



Optimum Allocation for Adaptive Multi-Wave Sampling in R: The R Package **optimall**

Jasper B. Yang 

University of Washington

Thomas Lumley 

University of Auckland

Bryan E. Shepherd 

Vanderbilt University

Pamela A. Shaw 

Kaiser Permanente Washington
Health Research Institute

Abstract

The R package **optimall** offers a collection of functions that efficiently streamline the design process of sampling in surveys ranging from simple to complex. The package's main functions allow users to interactively define and adjust strata cut points based on values or quantiles of auxiliary covariates, adaptively calculate the optimum number of samples to allocate to each stratum using Neyman or Wright allocation, and select specific units to sample based on a stratified sampling design. Using real-life epidemiological study examples, we demonstrate how **optimall** facilitates an efficient workflow for the design and implementation of surveys in R. Although tailored towards multi-wave sampling under two- or three-phase designs, the R package **optimall** may be useful for any sampling survey.

Keywords: optimal design, Neyman allocation, multi-wave sampling, R.

1. Introduction

Sample surveys are widely used by epidemiologists, political scientists, government organizations, and many other researchers to gather information about finite populations when complete enumerations of variables of interest are unreasonable. A common objective for these surveys is efficiency, meaning that the estimates generated from the budget-constrained number of samples aim to be as precise and accurate as possible. Maximizing efficiency under stratified sampling designs requires efficiently allocating samples to strata, but this task is rarely straightforward because the true optimal design frequently relies on unknown parameters. Traditional methods for approximating optimal design include proportional allocation

and x-optimal allocation, although the sampling designs that these strategies yield are only efficient under specific assumptions (Särndal, Swensson, and Wretman 2003).

Studies in epidemiology present an interesting case when certain variables such as laboratory tests are expensive to collect, but other variables such as demographic characteristics can easily be obtained for all or most of the population from electronic health records. In these conditions, a common approach is a two-phase sampling design, where readily available or inexpensive variables are obtained for all sampling units in phase one and used to inform the stratified sampling scheme for the collection of the expensive variables on a subsample in phase two (Neyman 1938; Tao, Zeng, and Lin 2020). Recent studies have demonstrated that even more precise estimates of the true population parameters can be obtained when the second phase is conducted over a series of adaptive waves (McIsaac and Cook 2015; Chen and Lumley 2020). Under this approach, which is often referred to as adaptive multi-wave sampling, estimates for the parameters specifying the underlying optimal design are updated after each wave and used to guide the design for the following round of expensive data collection. While theoretically appealing, a drawback of adaptive multi-wave sampling is that its implementation is more tedious and complicated compared to sampling in a single wave. In this article, we describe the R (R Core Team 2025) package **optimall**, which facilitates the processes of stratification, sample selection, optimum allocation, and workflow organization in surveys ranging from simple to complex. Package **optimall** (Yang and Shaw 2025) is available from the Comprehensive R Archive Network (CRAN) at <https://CRAN.R-project.org/package=optimall>.

A number of useful packages have been developed for the analysis of sampling survey data in R, notably the **survey** package (Lumley 2010), but fewer exist that are specifically tailored towards sampling design. Among the few are **surveyplanning** (Breidaks, Liberts, and Jukams 2020), which offers tools for the early stages of survey design including precision estimation and basic optimum allocation, and **stratifyR** (Reddy and Khan 2020), which uses a dynamic programming technique to estimate optimum strata boundaries and sample sizes for univariate populations under certain distributional assumptions. The packages **stratification** (Rivest and Baillargeon 2022) and **SamplingStrata** (Barcaroli, Ballin, Odendaal, Pagliuca, Willighagen, and Zardetto 2022) also implement optimal stratification under specific constraints for univariate and multivariate cases respectively. **PracTools** (Valliant and Zipf 2025) provides functions for sample size calculation and variance estimation in multistage sampling designs, and **sampling** (Tillé and Matei 2025) allows users to randomly draw samples based on various sampling schemes. These packages are useful for specific tasks under certain survey conditions, but none provide a framework for an entire multi-wave sampling workflow. Such a framework may be especially important for the iterative process of stratifying, optimally allocating samples, taking samples, and merging preliminary results, which can be cumbersome and error-prone. The **optimall** package seeks to fill this role. To the authors' knowledge, this is the first R package created specifically for multi-wave sampling design.

The **optimall** package offers a collection of functions that are relatively simple at their core, but they efficiently streamline the design process of sampling studies ranging from simple to complex when implemented together. The package is particularly tailored towards multi-wave sampling under two- or three-phase designs but may be useful for any sampling survey. In the following sections we introduce the methodological framework underlying the functions, describe the key features of **optimall** and how they are each implemented, and finally demonstrate their use through practical examples based on real-life epidemiological studies.

2. Methodological background

2.1. Multi-phase sampling designs

Two-phase (more generally, multi-phase) sampling designs were first introduced by [Neyman \(1938\)](#) as a strategy for developing an efficient sampling design when some variables of interest are expensive or difficult to collect, but others can be collected on a large portion of the population at a low cost. Such a scenario is common in modern epidemiological studies when biomarker test results or hand-validated electronic health records can only reasonably be measured for some patients, but other variables such as demographic information can be extracted for the entire population.

Under two-phase designs, which are detailed in [Särndal *et al.* \(2003\)](#), the objective of the first phase is to measure inexpensive variables on a large sample of the population. The results of this phase are then used to guide the sampling design for the next phase by defining strata boundaries or informing models for optimal selection. Assuming that the phase one variables are correlated at least partially with the phase two variables, this approach leads to a more efficient sampling design than simple random sampling for the collection of expensive variables in phase two.

2.2. Neyman allocation

Assuming that the per-unit sampling cost is the same in each stratum and that S_h , the standard deviation of the variable of interest in stratum h , can be estimated, [Neyman \(1934\)](#) presented the following solution to optimally allocate n samples among H strata:

$$n_h = n \frac{N_h S_h}{\sum_{i=1}^H N_i S_i},$$

where N_h is the number of population elements in stratum h and n_h is the number to be sampled from stratum h . This formula, known as Neyman allocation, is widely used as a robust method for minimizing the variance of an estimate for the population mean. Neyman allocation offers the advantage of producing exact sampling fractions that can later be multiplied by n or be taken on their own if n is not known. However, this theoretical advantage may sometimes lead to sub-optimal practical applications because it means that Neyman Allocation is not constrained to integer solutions. When taking a fraction of a sample is not reasonable, study designers are forced to stray from the theory by rounding stratum sample sizes to the nearest integer in ways that 1) may not be exactly optimal and 2) may no longer sum to n .

2.3. Wright's exact optimization algorithm

To address the issues involved with practical applications of Neyman allocation, [Wright \(2012\)](#) presented an algorithm for optimum sampling allocation based on the equal proportions method of allocating a fixed number of seats in the United States House of Representatives among states.

Like Neyman, the purpose of Wright's algorithm is to minimize the variance of the sample

mean, \hat{T} , from stratified random sampling, which can be written as

$$\text{VAR}(\hat{T}) = \sum_{h=1}^H \frac{N_h^2 S_h^2}{n_h} - \sum_{h=1}^H \frac{N_h^2 S_h^2}{N_h}. \quad (1)$$

When the total sample size, n , is fixed, Wright notes that the variance of \hat{T} can be minimized by focusing on

$$\sum_{h=1}^H \frac{N_h^2 S_h^2}{n_h}$$

because the second sum in (1) is a constant based on the known population sizes in each stratum, and both N_h and S_h must always be positive. Noting that $1/n_h$ can be written as a sum, we have

$$\sum_{h=1}^H \frac{N_h^2 S_h^2}{n_h} = \sum_{h=1}^H N_h^2 S_h^2 - \sum_{h=1}^H \left(\frac{N_h^2 S_h^2}{1 \cdot 2} + \frac{N_h^2 S_h^2}{2 \cdot 3} + \cdots + \frac{N_h^2 S_h^2}{(n_h - 1)(n_h)} \right). \quad (2)$$

Assuming that at least one sample should be allocated to each stratum, Wright points out that the integer sampling allocation among H strata that minimizes $\text{VAR}(\hat{T})$ is found by maximizing the sum in the second term of (2) under the constraint $\sum_{h=1}^H n_h = n$. This is the same as maximizing

$$\sum_{h=1}^H \left(\frac{N_h S_h}{\sqrt{1 \cdot 2}} + \frac{N_h S_h}{\sqrt{2 \cdot 3}} + \cdots + \frac{N_h S_h}{\sqrt{(n_h - 1)(n_h)}} \right). \quad (3)$$

Referring to each term in the sum above as a “priority value”, Wright points out that selecting the largest priority values across all h will maximize the sum. In what he calls Algorithm I, each stratum is required to have at least 1 sample, so the $n - H$ largest values of (3) across all strata determine the optimum allocation. Each stratum h starts with an allocation of 1 sample, and it receives an additional sample for every priority value of h that lies in the largest $n - H$ overall priority values.

Under this framework, any sample size constraints can be easily incorporated into the allocation process. For example, each stratum needs at least 2 samples in order to generate an unbiased estimate of $\text{VAR}(\hat{T})$. Wright incorporates this in his Algorithm II by assigning 2 samples to each stratum and selecting the remaining $n - 2H$ samples based on the largest priority values but now ignoring the largest priority value per stratum, the $\frac{N_h S_h}{\sqrt{1 \cdot 2}}$ term.

