



Faculty of science
Institute of Statistics

Survey Sampling Methods Applied to Experiments

Thesis submitted in fulfillment of the requirements for the degree
of Doctorat ès Sciences at the University of Neuchâtel

by

Ejub Talovic

The dissertation Committee:

Prof. Yves Tillé	Thesis director	University of Neuchâtel
Prof. Alina Matei	President of jury	University of Neuchâtel
Prof. David Haziza	Examiner	University of Ottawa
Prof. Camélia Goga	Examiner	University of Bourgogne Franche-Comté
Dr. Caren Hasler	Examiner	University of Neuchâtel and University of Zürich

Thesis defended on 24 June 2025

IMPRIMATUR POUR THESE DE DOCTORAT

La Faculté des sciences de l'Université de Neuchâtel autorise
l'impression de la présente thèse soutenue par

Monsieur Ejub TALOVIC

Titre :

“Survey Sampling Methods applied to Experiments”

sur le rapport des membres du jury composé comme suit :

- **Prof. Yves Tillé**, directeur de thèse, Université de Neuchâtel, Suisse
- **Prof. tit. Alina Matei**, Université de Neuchâtel, Suisse
- **Prof. David Haziza**, University of Ottawa, Canada
- **Prof. Camélia Goga**, Université Bourgogne-Franche Comté, France
- **Dr Caren Hasler**, Université de Neuchâtel, Suisse

Neuchâtel, le 27 août 2025

Le Doyen, Prof. P. Brunner



Acknowledgements

I want to thank first and foremost my supervisor, Professor Yves Tillé. His expertise in survey sampling was irreplaceable. I am grateful for his patience, for pushing me to think deeper and for his guidance. He was also a very approachable and friendly mentor, which allowed me to have a nice work environment where I could freely ask questions during those three years together.

I also want to thank Professor Louis-Paul Rivest for his comments, guidance and sharp feedback. His contribution was crucial for the article corresponding to Chapter 4. I would also like to express my gratitude to the members of the jury, Professor Alina Matei, Professor David Haziza, Professor Camélia Goga, Dr. Caren Hasler who will read my thesis and surely improve its quality. And finally, thanks to my former and remaining colleagues, Arnaud, Esther, Raphaël and Ziqing at the institute for the brainstorming sessions, the help with the thesis, for the laughs and extremely deep discussions. They are the reason the why coffee breaks were so fun.

Abstract

This thesis studies different aspects of the design of surveys and randomized experiments. Those two subjects have particular similarities and differences that allow us to establish links between their theories. An important part of this thesis is the trade-offs that occur when balancing covariates across groups to improve the efficiency of the estimators. We also adapt methods from survey sampling to experimental settings. We provide specific solutions to problems such as assigning participants to multiple treatment groups where the covariate means are balanced between groups even, with groups of unequal size, or assignment to group probabilities. We also propose a method to perform exact hypothesis tests for models like Poisson regression, which reduce reliance on large-sample approximations.

Keywords : Balanced sampling, Design of experiments, Design variance matrix, Exact inference, Mixture of designs.

Resumé

Cette thèse étudie différents aspects de la conception des enquêtes et des expériences randomisées. Ces deux sujets présentent des similitudes et des différences particulières qui nous permettent d'établir des liens entre leurs théories. Une partie importante de cette thèse porte sur les compromis associés à l'équilibrage des covariables entre les groupes dans le but d'améliorer l'efficacité des estimateurs. Nous adaptons également des méthodes d'échantillonnage à des contextes expérimentaux. Nous fournissons des solutions spécifiques à des problèmes tels que l'affectation des participants à plusieurs groupes de traitement où les moyennes des covariables sont équilibrées entre les groupes, même avec des groupes de taille inégale ou des probabilités d'affectation à un groupe. Nous proposons également une méthode permettant d'effectuer des tests d'hypothèse exacts pour des modèles tels que la régression de Poisson, ce qui réduit la dépendance à l'égard des approximations sur de grands échantillons.

Mots-clés: Échantillonnage équilibré, Inférence exacte, Matrice de variance du plan, Plan d'expériences.

Contents

1	Introduction	11
1.1	Overview of sampling in surveys	12
1.2	Overview of design of experiments	12
1.3	Outline	14
2	Risk Minimization Using Robust Experimental or Sampling Designs and Mixture of Designs	15
2.1	Introduction	15
2.2	Sampling designs	16
2.3	Design of experiments	18
2.4	Main designs	19
2.5	Upper bounds for the variance	21
2.5.1	General results	21
2.5.2	Upper bounds for probability designs	22
2.5.3	Case of balanced sampling	24
2.6	Lower bound for the largest eigenvalues	25
2.7	Robustification of the design by mixture	25
2.8	Simulations with mixtures of designs	27
2.9	Conclusions	29
2.10	Supplementary material for Risk Minimization Using Robust Experimental or Sampling Designs and Mixture of Designs	30
2.10.1	Proofs	30
2.10.2	Approximations for the largest eigenvalue	33
2.11	Model- and design-based approach in experimental designs	35
2.11.1	Complete simulations results	37
3	Balanced Random Assignments for Multiple Treatment Groups Using the Cube Method	39
3.1	Introduction	39
3.2	Cube method description	40
3.3	Notation for multiple treatment group assignment	41
3.4	Multiple treatment group Cube method	43
3.5	Simulations	45
3.5.1	Improvement in balance between groups	45
3.5.2	Variance reduction in the estimators	48
3.6	Inference	48
3.6.1	Recommendation	48
3.6.2	Randomization test	48
3.6.3	Approximation of the variance of estimators	49
3.7	Conclusion	51

3.8	Supplement for Balanced Random Assignments for Multiple Treatment Groups Using the Cube Method	52
3.8.1	Code	52
3.8.2	Additional information on the simulations	53
4	Exact Inference for Poisson Regression With Continuous Covariates using Survey Sampling Algorithms	55
4.1	Introduction	55
4.2	Notation and Problem Presentation	56
4.3	The Cube method applied for exact inference in logistic regressions	58
4.4	Exact Inference for Poisson Regression Models using a Survey Sampling Algorithm	60
4.4.1	Fundamental Results	60
4.4.2	Algorithm with K configurations	62
4.4.3	Application on Data Sets	63
4.5	Simulations: Poisson regression with $K = n$	64
4.6	Discussion	65
4.7	Appendix	67
4.7.1	A preliminary result on Poisson regression estimates	67
4.7.2	Proof of Proposition 1	67
5	The distribution of the Mahalanobis distance between two groups under co- variates multinormality assumption	71
5.1	Introduction	71
5.2	Chi-square approximation of the Mahalanobis distance	72
5.3	Distribution of the square of the Mahalanobis distance between two groups	73
5.4	Simulations to compare the Beta and Chi-Square approximations	75
5.5	Discussion	76
5.6	Supplementary material	77
5.6.1	Simulations to compare the chi-square and beta distribution approximations	77
5.6.2	Links to MANOVA	79
6	Conclusion	81

Chapter 1

Introduction

Survey sampling methods are techniques used to randomly select a subset of individuals or units from a larger population for the purpose of conducting surveys or research. The goal is usually to estimate the mean or total of a variable of interest from a population using a sample. The experimental design is a way to carry out an experiment in order to assess the effects of one or more variables. Both fields share similarities in their use of randomization for inference, in the construction of the estimators and in the common aim to obtain unbiased and low-variance estimates through an optimal selection schema. This thesis studied some aspects of designs that could be applied to both survey sampling and experiments. More specifically, we propose methods to generate balanced samples or assignments to groups under different settings and constraints. It is not difficult to see that generating samples that are balanced on the Horvitz-Thompson estimator (Horvitz and Thompson, 1952) of the total of the auxiliary variables is conceptually similar to generating group assignments that are balanced on the means of their covariates.

There are also other aspects of experiments that are very similar to surveys but will not be covered by the thesis. One of those aspects is non-response. In experiments, non-compliance, participant dropout, or deaths occur quite frequently. In parallel, non-response rates have been rising for decades in surveys, even in government-mandated ones. Currently, non-response and ways to deal with it are highly studied topics (see amongst other Chen and Haziza, 2019; Haziza and Beaumont, 2017; Lee and Kim, 2022). Both fields have studied and adopted similar methods to deal with it, such as the concept of double robustness for the imputation of missing values (Robins et al., 1994; Kang and Schafer, 2007; Han and Wang, 2013).

Another notable similarity between experiments and surveys are data that arrive sequentially. In both fields, not all surveys or experiments have data available at the outset of the study or experiment. In experiments, this often occurs in clinical trials where participants arrive sequentially over time (Chow and Chang, 2008; Kapelner and Krieger, 2014, 2023). In this setting, the total number of experimental units is usually unknown, so groups of equal size cannot even be guaranteed (Efron, 1971). Similarly, in survey sampling, it is not uncommon that data collection is staggered over time. In both cases, data arriving sequentially over time create some challenge for choosing the appropriate design. Additionally, there is a possibility that initial findings call for an adaptation of the trial or the survey.

As a more general similarity, both surveys and experiments try to use a randomly selected subgroup of a larger population in order to make inference about them. In experiments, it would involve assessing the effect of a treatment or an intervention on a larger group, which in the example of a clinical trial would be patients suffering from a specific condition. In surveys, it would typically involve making an estimate of the total of a variable in a population.

1.1 Overview of sampling in surveys

Survey sampling methods are techniques used to select a subset of individuals or items from a larger population for the purpose of conducting surveys or research. The goal is usually to estimate the mean or total of a variable of interest $\mathbf{y} = (y_1, \dots, y_N)^\top$ from a population, using a sample. We can define a sample as a subset, chosen randomly, of a certain population U of size N and defined as the set $= \{1, \dots, N\}$. Let $\mathbf{a} = (a_1, \dots, a_N)^\top$ be an indicator variable, where $a_i = 1$ when the unit $i \in U$ is included in the sample and 0 otherwise. Each unit in the population has a probability to be included in the sample. Let $\boldsymbol{\pi} = (\pi_1, \dots, \pi_n)^\top$ be the vector of inclusion probabilities. In surveys, the inclusion probabilities are often set to be unequal, in order to minimize the variance of certain estimators.

An example of an estimator is the Horvitz-Thompson estimator of the total (Horvitz and Thompson, 1952), whose formula is

$$t_{HT} = \sum_{i=1}^N \frac{a_i y_i}{\pi_i}.$$

The estimator of mean can be retrieved by dividing the Horvitz-Thompson estimator of the total by N , when it is known. The Horvitz-Thompson estimator is the basic estimator in surveys. It is designed to estimate population characteristics, usually the total or the mean, based on a sample. The basic idea is to assign sampling weights to each unit in the sample to obtain an unbiased estimate, which is simply the inverse of the probability of selection. It is used when the sampling frame is known and each unit in the population has a known and non-zero probability of being selected.

Another important concept in survey sampling is auxiliary variables. We assume that auxiliary variables are correlated with the variables of interest and are available for each unit in the population before the sampling procedure. We define the auxiliary variables for unit $i \in U$ as a vector of dimension $p \times 1$ and denote them as \mathbf{x}_i . One utility of the auxiliary variables is to help to determine potentially extreme units. This information can be used to reduce the variance of the Horvitz-Thompson estimator by sampling with higher probability the units in the population that we expect to have more extreme values for the response variable (see for instance Tillé, 2020a).

Auxiliary variables can also be used to create more efficient sampling designs. An aspect of that would be to ensure that the sample is balanced with respect to the population U . We define a sample \mathbf{a} as balanced when the Horvitz-Thompson estimator of the auxiliary variables is equal to the population total of the auxiliary variables. We obtain the following equation

$$\sum_{i=1}^N \frac{a_i \mathbf{x}_i}{\pi_i} = \sum_{i=1}^N \mathbf{x}_i. \quad (1.1)$$

Satisfying Equation (1.1), even approximately, is not a straightforward task. The sample still has to be random and respect the inclusion probabilities. Having a balanced sample is particularly useful in scenarios where the population is heterogeneous and the auxiliary information is strongly correlated with the main variable of interest. It improves the precision of the estimates by reducing the variance that arises from sampling.

1.2 Overview of design of experiments

We first clarify the terms used in this thesis. The term “design of experiments” will concern only the allocation to groups of randomized experiments whose units have fixed covariates or confounding variables, such as clinical trials that have the most similarities with survey sampling.

There exist cases in the literature with several independent variables and whose design is non-random and is simply optimized for a certain criterion (Atkinson and Donev, 1992). Agricultural block designs and industrial processes optimization are examples of applications where those type of experiments can occur. In this thesis, the term “balanced experimental design” refers to designs whose possible assignment to groups have covariates means that are close to each other, which is the counterpart of a balanced sampling design in surveys. In the literature of the design of experiments, a balanced experiment will often simply refer to experiments whose different treatment groups are of the same size.

We also clarify that we do not deal with observational studies (Cochran, 1965; Rubin, 2008) where the “assignment to the treatment group” is not controlled by the experimenter and is simply given by the data. There is no randomization controlled by the experimenter before the creation or collection of data. It creates problems in regard to inference due to higher potential of bias but also presents advantages such as less ethical problems with treatments that could be deleterious to the participant. Another advantage of observational studies is that they are generally cheaper than randomized experiments, since it simply collects data which is already available. We also do not deal with any type of self-controlled experiments where individuals receive both the control and the treatment, whose advantage is that it allows to easily control for time-invariant confounders.

Suppose we have a population U of size N and that we want to test a treatment T on the units $i \in U$. If unit i receives the treatment, its response variable becomes y_i^T . If unit i does not receive the treatment, it belongs to the control group and its response is denoted by y_i^C . This is a randomized experiment with two groups. The goal is to estimate the difference between the two groups

$$\tau = \frac{1}{N} \sum_{i \in U} y_i^T - \frac{1}{N} \sum_{i \in U} y_i^C. \quad (1.2)$$

Equation (1.2) is similar to the Neyman–Rubin causal model of potential outcomes (Rubin, 1974). The values of y_i^T and y_i^C cannot be observed at the same time for unit k , because all units either belong to the control group or the treatment group. We can assume that $a_i = 1$ when unit $i \in U$ is included in the treatment group and $a_i = 0$ when included in the group. The units have to be assigned randomly between two groups. The random aspect allows to make inference and to avoid the potential investigator bias in experiments. Often, those experiments would have confounding variables or covariates that are available and are suspected to have an effect on the response variable. The covariates for unit $i \in U$ can be denoted by the $N \times 1$ vector \mathbf{x}_i .

Then, τ can be estimated using the difference of the two the Horvitz-Thompson estimators:

$$\hat{\tau} = \frac{1}{N} \sum_{i \in U} \frac{y_i^T a_i}{\pi_i} - \frac{1}{N} \sum_{i \in U} \frac{y_i^C (1 - a_i)}{1 - \pi_i}.$$

When the sample size is fixed and the inclusion probabilities are all equal, $\hat{\tau}$ simply becomes equivalent to the widely known difference in means estimator. Later in the thesis, it will be shown mathematically that if the response variable follows a linear model with the covariates, the variance of the estimators is directly reduced when using a balanced design for both experiments and surveys. But it can also be inferred intuitively that comparing two groups that are similar is much easier than comparing groups that are quite different. An extreme scenario that can illustrate it would be if the control group contains participants of only one sex and the treatment group contains people of the other sex. In this case, it is impossible to determine whether it was the treatment or the sex of the participants that caused any difference in the response variable between the two groups, which is simply a consequence of multicollinearity in the matrix of the predictors.

There are different types of common experimental designs, such as rerandomization, the randomized block or stratified design or matching, that address the issue of covariate balance (Simon, 1979; Morgan and Rubin, 2012; Xu and Kalbfleisch, 2010; Li et al., 2018). Surveys often use sampling methods that are equivalent to the previous experimental designs that have been mentioned but under a different nomenclature.

1.3 Outline

Chapter 2 establishes a theoretical basis for evaluating the robustness and balance trade-off in sampling and experimental designs. By defining robustness through the spectral properties of the variance matrix of the design, we derive bounds and approximations for the largest eigenvalue of this matrix and use it as a measure of potential variance inflation in estimators like the Horvitz-Thompson total. The robustness of the design strongly correlates with the entropy or simply the number of possible samples of the design when inclusion probabilities are equal. We propose a method for mixing two designs in order to handle the robustness and balance trade-off. This method interpolates between two sampling strategies using a tunable parameter. Through simulations, we validate the accuracy of our eigenvalue approximations and compared our algorithm with the Gram-Schmidt walk design (Harshaw et al., 2024), we show that our method matches the ability to navigate the efficiency-robustness trade-off while being simpler and offering a more interpretable variance structure.

In Chapter 3, we introduce a modified Cube Method adapted for experimental covariate balance with multiple groups. Unlike existing approaches, our adaptation accommodates unequal treatment group sizes, heterogeneous inclusion probabilities, and balancing of the covariates. We also provide variance approximations for scenarios where randomization tests are computationally prohibitive.

Chapter 4 extends the utility of the Cube method into exact inference for generalized linear models. By using a stratified variant of the algorithm, we develop an exact testing procedure for Poisson regression with continuous covariates, reducing reliance on large-sample approximations. We generalize the central proposition of Rivest and Gaye (2023) to the Poisson setting by using a mild additional assumption, which justifies using the method for count data scenarios.

Finally, the thesis concludes with a note on the distribution of the Mahalanobis distance between two groups in Chapter 5, which can be applied for methods such as rerandomization and potential variance approximations under certain settings.

Chapter 2

Risk Minimization Using Robust Experimental or Sampling Designs and Mixture of Designs

Abstract

For both experimental and sampling designs, the efficiency or balance of designs has been extensively studied. There are many ways to incorporate auxiliary information into designs. However, when we use balanced designs to decrease the variance due to an auxiliary variable, the variance may increase due to an effect which we define as lack of robustness. This robustness can be written as the largest eigenvalue of the variance operator of a sampling or experimental design. If this eigenvalue is large, then it might induce a large variance in the Horvitz-Thompson estimator of the total. We calculate or estimate the largest eigenvalue of the most common designs. We determine lower, upper bounds and approximations of this eigenvalue for different designs. Then, we compare these results with simulations that show the trade-off between efficiency and robustness. Those results can be used to determine the proper choice of designs for experiments such as clinical trials or surveys. We also propose a new and simple method for mixing two sampling designs, which allows to use a tuning parameter between two sampling designs. This method is then compared to the Gram-Schmidt walk design, which also governs the trade-off between robustness and efficiency. A set of simulation studies shows that our method of mixture gives similar results to the Gram-Schmidt walk design while having an interpretable variance matrix. ¹

2.1 Introduction

Both experimental designs and sampling designs have a common feature: random variables are generated in order to select units either to constitute a sample or to create a test group and a control group. For survey sampling, the recommendations relate primarily to issues of design efficiency or balance under certain hypotheses relative to the variables of interest to be estimated. We can thus resort to designs of fixed size, stratified, balanced with equal or unequal probabilities (see amongst other Brewer and Hanif, 1983; Tillé, 2006, 2020a). For experimental designs, randomization between the test group and the control group is recommended. There are also stratification techniques, blocking, matching, rerandomization, balanced sampling that can make the two groups balanced to some degree (see amongst other Simon, 1979; Morgan and Rubin, 2012; Xu and Kalbfleisch, 2010; Li et al., 2018; Tillé, 2022). Unfortunately, everything has a price. It can be shown that if a design is more efficient for one variable, it is necessarily less efficient for another variable. Thus, one cannot rely solely on a single efficiency criterion to determine the appropriate design.

¹This chapter is a reprint of Talovic and Tillé (2025)

The notion of robustness for design of experiments has been introduced by Harshaw et al. (2024) and to some degree Kapelner et al. (2022). This notion is also valid for sampling designs. Robustness consists in identifying the eigenvector that can be associated to the largest eigenvalue of the variance matrix of the design. This eigenvector is associated with the largest variance in the estimators. We can thus evaluate the potential maximum price to pay for the use of a design that integrates auxiliary information. More details are given in Section 2.2.

In survey sampling, a few closely related topics have already been examined. Qualité (2009, p.45-57) studied the dispersion of the eigenvalues in sampling designs. The author also made links with the dispersion of the eigenvalues and the entropy of the sampling design. Stenger (1979) and Stenger and Gabler (1996) found some results related to minimax strategies in survey sampling.

Harshaw et al. (2024) propose to consider a trade-off between the efficiency and the robustness of a design. They also propose an algorithm called “Gram-Schmidt Walk design” (GS) in which there is a tuning parameter that allows us to oscillate between a design that would be totally robust and a balanced design that would integrate the auxiliary information as well as possible. Depending on the risk aversion, one can choose the appropriate value of the tuning parameter.

In this article, we review the different sampling and experimental designs. We evaluate the robustness of each of these designs. To do so, we give upper and lower bounds for the largest eigenvalues of the variance matrices of the sampling or experimental designs. We also give approximations for these eigenvalues in the supplementary material. We show the determining role of the Yates-Grundy condition which allows us to give a relatively low upper bound. Most of the designs are relatively robust except for systematic sampling which has a large maximal eigenvalue. We also study the potential effect of the choice of the sampling design on the variance of the Horvitz-Thompson estimator.

We propose a simple procedure that allows us to mix two designs. This method offers an alternative more flexible than the GS design of Harshaw et al. (2024). In addition, this method allows to mix almost any kind of designs, which opens the way to multiple combinations of designs. We next present a set of simulations that show that our mixing method gives similar results to the GS design while being more flexible and simpler. These simulations also show that the Cube method is actually quite robust in the sense that the largest eigenvalue of the variance matrix is relatively small.

We define survey designs in Section 2.2 and experimental designs with a design-based approach in Section 2.3. The main designs are briefly described in Section 2.4. Section 2.5 is dedicated to the identification of the upper bound of the largest eigenvalue of the variance matrix. In Section 2.6, a lower bound is sought for this matrix, which allows us to obtain a range of possible values for this eigenvalue. Section 2.7 is dedicated to the concept of robustification using mixture of designs. We present a new method for mixing two sampling designs. In Section 2.8, we perform simulations on the main sampling designs and compare them with the mixtures of designs. Finally, conclusions are drawn in Section 2.9.

In the supplementary material, we include all the proofs and we add approximations for the largest eigenvalue for different sampling designs. We also discuss the interest of mixing designs under model- and design-based approach for experimental designs. We show that the interest of mixing design is relatively limited under certain hypotheses. Finally, we give additional information regarding simulations.

2.2 Sampling designs

Consider a finite population $U = \{1, \dots, k, \dots, N\}$ of size N . A sample without replacement is a subset, $s \subset U$, of the population. A sampling design $p(\cdot)$ assigns to each sample s a probability

such that $p(s) \geq 0$ and

$$\sum_{s \subset U} p(s) = 1.$$

A random sample S takes as value s with probability $Pr(S = s) = p(s)$. For sampling with replacement, the s may include a unit of U more than once.

Let a_k denote the indicator random variable which takes as value 1 if unit k is in S and 0 otherwise. If the sampling is with replacement a_k is the number of times unit k is selected in the sample. The random sample can also be defined by the column vector of the indicator variables $\mathbf{a}^\top = (a_1, \dots, a_N)$. We define the vector $\boldsymbol{\pi}^\top = \mathbf{E}_a(\mathbf{a})^\top = (\pi_1, \dots, \pi_N)$, where $\mathbf{E}_a(\cdot)$ is the expectation under the sampling design. For sampling designs without replacement, $\boldsymbol{\pi}$ is the vector of first-order inclusion probabilities. In this article, when the order is not defined, it means we refer to the first-order inclusion probability. It is assumed that the sum of the inclusion probabilities is an integer number denoted by n .

We define the matrix $\mathbf{\Pi} = \mathbf{E}_a(\mathbf{a}\mathbf{a}^\top)$. For sampling designs without replacement, $\mathbf{\Pi}$ is the matrix of joint inclusion probabilities. The covariances between the indicator variables is defined by $\Delta_{k\ell} = \text{cov}_a(a_k a_\ell) = \mathbf{E}_a(a_k a_\ell) - \mathbf{E}_a(a_k)\mathbf{E}_a(a_\ell) = \pi_{k\ell} - \pi_k \pi_\ell$, for all $k, \ell \in U$. The variance-covariance matrix is thus $\mathbf{\Delta} = \text{var}_a(\mathbf{a}) = \mathbf{\Pi} - \boldsymbol{\pi}\boldsymbol{\pi}^\top$, where $\text{var}_a(\cdot)$ is the variance under the sampling design.

The aim is to estimate the total of a variable of interest $\mathbf{y}^\top = (y_1, \dots, y_N)$ given by

$$t_y = \sum_{k \in U} y_k.$$

If $\pi_k > 0$, for all $k \in U$, then the Horvitz-Thompson estimator (Horvitz and Thompson, 1952) given by

$$\hat{t}_y = \sum_{k \in S} \frac{y_k}{\pi_k}$$

is an unbiased estimator of t_y . If there is a unit k such that $\pi_k = 0$, the estimator is biased. The variance of this estimator is

$$\text{var}_a(\hat{t}_y) = \sum_{k \in U} \sum_{\ell \in U} \frac{y_k y_\ell}{\pi_k \pi_\ell} \Delta_{k\ell} = \mathbf{y}^\top \mathbf{D}^{-1} \mathbf{\Delta} \mathbf{D}^{-1} \mathbf{y},$$

where $\mathbf{D} = \text{diag}(\boldsymbol{\pi})$. If the sampling design has a fixed sample size, Yates and Grundy (1953) showed that the variance can also be written as

$$\text{var}_a(\hat{t}_y) = -\frac{1}{2} \sum_{k \in U} \sum_{\ell \in U} \left(\frac{y_k}{\pi_k} - \frac{y_\ell}{\pi_\ell} \right)^2 \Delta_{k\ell}.$$

If $\pi_{k\ell} > 0$, for $k, \ell \in U$, two unbiased estimators of variance can be constructed. The Horvitz-Thompson estimator of the variance is given by

$$\widehat{\text{var}}_a(\hat{t}_y) = \sum_{k \in S} \sum_{\ell \in S} \frac{y_k y_\ell}{\pi_k \pi_\ell} \frac{\Delta_{k\ell}}{\pi_{k\ell}}.$$

The Yates-Grundy estimator is only available for sampling designs with fixed sample size and is given by

$$\widehat{\text{var}}_a(\hat{t}_y) = -\frac{1}{2} \sum_{k \in S} \sum_{\ell \in S} \left(\frac{y_k}{\pi_k} - \frac{y_\ell}{\pi_\ell} \right)^2 \frac{\Delta_{k\ell}}{\pi_{k\ell}}. \quad (2.1)$$

Definition 1. A design satisfies the Yates-Grundy condition if $\Delta_{k\ell} \leq 0$ for all $k \neq \ell \in U$.

Yates and Grundy (1953) suggest using designs that satisfy this condition, because then the

estimator of the variance (2.1) cannot be negative. This condition is related to negative correlation, which is an important hypothesis to construct central limit theorems in finite population (Brändén and Jonasson, 2012; Bertail et al., 2017; Gerber et al., 2019). In this paper, we will show that the Yates-Grundy condition also contribute to robustify a sampling or experimental design.

2.3 Design of experiments

In randomized experiments, and more specifically for clinical trials, the problem is somewhat different. Suppose we have a population U of size N and that we want to test a treatment T on the units $k \in U$. If unit k receives the treatment, its response variable becomes y_k^T . If unit k does not receive the treatment, it belongs to the control group and its response is denoted by y_k^C . This is a randomized experiment with two groups. The goal is to estimate

$$\tau = \frac{1}{N} \sum_{k \in U} y_k^T - \frac{1}{N} \sum_{k \in U} y_k^C.$$

Here we adopt a purely design based approach as in Harshaw et al. (2024). The values of y_k^T and y_k^C cannot be observed at the same time for unit k , because no unit can belong to both the control and treatment groups. We select a sample with inclusion probabilities $\boldsymbol{\pi} = (\pi_1, \dots, \pi_N)$ and apply the treatment to the chosen units.

Then, τ can be estimated using the difference between the Horvitz-Thompson estimators of the two groups:

$$\begin{aligned} \hat{\tau} &= \frac{1}{N} \sum_{k \in U} \frac{y_k^T a_k}{\pi_k} - \frac{1}{N} \sum_{k \in U} \frac{y_k^C (1 - a_k)}{1 - \pi_k} \\ &= \frac{1}{N} \sum_{k \in U} \frac{y_k^T (1 - \pi_k) + y_k^C \pi_k}{(1 - \pi_k) \pi_k} a_k - \frac{1}{N} \sum_{k \in U} \frac{y_k^C}{(1 - \pi_k)}. \end{aligned}$$

If we define $\mathbf{z} = (z_1, \dots, z_N)^\top$, where

$$z_k = \frac{y_k^T (1 - \pi_k) + y_k^C \pi_k}{(1 - \pi_k) \pi_k},$$

the variance becomes

$$\text{var}_a(\hat{\tau}) = \frac{\mathbf{z}^\top \text{var}_a(\mathbf{a}) \mathbf{z}}{N^2} = \frac{\mathbf{z}^\top \boldsymbol{\Delta} \mathbf{z}}{N^2}. \quad (2.2)$$

If the inclusion probabilities are equal then

$$z_k = \frac{y_k^T (1 - \pi) + y_k^C \pi}{(1 - \pi) \pi}.$$

Moreover, if $\pi = 1/2$ then

$$z_k = 2(y_k^T + y_k^C).$$

In both experimental and survey designs, the variance operator $\boldsymbol{\Delta}$ plays a fundamental role in the precision of the estimates.

The variance given in (2.2) may seem similar to the variance in a sampling design. However, an important difference is that z_k is not known at the sample level. Estimating the variance is therefore a more difficult problem than for a sampling design. Tillé (2022) has proposed an estimator of this variance in the context of design-based inference. If the largest eigenvalue of $\boldsymbol{\Delta}$ is large, then there is a risk that the variance of the Horvitz-Thompson estimator becomes large

depending on the values of the response variable. In this article, we study only the eigenvalues of $\mathbf{\Delta}$ and not $\mathbf{D}^{-1}\mathbf{\Delta}\mathbf{D}^{-1}$, because the inclusion probabilities are often all equal in experiments and the inclusion probabilities are often to some degree proportional to the variable of interest in survey sampling.

2.4 Main designs

In this section, we briefly describe the main designs that can be used to select a sample or to randomize an experiment. The terms used for the designs are mainly coming from the field of sampling and will be freely used for their equivalent in experimental designs. In *Poisson sampling*, each unit of the population is subjected to an independent Bernoulli trial a_k with $E_a(a_k) = \pi_k, k \in U$. We thus have $\Delta_{kk} = \pi_k(1 - \pi_k), k \in U$ and $\Delta_{k\ell} = 0$ when $k \neq \ell$. Matrix $\mathbf{\Delta}$ is diagonal. The sample size is random and has a Poisson-Binomial distribution (Hodges and Le Cam, 1960; Stein, 1990; Deville, 2000; Chen, 1993). *Bernoulli sampling* is the special case of Poisson Sampling when the inclusion probabilities are equal and are denoted by π . The sample size is random and has a Binomial distribution with parameter π and exponent N .

Simple Random Sampling Without Replacement (SRSWOR) is the design with a fixed sample size n for which all samples of size n have the same probability $n!(N - n)!/N!$ of being selected. For this design,

$$\pi_k = \frac{n}{N}, \Delta_{kk} = \frac{n(N - n)}{N^2}, \pi_{k\ell} = \frac{n(n - 1)}{N(N - 1)},$$

and

$$\Delta_{k\ell} = -\frac{n(N - n)}{(N - 1)N^2}, k \neq \ell \in U.$$

Stratification consists of partitioning the population in H strata U_1, \dots, U_H of sizes N_1, \dots, N_H . Next, in each stratum h , a sample of size n_h is selected independently from the other strata with SRSWOR.

