Fast computation of exact confidence intervals for randomized experiments with binary outcomes

P. M. Aronow Haoge Chang Patrick Lopatto*

Abstract

Consider a randomized experiment with a binary outcome, such that half of units are randomly assigned to receive a treatment and the other half are assigned to a control group. In this setting, exact confidence intervals for the average causal effect of the treatment can be computed through a series of permutation tests. This approach requires minimal assumptions and is valid for all sample sizes, as it does not rely on large-sample approximations such as the central limit theorem. We show that these confidence intervals can be found in $O(n \log n)$ permutation tests. Prior to this work, the most efficient known construction was given by Li and Ding, who showed $O(n^2)$ permutation tests suffice. We also demonstrate how to construct confidence intervals using Monte Carlo permutation tests, instead of more computationally demanding exact tests, while retaining the guarantee that the intervals contain the true parameter at no less than the given rate. Taken together, the number of computations required by our method is $O\big(n^2\varepsilon^{-2}(\log n)\log(n\varepsilon^{-1}\log n)\big)$, with a precision loss proportional to ϵ, improved from the previous best of $O(2^n n^{5/2})$. Our results thus facilitate exact inference as a viable option for randomized experiments far larger than those accessible by previous methods.

*P.L. is partially supported by NSF postdoctoral fellowship DMS-220289.

Contents

1 Introduction

There is a vast literature on causal inference in randomized experiments with binary outcomes (see, for example, [IR15] and the references therein). This setting includes a large class of important examples, such as experiments probing the causal effect of an treatment for a dangerous disease, a design change in an online advertisement, or a reminder about the date of an upcoming election. In each case, the outcome for each subject – whether they are hospitalized during the study period, click a banner, or cast a ballot, respectively – is binary.

The study of causal inference for binary outcomes naturally leads to a variety of algorithmic questions, which are the focus of this article. To contextualize our results, we first review the relevant statistical background. We then describe previous work on algorithms for inference on randomized experiments with binary outcomes, and present our contributions. We conclude this introduction by discussing some open problems.

1.1 Statistical Background

A fundamental problem in causal inference is the construction of *confidence intervals* for an unknown causal parameter $\theta \in \mathbb{R}$. We observe data $\boldsymbol{D}$ generated by some random process, such as a randomized experiment, whose distribution is assumed to be function of the unknown θ. We would like to use the observed data $\boldsymbol{D}$ to provide an *interval estimate* for θ. Specifically, we fix some $\alpha \in (0,1)$ and aim to construct an interval $\mathcal{J}_\alpha(\boldsymbol{D}) = \big[L_\alpha(\boldsymbol{D}), U_\alpha(\boldsymbol{D})\big]$ such that

$$\mathbb{P}_\theta\big(\theta \in \mathcal{J}_\alpha(\boldsymbol{D})\big) \geq 1 - \alpha \tag{1.1}$$

holds for all possible θ, where $\mathbb{P}_\theta$ denotes the probability with respect to the random process arising from θ. A conventional choice is $\alpha = 0.05$, leading to the well-known 95% confidence interval. Among intervals that satisfying (1.1), shorter intervals $\mathcal{J}_\alpha$ more precisely estimate θ than longer ones, and are therefore more desirable. Crucially, $\mathcal{J}_\alpha$ contains θ with high probability, regardless of the true value of θ. This uniform coverage property guarantees that if many intervals $\mathcal{J}_\alpha$ are constructed from independent data sets (and α is small), then with high probability most of these intervals will cover their corresponding true parameters.[1] This is the standard justification for the use of confidence intervals in modern science, where they are enormously popular. Further details may be found in [Was04].

Many traditional approaches to constructing confidence intervals for experiments with binary outcomes are based on a binomial model for the outcome distribution. Examples include Wald intervals [Was04, Chapter 10], which are based on a normal approximation valid in large samples, as well as other interval estimators [BB19, SS80]. However, the assumptions underlying the binomial model are highly problematic in typical experimental designs. This was noted by Robins in [Rob88], whose discussion we now briefly summarize.

Consider an experiment with n subjects, where m subjects are assigned to treatment and the remaining $n - m$ are assigned to control, and the random assignment is such that all possible configurations are equally likely. Robins identifies two modeling assumptions under which binomial confidence intervals will cover the true average causal effect at the nominal rate.

1. If a subject is exposed to treatment, their outcome may be modeled as a Bernoulli random variable, and the treatment outcomes across subjects are independent with common mean

[1] Precise quantitative statements of this kind can easily be deduced from concentration of measure bounds. Specifically, one may apply Hoeffding's inequality to the indicator functions for the events $\{\theta \in \mathcal{J}_\alpha\}$.

p_1. Similarly, the control outcomes are independent Bernoulli random variables with common mean p_2. Define the average causal effect of treatment as $p_1 - p_2$.

2. The subjects in the study are drawn uniformly at random from some near-infinite "superpopulation." We let p_1 be the proportion of subjects in the superpopulation who whose outcome would be 1 under treatment, and similarly let p_2 be the proportion whose outcome would be 1 under control. Again define the average causal effect of treatment as $p_1 - p_2$.

Robins notes that Assumption (1) is untenable in most contexts, since it does not account for between-subject variation. For example, disease risk may vary among subjects in a clinical trial according to pre-existing medical conditions and demographic characteristics. Further, Assumption (2) is almost always false, since typical subject recruitment strategies do not result in a uniform sample from a well-defined superpopulation. Consider, for example, a clinical trial where subjects are recruited through newspaper advertisements. It is generally implausible that all members of the target population will read and respond to such advertisements at identical rates. Further, when both Assumption (1) and Assumption (2) fail, it is unclear what unknown parameter θ the binomial confidence intervals are supposed to estimate, and why they should satisfy (1.1).

Given these deficiencies of the binomial model, we adopt instead a finite population model, which we analyze using the *potential outcomes* framework. This perspective originated in work of Neyman and has subsequently been developed by many researchers (see [IR15] for details). Suppose we have a group of n subjects, and label them arbitrarily from 1 to n. Our fundamental assumption is that the observed outcome for subject i depends only on the index i and whether that subject is assigned to treatment or control. This is known as the *stable unit treatment value assumption*, and abbreviated SUTVA. In the particular, SUTVA rules out between-subject interaction effects. We let $\boldsymbol{y}_i = \big(y_i(0), y_i(1)\big)\big)$ with $y_i(0), y_i(1) \in \{0, 1\}$ denote the *potential outcomes* for the i-th subject, where $y_i(1)$ is the outcome that would be observed if the subject were assigned to treatment, and $y_i(0)$ is the outcome that would be observed if the subject were assigned to control. In the experiment, only *one* of these potential outcomes is observed; the other remains unknown. We aim to estimate the *sample average treatment effect*, defined as

$$\tau(\boldsymbol{y}) = \frac{1}{n} \sum_{j=1}^{n} \big(y_j(1) - y_j(0)\big). \tag{1.2}$$

The quantity $\tau(\boldsymbol{y})$ is defined using only the subjects in the experiment, and says nothing about the causal effect in any larger population. Such a generalization would require additional assumptions. However, precise estimates of $\tau(\boldsymbol{y})$ are still broadly useful for studying causal claims. For example, if a company asserts that their latest drug greatly reduces the chance of stroke in elderly people, and a study of the average causal effect in a group of such people produces a precise estimate centered at 0, it is reasonable to reject the company's claim.

Since we view the potential outcomes of the subjects as fixed, the only randomness in the experiment comes from the assignment of subjects to treatment and control. We retain our previous setup where m subjects are assigned to treatment, the remainder are assigned to control, and all such assignments are equally likely. Under this design, Robins gives a confidence interval for $\tau(\boldsymbol{y})$ that relies on a large-sample normal approximation, which is explicitly justified by randomization. While this interval is guaranteed to cover $\tau(\boldsymbol{y})$ at its nominal rate in the limit as n goes to infinity, there is no quantitative estimate available for its accuracy.

1.2 Algorithmic Perspective and Main Results

We now explain how the construction of confidence intervals for the causal effect in a randomized experiment with a binary outcome, where n subjects are randomized to groups of fixed sizes m and $n-m$, may be cast as a computational problem. Specifically, we seek confidence intervals that cover $\tau(\boldsymbol{y})$ with probability at least $1-\alpha$ regardless of the sample size or potential outcomes (in the sense of (1.1)).

Define a random vector $\boldsymbol{Z} \in \{0,1\}^n$ by letting $Z_i = 0$ if subject i is assigned to control, and $Z_i = 1$ if subject i is assigned to treatment. We let $\boldsymbol{Y}$ be the vector of observed outcomes, where

$$Y_j = Z_j y_j(1) + (1 - Z_j) y_j(0) \tag{1.3}$$

represents the outcome that the experimenter observes for subject j after the experiment is completed. Using the known vectors $\boldsymbol{Z}$ and $\boldsymbol{Y}$, the unknown parameter $\tau(\boldsymbol{y})$ may be estimated using the *Neyman estimator* $\boldsymbol{T}(\boldsymbol{Y}, \boldsymbol{Z})$. Informally speaking, this estimator subtracts the average of the observed outcomes in the control group from the average of the observed outcomes in the treatment group. The precise definition is

$$T(\boldsymbol{Y}, \boldsymbol{Z}) = \sum_{j=1}^{n} \frac{Z_j Y_j}{m} - \sum_{j=1}^{n} \frac{(1 - Z_j) Y_j}{n - m}. \tag{1.4}$$

An elementary calculation shows that $\mathbb{E}\big[T(\boldsymbol{Y}, \boldsymbol{Z})\big] = \tau(\boldsymbol{y})$, where the expectation is taken over the variable $\boldsymbol{Z}$. Hence $T(\boldsymbol{Y}, \boldsymbol{Z})$ is an *unbiased* estimator for $\tau(\boldsymbol{y})$.

As a preliminary step toward constructing a confidence interval for $\tau(\boldsymbol{y})$, we consider the probability that some possible configuration $\boldsymbol{w}$ of the potential outcomes could produce data at least as extreme as the observed data $(\boldsymbol{Y}, \boldsymbol{Z})$. To make this notion precise, we consider an arbitrary length n vector $\boldsymbol{w}$ of elements $\boldsymbol{w}_i$ such that $w_i(0), w_i(1) \in \{0,1\}$, and compute the *p-value* given by

$$p(\boldsymbol{w}, \boldsymbol{Y}, \boldsymbol{Z}) = \mathbb{P}\left(\big| T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}}) - \tau(\boldsymbol{w}) \big| \geq \big| T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau(\boldsymbol{w}) \big| \right). \tag{1.5}$$

Here $\widetilde{\boldsymbol{Z}}$ is independent from $\boldsymbol{Z}$ and has an identical distribution, and $\widetilde{\boldsymbol{Y}}$ is the vector of observed outcomes in a hypothetical experiment with potential outcomes $\boldsymbol{w}$ and randomization $\widetilde{\boldsymbol{Z}}$, so that $\widetilde{Y}_j = \widetilde{Z}_j w_j(1) + (1 - \widetilde{Z}_j) w_j(0)$. Calculating the p-value (1.5) is known as performing a *permutation test*, since the probability may be computed by averaging over all possible randomizations $\widetilde{\boldsymbol{Z}}$.

Recall that $\mathbb{E}\big[T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}})\big] = \tau(\boldsymbol{w})$. Then, informally speaking, the p-value measures how atypical (far from the center of the distribution) the observed data $T(\boldsymbol{Y}, \boldsymbol{Z})$ is, under the assumption that $\boldsymbol{w}$ is the true set of potential outcomes for the subjects in the experiment. The permutation test (1.5) was essentially known to Fisher in the 1930s [IR15, Chapter 5], and can be used to test the hypothesis that the true potential outcomes $\boldsymbol{y}$ are equal to some specific $\boldsymbol{w}$ (by rejecting this hypothesis if $p(\boldsymbol{w}, \boldsymbol{Y}, \boldsymbol{Z})$ is sufficiently small).

Recently, it was realized by Rigdon and Hudgens that exact confidence intervals for $\tau(\boldsymbol{y})$ could be constructed from a series of these permutation tests [RH15], in the following way. Given observed data $(\boldsymbol{Y}, \boldsymbol{Z})$, consider *all* potential outcome vectors $\boldsymbol{w}$ that could produce $(\boldsymbol{Y}, \boldsymbol{Z})$, if $\boldsymbol{w}$ were the true unknown set of potential outcomes.[2] (The binary outcome assumption ensures that this set is finite.) Fix $\alpha \in (0,1)$, and declare such a $\boldsymbol{w}$ *compatible* if $p(\boldsymbol{w}, \boldsymbol{Y}, \boldsymbol{Z}) \geq \alpha$, and *incompatible* otherwise. Let $L_\alpha(\boldsymbol{Y}, \boldsymbol{Z})$ equal the smallest value of $\tau(\boldsymbol{w})$ where $\boldsymbol{w}$ ranges over all compatible

[2] This already rules out many vectors $\boldsymbol{w}$. For example, if $n = 2$ and we randomize into equal groups of size 1, and the observed data yields $T(\boldsymbol{Y}, \boldsymbol{Z}) = 1$, then the configuration $\boldsymbol{w}_1 = \boldsymbol{w}_2 = (0,0)$ is not possible.

$\boldsymbol{w}$, and let $U_\alpha(\boldsymbol{Y}, \boldsymbol{Z})$ equal the largest value of $\tau(\boldsymbol{w})$ where $\boldsymbol{w}$ ranges over all compatible $\boldsymbol{w}$. Set $\mathcal{I}_\alpha = \big[L_\alpha(\boldsymbol{X}), U_\alpha(\boldsymbol{X})\big]$. Then it is straightforward to show that $\mathcal{I}_\alpha$ will contain $\tau(\boldsymbol{y})$ with probability at least $1 - \alpha$, regardless of the value of $\boldsymbol{y}$.[3]

Constructing $\mathcal{I}_\alpha$ through permutation tests eliminates any reliance on uncontrolled large-sample approximations, but has the drawback of requiring significant computational resources. While a naive construction of $\mathcal{I}_\alpha$ might proceed by directly computing (1.5) for all possible $\boldsymbol{w}$ to find L_α and U_α, this is extremely inefficient, and infeasible even for small n. The problem of efficiently determining $\mathcal{I}_\alpha$ then reduces to two subproblems. First, minimize the number of permutation tests required to find L_α and U_α. Second, minimize the computational effort required to perform a permutation test.

Regarding the first subproblem, Rigdon and Hudgens showed that $\mathcal{I}_\alpha$ can be constructing using at most $O(n^4)$ permutation tests [RH15]. Later, Li and Ding showed that the same intervals could be constructed in $O(n^2)$ permutation tests in the balanced case, where an equal number of subjects are assigned to treatment and control (that is, $n = 2m$), and conjectured that their algorithm also reproduces the Rigdon–Hudgens intervals in unbalanced trials [LD16]. Specifically, they showed that the interval their algorithm returns always contains the corresponding Rigdon–Hudgens interval $\mathcal{I}_\alpha$, but were unable to exclude the possibility that theirs are strictly larger for unbalanced designs.

We now consider the second subproblem. An *exact permutation test* enumerates all $\binom{n}{m}$ possible assignments of subjects, which grows exponentially in n and quickly becomes computationally infeasible. Rigdon and Hudgens therefore recommend performing approximate permutation tests, where the permutation p-value is approximated by Monte Carlo simulation with a fixed number of samples K from the distribution of $\widetilde{\boldsymbol{Z}}$ in (1.5). However, such intervals may fail to cover at the nominal rate due to Monte Carlo error. Further, previous literature has not provided any formal analysis regarding the question of how large K should be taken so that the resulting intervals $\mathcal{I}_\alpha$ cover the unknown parameter $\tau(\boldsymbol{y})$ with high probability. Such an analysis is not entirely straightforward, since if one uses approximate permutation tests in the Rigdon–Hudgens and Li–Ding algorithms described previously, the acceptance and rejection decisions made by successive tests may be highly correlated.

We make two contributions, corresponding to the two subproblems identified previously. We focus on the case of balanced experiments, where the treatment and control groups have equal size, as this setting is common in practice and it is minimax optimal in terms of efficiency in estimation.[4] First, we show that the exact confidence intervals $\mathcal{I}_\alpha$ from [RH15] can be constructed in $O(n \log n)$ permutation tests, improving on the $O(n^2)$ tests required by [LD16]. Our proof is based on a new monotonicity property for the permutation p-values (1.5), which permits us to greatly reduce the effective search space when finding L_α and U_α. Second, we show how to construct confidence intervals using approximate permutation tests that are guaranteed to cover the sample average treatment effect with probability $1 - \alpha$ for any $\alpha > 0$. We further show that by increasing K, these intervals will approximate those of [RH15] arbitrarily well, and quantify how K must scale in order to guarantee any desired coverage rate and approximation accuracy as n increases. The Monte Carlo guarantees are proved by a careful application of standard concentration inequalities, taking into account the details of our new algorithm for constructing $\mathcal{I}_\alpha$.