2.4. Optimum allocation for regression parameters

The previous sections describe well-known methods for minimizing the variance of an estimate for the population mean of a variable of interest given a fixed sample size. In practice however, research interest often lies in the relationship between covariates and an outcome of interest, which is explored through regression modeling. [Chen and Lumley \(2020\)](#) and [Amorim, Tao, Lotspeich, Shaw, Lumley, and Shepherd \(2021\)](#) point out that Neyman and Wright optimum allocation strategies are still useful in these cases because regression parameters can be considered as the sum of their influence functions. Thus, they show that the allocation of samples

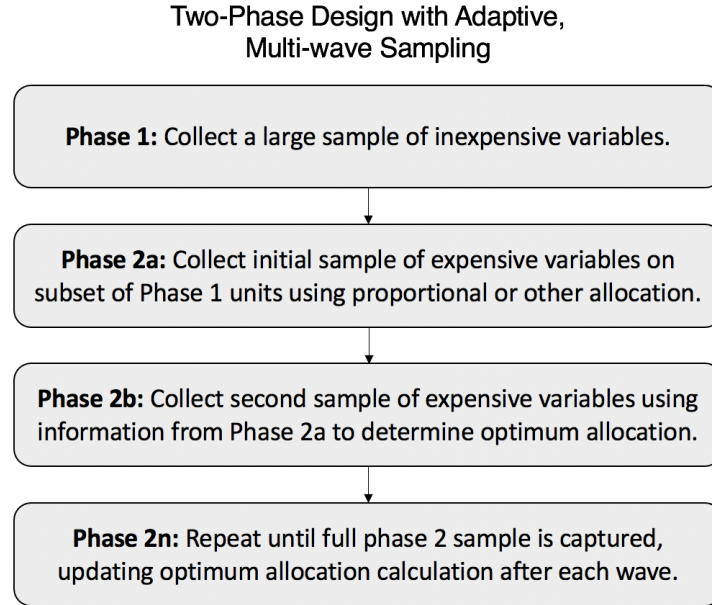


Figure 1: Flowchart of multi-wave sampling under a two-phase design.

to strata that minimizes the variance of the estimate of the sum of influence functions leads to the optimal sampling design, which can be written as:

$$n_h = n \frac{N_h S_k^*}{\sum_{i=1}^H N_i S_i^*},$$

where S_k^* is the standard deviation of influence functions for the regression coefficient of interest in stratum h .

2.5. Optimum allocation in multi-wave sampling

The optimum allocation algorithms presented by Neyman and Wright are useful in the design of stratified sampling surveys; however, they rely on standard deviation estimates for the variable of interest which are typically unknown at the design stage. One approach to this issue is to approximate the true standard deviations using information from previous surveys or using an inexpensive auxiliary variable strongly correlated with the variable of interest (Särndal *et al.* 2003). A second, and more effective, approach proposed by McIsaac and Cook (2015) is to conduct the sampling of the expensive variables over a series of waves, recalculating the optimum sampling proportion after each new wave of data collection. Under this multi-wave sampling approach, necessary within-strata standard deviation estimates, and thus the estimated optimal sampling proportions, are expected to be closer to their true value as the validated data accumulate. As such, the benefits of adaptive, multi-wave sampling better assure the cost-effective efficiency gains offered by two-phase (and, more generally, multi-phase) sampling designs.

In summary, when full data collection on all subjects is not feasible, we recommend a two-phase sampling design with adaptive, multi-wave sampling in the second phase (Figure 1).

This process is based upon the designs proposed by [McIsaac and Cook \(2015\)](#), [Chen and Lumley \(2020\)](#), and [Han, Lumley, Shepherd, and Shaw \(2021a\)](#). It is applicable across a wide range of epidemiological studies and can be easily implemented in R using the package **optimall**.

3. Features of **optimall**

The R package **optimall** offers a collection of functions that are designed to streamline the process of optimum sample allocation, specifically under an adaptive, multi-wave approach. Its main functions allow users to

- Define, split, or merge strata based on values or percentiles of variables.
- Determine the optimum number or fraction of samples to allocate to each stratum in a given study in order to minimize the variance of a sample mean of interest.
- Select specific IDs to sample based on a given stratified sampling design.
- Optimally allocate a fixed number of samples to a sampling wave based on results from a prior wave.

When used together, these functions can automate most of any stratified sampling workflow. This section describes in detail how each of these individual features is implemented, using examples generated with the **iris** dataset from the **datasets** package ([R Core Team 2025](#)). **optimall** generally works with data in this format, where each row represents one sampling unit and each column represents one variable.

3.1. Defining, splitting, and merging strata

In stratified sampling, strata are typically defined on values or quantiles of variables available for the entire population. Given a dataset with one row per sampling unit and at least one column containing a variable that can be used to construct strata, **optimall** allows users to easily define, split, and merge strata.

Suppose that we have defined three strata in the **iris** dataset based on Species. The dataset contains 50 observations per strata:

```
R> library("datasets")
R> table(iris$Species)
```

```
      setosa versicolor virginica
      50         50         50
```

Suppose that we want to split two out of the three strata, setosa and virginica, based on the within-stratum median of "Sepal.Width". In **optimall**, we can do this by calling the **split_strata()** function:

```
R> iris2 <- split_strata(data = iris, strata = "Species",
+   split = c("setosa", "virginica"), split_var = "Sepal.Width",
+   split_at = c(0.5), type = "local quantile")
```

`split_strata()` outputs a data frame that matches the input data frame, but contains a new column called "new_strata" that holds the new strata:

```
R> table(iris2$new_strata)

      setosa.Sepal.Width_(3.4,4.4]  setosa.Sepal.Width_[2.3,3.4]
                                22                             28
      versicolor virginica.Sepal.Width_(3,3.8]
                                50                             17
virginica.Sepal.Width_[2.2,3]
                                33
```

In this example, we split strata at the within-stratum median, as specified by `type = "local quantile"` and `split_at = c(0.5)`. The `type` argument also accepts "global quantile", "value", or "categorical", allowing for a wide range of split types, and `split_at` accepts vectors of cut points, so multiple splits can be performed with a single call to the function. Also note that the new stratum names display the boundaries on which each stratum has been defined. If the new names are too long, the `trunc` argument can be used to specify how they should be trimmed.

Similarly, we can merge strata using the `merge_strata()` function:

```
R> iris3 <- merge_strata(data = iris, strata = "Species",
+   merge = c("setosa", "versicolor"), name = "set_or_vers")

R> table(iris3$new_strata)

set_or_vers  virginica
           100         50
```

3.2. Optimally allocating samples

optimall allows users to perform Neyman allocation, Wright Algorithm I, or Wright Algorithm II on any stratified dataset with a continuous variable. The function defaults to using Wright Algorithm II, which requires that at least 2 samples are taken from each stratum. In **optimall**, stratum sampling sizes for both Wright algorithms are also constrained from above at N_h , the population stratum sizes.

Suppose we want to optimally allocate 40 samples among species in the iris dataset, minimizing the variance of the "Sepal.Width" sample mean. We can use `optimum_allocation()`:

```
R> sampling_design <- optimum_allocation(data = iris, strata = "Species",
+   y = "Sepal.Width", nsample = 40, method = "WrightII")

R> sampling_design
```

	strata	npop	sd	n_sd	stratum_fraction	stratum_size
1	setosa	50	0.38	18.95	0.38	15
2	versicolor	50	0.31	15.69	0.30	12
3	virginica	50	0.32	16.12	0.32	13

In the example above, `optimum_allocation()` determined the optimum sample allocation using standard deviation estimates generated from the full data, which held one row per population unit and contained a single column holding the variable of interest. If we instead have a data frame that holds the N_h and the within-stratum standard deviations, sd_h , for each stratum rather than data for each individual unit, we can still use `optimum_allocation()`:

```
R> iris_summary <- data.frame(strata = unique(iris$Species),
+   size = c(50, 50, 50), sd = c(0.3791, 0.3138, 0.3225))
R> optimum_allocation(data = iris_summary, strata = "strata",
+   sd_h = "sd", N_h = "size", nsample = 40, method = "WrightII")
```

	strata	npop	sd	n_sd	stratum_fraction	stratum_size
1	setosa	50	0.38	18.95	0.38	15
2	versicolor	50	0.31	15.69	0.30	12
3	virginica	50	0.32	16.12	0.32	13

This functionality may be especially useful for more complex optimization problems, such as allocation with respect to multiple parameters.

Regardless of input format, `optimum_allocation()` outputs a data frame with one row per stratum, and columns holding the specified total sample size, the stratum fraction, and the number of samples allocated across strata if `nsample` is given. We refer to this format as a “design data frame”. Assuming simple random sampling within strata, the design data frame specifies the sampling design for a given wave.

3.3. Sampling based on results

With the number of units to sample per stratum specified in a design data frame, **optimall** can select the IDs of the units to be sampled using simple random sampling without replacement within strata with the function `sample_strata()`:

```
R> iris$id <- 1:150
R> set.seed(743)
R> iris <- sample_strata(data = iris, strata = "Species", id = "id",
+   design_data = sampling_design, design_strata = "strata",
+   n_allocated = "stratum_size", probs = ~ stratum_size/npop)
```

The output of `sample_strata()` is the same input data frame with a new column called “`sample_indicator`” that holds a binary indicator for whether each unit should be sampled in the specified wave. This column can be given a more specific name by specifying a `wave` number in the function call. The output also includes a new column called “`sampling_prob`” holding the sampling probability attached to each selected sample, which is specified by the `probs` argument:

```
R> head(iris[, c("id", "Species", "Sepal.Width", "sample_indicator",
+   "sampling_prob")])
```

	id	Species	Sepal.Width	sample_indicator	sampling_prob
1	1	setosa	3.5	0	NA

2	2	setosa	3.0	0	NA
3	3	setosa	3.2	0	NA
4	4	setosa	3.1	1	0.3
5	5	setosa	3.6	0	NA
6	6	setosa	3.9	1	0.3

From this data frame, we can easily extract a vector of the ids to sample:

```
R> ids_to_sample <- iris$id[iris$sample_indicator == 1]
R> head(ids_to_sample)
```

```
[1] 4 6 8 11 14 15
```

```
R> length(ids_to_sample)
```

```
[1] 40
```

In cases where some units have already been sampled, `sample_strata()` will only select new units if an indicator for previously sampled is specified in the `already_sampled` argument. Note that the `design_data` argument refers to the output of `optimum_allocation()` in this example, but in practice it can be any data frame that has one row per stratum and one column specifying each stratum's desired sample size. As such, any method for allocating samples to strata can be implemented by `sample_strata()` as long as the design data frame is specified and simple random sampling within strata is used.