Maximum entropy design also called *Conditional Poisson Sampling* (CPS) is the design with fixed sample size n that maximizes the entropy

$$- \sum_{s \subset U | \#s=n} p(s) \log p(s),$$

given fixed inclusion probabilities π_k . The implementation of this design is quite complex and has been resolved relatively recently (Chen, 1993; Chen and Liu, 1997; Tillé, 2006). It is possible to compute the matrix of inclusion probabilities that satisfy the Yates-Grundy condition, for instance by using the R sampling package (Tillé and Matei, 2021). If the inclusion probabilities are equal, the maximum entropy design reduces to SRSWOR.

Sampford (1967) method is described in the book of Brewer and Hanif (1983) as follows “Select the first unit with probability proportional to measure of size π_k/n . At each subsequent draw, select with probability of selection proportional to $\pi_k(1 - \pi_k)$ with replacement. If any unit is selected twice, reject the whole sample selected and start again.” This design satisfy the Yates-Grundy condition. If the inclusion probabilities are equal, it also reduces to SRSWOR. Sampford sampling is also relatively close to CPS (Tillé, 2006, p.143-145).

Unequal probabilities systematic sampling (Madow and Madow, 1944; Madow, 1949) is defined as follows. First, compute the cumulated inclusion probabilities $V_k = \sum_{\ell=1}^k \pi_\ell$ with $V_0 = 0$ and $V_N = n$. Next generate a uniform continuous random variable u and select in the sample the units k such that $V_{k-1} \leq u + j < V_k$ for $j = 0, \dots, n - 1$. This sample has a fixed sample size n . Systematic sampling has no more than N samples with non-zero probability (Pea et al., 2007). This design has a very small entropy. The joint inclusion probabilities can be computed for instance using the R sampling package (Tillé and Matei, 2021). *Equal probability systematic*

sampling is the special case where the inclusion probabilities are equal. Moreover, if N/n is an integer, then only N/n samples have a non-null probability of being selected.

The *pivotal method* is one of the special cases of the splitting method proposed by Deville and Tillé (1998). The method was republished by Srinivasan (2001). At each step of the pivotal method, two units whose inclusion probabilities are not integers are randomly modified and one of these units is set to 0 or 1.

By choosing at each step a couple units close in a space, Grafström et al. (2012) have built a method to obtain well spread samples. The method satisfies the Yates-Grundy condition (Kramer et al., 2011; Chauvet, 2012). If the initial probabilities are all equal to $1/2$, the design reduces to a *matched pairs design* (Matched). The Matched design is also a stratified design where $N_h = 2$ and $n_h = 1$, for $h = 1, \dots, H$. In a Matched design, there are $H = N/2$ pairs on units and

$$\Delta_{k\ell} = \begin{cases} \frac{1}{4} & \text{if } k = \ell \in U_h \\ -\frac{1}{4} & \text{if } k \neq \ell \in U_h \\ 0 & \text{if } k \in U_h, \ell \in U_i, h \neq i. \end{cases} \quad (2.3)$$

Let matrix $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^\top$ denote the $N \times p$ matrix of p covariates. *Balanced sampling* refers to methods for selecting random samples \mathbf{a} that approximately satisfy the balancing equations:

$$\sum_{k \in U} \frac{a_k \mathbf{x}_k}{\pi_k} \approx \sum_{k \in U} \mathbf{x}_k. \quad (2.4)$$

Fixed sample size is a special case of balanced sampling when matrix \mathbf{X} contains a variable that is equal or proportional to the inclusion probabilities. The *cube method* (Cube) (Deville and Tillé, 2004a) selects samples that are approximately balanced and also respect the inclusion probabilities of each unit of the population. When all inclusion probabilities are equal to $\pi = n/N$, then the sample mean is approximately equal to the population mean

$$\frac{1}{n} \sum_{k \in U} a_k \mathbf{x}_k \approx \frac{1}{N} \sum_{k \in U} \mathbf{x}_k.$$

The main phase of the cube method consists of randomly modifying the inclusion probability vector at each step to obtain the sequence $\boldsymbol{\pi}(0) = \boldsymbol{\pi}, \boldsymbol{\pi}(1), \boldsymbol{\pi}(2), \dots, \boldsymbol{\pi}(t), \dots$ by setting one entry to 0 or 1 at each step while respecting the constraints that $\mathbb{E}\{\boldsymbol{\pi}(t) \mid \boldsymbol{\pi}(t-1)\} = \boldsymbol{\pi}(t-1)$ and that

$$\sum_{k \in U} \frac{\pi_k(t) \mathbf{x}_k}{\pi_k} \approx \sum_{k \in U} \mathbf{x}_k,$$

for all t . If the main phase of the algorithm does not terminate with a vector of integers, the second phase of the algorithm handles the few remaining non-integer entries. The method satisfies the inclusion probabilities exactly and is described in detail in Tillé (2006).

Tillé (2022) has proposed *Cube matched pairs design* (Cube-Matched) that combines the matched pairs design with balanced sampling by the Cube method. Pairs of similar units are constructed and only one unit is selected in each pair. Moreover, the design is balanced on covariates. This method is a simple application of the stratified-balanced sampling where each stratum contains only two units from which one is selected in the sample (Chauvet, 2009; Hasler and Tillé, 2014; Jauslin et al., 2021).

Multinomial sampling consists of selecting n times from the population a unit k with probability $p_k = \mathbb{E}_a(a_k)/n = \pi_k/n$ and with replacement. For sampling designs with replacement, the variable a_k contains the number of times the unit k is selected in the sample and has a binomial distribution with parameter $p_k = \pi_k/n$ and exponent n . To distinguish multinomial sampling from other sampling designs, we define the matrix $\boldsymbol{\Delta}^R = \text{var}_a(\mathbf{a})$, where $\Delta_{kk}^R = np_k(1 - p_k)$ and

$\Delta_{k\ell}^R = \text{cov}_a(a_k, a_\ell) = -np_k p_\ell = -\pi_k \pi_\ell / n$, for $k \neq \ell$. When $p_k = 1/N$, the multinomial sampling reduces to the *Simple Random Sampling With Replacement* (SRSWR). The properties of the main designs introduced are described in Table 2.1.

Table 2.1: Summary of the properties of some sampling methods. The maximum entropy column refers to maximum entropy sampling in the “obvious” support domain (e.g., fixed sample size) given certain inclusion probabilities.

	Fixed Size	Without Replacement	Equal Probabilities	Yates-Grundy Condition	Maximum Entropy
Poisson		✓		✓	✓
Bernoulli		✓	✓	✓	✓
SRSWOR	✓	✓	✓	✓	✓
Stratification	✓	✓	✓	✓	✓
CPS	✓	✓		✓	✓
Sampford	✓	✓		✓	
Unequal Systematic	✓	✓			
Equal Systematic	✓	✓	✓		
Pivotal	✓	✓		✓	
Cube	✓	✓			
Multinomial	✓			✓	✓
SRSWR	✓		✓	✓	✓

2.5 Upper bounds for the variance

2.5.1 General results

Matrix Δ is the heart of the variance for both the sampling and experimental designs. The analysis of this matrix is therefore the crucial question for the determination of the design that is unbiased and also has a low variance. The natural way to analyse the Δ operator is to perform a diagonalization. We can look for a vector \mathbf{u} which maximizes

$$\frac{\mathbf{u}^\top \Delta \mathbf{u}}{\mathbf{u}^\top \mathbf{u}}.$$

The solution to this maximization problem is an eigenvector \mathbf{u} of Δ associated to the largest eigenvalue λ_1 . This eigenvalue is associated with the greatest risk of the design, as this eigenvalue is directly associated with largest possible variance of the Horvitz-Thompson estimator.

A first result shows that there is no design that is uniformly better than the others for given inclusion probabilities. Consider two designs $p_1(\cdot)$ and $p_2(\cdot)$ with the same inclusion probabilities π_k and whose variance matrices are, respectively, Δ_1 and Δ_2 . Design $p_1(\cdot)$ would be uniformly better than design $p_2(\cdot)$ if $\mathbf{u}^\top \Delta_1 \mathbf{u} \leq \mathbf{u}^\top \Delta_2 \mathbf{u}$ for all $\mathbf{u} \in \mathbb{R}^N$ and if there is at least one vector \mathbf{y} such that $\mathbf{y}^\top \Delta_1 \mathbf{y} < \mathbf{y}^\top \Delta_2 \mathbf{y}$. We refer to the following result:

Result 1. *There exist no pair of design $p_1(\cdot)$ and $p_2(\cdot)$ without replacement, such that $p_2(s)$ is uniformly better than $p_1(\cdot)$ with the same first-order inclusion probability.*

We remind that the proofs of the results can be found in the supplement. The problem can be seen as explained in what follows. Since the inclusion probabilities determine the diagonal of Δ , two designs that have the same inclusion probabilities have the same trace for Δ and therefore they have the same sum of eigenvalues. Each eigenvalue gives a dispersion factor in the direction of an eigenvector. Therefore, if the dispersion is decreased in one direction, it is necessarily increased in another direction.

Among designs without replacement, Result 1 shows that there is no design which is better than all the others given certain inclusion probabilities. If additional auxiliary information is not

available, it is reasonable to try to equalize the eigenvalues and thus to take the most random design possible, that is to say, the design that maximizes the entropy. Additionally, we often measure several variables of interest so it is preferable to use a robust design.

For given inclusion probabilities, the design that is the most random is considered to be the one that maximizes the entropy. Poisson sampling maximizes the entropy

$$-\sum_{s \subset U} p(s) \log p(s),$$

subject to given inclusion probabilities. Then, we have

$$\Delta = \text{diag}(\pi_1(1 - \pi_1), \dots, \pi_N(1 - \pi_N)).$$

The largest eigenvalue is $\lambda_1 = \max_{k \in U} \pi_k(1 - \pi_k)$. If the inclusion probabilities are equal, i.e. if $\pi_k = \pi$, Poisson sampling reduces to Bernoulli sampling. Then, we have $\Delta = \pi(1 - \pi)\mathbf{I}$, where \mathbf{I} is an $N \times N$ identity matrix. This design can therefore be considered as the most robust design for given π , because the variance is the same in all directions.

The Poisson and Bernoulli designs have the disadvantage of having a random sample size. A fixed size allows to control the budget of the survey. Additionally, in our simulations, having fixed sample size greatly improves the balance of the design. However, fixed sample size implies that Δ has a null eigenvalue.

Result 2. *Let $p(\cdot)$ be a sampling design of fixed sample size n from a population U of size N , then Δ has a null eigenvalue associated to the vector of ones $\mathbf{1} \in \mathbb{R}^N$.*

In SRSWOR

$$\Delta = \frac{n(N - n)}{N(N - 1)}\mathbf{P},$$

where $\mathbf{P} = \text{diag}(\mathbf{1}) - \mathbf{1}\mathbf{1}^\top/N$ and $\mathbf{1}$ is a vector of ones of dimension N . Matrix \mathbf{P} is idempotent and has $N - 1$ eigenvalues equal to 1 and one eigenvalue equal to 0. Matrix Δ has thus $N - 1$ eigenvalues equal to $n(N - n)/\{N(N - 1)\}$ and one equal to 0. In simple random sampling with replacement, Δ^R is equal to $\mathbf{P}n/N$. Matrix Δ^R has thus $N - 1$ eigenvalues equal to n/N and one equal to 0. Simple random sampling without replacement is always more efficient than with replacement. The comparison of SRSROW and SRSWR is an example where Result 1 is not applicable, because it applies only to designs without replacement. SRSWOR is a robust design because the variance is the same for all centred vectors. The variance is zero when the variable of interest is a constant.

Result 3. *Let $p(s)$ be a sampling design of fixed sample size n from a population U of size N and with equal $\pi_k = \pi, k \in U$, then $\mathbf{1}$ is an eigenvector of $\mathbf{\Pi}$ and Δ . For any other centred eigenvector \mathbf{u} of Δ , \mathbf{u} is also an eigenvector of $\mathbf{\Pi}$ with the same eigenvalue.*

2.5.2 Upper bounds for probability designs

Upper bounds can be obtained for different types of designs.

Result 4. *For any sampling design on a finite population U of size N , the largest eigenvalue of Δ is smaller than or equal to $\max_{k \in U} (\sum_{\ell \in U} \pi_{k\ell})$. If the sampling design has a fixed sample size n , then the bound can be expressed as $n \max_{k \in U} (\pi_k)$.*

For designs without replacement, Δ is also the variance matrix of the complementary design selected with inclusion probabilities $1 - \pi_k, k \in U$. For sampling with fixed sample size and without replacement, an upper bound for λ_1 is

$$\lambda_1 \leq \min \left\{ n \max_{k \in U} (\pi_k), (N - n) \max_{k \in U} (1 - \pi_k) \right\}.$$

Result 4 induces directly a few corollary results for equal probability designs.

Corollary 1. *Let Δ be the variance matrix of a sampling design with fixed sample size n and equal $\pi_k = n/N, k \in U$, then the largest eigenvalue of Δ is smaller than or equal to $n\pi = n^2/N$.*

Using a similar reasoning with the complementary sample $\mathbf{1} - \mathbf{a}$, we obtain

Corollary 2. *Let Δ be the variance matrix of a sampling design without replacement with fixed sample size n and equal $\pi_k = n/N, k \in U$, then the largest eigenvalue of Δ is smaller than or equal to $(N - n)(1 - \pi) = (N - n)^2/N$.*

When $n > N/2$, the bound in Corollary 1 is worse than the bound deduced by the trace of Δ , which is equal to $N\pi(1 - \pi)$. When $n > N/2$, the bound from Corollary 2 becomes better. With equal probabilities, without replacement and fixed sample size, a general bound for the largest eigenvalue λ_1 of Δ is thus

$$\lambda_1 \leq \min \left(\frac{n^2}{N}, \frac{(N - n)^2}{N} \right).$$

This bound can be reached. Here is an example:

Example 1. Suppose that N is a multiple of n , i.e. $r = N/n \in \mathbb{N}$ and that we select a systematic sample with equal inclusion probabilities n/N . Then $\pi_{k\ell} = n/N$ if $(k - \ell \bmod r) = 0$ and 0 otherwise. We have $\pi_{k\ell} = n/N - n^2/N^2$ if $(k - \ell \bmod r) = 0$ and $-n^2/N^2$ otherwise. Define a vector $\mathbf{v}^\top = (v_1, \dots, v_N)$ where $v_k = 1$ if $k \bmod r = 1$, $v_k = -1$ if $k \bmod r = 2$ and $v_k = 0$ otherwise. We obtain $\Delta\mathbf{v} = (n^2/N)\mathbf{v}$. In this case $r - 1$ eigenvalues are equal to n^2/N and $N - r$ are equal to zero. If $r = 2$, $n = N/2$. All eigenvalues are equal to 0, except one that is equal to $N/4$.

Given $\boldsymbol{\pi}$, systematic sampling is the worst design in terms of robustness, because it concentrates all the dispersion in a very small number of directions or only one if N/n is an integer. Now, we introduce a new result for unequal probability designs.

Result 5. *For any sampling design on a finite population U of size N , the largest eigenvalue of Δ is smaller than or equal to $\max_{k \in U} (\sum_{\ell \in U} |\Delta_{k\ell}|)$.*

The bound of Result 5 will often be smaller than the one from Result 4 but it is not always the case. The sampling design that selects one unit out of a population of 3 units with equal probability is a simple example where $n \max_{k \in U} (\pi_k) = 1/3 < \max_{k \in U} (\sum_{\ell \in U} |\Delta_{k\ell}|) = 4/9$.

If the Yates-Grundy condition is satisfied, a simpler bound than the previous ones can be inferred.

Corollary 3. *For any sampling design without replacement that satisfies the Yates-Grundy condition and has fixed sample size n , the largest eigenvalue of Δ is smaller than or equal to $2 \max_{k \in U} \{\pi_k(1 - \pi_k)\}$.*

With a similar reasoning, we can obtain the following result.

Corollary 4. *For any sampling design with a fixed sample size n on a finite population U of size N , the largest eigenvalue of Δ is smaller than or equal to $2 \max_{k \in U} \left(\sum_{\ell \in U | \Delta_{k\ell} > 0} \Delta_{k\ell} \right)$.*

Corollary 4 shows that Yates-Grundy condition is desirable for fixed samples size designs as they give a lower theoretical maximal eigenvalue of Δ . It also indicates that sampling designs with fixed sample size that have a lot of units a_i that are strongly positively correlated may be very weakly robust. Systematic sampling, which gives very high or maximal possible eigenvalues, corresponds to such sampling design. The Yates-Grundy condition is satisfied in a number of

common sampling designs. This bound is reached in the Matched design. Indeed, the maximum eigenvalue of the matrix Δ given in (2.3) is equal to $\lambda_1 = 2\pi(1 - \pi) = 1/2$.

A small bound can also be identified by comparing unequal probability sampling without replacement with a multinomial design with parameters $\{p_k = \pi_k/n, k \in U\}$ and exponent n .

Result 6. *For the multinomial design, the largest eigenvalue of Δ^R is less than or equal to $\max_{k \in U} \pi_k$.*

Definition 2. A sampling design without replacement with inclusion probabilities π and variance operator Δ is said to be dominated by multinomial sampling with parameters $\{p_k = \pi_k/n, k \in U\}$ and the same sample size n if

$$\frac{\mathbf{u}^\top \Delta \mathbf{u}}{\mathbf{u}^\top \Delta^R \mathbf{u}} \leq 1, \text{ for all } \mathbf{u} \in R^N.$$

This definition also corresponds to a comparison with Δ^R in the Loewner order. Gabler (1981) proved that multinomial sampling dominates Sampford's sampling design (Sampford, 1967). Gabler (1984) gave a sufficient condition in order that a sampling design without replacement is dominated by multinomial sampling. Qualité (2008) proved that this condition is fulfilled for CPS.

Result 7. *For Sampford's design and CPS, the largest eigenvalue of Δ is less than or equal to*

$$\lambda_1 \leq \min \left\{ \max_{k \in U} \pi_k, \max_{k \in U} (1 - \pi_k) \right\}.$$

2.5.3 Case of balanced sampling

For sampling designs that select balanced samples, we deduce from Equation (2.4) that

$$\text{var}_a(\mathbf{X}^\top \mathbf{D}^{-1} \mathbf{a}) = \mathbf{X}^\top \mathbf{D}^{-1} \Delta \mathbf{D}^{-1} \mathbf{X} \approx 0.$$

In other words, balanced sampling designs will have a covariance matrix Δ with p eigenvalues that are equal or close to 0. If the inclusion probabilities are all equal and the design is perfectly balanced, the p covariates of \mathbf{X} generate the kernel of Δ , or at least a vector subspace from it. If we consider that the inclusion probabilities are fixed, then the trace of Δ is fixed. The trace is equal to the sum of the eigenvalues. Therefore, the use of a balanced sampling design means that some eigenvalues are equal or close to 0, but it will have to be compensated on the other non-zero eigenvalues of Δ . It is possible that a balanced sampling design can have a very large maximal eigenvalue due to this fact.

In balanced designs, it seems difficult to bound the eigenvalues more tightly than in Result 5. Indeed, balanced designs can respect the fixed sample size, but do not respect the Yates-Grundy condition. For algorithms that return balanced samples, the matrix Δ can usually only be estimated using simulations, which adds another layer of difficulty to bound the largest eigenvalue. Ideally, one would like balancing to cause approximately the same increase in all directions for the variables orthogonal to the balancing variables. Results from our simulation studies indicate that this idea is not far from reality. For balanced designs, we have an approximation for the largest eigenvalue given in the supplementary material. However, one may wonder if this approximation works in practice. For it to work, the increase in variance would have to be distributed more or less uniformly in all directions orthogonal to the balancing variables. Therefore, it is necessary to choose the balancing variables carefully and to avoid balancing on variables that are not correlated to the variables of interest.

The pivotal method respects the Yates-Grundy condition, so one can identify a bound for the largest eigenvalue of the variance matrix. Therefore, the pivotal method is a prudent way

to integrate auxiliary information into a design. Grafström and Lundström (2013) argue for the use of this method. We have here an additional argument of robustness. The Matched design is a special case of the pivotal method when $\pi = 1/2$. In this case, we can compute exactly the largest eigenvalue which is only twice the one of Bernoulli sampling. The risk is thus relatively contained.

2.6 Lower bound for the largest eigenvalues

In this section we find lower bounds for the largest eigenvalue.

Result 8. *For any design without replacement, $\lambda_1 \geq \max_{k \in U} \pi_k(1 - \pi_k)$.*

Result 9. *For any design without replacement with fixed sample size,*

$$\lambda_1 \geq \frac{N}{N-1} \max_{k \in U} \pi_k(1 - \pi_k).$$

Result 10. *Let \mathbf{X} be a $N \times p$ full rank matrix, $\mathbf{H} = \mathbf{X}(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top$, \mathbf{v}^k a vector of \mathbb{R}^N such that $v_k = 1$ and $v_\ell = 0$ if $\ell \neq k$ and $h_{kk} = \mathbf{v}^k \mathbf{H} \mathbf{v}^k$. For any design without replacement such that $\Delta \mathbf{X} = \mathbf{0}$, $\lambda_1 \geq \max_{k \in U} \pi_k(1 - \pi_k)/(1 - h_{kk})$.*

Unfortunately, the bound given by the Result 10 does not really seem applicable for balanced sampling designs. The reason is that an exactly balanced design does not exist in general because there is always a rounding problem which does not allow to obtain the exact balance. This bound works on stratified sampling designs because this is a particular case of balanced sampling that can be exactly balanced.

2.7 Robustification of the design by mixture

In this section, we try to introduce the notion of trade-off between balance/efficiency and robustness. This notion in experimental designs has been introduced by Harshaw et al. (2024).

Expression (2.2) shows that if the objective is to minimize the worst-case design-based variance, then the largest eigenvalue of Δ should simply be as small as possible. If we assume that the covariates $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^\top$ have an effect on the response, then the advantages of balanced sampling appear. Suppose that $y_k^C = \mathbf{x}_k \boldsymbol{\beta} + \varepsilon_k$, $y_k^T = \mathbf{x}_k \boldsymbol{\beta} + \tau + \varepsilon_k$ and that all inclusion probabilities are equal to π for all $k \in \{1, \dots, N\}$. We assume that the sample size is fixed in order to make the term τ in the following expression disappear. Using Expression (2.2), the design-based variance becomes

$$\text{var}_a(\hat{\tau}) = \left\{ \frac{1}{N\pi(1-\pi)} \right\}^2 \left\{ \boldsymbol{\beta}^\top \text{var}_a(\mathbf{X}^\top \mathbf{a}) \boldsymbol{\beta} + \boldsymbol{\varepsilon}^\top \Delta \boldsymbol{\varepsilon} + 2\boldsymbol{\beta}^\top \mathbf{X}^\top \Delta \boldsymbol{\varepsilon} \right\}, \quad (2.5)$$

where $\boldsymbol{\varepsilon}^\top = (\varepsilon_1, \dots, \varepsilon_N)$.

In order to minimise the variance under the design of $\hat{\tau}$, both the maximal eigenvalue of Δ and $\text{var}_a(\mathbf{X}^\top \mathbf{a})$ should be considered. We ignore the cross-term as both terms have an influence on it and usually the cross-term is relatively small. We consider that the maximal eigenvalue of Δ is the robustness of the sampling/experimental design and that $\text{var}_a(\mathbf{X}^\top \mathbf{a})$ is the term closely related to how balanced is the sampling or experimental design. Indeed, $\text{var}_a(\mathbf{X}^\top \mathbf{D}^{-1} \mathbf{a}) = \mathbf{X}^\top \mathbf{D}^{-1} \Delta \mathbf{D}^{-1} \mathbf{X} \approx 0$ and $\text{var}_a(\mathbf{X}^\top \mathbf{D}^{-1} \mathbf{a}) = \mathbf{X}^\top \Delta \mathbf{X} \approx 0$ when the inclusion probabilities are all equal. The terms $\boldsymbol{\beta}$ and $\boldsymbol{\varepsilon}$ are not known to the experimenter and for now, do not necessarily imply a model-based approach. However, the design-based variance of $\hat{\tau}$ given in expression (2.5) used the assumption that τ has the same effect on all the units. This is a pretty strong assumption that we already use for consistency with later notation. Harshaw

et al. (2024) obtain an expression similar to (2.5) without this assumption. They reach the same conclusion while having a purely design-based approach.

Harshaw et al. (2024) introduced the GS design that handles the trade-off between balance and robustness. The basic step of the GS design is very similar to the Cube method of Deville and Tillé (2004a). In the Cube method, at each step, the vector of inclusion probabilities is randomly changed in a direction that respects the balancing equations. However, Harshaw et al. (2024) uses the GS algorithm to balance an augmented covariate vector, which is a scaled concatenation of the covariates of the units and a unit-unique indicator variable. A tuning parameter ϕ allows to give more importance either to the robustness or to the balancing. It allows one to make a mixture of designs: a robust design like the Bernoulli design and a balanced design. When $\phi \in \{0, 1\}$, the GS algorithm will correspond either to a Bernoulli design or balanced design like the Cube method.

We propose another method which allows to mix two designs that is presented in Algorithm 1. This method is much simpler than the GS design. It requires only three lines of code to be implemented. Moreover, it allows us to mix any two designs, as long as at least one of the design is able to handle cases with unequal inclusion probabilities.

Algorithm 1 Simple method for mixing designs

- Select a sample \mathbf{a}_A with a design $p_A(\cdot)$, inclusion probabilities $\boldsymbol{\pi}$ and variance-covariance matrix $\boldsymbol{\Delta}_A$.
 - Compute $\boldsymbol{\pi}_B = \sqrt{\phi} \mathbf{a}_A + (1 - \sqrt{\phi})\boldsymbol{\pi}$.
 - Select a sample \mathbf{a} with a design $p_B(\cdot)$, inclusion probabilities $\boldsymbol{\pi}_B$ and variance-covariance matrix $\boldsymbol{\Delta}_B(\boldsymbol{\pi}_B)$.
-

Result 11. *With Algorithm 1:*

- (i) $E_a(\mathbf{a}) = \boldsymbol{\pi}$,
- (ii) $\boldsymbol{\Delta} = \text{var}_a(\mathbf{a}) = E_a\{\boldsymbol{\Delta}_B(\boldsymbol{\pi}_B)\} + \phi\boldsymbol{\Delta}_A$,
- (iii) $\frac{1}{1-\phi}\text{diag}\{E_a(\boldsymbol{\Delta}_B)\} = \text{diag}(\boldsymbol{\Delta}) = \text{diag}(\boldsymbol{\Delta}_A)$.

When mixing a design without replacement with a design with replacement, $p_B(\cdot)$ needs to be set as the design with replacement to avoid any error. Ideally, we would have that $E_a\{\boldsymbol{\Delta}_B(\boldsymbol{\pi}_B)\} \approx (1 - \phi)\boldsymbol{\Delta}_B(\boldsymbol{\pi})$ for better interpretability of $\boldsymbol{\Delta}$ and theoretical results. It would mean that $\boldsymbol{\Delta} \approx \phi\boldsymbol{\Delta}_A + (1 - \phi)\boldsymbol{\Delta}_B$. However, it is very difficult to derive analytically the term $E_a\{\boldsymbol{\Delta}_B(\boldsymbol{\pi}_B)\}$ and probably impossible with a mixture that uses a balanced sampling design. We carried out simulations of a mixture of designs where $p_A(\cdot)$ and $p_B(\cdot)$ correspond to, respectively, SRSWOR and a balanced sampling design based on the Cube method. In this case, $p_A(\cdot)$ corresponds to the robust design and $p_B(\cdot)$ corresponds to the balanced design. For $\boldsymbol{\Delta}_B(\boldsymbol{\pi}_B)$, the Cube method has been parametrized in a way to ensure that the inclusion probabilities are $\boldsymbol{\pi}_B$ while also having the propriety that $\mathbf{X}^\top \mathbf{D}^{-1} \mathbf{a} \approx \mathbf{X}^\top \mathbf{D}^{-1} \boldsymbol{\pi}_B$. The auxiliary variables are generated from a normal distribution, $\pi = 1/2$ and $N = 250$. The simulations, provided in the supplementary material, show that $E_a\{\boldsymbol{\Delta}_B(\boldsymbol{\pi}_B)\}$ is extremely close to $(1 - \phi)\boldsymbol{\Delta}_B(\boldsymbol{\pi})$.

Moreover if $\lambda_1, \lambda_{A1}, \lambda_{B1}$ are, respectively, the largest eigenvalues of $\boldsymbol{\Delta}, \boldsymbol{\Delta}_A$ and $E_a(\boldsymbol{\Delta}_B)/(1 - \phi)$, then

$$\lambda_1 \leq \phi\lambda_{A1} + (1 - \phi)\lambda_{B1}.$$

By taking a robust design for $p_A(\cdot)$ we can make the largest eigenvalue decrease. The parameter ϕ allows to set the robustness of the design. The method opens up the possibility of other

mixtures. For instance, we can mix a Matched design with a balanced design. In this way, we get a design that is not too risky, while combining two different uses of auxiliary information.

There is also another simple method to mix designs that we do not recommend. First, choose either the sampling design $p_A(\cdot)$ with probability ϕ and $p_B(\cdot)$ with probability $(1 - \phi)$. Then, select a sample with the chosen sampling design. This method creates a sampling design with variance $\phi\Delta_A + (1 - \phi)\Delta_B$, which is a good property. However, sampling designs are not uniquely defined by their variances. Using simulations, we clearly see that the distribution of the Horvitz-Thompson estimator of this method has a higher kurtosis than the distribution the Horvitz-Thompson estimator using Algorithm 1 when mixing the cube with SWSWOR even though their variances are similar. The consequence is that large deviations of the Horvitz-Thompson estimator are more likely to happen with this method. Additionally, this method cannot select new samples that are not contained in $p_A(\cdot)$ or $p_B(\cdot)$. Algorithm 1 creates new samples, unless some specific sampling designs are used.

2.8 Simulations with mixtures of designs

In order to evaluate the robustness of the designs, we performed simulations on a clinical trial data set. It contains information about 40 patients split into two groups before and after their corneal astigmatism surgery (Hayder, 2020). We are only interested in the variables measured before the surgeries. We retain five covariates: sex, eye, age, axis and topographic astigmatism. Two of those covariates are categorical. The size of the experiment is $N = 40$ and we chose equal inclusion probabilities $\pi = 1/2$ for all units, therefore $n = 20$.

We selected samples with four basic designs: CPS (SRSWOR when the inclusion probabilities equal), Matched design, Cube-Matched design, and systematic sampling. We also made simulations with mixture of designs using Algorithm 1 and the GS design. The simulations here do not include the Bernoulli design because it does not have fixed sample size. More details can be found in the supplement. For the Matched and Cube-Matched designs, the pairs have been constructed by discrete linear programming by minimizing the sum on the pairs of the squares of the Mahalanobis distances between the paired units. We define the squared Mahalanobis distance as follows: Let

$$\mathbf{t}_X = \mathbf{X}^\top \mathbf{1}, \hat{\mathbf{t}}_X = \mathbf{X}^\top \mathbf{D}^{-1} \mathbf{a}, \mathbf{V} = \frac{N}{N-1} \left(\frac{\mathbf{X}^\top \mathbf{X}}{N} - \frac{\mathbf{t}_X \mathbf{t}_X^\top}{N^2} \right),$$

the square of the Mahalanobis distance is

$$M = (\hat{\mathbf{t}}_X - \mathbf{t}_X)^\top \mathbf{V}^{-1} (\hat{\mathbf{t}}_X - \mathbf{t}_X).$$

This index assesses whether the samples are well balanced on the auxiliary variables.

The estimation of the largest eigenvalue by simulation is unstable, especially when the largest eigenvalues in a design are the same or similar. There are slight differences between the maximal eigenvalue estimated through simulation and the exact theoretical results. The maximal eigenvalue estimated by simulation will almost always be slightly inflated when all the inclusion probabilities are equal. Therefore, we performed a very large number of simulations (SIM=10,000,000).

