

Zach Branson* and Luke W. Miratrix

Randomization Tests that Condition on Non-Categorical Covariate Balance

<https://doi.org/10.1515/jci-2018-0004>

Received February 2, 2018; revised October 3, 2018; accepted December 27, 2018

Abstract: A benefit of randomized experiments is that covariate distributions of treatment and control groups are balanced on average, resulting in simple unbiased estimators for treatment effects. However, it is possible that a particular randomization yields covariate imbalances that researchers want to address in the analysis stage through adjustment or other methods. Here we present a randomization test that conditions on covariate balance by only considering treatment assignments that are similar to the observed one in terms of covariate balance. Previous conditional randomization tests have only allowed for categorical covariates, while our randomization test allows for any type of covariate. Through extensive simulation studies, we find that our conditional randomization test is more powerful than unconditional randomization tests and other conditional tests. Furthermore, we find that our conditional randomization test is valid (1) unconditionally across levels of covariate balance, and (2) conditional on particular levels of covariate balance. Meanwhile, unconditional randomization tests are valid for (1) but not (2). Finally, we find that our conditional randomization test is similar to a randomization test that uses a model-adjusted test statistic.

Keywords: conditional inference, covariate adjustment, statistical power, validity

1 After randomization: To adjust or not to adjust?

Randomized experiments are often considered the “gold standard” of statistical inference because randomization balances the covariate distributions of the treatment and control groups on average, which limits confounding between treatment effects and covariate effects. However, it is possible that a particular treatment assignment from a randomized experiment yields covariate imbalances that researchers wish to address. One option is to employ experimental design strategies such as blocking or rerandomization [21], which prevent substantial covariate imbalance from occurring before the experiment is conducted. When these strategies are not employed, covariate imbalance must be addressed in the analysis stage rather than the design stage. The analyst of such experiments must make a choice: to adjust or not to adjust for the covariate imbalance realized by a particular randomization. If adjustment is done, it is typically done via statistical models (e. g., regression adjustment); however, the results from such adjustment may be biased and/or sensitive to model specification [15, 7, 1]. Meanwhile, unadjusted estimators—though unbiased across randomizations—could be confounded by the realized covariate imbalance at hand. Lin [18] rigorously investigated these tradeoffs between unadjusted and adjusted estimators, noting that biases due to regression are often minimal, but also that unadjusted estimators are appealing for their simplicity and transparency. Regardless of where these works fall on the “to adjust or not to adjust” spectrum, they all agree that accounting for covariate balance is a key concern in randomized experiments.

*Corresponding author: Zach Branson, Department of Statistics, Harvard University, Cambridge, Massachusetts, United States, e-mail: zbranson@g.harvard.edu

Luke W. Miratrix, Graduate School of Education and Department of Statistics, Harvard University, Cambridge, Massachusetts, United States, e-mail: lmiratrix@stat.harvard.edu

1.1 Accounting for covariate balance in randomization tests

In addition to model-based testing, one can use randomization tests to account for covariate balance in experiments. Randomization tests are often considered minimal-assumption approaches in that they usually only require assuming a probability distribution on treatment assignment rather than structural modeling assumptions or central limit theorems [27]. In particular, a randomization test requires specifying only (1) the assumed assignment mechanism and (2) the test statistic. In this context, one can account for covariate balance by making particular choices for the assignment mechanism or the test statistic, but most have focused on the choice of the latter. For example, many have found that using model-adjusted estimators as test statistics to address covariate imbalances can result in statistically powerful randomization tests ([23], [26], [27] Chapter 2, [11], [16] Chapter 5). Meanwhile, practitioners typically use the assignment mechanism that was actually used in the design of the experiment when conducting a randomization test (e. g., if units were assigned completely at random, then this same assignment mechanism is used during the randomization test). However, by considering other choices for the assignment mechanism, one can also account for covariate balance.

In particular, a small strand of literature has explored randomization tests that restrict the assignment mechanism to only consider treatment assignments that are similar to the observed one in terms of covariate balance, even if such an assignment mechanism was not explicitly specified by design. This literature has focused on cases where all covariates are categorical, and thus treatment assignment is characterized by permutations within covariate strata. For example, Rosenbaum [24] proposed a conditional permutation test for observational studies that permutes the treatment indicator within groups of units with the same covariate values. This test assumes (1) the treatment assignment is strongly ignorable, (2) the true propensity score model is a logistic regression model, and (3) the collection of covariates is sufficient for the logistic regression model. More recently, Hennessy et al. [10] proposed a conditional randomization test for randomized experiments that is similar to Rosenbaum [24] in that it also permutes within groups of units with the same covariate values, but it does not require any kind of model specification. Rosenbaum [24] and Hennessy et al. [10] only consider cases with categorical covariates, and they make connections between their randomization tests and adjustment methods for categorical covariates, such as post-stratification [20].

1.2 Our contribution: Considering non-categorical covariates

We develop a randomization test that conditions on the realized covariate balance of an experiment for the more general case where covariates may be non-categorical. We demonstrate that our randomization test is more powerful than randomization tests that do not condition on covariate balance and is comparable to randomization tests that use model-adjusted estimators as test statistics. In general, we recommend the use of randomization tests that either condition on covariate balance through the assignment mechanism or utilize model-adjusted test statistics, instead of an unconditional randomization test that uses an unadjusted test statistic.