3.4. Allocating samples in waves

In a multi-wave sampling survey similar to the one described in Section 2.5, **optimal** allows users to determine the optimum allocation for a given wave, taking into consideration the samples from previous waves, with the function `allocate_wave()`. This function calculates the optimum sample allocation for the grand total sample size (cumulative of previous wave sizes and current wave size) to minimize the variance of the estimate for a mean of a given continuous variable using the Wright's allocation algorithm II for a fixed sample size. It then determines the allocation for the current wave by taking the difference between the previous stratum sample size and the total optimum sample size. If some strata have already been oversampled, it re-calculates the optimum allocation among the non-oversampled strata. In this case, the realized allocation will be sub-optimal.

For a simple example, suppose that we again want to allocate 40 samples to minimize the variance of "Sepal.Width" in the iris dataset, but 30 out of the 40 samples have already been collected in wave 1. We assume that the strata are still defined only by species, and that 16 of the first 30 samples were taken from the virginica species, 7 from setosa, and 7 from versicolor. Assuming we only have these 30 samples to base our within-stratum standard deviation estimates on, we can allocate the next 10 samples using `allocate_wave()`.

First, we set up wave 1 by collecting sepal width from the 30 samples using the `sample_strata` function and a manually specified design data frame:

```

R> wave1_design <- data.frame(
+   strata = c("setosa", "virginica", "versicolor"),
+   stratum_size = c(7, 16, 7))
R> phase1_data <- subset(datasets::iris, select = -Sepal.Width)
R> phase1_data$id <- 1:nrow(phase1_data)
R> set.seed(234)
R> phase1_data <- sample_strata(data = phase1_data, strata = "Species",
+   id = "id", design_data = wave1_design, design_strata = "strata",
+   n_allocated = "stratum_size")
R> wave1_ids <- iris$id[phase1_data$sample_indicator == 1]
R> wave1_sampled_data <- iris[iris$id %in% wave1_ids, c("id", "Sepal.Width")]
R> wave1_data <- merge(phase1_data, wave1_sampled_data, by = "id",
+   no.dups = TRUE, all.x = TRUE)

```

We have now collected the sepal width for our first 30 samples:

```
R> table(is.na(wave1_data$Sepal.Width), wave1_data$Species)
```

	setosa	versicolor	virginica
FALSE	7	7	16
TRUE	43	43	34

Next, we can run `allocate_wave()` to generate the wave 2 design:

```

R> wave2_design <- allocate_wave(data = wave1_data, strata = "Species",
+   y = "Sepal.Width", already_sampled = "sample_indicator", nsample = 10,
+   detailed = TRUE)
R> wave2_design[,1:6]

```

	strata	npop	nsample_optimal	nsample_actual	nsample_prior	n_to_sample
1	setosa	50	19	15	7	8
2	versicolor	50	12	9	7	2
3	virginica	50	9	16	16	0

Notice that in this case, the total optimum sample allocation, "nsample_optimal", does not match the realized allocation, "nsample_actual", for every stratum in the outputted wave 2 design data frame. This is because we oversampled from the virginica stratum in wave 1, meaning that the optimum stratum sample size among 40 total samples was smaller than the amount of samples already taken in that group. When no oversampling occurs, "nsample_optimal" will match "nsample_actual" for every stratum.

Finally, we can use `sample_strata` to select specific samples based on this design:

```

R> wave2_data <- sample_strata(data = wave1_data, strata = "Species",
+   id = "id", design_data = wave2_design, design_strata = "strata",
+   n_allocated = "n_to_sample")

```

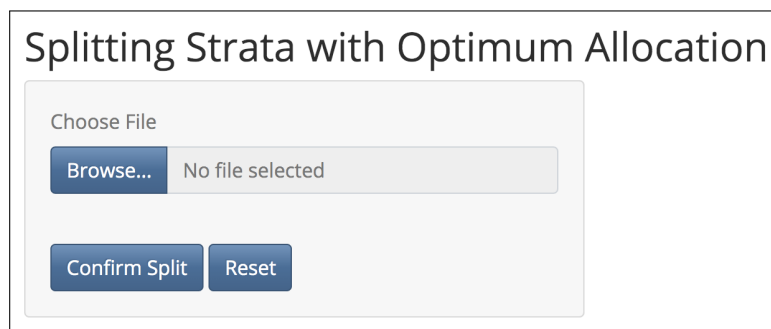


Figure 2: Uploading a file in the **shiny** application.

4. shiny app for interactively defining strata

In stratified sampling, efficiency is gained when strata cut points are chosen to minimize within-stratum variances and maximize the variance across strata. In some cases, it is also desirable for each of the strata to have similar sample sizes (Amorim *et al.* 2021). Determining the split points on which to define strata also typically relies on unknown parameters, but estimates for the within-stratum variances can be obtained through an auxiliary variable correlated with the variable of interest or from previous sampling waves in which the variable of interest was collected. The `split_strata()` function in **optimall** is a versatile tool for defining strata, but it is designed for situations where the strata cut points have already been decided by the user. Running it iteratively to experiment with different splits is possible yet tedious.

To aid with the process of defining strata, **optimall** offers a **shiny** (Chang *et al.* 2025) app that can be launched with the `optimall_shiny()` function. The app allows users to interactively select strata split points and view how the stratum sizes, standard deviations, and optimum allocations react. This feature may be useful as a tool for strata selection on its own or as a complement to more advanced algorithms such as stratification by k -means (MacQueen 1967). This section demonstrates the utility of `optimall_shiny()`.

4.1. Launching the app and uploading data

Launching the app is simple:

```
R> optimall_shiny()
```

Once open, the first step is uploading data (Figure 2). The data must be in the same format that the rest of the functions in **optimall** require, meaning that each row must correspond to one sampling unit. Any .csv, .rda, or .rds file can be uploaded in `optimall_shiny()`.

4.2. Reactive strata selection

After the dataset has been uploaded, options for the user to select inputs will appear. Each input corresponds to an argument of `split_strata()`, `optimum_allocation()`, or `allocate_wave()`. Values for each input can be selected from the list or slider of valid options (Figure 3):

Choose CSV File

Browse... iris.csv

Upload complete

Column Holding Strata

Species

Column Holding Variable to Split On

Sepal.Length

Column Holding Variable of Interest

Sepal.Width

Name of Strata to Split

setosa

Split Type

☐ global quantile

☒ local quantile

☐ value

☐ categorical

Split At

0 0.1 0.2 0.3 0.4 0.5 0.6 0.7 0.8 0.9 1

0.6

Include Information from Previous Wave

☒ optimum_allocation

☐ allocate_wave (have some units already been sampled?)

n to sample

30

Confirm Split Reset

Figure 3: User inputs in the **shiny** application.

	strata	npop	sd	n_sd	stratum_fraction	stratum_size
1	setosa.Sepal.Length_(5.1,5.8]	14	0.3215	4.5008	0.1	3
2	setosa.Sepal.Length_[4.3,5.1]	36	0.3074	11.0667	0.2333	7
3	versicolor	50	0.3138	15.6899	0.3333	10
4	virginica	50	0.3225	16.1248	0.3333	10

Showing 1 to 4 of 4 entries

Previous 1 Next

Figure 4: Reactive table output in the **shiny** application.

```
[1] "split_strata(data = 'df_name', strata = 'Species', split = 'setosa', split_var = 'Sepal.Length', type = 'local
quantile', split_at = '0.6')"
```

Figure 5: R code output by the **shiny** application after a split is confirmed.

In the example shown in Figure 4, we use the `iris` dataset again, splitting the species `setosa` at the within-stratum 60th percentile of sepal length. The data frame shown in the image below is the design data frame produced by `optimum_allocation()`. As we adjust the inputs, this data frame reactively updates to display the within-stratum standard deviations, population sizes, and optimal stratum sample size for our specified parameters. In this way, we can easily observe the impact of different stratum definitions on our sampling design:

When we are satisfied with the split parameters, we can select “Confirm Split” to run `split_strata()` with the specified arguments. The displayed data frame will then update to include the new strata, and the code to perform the split in R will appear below it (Figure 5). Each iteration of this process will add another line of code to the list:

After we have confirmed our desired strata definitions, we can leave the app and perform the splits in our R workflow by copying and pasting the text output of the app. The only change that needs to be added to the pasted code is an updated data frame name.

5. The multiwave object

Section 3 demonstrated how the base functions of **optimall** can be used to efficiently determine optimum allocation, select samples, and split strata during the design of a multi-wave stratified sampling survey. Despite these features, an efficient sampling workflow in R still requires the user to manually organize the many moving parts of the process including design data, a list of samples, data extracted from the samples, and merged data for each wave. When a sampling design involves many waves, these parts can be difficult to keep track of and may be prone to errors. Even more, it may be difficult to go back and reproduce results at the end of a long sampling process.

This section describes an additional feature of **optimall** called the multiwave object, which stores the metadata, design, samples, and merged data from each step of the multi-wave sampling process in an accessible format. It is optional for the user to use, but it offers the advantages of automatic organization, efficient compatibility with other functions in **optimall**, and an option for a summary of the sampling design to be printed at any point. The multiwave object contributes towards **optimall**'s goal of being a tool to streamline the often cumbersome aspects of the multi-wave sampling workflow.

5.1. Format

The multiwave object uses the S4 class system to store data from an entire multi-wave sampling survey. Its structure is illustrated in Figure 6.