Next, we ran the GS design of Harshaw et al. (2024) using the Julia implementation. We used a version of the algorithm that returns experimental designs with fixed size for each group. Unfortunately, we found that this algorithm does not provide comparable results depending on whether the variables have the same means or variances. We understand that this is due to the way the covariates are normalized in the GS algorithm. Therefore, we centred, reduced and made orthogonal all the variables using a singular value decomposition. Without this change of variables, the results obtained by the GS method are inconsistent.




Figure 2.1: Scatterplot of the estimated expectation of the squares of the Mahalanobis distances between the covariate Horvitz-Thompson estimators of the two groups by the largest eigenvalues. The numbers in brackets correspond to the tuning parameter ϕ which gives importance either to robustness or to balance depending on the value of ϕ . When the inclusion probabilities are equal, CPS is changed to SRSWOR.

Figure 2.1 contains the scatterplot of the estimated expectation of the squares of the Mahalanobis distances between the covariate Horvitz-Thompson estimators of the two groups by the largest eigenvalues. On the left side, one sees that the most robust is SRSWOR but it is not balanced. Not all the sampling or experimental designs we introduced in this article are included in the graph, as they reduce to SRSWOR or matching when all inclusion probabilities are equal to 0.5. Systematic sampling design lacks too much robustness to be shown in the graph.

The Cube method is a very good compromise between the two extremes as it is well balanced and relatively robust. The Matched method is less balanced and less robust than the Cube method. The Matched method can also have difficulties finding pairs that are sufficiently similar. However, one advantage of matching is that the samples or group generated from this method are relatively well-spread while samples generated from a balanced design like the Cube method could output samples that are not well-spread. Generating a mixture of design between the Cube method and the Matched method could give a good compromise in certain cases. Grafström

and Lundström (2013) describes advantages of well-spread samples. The designs resulting from mixtures with Algorithm 1 and the GS design approximately lie between the two designs from which they originate in the graph in function of their weights ϕ as expected. However, our method is simpler to carry out and has an interpretable output.

2.9 Conclusions

In this paper, we give lower and upper bounds of the largest eigenvalue for sample and experimental designs. For Poisson, Bernoulli and matched designs, the exact value of largest eigenvalue is given. We have also seen that the Yates-Grundy condition guarantees a relatively small upper bound of the largest eigenvalue. Our research indicates that there is a strong correlation between high/maximum entropy sampling and robust designs.

Balanced designs generated by algorithms such as the cube method or the GS design do not seem to create undue risks as systematic design does. However, it is important to avoid using balancing variables that are not correlated with the variable of interest.

Moreover, we proposed a simple method to mix designs that offers more possibilities than the GS design proposed by Harshaw et al. (2024). Simulations also show that our method creates mixtures of designs that are easily interpretable. If a design is too complex, it is always possible to estimate matrix Δ by means of simulations. This enables efficient risk assessment and correct inference.

Finally, we have doubts about the interest of mixing designs as a method to reduce the variance of $\hat{\tau}$ in a “robust” way. A simple calculation indicates that the optimal design is likely to not result from a mixture of designs if the hypotheses are correct. If we are convinced that the auxiliary variables are correlated to the variable of interest, then one should balance. If we are convinced that the auxiliary variables are uncorrelated with the variable of interest or simply unavailable, then one should use a robust design like SRSWOR or Bernoulli. Mixtures of designs can be used in the case where several variables of interest have to be estimated. In this case, a compromise between several designs must be found, each optimised for a particular variable.

2.10 Supplementary material for Risk Minimization Using Robust Experimental or Sampling Designs and Mixture of Designs

In the supplementary material, we include all the proofs and we add other approximations for the largest eigenvalue of different sampling designs. We also discuss the interest of mixing designs under model- and design-based approach for experimental designs. We show that the interest of mixing design is relatively limited under our hypotheses. Finally, we give additional information regarding simulations.

2.10.1 Proofs

Result 1. *There exist no pair of design $p_1(\cdot)$ and $p_2(\cdot)$ without replacement, such that $p_2(\cdot)$ is uniformly better than $p_1(\cdot)$ with the same first-order inclusion probability.*

Proof. (by contradiction) Suppose there is a design $p_1(\cdot)$ uniformly better than a design $p_2(\cdot)$ with the same first-order inclusion probabilities. In this case, $\mathbf{u}^\top (\mathbf{\Delta}_2 - \mathbf{\Delta}_1) \mathbf{u} \geq 0$, for all $\mathbf{u} \in \mathbb{R}^N$. This implies that the array $(\mathbf{\Delta}_2 - \mathbf{\Delta}_1)$ is positive semi-definite. Since the trace of $(\mathbf{\Delta}_2 - \mathbf{\Delta}_1)$ is zero and all eigenvalues of $(\mathbf{\Delta}_2 - \mathbf{\Delta}_1)$ are greater than or equal to zero, we deduce that all those eigenvalues are equal to zero. This is in contradiction with the fact that there is at least one vector \mathbf{y} such as $\mathbf{y}^\top \mathbf{\Delta}_1 \mathbf{y} < \mathbf{y}^\top \mathbf{\Delta}_2 \mathbf{y}$. \square

Result 2. *Let $p(\cdot)$ be a sampling design of fixed sample size n from a population U of size N , then $\mathbf{\Delta}$ has a null eigenvalue associated to the vector of ones $\mathbf{1} \in \mathbb{R}^N$.*

Proof. We have $\mathbf{\Delta} \mathbf{1} = E_a(\mathbf{a} \mathbf{a}^\top \mathbf{1}) - E_a(\mathbf{a}) E_a(\mathbf{a}^\top \mathbf{1}) = E_a(\mathbf{a} n) - E_a(\mathbf{a}) n = \mathbf{0}$. \square

Result 3. *Let $p(\cdot)$ be a sampling design of fixed sample size n from a population U of size N and with equal $\pi_k = \pi, k \in U$, then $\mathbf{1}$ is an eigenvector of $\mathbf{\Pi}$ and $\mathbf{\Delta}$. For any other centred eigenvector \mathbf{u} of $\mathbf{\Delta}$, \mathbf{u} is also an eigenvector of $\mathbf{\Pi}$ with the same eigenvalue.*

Proof. Since $\mathbf{1}$ is the eigenvector of $\mathbf{\Delta}$ associated with the null eigenvalues, all the other eigenvectors are orthogonal to $\mathbf{1}$ and are thus centred. We have $\mathbf{\Delta} \mathbf{1} = \mathbf{\Pi} \mathbf{1} - \boldsymbol{\pi} \boldsymbol{\pi}^\top \mathbf{1} = \mathbf{\Pi} \mathbf{1} - \boldsymbol{\pi} n = \mathbf{0}$. Thus, $\mathbf{\Pi} \mathbf{1} = \boldsymbol{\pi} n$ and $\mathbf{1}$ is an eigenvector of $\mathbf{\Pi}$ associated with the eigenvalue $n\pi$. If \mathbf{u} is a centred eigenvector of $\mathbf{\Delta}$, $\mathbf{\Delta} \mathbf{u} = \lambda \mathbf{u}$, thus $\mathbf{\Pi} \mathbf{u} - \boldsymbol{\pi} \boldsymbol{\pi}^\top \mathbf{u} = \lambda \mathbf{u}$. Since $\boldsymbol{\pi}^\top \mathbf{u} = 0$, vector \mathbf{u} is also an eigenvector of $\mathbf{\Pi}$ with the same eigenvalue as for $\mathbf{\Delta}$. \square

Result 4. *For any sampling design on a finite population U of size N , the largest eigenvalue of $\mathbf{\Delta}$ is smaller than or equal to $\max_{k \in U} (\sum_{\ell \in U} \pi_{k\ell})$. If the sampling design has a fixed sample size n , then the bound can be expressed as $n \max_{k \in U} (\pi_k)$.*

Proof. Using the Perron-Frobenius inequality (Horn and Johnson, 2013), we know that the largest eigenvalue of $\mathbf{\Pi}$ is smaller than or equal to

$$\max_{k \in U} \left(\sum_{\ell \in U} \pi_{k\ell} \right).$$

The matrix $\boldsymbol{\pi} \boldsymbol{\pi}^\top$ is a symmetric matrix of rank 1, whose only non-zero eigenvalue is equal to $\sum_{k \in U} \pi_k^2$ with the eigenvector $\boldsymbol{\pi}$. The largest eigenvalue of $-\boldsymbol{\pi} \boldsymbol{\pi}^\top$ is thus 0. Each eigenvalue of $\mathbf{\Delta}$ is non-negative, because it is a covariance matrix. Both $\mathbf{\Pi}$ and $\boldsymbol{\pi} \boldsymbol{\pi}^\top$ are Hermitian matrices. Therefore, from Weyl's inequality (Horn and Johnson, 2013), we deduce that the largest eigenvalue of $\mathbf{\Delta}$, denoted by λ_1 , is smaller than or equal to $\max_{k \in U} (\sum_{\ell \in U} \pi_{k\ell})$, i.e.,

$$\lambda_1 \leq \max_{k \in U} \left(\sum_{\ell \in U} \pi_{k\ell} \right).$$

If the sampling design has a fixed sample size n , then

$$\max_{k \in U} \left(\sum_{\ell \in U} \pi_{k\ell} \right) = \max_{k \in U} \left\{ \mathbb{E}_a \left(a_k \sum_{\ell \in U} a_\ell \right) \right\} = n \max_{k \in U} (\pi_k).$$

□

Corollary 1. *Let Δ be the variance matrix of a sampling design with fixed sample size n and equal $\pi = n/N$, then the largest eigenvalue of Δ is smaller than or equal to $n\pi = n^2/N$.*

Proof. The corollary is an application of Result 4. □

Corollary 2. *Let Δ be the variance matrix of a sampling design without replacement with fixed sample size n and equal $\pi = n/N$, then the largest eigenvalue of Δ is smaller than or equal to $(N - n)(1 - \pi) = (N - n)^2/N$.*

Proof. The random sample \mathbf{a} and $\mathbf{1} - \mathbf{a}$ have the same variance matrix Δ . The complementary sample is selected with inclusion probabilities $1 - \pi = (N - n)/N$ and fixed sample size $N - n$. Using corollary 1, we obtain that the maximal eigenvalue is $(N - n)(1 - \pi) = (N - n)^2/N$. □

Result 5. *For any sampling design on a finite population U of size N , the largest eigenvalue of Δ is smaller than or equal to $\max_{k \in U} \left(\sum_{\ell \in U} |\Delta_{k\ell}| \right)$.*

Proof. The proof is a simple corollary of the Gershgorin circle theorem (Horn and Johnson, 2013). However, the result can also be derived from Theorem 4.1 of Dol et al. (1996). □

Corollary 3. *For any sampling design without replacement that satisfies the Yates-Grundy condition and has fixed sample size n , the largest eigenvalue of Δ is smaller than or equal to $2 \max_{k \in U} \{ \pi_k(1 - \pi_k) \}$.*

Proof. We know that $\Delta_{kk} = \pi_k(1 - \pi_k)$ for all $i \in 1, \dots, N$. Every off-diagonal entry of the matrix Δ is negative because the selections of the units are all negatively correlated between each other due to the Yates-Grundy condition. Another property of Δ is that the sum of every row is equal to 0 due to fixed sample size. This is deduced from Result 2. We obtain that $\pi_k(1 - \pi_k) = \sum_{\ell \in U | \ell \neq k} \Delta_{k\ell}$ for all $k \in U$. Thus,

$$\begin{aligned} \max_{k \in U} \left(\sum_{\ell \in U} |\Delta_{k\ell}| \right) &= 2 \max_{k \in U} \left\{ \pi_k(1 - \pi_k) + \sum_{\ell \in U | \ell \neq k} \Delta_{k\ell} \right\} \\ &= 2 \max_{k \in U} \{ \pi_k(1 - \pi_k) \}. \end{aligned}$$

The result follows from Result 5. □

Corollary 4. *For any sampling design with a fixed sample size n on a finite population U of size N , the largest eigenvalue of Δ is smaller than or equal to $2 \max_{k \in U} \left(\sum_{\ell \in U | \Delta_{k\ell} > 0} \Delta_{k\ell} \right)$.*

Proof. The proof is the same as in Corollary 3. □

Result 6. *For the multinomial design, the largest eigenvalue of Δ^R is less than or equal to $\max_{k \in U} \pi_k$.*

Proof. Define $\mathbf{D} = \text{diag}(np_1, \dots, np_N)$. Matrix $\mathbf{D}^{-1/2} \Delta^R \mathbf{D}^{-1/2}$ is idempotent. All its eigenvalues are equal to 1 or 0. Thus

$$\frac{\mathbf{u}^\top \mathbf{D}^{-1/2} \Delta^R \mathbf{D}^{-1/2} \mathbf{u}}{\mathbf{u}^\top \mathbf{u}} \leq 1, \text{ for all } \mathbf{u} \in \mathbb{R}^N.$$

If we set $\mathbf{y} = \mathbf{D}^{-1/2}\mathbf{u}$

$$\frac{\mathbf{y}^\top \boldsymbol{\Delta}^R \mathbf{y}}{\mathbf{y}^\top \mathbf{y}} \leq \frac{\mathbf{y}^\top \mathbf{D} \mathbf{y}}{\mathbf{y}^\top \mathbf{y}} \leq \max_{k \in U} \mu_k, \text{ for all } \mathbf{y} \in \mathbb{R}^N. \quad (2.6)$$

□

Result 7. For Sampford's design and CPS, the largest eigenvalue of $\boldsymbol{\Delta}$ is less than or equal to

$$\lambda_1 \leq \min \left\{ \max_{k \in U} \pi_k, \max_{k \in U} (1 - \pi_k) \right\}.$$

Proof. Since these designs are dominated by multinomial sampling, the largest eigenvalue of $\boldsymbol{\Delta}$ is less than or equal to $\max_{k \in U} \pi_k$. Moreover, $\boldsymbol{\Delta}$ is also the variance operator matrix of the complementary design, where the sample is $S^c = U \setminus S$, the inclusion probabilities are $\pi_k^c = 1 - \pi_k$ and $a_k^c = 1 - a_k$, for all $k \in U$. Thus, the largest eigenvalue of $\boldsymbol{\Delta}$ is also less or equal to $\max_{k \in U} 1 - \pi_k$. □

Result 8. For any design without replacement, $\lambda_1 \geq \max_{k \in U} \pi_k(1 - \pi_k)$.

Proof. Let \mathbf{v}^k be a vector of \mathbb{R}^N such that $v_k = 1$ and $v_\ell = 0$ if $\ell \neq k$. Then $\mathbf{v}^{k\top} \boldsymbol{\Delta} \mathbf{v}^k / (\mathbf{v}^{k\top} \mathbf{v}^k) = \pi_k(1 - \pi_k)$. Thus, $\lambda_1 \geq \max_{k \in U} \pi_k(1 - \pi_k)$. □

Result 9. For any design without replacement with fixed sample size,

$$\lambda_1 \geq \frac{N}{N-1} \max_{k \in U} \pi_k(1 - \pi_k).$$

Proof. Let \mathbf{v}^k a vector of \mathbb{R}^N such that $v_k = 1$ and $v_\ell = 0$ if $\ell \neq k$ and $\mathbf{1}$ vector of N ones. Define

$$\mathbf{e}^k = \left(\mathbf{v}^k - \frac{\mathbf{1}}{N} \right) \sqrt{\frac{N}{N-1}}.$$

Vector \mathbf{e}^k is normed. Then

$$\begin{aligned} \frac{\mathbf{e}^{k\top} \boldsymbol{\Delta} \mathbf{e}^k}{\mathbf{e}^{k\top} \mathbf{e}^k} &= \left(\mathbf{v}^k - \frac{\mathbf{1}}{N} \right)^\top \boldsymbol{\Delta} \left(\mathbf{v}^k - \frac{\mathbf{1}}{N} \right) \frac{N}{N-1} \\ &= \left(\mathbf{v}^{k\top} \boldsymbol{\Delta} \mathbf{v}^k + \frac{\mathbf{1}^\top \boldsymbol{\Delta} \mathbf{1}}{N} - 2\mathbf{v}^{k\top} \boldsymbol{\Delta} \frac{\mathbf{1}}{N} \right) \frac{N}{N-1}. \end{aligned}$$

Since $\boldsymbol{\Delta} \mathbf{1} = \mathbf{0}$,

$$\frac{\mathbf{v}^{k\top} \boldsymbol{\Delta} \mathbf{v}^k}{\mathbf{v}^{k\top} \mathbf{v}^k} = \pi_k(1 - \pi_k) \frac{N}{N-1}.$$

Thus, $\lambda_1 \geq \max_{k \in U} \pi_k(1 - \pi_k)N/(N-1)$. □

Result 10. Let \mathbf{X} be a $N \times p$ non-singular matrix, $\mathbf{H} = \mathbf{X}(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top$, \mathbf{v}^k a vector of \mathbb{R}^N such that $v_k = 1$ and $v_\ell = 0$ if $\ell \neq k$ and $h_{kk} = \mathbf{v}^{k\top} \mathbf{H} \mathbf{v}^k$. For any design without replacement such that $\boldsymbol{\Delta} \mathbf{X} = \mathbf{0}$, $\lambda_1 \geq \max_{k \in U} \pi_k(1 - \pi_k)/(1 - h_{kk})$.

Proof. Define

$$\mathbf{e}^k = \left(\mathbf{v}^k - \mathbf{H} \mathbf{v}^k \right) \sqrt{\frac{1}{1 - h_{kk}}}.$$

Vector \mathbf{e}^k is normed. Then

$$\begin{aligned}\frac{\mathbf{e}^{k\top} \Delta \mathbf{e}}{\mathbf{e}^{k\top} \mathbf{e}^k} &= (\mathbf{v}^k - \mathbf{H}\mathbf{v}^k)^\top \Delta (\mathbf{v}^k - \mathbf{H}\mathbf{v}^k) \frac{1}{1 - h_{kk}} \\ &= (\mathbf{v}^{k\top} \Delta \mathbf{v}^k + \mathbf{v}^{k\top} \mathbf{H} \Delta \mathbf{H} \mathbf{v}^k - 2\mathbf{v}^{k\top} \Delta \mathbf{H} \mathbf{v}^k) \frac{1}{1 - h_{kk}}.\end{aligned}$$

Since $\Delta \mathbf{H} = \mathbf{0}$,

$$\frac{\mathbf{e}^{k\top} \Delta \mathbf{e}}{\mathbf{e}^{k\top} \mathbf{e}^k} = \mathbf{v}^{k\top} \Delta \mathbf{v}^k \frac{1}{1 - h_{kk}} = \pi_k (1 - \pi_k) \frac{1}{1 - h_{kk}}.$$

Thus $\lambda_1 \geq \max_{k \in U} \pi_k (1 - \pi_k) \frac{1}{1 - h_{kk}}$. □

Result 11. *With Algorithm 1:*

- (i) $\mathbb{E}_a(\mathbf{a}) = \boldsymbol{\pi}$,
- (ii) $\Delta = \text{var}_a(\mathbf{a}) = \mathbb{E}_a\{\Delta_B(\boldsymbol{\pi}_B)\} + \phi \Delta_A$,
- (iii) $\frac{1}{1-\phi} \text{diag}\{\mathbb{E}_a(\Delta_B)\} = \text{diag}(\Delta) = \text{diag}(\Delta_A)$.

Proof.

- (i) The final inclusion probabilities are

$$\mathbb{E}_a(\mathbf{a}) = \mathbb{E}_a \mathbb{E}_a(\mathbf{a}|\mathbf{a}_A) = \mathbb{E}_a\{\sqrt{\phi} \mathbf{a}_A + (1 - \sqrt{\phi})\boldsymbol{\pi}\} = \boldsymbol{\pi}.$$

- (ii) The final variance matrix is:

$$\begin{aligned}\Delta &= \text{var}_a(\mathbf{a}) = \mathbb{E}_a \text{var}_a(\mathbf{a}|\mathbf{a}_A) + \text{var}_a \mathbb{E}_a(\mathbf{a}|\mathbf{a}_A) \\ &= \mathbb{E}_a\{\Delta_B(\boldsymbol{\pi}_B)\} + \text{var}_a\{\sqrt{\phi} \mathbf{a}_A + (1 - \sqrt{\phi})\boldsymbol{\pi}\} \\ &= \mathbb{E}_a\{\Delta_B(\boldsymbol{\pi}_B)\} + \phi \Delta_A.\end{aligned}\tag{2.7}$$

- (iii) Vectors \mathbf{a} and \mathbf{a}_A have the same expectation. Thus $\text{diag}(\Delta) = \text{diag}(\Delta_A)$. From Equality (2.7), we have

$$\text{diag}(\Delta) = \text{diag}\{\mathbb{E}_a(\Delta_B)\} + \phi \text{diag}(\Delta_A),$$

and thus

$$\frac{1}{1 - \phi} \text{diag}\{\mathbb{E}_a(\Delta_B)\} = \text{diag}(\Delta).$$

□

2.10.2 Approximations for the largest eigenvalue

We have identified bounds for the largest eigenvalues. We can construct approximations using heuristic reasoning. Consider the matrix

$$\begin{aligned}\mathbf{M} &= \text{diag}\left(\frac{\pi_1(1 - \pi_1)}{np_1(1 - p_1)}, \dots, \frac{\pi_k(1 - \pi_k)}{np_k(1 - p_k)}, \dots, \frac{\pi_N(1 - \pi_N)}{np_N(1 - p_N)}\right) \\ &= \text{diag}\left(\frac{1 - \pi_1}{1 - p_1}, \dots, \frac{1 - \pi_k}{1 - p_k}, \dots, \frac{1 - \pi_N}{1 - p_N}\right).\end{aligned}$$

Then, $\mathbf{M}^{1/2} \Delta^R \mathbf{M}^{1/2}$ has the same diagonal as Δ . We expect that $\mathbf{M}^{1/2} \Delta^R \mathbf{M}^{1/2}$ is close to CPS because Multinomial sampling has fixed sample size and $\mathbf{M}^{1/2} \Delta^R \mathbf{M}^{1/2}$ has relatively well

spread off-diagonal values in its variance matrix, which we expect to be related to maximum entropy sampling. In this case, we can use the largest eigenvalue of $\mathbf{M}^{1/2}\mathbf{\Delta}^R\mathbf{M}^{1/2}$ to bound the largest eigenvalues of the CPS and Sampford sampling designs. Matrix $\mathbf{D}^{-1/2}\mathbf{\Delta}^R\mathbf{D}^{-1/2}$ is idempotent. All its eigenvalues are equal to 1 or 0. Thus

$$\frac{\mathbf{u}^\top \mathbf{D}^{-1/2} \mathbf{\Delta}^R \mathbf{D}^{-1/2} \mathbf{u}}{\mathbf{u}^\top \mathbf{u}} \leq 1, \text{ for all } \mathbf{u} \in \mathbb{R}^N.$$

If we set $\mathbf{y} = \mathbf{D}^{-1/2} \mathbf{u}$

$$\frac{\mathbf{y}^\top \mathbf{\Delta}^R \mathbf{y}}{\mathbf{y}^\top \mathbf{y}} \leq \frac{\mathbf{y}^\top \mathbf{D} \mathbf{y}}{\mathbf{y}^\top \mathbf{y}}, \text{ for all } \mathbf{y} \in \mathbb{R}^N. \quad (2.8)$$

If we set $\mathbf{v} \mathbf{M}^{1/2} = \mathbf{y}$, we have

$$\begin{aligned} \frac{\mathbf{v}^\top \mathbf{M}^{1/2} \mathbf{\Delta}^R \mathbf{M}^{1/2} \mathbf{v}}{\mathbf{v}^\top \mathbf{v}} &\leq \frac{\mathbf{v}^\top \mathbf{M}^{1/2} \mathbf{D} \mathbf{M}^{1/2} \mathbf{v}}{\mathbf{v}^\top \mathbf{v}} \\ &\leq \max \text{diag}(\mathbf{M}^{1/2} \mathbf{D} \mathbf{M}^{1/2}) = \max_{k \in U} \frac{\pi_k(1 - \pi_k)}{1 - \pi_k/n}. \end{aligned}$$

For Sampford and CPS, we can consider that λ_1 is very probably lower than

$$\max_{k \in U} \frac{\pi_k(1 - \pi_k)}{1 - \pi_k/n}. \quad (2.9)$$

Since the same reasoning can be done for the complementary design, we also have that λ_1 is very probably lower than

$$\max_{k \in U} \frac{\pi_k(1 - \pi_k)}{1 - (1 - \pi_k)/(N - n)}. \quad (2.10)$$

The approximation is exact when all the inclusion probabilities are equal as both CPS and Sampford sampling reduces to SRSWOR in this case. Using simulations, we notice that in practice those approximations should be used as very close approximations of λ_1 instead of approximate bounds for λ_1 in CPS and Sampford sampling designs.

Another approximation can be derived from Deville and Tillé (2005) who proposed an approximation of the variance for sampling designs with a large entropy and fixed sample size. CPS is a particular case of this approximation. Let $\mathbf{b} = (b_1, \dots, b_N)$ be a vector of weights and $\mathbf{B} = \text{diag}(\mathbf{b})$. Several possibilities exist and will be discussed later. Define an approximation of $\mathbf{\Delta}$

$$\mathbf{A} = \mathbf{B} - \frac{\mathbf{b} \mathbf{b}^\top}{\mathbf{b}^\top \mathbf{1}}.$$

Two main approximations are proposed. In the first approximation, we construct the b_k in such a way that the diagonal of \mathbf{A} is the same as that of $\mathbf{\Delta}$. The b_k must then be the solution of the nonlinear equation system

$$b_k - \frac{b_k^2}{\sum_{\ell \in U} b_\ell} = \pi_k(1 - \pi_k), \text{ for all } k \in U.$$

This system can be solved by a fixed-point method. A second possibility is to simply take

$$b_k = \frac{N}{N-1} \pi_k(1 - \pi_k).$$

For any value of b_k , matrix $\mathbf{B}^{-1/2} \mathbf{A} \mathbf{B}^{-1/2}$ is idempotent. Thus

$$\frac{\mathbf{u}^\top \mathbf{B}^{-1/2} \mathbf{A} \mathbf{B}^{-1/2} \mathbf{u}}{\mathbf{u}^\top \mathbf{u}} \leq 1, \text{ for all } k \in U.$$

If we pose $\mathbf{y}^\top = \mathbf{u}^\top \mathbf{B}^{-1/2}$, we obtain

$$\frac{\mathbf{y}^\top \mathbf{A} \mathbf{y}}{\mathbf{y}^\top \mathbf{y}} \leq \frac{\mathbf{y}^\top \mathbf{B} \mathbf{y}}{\mathbf{y}^\top \mathbf{y}} \leq \max_{k \in U} b_k.$$

Thus, we deduce that λ_1 is probably lower or equal to

$$\frac{N}{N-1} \max_{k \in U} \pi_k (1 - \pi_k),$$

which is close to the previous approximation for CPS and Sampford sampling. Not only that, but this approximation actually also corresponds to the minimal largest eigenvalue possible for fixed sample size sampling designs. It means that actually λ_1 will not be lower than this bound. But it also means that maximal entropy sampling is clearly close to the most robust sampling design. It also suggest that there is a link between the entropy of the sampling design and the robustness of the sampling design.

For balanced sampling, the development of Deville and Tillé (2005) implies that the approximation of the largest eigenvalues is

$$\lambda_1 \approx \frac{N}{N-p} \max_{k \in U} \pi_k (1 - \pi_k). \quad (2.11)$$

Given that using balanced sampling means that p eigenvalue are close or equal to 0 and that the trace is equal to $N\pi(1-\pi)$ in sampling designs with equal inclusion probabilities, the bound in (2.11) implies that the remaining non-null eigenvalues are very close or equal to each other. It becomes clear that this approximation is optimistic, but can be compensated by the fact that having a perfectly balanced sampling design is rare in practice. For a Cube-Matched design, each stratum can be considered as an additional balancing constraint. Therefore, a reasonable approximation for the largest eigenvalue is

$$\lambda_1 \approx \frac{N}{N - N/2 - p + 1} \max_{k \in U} \pi_k (1 - \pi_k).$$

2.11 Model- and design-based approach in experimental designs

Until now, the approach to the problem was design-based and only the variance due to the design was considered. However, for designs of experiments, it often makes sense to assume a model for the response that includes the covariates. Equation (2.5) gives the variance under the design when $y_k^C = \mathbf{x}_k \boldsymbol{\beta} + \varepsilon_k$, $y_k^T = \mathbf{x}_k \boldsymbol{\beta} + \tau + \varepsilon_k$ and all inclusion probabilities are equal to π . It also assumes that the sampling design used has fixed sample size. We obtain a model by keeping the same assumptions and adding the assumption that ε_k are random and independent variables of variance σ^2 and null expectation for all $k \in \{1, \dots, N\}$. Tillé (2022) showed that the total variance, which includes the variance due to the model, is equal to

$$\text{var}_T(\hat{\tau}) = \boldsymbol{\beta}^\top \text{var}_a(\bar{\mathbf{x}}_T - \bar{\mathbf{x}}_C) \boldsymbol{\beta} + \frac{\sigma^2}{N\pi(1-\pi)}, \quad (2.12)$$

where

$$\bar{\mathbf{x}}_T = \frac{1}{N\pi} \sum_{k \in U} a_k \mathbf{x}_k$$

and

$$\bar{\mathbf{x}}_C = \frac{1}{N(1-\pi)} \sum_{k \in U} (1 - a_k) \mathbf{x}_k$$

are the covariate means of the treatment group and the control group respectively. In Equation (2.12), the term related to the robustness of the design does not appear. On the other hand, the term with $\text{var}_a(\bar{\mathbf{x}}_T - \bar{\mathbf{x}}_C)$ suggests that balanced sampling designs will reduce the total variance of $\hat{\tau}$. Indeed, $\text{var}_a(\bar{\mathbf{x}}_T - \bar{\mathbf{x}}_C) \approx \mathbf{0}$ when a balanced sampling design is used. In conclusion, for the previous model, the notion of design robustness disappears and only balancing the covariates can help reduce the total variance. This is not necessarily the case with other more complex models.

We also investigate the design-based variance, under the worst-case scenario, of $\hat{\tau}$ given in (2.5) and try to minimize it. We expect that having a lower design-based variance might help reducing model dependence for causal inference. In Equation (2.5), if the covariates \mathbf{X} are highly correlated with the response variable, we expect to find $\boldsymbol{\beta}$ such that it explains a large amount of the variance of the y_k . In this case, the term $\|\mathbf{X}\boldsymbol{\beta}\|_2^2 = \boldsymbol{\beta}^\top \mathbf{X}^\top \mathbf{X} \boldsymbol{\beta}$ is expected to be much larger than $\|\boldsymbol{\varepsilon}\|_2^2 = \boldsymbol{\varepsilon}^\top \boldsymbol{\varepsilon}$. Therefore, it becomes reasonable to put more emphasis on balancing the covariates than on creating a robust design. The reverse is true when the covariates are weakly correlated with the response variable. The experimenter should not know how well his covariates are correlated with the response variable before the experiment. However, it is possible to imagine a case where similar previous experiments are carried out and the strength of the relationship between the response variable and the covariates can be estimated.

