{\alpha} \rfloor$ if $\xi > d\alpha$, or $\lceil \frac{x}{\alpha} \rceil$ if $\xi \leq d\alpha$, where $d \equiv \frac{x}{\alpha} - \lfloor \frac{x}{\alpha} \rfloor$, the “decimal part.”*

Proof of Lemma 1. If $\frac{x}{\alpha} \in \mathbb{Z}$, let $j = \frac{x}{\alpha}$, and we see that $\xi + (\frac{x}{\alpha} - 1)\alpha = \xi + x - \alpha \leq x$ because $\xi < \alpha$, so $\frac{x}{\alpha} \in A(x)$. Increasing j by 1 increases the lefthand side of the inequality by α , and since $\xi > 0$, we see this $\frac{x}{\alpha} + 1 \notin A(x)$. Therefore, $\frac{x}{\alpha}$ is the maximum.

If $\frac{x}{\alpha} \notin \mathbb{Z}$, let $d \equiv \frac{x}{\alpha} - \lfloor \frac{x}{\alpha} \rfloor$, the “decimal part.” We see that $\xi + (\lfloor \frac{x}{\alpha} \rfloor - 1)\alpha \underset{+d\alpha}{<} \xi + (\frac{x}{\alpha} - 1)\alpha = \xi + x - \alpha \underset{+a-\xi}{\leq} x$, so $\lfloor \frac{x}{\alpha} \rfloor \in A(x)$. Increasing j by 1 (to $\lceil \frac{x}{\alpha} \rceil$) increases the leftmost side of the inequality by α . If $\alpha \leq d\alpha + (\alpha - \xi)$, i.e. if $\xi \leq d\alpha$, then $\lceil \frac{x}{\alpha} \rceil \in A(x)$. \square


Proof of Claim 2. Consider household h , of size N_h . Let k^* be the last person before household h in the ordering used by the systematic sampling. Then $k^* + 1, \dots, k^* + N_h$ are the indices for all members of household h . In order to get the number sampled in household h , we consider the maximum of set $A(k^* + N_h)$ (the number of people sampled up through household h) and subtract from it the maximum of set $A(k^*)$ (the number of people sampled before household h). This gives, by Lemma 1, $\frac{k^* + N_h}{a} - \frac{k^*}{a} = \frac{N_h}{a}$. \square

Claim 3. *The sampling scheme in definition 1 is self-weighted.*

Proof.

$$\begin{aligned}
 P(\text{person } k \text{ is sampled}) &= P(\exists j \in \{1, \dots, n_{\text{target}}\} \text{ s.t. } k - 1 < \xi + (j - 1)a \leq k) \\
 &= P\left(\bigcup_{j=1}^{n_{\text{target}}} \{k - 1 - (j - 1)a < \xi \leq k - (j - 1)a\}\right) \\
 &\text{each interval is length } 1 = [k - (j - 1)a] - [k - 1 - (j - 1)a] \\
 &\text{space between intervals } j \text{ and } j+1 \text{ is } a - 1 = [k - 1 - j * a] - [k - (j - 1)a] \\
 &\text{so by the picture below, we see that } (0, a) \text{ has overlaps with the intervals} \\
 &\text{of length totaling } 1. \text{ So since } \xi \sim \mathbf{U}(0, a), \\
 &= \frac{1}{a}
 \end{aligned}$$

See below for a visual, in orange is the interval $(0, a)$, which can overlap at most 2 intervals of length 1 (shown as over-braces, with overlaps totaling a length 1):