The multiwave object contains three S4 classes:

- ‘Multiwave’: The ‘Multiwave’ class (dark blue in Figure 6) contains the metadata and

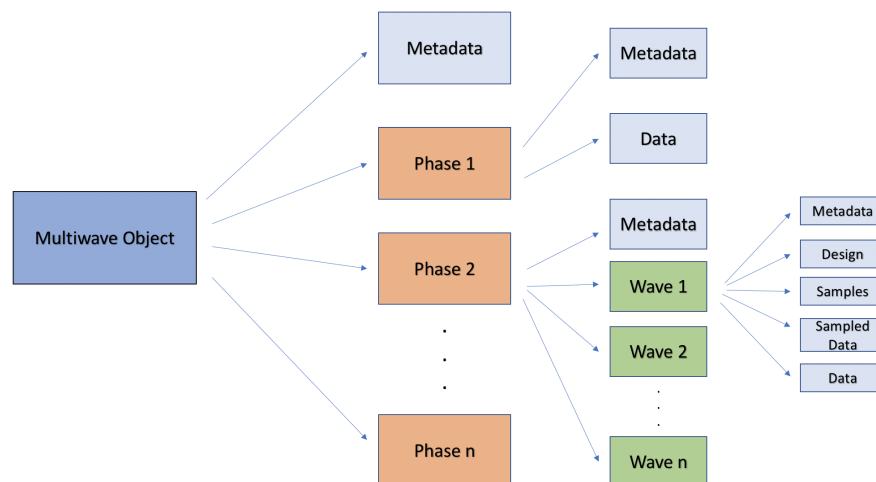


Figure 6: Multiwave object structure.

a list of Phase objects. It holds the entire sampling design and is the class that the user will interact with the most.

- ‘Phase’: The ‘Phase’ class (orange in Figure 6) holds the metadata and a list of Wave objects for each phase.
- ‘Wave’: The ‘Wave’ class (green in Figure 6) holds the metadata, design, samples, sampled data, and data for a single wave in multi-wave sampling.

The light blue objects in Figure 6 represent the slots that hold the survey information directly:

- Metadata: The metadata slot holds an initially empty list. Relevant information can be added as named elements, including titles, data dictionaries, and arguments for **optimall** functions.
- Design: The design slot holds the data frame specifying the sampling design of the current wave. In order for it to be used with `sample_strata()`, the design data frame must contain at least one column holding the strata names and one column holding the number of samples to be selected from each strata. It is typically the output of `optimum_allocation()` or `allocate_wave()`.
- Samples: The samples slot holds a list containing:
 - `ids`: A vector of the sample ids that were selected in a specific wave.
 - `probs`: Optionally, a numeric vector of sampling probabilities corresponding to each id in `ids`.
- Sampled Data: The sampled data slot contains a data frame holding the data collected in a specific wave.
- Data: The data slot contains a data frame with the full, accumulated data as it gets updated for each phase and wave. It is typically the sampled data merged with the (full) data of the previous wave. In this way, the data in the data slot of a study’s ultimate sampling wave is the full study data that can be used for analysis with the **survey** package.

5.2. Getting started with the multiwave object

To work with the multiwave object, users only need to know a few new functions. This section describes how to initialize an object, how to access and write slots of it, and how to deploy some of its useful features. It will generate examples from the `iris` dataset, again assuming that sepal width is the initially unknown variable of interest. More advanced examples based on real-life epidemiological studies will be presented in the following section.

A multiwave object is initialized by the function `multiwave()`:

```
R> MySurvey <- multiwave(phases = 2, waves = c(1, 3))
```

This code specifies that this survey will use a two-phase design with phase 2 divided into 3 sampling waves. Note that the length of the ‘waves’ argument must match the number of phases. Phase 1 will almost always have one wave.

5.3. Accessing and writing slots

As is standard for an S4 object in R, components of the multiwave object are stored in slots. The `get_mw()` function allows users to access these slots by specifying the phase, wave, and slot name. For example, the phase 2, wave 2 sample design (which in this case has yet to be determined) can be accessed with

```
R> get_mw(MySurvey, phase = 2, wave = 2, slot = "design")
```

```
Object of class "data.frame"
data frame with 0 columns and 0 rows
```

Calls to `get_mw()` for phase 1 do not require a wave to be specified, since phase 1 only consists of one wave:

```
R> get_mw(MySurvey, phase = 1, slot = "metadata")
```

```
list()
```

Note that metadata slots initialize as empty lists. The overall metadata is not phase nor wave specific, so neither needs to be specified in the call to `get_mw()`. For the overall metadata, it may be useful add a title. Slots can be written using `set_mw()`:

```
R> set_mw(MySurvey, phase = NA, slot = "metadata") <-  
+   list(title = "Sepal Width Survey")
```

```
R> get_mw(MySurvey, phase = NA, slot = "metadata")
```

```
$title  
[1] "Sepal Width Survey"
```

5.4. Call functions with fewer arguments

Another advantage of the multiwave object is that the primary functions of **optimall** including `optimum_allocation()`, `allocate_wave()`, and `sample_strata()` can be called on the object using the function `apply_multiwave()`. The `apply_multiwave()` function takes the standard arguments to the function as well as phase and wave, which are used to determine the input data frame(s) and the slot of the object where the output should be placed. If the arguments, including names of columns (which tend to be repetitive when used without the multiwave object framework), are specified in the metadata, the function will find them itself, allowing calls to the function to be much simpler.

To demonstrate, we return to the example of sampling sepal width from the `iris` dataset with an adaptive, multi-wave design. Suppose that we have collected data on sepal length

for all 150 iris plants in phase 1, but we have not yet collected any of the “expensive” sepal width variable. We have already placed the data in the corresponding slot of our multiwave object:

```
R> iris <- datasets::iris
R> iris$id <- 1:150
R> set_mw(MySurvey, phase = 1, slot = "data") <-
+   subset(iris, select = -Sepal.Width)
```

Now, we want to begin our first wave of sampling sepal width in phase 2. Since we expect that sepal length is correlated with our variable of interest, we decide to x-allocate the first wave of samples using the Wright II algorithm to determine the exact optimum allocation for a fixed sample size, using the already sampled sepal length variable as the design variable. Since we are working in the multiwave object framework, we can use `apply_multiwave()` to apply the `optimum_allocation()` function:

```
R> MySurvey <- apply_multiwave(MySurvey, phase = 2, wave = 1,
+   fun = "optimum_allocation", strata = "Species", y = "Sepal.Length",
+   nsample = 30, method = "WrightII")
```

Since "strata" will be "Species" for every wave, we may instead move that argument to the phase metadata so that we don't have to repetitively specify the same argument in every function call:

```
R> set_mw(MySurvey, phase = 2, slot = "metadata") <-
+   list(strata = "Species")
```

Now we no longer have to specify `strata` in the function call:

```
R> MySurvey <- apply_multiwave(MySurvey, phase = 2, wave = 1,
+   fun = "optimum_allocation", y = "Sepal.Length", nsample = 30,
+   method = "WrightII")
```

In the absence of a specific `strata` argument, `apply_multiwave()` turns to the wave, phase, and then overall metadata to look for the missing argument. In this case, it finds `strata = "Species"` in the phase metadata. By specifying the `phase` and `wave` in the function call, we are telling `optimum_allocation()` to use the full data from the prior wave as input and to output the results in the corresponding slot of the specified wave. As such, both calls to `apply_multiwave()` output an updated multiwave object with the results of `optimum_allocation()` in the phase 2, wave 1 design slot:

```
R> get_mw(MySurvey, phase = 2, wave = 1, slot = "design")
```

	strata	npop	sd	n_sd	stratum_fraction	stratum_size
1	setosa	50	0.35	17.62	0.23	7
2	versicolor	50	0.52	25.81	0.33	10
3	virginica	50	0.64	31.79	0.43	13

The `allocate_wave()` function can be applied to multiwave objects in the same manner and will be demonstrated in the following examples section. After these functions have been used to specify a “design” data frame (or a manually created design data frame has been placed in the design slot to implement a different allocation strategy), we can use `apply_multiwave()` to apply `sample_strata()` and select the ids to sample for the given sampling wave:

```
R> set.seed(340)
R> MySurvey <- apply_multiwave(MySurvey, phase = 2, wave = 1,
+   fun = "sample_strata", id = "id", design_strata = "strata",
+   n_allocated = "stratum_size", probs = ~ stratum_size/npop)
```

Note that we did not have to specify the `data` or `design_data` as we do in the standard version of `sample_strata()` because they are extracted using the `phase` and `wave` arguments. We also did not have to specify the `strata` argument again because it was available in the phase metadata. In this example, we specify the optional argument `probs` to attach the wave-specific inclusion probabilities to the selected samples. The result of this call to `apply_multiwave()` is an updated `MySurvey` with a vector of ids to sample and their inclusion probabilities in the `samples` slot:

```
R> get_mw(MySurvey, phase = 2, wave = 1, slot = "samples")

$ids
 [1] 1 4 22 26 28 30 39 51 53 66 70 73 79 80 88 95 99
[18] 101 111 119 120 121 130 138 139 140 143 144 148 150

$probs
 [1] 0.14 0.14 0.14 0.14 0.14 0.14 0.14 0.14 0.20 0.20 0.20 0.20 0.20 0.20
[14] 0.20 0.20 0.20 0.20 0.26 0.26 0.26 0.26 0.26 0.26 0.26 0.26 0.26
[27] 0.26 0.26 0.26 0.26
```

When working with a multiwave object, a new function called `merge_samples()` also becomes available. This function allows users to quickly and efficiently merge the sampled data with the accumulated data from the previous phases and waves. Suppose that we have collected the sepal width for these 30 plant ids and placed the data in the sampled data slot of phase 2, wave 1 of `MySurvey`:

```
R> set_mw(MySurvey, phase = 2, wave = 1, slot = "sampled_data") <-
+   iris[iris$id %in% get_mw(MySurvey, phase = 2, wave = 1,
+   slot = "samples")$ids, c("id", "Sepal.Width")]
```