Using the Algorithm 1 introduced in Section 2.8, we try to give a recommendation of a mixture of a design depending on the expected value of $\|\boldsymbol{\varepsilon}\|_2^2$ and $\|\mathbf{X}\boldsymbol{\beta}\|_2^2$. We do not assume that the experimenter knows the true value of $\boldsymbol{\beta}$ so we are going to develop the formulae with estimates. The following reasoning should work with any estimate $\hat{\boldsymbol{\beta}}$ of $\boldsymbol{\beta}$ even though the least-squares estimate would usually be used. Let $\mathbf{e} = \mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}$, where \mathbf{y} are the response variables. We assume, using the result of the simulations of Algorithm 1 in Section 2.7, that we can create a design which is a weighted sum of two designs:

$$\boldsymbol{\Delta} = \phi \boldsymbol{\Delta}_A + (1 - \phi) \boldsymbol{\Delta}_B$$

We also assume that $\boldsymbol{\Delta}_A$ is covariance matrix of a perfectly balanced design: $\boldsymbol{\Delta}_A \mathbf{X} = \mathbf{0}$. For $\boldsymbol{\Delta}_B$, we will use the covariance matrix of a SRSWOR design. We define λ_A and λ_B as the largest eigenvalues of $\boldsymbol{\Delta}_A$ and $\boldsymbol{\Delta}_B$, respectively. We obtain

$$\{N\pi(1 - \pi)\}^2 \text{var}_a(\hat{\tau}) = \left\{ \hat{\boldsymbol{\beta}}^\top \text{var}_a(\mathbf{X}^\top \mathbf{a}) \hat{\boldsymbol{\beta}} + \mathbf{e}^\top \boldsymbol{\Delta} \mathbf{e} + 2\hat{\boldsymbol{\beta}}^\top \mathbf{X}^\top \boldsymbol{\Delta} \mathbf{e} \right\} \quad (2.13)$$

We assume that we can ignore the cross-term in (2.13), as both \mathbf{e} and $\mathbf{X}\hat{\boldsymbol{\beta}}$ equally contribute to it and we expect this term to be small as \mathbf{e} and $\mathbf{X}\hat{\boldsymbol{\beta}}$ are orthogonal if $\boldsymbol{\beta}$ is fitted using least squares. Then, we obtain

$$\hat{\boldsymbol{\beta}}^\top \text{var}_a(\mathbf{X}^\top \mathbf{a}) \hat{\boldsymbol{\beta}} + \mathbf{e}^\top \boldsymbol{\Delta} \mathbf{e} \leq \{(1 - \phi)\lambda_B\} \|\mathbf{X}\hat{\boldsymbol{\beta}}\|_2^2 + \{\phi\lambda_A + (1 - \phi)\lambda_B\} \|\mathbf{e}\|_2^2. \quad (2.14)$$

We differentiate the bound in (2.14) by ϕ . We obtain

$$-\lambda_B \|\mathbf{X}\hat{\boldsymbol{\beta}}\|_2^2 + (\lambda_A - \lambda_B) \|\mathbf{e}\|_2^2.$$

The maximal eigenvalue of SRSWOR design is $N/\pi(1 - \pi)(N - 1)$. Given that (2.13) is linear in ϕ and that $\phi \in [0, 1]$, we obtain that

$$\phi_{opti} = \begin{cases} 1 & \text{if } \frac{\pi(1-\pi)(N-1)\lambda_A}{N} - 1 < \frac{\|\mathbf{X}\hat{\boldsymbol{\beta}}\|_2^2}{\|\mathbf{e}\|_2^2}, \\ 0 & \text{otherwise.} \end{cases} \quad (2.15)$$

is the optimal value of ϕ .

When it is not available, the largest eigenvalue λ_A can be estimated through simulations. We

recall that the recommended value of ϕ in (2.15) was calculated by ignoring the cross term and assuming that Δ_A was perfectly balanced, which is an assumption that is rarely true in practice. However, it indicates that doing mixtures of designs is not useful to reduce the design-based variance in (2.5) in worst-case scenarios. Choosing either a SRSWOR or a balanced sampling design depending on the strength of the relation between the response and the covariates seem to be more efficient than using a mixture in this case. If there is a model misspecification, using a mixture that puts more weight on the robust design might be preferable.

2.11.1 Complete simulations results

In this section, exact values of the simulations are shown on the tables for basic designs, and mixture of designs. The approximations for balanced designs are also compared with their estimated value.

The bounds for the largest eigenvalue of variance matrix show us that Poisson and Bernoulli sampling are the most robust. However, they do not incorporate any auxiliary information. It is almost always preferable to impose the fixed sample size. If the inclusion probabilities are equal, then SRSWOR is the most robust, because all non-zero eigenvalues of the variance matrix are equal. With unequal probabilities and fixed sample size, the CPS and Sampford designs should be used because the largest eigenvalue of the variance matrix can be more tightly bounded.

For balanced sampling, we cannot efficiently bound the eigenvalues. We know that each balancing variable increases the dispersion a little for the variables that are orthogonal to it. It is therefore necessary to choose the balancing variables carefully and to avoid balancing on variables that are not correlated to the variables of interest. For balanced designs, we have an approximation for the largest eigenvalue found in previous sections. However, one can wonder if this approximation works in practice. For it to work, the increase in variance would have to be distributed more or less uniformly in all directions orthogonal to the balancing variables.

Table 2.2: Simulations: $N = 40$, $n = 20$ and equal inclusion probabilities. Largest eigenvalue of Δ obtained by simulations and approximation or true value for this eigenvalue. Squares of the Mahalanobis distances between the estimators and the true totals of the auxiliary variables.

	Simulations	Approximations or true value	Mahalanobis
Bernoulli	0.253	$\lambda_1 = 1/4 = 0.25$	1718.846
SRSWOR	0.259	$\lambda_1 = N/\{4(N-1)\} = 0.2564$	199.875
Matched	0.504	$\lambda_1 = N/\{4(N-N/2)\} = 0.5$	50.012
Cube	0.351	$\lambda_1 \approx N/\{4(N-p)\} = 0.2857$	17.284
Cube-Matched	1.058	$\lambda_1 \approx N/\{4(N-N/2-p+1)\} = 0.625$	13.766
Systematic	10.000	$\lambda_1 = N/4 = 10$	264.408

The results of the Table 2.2 confirm that Bernoulli design is the most robust. However, SRSWOR is almost as robust while still showing the advantage in terms of balance of the fixed sample size. The Mahalanobis distance of the Bernoulli design could be significantly reduced by using the sample mean instead of the Horvitz-Thompson estimator in this case. But the comparison would be less general for survey sampling with a response variable with unequal inclusion probabilities. The Matched design has a relatively large loss of robustness as the largest eigenvalue is doubled compared to Bernoulli sampling. The Cube method appears more robust than the Matched design. The approximation of the eigenvalue underestimates the value a little in the simulations but the square root ratio between the result of the simulations and approximation is $(0.351/0.294)^{1/2} = 1.093$, which means that in the worst case the approximation of the standard deviation is only about 10% too low. The approximation is not accurate for the Cube-Matched method, where $(1.058/0.667)^{1/2} = 1.26$. It must be considered that the maximal eigenvalue estimated through simulation will almost always be slightly inflated when all the inclusion probabilities are equal. This effect can be observed in the maximal eigenvalue estimated through simulations of Bernoulli sampling, SRSWOR and Matching. For the Cube

method, we can also try to check the lower bound found in previous section. We obtain

$$\max_{k \in U} \pi_k (1 - \pi_k) \frac{1}{1 - h_{kk}} = 0.396,$$

which is greater than the largest eigenvalue. This is not necessarily surprising, as the Cube method does not perfectly balance those auxiliary variables.

Table 2.3: Mixtures of designs with algorithm 1. The values of ϕ are between the parentheses.

	Largest eigenvalue	Mahalanobis
SRSWOR/cube(0.25)	0.319	67.719
SRSWOR/cube(0.5)	0.294	117.188
SRSWOR/cube(0.75)	0.272	163.742
Matched/CPS(0.25)	0.320	162.375
Matched/CPS(0.5)	0.381	124.903
Matched/CPS(0.75)	0.443	87.368
Matched-cube/Cube(0.25)	0.476	17.991
Matched-cube/Cube(0.5)	0.641	18.474
Matched-cube/Cube(0.75)	0.809	18.207

Table 2.3 contains the results of simulations for mixture of designs with different values of ϕ . For example, SRSWOR/Cube(0.75) means that $p_A(\cdot)$ is simple random sampling and $p_B(\cdot)$ is the Cube method and $\phi = 0.75$. In Matched/CPS(0.5), the first design is the matched pairs design with equal inclusion probabilities and the second one is CPS that has unequal inclusion probabilities, also we have $\phi = 0.5$. We remind that CPS corresponds to SRSWOR when the inclusion probabilities are equal. When one of the two designs is Matched or Bal-Matched, the Matched design is applied in the first step of Algorithm 1. Indeed, it is not possible to apply strictly a Matched method when the inclusion probabilities are unequal on the second step.

Next, we ran the GS design of Harshaw et al. (2024) using the Julia implementation. We used a version of the algorithm that returns experimental designs with fixed size for each group. Unfortunately, we found that this algorithm does not provide comparable results depending on whether the variables have the same means or variances. We surmise this is due to the way the covariates are normalized in the GS algorithm. Therefore, we centred and reduced all variables and made them orthogonal using a singular value decomposition. Without this change of variables, the results obtained by the GS method are inconsistent. The results of the simulations are presented in Table 2.4.

Table 2.4: Mixtures of designs with Gram-Schmidt walk design. The values of ϕ are between the parentheses.

	Largest eigenvalue	Mahalanobis
GS(0)	0.370	17.187
GS(0.25)	0.330	60.027
GS(0.5)	0.305	99.703
GS(0.75)	0.282	145.087
GS(1)	0.260	199.469

Chapter 3

Balanced Random Assignments for Multiple Treatment Groups Using the Cube Method

Abstract

In this paper, we propose a new approach to achieving covariate balance in experimental settings by using a modified version of the Cube Method, which is typically used in survey sampling for the generation of balanced samples. The modified Cube Method allows participants to be allocated to several balanced treatment groups on the basis of several variables. The method we propose can also handle cases where treatment groups are of unequal size and where inclusion probabilities are unequal. Additionally, we provide variance approximations for cases where a randomization test is not computationally feasible with the modified Cube Method. ¹

3.1 Introduction

The balance of baseline covariates between treatment groups is a desirable outcome in the design of a clinical trial, as well as other types of randomized experiments, surveys or even observational studies. Balancing helps ensuring the validity of study results and reduces errors due to model misspecifications (Härdle and Bowman, 1988; Ho et al., 2007). It can also reduce or eliminate the need to use weighted adjustments or a parametric model in the analysis after the experiment has been carried out. In this article, high balance refer to group assignments whose covariate means are similar. The balance also has to be achieved randomly, as randomness enables us to carry out inference.

To address the issue of covariate balance in trials with two treatment groups, numerous forms of restricted randomization are used such as stratification, blocking, rerandomization, matching, and other more complex algorithms are available (see amongst other Simon, 1979; Morgan and Rubin, 2012; Xu and Kalbfleisch, 2010; Li et al., 2018; Tillé, 2022; Harshaw et al., 2024). However, the literature becomes more limited when dealing with multiple treatment groups. By adapting the principles of the Cube Method, traditionally used in survey sampling, we propose an algorithm that can generate balanced treatment groups. We define this algorithm as the Cube-Based Group Assignment (CBA) algorithm. Unlike existing methods such as some forms of rerandomization, the CBA can also consider groups of unequal size and unequal inclusion probabilities, which can be applied in more diverse settings. To the best of our knowledge, no other method offers this set of features.

The CBA uses the principles of the Cube method, therefore we give a reminder of this method

¹This chapter is a reprint of: E.Talovic and Y.Tillé. Balanced Random Assignments for Multiple Treatment Groups Using the Cube Method, submitted, 2025

and its application in the sampling literature. In practice, the CBA uses the enhanced Cube Method proposed by Jauslin et al. (2021), which effectively handles cases where the population is highly stratified, in a duplicated and highly stratified version of the initial data set. We carry out simulations that show the performance of the CBA in multiple settings compared to a fixed size randomization, which indirectly also compares the CBA to rerandomization. We recommend using randomization tests to obtain p -values and confidence intervals that take into account the complex randomization scheme used with the Cube method. In the case it is not computationally feasible, we provide a method to approximate the variance through simulations. The supplement gives an example of an implementation of the CBA.

3.2 Cube method description

In this section, we briefly describe the Cube method proposed by Deville and Tillé (2004a). In survey sampling, the goal is often to estimate a total or a mean of a variable y of a population $U = \{1, \dots, N\}$ of size N by using a sample $S \subset U$ selected randomly from the population with inclusion probabilities π_1, \dots, π_N . It is common to use auxiliary variables $\mathbf{x}_i = (x_1, \dots, x_p)^\top$ for unit i of the population U to make weighting adjustments to estimated totals by using the calibration proposed by Deville and Särndal (1992) for example. It is also possible to adjust the sampling design to obtain more accurate estimates of the total or the mean by balancing on the auxiliary variables.

The Cube method proposed by Deville and Tillé (2004a) is a sampling algorithm used in survey sampling that enables the selection of samples that are balanced. We define a sample \mathbf{a} as balanced when the Horvitz-Thompson estimator of the auxiliary variables is equal to the population total of the auxiliary variables. We obtain the following equation

$$\sum_{i \in U} \frac{a_i \mathbf{x}_i}{\pi_i} = \sum_{i \in U} \mathbf{x}_i, \quad (3.1)$$

where $a_i = 1$ if unit i is selected in sample S and 0 otherwise. Generating a sampling design that satisfies (3.1) and respects the inclusion probabilities is not an easy task. In most cases, an exact solution does not exist and exceptions to this are usually cases where the covariates are exclusively categorical.

The Cube method generates a sampling design that approximately respects the constraints in equation (3.1) with two phases: the flight phase and the landing phase. The flight phase consists of a random walk in the affine subspace denoted by $Q = \boldsymbol{\pi} + \text{Ker}\{(\mathbf{x}_1/\pi_1, \dots, \mathbf{x}_N/\pi_N)\}$, implied by the balancing equation given in (3.1). The algorithm initiates with $\boldsymbol{\pi}(0) = \boldsymbol{\pi}$. At each step, the algorithm chooses a direction for the walk, while remaining inside Q , and at each step, one dimension in $\boldsymbol{\pi}(t)$ is set to either 0 or 1. After each step, the dimension of the subspace is reduced by at least 1. The flight phase ends when the random walk cannot continue. We define $\boldsymbol{\pi}(T)$ as the last iteration of the random walk in the flight phase. At most p entries of $\boldsymbol{\pi}(T)$ are non-integers. The algorithm terminates if $a_i \in \{0, 1\}$ for all $i \in U$, otherwise the landing phase is initiated. Algorithm 2 shows the basic steps of the flight phase.

The random choice for the next step $\boldsymbol{\pi}(t+1)$ with probability of q and $(1-q)$ in the algorithm allows the random walk to be a martingale, which allows us to respect the inclusion probabilities. After the flight phase, the landing phase handles the remaining non-integer element of \mathbf{a} by linear programming or elimination of variable. The landing phase is still random. Elimination of variable simply consists in iteratively dropping a constraint and running the flight phase again. The linear programming solution is more complex (see Deville and Tillé, 2004a, for full development).

Figure 3.1 shows an example with a fixed sample size constraint $n = 2$ in a population of size 3. In this case, exact solutions can be found and the Cube algorithm terminates without initiating the landing phase.

Algorithm 2 Cube method: Flight phase

Require $\boldsymbol{\pi}$.

set $\boldsymbol{\pi}(0) \leftarrow \boldsymbol{\pi}$

For $t = 0, 1, 2, \dots$

Find a non-zero vector $\mathbf{v} \in \text{Ker}\{(\mathbf{x}_1/\pi_1, \dots, \mathbf{x}_N/\pi_N)\}$ such that $v_i = 0$ if a_i is an integer.

Find the largest λ_1 and λ_2 such that

$$0 \leq \boldsymbol{\pi}(t) + \lambda_1 \mathbf{v} \leq 1 \text{ and } 0 \leq \boldsymbol{\pi}(t) - \lambda_2 \mathbf{v} \leq 1.$$

set $q \leftarrow \frac{\lambda_1}{\lambda_1 + \lambda_2}$

set $\boldsymbol{\pi}(t+1) \leftarrow \begin{cases} \boldsymbol{\pi}(t) + \lambda_1 \mathbf{v}, & \text{with probability } q \\ \boldsymbol{\pi}(t) - \lambda_2 \mathbf{v}, & \text{otherwise.} \end{cases}$

Until There is no non-zero vector $\mathbf{v} \in \text{Ker}(Q)$ such that $v_i = 0$ if a_i is an integer for $\boldsymbol{\pi}(T)$

Output $\boldsymbol{\pi}(T)$




Figure 3.1: Representation of the flight phase in the algorithm. The vertices of the cube represent different samples. The arrow represents the initial inclusion probabilities, before the initiating the random walk. The darker subspace describe the subspace generated by the fixed size constraint and inclusion probabilities. In this example, the samples satisfy exactly the constraints as the vertices intersect with the subspace.

For the design of experiments in clinical trials, Tillé (2022) proposed the Cube method to create a test group and a control group with very similar means for their baseline covariates. While very different in principle, the final result in terms of balance is similar to a rerandomization with a low acceptance rate if the inclusion probabilities are equal and the two group are of the same size. In the following sections, we present a method that generalizes the Cube method to cases where the number of groups to be created exceeds 2.

3.3 Notation for multiple treatment group assignment

Let $U = \{1, \dots, N\}$ be the population of size N of the experiment. The goal is to assign the experimental units to K different treatment groups in a random and balanced manner. Treatments are randomly assigned to units in the population U . Let $\boldsymbol{\pi} = (\pi_1, \dots, \pi_N)$ be the vector of inclusion probabilities when there are two treatment groups. Otherwise, we use the vector $\boldsymbol{\pi} = (\pi_{11}, \dots, \pi_{N1}, \pi_{12}, \dots, \pi_{N2}, \dots, \pi_{1K}, \dots, \pi_{NK})^\top$ of length $K \times N$ and where π_{ik} denotes the inclusion probability of unit i to the treatment group k . We also define $\pi_{ik,jl}$ is the joint probability that unit i is included in group k and unit j is included in group l . We have

$\pi_{ik,ik} = \pi_{ik}$. Let $\mathbf{a} = (a_{11}, \dots, a_{N1}, a_{12}, \dots, a_{N2}, \dots, a_{1K}, \dots, a_{NK})^\top$ be the vector of indicator variables, where $a_{ik} = 1$ when the unit $i \in U$ is included in the group k and 0 otherwise. Since the number of assignments is fixed in each treatment group, we have

$$\sum_{i \in U} \pi_{ik} = \sum_{i \in U} a_{ik} = n_k, k = 1, \dots, K,$$

$$\sum_{k=1}^K \pi_{ik} = \sum_{k=1}^K a_{ik} = 1, i \in U,$$

and

$$\sum_{k=1}^K n_k = N,$$

where n_k is the number of units assigned in treatment group k .

The covariates or auxiliary variables are $\mathbf{x}_i = (x_1, \dots, x_p)^\top$ for unit i of the population, for $i \in U$. The value of the response variable y depends on the treatment group. Let y_{ik} denotes the value taken by variable y on unit k if it is assigned to treatment group k . In practice, most responses y_{ik} are not observed since the experiment is conducted only once, but this notation remains useful for theoretical purposes. Define also the covariates Hájek estimators of the means (Hájek, 1971) for group k as

$$\bar{\mathbf{x}}_k = \frac{\sum_{i \in U} a_{ik} \mathbf{x}_i / \pi_{ik}}{\sum_{i \in U} a_{ik} / \pi_{ik}}. \quad (3.2)$$

Define the square of the Mahalanobis distance between the group of indices k and ℓ as

$$M(k, \ell)^2 = (\bar{\mathbf{x}}_k - \bar{\mathbf{x}}_\ell)^\top \boldsymbol{\Sigma}^{-1} (\bar{\mathbf{x}}_k - \bar{\mathbf{x}}_\ell),$$

where $\boldsymbol{\Sigma}$ is the covariance matrix of the covariates $\mathbf{x}_1, \dots, \mathbf{x}_N$. We favor the use of the Mahalanobis distance because it takes into account the dispersion of the covariates and the correlation between the covariates. We use the average of the Mahalanobis distances of the covariates mean between all the groups to assess the balance of the group assignment. In mathematical terms, it can be written as

$$\bar{M} = \frac{\sum_{k < \ell} M(k, \ell)}{2K(K-1)} \quad \text{for } k, \ell \in \{1, \dots, K\}. \quad (3.3)$$

In experiments with multiple groups, the Horvitz-Thompson estimator is often used to account for unequal inclusion probabilities when estimating the effects of different treatments. One group is typically designated as the ‘‘control group,’’ which serves as the standard for comparison. The treatment effects are then assessed by measuring the differences between the outcomes of the control group and each treatment group k . This approach quantifies the treatment effect of group k relative to the control and ensures that these comparisons remain unbiased, even with unequal inclusion probabilities (Horvitz and Thompson, 1952; Deville and Tillé, 2004a).

An alternative to the Horvitz-Thompson estimator is the Hájek estimator, which is often preferred when inclusion probabilities are unequal and when some probabilities are very small in survey sampling. The Hájek estimator normalizes the sampling weights by their realized total, which both reduces the influence of extreme (very small) inclusion probabilities on variance and the biases that can arise when those extreme weights coincide with atypical response values. When the sample size is fixed and inclusion probabilities are equal, both the Horvitz-Thompson and Hájek estimators reduce to a simple difference-in-means estimator. The Hájek estimator is typically used in the field of survey sampling, but it can be applied in experiments.

Using the first group as the control group, we define the Horvitz-Thompson estimator of the

treatment effects in each group as

$$\hat{\tau}_k = \frac{1}{N} \sum_{i \in U} \frac{a_{ik} y_{ik}}{\pi_{ik}} - \frac{1}{N} \sum_{i \in U} \frac{a_{i1} y_{i1}}{\pi_{i1}} \quad k \in \{2, \dots, K\}.$$

We have that $\hat{\tau}_1 = 0$, as it is the control group. Using the first group as the control group again, we define the Hájek estimator of the treatment effects in each group as

$$\hat{\tau}_k^H = \frac{\sum_{i \in U} a_{ik} y_{ik} / \pi_{ik}}{\sum_{i \in U} a_{ik} / \pi_{ik}} - \frac{\sum_{i \in U} a_{i1} y_{i1} / \pi_{i1}}{\sum_{i \in U} a_{i1} / \pi_{i1}} \quad k \in \{2, \dots, K\}.$$

For simulations, we will use a simple regression model where the treatment effects are constant with

$$y_{ik} = \mathbf{x}_k^\top \boldsymbol{\beta} + \tau_k + \varepsilon_{ik}, \quad (3.4)$$

where $\boldsymbol{\beta}$ is the regression coefficient vector of dimension $p \times 1$, and ε_{ik} represents the variance term for unit i in group k . Applying a balanced sampling design can significantly reduce both the bias and variance associated with the term $\mathbf{x}_k^\top \boldsymbol{\beta}$ and allows for a better estimation of τ_k .

3.4 Multiple treatment group Cube method

In order to carry out the group assignment, we need to generate a random vector \mathbf{a} that respects the following properties. The inclusion probabilities must be respected:

$$E(a_{ik}) = \pi_{ik}, \quad i = 1, \dots, N, k = 1, \dots, K.$$

The sample must be balanced in each group:

$$\sum_{i \in U} \frac{a_{ik} \mathbf{x}_i}{\pi_{ik}} = \sum_{i \in U} \mathbf{x}_i, k = 1, \dots, K. \quad (3.5)$$

In order to obtain a fixed sample size, we also need

$$\sum_{i \in U} a_{ik} = n_k, k = 1, \dots, K. \quad (3.6)$$

Finally, as only one unit can be assigned per group:

$$\sum_{k \in U} a_{ik} = 1, i = 1, \dots, N. \quad (3.7)$$

The vector \mathbf{a} has dimensions $N \times K$. If p denotes the dimension of \mathbf{x}_i , the number of balancing constraints is $p \times K + K + N = N + (p + 1) \times K$, resulting in a large number of constraints. However, the constraints in Equation (3.7) can be viewed as a stratification, where only one unit is selected from each stratum. In order to solve this balanced sampling problem, we use Jauslin et al. (2021)'s enhanced algorithm, which allows us to deal with highly stratified populations.

This algorithm is implemented in the `stratifiedcube` function in the R package `StratifiedSampling` (Jauslin et al., 2022). It provides a balanced sample on auxiliary variables, while applying a very high level of stratification. In this case, the stratification constraint allows each individual to be allocated to exactly one treatment group, which means respecting the constraint given in Equation (3.7).

The implementation is structured as follows:

- Create a population U' that corresponds to K duplicates of the population $U = \{u_1, \dots, u_N\}$. We can write this new population as $U' = \{u'_{11}, \dots, u'_{N1}, u'_{12}, \dots, u'_{N2}, \dots, u'_{NK}\}$.

- Set as “stratum per individual” the sub-populations $\{u'_{i1}, \dots, u'_{iK}\}$ for all $i \in \{1, \dots, N\}$. We end up with a population U' of size $N \times K$ with N strata of size K . In each of those strata, the inclusion probabilities are $\{\pi_{i1}, \dots, \pi_{iK}\}$ for all $i \in U$. The sum of the inclusion of probability in each individual strata $\sum_{k=1}^K \pi_{ik}$ must be equal to 1, for all $i \in \{1, \dots, N\}$.
- For the matrix of auxiliary variables, we use the matrix resulting from a Kronecker product between an identity matrix of dimension $K \times K$ and the $N \times p$ covariates matrix \mathbf{X} .
- For each group, the vector of inclusion probability $\boldsymbol{\pi}_k = (\pi_{1k}, \dots, \pi_{Nk})$ is added in \mathbf{X} in order to have a fixed sample size for the groups. We end up with $K \times p + 1$ constraints or auxiliary variables in total, excluding the ones due to the stratification induced by Equation (3.7). In the case of sampling with unequal probability, it is strongly advised to include a column of 1 in the \mathbf{X} matrix. It would add one more variable to be balanced on if there was not already an intercept column but this variable improves the stability of the denominator the Hájek estimator given in Equation (3.2).
- The algorithm outputs a sample $\mathbf{a}' \in \{0, 1\}^{NK}$, where $a'_{ik} = 1$ implies that unit i is assigned to treatment group k , where $i \in \{1, \dots, N\}$ and $k \in \{1, \dots, K\}$ and $a'_{ik} = 0$ implies the opposite.

The matrix given in Expression (3.8) contains an example of the form of the matrix of balancing variables in the algorithm with 4 different groups. The $N \times p$ matrix \mathbf{X} is the covariates matrix, whose rows are the covariates $\mathbf{x}_1, \dots, \mathbf{x}_N$. A matrix of this form given to the enhanced algorithm enables us to obtain groups that are balanced on covariates means.

$$\begin{pmatrix} \boldsymbol{\pi}_1 & 0 & 0 & 0 & \mathbf{X} & 0 & 0 & 0 \\ 0 & \boldsymbol{\pi}_2 & 0 & 0 & 0 & \mathbf{X} & 0 & 0 \\ 0 & 0 & \boldsymbol{\pi}_3 & 0 & 0 & 0 & \mathbf{X} & 0 \\ 0 & 0 & 0 & \boldsymbol{\pi}_4 & 0 & 0 & 0 & \mathbf{X} \end{pmatrix}. \quad (3.8)$$

This matrix incorporates the constraints specified in (3.5) and (3.6), but does not fully capture all the constraints. The additional constraints in (3.7) are addressed through stratification in the enhanced Cube method. Figure 3.2 illustrates how the stratification is applied within the algorithm. The output is an allocation of the units to the groups that is generated randomly

u'_{11}	π_{11}	0	u'_{21}	π_{21}	0	...	u'_{N1}	π_{N1}	0
u'_{12}	π_{12}	0	u'_{22}	π_{22}	1		u'_{N2}	π_{N2}	0
u'_{13}	π_{13}	0	u'_{23}	π_{23}	0		u'_{N3}	π_{N3}	0
u'_{14}	π_{14}	1	u'_{24}	π_{24}	0		u'_{N4}	π_{N4}	0
u'_{15}	π_{15}	0	u'_{25}	π_{25}	0		u'_{N5}	π_{N5}	1
u'_{16}	π_{16}	0	u'_{26}	π_{26}	0		u'_{N6}	π_{N6}	0

Figure 3.2: Illustration of the stratification in the CBA with $K = 6$ treatment groups. Each column represents a “stratum per individual”, with its indices, its inclusion of probabilities to each group and the value assigned through the algorithm. Only one entry per column is set to 1, due to the stratification. This ensures that only each element is assigned to only 1 treatment group.

and approximately respects the given constraints and the stratification with the correct inclusion probabilities.

3.5 Simulations

3.5.1 Improvement in balance between groups

In this section, we compare the performance of the CBA with simple random sampling or fixed size randomization, which consists of randomly separating the experimental units into K groups of fixed size in this case. We specify that a randomization with a fixed sample size is used, as any comparison with the CBA would otherwise be unfair. The comparison with randomization also allows us to compare the CBA with rerandomization (Morgan and Rubin, 2012) for multiple treatment groups, as this method consists in selecting the most balanced assignment from a set of generated assignments. Alternatively, rerandomization consists in repeating simple random sampling until a satisfactory assignment to groups is generated, but the former definition of rerandomization is used in this article.

The main metric of interest is how balanced the group assignments are. In order to do this, the Mean Between groups Mahalanobis Distance \overline{M} , defined in (3.3), is used. Other criterion for the acceptance in rerandomization can be used, such as those typically from multivariate analysis of variance (Morgan, 2011). We also verify that the first-order inclusion probabilities are respected.

The first set of simulations is carried out on the data set of the experiment from Balcells et al. (2022). Its goal was to test the effect of a visit in a specific museum on the difference of political opinions of different university students from Chili, by comparing it to a control group that did not visit the museum. In this data set, there are 251 participants and 12 pre-treatment covariates, which are a mix of categorical and continuous variables. The columns of the data set denote the following variables recorded before the experiment: age, gender, political ideology, economic situation, interest in politics, religiosity, visits to other museums, trust in government, satisfaction in government, opinion on inequality and positive or negative emotions. The units with missing values are removed. This manipulation conveniently left 245 remaining participants, which is divisible by 7. In order to test the CBA, we separate the groups into 7 groups of 35 using the CBA and compare its performance relative to a fixed size randomization. Each participant has an equal probability of being assigned to any group.

Two sets of simulations *Sim1* and *Sim2* are with this data set. *Sim2* uses all 12 covariates and 10000 simulations are made. *Sim2* only uses 6 covariates to assess the effect of the number of covariates on the performance of the algorithm. The columns retained are simply the first six ones of the data set.

The second data set is drawn from Milosevic (2022). This a data set of patients' mild cognitive impairment and dementia due to Alzheimer's disease where the diagnostic accuracy is compared between two questionnaires. The data set has 75 participants and we retain 8 covariates, which are also a mix of categorical and continuous variables. Two acronyms are used in this data set. MMSE stands for mini-mental state examination, which corresponds to a cognitive screening made by a neurologist and the IRT-score corresponds to the score obtained from a scoring method a questionnaire related to cognitive impairment and dementia. The retained covariates are the following: Sex, Age, MMSE patient, years of education, IRT-score, diagnosis of patients' dementia and level of education. For our simulations, we create 3 treatment groups of equal size 25 and compare the performance of the CBA with fixed size randomization. We run two sets of simulations on this data set, one in which the inclusion probabilities are equal denoted by *Sim3* and the other in which they are randomly generated denoted by *Sim4*.

For the 10000 simulations on each set of simulations, we computed three indicators:

- Mean/Mean ratio: The mean of \overline{M} for balanced assignment divided by the mean of \overline{M} for an assignment by simple random sampling or fixed size randomization.
- Max/min ratio: The maximum of \overline{M} for balanced assignment divided by the minimum of \overline{M} for an assignment by simple random sampling.