\square

B Survey Sampling Simulation Results

B.1 Design-based results

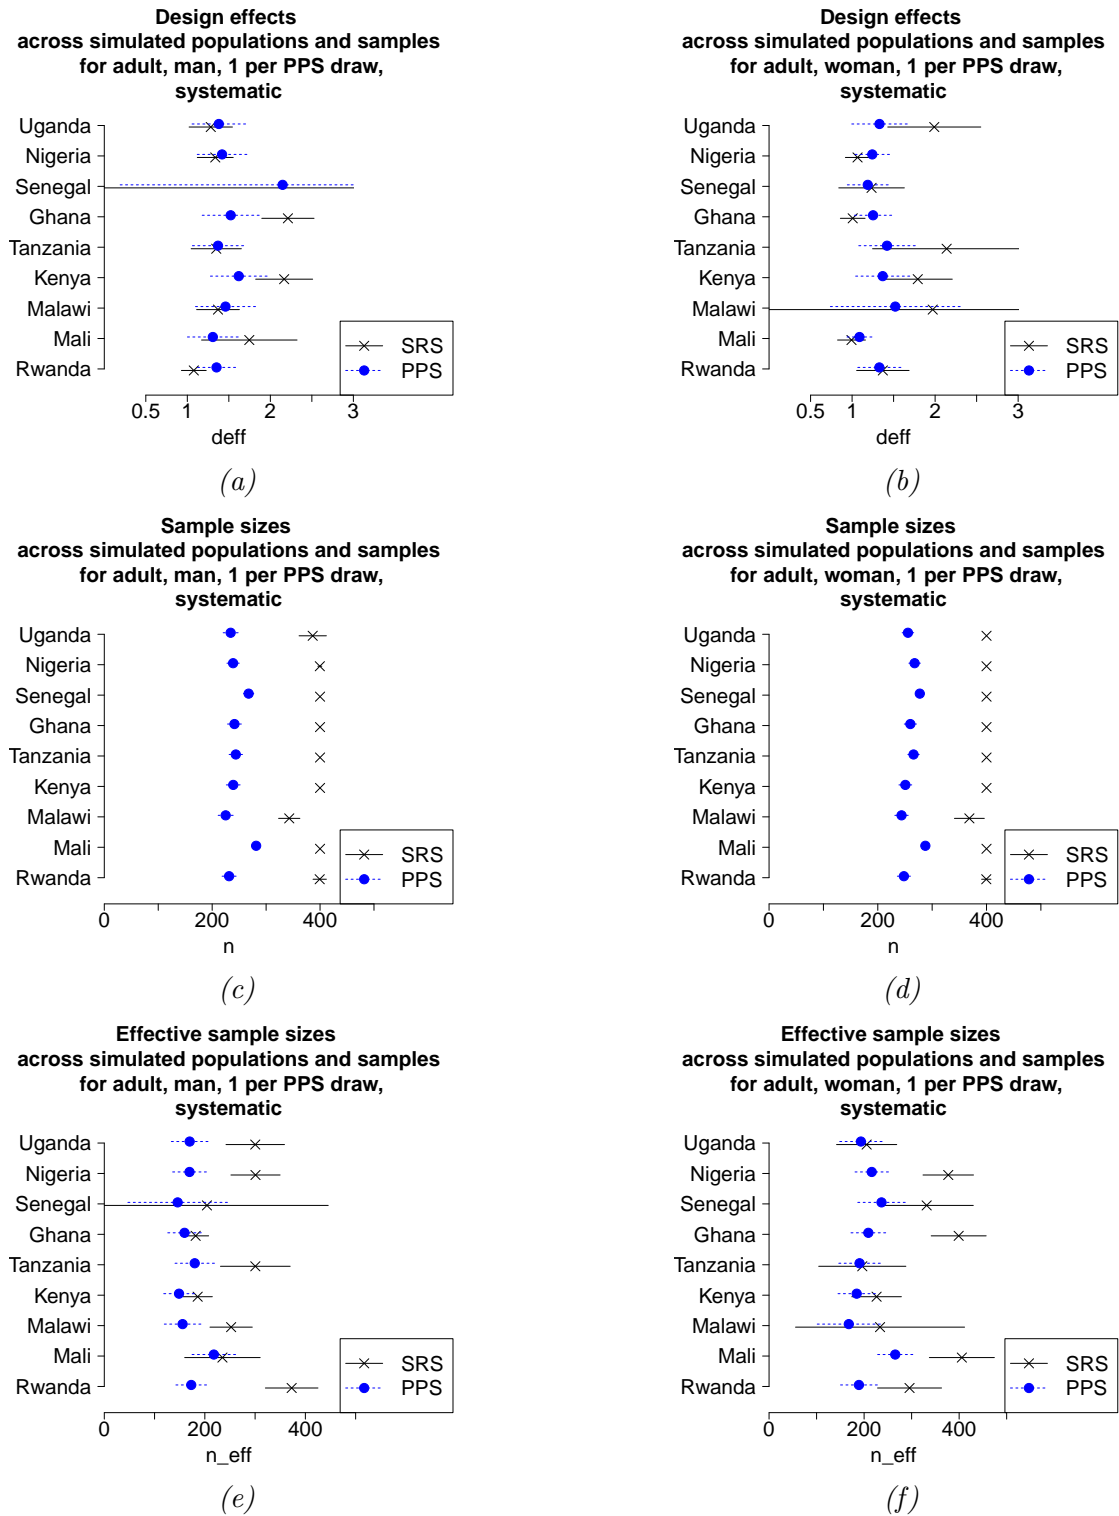




Figure 2: Design-based adult module results

Design effects
 across simulated populations and samples
 for blood, under_5, 1 per PPS draw,
 systematic




(a)

Design effects
 across simulated populations and samples
 for blood, school_age, 1 per PPS draw,
 systematic




(b)

Sample sizes
 across simulated populations and samples
 for blood, under_5, 1 per PPS draw,
 systematic




(c)

Sample sizes
 across simulated populations and samples