We can call `merge_samples()` to smoothly merge the sampled data of the current wave with the previously accumulated data (in this case, only the phase 1 data) and place the output in the data slot of the current wave:

```
R> MySurvey <- merge_samples(MySurvey, phase = 2, wave = 1, id = "id",
+   include_probs = TRUE)
```

In the data slot of phase 2, wave 1, we will now have an updated dataframe with all of the phase 1 data and a sepal width column that has the sampled data for the selected ids and NA values for the rest of the plants. There is also a new column called "sampled_phase2" that holds an indicator for which samples have been sampled in phase 2 thus far, a new column called "sampled_wave2.1" that holds an indicator for which samples were sampled in wave 1 of phase 2, and a column of inclusion probabilities because we specified `include_probs = TRUE`:

```
R> head(get_mw(MySurvey, phase = 2, wave = 1, slot = "data"))
```

	id	Species	Sepal.Length	Sepal.Width	sampled_phase2	sampled_wave2.1
1	1	setosa	5.1	3.5	1	1
2	2	setosa	4.9	NA	0	0
3	3	setosa	4.7	NA	0	0
4	4	setosa	4.6	3.1	1	1
5	5	setosa	5.0	NA	0	0
6	6	setosa	5.4	NA	0	0

	sampling_prob
1	0.14
2	NA
3	NA
4	0.14
5	NA
6	NA

Calls to `merge_samples()` in later waves of phase 2 will update the "sampled_phase2" column each time. The `sample_strata()` function will use this column to ensure that units sampled in previous waves are not selected again in the same phase. The utility of `merge_samples()` and `apply_multiwave()` are demonstrated further in Section 6.

5.5. View summary diagram of survey

At any point during the multi-wave sampling workflow, **optimall** allows users to view a diagram of the structure of their survey with `multiwave_diagram()`. This function produces a graphic illustrating the structure using the **DiagrammeR** package by [Iannone and Roy \(2024\)](#):

```
R> multiwave_diagram(MySurvey)
```

Figure 7 shows the output produced by this function call. We see that the title of the survey, "Sepal Width Survey", is found from the overall survey metadata and that boxes are colored depending on whether they have been filled yet. Filled slots are blue and contain a short description of their contents. This function enables users to track their progress during a multi-wave sampling survey. A simpler version of this figure can also be printed to the console using `summary()`.

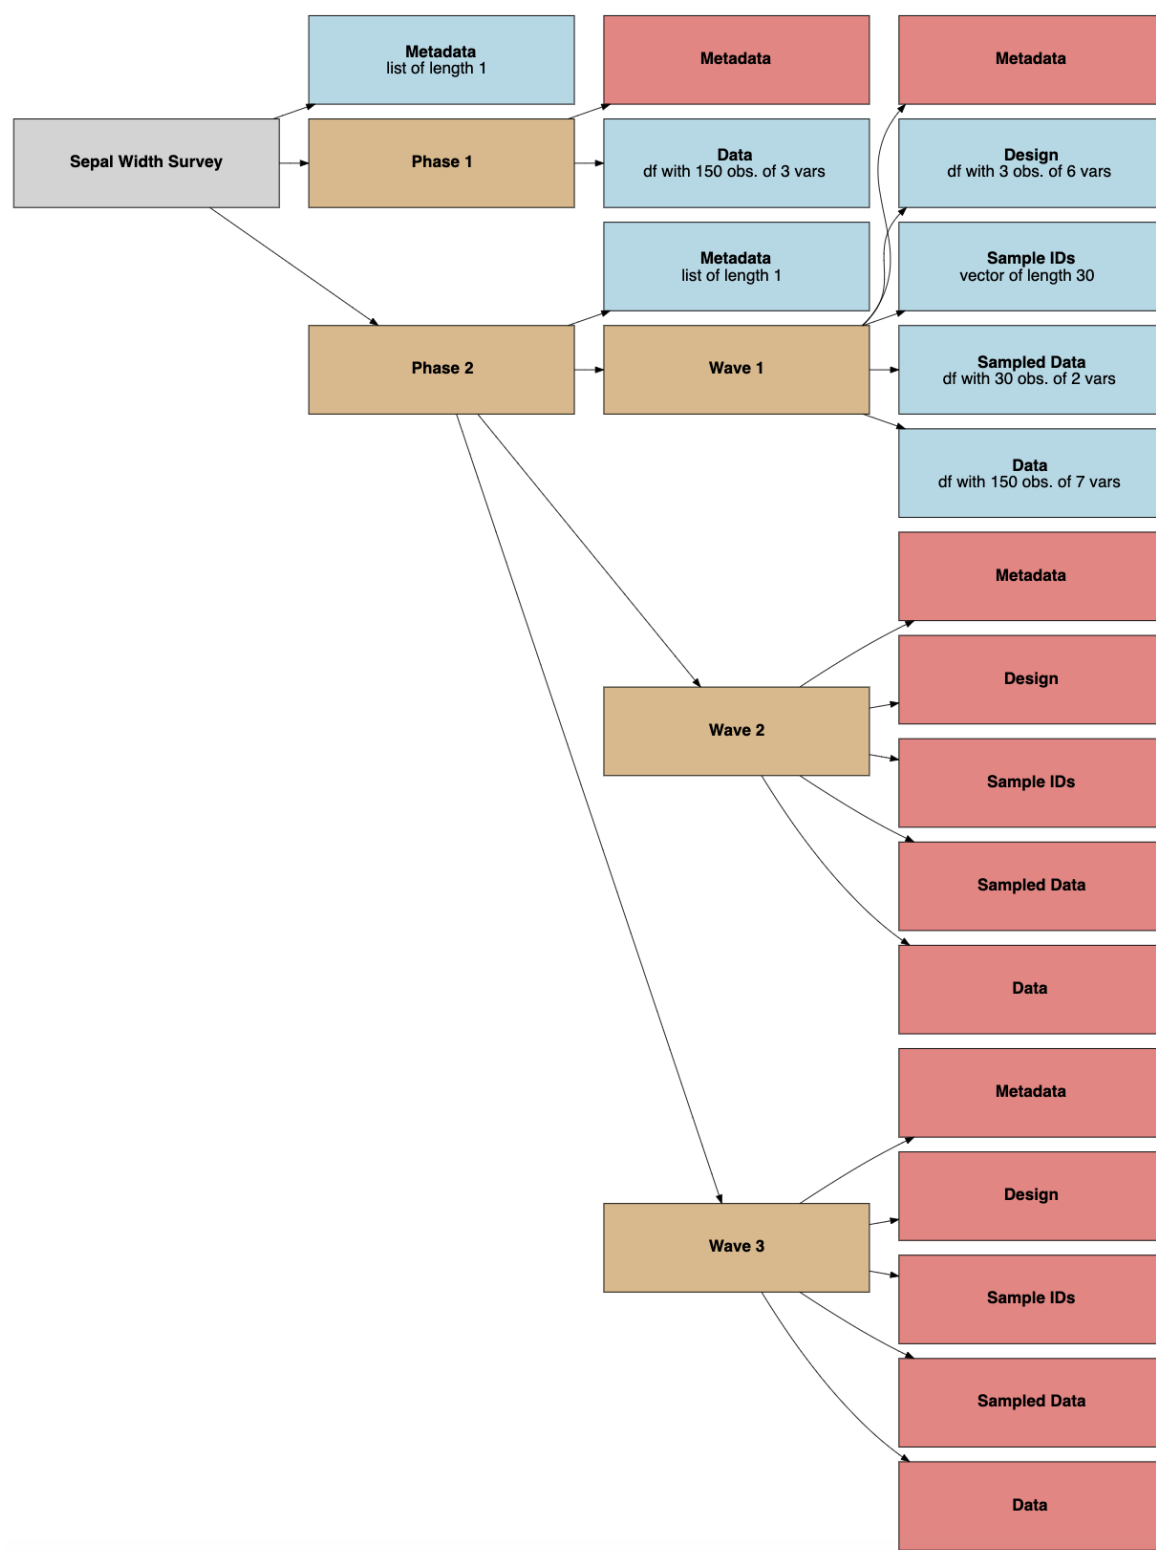


Figure 7: Multiwave diagram for Sepal Width Survey.

6. Examples

6.1. Introduction and data set up

In this section, we use the simulated dataset `MatWgt_Sim` to demonstrate how the functions of **optimall** can be used together to carry out the design of a sampling survey in R. This dataset is based on a real epidemiological study and includes measures of maternal weight gain during pregnancy, diabetes status, and race (Shepherd *et al.* 2023). The study illustrates another example of when multi-phase, multi-wave sampling is useful: data validation. Suppose that we have access to an inexpensive but error-prone measure of maternal weight change during pregnancy, but we are interested in the true weight change that can only be obtained through expensive data validation. In phase 1, we collect a simple random sample of error-prone maternal weight change, race, and diabetes status on all 10,335 patients in our target population. In phase 2, we can afford to validate 750 samples. We can use the adaptive, multi-wave workflow to efficiently allocate these validations.

The raw dataset is available in **optimall** and is organized as follows:

```
R> head(MatWgt_Sim)
```

	id	mat_weight_est	race	diabetes	obesity	mat_weight_true
1	7667	12.954711	Asian	0	0	10.768935
2	8554	9.607124	Asian	0	1	8.892049
3	6324	10.998284	Asian	0	0	12.075861
4	8320	10.702210	Asian	0	0	10.887584
5	8543	6.485988	Asian	1	0	7.331542
6	127	13.338445	Asian	0	1	12.936826

In our example, `mat_weight_true` is expensive to collect and unknown initially, so our phase 1 data excludes it:

```
R> phase1 <- subset(MatWgt_Sim, select = -mat_weight_true)
```

The following subsections demonstrate how **optimall** facilitates the sampling of the true maternal weight change under different conditions. We begin with a simple example before building up to more complex multi-wave designs.