In fact, we provide finite-sample bounds for both subproblems, with small, explicit constants. We are able to construct exact confidence intervals in $8n \log_2(n)$ permutation tests, and guarantee

[3] See Lemma 3.1 in Section 3, below.

[4] See also [Kal18, Section 2] for a justification of balanced randomization via its blinding properties. In general, [Kal18] shows that balanced complete randomization enjoys a minimax property even when potentially prognostic covariates are available to the designer of the experiment.

that our Monte Carlo intervals cover the unknown parameter $\tau(\boldsymbol{y})$ with probability $1-\alpha$ if

$$K \geq \frac{1}{\varepsilon^2} \log\left(\frac{16n\log_2(n)}{\varepsilon}\right), \tag{1.6}$$

where $\varepsilon < \alpha$ is an approximation parameter defined in Section 5, and smaller values of ε correspond to confidence intervals that are smaller on average. (For both of these bounds, we have assumed that $n \geq 10$ in order to state them simply.) Observe that a single Monte Carlo sample of $T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}})$ can be evaluated in $O(n)$ arithmetical operations, by using the Fisher–Yates shuffle algorithm to generate the randomization $\widetilde{\boldsymbol{Z}}$ into equal groups. Then the overall complexity of algorithm producing the Monte Carlo intervals is

$$O\left(\frac{n^2\log n}{\varepsilon^2}\log\left(\frac{n\log(n)}{\varepsilon}\right)\right). \tag{1.7}$$

We note that the previous best construction of permutation confidence intervals with guaranteed coverage in [LD16] requires $O(2^n n^{5/2})$ operations, due to the $O(n^2)$ permutation tests required, the $\binom{2m}{m} = O(2^n n^{-1/2})$ randomizations required for an exact permutation test, and the $O(n)$ operations required for a single randomization.

Our algorithmic contributions permit permutation inference to be applied to experiments far larger than those accessible by previous methods. Additionally, we also extend our analysis to experiments with missing data, a common statistical problem that has not been addressed previously in this context.

1.3 Open Problems

Many interesting questions remain. First, it is natural to ask whether an even more efficient construction of $\mathcal{I}_\alpha$ is possible, either by improving the number of number of permutation tests necessary or finding a more efficient way to compute (approximate) permutation p-values. We are unaware of any formal analysis of the computational hardness of such questions. Regarding the Monte Carlo permutation tests, we do not know how to implement any of the standard variance reduction techniques in this context, such as importance sampling, which would provide a practical (though non-asymptotic) improvement in run time. We note that these methods have been successfully applied to permutation inference in other contexts [BB19, MPS88].

Second, a number of questions arise from relaxing the various assumptions made earlier. For example, is efficient approximate inference possible in the case of a general bounded and discrete outcome, or a continuous outcome supported on a bounded interval? Or for unbalanced designs? Additionally, a more realistic setup is to suppose that each subject comes with additional data, such as demographic information and medical history, that is correlated with the potential outcomes. Rigdon and Hudgens show that their algorithm adapts to designs using stratified randomization, where subjects are independently randomized to treatment or control within subgroups formed using these additional covariates [RH15]. However, they observe that this approach quickly becomes computationally intractable, requiring roughly $O(n^{4k})$ permutation tests, where k is the number of strata. It would be interesting to find a fast method to analyze stratified designs, and to investigate more sophisticated randomization mechanisms involving clustering, rerandomization, and matching [MR12, GLSR04].

Finally, it is of practical importance to find an efficient algorithm for prospective sample size calculations for the exact intervals $\mathcal{I}_\alpha$. Researchers often wonder how many experimental subjects they need to detect an effect with high probability (that is, have $0 \notin \mathcal{I}_\alpha$), assuming the effect is larger than some given threshold. If too few subjects are used, the experiment is uninformative,

and if too many subjects are used, the experiment is potentially wasteful. While it is well known how to perform sample size estimation in simple settings such as the binomial model discussed earlier, we are unaware of any rigorous results for the Rigdon–Hudgens permutation intervals.

1.4 Outline

We begin in Section 2 by introducing our notation and reviewing previous work in detail. In Section 3, we present basic properties of the confidence intervals constructed in [RH15]. In Section 4, we show how to construct these intervals in $O(n \log n)$ permutation tests. In Section 5, we construct exact intervals using Monte Carlo permutation tests. We consider missing data in Section 6. The remaining sections contain proofs of the assertions in the previous sections.

2 Preliminaries

2.1 Notation

Let $\mathbb{Z}_{>0}$ denote the positive integers. We consider a group of $n \in \mathbb{Z}_{>0}$ subjects who undergo an experiment with a binary outcome. We always suppose that the experiment is balanced, so that $n = 2m$ for some $m \in \mathbb{Z}_{>0}$, m subjects are assigned to treatment, and the remaining m subjects are assigned to control.

A *potential outcome table* is any vector

$$\boldsymbol{w} \in \big\{(0,0),(0,1),(1,0),(1,1)\big\}^n. \tag{2.1}$$

Letting $\boldsymbol{w}_i = \big(w_i(0), w_i(1)\big)$ for $i \in [\![1,n]\!]$, the coordinate $w_i(1)$ denotes the outcome for the i-th subject under treatment, and $w_i(0)$ denotes the outcome under control. We use $\boldsymbol{y}$ to denote the potential outcome table for the subjects in the experiment, which is a fixed (but unknown) ground truth. We use $\boldsymbol{w}$ to denote an arbitrary vector in $\big\{(0,0),(0,1),(1,0),(1,1)\big\}^n$. Given some table $\boldsymbol{w}$, we let the corresponding count vector $\boldsymbol{v} = \boldsymbol{v}(\boldsymbol{w}) \in \mathbb{Z}_{\geq 0}^{\{0,1\}^2}$ be defined by

$$\boldsymbol{v} = (v_{11}, v_{10}, v_{01}, v_{00}), \qquad v_{ab} = \sum_{j=1}^{n} \mathbb{1}\{\boldsymbol{w}_j = (a,b)\}. \tag{2.2}$$

We now introduce notation for the randomization of the subjects. Consider the random vector

$$\boldsymbol{Z} = (Z_1, Z_1, \ldots, Z_n) \in \{0,1\}^n \tag{2.3}$$

whose distribution is uniform over the set

$$\mathcal{Z}(n) = \left\{ \boldsymbol{z} \in \{0,1\}^n : \sum_{i=1}^{n} z_i = m \right\}. \tag{2.4}$$

Here, the indices i such that $Z_i = 1$ represent the subjects assigned to treatment. The vector of observed experimental outcomes $\boldsymbol{Y} = (Y_j)_{j=1}^n$ is then given by

$$Y_j = Z_j y_j(1) + (1 - Z_j) y_j(0). \tag{2.5}$$

We say that a table $\boldsymbol{w}$ is *possible* given the observed data $\boldsymbol{Y}$ if for every $j \in [\![1,n]\!]$, we have $\boldsymbol{w}_j(1) = Y_j$ if $Z_j = 1$, and $\boldsymbol{w}_j(0) = Y_j$ if $Z_j = 0$. We say that a count vector $\boldsymbol{v}$ is possible if it arises

from some possible table $\boldsymbol{w}$.

The sample average treatment effect for a potential outcome table is defined by

$$\tau(\boldsymbol{w}) = \frac{1}{n}\sum_{j=1}^{n}\big(w_j(1) - w_j(0)\big). \tag{2.6}$$

The Neyman estimator of $\tau(\boldsymbol{y})$ is given by

$$T(\boldsymbol{Y}, \boldsymbol{Z}) = \sum_{j=1}^{n}\frac{Z_j Y_j}{m} - \sum_{j=1}^{n}\frac{(1-Z_j)Y_j}{n-m}. \tag{2.7}$$

This estimator depends on $(\boldsymbol{Y}, \boldsymbol{Z})$ only through the observed count vector $\boldsymbol{n} \in \mathbb{Z}_{\geq 0}^{\{0,1\}^2}$ defined by

$$\boldsymbol{n} = (n_{11}, n_{10}, n_{01}, n_{00}), \qquad n_{zy} = \sum_{j=1}^{n}\mathbb{1}\{Z_j = z, Y_j = y\}, \tag{2.8}$$

for $z \in \{0,1\}$ and $y \in \{0,1\}$. For example, n_{10} is the number of subjects in the treatment group with observed outcome 0. We overload the notation slightly and write $T(\boldsymbol{n})$ for the value of the Neyman estimator $T(\boldsymbol{Y}, \boldsymbol{Z})$, if $\boldsymbol{n}$ is the observed count vector for $(\boldsymbol{Y}, \boldsymbol{Z})$. We will do this without further comment for all other functions of $(\boldsymbol{Z}, \boldsymbol{Y})$ that depend only on the associated vector $\boldsymbol{n}$, and similarly for functions of tables $\boldsymbol{w}$ that depend only on the associated count vector $\boldsymbol{v}$.

Throughout the paper, $[\![a, b]\!]$ denotes the set $\{k \in \mathbb{Z} : a \leq k \leq b\}$. We write log for the natural logarithm. When we use the base 2 logarithm, it is always denoted by $\log_2$. Additionally, we often abbreviate probability mass function as pmf. Further notation, used only in the proofs, is introduced in Section 7.

2.2 Previous Work

2.2.1 Rigdon and Hudgens

We begin by recalling the confidence interval procedure proposed in [RH15]. As $\boldsymbol{w}$ ranges over all tables, $\tau(\boldsymbol{w})$ takes every value in the set

$$\mathcal{S}(n) = \left\{-\frac{n}{n}, -\frac{n-1}{n}, \ldots, \frac{0}{n}, \ldots, \frac{n-1}{n}, \frac{n}{n}\right\}, \tag{2.9}$$

which has $2n+1$ elements. After the experiment is completed, Rigdon and Hudgens note that the set of $\tau(\boldsymbol{w})$ for $\boldsymbol{w}$ that are possible given the observed data is further restricted to

$$\mathcal{C}(\boldsymbol{Y}, \boldsymbol{Z}) = \left\{\frac{1}{n}\left(\sum_{j=1}^{n} Y_j(2Z_j - 1) - m\right), \ldots, \frac{1}{n}\left(\sum_{j=1}^{n} Y_j(2Z_j - 1) - m + n\right)\right\}, \tag{2.10}$$

which has $n+1$ elements. We observe that $\mathcal{C}$ depends on $(\boldsymbol{Y}, \boldsymbol{Z})$ only through the associated observed count vector $\boldsymbol{n}$.

We now consider some $\alpha \in (0,1)$ and a realization of $(\boldsymbol{Y}, \boldsymbol{Z})$. We will construct a confidence interval for τ at level $1-\alpha$ (that is, one that contains the true parameter with probability at least $1-\alpha$). We first describe a permutation test for the sharp null hypothesis that $\boldsymbol{y} = \boldsymbol{w}$ for some given $\boldsymbol{w}$, then explain how to form a confidence interval from a series of these tests.

Definition 2.1. Let $\widetilde{\boldsymbol{Z}}$ be a random vector independent from $\boldsymbol{Z}$ with the same distribution, and let $\widetilde{\boldsymbol{Y}}$ be the vector with entries

$$\widetilde{Y}_j = \widetilde{Z}_j w_j(1) + (1 - \widetilde{Z}_j) w_2(0). \tag{2.11}$$

We say that $\boldsymbol{w}$ is compatible with data $(\boldsymbol{Y}, \boldsymbol{Z})$ at level $1 - \alpha$ if[5]

$$p(\boldsymbol{w}, \boldsymbol{Y}, \boldsymbol{Z}) = \mathbb{P}\left(\left| T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}}) - \tau(\boldsymbol{w}) \right| \geq \left| T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau(\boldsymbol{w}) \right| \right) \geq \alpha. \tag{2.12}$$

We say that $\boldsymbol{w}$ is compatible with an observed count vector $\boldsymbol{n}$ if (2.12) holds with $T(\boldsymbol{Y}, \boldsymbol{Z})$ replaced by $T(\boldsymbol{n})$, and define $p(\boldsymbol{w}, \boldsymbol{n})$ similarly.

Let $\boldsymbol{v}$ denote the count vector associated to $\boldsymbol{w}$, as defined in (2.2), and let $\boldsymbol{n}$ be as in (2.8). Then it is straightforward to see that $p(\boldsymbol{w}, \boldsymbol{Y}, \boldsymbol{Z}) = p(\boldsymbol{v}, \boldsymbol{n})$, so that the p-value depends on $(\boldsymbol{w}, \boldsymbol{Y}, \boldsymbol{Z})$ only through the count vectors $\boldsymbol{n}$ and $\boldsymbol{v}$.

Algorithm 2.2. A *permutation test* takes as inputs $\alpha \in (0, 1)$, a potential outcome table $\boldsymbol{w}$, and a vector $\boldsymbol{n} \in \mathbb{Z}_{\geq 0}^{\{0,1\}^2}$. It returns a binary decision, accept or reject, as follows. The algorithm computes $p(\boldsymbol{w}, \boldsymbol{n})$ by direct enumeration of all $\binom{n}{m}$ elements of the set $\mathcal{Z}(n)$ (defined in (2.4)), over which $\widetilde{\boldsymbol{Z}}$ is uniformly distributed. If $p(\boldsymbol{w}, \boldsymbol{n}) \geq \alpha$, it accepts. Otherwise, it rejects.

By using $p(\boldsymbol{v}, \boldsymbol{n})$ in place of $p(\boldsymbol{w}, \boldsymbol{n})$, we can generalize the previous definition to use count vectors $\boldsymbol{v}$ as inputs.

We now present the confidence interval procedure from [RH15]. Given an observed count vector $\boldsymbol{n}$, set

$$\mathcal{T}(\boldsymbol{n}) = \big\{ (i, j, k, l) : i \in [\![0, n_{11}]\!], j \in [\![0, n_{10}]\!], k \in [\![0, n_{01}]\!], l \in [\![0, n_{00}]\!] \big\}. \tag{2.13}$$

We define vectors $\boldsymbol{v}(i, j, k, l)$ in the following way. Given $(i, j, k, l) \in \mathcal{T}(\boldsymbol{n})$, we set

$$\boldsymbol{v}(i, j, k, l) = (i + k, n_{11} - i + l, n_{01} - k + j, n_{10} + n_{00} - j - l). \tag{2.14}$$

To motivate this definition, we think of starting with the observed count vector $\boldsymbol{n}$, and imputing various combinations of the unknown potential outcomes. In the n_{11} subjects with observed outcome 1 in the treatment group, we impute 1 for the unknown outcomes in exactly i of them, to get i potential outcome pairs $(1, 1)$. Similarly, j, k, l represent the number of 1's imputed in the other three combinations of group and outcome. It is straightforward to see that the collection of $\boldsymbol{v}(i, j, k, l)$ such that $(i, j, k, l) \in \mathcal{T}(\boldsymbol{n})$ is exactly the set of count vectors $\boldsymbol{v}$ possible given the observed count vector $\boldsymbol{n}$. (However, the map $(i, j, k, l) \mapsto \boldsymbol{v}(i, j, k, l)$ is not necessarily injective.)

Algorithm 2.3. The following algorithm takes as input $\alpha \in (0, 1)$ and a vector $\boldsymbol{n} \in \mathbb{Z}_{\geq 0}^{\{0,1\}^2}$. It returns an interval $\mathcal{I}_\alpha(\boldsymbol{n}) = \big[U_\alpha(\boldsymbol{n}), L_\alpha(\boldsymbol{n})\big]$.

By directly enumerating of all elements of $\mathcal{T}$, and applying Algorithm 2.2 to each one, compute the set

$$\mathcal{K}(\boldsymbol{n}) = \Big\{ \tau\big(\boldsymbol{v}(i, j, k, l)\big) : (i, j, k, l) \in \mathcal{T}(\boldsymbol{n}) \text{ and Algorithm 2.2 accepts } \boldsymbol{v}(i, j, k, l) \Big\}. \tag{2.15}$$

Return $U_\alpha(\boldsymbol{n}) = \max(\mathcal{K})$ and $L_\alpha = \min(\mathcal{K})$.

[5] In fact, the inequality $\geq \alpha$ can be replaced with the inequality $> \alpha$ in (2.12), which creates a stricter threshold for acceptance. This change leads to potentially shorter confidence intervals that still cover at the nominal rate, which can be seen by inspecting the proof of Proposition 3.2. Further, all of our results are still valid under the stricter definition. The definition (2.1) was used in the previous works [RH15, LD16], and we retain it for consistency.