 for blood, school_age, 1 per PPS draw,
 systematic



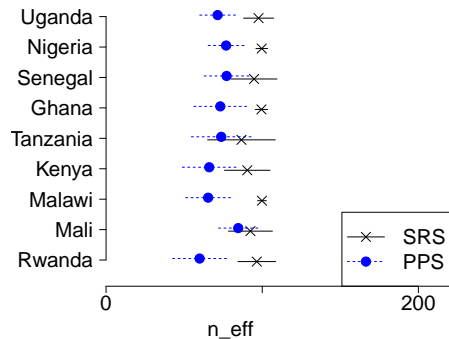
(d)

Effective sample sizes
 across simulated populations and samples
 for blood, under_5, 1 per PPS draw,
 systematic



(e)

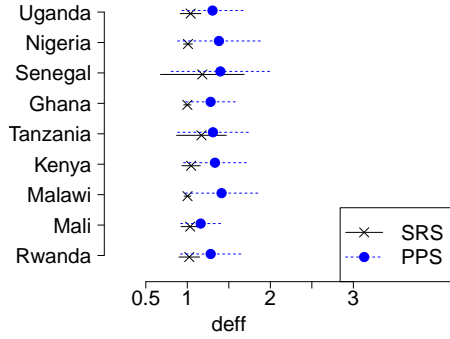
Effective sample sizes
 across simulated populations and samples
 for blood, school_age, 1 per PPS draw,
 systematic



(f)

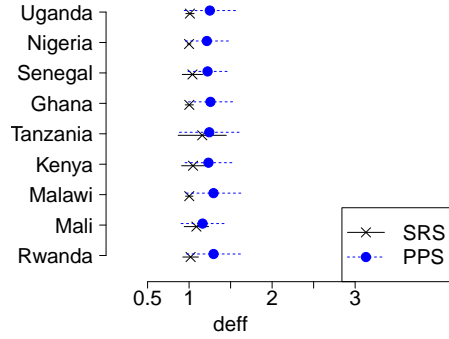
Figure 3: Design-based blood (malaria and anemia) module results: under 5 and school-age children