The estimates presented throughout this section are constructed using post-stratification, which is an efficient analysis method where the design weights are calculated from the final realized allocation of samples across strata (Valliant 1993; Lumley, Shaw, and Dai 2011). Notably, post-stratification does not require attachment of the wave-specific sampling probabilities to each element, and it can be easily implemented using the **Survey** package. The **optimall** package vignettes include more details on the rationale behind this choice of estimator, as well as examples with other reasonable estimator choices.

6.2. Estimating the mean of an expensive covariate in a single wave

Suppose that we are interested in estimating the population mean of the true maternal weight change with minimal variance, but we have to allocate all 750 validation samples in a single

wave. Stratified random sampling will be useful for efficiency in this scenario, and we can use the error-prone maternal weight change measure to approximate the optimum allocation for the true measure. To implement this in **optimall**, we first define the strata. We will split the population into non-overlapping strata using race and global percentiles of maternal weight gain. First, we initialize a strata column and combine the two smallest categories into a single stratum:

```
R> phase1$strata <- phase1$race
R> phase1 <- merge_strata(data = phase1, strata = "strata",
+   merge = c("Other", "Asian"), name = "Other")
```

Recall that **optimall** places the updated strata defined by `merge_strata()` and `split_strata()` in a new column called "new_strata". We next want to split this new column at the 25th and 75th percentiles of estimated maternal weight change:

```
R> phase1 <- split_strata(data = phase1, strata = "new_strata", split = NULL,
+   split_var = "mat_weight_est", type = "global quantile",
+   split_at = c(0.25, 0.75), trunc = "MWC_est")
```

Which gives us our final 9 strata:

```
R> table(phase1$new_strata)
```

Black.MWC_est_(15.06,38.46]	Black.MWC_est_(9.75,15.06]
628	1154
Black.MWC_est_[-30.21,9.75]	Other.MWC_est_(15.06,30.94]
745	325
Other.MWC_est_(9.75,15.06]	Other.MWC_est_[-5.39,9.75]
929	456
White.MWC_est_(15.06,51.69]	White.MWC_est_(9.75,15.06]
1631	3084
White.MWC_est_[-25.68,9.75]	
1383	

With the strata and phase 1 data defined, we can initialize a multiwave object to store the information from this survey:

```
R> Survey1 <- multiwave(phases = 2, waves = c(1,1))
R> set_mw(Survey1, phase = NA, slot = "metadata") <-
+   list(title = "Single Wave Mat Wgt Survey")
R> set_mw(Survey1, phase = 1, slot = "data") <- phase1
```

We can now use `apply_multiwave()` to generate the design for the single wave in phase 2. We have not yet sampled any values of the true maternal weight change (our variable of interest), but we do have access to the error-prone maternal weight change variable from phase 1 for every patient in our study. The true and error-prone versions are likely highly correlated, so we can approximate the optimum allocation of samples by applying Neyman (or Wright) allocation to the error-prone version. This approach is often referred to as x-optimal allocation (Särndal *et al.* 2003).

```
R> Survey1 <- apply_multiwave(Survey1, phase = 2, wave = 1,
+   fun = "optimum_allocation", strata = "new_strata",
+   y = "mat_weight_est", nsample = 750, method = "WrightII")
R> get_mw(Survey1, phase = 2, wave = 1, slot = "design")
```

	strata	npop	sd	n_sd	stratum_fraction	stratum_size
1	Black.MWC_est_(15.06,38.46]	628	3.63	2277.53	0.09	68
2	Black.MWC_est_(9.75,15.06]	1154	1.46	1688.46	0.07	51
3	Black.MWC_est_[-30.21,9.75]	745	4.09	3050.57	0.12	92
4	Other.MWC_est_(15.06,30.94]	325	2.41	781.93	0.03	24
5	Other.MWC_est_(9.75,15.06]	929	1.44	1334.35	0.05	40
6	Other.MWC_est_[-5.39,9.75]	456	2.46	1121.03	0.05	34
7	White.MWC_est_(15.06,51.69]	1631	3.51	5726.61	0.23	172
8	White.MWC_est_(9.75,15.06]	3084	1.44	4453.92	0.18	134
9	White.MWC_est_[-25.68,9.75]	1383	3.24	4485.44	0.18	135

We can then use this design and `sample_strata()`, implemented through the `apply_multiwave()` function, to select the specific ids to sample:

```
R> set.seed(435)
R> Survey1 <- apply_multiwave(Survey1, phase = 2, wave = 1,
+   fun = "sample_strata", strata = "new_strata", id = "id",
+   design_strata = "strata", n_allocated = "stratum_size")
```

We can simulate the data validation process by selecting the true maternal weight changes for these ids and placing the resulting data frame in the sampled data slot:

```
R> set_mw(Survey1, phase = 2, wave = 1, slot = "sampled_data") <-
+   MatWgt_Sim[MatWgt_Sim$id %in% get_mw(Survey1, phase = 2,
+   wave = 1, slot = "samples")$ids, c("id", "mat_weight_true")]
```

Finally, we can merge the sampled data with the phase 1 data using `merge_samples()`:

```
R> Survey1 <- merge_samples(Survey1, phase = 2, wave = 1, id = "id")
```

We have now completed validation of 750 samples! The data slot of the final wave (phase 2, wave 1 in this case) holds a data frame that can easily be used to generate estimates with the `survey` (Lumley 2010) package:

```
R> library("survey")
```

First, we extract the final data from our survey:

```
R> final_data <- get_mw(Survey1, phase = 2, wave = 1, slot = "data")
```

Then, we specify the two-phase design and use `svymean` to estimate the mean.

```
R> twophasedesign <- twophase(id = list(~ 1, ~ 1),
+   strata = list(NULL, ~ new_strata),
+   subset = ~ as.logical(sampled_phase2), data = final_data,
+   method = "approx")
R> survey::svymean(~mat_weight_true, design = twophasedesign)

              mean      SE
mat_weight_true 12.363 0.1003
```

Note that phase 1 is a simple random sample in this case, hence the phase 1 strata are NULL. More complex phase 1 designs can also be accommodated with **survey** and are equally compatible with the multiwave framework.

6.3. Estimating the mean of an expensive covariate over multiple waves

For this next example, we again suppose that we can afford to validate the true maternal weight change of 750 patients, but this time we will do it over a series of three adaptive waves. The phase 1 data is unchanged from the previous example, so initializing a multiwave object to capture this survey is straightforward:

```
R> Survey2 <- multiwave(phases = 2, waves = c(1, 3))
R> set_mw(Survey2, phase = NA, slot = "metadata") <-
+   list(title = "Multi-wave Mat Wgt Survey")
R> set_mw(Survey2, phase = 1, slot = "data") <- phase1
R> set_mw(Survey2, phase = 1, slot = "metadata") <- list(
+   title = "Phase 1", description = "This is simulated data
+   generated in December 2020 based on a real maternal weight study",
+   data_dict = data.frame(id = "unique identifier",
+   new_strata = "Name of updated strata",
+   race = "self-identified race of mother",
+   mat_weight_est = "error-prone estimate of
+   maternal weight change in lbs. during pregnancy",
+   obesity = "1/0 indicator for child obesity within 6 years of birth",
+   diabetes = "1/0 indicator for mother's diabetes status"))
```

We will now demonstrate the second phase of sampling by breaking it into waves.

Phase 2, wave 1

In the first wave of phase 2, we will allocate 250 samples using x-allocation with the error-prone phase 1 maternal weight change estimates as the design variable:

```
R> Survey2 <- apply_multiwave(Survey2, phase = 2, wave = 1,
+   fun = "optimum_allocation", strata = "new_strata",
+   y = "mat_weight_est", nsample = 250, method = "Neyman")
R> get_mw(Survey2, phase = 2, wave = 1, slot = "design")
```

	strata	npop	sd	n_sd	stratum_fraction	stratum_size
1	Black.MWC_est_(15.06,38.46]	628	3.63	2277.53	0.09	23
2	Black.MWC_est_(9.75,15.06]	1154	1.46	1688.46	0.07	17
3	Black.MWC_est_[-30.21,9.75]	745	4.09	3050.57	0.12	31
4	Other.MWC_est_(15.06,30.94]	325	2.41	781.93	0.03	8
5	Other.MWC_est_(9.75,15.06]	929	1.44	1334.35	0.05	13
6	Other.MWC_est_[-5.39,9.75]	456	2.46	1121.03	0.04	11
7	White.MWC_est_(15.06,51.69]	1631	3.51	5726.61	0.23	57
8	White.MWC_est_(9.75,15.06]	3084	1.44	4453.92	0.18	45
9	White.MWC_est_[-25.68,9.75]	1383	3.24	4485.44	0.18	45

Before we begin the sampling process, we can specify function arguments and other details in the metadata of phase 2:

```
R> set_mw(Survey2, phase = 2, slot = "metadata") <- list(
+   title = "Phase 2", description = "This phase will sample 750
+   patients across three waves of 250. The first wave will use
+   proportional allocation, and the next two will use Wright
+   Algorithm II based on the previous wave's samples",
+   strata = "new_strata", design_strata = "strata", id = "id",
+   y = "mat_weight_true")
```

We can now select the 250 wave 1 ids with `sample_strata()`:

```
R> set.seed(983)
R> Survey2 <- apply_multiwave(Survey2, phase = 2, wave = 1,
+   fun = "sample_strata", n_allocated = "stratum_size")
```

Then we can “collect” the sampled data:

```
R> set_mw(Survey2, phase = 2, wave = 1, slot = "sampled_data") <-
+   MatWgt_Sim[MatWgt_Sim$id %in% get_mw(Survey2, phase = 2, wave = 1,
+   slot = "samples")$ids, c("id", "mat_weight_true")]
```

and then merge it with the phase 1 data to get an updated dataset in the data slot using `merge_samples()`:

```
R> Survey2 <- merge_samples(Survey2, phase = 2, wave = 1)
```