As noted previously, any observed data $(\boldsymbol{Y}, \boldsymbol{Z})$ gives rise to an observed count vector $\boldsymbol{n}$. Then using Algorithm 2.3, we define

$$\mathcal{I}_\alpha(\boldsymbol{Y}, \boldsymbol{Z}) = \big[L_\alpha(\boldsymbol{Y}, \boldsymbol{Z}), U_\alpha(\boldsymbol{Y}, \boldsymbol{Z})\big] \tag{2.16}$$

by setting $\mathcal{I}_\alpha(\boldsymbol{Y}, \boldsymbol{Z}) = \mathcal{I}_\alpha(\boldsymbol{n})$. Rigdon and Hudgens observe that since $\mathcal{T}(\boldsymbol{n})$ has size $O(n^4)$, only $O(n^4)$ permutation tests are needed to construct the set $\mathcal{K}(\boldsymbol{n})$ in Algorithm 2.3. Further, this algorithm does not require our standing hypothesis that the treatment and control groups have an equal number of subjects, and may be applied to arbitrary unbalanced trials.

2.2.2 Li and Ding

As noted in the introduction, Li and Ding have proposed an algorithm that returns $\mathcal{I}_\alpha(\boldsymbol{n})$ in $O(n^2)$ permutation tests for balanced trials [LD16]. We do not give the details here, since they are not relevant to our work. However, we will make use of the following lemmas from [LD16]. The first provides a necessary and sufficient condition for checking whether a potential outcome table is possible given the observed data.

Lemma 2.4 ([LD16, Theorem 1])**.** *A potential outcome table $\boldsymbol{w}$ with count vector $\boldsymbol{v}$ is possible given observed data $\boldsymbol{n}$ if and only if*

$$\begin{aligned} &\max(0, n_{11} - v_{10}, v_{11} - n_{01}, v_{11} + v_{01} - n_{10} - n_{01}) \\ &\leq \min(v_{11}, n_{11}, v_{11} + v_{01} - n_{01}, n - v_{10} - n_{01} - n_{10}). \end{aligned} \tag{2.17}$$

The second lemma shows that in the balanced case, the confidence set that we would obtain if we tested *every* possible value of τ_0 is indeed an interval.

Lemma 2.5 ([LD16, Theorem A.4])**.** *Fix $\alpha \in (0, 1)$ and observed data $\boldsymbol{y}$ and $\boldsymbol{Z}$. For every $\tau_0 \in \mathcal{I}_\alpha(\boldsymbol{y}, \boldsymbol{Z}) \cap \mathcal{C}$, there exists a possible potential table $\boldsymbol{w}$ such that $\tau(\boldsymbol{w}) = \tau_0$ and $p(\boldsymbol{w}) \geq \alpha$.*

3 Basic Properties

We now note three basic properties of the confidence intervals $\mathcal{I}_\alpha$. While the first two have been mentioned informally in previous works, they have not been proved, and we feel there is value in giving precise statements and justifications. The third is new. The proofs appear in Section 7. The arguments for all results in this section extend to the unbalanced case, where the treatment group has size $m = cn$ for $c \in (0, 1)$, with only minor changes. We omit these extensions for brevity.

Our first lemma implies that the interval $\mathcal{I}_\alpha(\boldsymbol{n})$ always contains the estimate $T(\boldsymbol{n})$ of $\tau(y)$.

Lemma 3.1. *For any $\boldsymbol{n} \in \mathbb{Z}_{\geq 0}^{\{0,1\}^2}$, we have*

$$L_\alpha(\boldsymbol{n}) \leq T(\boldsymbol{n}) \leq U_\alpha(\boldsymbol{n}). \tag{3.1}$$

The next proposition states that the intervals $\mathcal{I}_\alpha$ are exact.

Proposition 3.2. *Fix potential outcomes $\boldsymbol{y}$ and $\alpha \in (0, 1)$. Then*

$$\mathbb{P}\big(\tau(\boldsymbol{y}) \in \mathcal{I}_\alpha(\boldsymbol{Y}, \boldsymbol{Z})\big) \geq 1 - \alpha, \tag{3.2}$$

where the probability is with respect to the variable $\boldsymbol{Z}$.

Finally, the following proposition states that the intervals $\mathcal{I}_\alpha$ converge at a $n^{-1/2}$ rate. This rate is characteristic of many confidence interval procedures based on central limit theorem asymptotics. We therefore see that the large n scaling behavior of the length of the interval $\mathcal{I}_\alpha$ is the same as that of the intervals produced by asymptotic methods.[6]

Proposition 3.3. *For any $\boldsymbol{n} \in \mathbb{Z}_{\geq 0}^{\{0,1\}^2}$, we have*

$$\left|I_\alpha(\boldsymbol{n})\right| \leq \sqrt{\frac{32\log(2/\alpha)}{n}}. \tag{3.3}$$

4 Fast Computation of Confidence Intervals

We now present our main result on the efficient computation of $\mathcal{I}_\alpha$. Its proof is given in Section 8.2.

Theorem 4.1. *Suppose $n \geq 10$. For any $\alpha \in (0,1)$ and observed count vector $\boldsymbol{n}$, the interval $\mathcal{I}_\alpha(\boldsymbol{n})$ can be constructed using at most $8n\log_2 n$ permutation tests.*

To prove Theorem 4.1, we exhibit an algorithm that constructs $\mathcal{I}_\alpha(\boldsymbol{n})$ in the required number of permutation tests.[7] We begin with a heuristic argument. We first observe that the problem of finding the upper bound $U_\alpha(\boldsymbol{n})$ reduces to finding a method to determine whether a given value $\tau_0 \in \mathcal{C}$ is incompatible with the observed data $\boldsymbol{n}$, in the sense that $p(\boldsymbol{v}, \boldsymbol{n}) < \alpha$ for all $\boldsymbol{v}$ possible given $\boldsymbol{n}$ such that $\tau(\boldsymbol{v}) = \tau_0$. If this can be done in $O(n)$ permutation tests, then we can perform a binary search on the set $\Big[T(\boldsymbol{n}), \max\big(C(\boldsymbol{n})\big)\Big] \cap \mathcal{S}(n)$ to find U_α by testing each value τ_0 considered by the binary search for compatibility with $\boldsymbol{n}$. Since Lemma 2.5 guarantees that every element $\mathcal{I}_\alpha(\boldsymbol{n})$ is compatible with the observed data, and Lemma 3.1 guarantees that $\mathcal{I}_\alpha(\boldsymbol{n})$ contains $T(\boldsymbol{n})$, the binary search will return the maximum of

$$\Big[T(\boldsymbol{n}), \max\big(C(\boldsymbol{n})\big)\Big] \cap \mathcal{S}(n) \cap \mathcal{I}_\alpha(\boldsymbol{n}), \tag{4.1}$$

which is equal to $U_\alpha(\boldsymbol{n})$. Since binary search on a set of size $O(n)$ can be completed in $O(\log n)$ time, in total $O(n\log n)$ permutation tests are needed. Analogous reasoning applies for finding the lower bound L_α.

Consequently, the proof of Theorem 4.1 focuses on showing that a given $\tau_0 \in \mathcal{C}$ can be checked for compatibility in $O(n)$ permutation tests. Observe that any possible count vector $\boldsymbol{v} = (v_{11}, v_{10}, v_{01}, v_{00})$ such that $\tau(\boldsymbol{v}) = \tau_0$ satisfies the equations

$$v_{11} + v_{10} + v_{01} + v_{00} = n, \qquad v_{10} - v_{01} = n\tau_0. \tag{4.2}$$

The following lemma allows us to efficiently search the space cut out by (4.2) by exploiting a certain monotonicity property of the permutation p-values defined in (2.12). It is proved in Section 8.3.

Lemma 4.2. *Fix observed data $\boldsymbol{n}$, and a count vector $\boldsymbol{v} = (v_{11}, v_{10}, v_{01}, v_{00})$. Suppose $\min(v_{10}, v_{01}) \geq 1$ and $\max(v_{10}, v_{01}) \geq 2$, and set*

$$\boldsymbol{v}' = (v_{11} + 1, v_{10} - 1, v_{01} - 1, v_{00} + 1). \tag{4.3}$$

[6] In Proposition 3.3, our focus is on establishing the correct scaling rate, not obtaining the optimal constant prefactor. This estimate is quite conservative compared to the interval lengths we find empirically through simulation. It should not be used in practice for prospective sample size calculations.

[7] Again, the estimate in Theorem 4.1 is conservative relative to the empirical performance of Algorithm 4.3 below. To streamline the arguments, we did not pursue the sharpest possible bounds.

Then $p(\boldsymbol{v}', \boldsymbol{n}) \geq p(\boldsymbol{v}, \boldsymbol{n})$.

The upshot of this lemma is that if $p(\boldsymbol{v}, \boldsymbol{n}) < \alpha$, meaning $\boldsymbol{v}$ is incompatible with the observed data, then all potential outcome tables on the line segment given by the translations (4.3), with endpoint $\boldsymbol{v}$, are also incompatible. Heuristically, the transformation in (4.3) causes the distribution of the variable $T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}}) - \tau(\boldsymbol{w})$ in the definition (2.12) to become more spread out about its mean (zero), increasing the p-value. This heuristic is made precise in the proof of Lemma 4.2.

Given this context, we now present our main algorithm. The additional check in Step (2c) exists to handle the case where Lemma 4.2 does not apply.

Algorithm 4.3. This algorithm takes as input $\alpha \in (0, 1)$ and a vector $\boldsymbol{n} \in \mathbb{Z}_{\geq 0}^{\{0,1\}^2}$. It returns an interval $\left[L_\alpha^{(1)}(\boldsymbol{n}), U_\alpha^{(1)}(\boldsymbol{n})\right]$. We give only the steps for finding $U_\alpha^{(1)}(\boldsymbol{n})$, since finding $L_\alpha^{(1)}(\boldsymbol{n})$ is analogous.

We perform a binary search using a function $f\colon [\![nT(\boldsymbol{n}), n\max(\mathcal{C}(\boldsymbol{n}))]\!] \to \{0, 1\}$ defined below. The precise definition of this binary search is given as Algorithm 8.1 in Section 8.1. Given an input x, the function f returns 0 if x/n is compatible with the vector $\boldsymbol{n}$, and 1 otherwise, where the compatibility of a given $\tau_0 \in \mathcal{C}(\boldsymbol{n})$ is determined through the following steps.

1. Initialize $j \leftarrow 0$.
2. If $j \leq n$, perform the following steps.
   (a) Construct a count vector $\boldsymbol{v}$ as follows (recall (2.2)). Let v_{10} be a free parameter, and set $v_{11} = j - v_{10}$. Then (solving (4.2)) we set
   $$v_{01} = v_{10} - n\tau_0, \qquad v_{00} = n - j - v_{10} + n\tau_0. \tag{4.4}$$
   (b) Find the smallest v_{10} that leads to a possible table $\boldsymbol{v}$ given the observed data $\boldsymbol{n}$. This can be done in constant time using Lemma 8.3, stated in Section 8.1 below. If no value of v_{10} leads to a possible table, set $j \leftarrow j + 1$ and go to (2).
   (c) Permutation test the count vector $\boldsymbol{v}$ coming from the choice of v_{10} in the previous step using Algorithm 2.2 with the given value of α. If $\boldsymbol{v}$ is accepted, return that τ_0 is compatible. If $\boldsymbol{v}$ is rejected and $v_{10} \geq 1$ or $v_{01} \geq 1$, set $j \leftarrow j + 1$ and go to (2).
   If $\boldsymbol{v}$ is rejected and $v_{10} = v_{01} = 0$, additionally permutation test the count vector given by $v_{10} = 1$. If it is possible and accepted, terminate the algorithm and declare τ_0 compatible. Otherwise, set $j \leftarrow j + 1$ and go to (2).
3. Return that τ_0 is incompatible.

Our simulations using Algorithm 5.2 show significant improvements over the previous algorithms of Li–Ding and Rigdon–Hudgen, even in small samples [RH15, LD16]. See Table 1 for examples.

5 Monte Carlo Intervals

In Section 4, we constructed the interval $\mathcal{I}_\alpha$ using exact permutation tests (Algorithm 2.2). In this section, we show how to construct confidence intervals using approximate Monte Carlo permutation tests, which are far less computationally demanding.

We begin by defining an approximate permutation test.

$\boldsymbol{n}$	Confidence Interval	Algorithm 4.3	[LD16]	[RH15]
(2,6,8,0)	$[-14,-5]$	24	113	189
(6,4,4,6)	$[-4,10]$	16	308	1225
(8,4,5,7)	$[3,13]$	26	421	2160

Table 1: 95% confidence intervals for $n\tau(\boldsymbol{n})$. The second column gives the confidence interval $\mathcal{I}(\boldsymbol{n})$, after scaling by n. The remaining columns indicate the number of permutation tests required to return $\mathcal{I}_\alpha$ for Algorithm 4.3 and the algorithms from [RH15, LD16]. Data for the last two columns is from [LD16, Table I].

Algorithm 5.1. Take as input $\alpha \in (0,1)$, $\varepsilon \in (0,\alpha)$, $K \in \mathbb{Z}_{>0}$, a potential outcome table $\boldsymbol{w}$, and a vector $\boldsymbol{n} \in \mathbb{Z}_{\geq 0}^{\{0,1\}^2}$. The algorithm will return a binary decision, either accept or reject.

Let $\boldsymbol{Z}^{(1)}, \ldots, \boldsymbol{Z}^{(K)}$ be a sequence of independent, identically distributed random vectors, with common distribution equal to that of $\boldsymbol{Z}$. Define random variables $V_1, \ldots, V_K$ by

$$V_i = \mathbb{1}\Big(\big|T(\boldsymbol{w}, \boldsymbol{Z}^{(i)}) - \tau(\boldsymbol{w})\big| \geq \big|T(\boldsymbol{n}) - \tau(\boldsymbol{w})\big|\Big), \qquad i \in [\![1, K]\!], \tag{5.1}$$

where $T(\boldsymbol{w}, \boldsymbol{Z}^{(i)})$ is to defined to be $T(\widetilde{\boldsymbol{Y}}, \boldsymbol{Z}^{(i)})$ and $\widetilde{\boldsymbol{Y}}$ is the observed data arising from $\boldsymbol{w}$ and $\boldsymbol{Z}$. Set

$$S = \frac{1}{K}\sum_{i=1}^{K} V_i. \tag{5.2}$$

Simulate a realization of S by simulating realizations of the K variables $\boldsymbol{Z}^{(1)}, \ldots, \boldsymbol{Z}^{(K)}$, and constructing the V_i and S accordingly. Return the decision accept if $S + \varepsilon \geq \alpha$, and return the decision reject otherwise.

Algorithm 5.2. This algorithm takes as input $\alpha \in (0,1)$, $\varepsilon \in (0,\alpha)$, and a vector $\boldsymbol{n} \in \mathbb{Z}_{\geq 0}^{\{0,1\}^2}$. It returns an interval $\mathcal{I}_{\alpha,\varepsilon,K}(\boldsymbol{n}) = \big[L_{\alpha,\varepsilon,K}(\boldsymbol{n}), U_{\alpha,\varepsilon,K}(\boldsymbol{n})\big]$. The algorithm is identical to Section 4, except every use of an exact permutation test (Algorithm 2.2) is replaced with an approximate permutation test with the given parameters (Algorithm 5.1).

Given observed data $(\boldsymbol{Y}, \boldsymbol{Z})$ we define the interval $\mathcal{I}_{\alpha,\varepsilon,K}(\boldsymbol{Y}, \boldsymbol{Z})$ to be the interval $\mathcal{I}_{\alpha,\varepsilon,K}(\boldsymbol{n})$ returned by Algorithm 5.2 for the observed count vector $\boldsymbol{n}$ corresponding to $(\boldsymbol{Y}, \boldsymbol{Z})$.

The first part of the next proposition shows that the intervals generated by Algorithm 5.2 are guaranteed to cover $\tau(\boldsymbol{y})$ with probability $1-\alpha$, if the parameters of the algorithm are chosen correctly. The second part shows that these intervals can be made to approximate the deterministic intervals $\mathcal{I}_\alpha$ arbitrarily well.

Proposition 5.3. *Fix a potential outcome table $\boldsymbol{y}$ and $\alpha \in (0,1)$. Suppose $n \geq 0$.*

1. *Fix $\varepsilon \in (0,\alpha)$. If*

$$K \geq \frac{1}{\varepsilon^2}\log\left(\frac{16n\log_2(n)}{\varepsilon}\right), \tag{5.3}$$

then

$$\mathbb{P}\big(\tau(\boldsymbol{y}) \in \mathcal{I}_{\alpha-\varepsilon,\varepsilon,K}(\boldsymbol{Y}, \boldsymbol{Z})\big) \geq 1-\alpha, \tag{5.4}$$

where the probability is with respect to the variable $\boldsymbol{Z}$.