Design effects
 across simulated populations and samples
 for blood, man, 1 per PPS draw,
 systematic




(a)

Design effects
 across simulated populations and samples
 for blood, woman, 1 per PPS draw,
 systematic




(b)

Sample sizes
 across simulated populations and samples
 for blood, man, 1 per PPS draw,
 systematic




(c)

Sample sizes
 across simulated populations and samples
 for blood, woman, 1 per PPS draw,
 systematic




(d)

Effective sample sizes
 across simulated populations and samples
 for blood, man, 1 per PPS draw,
 systematic



(e)


Effective sample sizes
 across simulated populations and samples
 for blood, woman, 1 per PPS draw,
 systematic



(f)

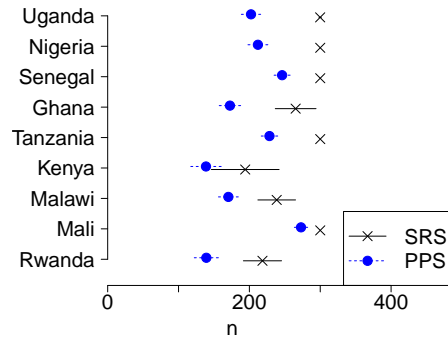
Figure 4: Design-based blood (malaria and anemia) module results: men and women

Design effects
across simulated populations and samples
for anthro, under_5, 1 per PPS draw,
systematic




(a)

Sample sizes
across simulated populations and samples
for anthro, under_5, 1 per PPS draw,
systematic



(b)


Effective sample sizes
across simulated populations and samples
for anthro, under_5, 1 per PPS draw,
systematic



(c)


Figure 5: Design-based anthro module results

Design effects
 across simulated populations and samples
 for adult, man, 1 per PPS draw,
 stratified



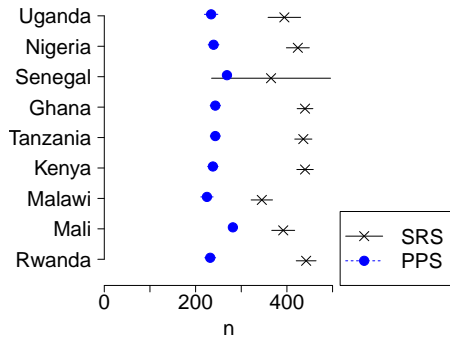
(a)

Design effects
 across simulated populations and samples
 for adult, woman, 1 per PPS draw,
 stratified




(b)

Sample sizes
 across simulated populations and samples
 for adult, man, 1 per PPS draw,
 stratified




(c)

Sample sizes
 across simulated populations and samples
 for adult, woman, 1 per PPS draw,
 stratified




(d)

Effective sample sizes
 across simulated populations and samples
 for adult, man, 1 per PPS draw,
 stratified



(e)


Effective sample sizes
 across simulated populations and samples
 for adult, woman, 1 per PPS draw,
 stratified



(f)


Figure 6: Design-based adult module results

Design effects
 across simulated populations and samples
 for blood, under_5, 1 per PPS draw,
 stratified




(a)

Design effects
 across simulated populations and samples
 for blood, school_age, 1 per PPS draw,
 stratified




(b)

Sample sizes
 across simulated populations and samples
 for blood, under_5, 1 per PPS draw,
 stratified




(c)

Sample sizes
 across simulated populations and samples
 for blood, school_age, 1 per PPS draw,
 stratified



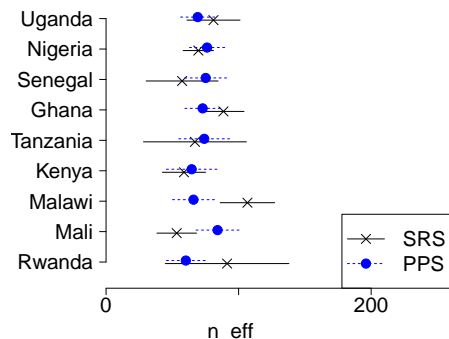
(d)