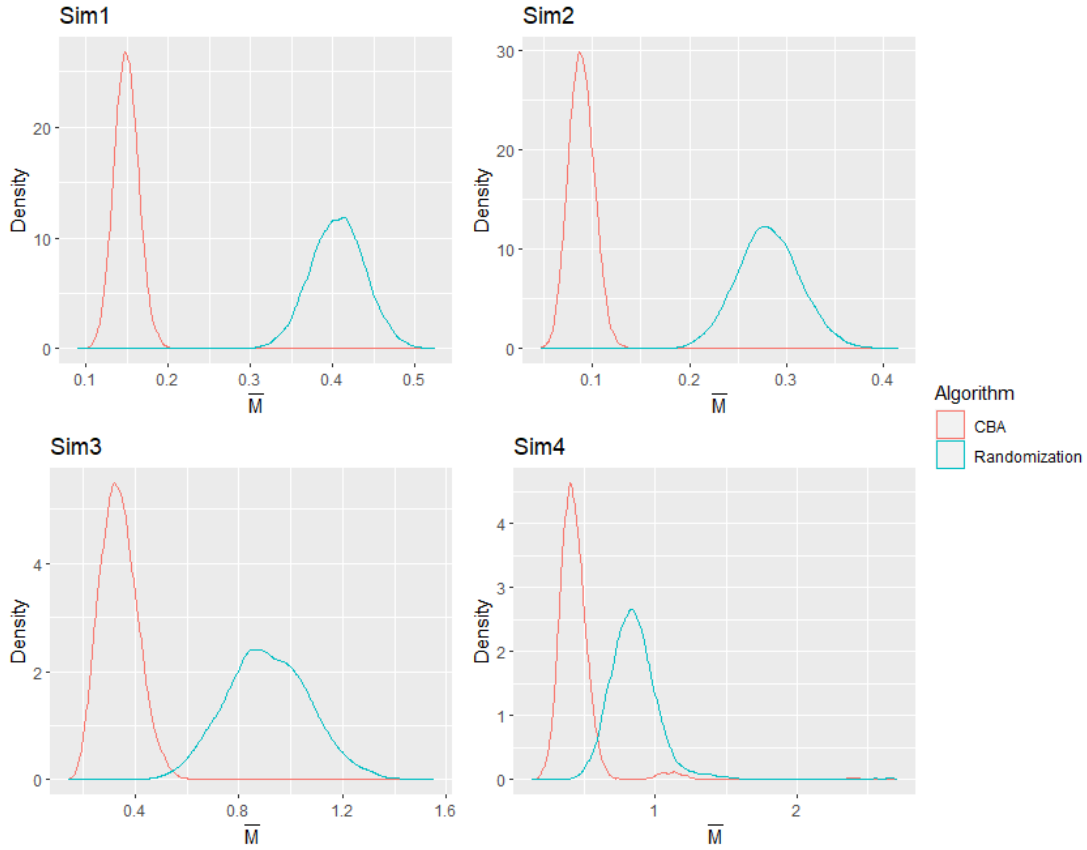


Figure 3.3: Density plots of \overline{M} of the CBA and randomization under varying simulation settings.

- Mean/Min ratio: The mean of \overline{M} for balanced assignment divided by the minimum of \overline{M} for an assignment by simple random sampling.

Those indicators are useful for making a comparison between the CBA and rerandomization algorithms that take the best assignment out of a certain number of randomizations.

Figure 3.3 shows the density plots from our simulations of \overline{M} , allowing one to compare the CBA and randomization methods. The CBA algorithm performs significantly better, and in some cases, the distributions are distinct.

In Table 3.1, the top-left case shows that the CBA produces assignments with an average value of \overline{M} significantly lower than those generated through randomization in the first set of simulations *Sim1*. The first entry in the second column indicates that even the worst-case group assignment by the CBA outperforms the best group assignment generated through randomization. This also holds true for the average assignments produced by the CBA in the first dataset. Such metrics are interesting, as they show that even a rerandomization algorithm that accepts only the best assignment out 10000 is not preferable to the CBA in terms of the expected \overline{M} . The results from *Sim2* suggest that reducing the number of covariates to 6 does not change any of the outcomes significantly.

The set of simulation *Sim3* is carried on the second dataset, which is smaller and has fewer covariates to balance with equal probabilities. In *Sim3*, the relative average performance of the CBA is similar to the one observed in *Sim1*. On the other hand, the best assignment to groups of randomizations had lower \overline{M} than the worst assignment to groups of the CBA in a set of 10000 simulations. On average, the CBA still outperforms the most balanced randomization out of 10000 of trials.

In a context outside of simulations, the unequal inclusion probabilities would usually be pro-

Table 3.1: Relative performance of different metrics between the CBA algorithm and a fixed size randomization using the Mahalanobis distance between groups. For *Sim4*, the CBA is compared with an unequal inclusion probability sampling designs whose only constraint is group size. The ratios are drawn from a set of 10000 simulations. *Sim1* represents the first data set with $N = 245$ units, $p = 12$ covariates, $K = 7$ treatment groups and fully equal inclusion probabilities. *Sim2* is the same set of simulations, with only 6 covariates to balance. *Sim3* represents the second data set with $N = 75$ units, $p = 8$, $K = 3$ treatment groups and fully equal inclusion probabilities. *Sim4* is the same data set as *Sim3* with unequal inclusion probabilities that are randomly generated.

	Mean/Mean ratio	Max/Min ratio	Mean/Min ratio
<i>Sim1</i>	0.36	0.72	0.52
<i>Sim2</i>	0.32	0.84	0.54
<i>Sim3</i>	0.36	1.75	0.86
<i>Sim4</i>	0.51	6.74	1.19

portional to one of the covariates that is expected to have a large effect on the response variable, which is absent in our data set. For the simulation set *Sim4*, unequal inclusion probabilities are partially based on the first variable (years of education) and partially generated randomly using a uniform distribution, then adjusted for the margins through raking (Deville et al., 1993). Further details are provided in the supplementary material. The simulations with unequal inclusion probabilities (*Sim4*) seem to perform well on average, but the lowest \bar{M} from the randomization was lower than the average \bar{M} from the CBA. The equivalent of a fixed size randomization but with unequal probability has been made using the Cube method for simplicity with the correct stratification and inclusion probabilities but without balancing variables. It could also have been solved using a form of unbiased Controlled Rounding (Cox, 1987; Friedrich et al., 2007).

The simulations demonstrate that the CBA remains effective even in scenarios with arbitrary inclusion probabilities. For all the set of simulations, the largest difference between the simulated inclusion probabilities and the theoretical ones was strictly lower 0.013 and the average difference was strictly lower than 0.0037 for the CBA. Exact test based on the binomial distribution indicates that the individual inclusion probabilities are respected on all the sets of simulations. In comparison, when taking the 1000 most balanced randomizations out of 10^6 with the set of simulation with unequal probabilities (*Sim4*), the largest difference absolute difference was 0.04957. When using exact binomial tests, 3 entries out of 225 (3 groups of size 25 for a total of 75 units implies there are 225 inclusion probabilities to be measured.) had values so extreme that they had p -values lower than 0.0017. It indicates that there is strong evidence of bias in the inclusion probabilities of certain units when using rerandomization. Biased inclusion probabilities with rerandomization typically occur with units that can be considered outliers. It can also be expected that the bias would tend to be larger when the groups are of unequal size, smaller or when the sizes of the groups are not fixed.

Another important point is the computation time. On the set of simulations *Sim1*, *Sim2* and *Sim3*, the CBA algorithm terminated respectively in 7.2, 2.2 and 0.12 seconds for a single computation on average. The sample size, number of groups and number of variables to balance greatly influence the computation time. In cases with equal group size and inclusion probabilities, one can run the CBA multiple times and select the best group assignment. In other words, it involves combining rerandomization with the CBA, which, given the computation times, is feasible unless the dataset is very large.

Using the dataset from U.S. Census Bureau (2024) that includes various variables for cities in the United States and contains 31120 rows, we retain the half most populated cities. The selected covariates are population, density, latitude, and longitude. Amongst the remaining cities, we simulate an experiment that uses half of those and ignores the rest. There are four treatment

groups and each unit has an inclusion probability of being included in each group of $1/8$. The computation time is approximately 5 minutes, significantly longer than what is required for a randomization. Therefore, only a single computation using the CBA is performed. This computation yields a result that outperforms 10000 simulations generated through randomization. The ratio of \bar{M} between the CBA algorithm and the average of the randomization simulations is 0.16, while the best result from randomization yields a ratio of 0.78.

3.5.2 Variance reduction in the estimators

In this section, we showcase the reduction in variance in the Horvitz-Thompson or the Hájek estimator of the treatment effect.

We generate a response variable \mathbf{y} using the datasets and simulations from Table 3.1. For *Sim1* and *Sim2*, the response is generated through a linear model as in (3.4) with an intercept of 2, regression coefficients of 0.5 for the age and the 2 for the level of political interest, and a variance term drawn from a centered normal distribution with variance 5. We define it as Model 1. There are seven groups: we assume a null treatment effect for all groups. For *Sim3* and *Sim4*, the response is generated through a linear model with an intercept of 2, regression coefficients of 0.5 for both age and the IRT score, and a variance term drawn from a centered normal distribution with variance 5. We define it as Model 2. There are three groups: we assume a null treatment effect for group 2 and a constant treatment effect of 4 for group 3.

Our simulations demonstrate that using the CBA instead of a pure randomization significantly reduces the variance of $\hat{\tau}_k$. For the first data set (*Sim1* and *Sim2*), the variance ratio between the CBA, which balances even on covariates that have no effect on \mathbf{y} , and randomization is 0.075. For the second data set under equal inclusion probabilities (*Sim3*), the variance ratio between the CBA algorithm and randomization is 0.04 when using the Horvitz-Thompson estimator. We obtain a similar improvement when using the Hájek estimator with the unequal inclusion probabilities (*Sim4*).

In the simulated example, the variance reduction achieved by the CBA in comparison to a pure randomization is very large, even though only two variables were relevant for balancing. It is important to note that the reduction in the variance may be less pronounced if the response variable does not follow a linear model or if the covariates are weakly related to the response variable.

3.6 Inference

3.6.1 Recommendation

We give some indications for performing inference after using the CBA or other designs that balance on covariates. We recommend using a randomization test whenever it is feasible. We also give an approximation of the variance for Horvitz-Thompson estimators in order to obtain confidence intervals for cases where a randomization test is not feasible. Covariates adjustments or calibration weighting are also a valid approach, but using the CBA reduces the need for such adjustments by ensuring better initial balance in the design.

3.6.2 Randomization test

After applying the assignment algorithm to allocate the experimental units into multiple treatment groups, it is performing statistical inference can be slightly more complicated, as the assignment of each unit is not independent from the other units. An effective approach is to utilize a method akin to a randomization test (Good, 1994; Rubin, 1980; Basu, 1980; Cox, 2009). This method allows for the calculation of p -values and confidence intervals. The null hypothesis

is often that the treatments have no effect but it can also be assumed that those have a non-null fixed effect.

With the null hypothesis in place, rerun the assignment algorithm a large number of times. In each iteration, randomly assign, according to the assignment scheme of the algorithm, the subjects to treatment groups. This process simulates the distribution of the response variable under the null hypothesis. If the null hypothesis is $\tau_k = 0$ for all $k \in \{1, \dots, K\}$, only the assignments vary and the response variable remains unchanged. For each randomization, calculate a test statistic of interest, such as the average absolute difference in response means or Horvitz-Thompson estimators of the mean between each group, a Mahalanobis distance, or an F-statistics in some cases. This process also simulates the distribution of the test statistic under the null hypothesis. Compare the observed test statistic from the original data to the distribution of test statistics obtained from the randomizations. The p -value is calculated as the proportion of randomizations where the test statistic is as extreme or more extreme than the observed test statistic. The confidence intervals can also be deduced using the simulated distribution. This method takes into account the randomization scheme that generates the assignments and does not depend on a parametric model for inference.

3.6.3 Approximation of the variance of estimators

In this subsection, we give an approximation of the variance that is applicable when repeating a large number of simulations is not feasible due to computation time. The variance of the Horvitz-Thompson estimator of the mean of group k is

$$\text{Var}(\bar{y}_k) = \frac{1}{N^2} \sum_{i \in U} \sum_{j \in U} \left(\frac{\pi_{ik,jk} - \pi_{ik}\pi_{jk}}{\pi_{ik}\pi_{jk}} \right) y_{ik}y_{jk},$$

where \bar{y}_k is the Horvitz-Thompson estimator of mean of the response variable in group k . The variance of $\hat{\tau}_k$ is

$$\text{Var}(\hat{\tau}_k) = \text{Var}(\bar{y}_k) + \text{Var}(\bar{y}_0) - 2\text{Cov}(\bar{y}_k, \bar{y}_0).$$

Thus, $\text{Var}(\hat{\tau}_k)$ becomes

$$\begin{aligned} \text{Var}(\hat{\tau}_k) &= \frac{1}{N^2} \sum_{i \in U} \sum_{j \in U} \left(\frac{\pi_{ik,jk} - \pi_{ik}\pi_{jk}}{\pi_{ik}\pi_{jk}} \right) y_{ik}y_{jk} \\ &\quad + \frac{1}{N^2} \sum_{i \in U} \sum_{j \in U} \left(\frac{\pi_{i1,j1} - \pi_{i1}\pi_{j1}}{\pi_{i1}\pi_{j1}} \right) y_{i1}y_{j1} \\ &\quad - \frac{2}{N^2} \sum_{i \in U} \sum_{j \in U} \left(\frac{\pi_{ik,j1} - \pi_{ik}\pi_{j1}}{\pi_{ik}\pi_{j1}} \right) y_{ik}y_{j1} \end{aligned}$$

In practice, when conducting experiments with multiple groups, the true response y_{ik} for each unit i in group k is often not observable, as the experiment can typically be carried out only once. To address this limitation, we assume that treatment effects are constant, allowing us to obtain unbiased estimates of the unobserved responses based on the real observations. Specifically, we introduce the notation $\hat{y}_{ik} = y_{iR_i} - \hat{\tau}_{R_i} + \hat{\tau}_k$, where R_i is the group to which unit i was assigned, which implies that y_{iR_i} is the observed response variable of unit i . This formulation allows us to construct the estimated responses for all group k and to calculate the estimated variance $\widehat{\text{var}}(\hat{\tau}_k)$ for the treatment effect. Alternatively, it is possible to calculate regression-based estimator for the difference between groups.

The estimate of the variance of $\hat{\tau}_k$ is

$$\begin{aligned}\widehat{\text{Var}}(\hat{\tau}_k) &= \frac{1}{N^2} \sum_{i \in U} \sum_{j \in U} \left(\frac{\pi_{ik,jk} - \pi_{ik}\pi_{jk}}{\pi_{ik}\pi_{jk}} \right) \hat{y}_{ik}\hat{y}_{jk} \\ &+ \frac{1}{N^2} \sum_{i \in U} \sum_{j \in U} \left(\frac{\pi_{i1,j1} - \pi_{i1}\pi_{j1}}{\pi_{i1}\pi_{j1}} \right) \hat{y}_{i1}\hat{y}_{j1} \\ &- \frac{2}{N^2} \sum_{i \in U} \sum_{j \in U} \left(\frac{\pi_{ik,j1} - \pi_{ik}\pi_{j1}}{\pi_{ik}\pi_{j1}} \right) \hat{y}_{ik}\hat{y}_{j1}\end{aligned}$$

In specific sampling designs, the variances $\text{Var}(a_{ik}a_{jk})$ or the joint inclusion probabilities $\pi_{ik,jk}$ are known analytically. However, it is not the case for the Cube method, whose variances can only be accurately approximated through simulations. Deville and Tillé (2005) gives a good approximation of the variance matrix of balanced sampling designs. The paper approximates the variance by assuming that balanced sampling is close to Poisson sampling, whose design variance matrix is known, conditioned on certain constraints related to balancing. It resembles a projection onto to a modified subspace generated by the covariates \mathbf{X} . Using the same reasoning by replacing the Poisson sampling by a stratified Poisson sampling, whose design variance matrix is known, we give an approximation applied to our case with multiple groups and large sample size. The joint inclusion probabilities can easily be deduced the first-order inclusion probabilities and the variance matrix of a sampling design.

First, define \mathbf{S} as the matrix described in Expression (3.8) whose rows corresponding to unit i in group k are then divided by π_{ik} respectively, for $i \in \{1, \dots, N\}$ $k \in \{1, \dots, K\}$. Define Δ_{Poiss} as the variance matrix of a stratified Poisson sampling design with the same inclusion probabilities and stratification as defined in the CBA. All the entries of this matrix are known. The diagonal entries equal $\pi_{ik}(1 - \pi_{ik})$. The entry corresponding to unit i in different groups k and ℓ equals $-\pi_{ik}\pi_{i\ell}$ and 0 for the remaining cross terms. In other words, if i', j' corresponds to respectively unit i in group k and unit j in group ℓ ($i' \bmod N = i$ and $j' \bmod N = j$) then

$$\Delta_{Poiss_{i'j'}} = \begin{cases} \pi_{ik}(1 - \pi_{ik}), & \text{if } i' = j' \\ -\pi_{ik}\pi_{i\ell}, & \text{otherwise if } (i' \bmod N) = (j' \bmod N) \\ 0, & \text{otherwise.} \end{cases} \quad (3.9)$$

The approximation of the variance matrix of the balanced sampling design is

$$\Delta_{bal} = r \Delta_{Poiss} \left[\mathbf{I} - \mathbf{S}(\mathbf{S}^\top \Delta_{Poiss} \mathbf{S})^{-1} \mathbf{S}^\top \Delta_{Poiss} \right], \quad (3.10)$$

where r is a constant. Δ_{bal} must be adjusted by a constant factor r to account for the loss of dimensionality caused by the projection onto the subspace \mathbf{X} . We propose defining r as the ratio between the sum of the diagonal entries of Δ_{bal} and the sum of the first-order inclusion probabilities. A close alternative to this is $r = N/(N - p)$. An approximation of $\pi_{ik,j\ell}$ can be deduced from Δ_{bal} . We have $\pi_{ik,j\ell} = \Delta_{bal_{ik,j\ell}} + \pi_{ik}\pi_{j\ell}^\top$. This allows us to compute $\widehat{\text{Var}}(\hat{\tau}_k)$ for large data sets without simulations. The variance of the Hájek estimator can be approximated with the variance of the Horvitz-estimator but corrected by linearization (Tillé, 2020b).

Using the sets of simulation *Sim1*, *Sim2*, *Sim3* and *Sim4*, each comprising 10,000 simulations, we compute $\hat{\tau}$ to obtain reliable confidence intervals. We also compute confidence intervals based on our variance approximations and compare them to the confidence intervals found by simulations. The coverage probability for of the true parameter τ_2 of the approximated 95% confidence intervals are given in Table 3.2. The approximation tends to underestimate the size of the confidence intervals, but the approximation improves when the sample size gets larger and uses equal probabilities, as the Cube method is less affected by rounding issues under these

conditions. This improvement in performance with larger sample sizes is advantageous, as the variance approximation becomes a practical alternative when conducting a randomization test would be computationally too demanding. It also shows that the behavior of the variance is relatively predictable when the sample size tends to infinity.

Table 3.2: Coverage probability of the true parameter τ_2 within the approximated confidence interval. Model 1 has been used for *Sim1* and *Sim2* and Model 2 for *Sim3* and *Sim4*.

	95%	90%	50%
<i>Sim1</i>	0.9537	0.899	0.4983
<i>Sim2</i>	0.9454	0.8914	0.4949
<i>Sim3</i>	0.9245	0.862	0.4577
<i>Sim4</i>	0.9002	0.8342	0.4289

3.7 Conclusion

The CBA outperforms randomization or rerandomization in many of the simulations. In some instances, even the most unbalanced group assignment produced by the CBA is more balanced than those achieved through rerandomization. There are cases where rerandomization, given a high number of iterations, performs better on average than a single computation of the CBA. However, it is also possible to combine the CBA algorithm and rerandomization if one needs very balanced group assignments when the inclusion probabilities and group size are equal.

Another important point of the CBA is that it enables the creation of groups that are balanced even with unequal inclusion probabilities. Rerandomization does not allow for that, because the most balanced assignments to groups often tend to include certain units in specific groups more frequently than expected when dealing with groups of unequal size or unequal inclusion probabilities. The consequence is that the inclusion probabilities would be biased in this case. In fact, carrying out the equivalent of a fixed size randomization for multiple groups with fixed sample sizes and with unequal inclusion probabilities is not necessarily a common and easy task.

One particular application of the CBA is in the design of experiments of clinical trials. However, it is not common for all participants to be available at the start of the trial. In many such cases, there is a challenge in balancing the covariates between treatment groups when not all data is available from the outset, and even ensuring equal group sizes cannot be guaranteed. An adaptive method similar to the one introduced by Kapelner and Krieger (2023) can be used. A potential solution to this problem based on the Cube method is to adopt a method similar to that proposed by Jauslin et al. (2022). The core idea of the method is to have participants wait in a “reservoir” group. Once the reservoir group reaches a sufficient size, the Cube method can be applied to assign the participants to the trial in a balanced manner.

3.8 Supplement for Balanced Random Assignments for Multiple Treatment Groups Using the Cube Method

3.8.1 Code

In this section, we provide an example of an implementation of the Cube-based Assignment to Groups Algorithm using the *R* programming language (R Core Team, 2024).

```
library(StratifiedSampling)
library(BalancedSampling)
library(sampling)

set.seed(12345)
N=200
K=5
p=7
X=array(rnorm(N*p),c(N,p))
X=t(t(X)+2*(1:p))

# Cube
XX=kronecker(diag(rep(1,K)), X)
XX=cbind(kronecker(diag(rep(1,K)),rep(1,N)),XX)
stra=rep(1:N,K)
pik=rep(1/K,N*K)
s=stratifiedcube(XX ,stra, pik, EPS = 1e-07, rand=TRUE, landing=TRUE)
S=array(s,c(N,K))
E=t(X*K)%*%S/N

# Simple random sampling
rsums <- rep(1, 200)
csums <- c(40,40,40,40,40)
S_simple <- r2dtable(1, rsums, csums)[[1]]
E_simple=t(X*K)%*%S_simple/N

# Comparison
mal_dist<-function(col1,col2){
  mal=c(t(col1-col2)%*%Sig1%*(col1-col2))
  return(sqrt(mal))
}

pairwise.mal<-function(E,Sig1){
  sapply(1:ncol(E), function(i) sapply(1:ncol(E), function(j){
    mal_dist(E[,c(i)],E[,c(j)]) #E[,c(j)]
  })))
}
Sig1=solve(var(X))

sum(pairwise.mal(E_simple,Sig1))/(K*(K-1))
sum(pairwise.mal(E,Sig1))/(K*(K-1))
```

This is an example of an implementation of the CBA with *R* code on equal inclusion probabilities on a simulated data set. The argument *landing=FALSE* in the *stratifiedcube* can be changed to true to use the linear programming for the landing phase of the Cube method instead of the elimination of variables. With large data sets, it is preferable to keep this argument as false in order to reduce the computation time.

3.8.2 Additional information on the simulations

The general description of the data set can be found in the main document. For the first set of simulations *Sim1* on the Balcells et al. (2022) data set, the columns were all retained with the exception of the last two that correspond to the treatment groups. Rows containing missing values were dropped.

On the second set of simulations *Sim2* that uses the same data set, similar manipulations are made except that more columns were dropped. The first 6 columns were retained. Those variables correspond to age, gender, ideology, economic situation, level of political interest and level of importance of religion.

The second data set, from Milosevic (2022), originally has 13 variables but only 6 were retained in order to avoid having too many variables for this sample size. One variable was a categorical variable with 3 possible different entries. Columns that were expected to be highly correlated with others were removed. The third set of simulation *Sim3* uses equal inclusion probabilities.

For the simulation set *Sim4*, the same data set is used. The unequal inclusion probabilities of *Sim4* are partially based on the first variable (years of education) and partially generated randomly using a uniform distribution, then adjusted for the margins through raking. The formula to generate the initial inclusion probabilities matrix is the following

$$\left(0.1 \cdot \mathbf{U} + \frac{\mathbf{x}}{2\bar{x}}, \frac{\mathbf{x}}{3\bar{x}}, 1 - \frac{\mathbf{x}}{2\bar{x}}\right)_{75 \times 3},$$

where \mathbf{x} is the first variable of the data set, \bar{x} the mean of this variable and \mathbf{U} is a randomly generated vector through a uniform distribution between 0 and 1. This matrix is adjusted by raking to obtain the correct row and column sums.

Chapter 4

Exact Inference for Poisson Regression With Continuous Covariates using Survey Sampling Algorithms

Abstract

Exact inference algorithms allow us to perform inference when approximation methods for large samples, such as Wald tests and likelihood ratio tests, fail or are unreliable. Exact inference methods for discrete models are generally applied to data sets that can be presented as contingency tables or whose covariates are grouped together to allow such a representation. A recent development in this area concerns a method of exact inference with continuous covariates using the Cube algorithm for logistic regression models. Using a version of the Cube method designed for the case of highly stratified populations, we propose a new algorithm that allows us to perform exact inference for Poisson regression with continuous covariates. ¹

4.1 Introduction

Exact inference approaches can be used for logistic regression and Poisson regression to perform inference in the presence of nuisance parameters. To allow inference of a parameter of interest, the effect of nuisance parameters is eliminated by conditioning on their sufficient statistics. Exact inference conditional on the sufficient statistics of nuisance parameters is mainly used for contingency tables or datasets that can be presented as such. To motivate these approaches, Agresti (2003, p. 158) points out that “when the sample size is small, conditional likelihood-based exact inference in logistic regression is more reliable than the ordinary large-sample inferences”. In addition, an exact inference approach is very useful when the data are separable or quasi-separable, as tests based on large sample approximations perform poorly under these conditions. Even when the sample size is large, large sample approximations can give poor results on contingency tables that contain both small and large expected frequencies (Haberman 1988, Agresti 1992). Methods based on exact inference can be used in parallel with classical methods based on large sample approximations methods at least for the potential insight it might give about the data set and, if needed, replace them.

In this paper, unless otherwise stated, exact inference refers to exact inference conditional on sufficient nuisance parameter statistics for parametric models such as logistic regression and Poisson regression.

Using a version of the Cube method designed to deal with highly stratified populations (see

¹This chapter is a working paper with Prof. Yves Tillé. and Prof. Louis-Paul Rivest

Chauvet, 2009; Hasler and Tillé, 2014; Jauslin et al., 2021), we propose an algorithm to perform exact inference for Poisson regression. This method makes it possible to obtain empirical values p close to the nominal values p under the null hypothesis, even with a relatively small sample size, and under the alternative hypothesis. The statistical power of our method is close to that of asymptotic tests such as the Likelihood Ratio Test (LRT). Our algorithm is also capable of handling continuous covariates.

Let us outline the structure of this article. In Section 4.2, we introduce the basic notation and theory for exact inference. Then we explain the method for exact inference on logistic regression models developed by Rivest and Gaye (2023) in Section 4.3, which should help to understand the algorithm that we propose in Section 4.4. In Section 4.5, we present some simulations with continuous variables. A discussion follows in Section 4.6.

4.2 Notation and Problem Presentation

The notation is based on the one used in Rivest and Gaye (2023). Exact inference procedures apply to models that have sufficient statistics for the nuisance parameters such as the logistic or Poisson regression. We first introduce exact inference with the Poisson distribution. When it comes to exact inference with the Poisson distribution, we can slightly distinguish the problem into two categories. The first is more focused on contingency tables, where we typically perform goodness-of-fit tests using column and row sums. Each cell in the table is assumed to have a Poisson distribution. The row and column sums may or may not be fixed depending on the assumptions or the way the data set was collected, which changes the conditional distribution of the data.

To illustrate, assume we have a contingency table of dimension 2×2 whose first row and column sum are fixed. Assume that each cell of the table is distributed according to a Poisson distribution with mean λ_{ij} . We can assume a model where

$$\log(\lambda_{ij}) = \mu + \mu_i^X + \mu_j^Y + \mu_{ij}, \quad \text{for } i, j \in \{1, 2\}. \quad (4.1)$$

An independence test would test $H_0 : \mu_{ij} = 0$ or more precisely test whether the odds-ratio

$$\theta = \frac{\lambda_{11}\lambda_{22}}{\lambda_{12}\lambda_{21}}$$

is equal to 1. An exact test would reduce to the hyper-geometric distribution (Agresti, 1992).

If we set the parameter of interest to μ_{11} , then its sufficient statistic is the upper left cell. The other parameters (μ, μ_1^X, μ_1^Y) are nuisance parameters whose sufficient statistics are respectively the total of all cells, the total of cells of the first row and the total of the cells of the first column. The conditional distribution of upper left cell follows a central hyper-geometric distribution under the null hypothesis. This problem can be generalized for contingency table of dimension $I \times J$ (Agresti, 1992). In those cases, we would often test the independence model against the saturated model and those tests often reduce to the multivariate hyper-geometric distribution. Diaconis and Sturmfels (1998) describe other more complex problems that the ones shown above, such as three-way contingency tables.

The other case focuses more on a regression model and aims to determine the significance of a parameter/variable in the Poisson regression model. Our algorithm is adapted to this form of problem, even though the first case we have introduced is a special case of “regression-centred” exact inference. Let n be sample size, \mathbf{y} be the vector of the response variable of dimension n , $\mathbf{z} = (z_1, \dots, z_n)^\top$ the vector of the variable of interest and $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^\top$ the confounding variables, where $\mathbf{x}_i \in \mathbb{R}^q$. For convenience, we assume that $\mathbf{z} \in \mathbb{R}^n$. Assume that

$$y_i \sim \text{Pois}(\lambda_i), \quad \lambda_i = \exp\left(\boldsymbol{\beta}^\top \mathbf{x}_i + \gamma z_i\right), \quad \text{for } i \in \{1, \dots, n\}.$$

The aim is to determine whether the parameter of interest γ is significant. We can use the sufficient statistic for the nuisance parameters β , to carry out an exact test by conditioning on $\mathbf{X}^\top \mathbf{y}$ (Caffo and Booth, 2003). This type of approach needs discrete or categorical data, as it has the problem of being dependent on the size of the conditioning set. It is therefore convenient to present exact methods for Poisson regression in terms of an aggregated data set. Let K be the number of different configurations for the explanatory variables \mathbf{x}_i . Let m_k and $y_{\bullet k}^{(a)}$ be respectively the sample size and the sum of the Poisson variables for configuration $\mathbf{x}_k^{(a)}$, $k = 1, \dots, K$. We define the response variable and variable of interest of the unit i in configuration k as respectively $y_{ik}^{(a)}$ and $z_{ik}^{(a)}$, where $i \in \{1, \dots, m_k\}$. We have $y_{\bullet k}^{(a)} = \sum_{i=1}^{m_k} y_{ik}^{(a)}$. The sufficient statistics for the null model ($\gamma = 0$) is a vector of length q , $\sum_{k=1}^K \mathbf{x}_k^{(a)} y_{\bullet k}^{(a)} = \sum_{k=1}^K \sum_{i=1}^{m_k} \mathbf{x}_k^{(a)} y_{ik}^{(a)}$. Exact methods use the conditional distribution of $\{y_{ik}^{(a)} : k = 1, \dots, K; i = 1, \dots, m_k\}$ given that the sufficient statistics for β are fixed and equal to their observed values, $\sum_{k=1}^K \sum_{i=1}^{m_k} \mathbf{x}_k^{(a)} y_{ik}^{(a)}$.

In other words, we assess whether $\mathbf{y}^{(a)} = (y_{11}^{(a)}, \dots, y_{m_1 1}^{(a)}, \dots, y_{1K}^{(a)}, \dots, y_{m_K K}^{(a)})$ or $\sum z_{ik}^{(a)} y_{ik}^{(a)}$ is extreme with respect to the conditional distribution of $\mathbf{y}^{(b)}$ or $\sum z_{ik}^{(a)} y_{ik}^{(b)}$ under the null hypothesis ($\gamma = 0$), where $\mathbf{y}^{(b)}$ is random vector of \mathbb{R}^n distributed according to the conditional distribution of the outcomes given the sufficient statistics for the nuisance parameters. On small data sets, making exact inference based on the conditional distribution $\mathbf{y}^{(b)}$ or $\sum z_{ik}^{(a)} y_{ik}^{(b)}$ can be done by enumerating all the possible cases. When this method is not available because the sample size n is too large, it is possible to perform an ‘‘approximate’’ exact inference using algorithms based on Monte Carlo methods (MCMC) such as the algorithms given in Mehta et al. (2000) or Forster et al. (2003). These methods can handle larger problems than simple enumeration, but the chain may consume too much memory, be reducible or have a very high rejection rate, depending on the method used. An overview of MCMC-based conditional inference methods for logistic regression and Poisson regression is presented in Caffo and Booth (2003). Our proposed algorithm also performs ‘approximate’ exact inference.