2. Fix $\varepsilon \in (0, \alpha/3)$. We have

$$\mathbb{P}\big(\mathcal{I}_{\alpha-\varepsilon,\varepsilon,K}(\boldsymbol{Y}, \boldsymbol{Z}) \subset \mathcal{I}_{\alpha-3\varepsilon}\big) \geq 1 - 16\big(n \log_2(n)\big) \exp(-K\varepsilon^2). \tag{5.5}$$

Proposition 5.3 clarifies two important points. First, if n is held fixed, guaranteeing coverage at level $1 - \alpha$ for the ε-approximate interval $\mathcal{I}_{\alpha-\varepsilon,\varepsilon,K}$ requires a number of Monte Carlo samples K that scales asymptotically as $\varepsilon^{-2} \log(\varepsilon^{-1})$ as ε tends to zero. Second, if ε is fixed, the number of samples K should increase asymptotically as $\log(16n \log_2(n))$ as n grows to guarantee that the Monte Carlo intervals are contained in the exact interval $\mathcal{I}_{\alpha-3\varepsilon}$.

We close with two notes on the practical implementation of Algorithm 5.2. First, it is trivially parallelizable, for example by dividing the computation of the V_i in (5.1) across distinct processors. Second, for clarity we have stated Algorithm 5.2 using a fixed number of K samples for each potential outcome table $\boldsymbol{w}$ to be tested using Algorithm 5.1. A simple modification is to stop sampling variables V_i and reject $\boldsymbol{w}$ when the probability that $S + \varepsilon > \alpha$ becomes sufficiently small in Algorithm 5.1. In general, this will require fewer samples than the lower bound for K given in Proposition 5.3, which is essentially a worst-case bound for tables $\boldsymbol{w}$ such that $p(\boldsymbol{w}, \boldsymbol{Y}, \boldsymbol{Z}) \approx \alpha$. The validity of this adaptive algorithm can be justified by replaced the concentration inequalities used in the proof of Proposition 5.3 with more general ones that account for the correlations induced by the early stopping. Such inequalities are well-known in the bandit algorithm literature; see for example [LS20, Exercise 7.1]. For brevity, we do not give the details here.

6 Missing Data

We now consider a more general situation in which the realized outcomes Y_i may not be observed for some subjects, with no restriction on the distribution of the unobserved Y_i. More precisely, we suppose that instead of observing $(\boldsymbol{Y}, \boldsymbol{Z})$, the experimenter observes $(\boldsymbol{M}, \boldsymbol{Z})$, where

$$\boldsymbol{M} \in \{-1, 0, 1\}^n, \qquad M_i = -J_i + (1 - J_i)Y_i, \tag{6.1}$$

and $\boldsymbol{J} \in \{1, 0\}^n$. Here $\boldsymbol{J}$ is a random variable which may have an arbitrary distribution, and in particular may depend on $\boldsymbol{Y}$ and $\boldsymbol{Z}$. The indices i such that $J_i = 0$ denote realized outcomes that are observed, while those such that $J_i = 1$ indicate missing data. Hence, we have $M_i = -1$ if the realized outcome for the i-th subject is not observed.

We can construct exact confidence intervals for $\tau(\boldsymbol{y})$ given $(\boldsymbol{M}, \boldsymbol{Z})$ by considering, roughly speaking, the two extremal imputations of the missing data (leading to the smallest and largest estimates of $\tau(\boldsymbol{y})$).

Definition 6.1. Given a observed data $(\boldsymbol{M}, \boldsymbol{Z})$, we define the vector $\boldsymbol{Y}^{(+)} \in \{0, 1\}^n$ by

$$Y_i^{(+)} = 1 \text{ if } J_i = 1 \text{ and } Z_i = 1, \tag{6.2}$$

$$Y_i^{(+)} = 0 \text{ if } J_i = 1 \text{ and } Z_i = 0, \tag{6.3}$$

$$Y_i^{(+)} = M_i \text{ if } J_i = 0. \tag{6.4}$$

Similarly, we define $\boldsymbol{Y}^{(-)} \in \{0,1\}^n$ by

$$Y_i^{(-)} = 0 \text{ if } J_i = 1 \text{ and } Z_i = 1, \tag{6.5}$$

$$Y_i^{(-)} = 1 \text{ if } J_i = 1 \text{ and } Z_i = 0, \tag{6.6}$$

$$Y_i^{(-)} = M_i \text{ if } J_i = 0. \tag{6.7}$$

Finally, for $\alpha \in (0,1)$, we set

$$\mathcal{I}_\alpha^\circ(\boldsymbol{M}, \boldsymbol{Z}) = \big[L_\alpha(\boldsymbol{Y}^{(-)}, \boldsymbol{Z}), U_\alpha(\boldsymbol{Y}^{(+)}, \boldsymbol{Z})\big]. \tag{6.8}$$

The following proposition is proved in Section 10.

Proposition 6.2. *Fix potential outcomes $\boldsymbol{y}$ and $\alpha \in (0,1)$. Then*

$$\mathbb{P}\big(\tau(\boldsymbol{y}) \in \mathcal{I}_\alpha^\circ(\boldsymbol{M}, \boldsymbol{Z})\big) \geq 1 - \alpha, \tag{6.9}$$

where the probability is with respect to the variable $\boldsymbol{Z}$.

Remark 6.3. The intervals $\mathcal{I}_\alpha^\circ$ can be used to lift our hypothesis that n is even, in the following way. Given a group of $2m-1$ subjects, insert a fictitious subject to create a group of $2m$ subjects, randomize to equal groups, and construct $\mathcal{I}_\alpha^\circ$ by treating the data from the fictitious subject as missing. An argument similar to the one that proves Proposition 6.2 shows that this covers the sample average treatment effect for the original group of $2m-1$ subjects with probability at least $1-\alpha$. Such intervals will slightly sacrifice precision relative to the intervals $\mathcal{I}_\alpha$ applied to an unbalanced design with groups of m and $m-1$. (This loss of precision is straightforwardly quantified and decays with rate n^{-1}.) However, the only provably correct construction of $\mathcal{I}_\alpha$ in the unbalanced setting requires $O(n^4)$ permutation tests (as discussed in Section 1), while $\mathcal{I}_\alpha^\circ$ can be constructed in $O(n \log n)$ permutation tests by Theorem 4.1. Hence, there is a trade-off between precision and computational practicality, which may favor using the $\mathcal{I}_\alpha^\circ$ construction when n is large.

7 Proofs for Section 3

The following definition will be used in the proofs of our results. For $k \in \mathbb{Z}_{>0}$, we denote

$$k^{-1}\mathbb{Z} = \left\{\frac{j}{k} : j \in \mathbb{Z}\right\}, \quad k^{-1}(2\mathbb{Z}) = \left\{\frac{2j}{k} : j \in \mathbb{Z}\right\}, \quad k^{-1}(2\mathbb{Z}+1) = \left\{\frac{2j+1}{k} : j \in \mathbb{Z}\right\}. \tag{7.1}$$

Definition 7.1. Let $f\colon \mathbb{Z} \to \mathbb{R}_{\geq 0}$ be a probability mass function. We say that f is *symmetric about* k_0 for some $k_0 \in 2^{-1}\mathbb{Z}$ if

$$f(k_0 - x) = f(k_0 + x) \text{ for all } x \in 2^{-1}\mathbb{Z} \text{ such that } k_0 + x \in \mathbb{Z}. \tag{7.2}$$

If f is symmetric about k_0, we say that f is *decreasing away from* k_0 if

$$f(k_0 + y) \leq f(k_0 + x) \text{ for all } x, y \in 2^{-1}\mathbb{Z} \text{ such that } y > x \text{ and } k_0 + x, k_0 + y \in \mathbb{Z}. \tag{7.3}$$

We now suppose that $\sum_{k=-\infty}^{\infty} kf(k)$ converges absolutely, so that the mean of the distribution corresponding to f exists. In this case, we say that f is *symmetric–decreasing* (abbreviated SD) if f is symmetric about $k_0 = \sum_{k=-\infty}^{\infty} kf(k)$, and decreasing away from k_0.

Given $k \in \mathbb{Z}_{>0}$ and $k_0 \in 2^{-1}\mathbb{Z}$, we say that a probability mass function $f\colon k^{-1}\mathbb{Z} \to \mathbb{R}_{\geq 0}$ is symmetric about $k^{-1}k_0 \in (2k)^{-1}\mathbb{Z}$ if the function $g\colon \mathbb{Z} \to \mathbb{R}_{\geq 0}$ defined by $g(x) = f(kx)$ is symmetric about k_0. Similarly, we say that f is SD if $g(x) = f(kx)$ is.

7.1 Proof of Lemma 3.1

We begin with two preliminary lemmas.

Lemma 7.2. *For any potential outcome table $\boldsymbol{w}$, the pmf of $T(\boldsymbol{w}, \boldsymbol{Z})$ is symmetric about $\mathbb{E}\big[T(\boldsymbol{w}, \boldsymbol{Z})\big]$.*

Proof. Let $\boldsymbol{Z}^\dagger$ be the random vector defined by $Z_i^\dagger = 1 - Z_i$. Then for any randomization given by $\boldsymbol{Z}$, the vector $\boldsymbol{Z}^\dagger$ represents a randomization that exchanges the control and treatment groups. We have

$$\frac{1}{2}\Big(T(\boldsymbol{w}, \boldsymbol{Z}) + T(\boldsymbol{w}, \boldsymbol{Z}^\dagger)\Big) = \frac{1}{2m}\sum_{i=1}^{n}\big(w_j(1) - w_j(0)\big) = \mathbb{E}\big[T(\boldsymbol{w}, \boldsymbol{Z})\big]. \tag{7.4}$$

Let f be the pmf of $T(\boldsymbol{w}, \boldsymbol{Z})$, which is supported on $m^{-1}\mathbb{Z}$. Since $\boldsymbol{Z}$ and $\boldsymbol{Z}^\dagger$ have the same distribution, we deduce from (7.4) that

$$f\Big(\mathbb{E}\big[T(\boldsymbol{w}, \boldsymbol{Z})\big] - j\Big) = f\Big(\mathbb{E}\big[T(\boldsymbol{w}, \boldsymbol{Z})\big] + j\Big) \tag{7.5}$$

for all $j \in n^{-1}\mathbb{Z}$ such that $\mathbb{E}\big[T(\boldsymbol{w}, \boldsymbol{Z})\big] + j \in m^{-1}\mathbb{Z}$, which completes the proof. □

Lemma 7.3. *Let $n = 2m$ for some $m \in \mathbb{Z}_{>0}$, and fix observed data $\boldsymbol{Y}, \boldsymbol{Z}$. Then then there exists a potential outcome table $\boldsymbol{w}$ such that $\tau(\boldsymbol{w}) = T(\boldsymbol{Y}, \boldsymbol{Z})$ and $\boldsymbol{w}$ is possible given $\boldsymbol{Y}, \boldsymbol{Z}$.*

Proof. We have $T(\boldsymbol{Y}, \boldsymbol{Z}) = m^{-1}(n_{11} - n_{01})$. Let $\boldsymbol{w}$ be a potential outcome table with count vector $\boldsymbol{v} = (0, n_{11} + n_{00}, n_{10} + n_{01}, 0)$. Then using

$$n_{11} + n_{10} = n_{01} + n_{00} = m, \tag{7.6}$$

we compute

$$\tau(\boldsymbol{w}) = \frac{1}{n}(n_{11} + n_{00} - n_{10} - n_{01}) \tag{7.7}$$

$$= \frac{1}{n}\big(n_{11} + (m - n_{01}) - (m - n_{11}) - n_{01}\big) = \frac{2n_{11} - 2n_{01}}{n} = T(\boldsymbol{Y}, \boldsymbol{Z}). \tag{7.8}$$

This completes the proof. □

Proof of Lemma 3.1. Let $\boldsymbol{w}$ be the potential outcome table provided by Lemma 7.3, so that $\tau(\boldsymbol{w}) = T(\boldsymbol{Y}, \boldsymbol{Z})$ and $\mathcal{I}_\alpha(\boldsymbol{n}) = \mathcal{I}_\alpha(\boldsymbol{Y}, \boldsymbol{Z})$. Then the distribution of $T(\boldsymbol{w}, \boldsymbol{Z})$ is symmetric about $T(\boldsymbol{Y}, \boldsymbol{Z})$, so $T(\boldsymbol{Y}, \boldsymbol{Z})$ is the median of the distribution. Then by (2.12), we have $p(\boldsymbol{w}, \boldsymbol{Y}, \boldsymbol{Z}) = 1$, which shows that $\tau(\boldsymbol{w})$ will always be accepted in Algorithm 2.3. This completes the proof. □

7.2 Proof of Proposition 3.2

Proof of Proposition 3.2. The event $\big\{\tau(\boldsymbol{y}) \notin \mathcal{I}_\alpha(\boldsymbol{y}, \boldsymbol{Z})\big)\big\}$ is contained in the event where Algorithm 2.3 rejects $\tau(\boldsymbol{y})$, which is in turn contained in the event where

$$\mathbb{P}\left(\big|T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}}) - \tau(\boldsymbol{y})\big| \geq \big|T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau(\boldsymbol{y})\big|\right) < \alpha. \tag{7.9}$$

Here probability is with respect to $\widetilde{\boldsymbol{Z}}$, and we consider the left side of (7.9) as a function of $T(\boldsymbol{Y}, \boldsymbol{Z})$. Let x denote the smallest real number such that

$$\mathbb{P}\left(\left|T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}}) - \tau(\boldsymbol{y})\right| \geq x\right) < \alpha, \tag{7.10}$$

which exists because $T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}})$ is a discrete random variable. Then the event where (7.9) occurs is the same as the event where $\left|T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau(\boldsymbol{y})\right| \geq x$. By the definition of x, and the fact that $(\boldsymbol{Y}, \boldsymbol{Z})$ and $(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}})$ are identically distributed, this event has probability at most α. This completes the proof. □

7.3 Proof of Proposition 3.3

Proof of Proposition 3.3. Let $\mathcal{P}$ be the set of all sequences of pairs $\big((q_i, r_i)\big)_{i=1}^m$ such that $q_i, r_i \in [\![1, n]\!]$ for all $i \in [\![1, m]\!]$ and

$$\{q_1, q_2, \ldots, q_m, r_1, r_2, \ldots r_m\} = \{1, 2, \ldots, n\}. \tag{7.11}$$

Observe that every element of $[\![1, n]\!]$ appears exactly once in such a sequence as some q_i or r_i.

We now perform complete randomization on the group of n subjects in the following way. First, let $P = \big((Q_i, R_i)\big)_{i=1}^m$ be a random variable uniformly distributed on the set $\mathcal{P}$. Hence, for all $i \in [\![1, m]\!]$, Q_i and R_i are random variables taking values in $[\![1, n]\!]$. We define a random vector $\widetilde{\boldsymbol{Z}} \in \{0, 1\}^n$ as follows. For each $i \in [\![1, N]\!]$, set

$$(\widetilde{Z}_{Q_i}, \widetilde{Z}_{R_i}) = (1, 0) \text{ if } B_i = 1, \qquad (\widetilde{Z}_{Q_i}, \widetilde{Z}_{R_i}) = (0, 1) \text{ if } B_i = 0. \tag{7.12}$$

One straightforwardly checks that the resulting $\widetilde{\boldsymbol{Z}}$ has a uniform distribution on randomizations of the n subjects to two equal groups. That is, $\widetilde{\boldsymbol{Z}}$ has the same distribution as $\boldsymbol{Z}$.

Next, let $\boldsymbol{w}$ by any potential outcome table. We will condition on P and derive a concentration inequality for the resulting conditional distribution of $T(\boldsymbol{w}, \widetilde{\boldsymbol{Z}})$. First note that the conditional mean of $T(\boldsymbol{w}, \widetilde{\boldsymbol{Z}})$ is

$$\mathbb{E}\big[T(\boldsymbol{w}, \widetilde{\boldsymbol{Z}})\big|P\big] = \tau(\boldsymbol{w}).$$

Next, observe that the variables

$$A_i = Z_{Q_i} y_{Q_i}(1) - (1 - Z_{Q_i}) y_{Q_i}(0) + Z_{R_i} y_{R_i}(1) - (1 - Z_{R_i}) y_{R_i}(0) \tag{7.13}$$

defined for $i \in [\![1, m]\!]$ are conditionally independent, after conditioning on P, and the conditional distribution of $T(\boldsymbol{w}, \widetilde{\boldsymbol{Z}})$ is given by

$$\frac{1}{m} \sum_{i=1}^{m} A_i. \tag{7.14}$$

Observe that we have $A_i \in \{-1, 0, 1\}$.