Effective sample sizes
 across simulated populations and samples
 for blood, under_5, 1 per PPS draw,
 stratified



(e)


Effective sample sizes
 across simulated populations and samples
 for blood, school_age, 1 per PPS draw,
 stratified



(f)

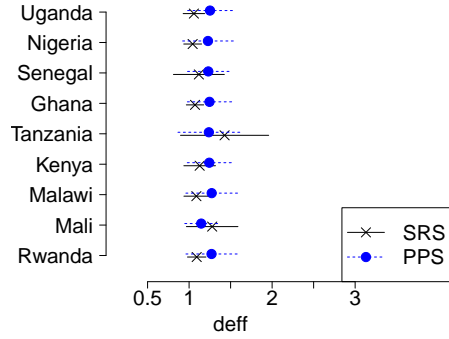
Figure 7: Design-based blood (malaria and anemia) module results: under 5 and school-age children

Design effects
 across simulated populations and samples
 for blood, man, 1 per PPS draw,
 stratified




(a)

Design effects
 across simulated populations and samples
 for blood, woman, 1 per PPS draw,
 stratified




(b)

Sample sizes
 across simulated populations and samples
 for blood, man, 1 per PPS draw,
 stratified




(c)

Sample sizes
 across simulated populations and samples
 for blood, woman, 1 per PPS draw,
 stratified




(d)

Effective sample sizes
 across simulated populations and samples
 for blood, man, 1 per PPS draw,
 stratified



(e)


Effective sample sizes
 across simulated populations and samples
 for blood, woman, 1 per PPS draw,
 stratified



(f)


Figure 8: Design-based blood (malaria and anemia) module results: men and women

Design effects
across simulated populations and samples
for anthro, under_5, 1 per PPS draw,
stratified




(a)

Sample sizes
across simulated populations and samples
for anthro, under_5, 1 per PPS draw,
stratified



(b)

Effective sample sizes
across simulated populations and samples
for anthro, under_5, 1 per PPS draw,
stratified



(c)

Figure 9: Design-based anthro module results

**Variance of the estimate of consumption
across simulated populations and samples
(number of sampled households equal for PPS & SRS)**