We give the conditional distribution of the response variable under the null hypothesis. The support \mathcal{F} of the random vector $\mathbf{y}^{(b)}$ is

$$\mathcal{F} = \left\{ y_{ik}^{(c)}, k = 1, \dots, K \ i = 1, \dots, m_k : y_{ik}^{(c)} \in \{0, 1, \dots\} \right. \\ \left. \text{and } \sum_{k=1}^K \sum_{i=1}^{m_k} \mathbf{x}_k y_{ik}^{(a)} = \sum_{k=1}^K \mathbf{x}_k \sum_{i=1}^{m_k} y_{ik}^{(c)} \right\}.$$

The conditional probability of $\mathbf{y}^{(b)}$ being equal to vector $\mathbf{y}^{(c)} \in \mathcal{F}$ under H_0 is then

$$\Pr \left\{ \mathbf{y}^{(b)} = \mathbf{y}^{(c)} \right\} = \frac{\prod_k \prod_i 1/y_{ik}^{(c)}!}{\sum_{\mathbf{y}^{(a)} \in \mathcal{F}} \prod_k \prod_i 1/y_{ik}^{(a)}!}. \quad (4.2)$$

The distribution described in Equation (4.2) quickly becomes intractable, even with a relatively small sample size n with a non-trivial model. It is also often useful to consider the conditional distribution of the vector of \mathbb{R}^K of the aggregated data for the K configurations of the explanatory variable. The k entry of this vector is $y_{\bullet k}^{(b)} = \sum_{i=1}^{m_k} y_{ik}^{(b)}$. The probability that the vector $\mathbf{y}_{\bullet}^{(b)} \in \mathbb{R}^K$ is equal to $\mathbf{y}_{\bullet}^{(a)}$ can be obtained by summing the probabilities (4.2) of the vector $\mathbf{y}^{(a)}$ whose aggregated version yields $y_{\bullet}^{(a)}$. Indeed one has

$$\sum_{\sum_i y_{ik}^{(b)} = y_{\bullet k}^{(a)}} \frac{1}{\prod_i y_{ik}^{(b)}!} = \frac{m_k^{y_{\bullet k}^{(a)}}}{y_{\bullet k}^{(a)}!},$$

which is easily deduced from the multinomial expansion of $(1/m_k + \dots + 1/m_k)^{y_{\bullet k}^{(a)}}$. Thus the conditional probability of the aggregated vector $\mathbf{y}_{\bullet}^{(b)} \in \mathbb{R}^K$ being equal to $\mathbf{y}_{\bullet}^{(c)} \in \mathcal{F}_{\bullet}$ is

$$\Pr \left\{ \mathbf{y}_{\bullet}^{(b)} = \mathbf{y}_{\bullet}^{(c)} \right\} = \frac{\prod_k m_k^{y_{\bullet k}^{(c)}} / y_{\bullet k}^{(c)}!}{\sum_{\mathbf{y}_{\bullet}^{(a)} \in \mathcal{F}_{\bullet}} \prod_k m_k^{y_{\bullet k}^{(a)}} / y_{\bullet k}^{(a)}!}, \quad (4.3)$$

where \mathcal{F}_{\bullet} is the support of $\mathbf{y}_{\bullet}^{(b)}$ that is deduced from the set \mathcal{F} .

Apart from the assumption of a model and hence the risk of model misspecification, one of the main drawbacks of conditional exact inference approaches is when $K = n$. In this case, the sufficient statistic of $\boldsymbol{\beta}$ completely determines the distribution of $\mathbf{y}^{(b)}$. This situation generally occurs when the covariates are continuous. Using too many nuisance parameters in the model with an insufficient sample size can also lead to the same problem. It can also lead to a sample space that satisfies the constraints of sufficient statistics, but is too small to allow reliable inference. The distribution of covariates, even when they are discrete or categorical, can also pose difficulties. This is a common problem for all conditional exact inference methods. To counteract the problem of continuous covariates, there is the possibility of grouping the values of continuous variables corresponding to nuisance parameters (Potter, 2005). However, changing the covariates by rounding them to the closest integer or decimal or any perturbation of the covariates can greatly change the conditional sample space which satisfy the sufficient statistic. Brazzale and Davison (2008) gives a good example of the sensitivity of the sample space which satisfy the sufficient statistic constraint. As far as confidence intervals are concerned, Agresti and Coull (1998) and Brazzale and Davison (2008, p. 8) state that the use of the Wald confidence interval is preferable as they are closer to nominal levels than confidence intervals derived from exact procedures which tend to be far too conservative.

A recent development in the field concerns an exact method for logistic regression that has made it possible to perform exact inference in the presence of continuous covariates without clustering the continuous variables. This method is based on a survey sampling algorithm which is described in Section 4.3.

4.3 The Cube method applied for exact inference in logistic regressions

In this section, we introduce the Cube method applied for exact inference in logistic regressions, which should make the extension we describe in the next section easier to understand. Here, the response variable corresponds to a binary variable distributed according to a logistic regression model. Define the following model:

$$y_i \sim \text{Bern}(\pi_i), \quad \pi_i = \frac{\exp(\mathbf{x}_i^{\top} \boldsymbol{\beta} + z_i \gamma)}{1 + \exp(\mathbf{x}_i^{\top} \boldsymbol{\beta} + z_i \gamma)}, \quad i = 1, \dots, n,$$

where \mathbf{x}_i and $\boldsymbol{\beta}$ are vectors of \mathbb{R}^q . Usually, the goal is to perform inference on the parameter of interest γ . Let $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^{\top}$ be the confounding variables. The null hypothesis is $H_0 : \gamma = 0$ and the alternative is $H_1 : \gamma \neq 0$. Methods based on conditional exact inference are usually used for contingency tables and that logistic regression models have a binary response variable. Therefore, let us represent our data set as a $K \times 2$ contingency table, where K is the number of possible configurations. For $k \in \{1, \dots, K\}$, let $\mathbf{x}_k^{(a)}$ represent the different possible configurations of the confounding variables. Let $z_{ik}^{(a)}$ be the variable of interest of unit i in configuration k . Define m_k and $y_{\bullet k}^{(a)}$ respectively as the number of Bernoulli trials and the number of successes for configuration $\mathbf{x}_k^{(a)}$. Consequently, we have $y_{\bullet k}^{(a)} \in \{1, \dots, m_k\}$ and

$y_{ik}^{(a)} \in \{0, 1\}$ for the corresponding unit i in the configuration. Essentially, the conditional exact inference methods consist in analysing the distribution of the sufficient statistic of $\gamma : \sum z_k^{(a)} y_{\bullet k}^{(a)}$ conditioned on the sufficient statistic for β , the nuisance parameter, being equal to $\sum \mathbf{x}_k^{(a)} y_{\bullet k}^{(a)}$. More precisely, let $\mathbf{y}_{\bullet}^{(b)} = (y_{\bullet 1}^{(b)}, \dots, y_{\bullet K}^{(b)})^\top$ be a random vector with the distribution written in (4.3) under the null hypothesis but with the restriction that $\sum_{k=1}^K \mathbf{x}_k^{(a)} y_{\bullet k}^{(b)} = \sum_{k=1}^K \mathbf{x}_k^{(a)} y_{\bullet k}^{(a)}$. As explained in Caffo and Booth (2003); Zamar et al. (2007), this amounts to analysing the distribution of $\mathbf{y}_{\bullet}^{(b)}$ given the sufficient statistic for β , where the density is

$$f\left(\mathbf{y}_{\bullet}^{(b)} \mid \sum_{k=1}^K \mathbf{x}_k^{(a)} y_{\bullet k}^{(b)} = \sum_{k=1}^K \mathbf{x}_k^{(a)} y_{\bullet k}^{(a)}\right) \propto \prod_{k=1}^K \binom{m_k}{y_{\bullet k}^{(b)}} \exp(\gamma \mathbf{z}^\top \mathbf{y}_{\bullet}^{(b)}). \quad (4.4)$$

Again, we assess whether $\mathbf{y}_{\bullet}^{(a)}$ or $\sum z_{ik}^{(a)} y_{ik}^{(a)}$ is extreme with respect to the conditional distribution of $\mathbf{y}_{\bullet}^{(b)}$ or $\sum z_{ik}^{(a)} y_{ik}^{(b)}$ under the null hypothesis ($\gamma = 0$). The problems of the conditioning set is the same as with the Poisson regression. When $K = n$, which usually happens with continuous variables, inference is not possible as the sufficient statistic of nuisance parameters determines the conditional distribution. Also, calculating the probability in (4.4) by enumeration is often not possible when the sample size is large.

Rivest and Gaye (2023) proved that, under a few assumptions, $y_{\bullet k}^{(b)}/m_k$ converges in probability to the logistic regression fitted probabilities for the configuration k when n tends to infinity for $k \in \{1, \dots, K\}$. In other words, the proportion of positive responses $y_{\bullet k}^{(b)}/m_k$ in the configuration $\mathbf{x}_k^{(a)}$ converges to the fitted logistic regression probabilities. The authors used this result to justify the use of the Cube method developed by Deville and Tillé (2004b) in the context of conditional exact inference in logistic regression. The main advantage of this method is that it can be used even when the covariates are continuous. This method is limited to trying to approximately satisfy the sufficient statistic constraints. The Cube method is already used for the Horvitz–Thompson estimator (Horvitz and Thompson, 1952) of the totals in the context of survey sampling. Rivest and Gaye (2023) has used it for exact inference in logistic regression. The authors also adapted this approach for tests of partial association.

Rivest and Gaye (2023) reframed the exact inference approach as a sampling problem with a population of n units, unequal selection probabilities and balancing constraints. For a population of size n , a random sample can be denoted by a vector of Bernoulli random variables $\mathbf{a} = (a_1, \dots, a_n)^\top \in \{0, 1\}^n$. Let $\pi_i = \Pr(a_i = 1)$ denote inclusion probability for a unit i . Let also $\mathbf{x}_i = (\mathbf{x}_{i1}, \dots, \mathbf{x}_{ip})^\top$ be the vector of auxiliary variables associated for unit i . If $\pi_i > 0$, for all $i \in \{1, \dots, n\}$, the Horvitz–Thompson estimator

$$\sum_{i=1}^n \frac{a_i \mathbf{x}_i}{\pi_i}$$

is an unbiased estimate of the vector of the total of the auxiliary variables

$$\sum_{i=1}^n \mathbf{x}_i.$$

A random sample \mathbf{a} is said to be balanced if it satisfies the balancing equations

$$\sum_{i=1}^n \frac{a_i \mathbf{x}_i}{\pi_i} \approx \sum_{i=1}^n \mathbf{x}_i. \quad (4.5)$$

The Cube method allows one to select a sample \mathbf{a} which satisfy Equation (4.5). Generally, it is not possible to satisfy exactly those constraints and so those are only approximately satisfied

by the Cube method.

Using the notation from Section 4.1, the authors proposed to use the sample \mathbf{a} obtained through the Cube algorithm to generate a new response variable $\mathbf{y}_b = (y_{b_1}, \dots, y_{b_n})^\top \in \{0, 1\}^n$, use the fitted logistic regression probabilities $\hat{\pi}_1, \dots, \hat{\pi}_n$ as inclusion probabilities in the algorithm and $\mathbf{x}_i^\pi = \hat{\pi}_i \mathbf{x}_i$ as auxiliary variables in the Cube algorithm. This means that the Cube algorithm satisfy the following constraints:

$$\sum_{i=1}^n \frac{y_{b_i} \mathbf{x}_i^\pi}{\hat{\pi}_i} \approx \sum_{i=1}^n \mathbf{x}_i^\pi. \quad (4.6)$$

Simplifying Equation (4.6), we obtain $\sum_{i=1}^n \mathbf{x}_i (y_{b_i} - \hat{\pi}_i) \approx 0$. When we simplify the problem by adopting the “ K different configurations” point of view, this is also equivalent to

$$\sum_{k=1}^K \mathbf{x}_k^{(a)} (y_{\bullet k}^{(b)} - m_k \hat{\pi}_k^{(a)}) \approx 0, \quad (4.7)$$

where $\hat{\pi}_k^{(a)}$ corresponds to the fitted logistic regression success probability for the samples in the configuration $\mathbf{x}_k^{(a)}$.

Then, given that n is sufficiently large, Proposition 1 from Rivest and Gaye (2023) ensures that the generated $\mathbf{y}^{(b)}$ have a sufficient statistic for the nuisance parameters $\mathbf{X}^\top \mathbf{y}_b$ that is approximately equal to the sufficient statistic of the nuisance parameters $\mathbf{X}^\top \mathbf{y}$ of the original observed \mathbf{y} . However, this proposition also implies that the method proposed is unreliable if the sample size n is too small. To summarize, the method allows us to sample from the conditional distribution of \mathbf{y}_b if n is sufficiently large. The generated sample \mathbf{y}_b does not sufficiently satisfy Equation (4.6) in the context of exact inference. To handle this problem, Rivest and Gaye (2023) uses a tolerance parameter tol and rerun the algorithm when the L1 norm of $\mathbf{X}^\top (\mathbf{y} - \mathbf{y}_b)$ is greater than tol .

To perform the inference, we calculate the sufficient statistic of γ ; $T = \mathbf{z}^\top \mathbf{y}$ and we generate B times a new response variable \mathbf{y}_b under the null hypothesis to record the sufficient statistic T_{bi} for the parameter of interest with $i \in \{1, \dots, B\}$. Then, let n_1 be the cardinal of the set $\{i : T_{bi} < T\}$ and n_2 be the cardinal of the set $\{i : T_{bi} > T\}$ to determine how extreme T is under the null hypothesis. We obtain a test statistic whose p -value is $P = 2 \min(n_1, n_2)/B$.

4.4 Exact Inference for Poisson Regression Models using a Survey Sampling Algorithm

4.4.1 Fundamental Results

We extend exact inference based on the Cube method for Poisson regressions. The extension we propose is developed under the same principles as the exact logistic regression from Rivest and Gaye (2023). The authors give a theoretical justification for his approach through a proof of convergence in probability of the fitted logistic regression under a few assumptions. For the Poisson distribution, we will use Proposition 1, which mainly states that the fitted Poisson means, given the sufficient statistic, asymptotically converge to the mean of each configuration.

Let $\hat{\boldsymbol{\tau}}^{(n)} \in \mathbb{R}^K$ be the vector with entries equal to $\hat{\tau}_k = y_{\bullet k}^{(b)}/m_k$, $k = 1, \dots, K$. The number of configurations K for the explanatory variables is constant and $m_k^{(n)}/n$ is assumed to converge to a positive number as n goes to ∞ for $k = 1, \dots, K$. The vector of the independent variables $\mathbf{x}_k^{(a)}$ does not depend on n and the vector of relative sufficient statistics $\sum_{k=1}^K \mathbf{x}_k^{(a)} y_{\bullet k}^{(a)}/n \in \mathbb{R}^q$ is assumed to converge to a fixed vector of constant as n goes to ∞ . The intercept is included in the covariates $\mathbf{x}_k^{(a)}$. We assume that, without conditioning on the sufficient statistic, $y_{ik}^{(b)}$ follows

a Poisson distribution of mean

$$\lambda_k^{(n)} = \exp\left(\boldsymbol{\beta}^\top \mathbf{x}_k\right), \quad \text{for any } i = 1, \dots, m_k.$$

Let $\{\hat{\lambda}_k^{(n)} : k = 1, \dots, K\}$ be the fitted Poisson parameters, defined such that $\hat{\lambda}_k^{(n)}$ is the estimated Poisson mean for $y_{ik}^{(b)}$, for any $i = 1, \dots, m_k$. The Poisson estimates of $\boldsymbol{\beta}$ only depend on the vector of sufficient statistics $\sum_{k=1}^K \mathbf{x}_k^{(a)} y_{\bullet k}^{(a)}$, for a given value of n .

Proposition 1. *In addition to the assumptions above, as the sample size n goes to ∞ in such a way that $\lim_n m_k^{(n)}/n > 0$ and $\lim_n \hat{\lambda}_k^{(n)}$ converge to a positive value, for $k = 1, \dots, K$, one has the following convergence in probability, for any $\varepsilon > 0$,*

$$\lim_{n \rightarrow \infty} Pr \left\{ \sum_{k=1}^K \left(\hat{\tau}_k - \hat{\lambda}_k^{(n)} \right)^2 < \varepsilon \right\} = 1.$$

where Pr refers to the discrete probability defined by (4.3).

Proposition 1 is proved in the Appendix. The proof and assumptions are very similar to the proof and assumptions used by Rivest and Gaye (2023) for the equivalent result for the logistic regression. The additional assumption is the presence of an intercept to bound the size of the support and some other term. Proposition 1 is useful because we use this Poisson mean for the conditional distribution of the marginals in each configuration. We remind that Proposition 1 only states that the mean of each configuration will converge to the estimated Poisson means under a few assumptions. It does not state that the marginal conditional distribution of a unit in configuration k will have a Poisson distribution.

With correct input parameters, the Cube algorithm allows us to generate response variables that satisfy the sufficient statistic constraints while respecting the distribution of certain marginal distribution. The procedure we propose in the next section carries out exact inference for Poisson regression by generating the sum of the response variables in each configuration $\mathbf{y}_{\bullet}^{(b)}$.

The main problem is the choice of the conditional distribution of the marginals of $\mathbf{y}_{\bullet}^{(b)}$, which should be unknown. A candidate to approximate this distribution would be the Poisson distribution, as the sum of independent Poisson variables also has a Poisson distribution. However, a Poisson distribution has an unbounded support while the conditional distribution of $\mathbf{y}_{\bullet k}^{(b)}$ has finite support. The vector $\mathbf{y}_{\bullet k}^{(b)}$ has a distribution that exactly corresponds to a binomial distribution when the only explanatory variable is the intercept. For this reason, the binomial distribution with proper parameters is a better approximation than a Poisson distribution. It also converges to a Poisson distribution as the sample size tends infinity.

The classic example of a 2×2 contingency table under the independence model (4.1) corresponds to a case with $K = 4$ configurations. Each cell of the table is a configuration of the explanatory variables of size 1. This is an example where our algorithm would not output the exact distribution if we use a Poisson or binomial distribution for the marginals. It is known that the conditional distribution of the marginal in each configuration has a hyper-geometric distribution. However, in this example, it is possible to increase the number of units in each configuration. If we increase each configuration size equally, we can find with quite easily that again the conditional distribution of the marginals is hyper-geometric. However, as the configuration sizes increase to infinity, the distribution of the marginals will converge to a binomial distribution.

4.4.2 Algorithm with K configurations

We represent our data into $K \leq n$ configurations. Define the following model:

$$y_{ik}^{(a)} \sim \text{Pois}(\lambda_{ik}), \quad \lambda_{ik} = \exp \left\{ \log(t_{ik}) + \boldsymbol{\beta}^\top \mathbf{x}_k + \gamma z_{ik} \right\},$$

for $k \in \{1, \dots, K\}, i \in \{1, \dots, m_k\}$. (4.8)

We assume that $\mathbf{z} = (z_{11}, \dots, z_{m_1 1}, \dots, z_{m_k K}) \in \mathbb{R}^n$. We set γ as the parameter of interest and $\boldsymbol{\beta}$ as the nuisance parameters. Let K be the number of configuration for the explanatory variables, where \mathbf{x}_k corresponds to the configuration k and $y_{ik}^{(a)}$ is the response variable of the i -th the unit in configuration k . When we model rates rather than count data, using an offset is useful. We denote the offset by t_{ik} . If there is no offset, we set t_k to 1, for all $k \in \{1, \dots, K\}, i \in \{1, \dots, m_k\}$. In this section, we used the notation λ_{ik} instead of λ_k to accommodate for the presence of offset effects. We generate $\mathbf{y}_\bullet^{(b)}$ in order to satisfy

$$\mathbf{x}_1^{(a)} y_{\bullet 1}^{(b)} + \dots + \mathbf{x}_K^{(a)} y_{\bullet K}^{(b)} = \mathbf{x}_1^{(a)} y_{\bullet 1}^{(a)} + \dots + \mathbf{x}_K^{(a)} y_{\bullet K}^{(a)}. \quad (4.9)$$

If the set \mathcal{F}_\bullet is easily computable, then we can just use the equation (4.3) to calculate the distribution of $\mathbf{y}_\bullet^{(b)}$. Otherwise, we propose to use the Cube method to generate $\mathbf{y}_\bullet^{(b)}$ that satisfy the constraints in (4.9).

The objective is to test whether the parameter γ is significant. We carry out an exact inference procedure using a modified version Cube method designed to select samples from stratified populations. The Cube function used is *stratifiedcube* from the *StratifiedSampling* package (Jauslin et al., 2022). This function allows to select samples with a large amount of strata and constraints. It is useful, for clarity, to define the input parameters that are to be chosen by the user:

- S : The maximum strata size. It is preferable to set it by default to $\sum_{k=1}^K y_{\bullet k}^{(a)}$.
- r : The number of simulated by dis-aggregating $\mathbf{y}_\bullet^{(b)}$ using a multinomial distribution.
- B : The total number of simulated $\mathbf{y}^{(b)}$.
- F_λ : A positive discrete distribution whose mean is λ and whose support is $\{0, 1, \dots, S\}$. By default, we propose a binomial distribution.
- tol : The tolerance criterion for acceptance of a sample.

The procedure works as follows:

Step 1: Estimate the Poisson means $\hat{\lambda}_{\bullet 1}, \dots, \hat{\lambda}_{\bullet k}$ under the null hypothesis $\gamma = 0$ using a maximum likelihood estimator, where

$$\hat{\lambda}_{\bullet k} = \exp \left(\sum_{i=1}^{m_k} \log(t_{ik}) + m_k \hat{\boldsymbol{\beta}}^\top \mathbf{x}_k \right), \quad \text{for } k \in \{1, \dots, K\}.$$

Step 2: For $k \in \{1, \dots, K\}$ and $i \in \{1, \dots, S + 1\}$, calculate $p_{ik} = \Pr(u_k = i)$ with $u_k \sim F_{\hat{\lambda}_{\bullet k}}$. If we use a binomial distribution for $F_{\hat{\lambda}_{\bullet k}}$, the parameters are $\sum_{k=1}^K y_{\bullet k}^{(a)}$ for the size and

$$\frac{\hat{\lambda}_{\bullet k}}{\sum_{k'=1}^K \hat{\lambda}_{\bullet k'}}$$

for the mean.

Step 3: Define K strata in the following manner: for $k \in \{1, \dots, K\}$ set the inclusion probabilities $\mathbf{p}_k = (p_{1k}, \dots, p_{(S+1)k})^\top$. Figure 4.1 illustrates the assignment of the strata.

Step 4: Define $\mathbf{v}_k = (0, \dots, S)^\top$ and set $\mathbf{D}_k = \mathbf{v}_k \mathbf{x}_k^\top$, for $k \in \{1, \dots, K\}$. Define \mathbf{D}_k^π whose column j is equal to the element-wise product of \mathbf{D}_k j -th column and \mathbf{p}_k . Set \mathbf{D}_i^π as auxiliary variables or constraints for each stratum $k \in \{1, \dots, K\}$ in the *stratifiedcube* algorithm. Also, add a column of inclusion probabilities as a constraint to ensure we sample exactly one element per strata. Additionally, this ensures that the generated response variable has length K . This step mirrors the use of the Cube algorithm for logistic regression in equation (4.6).

Step 5: Repeat the following until B simulations have been generated and accepted. When a sample is accepted, record the generated sufficient statistic of the parameter of interest T_{bj} , for $j \in \{1, \dots, B\}$.

- Generate $\mathbf{y}_\bullet^{(b)} = y_{\bullet 1}^{(b)}, \dots, y_{\bullet K}^{(b)}$ by following the process described in Step 1 to Step 4, which samples one element per stratum with the Cube method respecting the inclusion probabilities and balancing constraints.
- Accept a generated sample $\mathbf{y}_\bullet^{(b)}$ if the L1 norm of $\mathbf{X}_K^\top (\mathbf{y}_\bullet^{(b)} - \mathbf{y}_\bullet^{(a)})$ is smaller than *tol*. Repeat until a sample $\mathbf{y}_\bullet^{(b)}$ is accepted. When a sample is accepted, dis-aggregate $\mathbf{y}_{\bullet k}^{(b)}$ using a multinomial distribution of size m_k and with parameters

$$\frac{\hat{\lambda}_{ik}}{\sum_{i=1}^{m_k} \hat{\lambda}_{ik}}, \quad (4.10)$$

for $k \in \{1, \dots, K\}$. The parameters in (4.10) reduce to $1/m_k$ if there is no offset. Repeat it r times, to generate r simulated $\mathbf{y}^{(b)}$.

- Record as T_{bj} the corresponding sufficient statistic for every simulated $\mathbf{y}^{(b)}$ the parameter of interest: $\sum z_{ik}^{(a)} y_{ik}^{(b)}$.

Step 6: Let n_1 be the cardinal of the set $\{j : t < T_{bj}\}$ and n_2 be the cardinal of the set $\{j : t > T_{bj}\}$. Calculate the p -value $P = 2 \min(n_1, n_2)/B$.

A trade-off between the accuracy of the p -value and the computation time can be made by changing the values of *tol*, B . The *stratifiedcube* function does not always output samples that are sufficiently balanced to get accurate results, which induces the necessity to reject the samples which are not balanced enough. In our simulations, we have used a tolerance of $tol = 0.15 \sqrt{\sum_{i=1}^n |\mathbf{x}_i^\top y_i|}$, and around 25% are accepted with this level of tolerance. Using rejective sampling after the use of the Cube method improves the fit of the constraints of the sufficient statistic but might not be ideal because the inclusion probabilities might not be exactly respected. The choice of F_λ and S can also change the accuracy and computation time. The reasoning of the procedure also functions when $K = n$. In this case, the step of dis-aggregating $\mathbf{y}^{(b)}$ with a multinomial distribution becomes trivial.

4.4.3 Application on Data Sets

To check that the procedure agrees with existing exact Poisson regression algorithms, we first compare the p -values obtained with *expoisson* from the statistical programming software Stata on the *Cerebacc* data set that contains “the number of cerebrovascular accidents experienced by 41 men during a fixed period, each of whom had recovered from a previous cerebrovascular accident and was hypertensive” (StataCorp, 2023). The number of cerebrovascular accidents is the response variable and is assumed to follow a Poisson distribution. We also have two

\mathbf{p}_1	$y_{\bullet 1}^{(b)}$	\mathbf{p}_2	$y_{\bullet 2}^{(b)}$	\dots	\mathbf{p}_3	$y_{\bullet K}^{(b)}$
p_{11}	0	p_{12}	0		p_{1K}	0
p_{21}	1	p_{22}	1		p_{2K}	1
p_{31}	2	p_{32}	2		p_{3K}	2
p_{41}	3	p_{42}	3		p_{4K}	3
p_{51}	4	p_{52}	4		p_{5K}	4
p_{61}	5	p_{62}	5		p_{6K}	5

Figure 4.1: Illustration of the stratification in the algorithm. Each column represents a stratum with different auxiliary variables and their inclusion probabilities. One unit from each stratum is selected with the sufficient statistic constraints considered.

indicator variables for the treatment group and the age group (> 60). When the age is set as the confounding variable and the treatment as the variable of interest, our algorithm outputs a p -value of 0.002388297 and the *expoisson* outputs 0.0026. When the treatment is set as the confounding variable and age as the variable of interest, our algorithm outputs a p -value of 0.2822434 and the *expoisson* outputs a p -value of 0.2794.

We also compare the p -values which are obtained by the procedure and the ones obtained with the *GENMOD* procedure from SAS. The *GENMOD* allows us to carry out exact inference on Poisson regression models. The *GENMOD* exact procedure uses the multivariate shift algorithm from Hirji et al. (1987) We use the ingots dataset, which is displayed on Table 4.1, is taken from Cox and Snell (1989) and is used in the SAS documentation related to exact inference with the *GENMOD* procedure. This is a relatively small data set with $n = 19$ which concerns groups of ingots that are prepared with different heat and soaking times. The objective is to determine the number of ingots which are not ready using a Poisson regression model with the total per group of ingots as offset. In other words, we model the ratio of the number of ingots which are not ready divided by the total.

The confounding variable is set as a categorical variable and the variable of interest as a quantitative variable to make sure the *GENMOD* procedure works correctly and also for simplicity. When the soaking time variable is used as a confounding variable and the heating time p -value is calculated using exact methods, the p -value using *GENMOD* is 0.0004 and for our algorithm the p -value is 0.00095 with $B = 10000$. In the reversed situation, with the heating time as the confounding variable and soaking time as variable of interest, the p -value using *GENMOD* is 0.8645 and for our algorithm the p -value obtained is 0.8784 with $B = 10000$. Those results show that our algorithm and the *GENMOD* procedure seem to generally agree even though our algorithm relies on the sample size n not being too small.

4.5 Simulations: Poisson regression with $K = n$

In this section, we compare our algorithms to the LRT and the Wald test in setting with continuous variables, with $K = n$. We observe at the empirical p -values against the nominal p -values in the null hypothesis case. We also compare the power of the different tests. For this set of simulations, we generate the nuisance variables with \mathbf{X} of dimension $n \times p$ where $n = 75$ or 25 and $p = 2$. For the matrix \mathbf{X} , the first column has only ones as entries for the intercept and the second column generated from a normal distribution. The variable of interest \mathbf{z} is generated with a normal distribution independent of \mathbf{X} . On average, the size of the strata is around 11 and the maximum strata size would almost always go over 100 but the computation time would

Table 4.1: Ingots data set

Heat	Soak	NotReady	Total	lnTotal
7	1	0	10	2.30
14	1	0	31	3.43
27	1	1	56	4.03
51	1	3	13	2.56
7	1.7	0	17	2.83
14	1.7	0	43	3.76
27	1.7	4	44	3.78
51	1.7	0	1	0
7	2.2	0	7	1.95
14	2.2	2	33	3.50
27	2.2	0	21	3.04
51	2.2	0	1	0
7	2.8	0	12	2.48
14	2.8	0	31	3.43
27	2.8	1	22	3.09
51	4	0	1	0
7	4	0	9	2.20
14	4	0	19	2.94
27	4	1	16	2.77

still be low for a single simulation. The computation time still low, even with more than 50000 entries. We simulate $NB_{sim} = 2500$ data sets with $B = 500$ and $tol = 0.15\sqrt{\sum_{i=1}^n |\mathbf{x}_i^\top y_i|}$. Table 4.2 show that under the null hypothesis $H_0 : \gamma = 0$, the empirical p -values are close to the nominal p -value when $n = 75$, and are less close when $n = 25$. But our algorithm seem to perform better than the LRT and Wald test. In the alternative hypothesis $H_1 : \gamma \neq 0$, the power of the Cube method is slightly weaker than the LRT.

4.6 Discussion

The algorithm that we propose is an extension of the algorithm proposed in Rivest and Gaye (2023), which enables us to handle Poisson regression models with exact methods based on the Cube method. With this extension, it should be possible to perform exact conditional inference on any discrete model whose sufficient statistic is $\mathbf{X}^\top \mathbf{y}$ using the Cube method, and with higher accuracy for the case when the number of configurations K is lower than n if the distribution of the marginal is known or well approximated. Exact inference based on the Cube method has a longer computation time than LRT or Wald tests, which is to be expected as exact methods are simulation based (or require to enumerate all the samples which satisfy the sufficient statistic).