Now recall that Hoeffding's inequality states that given independent random variables $\{X_i\}_{i=1}^n$ such that $a \leq X \leq b$, we have

$$\mathbb{P}\Big(\big|S_n - \mathbb{E}[S_n]\big| \geq t\Big) \leq 2 \exp\left(-\frac{2t^2}{n(b-a)^2}\right) \tag{7.15}$$

for all $t > 0$, where

$$S_n = X_1 + \cdots + X_n. \tag{7.16}$$

We apply Hoeffding's inequality to the sequence $\{m^{-1}A_i\}_{i=1}^m$ with $a = -m^{-1}$ and $b = m^{-1}$, after conditioning on P, yielding

$$\mathbb{P}\Big(\big|T(\boldsymbol{w}, \widetilde{\boldsymbol{Z}}) - \tau(\boldsymbol{w})\big| \geq t \Big| P\Big) \leq 2\exp\left(-\frac{t^2 n}{8}\right). \tag{7.17}$$

Because this inequality is true conditional on every realization of $\mathcal{P}$, we have

$$\mathbb{P}\Big(\big|T(\boldsymbol{w}, \widetilde{\boldsymbol{Z}}) - \tau(\boldsymbol{w})\big| \geq t\Big) \leq 2\exp\left(-\frac{t^2 n}{8}\right). \tag{7.18}$$

Inserting $t = \big|T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau(\boldsymbol{w})\big|$, we get

$$\mathbb{P}\Big(\big|T(\boldsymbol{w}, \widetilde{\boldsymbol{Z}}) - \tau(\boldsymbol{w})\big| \geq \big|T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau(\boldsymbol{w})\big|\Big) \leq 2\exp\left(-\frac{\big|T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau(\boldsymbol{w})\big|^2 n}{8}\right). \tag{7.19}$$

To reject $\boldsymbol{w}$, we must have

$$\mathbb{P}\Big(\big|T(\boldsymbol{w}, \widetilde{\boldsymbol{Z}}) - \tau(\boldsymbol{w})\big| \geq \big|T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau(\boldsymbol{w})\big|\Big) < \alpha. \tag{7.20}$$

We introduce the shorthand

$$\delta = \big|T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau(\boldsymbol{w})\big|. \tag{7.21}$$

Hence, we solve for

$$2\exp\left(-\frac{\delta^2 n}{8}\right) < \alpha, \tag{7.22}$$

giving

$$\delta \geq \sqrt{\frac{8\log(2/\alpha)}{n}}. \tag{7.23}$$

Therefore the permutation test will reject all values of τ_0 with distance at least

$$\sqrt{\frac{8\log(2/\alpha)}{n}} \tag{7.24}$$

from $T(\boldsymbol{Y}, \boldsymbol{Z})$. This implies the conclusion. □

8 Proofs for Section 4

8.1 Preliminary Results

For completeness, and to facilitate the proofs of our finite-sample bounds, we begin by specifying the version of binary search that we use.

Algorithm 8.1. This *binary search algorithm* takes as input $k_1, k_2 \in \mathbb{Z}_{>0}$ such that $k_2 > k_1$, and a function $f\colon [\![k_1, k_2]\!] \to \{0, 1\}$ such that $f(x) = 0$ if $x \leq r$ and $f(x) = 1$ if $x > r$, where $r \in [\![k_1 - 1, k_2]\!]$ is unknown. The algorithm returns r through the following steps.

1. Initialize $a \leftarrow k_1$ and $b \leftarrow k_2$.
2. Set $c \leftarrow \lfloor (a + b)/2 \rfloor$ and evaluate $f(c)$. If $f(c) = 0$, set $a \leftarrow c$. If $f(c) = 1$, set $b \leftarrow c$.

3. Repeat the previous step until $b = a + 1$. In this case, return a if $a > k_1$ and $b < k_2$. If $a = k_1$, evaluate $f(k_1)$ and return $k_1 - 1$ if $f(k_1) = 1$, and k_1 if $f(k_1) = 0$. If $b = k_2$, evaluate $f(k_2)$, and return k_2 if $f(k_2) = 0$ and $k_2 - 1$ if $f(k_2) = 1$.

The following fact is well known.

Lemma 8.2. *Algorithm 8.1 terminates in at most $\lfloor \log_2(k_2 - k_1 + 1) + 2 \rfloor$ evaluations of f.*

Proof. Without loss of generality, we may suppose that $k_1 = 1$. Further, by setting $f(x) = 1$ for $x > k_2$, we may consider f as a function on $[\![1, 2^d + 1]\!]$, where $d = \lfloor \log_2(k_2) + 1 \rfloor$ satisfies $2^d \geq k_2$. Then the number of integers in the interval $[a, b]$ after each evaluation of f follows the sequence $2^{d-1} + 1, \ldots, 2^1 + 1, 2$ as the algorithm progresses, with a possible additional evaluation of f if $a = 1$ or $b = 2^d + 1$ at the final step. Then at at worst $d + 1$ evaluations are needed in total, which proves the theorem. □

The following lemma determines the possible potential outcome tables, given some observed data, in a certain one-parameter family of tables. We remark that it shows Step (2b) of Algorithm 4.3 can be completed in constant time.

Lemma 8.3. *Fix $j, v_{10} \in \mathbb{Z}_{\geq 0}$ and $\tau_0 \in \mathcal{C}$, and consider the potential outcome vector*

$$\boldsymbol{v} = (j - v_{10}, v_{10}, v_{10} - n\tau_0, n - j - v_{10} + n\tau_0). \tag{8.1}$$

Given observed data $\boldsymbol{n}$, a necessary condition for $\boldsymbol{v}$ to be possible is that all of the following inequalities hold:

$$j \geq n\tau_0 + n_{01}, \quad j \geq n_{11}, \quad n \geq j + n_{10}, \quad n_{11} + n\tau_0 + n_{10} + n_{01} \geq j. \tag{8.2}$$

In this case, the possible values of v_{10} are given by the interval

$$\max(0, n\tau_0, j - n_{11} - n_{01}, n_{11} + n_{01} + n\tau_0 - j) \leq v_{10} \leq \min(j, n_{11} + n_{00}, n_{10} + n_{01} + n\tau_0, n + n\tau_0 - j), \tag{8.3}$$

which may be empty.

Proof of Lemma 8.3. We apply the criterion of Lemma 2.4. The maximum in that statement becomes

$$\max(0, n_{11} - v_{10}, j - v_{10} - n_{01}, j - n\tau_0 - n_{10} - n_{01}). \tag{8.4}$$

Similarly, the minimum becomes

$$\min(j - v_{10}, n_{11}, j - n\tau_0 - n_{01}, n - v_{10} - n_{10} - n_{01}). \tag{8.5}$$

Each argument of the maximum function must be less than each argument of the minimum function. We test the arguments of the maximum from left to right. Starting with 0, we must have

$$j \geq v_{10}, \quad n_{11} \geq 0, \quad j \geq n\tau_0 + n_{01}, \quad n - n_{01} - n_{10} \geq v_{10}, \tag{8.6}$$

where we note that the condition $n_{11} \geq 0$ is always true. Considering $n_{11} - v_{10}$, we get

$$j \geq n_{11}, \quad v_{10} \geq 0, \quad v_{10} \geq n_{11} + n_{01} + n\tau_0 - j, n \geq n_{11} + n_{10} + n_{01}. \tag{8.7}$$

Note that the last condition is always satisfied. Considering $j - v_{10} - n_{01}$, we get

$$n_{10} \geq 0, \quad v_{10} \geq j - n_{11} - n_{01}, \quad v_{10} \geq n\tau_0, \quad n \geq j + n_{10}, \tag{8.8}$$

and the first condition is always true. Finally, considering $j - n\tau_0 - n_{10} - n_{01}$,

$$n\tau_0 + n_{10} + n_{01} \geq v_{10}, \quad n_{11} + n\tau_0 + n_{10} + n_{01} \geq j, \quad n_{10}, \quad n + n\tau_0 - j \geq v_{10}. \tag{8.9}$$

Collecting the previous four displays completes the proof. □

8.2 Proof of Theorem 4.1

We now prove Theorem 4.1 assuming Lemma 4.2, which is proved below.

Proof of Theorem 4.1. We will show that Algorithm 4.3 returns $\mathcal{I}_\alpha$ in the required number of permutation tests. It suffices to show that Algorithm 4.3 returns $U_\alpha^{(1)}(\boldsymbol{n}) = U_\alpha(\boldsymbol{n})$ for the upper bound in $2(n+1)\lfloor \log_2(n+1) + 2 \rfloor$ permutation tests, since the same estimate holds for the lower bound L_α by analogous reasoning. Then the conclusion follows after using

$$8n \log_2(n) \geq 4(n+1)\lfloor \log_2(n+1) + 2 \rfloor \tag{8.10}$$

for $n \geq 10$.

Let $C_+ = \max\big(\mathcal{C}(\boldsymbol{n})\big)$, where we recall that $\mathcal{C}(\boldsymbol{n})$ was defined in (2.10). By Lemma 3.1, we have $U_\alpha(\boldsymbol{n}) \in \big[T(\boldsymbol{n}), C_+\big]$. Recall that by Lemma 2.5, the set of elements

$$\mathcal{J}_\alpha(\boldsymbol{n}) = \{\tau_0 \in \mathcal{C}(\boldsymbol{n}) : \text{there exists a table } \boldsymbol{u} \text{ such that } \tau(\boldsymbol{u}) = \tau_0 \text{ and } p(\boldsymbol{n}, \boldsymbol{u}) \geq \alpha\} \tag{8.11}$$

is an interval equal to $\mathcal{I}_\alpha(\boldsymbol{n})$, so that (by Lemma 3.1)

$$\big[T(\boldsymbol{n}), C_+\big] \cap J_\alpha(\boldsymbol{n}) = \big[T(\boldsymbol{n}), U_\alpha(\boldsymbol{n})\big]. \tag{8.12}$$

Then it suffices to show that the f given in Algorithm 4.3 satisfies $f(x) = 0$ if $x/n \in \mathcal{J}_\alpha$ and $f(x) = 1$ otherwise, and that $f(x)$ can be computed in at most $2(n+1)$ permutation tests for every x. In this case, the binary search indicated in Algorithm 4.3 will return $U_\alpha(\boldsymbol{n})$ in a total $2(n+1)\lfloor \log_2(n+1) + 2 \rfloor$ permutation tests.

We begin with the claim that $f(x) = 0$ if $x/n \in \mathcal{J}_\alpha$ and $f(x) = 1$ otherwise. Set $\tau_0 = x/n$. If a vector $\boldsymbol{v}$ permutation tested in Step (2c) satisfies $p(\boldsymbol{n}, \boldsymbol{v}) \geq \alpha$, then $\tau_0 = \tau(\boldsymbol{v})$ lies in $\mathcal{J}_\alpha(\boldsymbol{n})$, and f correctly declares τ_0 compatible.

Otherwise, if $j = n+1$ and the last step declares τ_0 incompatible, we must show that every possible potential outcome count vector $\boldsymbol{u}$ with $\tau(\boldsymbol{u}) = \tau_0$ satisfies $p(\boldsymbol{n}, \boldsymbol{u}) < \alpha$. To this end, fix an arbitrary $\boldsymbol{u}$ such that $\tau(\boldsymbol{u}) = \tau_0$. If $\boldsymbol{u}$ was tested for compatibility in Step (2c), then we must have $p(\boldsymbol{n}, \boldsymbol{u}) < \alpha$, since the permutation test of $\boldsymbol{u}$ must have rejected for f to have declared τ_0 incompatible.

Therefore, we may suppose $\boldsymbol{u}$ was not tested, let $j = u_{11} + u_{01}$, and let v_{01} be the value determined in Step (2b). Let $\boldsymbol{v}$ be the vector determined by v_{01} through the equalities in (2a). If τ_0 was tested declared incompatible by f, then $\boldsymbol{v}$ was permutation tested and rejected. Hence $\boldsymbol{u} \neq \boldsymbol{v}$ and $p(\boldsymbol{n}, \boldsymbol{v}) < \alpha$. Suppose $v_{01} + v_{10} \geq 1$. Using this inequality and Lemma 4.2, we deduce that

$$p(\boldsymbol{n}, \boldsymbol{u}) \leq p(\boldsymbol{n}, \boldsymbol{v}) < \alpha, \tag{8.13}$$

since $\boldsymbol{v} = \boldsymbol{u} + k(1, -1, -1, 1)$ for some $k \in \mathbb{Z}_{>0}$, as desired. Now suppose that $v_{10} = v_{01} = 0$. Let $\boldsymbol{v}'$ be the vector constructed from $v'_{10} = 1$ using (4.4) and $j = u_{11} + u_{01}$. We claim that $\boldsymbol{v}'$ is possible, and hence tested and rejected. Indeed, if $\boldsymbol{v}'$ is not possible given the observed data, then

by Lemma 8.3 we find that $\boldsymbol{u} \neq \boldsymbol{v}$ is not possible, since since the set of possible coordinates in (8.3) is an interval. Then $p(\boldsymbol{v}', \boldsymbol{n}) < \alpha$, and as before we have $p(\boldsymbol{n}, \boldsymbol{u}) \leq p(\boldsymbol{n}, \boldsymbol{v}') < \alpha$.

Hence, $\boldsymbol{u}$ is rejected if $\boldsymbol{v}$ is, and we have shown that evaluating f provides a valid method for checking the compatibility of τ_0. Finally, the claim that f can be evaluated using $2(n+1)$ permutation tests follows from the fact that examining a given $j \in [\![0, n]\!]$ requires at most two permutation tests, as described in Step (2c). This completes the proof. □

8.3 Proof of Lemma 4.2

We first provide some intuition for the proof of Lemma 4.2. We will see that the sampling distribution for the Neyman estimator $T(\boldsymbol{w})$ of a generic table $\boldsymbol{w}$ is SD on $m^{-1}\mathbb{Z}$. If all tables $\boldsymbol{w}$ had this property, the proof of Lemma 4.2 would be fairly straightforward (by applying Lemma 8.4 below). However, there are exceptional tables that do not, due to parity issues. Whenever we have a table of the form $\boldsymbol{v} = (a, 0, 0, b)$ with $a + b = n = 2m$, the estimator $T(\boldsymbol{w}, \boldsymbol{Z})$ is supported on $m^{-1}(2\mathbb{Z})$ or $m^{-1}(2\mathbb{Z}+1)$ instead of $m^{-1}\mathbb{Z}$. Most of our effort goes into showing that this is the only such bad case.

Lemma 8.4. *Fix observed data $\boldsymbol{n}$, and a count vector*

$$\boldsymbol{v} = (v_{11}, v_{10}, v_{01}, v_{00}). \tag{8.14}$$

Suppose $\min(v_{10}, v_{01}) \geq 1$*, and define*

$$\boldsymbol{v}' = (v_{11}+1, v_{10}-1, v_{01}-1, v_{00}+1). \tag{8.15}$$

Set

$$\boldsymbol{v}^\circ = (v_{11}, v_{10}-1, v_{01}-1, v_{00}), \tag{8.16}$$

and suppose that the pmf of $T(\boldsymbol{v}^\circ, \boldsymbol{Z})$ is SD on the lattice $(m-1)^{-1}\mathbb{Z}$, where here $\boldsymbol{Z}$ is uniformly distributed over $\mathcal{Z}(n-2)$. Then

$$p(\boldsymbol{n}, \boldsymbol{v}') \geq p(\boldsymbol{n}, \boldsymbol{v}). \tag{8.17}$$

Proof. We couple the distributions of $T(\boldsymbol{v}, \boldsymbol{Z})$ and $T(\boldsymbol{v}', \boldsymbol{Z})$ in the following way. Let $\boldsymbol{w}$ be a potential outcome table with count vector $\boldsymbol{v}$, and label the subjects so that $\boldsymbol{w}_1 = (1, 0)$ and $\boldsymbol{w}_2 = (0, 1)$. Let $\boldsymbol{w}'$ be a potential outcome table such that $\boldsymbol{w}'_1 = (1, 1)$ and $\boldsymbol{w}'_2 = (0, 0)$, and $\boldsymbol{w}'_i = \boldsymbol{w}_i$ for $i \geq 3$. Then $\boldsymbol{w}'$ has the count vector $\boldsymbol{v}'$. Let $\widetilde{\boldsymbol{Y}}$ be the observed outcome vector for $\boldsymbol{w}$, let $\widetilde{\boldsymbol{Y}}'$ be the observed outcome vector for $\boldsymbol{w}'$, and recall from (2.12) that

$$p(\boldsymbol{w}, \boldsymbol{Y}, \boldsymbol{Z}) = \mathbb{P}\left(\left|T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}}) - \tau_0\right| \geq \left|T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau_0\right|\right). \tag{8.18}$$

by definition.