Figure 10: Design-based consumption module results

B.2 Model-based Results

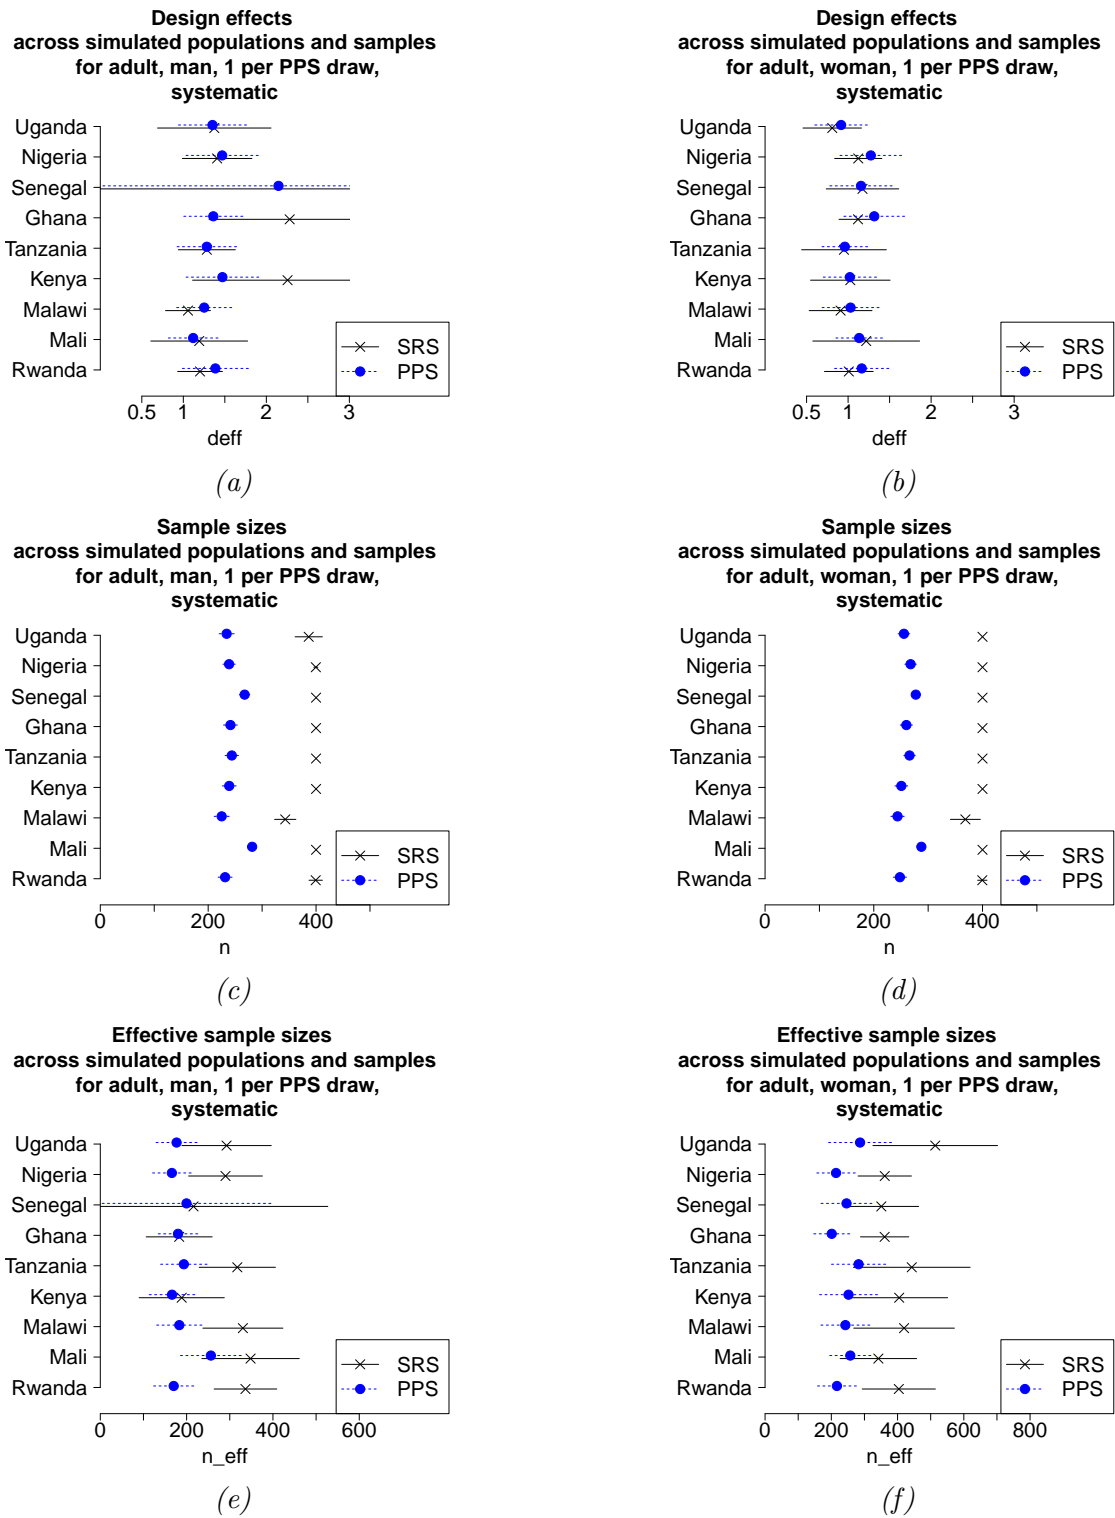




Figure 11: Model-based adult module results

Design effects
 across simulated populations and samples
 for blood, under_5, 1 per PPS draw,
 systematic