With that, wave 1 is complete!

```
R> dim(get_mw(Survey2, phase = 2, wave = 1, slot = "data"))
```

```
[1] 10335    10
```

```
R> table(is.na(get_mw(Survey2, phase = 2, wave = 1, "data")$mat_weight_true))
```

```
FALSE  TRUE
 250 10085
```


Phase 2, wave 2

For the next wave, we essentially repeat the process of wave 1. This time, though, we use the 250 samples from wave 1 to estimate the optimal design with `allocate_wave()`. Recall from Section 3.4 that `allocate_wave()` will output a data frame holding the optimum allocation based on the previous samples of `y` using Wright algorithm II. When used within `apply_multiwave`, the output design data frame will be automatically placed in the design slot of the specified wave:

```
R> Survey2 <- apply_multiwave(Survey2, phase = 2, wave = 2,
+   fun = "allocate_wave", already_sampled = "sampled_phase2",
+   nsample = 250)
```

Looking at the design, we see that the proportional allocation in wave 1 did not appear to oversample from any strata:

```
R> get_mw(Survey2, phase = 2, wave = 2, slot = "design")
```

	strata	npop	nsample_optimal	nsample_actual
1	Black.MWC_est_(15.06,38.46]	628	37	37
2	Black.MWC_est_(9.75,15.06]	1154	40	40
3	Black.MWC_est_[-30.21,9.75]	745	69	69
4	Other.MWC_est_(15.06,30.94]	325	14	14
5	Other.MWC_est_(9.75,15.06]	929	27	27
6	Other.MWC_est_[-5.39,9.75]	456	20	20
7	White.MWC_est_(15.06,51.69]	1631	121	121
8	White.MWC_est_(9.75,15.06]	3084	99	99
9	White.MWC_est_[-25.68,9.75]	1383	73	73

	nsample_prior	n_to_sample	sd
1	23	14	2.97
2	17	23	1.79
3	31	38	4.74
4	8	6	2.18
5	13	14	1.47
6	11	9	2.30
7	57	64	3.80
8	45	54	1.64
9	45	28	2.70

We use this design to inform the selection of ids to sample:

```
R> set.seed(584)
R> Survey2 <- apply_multiwave(Survey2, phase = 2, wave = 2,
+   fun = "sample_strata", already_sampled = "sampled_phase2",
+   design_strata = "strata", n_allocated = "n_to_sample")
```

then we “collect” the data from these samples:

```
R> set_mw(Survey2, phase = 2, wave = 2, slot = "sampled_data") <-
+   MatWgt_Sim[MatWgt_Sim$id %in% get_mw(Survey2, phase = 2, wave = 2,
+     slot = "samples")$ids, c("id", "mat_weight_true")]
```

and merge it back with the data from the previous wave to get the updated data in the data slot of wave 2:

```
R> Survey2 <- merge_samples(Survey2, phase = 2, wave = 2)
```

With that, wave 2 is complete. The updated data frame now has columns indicating which elements have been sampled in Phase 2, as well all columns indicating the specific wave in which each element was sampled.

```
R> names(get_mw(Survey2, phase = 2, wave = 2, slot = "data"))

[1] "new_strata"      "old_strata"      "id"              "mat_weight_est"
[5] "race"           "diabetes"        "obesity"         "sampled_phase2"
[9] "sampled_wave2.1" "mat_weight_true" "sampled_wave2.2"

R> table(get_mw(Survey2, phase = 2, wave = 2, "data")$sampled_phase2)

    0     1
9835  500
```

Phase 2, wave 3

We proceed with wave 3 in the exact same manner as in wave 2. This subsection leaves out some of the specific details covered in the previous waves for brevity.

Allocate the samples:

```
R> Survey2 <- apply_multiwave(Survey2, phase = 2, wave = 3,
+   fun = "allocate_wave", already_sampled = "sampled_phase2",
+   nsample = 250, detailed = FALSE)
```

Select the samples:

```
R> set.seed(584)
R> Survey2 <- apply_multiwave(Survey2, phase = 2, wave = 3,
+   fun = "sample_strata", already_sampled = "sampled_phase2",
+   design_strata = "strata", n_allocated = "n_to_sample")
```

“Collect” the data:

```
R> set_mw(Survey2, phase = 2, wave = 3, slot = "sampled_data") <-
+   MatWgt_Sim[MatWgt_Sim$id %in% get_mw(Survey2, phase = 2, wave = 3,
+     slot = "samples")$ids, c("id", "mat_weight_true")]
```

Merge the sampled data:

```
R> Survey2 <- merge_samples(Survey2, phase = 2, wave = 3)
```

We have now collected all 750 samples and can apply the same estimation procedure as in the single-wave example of Section 6.3 to compute a post-stratified estimate for the mean in this multi-wave case. This method calculates sampling weights using the realized sample allocation in each stratum (Holt and Smith 1979; Valliant 1993).

6.4. A multi-wave sampling design for regression modeling

Suppose that we have the same phase 1 dataset for the maternal weight change study, but we are now interested in modeling the relationship between the odds of childhood obesity and true maternal weight change during pregnancy, with diabetes and race as covariates. We will consider childhood obesity as the binary outcome variable and regress it on the covariates using logistic regression. This problem closely resembles those described in Chen and Lumley (2020) and McIsaac and Cook (2015). In this example, we will conduct the sampling over two waves, with 250 samples in the first wave and 500 in the second:

```
R> Survey3 <- multiwave(phases = 2, waves = c(1,2))
R> set_mw(Survey3, phase = NA, slot = "metadata") <-
+   list(title = "Survey for Regression of Childhood Obesity")
R> set_mw(Survey3, phase = 1, slot = "data") <- phase1
```

Phase 2, wave 1

Optimum allocation for a regression covariate is slightly different from the simple mean estimation from the previous examples. In the two-phase sampling setting, estimation of a weighted regression model from the complete case phase 2 data may become more efficient by making use of auxiliary phase 1 variables. The technique, known as survey calibration or raking, uses an auxiliary variable available on the whole cohort to adjust (calibrate) the usual Horwitz-Thompson weights (Deville and Särndal 1992). This approach can achieve the optimum minimum variance in the class of “augmented” inverse probability weighted estimates in the case of the ideal auxiliary variable, which is the expected value of the target parameter’s influence function given the observed data (Lumley *et al.* 2011; Robins, Rotnitzky, and Zhao 1994). A list of steps to implement this method in practice was provided by Breslow, Lumley, Ballantyne, Chambless, and Kulich (2009a). This approach, which was first described in Kulich and Lin (2004), involves approximating the ideal raking variables by first imputing the unobserved phase 2 variables on the whole cohort and then obtaining the influence function for all model parameters from the model fit with the imputed data. This approach was used in the setting of case-cohort designs in Breslow, Lumley, Ballantyne, Chambless, and Kulich (2009b) and in case-matching in Rivera and Lumley (2016). Others suggest repeatedly estimating this influence function with multiple imputation to estimate the expected value of the influence function given the observed data (Han, Shaw, and Lumley 2021b; Han 2016). If correctly modeled, this approach will achieve maximal efficiency.

For wave 1 of this example, we employ a simplified version of the Kulich and Lin approach, which uses the influence function from the naive logistic regression model fit with the inexpensive, error-prone phase 1 data for maternal weight to approximate the desired optimal design to estimate this regression coefficient. We can use `apply_multiwave()` to apply `optimum_allocation()` to the estimated influence functions to allocate the first 250 samples:

```

R> fit <- glm(obesity ~ mat_weight_est + diabetes + race,
+   data = get_mw(Survey3, phase = 1, slot = "data"),
+   family = "quasibinomial")
R> design_mat <- model.matrix(fit)
R> I_hat <- (t(design_mat) %*% (design_mat * fit$fitted.values *
+   (1 - fit$fitted.values))) / nrow(design_mat)
R> influence <- (design_mat * resid(fit, type = "response")) %*%
+   solve(I_hat)
R> influence <- as.data.frame(influence)

```

We can then add the influence functions as a new column in our phase 1 data:

```

R> set_mw(Survey3, phase = 1, slot = "data") <- cbind(
+   get_mw(Survey3, phase = 1, slot = "data"),
+   naive_influence = influence[, "mat_weight_est"])

```

Note that we used the Fisher information matrix and the generalized linear model residuals to estimate the influence function for our parameter of interest. As discussed by [Breslow *et al.* \(2009b\)](#), influence functions of the Cox model can be readily approximated with the delta-beta residuals in the Cox model fit using the `coxph()` function in the **survey** R package ([Therneau 2024](#)). We can now generate the wave 1 design with `optimum_allocation`:

```

R> Survey3 <- apply_multiwave(Survey3, phase = 2, wave = 1,
+   fun = "optimum_allocation", strata = "new_strata",
+   y = "naive_influence", nsample = 250, method = "WrightII")

```

The rest of Phase 2 proceeds as demonstrated in Section 6.3, only now with the influence function as the variable of interest. We thus leave out the details describing the process in **optimall** to avoid repetition.

7. Summary

This article has demonstrated the utility of the **optimall** package in stratified surveys ranging from simple to complex. For simple stratified surveys, `split_strata()`, `merge_strata()`, and the **shiny** app launched by `optimall_shiny()` are straightforward tools for making and implementing stratification decisions, while `sample_strata()` is useful for sample selection when simple random sampling is used within strata. For complex multi-phase and multi-wave survey designs, the multiwave object framework streamlines the previously cumbersome implementation process. The functions `optimum_allocation()` and `allocate_wave()` facilitate Neyman and Wright allocation, which are broadly applicable, but the framework supports the manual implementation of other allocation strategies. Further details on **optimall**'s features and guidance on analytical approaches following a multi-wave design are provided in the package vignettes.