We may suppose $T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau_0 \neq 0$, otherwise $p(\boldsymbol{n}, \boldsymbol{v}') = p(\boldsymbol{n}, \boldsymbol{v}) = 1$ and the claim (8.17) is trivially true. Then (8.18) implies

$$p(\boldsymbol{w}, \boldsymbol{Y}, \boldsymbol{Z}) = \mathbb{P}\left(T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}}) - \tau_0 \geq \left|T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau_0\right|\right) \tag{8.19}$$

$$+ \mathbb{P}\left(T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}}) - \tau_0 \leq -\left|T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau_0\right|\right). \tag{8.20}$$

To complete the proof, it suffices to show that each of the probabilities (8.19) and (8.20) increases

when $\boldsymbol{w}$ is replaced $\boldsymbol{w}'$. Observe that

$$\mathbb{E}\big[T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}})\big] = \tau(\boldsymbol{w}) = \tau_0, \tag{8.21}$$

and that $T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}})$ is symmetric about its mean, by Lemma 7.2. Then by symmetry, it suffices to show that (8.19) increases when $\boldsymbol{w}$ is replaced with $\boldsymbol{w}'$.

As a preliminary observation, note that $\tau_0 \in \mathcal{C}$ takes on values in the lattice $n^{-1}\mathbb{Z}$, while T takes on values in $m^{-1}\mathbb{Z} = 2n^{-1}\mathbb{Z}$. However, we are considering the probability

$$\mathbb{P}\left(T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}}) \geq \big|T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau_0\big| + \tau_0\right), \tag{8.22}$$

from (8.19), and one checks that $\big|T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau_0\big| + \tau_0$ is an element of $m^{-1}\mathbb{Z}$ regardless of the value of τ_0, since $T \in m^{-1}\mathbb{Z}$.

Now, given $T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}})$ for some fixed realization of $\widetilde{\boldsymbol{Z}}$, we consider what happens to its value as $\boldsymbol{w}$ changes to $\boldsymbol{w}'$, so that $\widetilde{\boldsymbol{Y}}$ is replaced by $\widetilde{\boldsymbol{Y}}'$. There are three cases.

1. $\widetilde{Z}_1 = \widetilde{Z}_2$. Then the first two subjects are in the same group. The change $(1,0) \mapsto (1,1)$ and $(0,1) \mapsto (0,0)$ leaves T invariant in this case.
2. $\widetilde{Z}_1 = 1$ and $\widetilde{Z}_2 = 0$. The change of $(1,0) \mapsto (1,1)$ for the first subject does not affect T. The change $(0,1) \mapsto (0,0)$ for the second subject increases T by m^{-1}.
3. $\widetilde{Z}_1 = 0$ and $\widetilde{Z}_2 = 1$. The change of $(1,0) \mapsto (1,1)$ for the first subject decreases T by m^{-1}. The change $(0,1) \mapsto (0,0)$ for the second subject does not affect T.

In the conclusion, the value of T changes by 0 or $\pm m^{-1}$. Then for $a \in \mathbb{Z}$, we have

$$\begin{aligned}
\mathbb{P}\left(T(\widetilde{\boldsymbol{Y}}', \widetilde{\boldsymbol{Z}}) = \frac{a}{m}\right) &= \mathbb{P}\left(T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}}) = \frac{a-1}{m} \text{ and } (\widetilde{Z}_1, \widetilde{Z}_2) = (1,0)\right) && (8.23)\\
&+ \mathbb{P}\left(T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}}) = \frac{a}{m} \text{ and } \widetilde{Z}_1 = \widetilde{Z}_2\right) && (8.24)\\
&+ \mathbb{P}\left(T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}}) = \frac{a+1}{m} \text{ and } (\widetilde{Z}_1, \widetilde{Z}_2) = (0,1)\right). && (8.25)
\end{aligned}$$

As noted below (8.20), it suffices to show that

$$\mathbb{P}\left(T(\widetilde{\boldsymbol{Y}}', \widetilde{\boldsymbol{Z}}) \geq \big|T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau_0\big| + \tau_0\right) \geq \mathbb{P}\left(T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}}) \geq \big|T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau_0\big| + \tau_0\right), \tag{8.26}$$

since this implies that (8.19) increases if $\boldsymbol{v}$ is replaced by $\boldsymbol{v}'$. Summing (8.23), we get

$$\begin{aligned}
&\mathbb{P}\left(T(\widetilde{\boldsymbol{Y}}', \widetilde{\boldsymbol{Z}}) \geq \big|T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau_0\big| + \tau_0\right) && (8.27)\\
&= \mathbb{P}\left(T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}}) \geq \big|T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau_0\big| + \tau_0 + m^{-1}\right) && (8.28)\\
&+ \mathbb{P}\left(T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}}) = \big|T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau_0\big| + \tau_0 \text{ and } \widetilde{Z}_1 = \widetilde{Z}_2\right) && (8.29)\\
&+ \mathbb{P}\left(T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}}) = \big|T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau_0\big| + \tau_0 \text{ and } (\widetilde{Z}_1, \widetilde{Z}_2) = (1,0)\right) && (8.30)\\
&+ \mathbb{P}\left(T(\widetilde{\boldsymbol{Y}}, \widetilde{\boldsymbol{Z}}) = \big|T(\boldsymbol{Y}, \boldsymbol{Z}) - \tau_0\big| + \tau_0 - m^{-1} \text{ and } (\widetilde{Z}_1, \widetilde{Z}_2) = (1,0)\right). && (8.31)
\end{aligned}$$

Using the previous equality, we see that (8.26) is equivalent to the statement that

$$\mathbb{P}\left(T(\widetilde{\boldsymbol{Y}},\widetilde{\boldsymbol{Z}}) = \big|T(\boldsymbol{Y},\boldsymbol{Z})-\tau_0\big|+\tau_0 \text{ and } \widetilde{Z}_1=\widetilde{Z}_2\right) \tag{8.32}$$

$$+\,\mathbb{P}\left(T(\widetilde{\boldsymbol{Y}},\widetilde{\boldsymbol{Z}}) = \big|T(\boldsymbol{Y},\boldsymbol{Z})-\tau_0\big|+\tau_0 \text{ and } (\widetilde{Z}_1,\widetilde{Z}_2)=(1,0)\right) \tag{8.33}$$

$$+\,\mathbb{P}\left(T(\widetilde{\boldsymbol{Y}},\widetilde{\boldsymbol{Z}}) = \big|T(\boldsymbol{Y},\boldsymbol{Z})-\tau_0\big|+\tau_0-m^{-1} \text{ and } (\widetilde{Z}_1,\widetilde{Z}_2)=(1,0)\right) \tag{8.34}$$

$$\geq \mathbb{P}\left(T(\widetilde{\boldsymbol{Y}},\widetilde{\boldsymbol{Z}}) = \big|T(\boldsymbol{Y},\boldsymbol{Z})-\tau_0\big|+\tau_0\right), \tag{8.35}$$

which rearranges to

$$\mathbb{P}\left(T(\widetilde{\boldsymbol{Y}},\widetilde{\boldsymbol{Z}}) = \big|T(\boldsymbol{Y},\boldsymbol{Z})-\tau_0\big|-m^{-1}+\tau_0 \text{ and } (\widetilde{Z}_1,\widetilde{Z}_2)=(1,0)\right) \tag{8.36}$$

$$\geq \mathbb{P}\left(T(\widetilde{\boldsymbol{Y}},\widetilde{\boldsymbol{Z}}) = \big|T(\boldsymbol{Y},\boldsymbol{Z})-\tau_0\big|+\tau_0 \text{ and } (\widetilde{Z}_1,\widetilde{Z}_2)=(0,1)\right). \tag{8.37}$$

Let $\widetilde{\boldsymbol{Z}}^\circ = (\widetilde{Z}_3^\circ,\dots\widetilde{Z}_n^\circ)$ be uniformly distributed on $\mathcal{Z}(n-2)$ (recall (2.4)), let $\boldsymbol{w}^\circ=(\boldsymbol{w}_3,\dots,\boldsymbol{w}_n)$, and let $\widetilde{\boldsymbol{Y}}^\circ$ be the vector of observed outcomes corresponding to $\widetilde{\boldsymbol{Z}}^\circ$ and $\boldsymbol{w}^\circ$. Observe that $\boldsymbol{v}^\circ$ is the count vector for $\boldsymbol{w}^\circ$. Then (8.36) may be rewritten as

$$\mathbb{P}\left(T(\widetilde{\boldsymbol{Y}},\widetilde{\boldsymbol{Z}}) = \big|T(\boldsymbol{Y},\boldsymbol{Z})-\tau_0\big|-m^{-1}+\tau_0 \;\Big|\; (\widetilde{Z}_1,\widetilde{Z}_2)=(1,0)\right)\mathbb{P}\big((\widetilde{Z}_1,\widetilde{Z}_2)=(1,0)\big) \tag{8.38}$$

$$\geq \mathbb{P}\left(T(\widetilde{\boldsymbol{Y}},\widetilde{\boldsymbol{Z}}) = \big|T(\boldsymbol{Y},\boldsymbol{Z})-\tau_0\big|+\tau_0 \;\Big|\; (\widetilde{Z}_1,\widetilde{Z}_2)=(0,1)\right)\mathbb{P}\big((\widetilde{Z}_1,\widetilde{Z}_2)=(0,1)\big), \tag{8.39}$$

and further as

$$\mathbb{P}\left(T(\widetilde{\boldsymbol{Y}}^\circ,\widetilde{\boldsymbol{Z}}^\circ)=A-(m-1)^{-1}\right)\geq \mathbb{P}\left(T(\widetilde{\boldsymbol{Y}}^\circ,\widetilde{\boldsymbol{Z}}^\circ)=A\right), \tag{8.40}$$

where

$$A=\frac{m}{m-1}\cdot\Big(\big|T(\boldsymbol{Y},\boldsymbol{Z})-\tau_0\big|+\tau_0\Big)\in(m-1)^{-1}\mathbb{Z}. \tag{8.41}$$

In the previous computations, we used the fact that the net contribution to $T(\widetilde{\boldsymbol{Y}},\widetilde{\boldsymbol{Z}})$ from the first two subjects is zero, since we assumed that $\boldsymbol{w}_1=(1,0)$ and $\boldsymbol{w}_2=(0,1)$. We also used that the conditional distribution of $(\widetilde{Z}_3,\dots\widetilde{Z}_n)$, conditional on either $(\widetilde{Z}_1,\widetilde{Z}_2)=(1,0)$ or $(\widetilde{Z}_1,\widetilde{Z}_2)=(0,1)$, is uniform on the set $\mathcal{Z}(n-2)$, and that the probabilities of these two events are equal. Observe that

$$A>\frac{m}{m-1}\cdot\tau_0=\mathbb{E}\Big[T\big(\widetilde{\boldsymbol{Y}}^\circ,\widetilde{\boldsymbol{Z}}^\circ\big)\Big], \tag{8.42}$$

since we assumed earlier that $T(\boldsymbol{Y},\boldsymbol{Z})-\tau_0\neq 0$. By hypothesis, the pmf of $T\big(\widetilde{\boldsymbol{Y}}^\circ,\widetilde{\boldsymbol{Z}}^\circ\big)$ is SD, since it corresponds to the count vector $\boldsymbol{v}^\circ$. This fact, and the fact that A is strictly greater than the mean by (8.42), together imply that (8.40) holds. Since (8.40) is equivalent to (8.26), we see that (8.26) holds, which completes the proof. □

We next prove some lemmas to set up an induction.

Lemma 8.5. *Let X_1 be SD on $m^{-1}\mathbb{Z}$, and let X_2 be uniformly distributed on $\{0,m^{-1}\}$. Then X_1+X_2 is SD on $m^{-1}\mathbb{Z}$.*

Proof. It suffices to show that $m(X_1+X_2)$ is SD on $\mathbb{Z}$. Denote the pdf of mX_1 by f_1. Then the pdf g of $m(X_1+X_2)$ is given by $g(x)=\frac{1}{2}\big(f_1(x)+f_1(x-1)\big)$. Suppose that mX_1 is symmetric

about $a \in \mathbb{Z}$. Then it is clear that $g(x)$ is symmetric about $a + \frac{1}{2}$. To show that $g(x)$ is decreasing, it suffices to show that for any $b \geq a + 1/2$, we have $g(b) \geq g(b+1)$. We find

$$g(b) = \frac{1}{2}\big(f_1(b) + f_1(b-1)\big) \geq \frac{1}{2}\big(f_1(b+1) + f_1(b)\big) = g(b+1), \tag{8.43}$$

as desired. In the previous equation, we used $f(b) \geq f(b+1)$, which is true since f is SD, and $f_1(b-1) \geq f(b)$. The latter is true by the SD property of f_1 if $b \geq a+1$. If $b = a + 1/2$, then this follows from the symmetric of f about a. □

Lemma 8.6. *Fix a potential outcome count vector*

$$\boldsymbol{v} = (v_{11}, v_{10}, v_{01}, v_{00}), \tag{8.44}$$

and suppose at least one of the following conditions holds.

1. *We have $v_{10} \geq 2$, and the pmf for $T(\boldsymbol{v} - \boldsymbol{a}, \boldsymbol{Z}')$ is SD on $(m-1)^{-1}\mathbb{Z}$ for every choice of*

$$\boldsymbol{a} \in \big\{(1,1,0,0), (0,2,0,0), (0,1,1,0), (0,1,0,1)\big\} \tag{8.45}$$

   such that $\boldsymbol{v} - \boldsymbol{a}$ has nonnegative entries. Here, $\boldsymbol{Z}'$ denotes a random variable uniformly distributed on $\mathcal{Z}(n-2)$.

2. *We have $v_{01} \geq 2$, and the pmf for $T(\boldsymbol{v} - \boldsymbol{a}, \boldsymbol{Z}')$ is SD on $(m-1)^{-1}\mathbb{Z}$ for every choice of*

$$\boldsymbol{a} \in \big\{(1,0,1,0), (0,1,1,0), (0,0,2,0), (0,0,1,1)\big\} \tag{8.46}$$

   such that $\boldsymbol{v} - \boldsymbol{a}$ has nonnegative entries.

Then the pmf of $T(\boldsymbol{v}, \boldsymbol{Z})$ is SD on $m^{-1}\mathbb{Z}$.

Proof. We prove the claim only for the case $v_{10} \geq 2$, since the proof in the case that $v_{01} \geq 2$ is similar. By relabeling, we may suppose that the first subject has the potential outcome vector $\boldsymbol{w}_1 = (1,0)$. We sample $\boldsymbol{Z}$ in the following way. First, choose subject j from $\{2, \ldots, n\}$ uniformly at random and consider the pair $(\boldsymbol{w}_1, \boldsymbol{w}_j)$. Then sample (Z_1, Z_j) uniformly at random from the assignments $(Z_1 = 1, Z_j = 0)$ and $(Z_1 = 0, Z_j = 1)$. Finally, assign the remaining $n-2$ subjects to groups by an independent randomization to equal groups (chosen uniformly at random). We will show that conditional distribution of $T(\boldsymbol{v}, \boldsymbol{Z})$, conditional on the choice of j, is SD with a mean of $\mathbb{E}\big[T(\boldsymbol{v}, \boldsymbol{Z})\big]$. This suffices to prove the theorem, since the unconditional pmf for $T(\boldsymbol{v}, \boldsymbol{Z})$ is the weighted sum of such conditional pmfs. There are four cases.

1. $\boldsymbol{w}_j = (1,1)$. Then $w_1(1) - w_j(0) = 0$, and $w_j(1) - w_1(0) = 1$, so the distribution of

$$Z_1 w_1(1) + (1 - Z_1) w_1(0) + Z_j w_j(1) + (1 - Z_j) w_j(0) \tag{8.47}$$

   is uniform on the set $\{0,1\}$. Further, the contribution to T from the other $n-2$ subjects, after the independent randomization is also symmetric decreasing on $m^{-1}\mathbb{Z}$, by assumption. Then the sum of these (conditionally) independent variables is SD with mean $\mathbb{E}\big[T(\boldsymbol{v}, \boldsymbol{Z})\big]$, by Lemma 8.5.

2. $\boldsymbol{w}_j = (1,0)$. One checks that the distribution (8.47) is constant and equal to 1, and that the same reasoning as in the previous case applies, since the sum of a random variable that is SD and the constant 1 is still SD.

3. $\boldsymbol{w}_j = (0,1)$. The distribution of (8.47) is constant and equal to 0, and the same reasoning as in the previous point applies.