(a)

Design effects
 across simulated populations and samples
 for blood, school_age, 1 per PPS draw,
 systematic




(b)

Sample sizes
 across simulated populations and samples
 for blood, under_5, 1 per PPS draw,
 systematic




(c)

Sample sizes
 across simulated populations and samples
 for blood, school_age, 1 per PPS draw,
 systematic



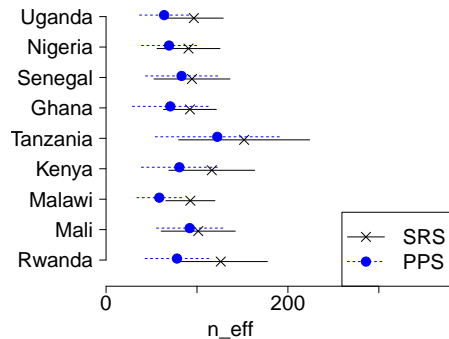
(d)

Effective sample sizes
 across simulated populations and samples
 for blood, under_5, 1 per PPS draw,
 systematic



(e)


Effective sample sizes
 across simulated populations and samples
 for blood, school_age, 1 per PPS draw,
 systematic



(f)


Figure 12: Model-based blood (malaria and anemia) module results: under 5 and school-age children

Design effects
 across simulated populations and samples
 for blood, man, 1 per PPS draw,
 systematic




(a)

Design effects
 across simulated populations and samples
 for blood, woman, 1 per PPS draw,
 systematic




(b)

Sample sizes
 across simulated populations and samples
 for blood, man, 1 per PPS draw,
 systematic




(c)

Sample sizes
 across simulated populations and samples
 for blood, woman, 1 per PPS draw,
 systematic




(d)

Effective sample sizes
 across simulated populations and samples
 for blood, man, 1 per PPS draw,
 systematic



(e)


Effective sample sizes
 across simulated populations and samples
 for blood, woman, 1 per PPS draw,
 systematic



(f)


Figure 13: Model-based blood (malaria and anemia) module results: men and women

Design effects
across simulated populations and samples
for anthro, under_5, 1 per PPS draw,
systematic




(a)

Sample sizes
across simulated populations and samples
for anthro, under_5, 1 per PPS draw,
systematic



(b)


Effective sample sizes
across simulated populations and samples
for anthro, under_5, 1 per PPS draw,
systematic



(c)


Figure 14: Model-based anthropometry module results

Design effects
 across simulated populations and samples
 for adult, man, 1 per PPS draw,
 stratified




(a)

Design effects
 across simulated populations and samples
 for adult, woman, 1 per PPS draw,
 stratified




(b)

Sample sizes
 across simulated populations and samples
 for adult, man, 1 per PPS draw,
 stratified




(c)

Sample sizes
 across simulated populations and samples
 for adult, woman, 1 per PPS draw,
 stratified




(d)

Effective sample sizes
 across simulated populations and samples
 for adult, man, 1 per PPS draw,
 stratified



(e)


Effective sample sizes
 across simulated populations and samples
 for adult, woman, 1 per PPS draw,
 stratified



(f)


Figure 15: Model-based adult module results

Design effects
 across simulated populations and samples
 for blood, under_5, 1 per PPS draw,
 stratified




(a)

Design effects
 across simulated populations and samples
 for blood, school_age, 1 per PPS draw,
 stratified




(b)

Sample sizes
 across simulated populations and samples
 for blood, under_5, 1 per PPS draw,
 stratified




(c)

Sample sizes
 across simulated populations and samples
 for blood, school_age, 1 per PPS draw,
 stratified




(d)