The advantages of classical exact methods on contingency tables should be retained with the exact approaches based on the Cube method while also being applicable on data sets with continuous covariates. In our simulations, the algorithm that we proposed yielded empirical p -values that were closer to their nominal levels than the Wald test or the LRT.

The method assumes that the Cube algorithm generates the sufficient statistics according to their conditional distribution sufficiently closely and further investigations might be needed to tell whether such method provides reliable parameter estimates. Further investigation might also help us choose a more suitable or perhaps the correct marginal distribution in the algorithm we propose.

Table 4.2: Empirical p -values against nominal p -values for each method in a Poisson regression model for $n = 75$ and $n = 25$. Wald corresponds to the Wald test, LRT represents the likelihood ratio test and Scube represents our algorithm based on the *stratifiedcube* function.

$n = 75$

$H_0 : \gamma = 0$	0.5	0.1	0.05	0.01
LRT	0.508	0.0988	0.0468	0.0088
Wald	0.5052	0.092	0.042	0.0068
Scube	0.4952	0.0944	0.0504	0.01
$H_1 : \gamma \neq 0$	0.5	0.1	0.05	0.01
LRT	0.764	0.3724	0.272	0.1216
Wald	0.762	0.3652	0.2596	0.1088
Scube	0.7624	0.3644	0.2696	0.1216

$n = 25$

$H_0 : \gamma = 0$	0.5	0.1	0.05	0.01
LRT	0.5116	0.1164	0.0572	0.008
Wald	0.5036	0.0988	0.0412	0.0032
Scube	0.4908	0.1032	0.0552	0.012
$H_1 : \gamma \neq 0$	0.5	0.1	0.05	0.01
LRT	0.7472	0.3776	0.2616	0.1156
Wald	0.742	0.3492	0.2248	0.078
Scube	0.7416	0.3476	0.2436	0.1112

4.7 Appendix

4.7.1 A preliminary result on Poisson regression estimates

The goal of this section is to show that the estimated predicted values for the aggregated totals,

$$\hat{\lambda}_{\bullet k} = m_k^{(n)} \exp\left(\mathbf{x}_k^\top \hat{\boldsymbol{\beta}}\right), \quad k = 1, \dots, K, \quad (4.11)$$

where $\hat{\boldsymbol{\beta}}$ is the maximum likelihood estimator of $\boldsymbol{\beta}$, maximizes an entropy like function under the q constraints $\sum_{k=1}^K y_{\bullet k}^{(c)} \mathbf{x}_k = \sum_{k=1}^K y_{\bullet k}^{(a)} \mathbf{x}_k$. The function depends on $\mathbf{y}_{\bullet}^{(c)} = (y_{\bullet 1}^{(c)}, \dots, y_{\bullet K}^{(c)})^\top$; it is defined by

$$\mathcal{I}_0(\mathbf{y}_{\bullet}^{(c)}) = -\frac{1}{n} \sum_{k=1}^K \left\{ y_{\bullet k}^{(c)} \log \left(\frac{y_{\bullet k}^{(c)}}{m_k^{(n)}} \right) - y_{\bullet k}^{(c)} \right\}.$$

This function is concave and its maximum is determined by considering the following Lagrange function,

$$\mathcal{L}(\mathbf{y}_{\bullet}^{(c)}, \boldsymbol{\beta}) = -\frac{1}{n} \sum_{k=1}^K \left\{ y_{\bullet k}^{(c)} \log \left(\frac{y_{\bullet k}^{(c)}}{m_k^{(n)}} \right) - y_{\bullet k}^{(c)} \right\} + \frac{1}{n} \sum_{k=1}^K \boldsymbol{\beta}^\top \mathbf{x}_k (y_{\bullet k}^{(c)} - m_k^{(n)} \lambda_k),$$

where $\boldsymbol{\beta} \in \mathbb{R}^q$ is a vector of Lagrange multipliers. The partial derivative of $\mathcal{L}(\mathbf{y}_{\bullet}^{(c)}, \boldsymbol{\beta})$ with respect to $y_{\bullet k}^{(c)}$ is

$$\frac{\partial \mathcal{L}(\mathbf{y}_{\bullet}^{(c)}, \boldsymbol{\beta})}{\partial y_{\bullet k}^{(c)}} = -\frac{1}{n} \log \left(\frac{y_{\bullet k}^{(c)}}{m_k^{(n)}} \right) + \frac{\boldsymbol{\beta}^\top \mathbf{x}_k}{n}.$$

By setting this expression to 0, we can see that the maximum is attained as $y_{\bullet k}^{(c)} = m_k^{(n)} \exp(\boldsymbol{\beta}^\top \mathbf{x}_k)$, where the parameter vector $\boldsymbol{\beta}$ is determined by the constraints $\sum_{k=1}^K y_{\bullet k}^{(c)} \mathbf{x}_k = \sum_{k=1}^K y_{\bullet k}^{(a)} \mathbf{x}_k$. These are the score functions for the parameter $\boldsymbol{\beta}$ in the Poisson regression of $\{y\}$ on $\{x\}$ and the maximum likelihood estimators $\hat{\mathbf{y}}_{\bullet} = (\hat{\lambda}_{\bullet 1}, \dots, \hat{\lambda}_{\bullet K})$ given in (4.11) maximizes $\mathcal{I}_0(\mathbf{y}_{\bullet}^{(c)})$.

For the proof in the next section, it is convenient to consider a modified function,

$$\mathcal{I}_1(\mathbf{y}_{\bullet}^{(c)}) = \mathcal{I}_0(\mathbf{y}_{\bullet}^{(c)}) + c_0 \sum_{k=1}^K \left(\frac{y_{\bullet k}^{(c)}}{m_k^{(n)}} - \hat{\lambda}_k \right)^2. \quad (4.12)$$

where $0 < c_0$ is a small positive constant. For n large enough, the matrix of second order partial derivatives of $\mathcal{I}_1(\mathbf{y}_{\bullet}^{(c)})$, that is the diagonal matrix of $-1/\{ny_{\bullet k}^{(c)}\} + 2c_0/(m_k^{(n)})^2$, $k = 1, \dots, K$, is negative definite as its diagonal entries are negative. Thus, $\mathcal{I}_1(\mathbf{y}_{\bullet}^{(c)})$ is a concave function that can be maximized using Lagrange multipliers as outlined in this section. Its maximum is obtained at $\hat{\mathbf{y}}_{\bullet} = (\hat{\lambda}_{\bullet 1}, \dots, \hat{\lambda}_{\bullet K})$, defined in (4.11).

4.7.2 Proof of Proposition 1

To approximate (4.2), we use Stirling's approximation that provides the following bounds for $x!$, valid for any $x > 0$,

$$e\sqrt{xx}e^{-x} \geq x! \geq \sqrt{2\pi xx}e^{-x}. \quad (4.13)$$

This implies the conditional probability function (4.3) satisfies

$$\begin{aligned}
\Pr\{\mathbf{y}_{\bullet}^{(b)} = \mathbf{y}_{\bullet}^{(c)}\} &\sim \prod_k \binom{m_k^{(n)} y_{\bullet k}^{(c)}}{y_{\bullet k}^{(c)}} \\
&\approx \exp \sum_{k=1}^K \left\{ -y_{\bullet k}^{(c)} \log \left(\frac{y_{\bullet k}^{(c)}}{m_k^{(n)}} \right) + y_{\bullet k}^{(c)} - \frac{\log(y_{\bullet k}^{(c)})}{2} \right\} \\
&= \exp \left\{ n\mathcal{I}_1(\mathbf{y}_{\bullet}^{(c)}) - nc_0 \sum_{k=1}^K \left(\frac{y_{\bullet k}^{(c)}}{m_k^{(n)}} - \hat{\lambda}_k^{(n)} \right)^2 - \sum_{k=1}^K \frac{\log(y_{\bullet k}^{(c)})}{2} \right\},
\end{aligned}$$

where \approx means that the ratio of the left hand side over the hand side belongs to a fixed positive interval (C_1, C_2) that does not depend on $y_{\bullet k}^{(c)}$.

Consider now

$$\Pr \left\{ \sum_{k=1}^K (\hat{\tau}_k - \hat{\lambda}_k^{(n)})^2 > \varepsilon \right\}, \quad (4.14)$$

where $\hat{\tau}_k = y_{\bullet k}^{(b)}/m_k^{(n)}$. The numerator for the probability in (4.14), given by (4.3), involves at most $\left(\sum_{k=1}^K y_{\bullet k}^{(a)} \right)^K$ terms because the sufficient statistic is also conditioned on the intercept. Given that we also have the assumption that the vector of sufficient statistics divided by n converges to a constant, we deduce that $\left(\sum_{k=1}^K y_{\bullet k}^{(a)} \right)$ is $O(n)$. An upper bound for the numerator is

$$\begin{aligned}
&C_2^K \exp \left\{ K \log \left(\sum_{k=1}^K y_k^{(a)} \right) + n \max \mathcal{I}_1(\mathbf{y}_{\bullet}^{(c)}) - \min \left[nc_0 \sum_{k=1}^K \left(\frac{y_{\bullet k}^{(c)}}{m_k^{(n)}} - \hat{\lambda}_k^{(n)} \right)^2 \right] \right\} \\
&< C_2^K \exp \left\{ K \log \left(\sum_{k=1}^K y_k^{(a)} \right) + n\mathcal{I}_0(\hat{\mathbf{y}}_{\bullet}) - nc_0\varepsilon \right\}
\end{aligned} \quad (4.15)$$

where the max and the min are taken on the set of $\mathbf{y}_{\bullet k}$ satisfying

$$\sum_{k=1}^K \left(\frac{y_{\bullet k}^{(c)}}{m_k^{(n)}} - \hat{\lambda}_k^{(n)} \right)^2 > \varepsilon.$$

Some terms $\{\mathbf{y}_{\bullet}^{(c)}\}$ that satisfy

$$\sum_{k=1}^K \left(\frac{y_{\bullet k}^{(c)}}{m_k^{(n)}} - \hat{\lambda}_k^{(n)} \right)^2 < \varepsilon$$

appear in the denominator for (4.14). There is at least one vector $\mathbf{y}_{\bullet}^{(c)}$, close to $\hat{\lambda}_k^{(n)} m_k^{(n)}$, such that $\sum_{k=1}^K \left(\frac{y_{\bullet k}^{(c)}}{m_k^{(n)}} - \hat{\lambda}_k^{(n)} \right)^2 < C_3/n^2$, for some constant C_3 independent of n .

As \mathcal{I}_0 has bounded partial derivatives, $\mathcal{I}_0(\mathbf{y}_{\bullet}^{(c)}) > \mathcal{I}_0(\hat{\mathbf{y}}_{\bullet}) - C_4/n$ for a constant C_4 and a lower bound for the contribution of that particular $\{\mathbf{y}_{\bullet}^{(c)}\}$ with

$$\sum_{k=1}^K \left(\frac{y_{\bullet k}^{(c)}}{m_k^{(n)}} - \hat{\lambda}_k^{(n)} \right)^2 < \varepsilon,$$

to the denominator of (4.14) is

$$C_1^K \exp \left\{ n\mathcal{I}_0(\hat{y}_{\bullet}) - C_4 - \sum_{k=1}^K \frac{\log(y_{\bullet k}^{(c)})}{2} \right\} \quad (4.16)$$

An upper bound for (4.14) is the ratio defined by (4.15) over (4.16) plus (4.15). We have that $\sum_{k=1}^K \log(y_{\bullet k}^{(c)})/2$ is $O\{\log(n)\}$ because

$$\sum_{k=1}^K \log(y_{\bullet k}^{(c)}) < \sum_{k=1}^K \log \left(\sum_{k=1}^K y_k^{(c)} \right)$$

which is also $O\{\log(n)\}$ using the same justifications used above (4.15). Therefore, the upper bound for (4.14) defined (4.15) over (4.16) plus (4.15) is of $O\{\exp(-nc_0\epsilon)\}$. It goes to 0 as n goes to ∞ .

The proof can be accommodated to include offsets t_{ik} by multiplying (4.11) by

$$\exp(t_{\bullet k}) = \exp \left(\sum t_{ik} \right)$$

and adding $t_{\bullet k}$ to $\mathcal{I}_0(\cdot)$. The remaining equations can be modified accordingly.

Chapter 5

The distribution of the Mahalanobis distance between two groups under covariates multinormality assumption

Abstract

In this paper, we calculate the distribution of the square of the Mahalanobis distance between two groups under the assumption of normally distributed covariates. We show how this result can be used for algorithms like rerandomization in the context of experimental design, as it allows one to choose a threshold corresponding to a quantile without the use of simulations. The distribution found can still be used as an approximation of the distribution of the square of the Mahalanobis distance between two groups without the assumption of covariates normality unless the number of experimental units is low and there are categorical covariates. Using simulations, we show that the distribution we calculate induce an approximation which seems to perform better in approximating the squared Mahalanobis distance between two groups than an existing approximation based on the chi-square distribution. ¹

5.1 Introduction

In the theory of design of experiments, randomization defines the act of allocating the subjects of the experiment randomly to the different treatment groups. For example, if an experimenter is comparing a new drug to a standard drug, patients can be randomly assigned to either the group using the new drug or the group using the control drug. However, it is not uncommon for this method to result in an unbalanced allocation to groups, i.e., groups with significant differences in the means of one or more covariates. Amongst other critical problems, this imbalance can induce higher variance of regression coefficient estimates Tillé (2022, p. 7) or that the confidence intervals and significance tests are sensitive to those imbalances Morgan and Rubin (2012), which is not desirable for inference. Several methods to get balanced groups in experiments such as clinical trials or balanced samples for survey sampling have been proposed, popular ones being stratification, rerandomization and multivariate matching. Morgan and Rubin (2012) recommends the use of rerandomization for experiments. Tillé (2022) adapted survey sampling methods to get balanced groups in clinical trials.

In this paper, we show that if the covariates are generated from a multinormal distribution, the distribution of the square of the Mahalanobis distance between two groups under randomization follows a beta distribution using the proof from a conference article Djauhari (2004). We

¹E.Talovic and Y.Tillé. Note on the distribution of the Mahalanobis distance between two groups under covariates multinormality assumption, submitted, 2025

make some links of this result to the field of MANOVA in the supplementary material, where we show that there is also an indirect proof of the result. Also, we show an application of this result by extracting from this result an approximation of the squared Mahalanobis distance between two groups which is useful rerandomization algorithms in an experimental setting. Finally, using simulations, we show that the approximation we introduce seems to perform better than an existing approximation based on the chi-square distribution.

5.2 Chi-square approximation of the Mahalanobis distance

In this section, we calculate the squared Mahalanobis distance under the assumption that the covariates are non-random and that the randomness comes from the act of randomization (experimental units randomly assigned to one group or the other). We also introduce some notations. Assume that the number of groups is two, and that the number of experimental units per group is non-random: n_0 and n_1 units respectively belong to group 0 and 1. Let $\mathbf{x}_1, \dots, \mathbf{x}_n$ be covariates vectors of dimension p and $\mathbf{a} = (a_1, \dots, a_n)$ a random vector of indicator variable, where $a_k = 1$ if the experimental unit k is in group 1 and 0 otherwise. The proportion of $a_k = 1$ is n_1/n . Define \mathbf{V}_{xx} as the sample covariance matrix of $\mathbf{x}_1, \dots, \mathbf{x}_n$, and let $\mathbf{V}_{x_0x_0}$ and $\mathbf{V}_{x_1x_1}$ be the sample covariance matrices of experimental units belonging respectively to group 0 and 1. Finally, define the square of the Mahalanobis distance between the group 0 and 1 by

$$d_M^2(\bar{\mathbf{x}}_0, \bar{\mathbf{x}}_1) = (\bar{\mathbf{x}}_0 - \bar{\mathbf{x}}_1)^\top \mathbf{V}_{xx}^{-1} (\bar{\mathbf{x}}_0 - \bar{\mathbf{x}}_1),$$

where $\bar{\mathbf{x}}_0$ and $\bar{\mathbf{x}}_1$ correspond to the sample mean of the covariates of group 0 and 1 respectively. This distance allows one to assess how balanced an allocation to groups is. This distance takes into account both the dispersion of covariates and correlation between covariates. For convenience, the notation d_M^2 for $d_M^2(\bar{\mathbf{x}}_0, \bar{\mathbf{x}}_1)$ will be used in the text. Note that under the assumption of this section, d_M^2 , $\mathbf{V}_{x_0x_0}$ and $\mathbf{V}_{x_1x_1}$ depend on \mathbf{a} as $\mathbf{x}_1, \dots, \mathbf{x}_n$ are given.

Rerandomization is an algorithm which usually consists either in performing randomizations until the covariates are balanced according to a threshold chosen beforehand for the squared Mahalanobis distance between the two groups d_M^2 or performing a fixed number N_g of randomizations and using the most balanced one for the allocation of the experimental units to each group. The choice of this threshold is important. If the chosen threshold is too low, the number of randomizations needed could be very high or infinite, and if is too high, the rerandomization algorithm is likely to end with only one or very few randomizations and would offer little improvement over using a simple randomization. Therefore calculating the distribution of d_M^2 under randomization can be useful, as it allows one to choose the threshold as a quantile for rerandomization. Also, it enables one to calculate the distribution of the minimal d_M^2 out of N_g allocations.

For now, we assume that covariates $\mathbf{x}_1, \dots, \mathbf{x}_n$ are *not* generated from a multinormal distribution. Morgan and Rubin (2012, p. 9–10) found that,

$$\lim_{n \rightarrow \infty} \frac{n_0 n_1}{n} d_M^2(\bar{\mathbf{x}}_0, \bar{\mathbf{x}}_1) \sim \chi_p^2, \quad (5.1)$$

under randomization with non-random n_0 and n_1 . This result is derived from the finite population central limit theorem which implies that $\bar{\mathbf{x}}_0 - \bar{\mathbf{x}}_1$ has a multivariate normal distribution under the assignment scheme for the groups (Hájek, 1960). In practice, this theorem uses slight distributional assumptions for $\mathbf{x}_1, \dots, \mathbf{x}_n$ and n_0 and n . In our analysis the source of randomness comes from the randomization, in other words from \mathbf{a} . Considering that $E_a(\bar{\mathbf{x}}_0 - \bar{\mathbf{x}}_1) = 0$ and

$$\text{Var}_a(\bar{\mathbf{x}}_0 - \bar{\mathbf{x}}_1) = \frac{n_0 n_1}{n} \mathbf{V}_{xx},$$

where E_a and Var_a respectively correspond to the expectation and variance under the assignment

scheme for the groups, it is possible to deduce the result in Equation (5.1).

However, Tillé (2022, p. 5) showed that

$$d_M^2(\bar{\mathbf{x}}_0, \bar{\mathbf{x}}_1) \leq \frac{n(n-1)}{n_0 n_1}, \quad (5.2)$$

which seems inconsistent with the fact that a chi-square random variable has a density with unbounded support. Approximating d_M^2 with a distribution whose support is bounded is more appropriate. In the next section, by assuming that the covariates are multinormal, we calculate a distribution for the squared Mahalanobis distance between two groups which seems to perform much better than the distribution in (5.1) from our simulation study. Also, this approximation has a bounded support which is coherent with inequality (5.2).

5.3 Distribution of the square of the Mahalanobis distance between two groups

In this section, we assume that the covariates $\mathbf{x}_1, \dots, \mathbf{x}_n$ are independently and identically distributed random vectors from a p -variate normal distribution $\mathcal{N}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. While it would be more realistic not to view the covariates as normally distributed, we can obtain a distribution function of the squared Mahalanobis distance for every experiment sample size n . Also, as the covariates are independent and identically distributed and \mathbf{a} is independent of them, we can do the calculations as if \mathbf{a} is fixed and then generalize the results for the case where the components of \mathbf{a} are generated from a simple randomization with fixed group sizes. Several results and definitions are presented before calculating the distribution of the square of the Mahalanobis distance.

Definition 3. If $\mathbf{x}_1, \dots, \mathbf{x}_m$ are independently and identically distributed random vectors from a p -variate normal distribution $\mathcal{N}_p(\mathbf{0}, \boldsymbol{\Sigma})$, then $\mathbf{W} = \sum_{i=1}^m \mathbf{x}_i \mathbf{x}_i^\top$ is said to have a Wishart distribution $\mathbf{W}_p(m, \boldsymbol{\Sigma})$ with m degrees of freedom.

Theorem 1. Let $\mathbf{H} \sim \mathbf{W}_p(m_H, \boldsymbol{\Sigma})$ and $\mathbf{E} \sim \mathbf{W}_p(m_E, \boldsymbol{\Sigma})$ be two independent Wishart matrices. If $m_H < p$ and $m_E \geq p$, then $\mathbf{V} = (\mathbf{H} + \mathbf{E})^{-1/2} \mathbf{H} (\mathbf{H} + \mathbf{E})^{-1/2}$ has rank m_H with probability one, and the m_H non-zero eigenvalues $\theta_1 > \theta_2 > \dots > \theta_{m_H}$ the following joint density

$$f(\theta_1, \dots, \theta_{m_H}) = a^{-1} \left(\prod_{j=1}^{m_H} \theta_j \right)^{(p-m_H-1)/2} \left\{ \prod_{j=1}^{m_H} (1 - \theta_j) \right\}^{(m_E-p-1)/2} \prod_{j < k}^{m_H} (\theta_j - \theta_k)$$

where

$$a = \pi^{m_H^2/2} B_{m_H} \left\{ \frac{1}{2}p, \frac{1}{2}(m_E + m_H - p) \right\} \Gamma_{m_H} \left(\frac{1}{2}m_H \right),$$

with Γ_m the multivariate gamma function and

$$B_m(x, y) = \frac{\Gamma_m(a) \Gamma_m(y)}{\Gamma_m(x+y)}.$$

Also if $m_H = 1$ then the eigenvalues of \mathbf{V} and $\mathbf{H}(\mathbf{H} + \mathbf{E})^{-1}$ are the same.

Proof of Theorem 1. See Seber (1984, p. 35–38) or Anderson (1958, p. 537). □

Corollary 9. Define two independent Wishart matrices $\mathbf{H} \sim \mathbf{W}_p(m_H, \boldsymbol{\Sigma})$ and $\mathbf{E} \sim \mathbf{W}_p(m_E, \boldsymbol{\Sigma})$. If $m_H = 1$ and $m_E \geq p$, then $\mathbf{V} = \mathbf{H}(\mathbf{H} + \mathbf{E})^{-1}$ has only one non-zero eigenvalue with probability one and follows a beta distribution:

$$f(u) = \frac{\Gamma\left(\frac{n-1}{2}\right)}{\Gamma\left(\frac{p}{2}\right)\Gamma\left(\frac{n-p-1}{2}\right)} u^{p/2-1} (1-u)^{(n-p-1)/2-1}, \quad 0 < u < 1.$$

Let us also introduce the following lemmas which are also useful for the proof of the next result.

Lemma 1. Let $\mathbf{v} = (v_1, \dots, v_n)^\top$ and \mathbf{M} be $n \times n$ matrix. Then

$$\mathbf{v}^\top \mathbf{M} \mathbf{v} = \text{Tr}(\mathbf{v} \mathbf{v}^\top \mathbf{M}).$$

Lemma 2. Define the sample mean of the covariates as $\bar{\mathbf{x}}$, then

$$n(\bar{\mathbf{x}}_0 - \bar{\mathbf{x}})/n_1 = (\bar{\mathbf{x}}_0 - \bar{\mathbf{x}}_1).$$

Proof of Lemma 2.

$$(\bar{\mathbf{x}}_0 - \bar{\mathbf{x}}_1) = \left(\bar{\mathbf{x}}_0 - \frac{n}{n_1} \bar{\mathbf{x}} + \frac{n_0}{n_1} \bar{\mathbf{x}}_0 \right) = \left\{ \left(1 + \frac{n_0}{n_1} \right) \bar{\mathbf{x}}_0 - \frac{n}{n_1} \bar{\mathbf{x}} \right\} = \frac{n}{n_1} (\bar{\mathbf{x}}_0 - \bar{\mathbf{x}}).$$

□

Using the previous results, we can calculate the distribution of the square of the Mahalanobis distance between the two groups. The proof is a rigorous and complete version of the sketch of a proof found in the conference paper in Djauhari (2004).

Result 23. Let $\mathbf{x}_1, \dots, \mathbf{x}_n$ be a random sample of size n from a p -variate normal distribution $\mathcal{N}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Let $a_k \in \{0, 1\}$ for $k \in \{1, \dots, n\}$ be fixed. If $p > 1$ and $n - 2 \geq p$, then the square of the Mahalanobis distance d_M^2 between $\bar{\mathbf{x}}_0$ and $\bar{\mathbf{x}}_1$ follows a beta distribution:

$$\frac{n_0 n_1}{(n-1)n} d_M^2(\bar{\mathbf{x}}_1, \bar{\mathbf{x}}_0) \sim \text{Beta}\left(\frac{p}{2}, \frac{n-p-1}{2}\right). \quad (5.3)$$

Proof of Result 23. Let \mathbf{S}_B , \mathbf{S}_W and \mathbf{S}_T be defined respectively as the scatter matrices between, within and the total of the two complementary groups:

$$\begin{aligned} \mathbf{S}_B &= n_0(\bar{\mathbf{x}}_0 - \bar{\mathbf{x}})(\bar{\mathbf{x}}_0 - \bar{\mathbf{x}})^\top + n_1(\bar{\mathbf{x}}_1 - \bar{\mathbf{x}})(\bar{\mathbf{x}}_1 - \bar{\mathbf{x}})^\top, \\ \mathbf{S}_W &= \sum_{k=1}^n (1 - a_k)(\mathbf{x}_k - \bar{\mathbf{x}}_0)(\mathbf{x}_k - \bar{\mathbf{x}}_0)^\top + \sum_{k=1}^n a_k(\mathbf{x}_k - \bar{\mathbf{x}}_1)(\mathbf{x}_k - \bar{\mathbf{x}}_1)^\top, \\ \mathbf{S}_T &= \sum_{k=1}^n (\mathbf{x}_k - \bar{\mathbf{x}})(\mathbf{x}_k - \bar{\mathbf{x}})^\top = \mathbf{S}_B + \mathbf{S}_W = (n-1)\mathbf{V}_{xx}. \end{aligned}$$

Using the relation $n(\bar{\mathbf{x}}_0 - \bar{\mathbf{x}})/n_1 = (\bar{\mathbf{x}}_0 - \bar{\mathbf{x}}_1)$ from Lemma 2, the square of the Mahalanobis distance can be expressed as

$$d_M^2(\bar{\mathbf{x}}_1, \bar{\mathbf{x}}_0) = \frac{(n-1)n^2}{n_1^2} (\bar{\mathbf{x}}_0 - \bar{\mathbf{x}})^\top \mathbf{S}_T^{-1} (\bar{\mathbf{x}}_0 - \bar{\mathbf{x}}) = \frac{(n-1)n}{n_1 n_0} U,$$

where

$$U = \frac{nn_0}{n_1} (\bar{\mathbf{x}}_0 - \bar{\mathbf{x}})^\top \mathbf{S}_T^{-1} (\bar{\mathbf{x}}_0 - \bar{\mathbf{x}}).$$

From Lemma 1, it is possible to write U

$$U = \text{Tr} \left\{ \frac{nn_0}{n_1} (\bar{\mathbf{x}}_0 - \bar{\mathbf{x}})(\bar{\mathbf{x}}_0 - \bar{\mathbf{x}})^\top \mathbf{S}_T^{-1} \right\},$$

and using Lemma 2 , \mathbf{S}_B can be rewritten as

$$\mathbf{S}_B = \frac{nn_0}{n_1}(\bar{\mathbf{x}}_0 - \bar{\mathbf{x}})(\bar{\mathbf{x}}_0 - \bar{\mathbf{x}})^\top. \quad (5.4)$$

So, $U = \text{Tr}(\mathbf{S}_B \mathbf{S}_T^{-1}) = \text{Tr}\{\mathbf{S}_B(\mathbf{S}_B + \mathbf{S}_W)^{-1}\}$, which is Pillai's trace statistic where \mathbf{S}_B follows a Wishart distribution with one degree of freedom and \mathbf{S}_W a Wishart distribution with $n - 2$ degrees of freedom.

To finish the proof, we use Corollary 2. To do this, the independence of \mathbf{S}_B and \mathbf{S}_W must be established to use Corollary 9. Using Lemma 2 again, we rewrite

$$\mathbf{S}_B = \frac{n_0 n_1}{n}(\bar{\mathbf{x}}_0 - \bar{\mathbf{x}}_1)(\bar{\mathbf{x}}_0 - \bar{\mathbf{x}}_1)^\top.$$

and $\mathbf{S}_W = (n_0 - 1)\mathbf{V}_{x_0 x_0} + (n_1 - 1)\mathbf{V}_{x_1 x_1}$. Mardia et al. (1979, p. 69) show that $\bar{\mathbf{x}}_i$ is independent of $\mathbf{V}_{x_i x_i}$ for $i \in \{0, 1\}$ and samples from different groups are independent, so \mathbf{S}_B and \mathbf{S}_W are independent. Therefore, from Corollary 9, $\mathbf{S}_B \mathbf{S}_T^{-1}$ has only one non-zero eigenvalue with probability one and this eigenvalue has density

$$f(u) = \frac{\Gamma\left(\frac{n-1}{2}\right)}{\Gamma\left(\frac{p}{2}\right)\Gamma\left(\frac{n-p-1}{2}\right)} u^{p/2-1} (1-u)^{(n-p-1)/2-1}, \quad 0 < u < 1. \quad (5.5)$$

Finally, U is equal to the only non-zero eigenvalue of $\mathbf{S}_B \mathbf{S}_T^{-1}$ because the trace is equal to the sum of the eigenvalues of a matrix, which means the density function of U is the same density function as in (5.5). This concludes the proof. \square

Again, if \mathbf{a} is independent of the covariates $\mathbf{x}_1, \dots, \mathbf{x}_n$, then we can generalize Result 23 to the case where \mathbf{a} is generated from a simple randomization with fixed group sizes. Additionally, the support of the density function of the beta distribution is $[0, 1]$, which means that Result 23 is consistent with inequality (5.2).

5.4 Simulations to compare the Beta and Chi-Square approximations

In this section, we briefly present simulations to compare the Beta and Chi-Square distribution in different settings using different distributions such the Extended skew-normal distribution (ESN) Azzalini and Valle (1996); Canale (2011), the Weibull distribution and discrete distributions. The detailed simulation results can be found in the supplementary material.

From Figure 5.1, the overall conclusion is that the beta approximation is better than the chi-square in any situation, even when the data are generated from a non-normal distribution. The chi-square approximation comes from an asymptotic result, so it is not surprising that it performs worse than the beta approximation that also uses the number of covariates as an additional parameter. The beta approximation seems to be relatively robust against some degree of skewness in the data even though high skewness can greatly reduce the quality of the fit. However, the multivariate kurtosis and skewness statistics seem to vary quite a lot even with the same parameter choices, which is probably due to low sample size, and do not seem to have a clear correlation in the outcome of the graph. Therefore making recommendations using those statistics is difficult.




Figure 5.1: Histograms of four set of simulations. In panel (a), the data follows an ESN distribution with $n = 16$, $p = 4$, $\boldsymbol{\alpha} = (12, -12, 20, 20)$ and $\tau = 0$. In panel (b), the data follows an ESN distribution with $n = 16$, $p = 4$, $\boldsymbol{\alpha} = (120, 120, 120, 120)$ and $\tau = -1$. In panel (c), the data follows a Weibull distribution with $n = 16$, $p = 4$, with a shape and scale parameters respectively equal to 0.5 and 1. In panel (d), $n = 16$ with two covariates. The first component is generated from a normal distribution and the second from a Bernoulli distribution with mean 0.5.