4. $\boldsymbol{w}_j = (0,0)$. The distribution of (8.47) is uniform on $\{0,1\}$, and we can apply Lemma 8.5.

This completes the proof. □

Lemma 8.7. *Fix any potential outcome count vector of the form*

$$\boldsymbol{v} = (v_{11}, 0, 0, v_{00}). \tag{8.48}$$

Then the pmf of $T(\boldsymbol{v}, \boldsymbol{Z})$ is SD on $m^{-1}(2\mathbb{Z})$ if v_{11} is even. It is SD on $m^{-1}(2\mathbb{Z}+1)$ if v_{11} is odd.

Remark 8.8. Note that this statement concerns the lattices $m^{-1}(2\mathbb{Z})$ and $m^{-1}(2\mathbb{Z}+1)$, instead of the lattice $m^{-1}\mathbb{Z}$, due to the parity issue mentioned in the introduction to this subsection.

Proof. We prove the claim for all even $n \in \mathbb{Z}_{\geq 0}$ by induction. The base case $n = 0$ is trivial. For the induction step, we suppose that the claim holds for $n-2$ and will show it is true for n. Fix an arbitrary $\boldsymbol{v}$ of the form (8.48) with $v_{11} + v_{00} = n$. We suppose that $v_{11} \geq 2$. When $v_{11} = 0$, the conclusion is trivial, and when $v_{11} = 1$, we can apply the reasoning below with v_{00} instead of v_{11}.

By relabeling, we may suppose $\boldsymbol{w}_1 = (1,1)$. We sample $\boldsymbol{Z}$ as in the previous proof by choosing a partner $\boldsymbol{w}_j$ for $\boldsymbol{w}_1$ uniformly at random, randomizing the pair $\boldsymbol{w}_1$ and $\boldsymbol{w}_j$ to treatment and control, and independently randomizing the other $n-2$ subjects to equal groups.

We consider first the case that $\boldsymbol{w}_2 = (1,1)$. Then the distribution of

$$Z_1 w_1(1) + (1 - Z_1) w_1(0) + Z_j w_j(1) + (1 - Z_j) w_j(0) \tag{8.49}$$

conditional on the choice of partner and that fact that subjects 1 and 2 are assigned to different groups is constant and equal to 0. This implies we are done by induction, since $v_{11} - 2$, the number of $(1,1)$ potential outcome vectors in the remaining $n-2$ subjects, has the same parity as v_{11}, and the sampling distributions of

$$\sum_{i=1}^{n} Z_i w_i(1) + (1 - Z_i) w_i(0) \tag{8.50}$$

and

$$\sum_{i=3}^{n} Z_i w_i(1) + (1 - Z_i) w_i(0) \tag{8.51}$$

are then both supported and SD on either $m^{-1}(2\mathbb{Z})$ or $m^{-1}(2\mathbb{Z}+1)$, as desired.

We next consider the case where $\boldsymbol{w}_2 = (0,0)$. Then the conditional distribution of (8.49) is uniform on $\{-1,1\}$. We are again done by induction, since there are $v_{11} - 1$ potential outcome vectors of the form $(1,1)$ in the remaining $n-2$ subjects, and this number has the opposite parity to v_{11}. Reasoning as in the proof of Lemma 8.5, we see that the sums (8.50) and (8.51) are such that one is supported and SD on $m^{-1}(2\mathbb{Z})$ and one is supported and SD on $m^{-1}(2\mathbb{Z}+1)$, as desired.

Combining the conclusions of these two cases completes the proof. □

Lemma 8.9. *Let $n = 2m$ for some $m \in \mathbb{Z}_{>0}$. Fix a potential outcome count vector*

$$\boldsymbol{v} = (v_{11}, 1, 0, v_{00}) \text{ or } (v_{11}, 0, 1, v_{00}). \tag{8.52}$$

Then the pmf of $T(\boldsymbol{v}, \boldsymbol{Z})$ is SD on $m^{-1}\mathbb{Z}$.

Proof. We consider only the first case, where $v_{10} = 1$ and $v_{01} = 0$, since the argument for the other case is similar. By relabeling, we may suppose $\boldsymbol{w}_1 = (1, 0)$, and we resample $\boldsymbol{Z}$ as in the proof of Lemma 8.6 by first choosing uniformly at random a partner $\boldsymbol{w}_j$ for $\boldsymbol{w}_1$ from the set $\{\boldsymbol{w}_2, \dots, \boldsymbol{w}_n\}$, sampling (Z_1, Z_j) uniformly at random from the assignments $(Z_1 = 1, Z_j = 0)$ and $(Z_1 = 0, Z_j = 1)$, and assigning the remaining $n - 2$ subjects to groups by an independent, uniform randomization to equal groups. Let $\boldsymbol{v}^\circ$ denote the (random) count vector for the potential outcome table given by $\{\boldsymbol{w}_2, \dots, \boldsymbol{w}_n\} \setminus \{\boldsymbol{w}_j\}$, and let $\boldsymbol{Z}'$ denote a vector uniformly distributed on $\mathcal{Z}(n-2)$.

We now condition on the choice of $\boldsymbol{w}_j$. By Lemma 8.7, the conditional distribution of $T(\boldsymbol{v}^\circ, \boldsymbol{Z}')$ for the remaining $n - 2$ people is SD on $(m-1)^{-1}(2\mathbb{Z})$ or $(m-1)^{-1}(2\mathbb{Z}+1)$. Similarly, the conditional distribution of (8.47) for the partnership is always uniform on $\{0, 1\}$, since $\boldsymbol{w}_j = (0, 0)$ or $\boldsymbol{w}_j = (1, 1)$. Then the conditional pmf for $T(\boldsymbol{v}, \boldsymbol{Z})$ is SD on $m^{-1}\mathbb{Z}$ by Lemma 8.5, with mean independent of the choice of $\boldsymbol{w}_j$ and equal to the unconditional mean $\mathbb{E}\big[T(\boldsymbol{v}, \boldsymbol{Z})\big]$. Since this holds for every choice of $\boldsymbol{w}_j$, it holds for the unconditional distribution of $T(\boldsymbol{v}, \boldsymbol{Z})$, as desired. □

The proof of the following lemma is somewhat computational, so we defer it to the next subsection.

Lemma 8.10. *Let $n = 2m$ for some $m \in \mathbb{Z}_{>0}$. Fix a potential outcome count vector*

$$\boldsymbol{v} = (v_{11}, 2, 0, v_{00}). \tag{8.53}$$

Then the pmf of $T(\boldsymbol{v}, \boldsymbol{Z})$ is SD on $m^{-1}\mathbb{Z}$.

Corollary 8.11. *Let $n = 2m$ for some $m \in \mathbb{Z}_{>0}$. Fix a potential outcome count vector*

$$\boldsymbol{v} = (v_{11}, 1, 1, v_{00}) \text{ or } (v_{11}, 0, 2, v_{00}). \tag{8.54}$$

Then the pmf of $T(\boldsymbol{v}, \boldsymbol{Z})$ is SD on $m^{-1}\mathbb{Z}$.

Proof. The pmfs of $T(\boldsymbol{v}, \boldsymbol{Z})$ for the given vectors are translates of the pmf corresponding to $T(\boldsymbol{v}, \boldsymbol{Z})$ for

$$\boldsymbol{v} = (v_{11}, 2, 0, v_{00}), \tag{8.55}$$

so we are done by the previous lemma. □

Lemma 8.12. *Let $n = 2m$ for some $m \in \mathbb{Z}_{>0}$. Fix any potential outcome count vector*

$$\boldsymbol{v} = (v_{11}, v_{10}, v_{01}, v_{00}), \tag{8.56}$$

and suppose $v_{10} + v_{01} \geq 1$. Then the pmf of $T(\boldsymbol{v}, \boldsymbol{Z})$ is SD on $m^{-1}\mathbb{Z}$.

Proof. We induct on n. The base case $n = 0$ is trivial. For the induction step, suppose that the claim is true for $n - 2$, and fix any $\boldsymbol{v}$ corresponding to a potential outcome table with n subjects such that $\min(v_{10}, v_{01}) \geq 1$. If $(v_{10}, v_{01}) = (1, 0)$ or $(v_{10}, v_{01}) = (0, 1)$, we are done by Lemma 8.9. If $v_{01} + v_{01} = 2$, we are done by Lemma 8.10 and Corollary 8.11.

In the case $v_{01} + v_{01} > 2$, we will use Lemma 8.6 to show the claim is true for the given $\boldsymbol{v}$. In this case, suppose first that $v_{10} \geq 2$ and $v_{01} \geq 1$. Then by the induction hypothesis, Lemma 8.9, Lemma 8.10, and Corollary 8.11, the first condition given in the statement of Lemma 8.6 holds. The case where $v_{10} \geq 1$ and $v_{01} \geq 2$ is analogous. This completes the proof. □

Proof of Lemma 4.2. This is an immediate consequence of combining Lemma 8.12 and Lemma 8.4. □

8.4 Proof of Lemma 8.10

We first recall a formula of Copas [Cop73]. Given a potential outcome table $\boldsymbol{w}$ and a randomization $\boldsymbol{Z}$, we define a *treatment count vector* $\boldsymbol{x}$ by

$$\boldsymbol{x} = (x_{11}, x_{10}, x_{01}, x_{00}), \qquad x_{ab} = \sum_{j=1}^{n} Z_j \mathbb{1}\{\boldsymbol{w}_j = (a, b)\}. \tag{8.57}$$

Note that

$$x_{11} + x_{10} + x_{01} + x_{00} = m, \tag{8.58}$$

if $n = 2m$ and we randomize into equal groups.

Lemma 8.13 ([Cop73])**.** *Fix a potential outcome count vector $\boldsymbol{v}$. For any $s_0, s_1 \in \mathbb{Z}_{\geq 0}$, the probability of observing a treatment count vector $\boldsymbol{x}$ such that*

$$s_1 = x_{11} + x_{10}, \qquad s_0 = (v_{11} - x_{11}) + (v_{01} - x_{01}) \tag{8.59}$$

is

$$p(s_1, s_0) = C_{\boldsymbol{v}} \sum_{x=-\infty}^{\infty} \binom{v_{11}}{x} \binom{v_{10}}{s_1 - x} \binom{v_{01}}{v_{11} + v_{01} - s_0 - x} \binom{v_{00}}{m - v_{11} - s_1 - v_{01} + s_0 + x}, \tag{8.60}$$

where $C_{\boldsymbol{v}} > 0$ is a constant depending only on $\boldsymbol{v}$.

Proof of Lemma 8.10. We explicitly compute the distribution of $T(\boldsymbol{v}, \boldsymbol{Z})$. We first note that direct computation shows that $m\mathbb{E}\big[T(\boldsymbol{v}, \boldsymbol{Z})\big] = 1$, and we know that the distribution of $T(\boldsymbol{v}, \boldsymbol{Z})$ is symmetric about its mean by Lemma 7.2.

Adopt the notation of Lemma 8.13. Then

$$s_1 = x_{11} + x_{10}, \qquad s_0 = v_{11} - x_{11}. \tag{8.61}$$

With p defined as in (8.60), and using the assumed form of $\boldsymbol{v}$, we have

$$p(s_1, s_0) = C_{\boldsymbol{v}} \sum_{x} \binom{v_{11}}{x} \binom{2}{s_1 - x} \binom{0}{v_{11} - s_0 - x} \binom{v_{00}}{m - v_{11} - s_1 - v_{01} + s_0 + x} \tag{8.62}$$

$$= C_{\boldsymbol{v}} \sum_{x} \binom{v_{11}}{x} \binom{2}{x_{11} + x_{10} - x} \binom{0}{x_{11} - x} \binom{v_{00}}{m - v_{11} - s_1 + s_0 + x}. \tag{8.63}$$

For this to be nonzero, we must have $x = x_{11}$. Using this and (8.61), we get

$$p(s_1, s_0) = p(x_{11}, x_{10}) = C_{\boldsymbol{v}} \binom{v_{11}}{x_{11}} \binom{2}{x_{10}} \binom{v_{00}}{m - x_{11} - x_{10}}. \tag{8.64}$$

We abbreviate $j = mT(\boldsymbol{v}, \boldsymbol{Z}) - 1$ (centering by subtracting the expectation). Then (using $v_{01} = 0$)

$$j = s_1 - s_0 - 1 = x_{11} + x_{10} - v_{11} + x_{11} - 1 = 2x_{11} + x_{10} - v_{11} - 1. \tag{8.65}$$

This yields

$$x_{11} = \frac{j + v_{11} - x_{10} + 1}{2}. \tag{8.66}$$

We also have

$$\frac{v_{00}+v_{11}}{2}+1=m. \tag{8.67}$$

Using (8.66) and (8.67) in (8.64), we get

$$p(x_{11},x_{10})=C_{\boldsymbol{v}}\binom{2}{x_{10}}\binom{v_{11}}{\frac{v_{11}}{2}+\frac{1-x_{10}+j}{2}}\binom{v_{00}}{\frac{v_{00}}{2}+\frac{1-x_{10}-j}{2}}. \tag{8.68}$$

We now consider two cases, depending on the parity of v_{11}.

Case I: v_{11} is even. Then v_{00} is also even, since n is even. There are two subcases, depending on the parity of j. Suppose that j is even. Then parity considerations from (8.66) force $x_{10}=1$. Then the probability mass function becomes

$$p(j)=2C_{\boldsymbol{v}}\binom{v_{11}}{\frac{v_{11}}{2}+\frac{j}{2}}\binom{v_{00}}{\frac{v_{00}}{2}-\frac{j}{2}}. \tag{8.69}$$

When j is odd, both $x_{10}=0$ and $x_{10}=2$ are possible, and the pmf is

$$p(j)=C_{\boldsymbol{v}}\binom{v_{11}}{\frac{v_{11}}{2}+\frac{1+j}{2}}\binom{v_{00}}{\frac{v_{00}}{2}+\frac{1-j}{2}}+C_{\boldsymbol{v}}\binom{v_{11}}{\frac{v_{11}}{2}+\frac{j-1}{2}}\binom{v_{00}}{\frac{v_{00}}{2}+\frac{-1-j}{2}}. \tag{8.70}$$

Now it is a straightforward computation to show that $p(j)$ is decreasing for $j\geq 0$. We will show that $p(j)\geq p(j+1)$ for all $j\geq 0$. First, suppose that j is even. Then, dividing through by $C_{\boldsymbol{v}}$. we must show that

$$2\binom{v_{11}}{\frac{v_{11}}{2}+\frac{j}{2}}\binom{v_{00}}{\frac{v_{00}}{2}-\frac{j}{2}}\geq\binom{v_{11}}{\frac{v_{11}}{2}+\frac{j+2}{2}}\binom{v_{00}}{\frac{v_{00}}{2}-\frac{j}{2}}+\binom{v_{11}}{\frac{v_{11}}{2}+\frac{j}{2}}\binom{v_{00}}{\frac{v_{00}}{2}+\frac{-2-j}{2}}. \tag{8.71}$$

This is clear, since $\binom{n}{k}$ is symmetric and decreasing away from $k=n/2$.

Next, we suppose that j is odd. We want to show

$$\binom{v_{11}}{\frac{v_{11}}{2}+\frac{1+j}{2}}\binom{v_{00}}{\frac{v_{00}}{2}+\frac{1-j}{2}}+\binom{v_{11}}{\frac{v_{11}}{2}+\frac{j-1}{2}}\binom{v_{00}}{\frac{v_{00}}{2}+\frac{-1-j}{2}}\geq 2\binom{v_{11}}{\frac{v_{11}}{2}+\frac{j+1}{2}}\binom{v_{00}}{\frac{v_{00}}{2}-\frac{j+1}{2}}. \tag{8.72}$$

This is again clear, for the same reason.