Effective sample sizes
 across simulated populations and samples
 for blood, under_5, 1 per PPS draw,
 stratified



(e)


Effective sample sizes
 across simulated populations and samples
 for blood, school_age, 1 per PPS draw,
 stratified



(f)


Figure 16: Model-based blood (malaria and anemia) module results: under 5 and school-age children

Design effects
 across simulated populations and samples
 for blood, man, 1 per PPS draw,
 stratified




(a)

Design effects
 across simulated populations and samples
 for blood, woman, 1 per PPS draw,
 stratified




(b)

Sample sizes
 across simulated populations and samples
 for blood, man, 1 per PPS draw,
 stratified




(c)

Sample sizes
 across simulated populations and samples
 for blood, woman, 1 per PPS draw,
 stratified




(d)

Effective sample sizes
 across simulated populations and samples
 for blood, man, 1 per PPS draw,
 stratified



(e)


Effective sample sizes
 across simulated populations and samples
 for blood, woman, 1 per PPS draw,
 stratified



(f)


Figure 17: Model-based blood (malaria and anemia) module results: men and women

Design effects
across simulated populations and samples
for anthro, under_5, 1 per PPS draw,
stratified




(a)

Sample sizes
across simulated populations and samples
for anthro, under_5, 1 per PPS draw,
stratified



(b)

Effective sample sizes
across simulated populations and samples
for anthro, under_5, 1 per PPS draw,
stratified



(c)

Figure 18: Model-based anthropometry module results

**Variance of the estimate of consumption
across simulated populations and samples
(number of sampled households equal for PPS & SRS)**




Figure 19: Model-based consumption module results

References

- A Gelman, J B Carlin, H S Stern, D B Dunson, A Vehtari, and D B Rubin. *Bayesian Data Analysis*. Chapman & Hall/CRC texts in statistical science, third edition, 2014.
- L Kish. Weighting for unequal p_i . *Journal of Official Statistics*, 8(2):183–200, 1992.
- S L Lohr. *Sampling: Design and Analysis*. Cengage Learning, 2 edition, 2010.
- T Lumley. Analysis of complex survey samples. *Journal of Statistical Software*, 9(8), April 2004.
- S Mitchell, A Gelman, R Ross, U Kim Huynh, L McClellan, M Harris, S Bari, J Chen, S Ohemeng-Dapaah, P Namakula, S E Sachs C Palm, and J D Sachs. The Millennium Villages Project: A protocol for the final evaluation. *Protocol in The Lancet*, 2015a. URL <http://www.thelancet.com/doi/story/10.1016/html.2015.07.03.2167>.
- S Mitchell, R Ross, S Makela, E A Stuart, A Feller, A M Zaslavsky, and A Gelman. Causal inference with small samples and incomplete baseline for the Millennium Villages Project. Working Paper, 2015b. URL http://www.stat.columbia.edu/~gelman/research/unpublished/MVP_paper_technical_JRSSA_short.pdf.
- R Development Core Team. The R project for statistical computing, February 2014. URL <http://www.r-project.org/>.
- J D Sachs and J W McArthur. The millennium project: a plan for meeting the millennium development goals. *The Lancet*, 365(9456):347–353, January 2005.
- P Sanchez, C Palm, J D Sachs, G Denning, R Flor, R Harawa, B Jama, T Kifleariam, B Konecky, R Kozar, E Lelerai, A Malik, P Mutuo, A Niang, H Okoth, F Place, S E Sachs, A Said, D Siriri, A Teklehaimanot, K Wang, J Wangila, and C Zamba. The african millennium villages. *Proceedings of the National Academy of Sciences*, 104(43):6775–80, 2007.
- C E Särndal, B Swensson, and J Wretman. *Model Assisted Survey Sampling*. Springer-Verlag, New York, 1992.
- Stan Development Team. Stan: A c++ library for probability and sampling, version 1.3, February 2013. URL <http://mc-stan.org/>.