5.5 Discussion

In experimental settings, the covariates $\mathbf{x}_1, \dots, \mathbf{x}_n$ are fixed values that we cannot generate with simulations based on a normal distribution, therefore we can only use Result 23 to get an approximation based on the beta distribution to estimate the distribution of the square of the Mahalanobis distance between two groups d_M^2 . We can also use Equation (5.1) to get a chi-square approximation for the square of the Mahalanobis distance between the two groups. The beta approximation assumes that the covariates are close to a multivariate normal distribution and the chi-square approximation is derived from an asymptotic result. When n tends to infinity, both approximations converge to the same distribution. But those approximations are different when n is small. Simulations show that the beta approximation seems to be always better than the chi-square approximation, even when the covariates $\mathbf{x}_1, \dots, \mathbf{x}_n$ are generated from a

skewed and heavy- or light-tailed distribution. In other words, even with strong departure from multinormality of the covariates, the beta approximation performs better. These simulations combined with the apparent inconsistency with the support of the chi-square density function seem to indicate that the beta approximation is better overall. If n is small and there are categorical covariates, which is often the case in many experiments, both the beta and chi-square approximations fit the square of the Mahalanobis distribution between the two groups d_M^2 very poorly. This problem is highlighted when there is only one covariate which is categorical. In this case, we strongly recommend to use simulations to obtain a distribution of the square of the Mahalanobis distance. Additionally, if the covariates follow a distribution with undefined second moments, for example the Cauchy distribution, both approximations fail to correctly estimate the distribution of d_M^2 , even when n tends to infinity. The finite population central limit theorem is not applicable if the variances of $\mathbf{x}_1, \dots, \mathbf{x}_n$ are not defined.

5.6 Supplementary material

5.6.1 Simulations to compare the chi-square and beta distribution approximations

Definition 4. A random variable follows X an extended skew-normal distribution with location parameter ξ , scale parameter ω , shape parameter α and truncation parameter τ if its density function is

$$f(x) = \frac{1}{\omega} \phi\left(\frac{x - \xi}{\omega}\right) \Phi\left\{\tau\sqrt{\alpha^2 + 1} + \alpha\left(\frac{x - \xi}{\omega}\right)\right\} / \Phi(\tau), \quad \xi, \alpha, \tau \in \mathbb{R}, \quad \omega > 0$$

where ϕ denotes the standard normal probability density function and Φ its cumulative distribution function.

The skew-normal distribution allows one to simulate skewed data while also giving normal data, depending on the choice of the parameters. The extended skew-normal distribution can also model data which is skewed and heavy- or light-tailed. Increasing the shape parameter α induces a more right-skewed distribution. Increasing the truncation parameter τ create an heavier tailed distribution but also reduces the skewness. The effect of the τ parameter depends on α . If $\alpha = 0$, then changing τ has no effect on the distribution. The location parameter ξ and ω can be considered as the mean and variance parameters from the normal distribution, even if ξ does not correspond to the mean or median of the skew-normal distribution if $\alpha \neq 0$. Figure 5.2 displays the density of the extended skew-normal with different parameter choices. Multivariate simulations for this distribution are made using Azzalini (2020).

To illustrate our main findings, four density histograms of the squared Mahalanobis distance multiplied by $(n_0 n_1)/n$ for different series of simulations are displayed in Figure 5.1. The data are generated from a certain distribution once at the start, and then the simulations consist in doing a fixed number of allocations to groups using a simple random sampling without replacement and calculating the square of the Mahalanobis distance after every allocation to groups to generate the histogram. In every series of simulations, 10000 simulations are done and $n_0 = n_1 = n/2$. The multivariate skewness and kurtosis are assessed using the statistics developed by Mardia (1970).

The panel (a) of Figure 5.1 shows the histogram of simulations with $n = 16$ and $p = 4$. The distribution of the simulation is generated from a extended skew-normal distribution with parameters chosen such that the vector of marginal shapes is $\alpha = (12, -12, 20, 20)$ with truncation parameter $\tau = 0$, ω , which is a positive definite matrix, is randomly generated from a Wishart distribution. The beta approximation fits the data, is relatively robust to some degree of skewness and is much better than the chi-square approximation.




Figure 5.2: Extended skew-normal densities with different parameter choices.

In panel (b) of Figure 5.1, the data is generated with a similar distribution as in panel (a) but the shape parameter vector is equal to $\alpha = (120, 120, 120, 120)$ and $\tau = -1$ to try to get very skewed data. The fit of the beta approximation is slightly worse than in the previous set of simulations and the chi-square approximation is still too left-skewed.

In panel (c) of Figure 5.1, the simulations uses same n and p as before but the data is generated from a Weibull distribution, which allows us to generate a distribution with slightly more extreme multivariate skewness and kurtosis values than the skew-normal distribution. The shape and scale parameters are respectively equal to 0.5 and 1. The resulting extreme departure from normality makes the beta approximation hardly usable. With data generated from a Weibull distribution, the shape of the histogram varies more than it does with the skew-normal distribution. One of the poorer fits is displayed.

The previous simulations were generated from continuous distributions. The panel (d) displays the histogram from simulations with $p = 2$, the first component being generated from a normal distribution and the second from a Bernoulli distribution with mean 0.5. Due to the discrete nature of the Bernoulli distribution, the Mahalanobis distance between the two groups tends to cluster around certain values. Of course, this problem is highlighted when there is only one covariate from a Bernoulli distribution.

The overall conclusion is that the beta approximation is better than the chi-square in any situation, even when the data are generated from a non-normal distribution. The chi-square approximation comes from an asymptotic result, so it is not surprising that it performs worse than the beta approximation. Also, the beta approximation seems to be relatively robust against some degree of skewness in the data even though high skewness can greatly reduce the quality of the fit. However, the multivariate kurtosis and skewness statistics seem to vary quite a lot even with the same parameter choices, which is probably due to low sample size, and appear to

not have a clear correlation in the outcome of the graph. Therefore making recommendations using those statistics is difficult.

If neither the beta nor chi-square approximation fits the distribution of d_M^2 , the rerandomization approach is still valid. This simply implies that approximating the distribution of d_M^2 using simulations or taking the assignment to groups that induce the lowest d_M^2 out of large number of simulations is more advisable in this case

5.6.2 Links to MANOVA

The results shown in this paper are often derived from results found in the field of multivariate analysis. MANOVA stands for multivariate analysis of variance and, as the name implies, it is a multivariate generalization of the analysis of variance (ANOVA). Unlike with ANOVA, there are multiple response variables and MANOVA uses the covariance between the response variables to test the statistical significance of the mean differences instead of the sums of squares. Therefore under normality assumptions about error distributions, the counterpart of the sum of squares due to error has a Wishart distribution. Also despite ANOVA having a simple theory with only one test statistic generally used, the F-test, MANOVA typically uses four test statistics.

The most common statistics are summaries based on the roots (or eigenvalues) $\lambda_1, \dots, \lambda_p$ of the $\mathbf{A} = \mathbf{H}\mathbf{E}^{-1}$ matrix, where \mathbf{H} and \mathbf{E} are respectively “error” and “hypothesis” matrices which are assumed to have Wishart distributions with m_H and m_E degrees of freedom under the null hypothesis. The four test statistics are the following:

- Wilks’s Lambda:

$$\Lambda_{\text{Wilks}} = \prod_{i=1}^p \frac{1}{1 + \lambda_i} = \frac{\det(\mathbf{E})}{\det(\mathbf{E} + \mathbf{H})};$$

- Pillai’s Trace:

$$\Lambda_{\text{Pillai}} = \sum_{i=1}^p \frac{1}{1 + \lambda_i} = \text{tr}\{\mathbf{A}(\mathbf{I} + \mathbf{A})^{-1}\};$$

- Lawley–Hotelling Trace:

$$\Lambda_{\text{LH}} = \sum_{i=1}^p \lambda_i = \text{tr}(\mathbf{A});$$

- Roy’s Largest Root:

$$\Lambda_{\text{Roy}} = \max_i(\lambda_i).$$

For most tests, the distribution under H_0 is very difficult to calculate and only approximations from the F-distribution are used, except in cases with low dimensionality where the distributions can be calculated exactly, for example with $m_H = 1$. The following approximation of the Pillai’s trace statistic distribution is generally used Seber (1984, p. 564),

$$\frac{(2v_2 + s + 1)\Lambda_{\text{Pillai}}}{(2v_1 + s + 1)(s - \Lambda_{\text{Pillai}})} \sim F_{s(2v_1+s+1), s(2v_2+s+1)}, \quad (5.6)$$

where $s = \min(m_H, p)$, $v_1 = (|p - m_H| - 1)/2$ and $v_2 = (m_E - p - 1)/2$. By setting $m_H = 1$ and $m_E = n - 2$, we get that Λ_{Pillai} corresponds to U from the proof of Result 1. From Equation (5.6), the final result of Result 1 can be deduced using the relation between the F-distribution and the beta distribution.

Chapter 6

Conclusion

This conclusion explains the potential caveats, problems, improvements and additional directions for further research of some the concepts introduced in this thesis.

The first part of the thesis was focuses on the efficiency-robustness trade-offs in experimental and sampling designs. We also observed that robustness concept is strongly correlated with entropy, which gives some additional justification for older sampling literature focusing on carrying out maximum entropy sampling. When the probabilities of drawing each possible sample or assignment are equal, maximum entropy sampling corresponds to the sampling design with highest number of possible samples. In experiments, designs with a high number of possible samples are useful to carry out randomization tests. If this number is too low, the resolution of the test becomes too low to draw conclusive inference.

The robustness/balance trade-off and the mixing of a balanced design with the SRSWR design proposed in this thesis is similar to a mix of two minimax principles under different models. We quickly introduce the concept of “a priori balance” from Kallus (2018), which employs a minimax paradigm to minimize the maximum potential variance of the sample average treatment effect estimator $\hat{\tau}$. If $Y = f(X) + \epsilon$ and $E(Y|X) = f(X)$. Define

$$B_{12}(W, f) = \sum_{i:W_i=1} f(X_i) - \sum_{i:W_i=2} f(X_i).$$

For the “pure strategy optimal design”, we choose W uniformly among the set that minimizes

$$W \in \max_{\|f\| \leq 1} B_{12}^2(W, f). \tag{6.1}$$

Depending on f and the norm $\|f\|$ used, certain common designs such as exact matching, rerandomization or blocking are the optimal minimax designs. If there is no structure in the dependence of outcomes, then complete randomization is optimal. In this framework, when f corresponds to a regression with covariates $\mathbf{x}_1, \dots, \mathbf{x}_N$ as explanatory variables, the optimal design corresponds to designs exactly balanced on the mean of the covariates. In this setting, the Cube method or rerandomization gives assignments to groups that approximately satisfy the (6.1). We can see that the optimal design is dependent on the underlying model of the relationship between the response or variable of interest and the covariates. Other concepts such as optimality of sampling plans under certain models or other forms of minimax principles were also studied in the field of sampling (Tillé and Wilhelm, 2017; Scott and Smith, 1975; Nedyalkova and Tillé, 2008). Potential axis of research could involve unifying those minimax concepts or expanding them to include risks due to missing data, non-response or attrition in clinical trials.

Our analysis implies that primarily focusing on balance on the covariate means usually does not involve much risk when we look at the model-based variance and total variance, so there is room to add other metrics to balance on without unduly reducing the robustness. We give a

specific example of potential new methods to be used. Assume that some variables are assumed to be particularly important in the model and at the same time to not have a simple linear relationship with the response. Given that finding pairs that are sufficiently similar, it might be better to try to use a mixed strategy of matching only for the specific covariates assumed to be important and have the remaining covariates be used for balancing on the means between the groups. It could be done with the help of the design mixing algorithm from Chapter 2 or a special implementation of the matched-Cube algorithm Tillé (2022).

Our extension of the Cube Method to exact inference for Poisson regression does not work optimally. One can easily address the concern that the conditional probabilities/distribution are not exactly the correct one and also that the Cube method only approximately satisfies the constraints from the sufficient statistic. On top of that there is the critic which concern all “approximately exact” methods, which is that the exact Cube algorithm may not necessarily output the all the possible samples that satisfy exactly the constraints even after a high number of simulations. There is also the general criticism against exact inference in general that could be made: This method works with parametric models that are very likely to be incorrect and often applied on small data sets. Therefore the practical utility of exactness in this setting is less clear when ignoring cases where the data is separable. On a more positive note, the beta distribution used for the univariate conditional probabilities was a fairly good approximation and the algorithm outputs p-values that are close to exact when the model is correct. Also, finding the exact conditional probabilities could significantly reduce the concerns of potential lack of accuracy.

List of Figures

- 2.1 Scatterplot of the estimated expectation of the squares of the Mahalanobis distances between the covariate Horvitz-Thompson estimators of the two groups by the largest eigenvalues. The numbers in brackets correspond to the tuning parameter ϕ which gives importance either to robustness or to balance depending on the value of ϕ . When the inclusion probabilities are equal, CPS is changed to SRSWOR. 28
- 3.1 Representation of the flight phase in the algorithm. The vertices of the cube represent different samples. The arrow represents the initial inclusion probabilities, before the initiating the random walk. The darker subspace describe the subspace generated by the fixed size constraint and inclusion probabilities. In this example, the samples satisfy exactly the constraints as the vertices intersect with the subspace. 41
- 3.2 Illustration of the stratification in the CBA with $K = 6$ treatment groups. Each column represents a “stratum per individual”, with its indices, its inclusion of probabilities to each group and the value assigned through the algorithm. Only one entry per column is set to 1, due to the stratification. This ensures that only each element is assigned to only 1 treatment group. 44
- 3.3 Density plots of \bar{M} of the CBA and randomization under varying simulation settings. 46
- 4.1 Illustration of the stratification in the algorithm. Each column represents a stratum with different auxiliary variables and their inclusion probabilities. One unit from each stratum is selected with the sufficient statistic constraints considered. . 64
- 5.1 Histograms of four set of simulations. In panel (a), the data follows an ESN distribution with $n = 16$, $p = 4$, $\alpha = (12, -12, 20, 20)$ and $\tau = 0$. In panel (b), the data follows an ESN distribution with $n = 16$, $p = 4$, $\alpha = (120, 120, 120, 120)$ and $\tau = -1$. In panel (c), the data follows a Weibull distribution with $n = 16$, $p = 4$, with a shape and scale parameters respectively equal to 0.5 and 1. In panel (d), $n = 16$ with two covariates. The first component is generated from a normal distribution and the second from a Bernoulli distribution with mean 0.5. 76
- 5.2 Extended skew-normal densities with different parameter choices. 78

List of Tables

- 2.1 Summary of the properties of some sampling methods. The maximum entropy column refers to maximum entropy sampling in the “obvious” support domain (e.g., fixed sample size) given certain inclusion probabilities. 21
- 2.2 Simulations: $N = 40$, $n = 20$ and equal inclusion probabilities. Largest eigenvalue of Δ obtained by simulations and approximation or true value for this eigenvalue. Squares of the Mahalanobis distances between the estimators and the true totals of the auxiliary variables. 37
- 2.3 Mixtures of designs with algorithm 1. The values of ϕ are between the parentheses. 38
- 2.4 Mixtures of designs with Gram-Schmidt walk design. The values of ϕ are between the parentheses. 38

- 3.1 Relative performance of different metrics between the CBA algorithm and a fixed size randomization using the Mahalanobis distance between groups. For *Sim4*, the CBA is compared with an unequal inclusion probability sampling designs whose only constraint is group size. The ratios are drawn from a set of 10000 simulations. *Sim1* represents the first data set with $N = 245$ units, $p = 12$ covariates, $K = 7$ treatment groups and fully equal inclusion probabilities. *Sim2* is the same set of simulations, with only 6 covariates to balance. *Sim3* represents the second data set with $N = 75$ units, $p = 8$, $K = 3$ treatment groups and fully equal inclusion probabilities. *Sim4* is the same data set as *Sim3* with unequal inclusion probabilities that are randomly generated. 47
- 3.2 Coverage probability of the true parameter τ_2 within the approximated confidence interval. Model 1 has been used for *Sim1* and *Sim2* and Model 2 for *Sim3* and *Sim4*. 51

- 4.1 Ingots data set 65
- 4.2 Empirical p -values against nominal p -values for each method in a Poisson regression model for $n = 75$ and $n = 25$. Wald corresponds to the Wald test, LRT represents the likelihood ratio test and Scube represents our algorithm based on the *stratifiedcube* function. 66

Bibliography

- Agresti, A. (1992). A survey of exact inference for contingency tables. *Statistical Science* 7(1), 131–153.
- Agresti, A. (2003). *Categorical Data Analysis* (2 ed.). Hoboken, New Jersey: John Wiley & Sons.
- Agresti, A. and B. A. Coull (1998, 05). Approximate is better than 'exact' for interval estimation of binomial proportions. *The American Statistician* 52, 119–126.
- Anderson, T. W. (1958). *An Introduction to Multivariate Statistical Analysis*. Wiley, New York.
- Atkinson, A. C. and A. N. Donev (1992). *Optimum Experimental Designs*. Oxford: Oxford University Press.
- Azzalini, A. (2020). *The R package sn: The Skew-Normal and Related Distributions such as the Skew-t (version 1.6-2)*. Università di Padova, Italia.
- Azzalini, A. and A. D. Valle (1996, 12). The multivariate skew-normal distribution. *Biometrika* 83(4), 715–726.
- Balcells, L. et al. (2022). Do transitional justice museums persuade visitors? evidence from a field experiment. *The Journal of Politics* 84(1), 496–510.
- Basu, D. (1980). Randomization analysis of experimental data: The fisher randomization test. *Journal of the American Statistical Association* 75(371), 575–582.
- Bertail, P., E. Chautru, and S. Cléménçon (2017). Empirical processes in survey sampling with (conditional) Poisson designs. *Scandinavian Journal of Statistics* 44(1), 97–111.
- Brändén, P. and J. Jonasson (2012). Negative dependence in sampling. *Scandinavian Journal of Statistics* 39(4), 830–838.
- Brazzale, A. and A. Davison (2008, 11). Accurate parametric inference for small samples. *Statistical Science* 23, pp. 465–484.
- Brewer, K. R. W. and M. Hanif (1983). *Sampling with Unequal Probabilities*. New York: Springer.
- Caffo, B. S. and J. G. Booth (2003). Monte carlo conditional inference for log-linear and logistic models: a survey of current methodology. *Statistical Methods in Medical Research* 12(2), 109–123.
- Canale, A. (2011, 12). Statistical aspects of the scalar extended skew-normal distribution. *Metron* LXIX, 279–295.
- Chauvet, G. (2009). Stratified balanced sampling. *Survey Methodology* 35, 115–119.
- Chauvet, G. (2012). On a characterization of ordered pivotal sampling. *Bernoulli* 18(4), 1320–1340.

- Chen, S. and D. Haziza (2019). Recent developments in dealing with item non-response in surveys: A critical review. *International Statistical Review* 87, S192–S218.
- Chen, X.-H. (1993). Poisson-binomial distribution, conditional Bernoulli distribution and maximum entropy. Technical report, Department of Statistics, Harvard University.
- Chen, X.-H. and J. S. Liu (1997). Statistical applications of the Poisson-binomial and conditional Bernoulli distributions. *Statistica Sinica* 7, 875–892.
- Chow, S. C. and M. Chang (2008). Adaptive design methods in clinical trials - a review. *Orphanet Journal of Rare Diseases* 3, 11.
- Cochran, W. G. (1965). The planning of observational studies of human populations. *J. Roy. Statist. Soc. Ser. A* 128(2), 234–265.
- Cox, D. R. (2009). Randomization in the design of experiments. *International Statistical Review / Revue Internationale de Statistique* 77(3), 415–429.
- Cox, D. R. and E. J. Snell (1989). Analysis of binary data (vol. 32). *Monographs on Statistics and Applied Probability*.
- Cox, L. H. (1987). A constructive procedure for unbiased controlled rounding. *Journal of the American Statistical Association* 82(398), 520–524.
- Deville, J.-C. (2000). Note sur l’algorithme de Chen, Dempster et Liu. Note interne manuscrite, CREST-ENSAI, Rennes.
- Deville, J.-C. and C.-E. Särndal (1992). Calibration estimators in survey sampling. *Journal of the American Statistical Association* 87(418), 376–382.
- Deville, J.-C., C.-E. Särndal, and O. Sautory (1993). Generalized raking procedures in survey sampling. *Journal of the American Statistical Association* 88(423), 1013–1020.
- Deville, J.-C. and Y. Tillé (1998). Unequal probability sampling without replacement through a splitting method. *Biometrika* 85, 89–101.
- Deville, J.-C. and Y. Tillé (2004a). Efficient balanced sampling: The cube method. *Biometrika* 91, 893–912.
- Deville, J.-C. and Y. Tillé (2004b). Efficient balanced sampling: The cube method. *Biometrika* 91, 893–912.
- Deville, J.-C. and Y. Tillé (2005). Variance approximation under balanced sampling. *Journal of Statistical Planning and Inference* 128, 569–591.
- Deville, J.-C. and Y. Tillé (2005). Variance approximation under balanced sampling. *Journal of Statistical Planning and Inference* 128(2), 569–591.
- Diaconis, P. and B. Sturmfels (1998). Algebraic algorithms for sampling from conditional distributions. *The Annals of Statistics* 26(1), 363 – 397.
- Djauhari, M. A. (2004). Mahalanobis distance of two complementary groups of observations. In *Symposium Kebangsaan Sains Matematik ke 12*, pp. 1–9.
- Dol, W., T. Steerneman, and T. Wansbeek (1996). Matrix algebra and sampling theory: The case of the Horvitz-Thompson estimator. *Linear algebra and its applications* 237, 225–238.
- Efron, B. (1971). Forcing a sequential experiment to be balanced. *Biometrika* 58(3), 403–417.

- Forster, J. J., J. W. McDonald, and P. W. F. Smith (2003). Markov chain monte carlo exact inference for binomial and multinomial logistic regression models. *Statistics and Computing* 13, 169–177.
- Friedrich, T., B. Doerr, C. Klein, and R. Osbild (2007). Unbiased matrix rounding. *Electronic Notes in Discrete Mathematics* 28, 41–46. 6th Czech-Slovak International Symposium on Combinatorics, Graph Theory, Algorithms and Applications.
- Gabler, S. (1981). A comparison of Sampford’s sampling procedure versus unequal probability sampling with replacement. *Biometrika* 68, 725–727.
- Gabler, S. (1984). On unequal probability sampling: sufficient conditions for the superiority of sampling without replacement. *Biometrika* 71, 171–175.
- Gerber, M., N. Chopin, and N. Whiteley (2019). Negative association, ordering and convergence of resampling methods. *The Annals of Statistics* 47(4), 2236–2260.
- Good, P. (1994). *Experimental Designs*. New York, NY: Springer New York.
- Grafström, A. and N. L. P. Lundström (2013). Why well spread probability samples are balanced? *Open Journal of Statistics* 3(1), 36–41.
- Grafström, A., N. L. P. Lundström, and L. Schelin (2012). Spatially balanced sampling through the pivotal method. *Biometrics* 68(2), 514–520.
- Haberman, S. J. (1988). A warning on the use of chi-squared statistics with frequency tables with small expected cell counts. *Journal of the American Statistical Association* 83(402), 555–560.
- Hájek, J. (1960). Limiting distributions in simple random sampling from a finite population. *Publ. Math. Inst. Hung. Acad. Sci.* 5, 361–374.
- Hájek, J. (1971). Discussion of an essay on the logical foundations of survey sampling, part on by D. Basu. In V. P. Godambe and D. A. Sprott (Eds.), *Foundations of Statistical Inference*, Toronto, pp. 326. Holt, Rinehart, Winston.
- Han, P. and L. Wang (2013). Estimation with missing data: beyond double robustness. *Biometrika* 100(2), 417–430.
- Härdle, W. and A. Bowman (1988). Robust estimation in finite populations I. *Journal of the American Statistical Association* 83(401), 673–679.
- Harshaw, C., F. Sävje, D. Spielman, and P. Zhang (2024). Balancing covariates in randomized experiments with the gram–schmidt walk design. *Journal of the American Statistical Association* 119(548), 2934–2946.
- Hasler, C. and Y. Tillé (2014). Fast balanced sampling for highly stratified population. *Computational Statistics and Data Analysis* 74, 81–94.
- Hayder, F. (2020, june). Corneal astigmatism. Zenodo.
- Haziza, D. and J.-F. Beaumont (2017). Construction of weights in surveys: A review. *Statistical Science* 32(2), 206–226.
- Hirji, K. F., C. R. Mehta, and N. R. Patel (1987). Computing distributions for exact logistic regression. *Journal of the American Statistical Association* 82(400), 1110–1117.

- Ho, D., K. Imai, G. King, and E. Stuart (2007). Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Political Analysis* 15, 199–236.
- Hodges, Jr., J. L. and L. Le Cam (1960). The Poisson approximation to the Poisson binomial distribution. *Annals of Mathematical Statistics* 31, 737–740.
- Horn, R. A. and C. R. Johnson (2013). *Matrix Analysis* (2 ed.). Cambridge: Cambridge University Press.
- Horvitz, D. G. and D. J. Thompson (1952). A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association* 47, 663–685.
- Jauslin, R., E. Eustache, B. Panahbehagh, and Y. Tillé (2022). *StratifiedSampling: Different Methods for Stratified Sampling*. R package version 0.3.0.
- Jauslin, R., E. Eustache, and Y. Tillé (2021). Enhanced cube implementation for highly stratified population. *Japanese Journal of Statistics and Data Science* 4, 783–795.
- Jauslin, R., B. Panahbehagh, and Y. Tillé (2022). Sequential spatially balanced sampling. *Environmetrics* 33(8), e2776.
- Kallus, N. (2018, 01). Optimal a priori balance in the design of controlled experiments. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 80, 85–112.
- Kang, J. D. Y. and J. L. Schafer (2007, November). Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. *Statist. Sci.* 22(4), 523–539.
- Kapelner, A. and A. Krieger (2014, june). Matching on-the-fly: sequential allocation with higher power and efficiency. *Biometrics* 70(2), 378–388.
- Kapelner, A. and A. Krieger (2023, 07). A matching procedure for sequential experiments that iteratively learns which covariates improve power. *Biometrics* 79(1), 216–229.
- Kapelner, A., A. M. Krieger, M. Sklar, and D. Azriel (2022). Optimal rerandomization designs via a criterion that provides insurance against failed experiments. *Journal of Statistical Planning and Inference* 219, 63–84.
- Kramer, J. B., J. Cutler, and A. Radcliffe (2011). Negative dependence and Srinivasan’s sampling process. *Combinatorics, Probability and Computing* 20(3), 347–361.
- Lee, D. and J. K. Kim (2022). Semiparametric imputation using conditional gaussian mixture models under item nonresponse. *Biometrics* 78(1), 227–237.
- Li, X., P. Ding, and D. B. Rubin (2018). Asymptotic theory of rerandomization in treatment–control experiments. *Proceedings of the National Academy of Sciences* 115(37), 9157–9162.
- Madow, L. H. and W. G. Madow (1944). On the theory of systematic sampling. *Annals of Mathematical Statistics* 15, 1–24.
- Madow, W. G. (1949). On the theory of systematic sampling, II. *Annals of Mathematical Statistics* 20, 333–354.
- Mardia, K. V. (1970). Measures of multivariate skewness and kurtosis with applications. *Biometrika* 57(3), 519–530.
- Mardia, K. V., J. T. Kent, and J. Bibby (1979). *Multivariate Analysis*. Academic Press Inc., San Diego.

- Mehta, C. R., N. R. Patel, and P. Senchaudhuri (2000). Efficient monte carlo methods for conditional logistic regression. *Journal of the American Statistical Association* 95(449), 99–108.
- Milosevic, V. (2022, october). A-IADL-Q - Serbian version in patients with dementia and MCI - dataset.
- Morgan, K. L. (2011, May). *Rerandomization to Improve Covariate Balance in Randomized Experiments*. Ph. D. thesis, Harvard University, Cambridge, MA.
- Morgan, K. L. and D. B. Rubin (2012). Rerandomization to improve covariate balance in experiments. *The Annals of Statistics* 40(2), 1263–1282.
- Nedyalkova, D. and Y. Tillé (2008). Optimal sampling and estimation strategies under linear model. *Biometrika* 95, 521–537.
- Pea, J., L. Qualité, and Y. Tillé (2007). Systematic sampling is a minimal support design. *Computational Statistics & Data Analysis* 51, 5591–5602.
- Potter, D. (2005, 03). A permutation test for inference in logistic regression with small- and moderate-sized data sets. *Statistics in medicine* 24, 693–708.
- Qualité, L. (2008). A comparison of conditional Poisson sampling versus unequal probability sampling with replacement. *Journal of Statistical Planning and Inference* 138, 1428–1432.
- Qualité, L. (2009). *Unequal probability sampling and repeated surveys*. Ph. D. thesis, Université de Neuchâtel.
- R Core Team (2024). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Rivest, L. and S. A. Gaye (2023, April). Using Survey Sampling Algorithms For Exact Inference in Logistic Regression. *International Statistical Review* 91(1), 18–34.
- Robins, J. M., A. Rotnitzky, and L. P. Zhao (1994). Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association* 89(427), 846–866.
- Rubin, D. (1974, 10). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* 66, 688–701.
- Rubin, D. B. (1980). Randomization analysis of experimental data: The fisher randomization test comment. *Journal of the American Statistical Association* 75(371), 591–593.
- Rubin, D. B. (2008). For objective causal inference, design trumps analysis. *Ann. Appl. Stat.* 2(3), 808–840.
- Sampford, M. R. (1967). On sampling without replacement with unequal probabilities of selection. *Biometrika* 54, 499–513.
- Scott, A. J. and T. M. F. Smith (1975). Minimax designs for sample surveys. *Biometrika* 62(2), 353–357.
- Seber, G. (1984). *Multivariate Observations*. Wiley, Hobokon, New Jersey.
- Simon, R. (1979). Restricted randomization designs in clinical trials. *Biometrics* 35, 503–512.

- Srinivasan, A. (2001). Distributions on level-sets with applications to approximation algorithms. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pp. 588–597. IEEE.
- StataCorp (2023). `expoisson` — exact poisson regression, stata 18 base reference manual. *College Station*.
- Stein, C. (1990). Application of Newton’s identities to a generalized birthday problem and to the Poisson-Binomial distribution. Technical Report TC 354, Department of Statistics, Stanford University.
- Stenger, H. (1979). A minimax approach to randomization and estimation in survey sampling. *Annals of Statistics* 7, 395–399.
- Stenger, H. and S. Gabler (1996). A minimax property of Lahiri-Midzuno-Sen’s sampling scheme. *Metrika* 43, 213–220.
- Talovic, E. and Y. Tillé (2025). Risk minimization using robust experimental or sampling designs and mixture of designs. *Journal of Statistical Planning and Inference* 236, 106241.
- Tillé, Y. (2006). *Sampling Algorithms*. New York: Springer.
- Tillé, Y. (2020a). *Sampling and Estimation From Finite Populations*. Hoboken: Wiley.
- Tillé, Y. (2020b). *Sampling and Estimation from Finite Populations*, Chapter 15, pp. 295–332. Wiley.
- Tillé, Y. (2022). Some solutions inspired by survey sampling theory to build effective clinical trials. *International Statistical Review* 90(3), 481–498.
- Tillé, Y. and A. Matei (2021). *sampling: Survey Sampling*. R package version 2.9.
- Tillé, Y. and M. Wilhelm (2017, May). Probability sampling designs: Principles for choice of design and balancing. *Statist. Sci.* 32(2), 176–189.
- U.S. Census Bureau (2024). U.S. Census Bureau cities data set. <https://simplemaps.com/data/us-cities>. Accessed: 2024.
- Xu, Z. and J. D. Kalbfleisch (2010). Propensity score matching in randomized clinical trials. *Biometrics* 66(3), 813–823.
- Yates, F. and P. M. Grundy (1953). Selection without replacement from within strata with probability proportional to size. *Journal of the Royal Statistical Society B15*, 235–261.
- Zamar, D., B. McNeney, and J. Graham (2007). `elrm`: Software implementing exact-like inference for logistic regression models. *Journal of Statistical Software* 21(3), 1–18.