Case II: v_{11} is odd. Then v_{11} is odd, since n is even. There are two subcases, depending on the parity of j. Suppose that j is odd. Then parity considerations from (8.66) force $x_{10}=1$. Then the probability mass function becomes

$$p(j)=2C_{\boldsymbol{v}}\binom{v_{11}}{\frac{v_{11}}{2}+\frac{j}{2}}\binom{v_{00}}{\frac{v_{00}}{2}-\frac{j}{2}}. \tag{8.73}$$

When j is even, both $x_{10}=0$ and $x_{10}=2$ are possible, and the pmf is

$$p(j)=C_{\boldsymbol{v}}\binom{v_{11}}{\frac{v_{11}}{2}+\frac{1+j}{2}}\binom{v_{00}}{\frac{v_{00}}{2}+\frac{1-j}{2}}+C_{\boldsymbol{v}}\binom{v_{11}}{\frac{v_{11}}{2}+\frac{j-1}{2}}\binom{v_{00}}{\frac{v_{00}}{2}+\frac{-1-j}{2}}. \tag{8.74}$$

The same verification as in the previous case shows that $p(j)$ is decreasing for $j\geq 0$. □

9 Proof of Proposition 5.3

Lemma 9.1. *Fix $\varepsilon > 0$ and $K \in \mathbb{Z}_{>0}$. Then*

$$\mathbb{P}\Big(\big|p(\boldsymbol{w}, \boldsymbol{Y}, \boldsymbol{Z}) - S\big| > \varepsilon\Big) \leq 2\exp\left(-K\varepsilon^2\right). \tag{9.1}$$

Proof. We have $\mathbb{E}[V_i] = p(\boldsymbol{w}, \boldsymbol{Y}, \boldsymbol{Z})$ and $0 \leq V_i \leq 1$ by definition, so Hoeffding's inequality yields

$$\mathbb{P}\left(\left|K \cdot p(\boldsymbol{w}, \boldsymbol{Y}, \boldsymbol{Z}) - \sum_{i=1}^{K} V_i\right| > K\varepsilon\right) \leq 2\exp\left(-\frac{(K\varepsilon)^2}{K}\right) = 2\exp\left(-K\varepsilon^2\right), \tag{9.2}$$

as desired. □

Lemma 9.2. *Fix a potential outcome table $\boldsymbol{y}$ and $\alpha \in (0, 1)$.*

1. *With probability at least $1 - 8(n+1)\lfloor \log_2(n+1) + 2 \rfloor \exp(-K\varepsilon^2)$, we have*

$$\mathcal{I}_\alpha(\boldsymbol{Y}, \boldsymbol{Z}) \subset \mathcal{I}_{\alpha,\varepsilon,K}(\boldsymbol{Y}, \boldsymbol{Z}). \tag{9.3}$$

2. *With probability at least $1 - 8(n+1)\lfloor \log_2(n+1) + 2 \rfloor \exp(-K\varepsilon^2)$, we have*

$$\mathcal{I}_{\alpha,\varepsilon,K}(\boldsymbol{Y}, \boldsymbol{Z}) \subset \mathcal{I}_{\alpha-2\varepsilon}(\boldsymbol{Y}, \boldsymbol{Z}). \tag{9.4}$$

Proof. We begin with the first claim. The event where $\mathcal{I}_\alpha(\boldsymbol{Y}, \boldsymbol{Z}) \subset \mathcal{I}_{\alpha,\varepsilon,K}(\boldsymbol{Y}, \boldsymbol{Z})$ is contained in the event

$$\big\{U_\alpha(\boldsymbol{Y}, \boldsymbol{Z}) \leq U_{\alpha,\varepsilon,K}(\boldsymbol{Y}, \boldsymbol{Z})\big\} \cap \big\{L_\alpha(\boldsymbol{Y}, \boldsymbol{Z}) \geq L_{\alpha,\varepsilon,K}(\boldsymbol{Y}, \boldsymbol{Z})\big\}. \tag{9.5}$$

We begin by bounding the probability that $U_\alpha(\boldsymbol{Y}, \boldsymbol{Z}) \leq U_{\alpha,\varepsilon,K}(\boldsymbol{Y}, \boldsymbol{Z})$ does not hold. Let N be the total number of permutation tests performed by Algorithm 4.3 in the course of finding $U_\alpha(\boldsymbol{Y}, \boldsymbol{Z})$. Letting 0 represent acceptance and 1 represent rejection, we let $\boldsymbol{a} = (a_i)_{i=1}^N \in \{0,1\}^N$ denote the sequences of acceptances and rejections made by the permutation tests in Algorithm 4.3. We condition on the observed data $(\boldsymbol{Y}, \boldsymbol{Z})$, so that $\boldsymbol{a}$ becomes deterministic. Given this observed data, let M denote the (random) number of approximate permutations performed by Algorithm 5.2 before returning $U_{\alpha,\varepsilon,K}(\boldsymbol{Y}, \boldsymbol{Z})$. Let $\boldsymbol{b} = (b_i)_{i=1}^M \in \{0,1\}^M$ denote the corresponding (random) sequence of acceptances and rejections. Let J be a random variable denoting the the smallest positive integer such that $a_J \neq b_J$. If $a_i = b_i$ for all $i \in [\![1, N]\!]$, in which case $U_\alpha = U_{\alpha,\varepsilon,K}$, we set $J = N + 1$ and define $b_J = 0$. Note that J is well defined, since it is not possible for Algorithm 5.2 to terminate unless $\boldsymbol{b}$ differs from $\boldsymbol{a}$ at some location, or $M = N$ and $\boldsymbol{b} = \boldsymbol{a}$. Further, we have the deterministic bound $J \leq N + 1$.

The event $\big\{U_\alpha(\boldsymbol{Y}, \boldsymbol{Z}) \leq U_{\alpha,\varepsilon,K}(\boldsymbol{Y}, \boldsymbol{Z})\big\}$ is equal to the event that $b_J = 0$. That is, in order that $U_\alpha \leq U_{\alpha,\varepsilon,K}$, at the first time the result of an approximate permutation test in Algorithm 5.2 differs from the result of an exact permutation test, the approximate test must accept. We therefore

compute

$$\mathbb{P}\Big(\big\{U_\alpha(\boldsymbol{Y},\boldsymbol{Z}) \le U_{\alpha,\varepsilon,K}(\boldsymbol{Y},\boldsymbol{Z})\big\}^c\Big) = \mathbb{P}(b_J = 1) \tag{9.6}$$

$$= \sum_{i=1}^{N} \mathbb{P}(b_i = 1 \text{ and } J = i) \tag{9.7}$$

$$\le \sum_{i=1}^{N} \mathbb{P}(b_i = 1 \text{ and } a_i = 0) \tag{9.8}$$

$$\le N \cdot 2\exp(-K\varepsilon^2) \tag{9.9}$$

$$\le 4(n+1)\lfloor \log_2(n+1) + 2 \rfloor \exp(-K\varepsilon^2). \tag{9.10}$$

To justify (9.9), we note that $a_i = 0$ implies $\boldsymbol{p}(\boldsymbol{w},\boldsymbol{Y},\boldsymbol{Z}) \ge \alpha$ for the table $\boldsymbol{w}$ considered at the i-th step of Algorithm 4.3. For Algorithm 5.2 to reject $\boldsymbol{w}$, it must be rejected by Algorithm 5.1. In the notation of Algorithm 5.2, this happens when $S + \varepsilon < \alpha$. Then by Lemma 9.1, Algorithm 5.1 rejects $\boldsymbol{w}$ with probability at most $2\exp(-K\varepsilon^2)$. In (9.10), we used the fact that at most

$$2(n+1)\lfloor \log_2(n+1) + 2 \rfloor \tag{9.11}$$

permutation tests are required in Algorithm 5.2 to find U_α, which was shown in the proof of Theorem 4.1. The bound on the probability that $L_\alpha \ge L_{\alpha,\varepsilon,K}$ fails to holds follows by a nearly identical argument. From a union bound over the complements of the two events in (9.5), we deduce (9.3).

For the proof of the second claim, we keep the notation $\boldsymbol{a}$ and $\boldsymbol{a}$, but now define $\boldsymbol{a}$ using the interval $\mathcal{I}_{\alpha-2\varepsilon}(\boldsymbol{Y},\boldsymbol{Z})$. The event $\mathcal{I}_{\alpha,\varepsilon,K}(\boldsymbol{Y},\boldsymbol{Z}) \subset \mathcal{I}_{\alpha-2\varepsilon}(\boldsymbol{Y},\boldsymbol{Z})$ is contained in the event

$$\big\{U_{\alpha-2\varepsilon}(\boldsymbol{Y},\boldsymbol{Z}) \ge U_{\alpha,\varepsilon,K}(\boldsymbol{Y},\boldsymbol{Z})\big\} \cap \big\{L_{\alpha-2\varepsilon}(\boldsymbol{Y},\boldsymbol{Z}) \le L_{\alpha,\varepsilon,K}(\boldsymbol{Y},\boldsymbol{Z})\big\}. \tag{9.12}$$

The event $\big\{U_{\alpha-2\varepsilon}(\boldsymbol{Y},\boldsymbol{Z}) \ge U_{\alpha,\varepsilon,K}(\boldsymbol{Y},\boldsymbol{Z})\big\}$ is equal to the event

$$\{b_J = 1\} \cup \{J = N+1\}. \tag{9.13}$$

We compute

$$\mathbb{P}\Big(\big\{U_{\alpha-2\varepsilon}(\boldsymbol{Y},\boldsymbol{Z}) \ge U_{\alpha,\varepsilon,K}(\boldsymbol{Y},\boldsymbol{Z})\big\}^c\Big) = \mathbb{P}(b_J = 0 \text{ and } J \le N) \tag{9.14}$$

$$= \sum_{i=1}^{N} \mathbb{P}(b_i = 0 \text{ and } J = i) \tag{9.15}$$

$$\le \sum_{i=1}^{N} \mathbb{P}(b_i = 0 \text{ and } a_i = 1) \tag{9.16}$$

$$\le N \cdot 2\exp(-K\varepsilon^2) \tag{9.17}$$

$$\le 4(n+1)\lfloor \log_2(n+1) + 2 \rfloor \exp(-K\varepsilon^2). \tag{9.18}$$

To justify (9.17), we note that $a_i = 1$ implies $\boldsymbol{p}(\boldsymbol{w},\boldsymbol{Y},\boldsymbol{Z}) < \alpha - 2\varepsilon$ for the table $\boldsymbol{w}$ considered at the i-th step of Algorithm 4.3. For Algorithm 5.2 to accept $\boldsymbol{w}$, it must be accepted by Algorithm 5.1.

We bound the acceptance probability as

$$\mathbb{P}(S+\varepsilon \geq \alpha) \leq \mathbb{P}\big(|p(\boldsymbol{w},\boldsymbol{Y},\boldsymbol{Z}) - S| > \varepsilon\big) \leq 2\exp(-K\varepsilon^2) \tag{9.19}$$

using Lemma 9.1. The estimate for the event $\big\{L_{\alpha-2\varepsilon}(\boldsymbol{Y},\boldsymbol{Z}) \leq L_{\alpha,\varepsilon,K}(\boldsymbol{Y},\boldsymbol{Z})\big\}$ is identical, and is shown by essentially the same argument. This completes the proof of (9.4) after using a union bound. □

Proof of Proposition 5.3. For the first part, it suffices to choose K large enough so that

$$\mathbb{P}(\mathcal{I}_{\alpha-\varepsilon} \subset \mathcal{I}_{\alpha-\varepsilon,\varepsilon,K}) \geq 1-\varepsilon. \tag{9.20}$$

Then the claim follows from Lemma 9.2 and the fact that

$$16n\log_2(n) \geq 8(n+1)\lfloor \log_2(n+1)+2\rfloor \tag{9.21}$$

for $n \geq 10$. The second claim is immediate from Lemma 9.2 and the previous inequality. □

10 Proof of Proposition 6.2

Proof of Proposition 6.2. First, note that by Proposition 3.2, we have

$$\mathbb{P}\big(\tau(\boldsymbol{y}) \in \mathcal{J}_\alpha(\boldsymbol{M},\boldsymbol{Z})\big) \geq 1-\alpha, \tag{10.1}$$

where

$$J_\alpha(\boldsymbol{M},\boldsymbol{Z}) = \bigcup_{\boldsymbol{Y}'\in\mathcal{Y}(\boldsymbol{M})} \mathcal{I}_\alpha(\boldsymbol{Y}',\boldsymbol{Z}), \tag{10.2}$$

and $\mathcal{Y}(\boldsymbol{M})$ consists of all vectors of realized outcomes $\boldsymbol{Y}'$ that could arise from the partially observed vector $\boldsymbol{M}$. To complete the proof of the theorem it then suffices to show that

$$J_\alpha(\boldsymbol{M},\boldsymbol{Z}) \subset I^\circ_\alpha(\boldsymbol{M},\boldsymbol{Z}), \tag{10.3}$$

or equivalently,

$$I^\circ_\alpha(\boldsymbol{M},\boldsymbol{Z})^c \subset J_\alpha(\boldsymbol{M},\boldsymbol{Z})^c = \bigcap_{\boldsymbol{Y}'\in\mathcal{Y}(\boldsymbol{M})} \mathcal{I}_\alpha(\boldsymbol{Y}',\boldsymbol{Z})^c. \tag{10.4}$$

Fix some $\tau_0 \in I^\circ_\alpha(\boldsymbol{M},\boldsymbol{Z})^c$. We may suppose that $\tau_0 > U_\alpha(\boldsymbol{Y}^{(+)},\boldsymbol{Z})$, since the analogous argument in the complementary case is similar. Then by definition, for all $\boldsymbol{w}$ such that $\tau(\boldsymbol{w}) = \tau_0$, $\boldsymbol{w}$ is incompatible with the data $(\boldsymbol{Y}^{(+)},\boldsymbol{Z})$. We must show that all such $\boldsymbol{w}$ are incompatible with all choices of $(\boldsymbol{Y}',\boldsymbol{Z})$ with $\boldsymbol{Y}' \in \mathcal{Y}(\boldsymbol{M})$. Fix such a $\boldsymbol{w}$ and a choice of $\boldsymbol{Y}'$.

By the definition of $\boldsymbol{Y}^{(+)}$, and since $T(\boldsymbol{Y}^{(+)},\boldsymbol{Z})$ and $T(\boldsymbol{Y}',\boldsymbol{Z})$ are coupled through $\boldsymbol{Z}$, we have

$$\tau_0 \geq T(\boldsymbol{Y}^{(+)},\boldsymbol{Z}) \geq T(\boldsymbol{Y}',\boldsymbol{Z}), \tag{10.5}$$

where the first inequality follows from Lemma 3.1 and our assumption that $\tau_0 > U_\alpha(\boldsymbol{Y}^{(+)},\boldsymbol{Z})$. Then

$$\mathbb{P}\left(\big|T(\widetilde{\boldsymbol{Y}},\widetilde{\boldsymbol{Z}}) - \tau_0\big| \geq \big|T(\boldsymbol{Y}^{(+)},\boldsymbol{Z}) - \tau_0\big|\right) \leq \mathbb{P}\left(\big|T(\widetilde{\boldsymbol{Y}},\widetilde{\boldsymbol{Z}}) - \tau_0\big| \geq \big|T(\boldsymbol{Y}',\boldsymbol{Z}) - \tau_0\big|\right) < \alpha, \tag{10.6}$$

so $\boldsymbol{w}$ is incompatible with $(\boldsymbol{Y}',\boldsymbol{Z})$, as desired. □

References

[BB19] Zach Branson and Marie-Abèle Bind. Randomization-based inference for Bernoulli trial experiments and implications for observational studies. *Statistical Methods in Medical Research*, 28(5):1378–1398, 2019.

[Cop73] J. B. Copas. Randomization models for the matched and unmatched 2×2 tables. *Biometrika*, 60(3):467–476, 1973.

[GLSR04] Robert Greevy, Bo Lu, Jeffrey Silber, and Paul Rosenbaum. Optimal multivariate matching before randomization. *Biostatistics*, 5(2):263–275, 2004.

[IR15] Guido Imbens and Donald Rubin. *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press, 2015.

[Kal18] Nathan Kallus. Optimal a priori balance in the design of controlled experiments. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(1):85–112, 2018.

[LD16] Xinran Li and Peng Ding. Exact confidence intervals for the average causal effect on a binary outcome. *Statistics in Medicine*, 35(6):957–960, 2016.

[LS20] Tor Lattimore and Csaba Szepesvári. *Bandit Algorithms*. Cambridge University Press, 2020.

[MPS88] Cyrus R. Mehta, Nitin R. Patel, and Pralay Senchaudhuri. Importance sampling for estimating exact probabilities in permutational inference. *Journal of the American Statistical Association*, 83(404):999–1005, 1988.

[MR12] Kari Lock Morgan and Donald Rubin. Rerandomization to improve covariate balance in experiments. *The Annals of Statistics*, 40(2):1263–1282, 2012.

[RH15] Joseph Rigdon and Michael Hudgens. Randomization inference for treatment effects on a binary outcome. *Statistics in Medicine*, 34(6):924–935, 2015.

[Rob88] James Robins. Confidence intervals for causal parameters. *Statistics in Medicine*, 7(7):773–785, 1988.

[SS80] Thomas Santner and Mark Snell. Small-sample confidence intervals for $p_1 - p_2$ and p_1/p_2 in 2×2 contingency tables. *Journal of the American Statistical Association*, 75(370):386–394, 1980.

[Was04] Larry Wasserman. *All of Statistics: A Concise Course in Statistical Inference*. Springer, 2004.