The overarching goal of **optimall** is to facilitate a streamlined design process for stratified sampling surveys in R, including those that employ multi-phase and multi-wave designs. While the current version closely aligns with this goal, user feedback will help guide improvements over time. We encourage these comments to be provided at the package's GitHub page

(<https://github.com/yangjasp/optimall>). For now, future directions include adding functions to implement other methods for optimum allocation and addressing cases where sampling cost is to be minimized for a given precision level for estimation.

Acknowledgments

We would like to thank Gustavo Amorim, PhD for providing helpful comments on the paper and package vignettes. This work was supported in part by the U.S. National Institutes of Health (NIH) grant R01-AI131771 and Patient Centered Outcomes Research Institute (PCORI) Award R-1609-36207. The statements in this manuscript are solely the responsibility of the authors and do not necessarily represent the views of PCORI or NIH.

References

- Amorim G, Tao R, Lotspeich S, Shaw PA, Lumley T, Shepherd BE (2021). “Two-Phase Sampling Designs for Data Validation in Settings with Covariate Measurement Error and Continuous Outcome.” *Journal of the Royal Statistical Society A*, **184**(4), 1368–1389. doi:[10.1111/rssa.12689](https://doi.org/10.1111/rssa.12689).
- Barcaroli G, Ballin M, Odendaal H, Pagliuca D, Willighagen E, Zardetto D (2022). **SamplingStrata**: *Optimal Stratification of Sampling Frames for Multipurpose Sampling Surveys*. R package version 1.5-4, URL <https://CRAN.R-project.org/package=SamplingStrata>.
- Breidaks J, Liberts M, Jukams J (2020). **surveyplanning**: *Survey Planning Tools*. Riga, Latvia. R package version 4.0, URL <https://csblatvia.github.io/surveyplanning/>.
- Breslow NE, Lumley T, Ballantyne CM, Chambless LE, Kulich M (2009a). “Improved Horvitz-Thompson Estimation of Model Parameters from Two-Phase Stratified Samples: Applications in Epidemiology.” *Statistics in Biosciences*, **1**(1), 32–49. doi:[10.1007/s12561-009-9001-6](https://doi.org/10.1007/s12561-009-9001-6).
- Breslow NE, Lumley T, Ballantyne CM, Chambless LE, Kulich M (2009b). “Using the Whole Cohort in the Analysis of Case-Cohort Data.” *American Journal of Epidemiology*, **169**(11), 1398–1405. doi:[10.1007/s12561-013-9080-2](https://doi.org/10.1007/s12561-013-9080-2).
- Chang W, Cheng J, Allaire JJ, Sievert C, Schloerke B, Xie Y, Allen J, McPherson J, Dipert A, Borges B (2025). **shiny**: *Web Application Framework for R*. doi:[10.32614/CRAN.package.shiny](https://doi.org/10.32614/CRAN.package.shiny). R package version 1.11.1.
- Chen T, Lumley T (2020). “Optimal Multi-Wave Sampling for Regression Modelling in Two-Phase Designs.” *Statistics in Medicine*, **39**(30), 4912–4921. doi:[10.1002/sim.8760](https://doi.org/10.1002/sim.8760).
- Deville JC, Särndal CE (1992). “Calibration Estimators in Survey Sampling.” *Journal of the American Statistical Association*, **87**(418), 376–382. doi:[10.2307/2290268](https://doi.org/10.2307/2290268).
- Han K, Lumley T, Shepherd BE, Shaw PA (2021a). “Two-Phase Analysis and Study Design for Survival Models with Error-Prone Exposures.” *Statistical Methods in Medical Research*, **30**(3), 857–874. doi:[10.1177/0962280220978500](https://doi.org/10.1177/0962280220978500).

- Han K, Shaw PA, Lumley T (2021b). “Combining Multiple Imputation with Raking of Weights: An Efficient and Robust Approach in the Setting of Nearly-True Models.” *Statistics in Medicine*, **40**(30), 6777–6791. doi:[10.1002/sim.9210](https://doi.org/10.1002/sim.9210).
- Han P (2016). “Combining Inverse Probability Weighting and Multiple Imputation to Improve Robustness of Estimation.” *Scandinavian Journal of Statistics*, **43**(1), 246–260. doi:[10.1111/sjos.12177](https://doi.org/10.1111/sjos.12177).
- Holt D, Smith TMF (1979). “Post Stratification.” *Journal of the Royal Statistical Society A*, **142**(1), 33–46. doi:[10.2307/2344652](https://doi.org/10.2307/2344652).
- Iannone R, Roy O (2024). **DiagrammeR**: *Graph/Network Visualization*. doi:[10.32614/CRAN.package.DiagrammeR](https://doi.org/10.32614/CRAN.package.DiagrammeR). R package version 1.0.11.
- Kulich M, Lin DY (2004). “Improving the Efficiency of Relative-Risk Estimation in Case-Cohort Studies.” *Journal of the American Statistical Association*, **99**(467), 832–844. doi:[10.1198/016214504000000584](https://doi.org/10.1198/016214504000000584).
- Lumley T (2010). *Complex Surveys: A Guide to Analysis Using R*. John Wiley & Sons.
- Lumley T, Shaw PA, Dai JY (2011). “Connections between Survey Calibration Estimators and Semiparametric Models for Incomplete Data.” *International Statistical Review*, **79**(2), 200–220. doi:[10.1111/j.1751-5823.2011.00138.x](https://doi.org/10.1111/j.1751-5823.2011.00138.x).
- MacQueen J (1967). “Some Methods for Classification and Analysis of Multivariate Observations.” In LM Le Cam, J Neyman (eds.), *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, pp. 281–297.
- McIsaac MA, Cook RJ (2015). “Adaptive Sampling in Two-Phase Designs: A Biomarker Study for Progression in Arthritis.” *Statistics in Medicine*, **34**(21), 2899–2912. doi:[10.1002/sim.6523](https://doi.org/10.1002/sim.6523).
- Neyman J (1934). “On the Two Different Aspects of the Representative Method: The Method of Stratified Sampling and the Method of Purposive Selection.” *Journal of the Royal Statistical Society*, **97**(4), 558–625. doi:[10.2307/2342192](https://doi.org/10.2307/2342192).
- Neyman J (1938). “Contribution to the Theory of Sampling Human Populations.” *Journal of the American Statistical Association*, **33**(201), 101–116. doi:[10.2307/2279117](https://doi.org/10.2307/2279117).
- R Core Team (2025). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. doi:[10.32614/R.manuals](https://doi.org/10.32614/R.manuals). URL <https://www.R-project.org/>.
- Reddy KG, Khan MGM (2020). “**stratifyR**: An R Package for Optimal Stratification and Sample Allocation for Univariate Populations.” *Australian & New Zealand Journal of Statistics*, **62**(3), 383–405. doi:[10.1111/anzs.12301](https://doi.org/10.1111/anzs.12301).
- Rivera C, Lumley T (2016). “Using the Whole Cohort in the Analysis of Countermatched Samples.” *Biometrics*, **72**(2), 382–391. doi:[10.1111/biom.12419](https://doi.org/10.1111/biom.12419).
- Rivest LP, Baillargeon S (2022). **stratification**: *Univariate Stratification of Survey Populations*. R package version 2.2-7, URL <https://CRAN.R-project.org/package=stratification>.

- Robins JM, Rotnitzky A, Zhao LP (1994). “Estimation of Regression Coefficients When Some Regressors Are Not Always Observed.” *Journal of the American Statistical Association*, **89**(427), 846–866. doi:[10.1080/01621459.1994.10476818](https://doi.org/10.1080/01621459.1994.10476818).
- Särndal CE, Swensson B, Wretman J (2003). *Model Assisted Survey Sampling*. Springer-Verlag. doi:[10.1007/978-1-4612-4378-6](https://doi.org/10.1007/978-1-4612-4378-6).
- Shepherd BE, Han K, Chen T, Bian A, Pugh S, Duda S, Lumley T, Heerman WJ, Shaw PA (2023). “Multiwave Validation Sampling for Error-Prone Electronic Health Records.” *Biometrics*, **79**(3), 2649–2663. doi:[10.1111/biom.13713](https://doi.org/10.1111/biom.13713).
- Tao R, Zeng D, Lin DY (2020). “Optimal Designs of Two-Phase Studies.” *Journal of the American Statistical Association*, **115**(532), 1946–1959. doi:[10.1080/01621459.2019.1671200](https://doi.org/10.1080/01621459.2019.1671200).
- Therneau TM (2024). *survival: A Package for Survival Analysis in R*. R package version 3.8-3, URL [10.32614/CRAN.package.survival](https://CRAN.R-project.org/package=survival).
- Tillé Y, Matei A (2025). *sampling: Survey Sampling*. R package version 2.11, URL <https://CRAN.R-project.org/package=sampling>.
- Valliant R (1993). “Poststratification and Conditional Variance Estimation.” *Journal of the American Statistical Association*, **88**(421), 89–96. doi:[10.1080/01621459.1993.10594298](https://doi.org/10.1080/01621459.1993.10594298).
- Valliant R, Zipf G (2025). *PracTools: Tools for Designing and Weighting Survey Samples*. R package version 1.6.1, URL <https://CRAN.R-project.org/package=PracTools>.
- Wright T (2012). “The Equivalence of Neyman Optimum Allocation for Sampling and Equal Proportions for Apportioning the US House of Representatives.” *The American Statistician*, **66**(4), 217–224. doi:[10.1080/00031305.2012.733679](https://doi.org/10.1080/00031305.2012.733679).
- Yang J, Shaw P (2025). *optimall: Allocate Samples Among Strata*. doi:[10.32614/CRAN.package.optimall](https://doi.org/10.32614/CRAN.package.optimall). R package version 1.3.0.

Affiliation:

Jasper B. Yang
University of Washington
Department of Biostatistics
Seattle, WA 98195, United States of America
E-mail: jbyang@uw.edu