Our main contribution is outlining a randomization test that conditions on covariate balance through the assignment mechanism for the general case of non-categorical covariates. Unlike the case where only categorical covariates are present, samples from the conditional randomization distribution cannot be obtained via permutations of the treatment indicator when there are non-categorical covariates. In response to this complication, we develop a rejection-sampling algorithm to sample from the conditional randomization distribution.

We find that our conditional randomization test appears to be equivalent to randomization tests that use regression-based test statistics. This contribution is particularly notable because most have characterized the choice of test statistic as the main avenue for increasing the power of a randomization test and for adjusting for imbalance in an experiment. Our work suggests how the choice of assignment mechanism can be an analogous avenue for obtaining statistically powerful randomization tests that appropriately adjust for imbalance.

Furthermore, through simulation, we also find that our conditional randomization test is valid across randomizations conditional on a particular level of covariate balance, while unconditional randomization tests are often not valid across such randomizations. This suggests that our conditional randomization test can be used to ensure that statistical inferences are valid for the observed data at hand; meanwhile, unconditional randomization tests do not provide this benefit. Overall this suggests that practitioners using randomization tests should either condition on observed imbalance or use adjusted test statistics rather than the traditional randomization procedures usually seen in the literature.

To build intuition for our conditional randomization test, in Section 2 we review randomization tests for Fisher’s Sharp Null and review the conditional randomization test of [10]. In Section 3 we outline our conditional randomization test, which can flexibly condition on multiple levels of balance for non-categorical covariates. In Section 4 we provide simulation evidence that our conditional randomization test (1) is more powerful than unconditional and other conditional randomization tests, and (2) is approximately equivalent to an unconditional randomization test that uses a regression-based test statistic. In Section 5 we conclude by discussing how confidence intervals can be constructed from our conditional randomization test and the extent to which our conditional randomization test can be used for observational studies.

2 Review of randomization tests for Fisher’s sharp null

We focus on randomization tests for Fisher’s Sharp Null. While conclusions from such tests are limited—the only conclusion that can be made is whether or not there is any treatment effect among the experimental units—in Section 5 we discuss how such tests can be inverted to yield uncertainty intervals as well.

First we review a general framework for randomization tests for Fisher’s Sharp Null. We then review the unconditional randomization test typically discussed in the literature under this framework. Finally, we review the conditional randomization test of Hennessy et al. [10] that conditions on categorical covariate balance.

2.1 Setup and randomization test procedure

Consider N units to be allocated to treatment and control in a randomized experiment. Following [29], let $Y_i(1)$ and $Y_i(0)$ denote the treatment and control potential outcomes, respectively, for unit $i = 1, \dots, N$, and let \mathbf{x}_i denote a p -dimensional vector of pre-treatment covariates. Let $W_i = 1$ if unit i is assigned to treatment and 0 otherwise. Furthermore, define $\mathbf{X} \equiv (\mathbf{x}_1, \dots, \mathbf{x}_N)^T$ and $\mathbf{W} \equiv (W_1, \dots, W_N)$ as the covariate matrix and vector of treatment assignments, respectively. The observed outcomes are $y_i = W_i Y_i(1) + (1 - W_i) Y_i(0)$. Importantly, the potential outcomes ($Y_i(1), Y_i(0)$) and covariates \mathbf{x}_i are fixed; the only stochastic element of the observed outcomes y_i is the treatment assignment W_i .

Throughout, we assume a completely randomized experiment, where the true distribution of the treatment assignment \mathbf{W} is:

$$P(\mathbf{W} = \mathbf{w}) = \begin{cases} \binom{N}{N_T}^{-1} & \text{if } \sum_{i=1}^N w_i = N_T \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

with the number of treated units, N_T , fixed. Many causal estimands can be considered in this framework, but we focus on the average treatment effect

$$\tau = \frac{1}{N} \sum_{i=1}^N (Y_i(1) - Y_i(0)) \quad (2)$$

because it is the most common estimand in the causal inference literature. The potential outcomes $Y_i(1)$ and $Y_i(0)$ are never both observed, so (2) needs to be estimated. One common estimator is the mean-difference

estimator

$$\hat{\tau}_{sd} = \frac{\sum_{i=1}^N W_i Y_i(1)}{\sum_{i=1}^N W_i} - \frac{\sum_{i=1}^N (1 - W_i) Y_i(0)}{\sum_{i=1}^N (1 - W_i)} = \frac{\sum_{i:W_i=1} Y_i}{N_T} - \frac{\sum_{i:W_i=0} Y_i}{N_C} = \bar{y}_T - \bar{y}_C \quad (3)$$

where $N_T \equiv \sum_{i=1}^N W_i$ and $N_C \equiv \sum_{i=1}^N (1 - W_i)$ are the number of units that receive treatment and control, respectively.

A common test for assessing if an estimate for the average treatment effect is statistically significant is to test for Fisher's Sharp Null:

$$H_0 : Y_i(1) = Y_i(0), \quad \forall i = 1, \dots, N \quad (4)$$

which states that there is no treatment effect for any of the N units. A rejection of Fisher's Sharp Null implies that a treatment effect is present. We focus on testing Fisher's Sharp Null because it is the most common hypothesis to assess using randomization tests in the causal inference literature [27, 16]. See [5] and the ensuing comments for a discussion of how testing Fisher's Sharp Null compares to testing Neyman's Weak Null within the context of randomization-based causal inference.

Under Fisher's Sharp Null, the outcomes for any particular randomization will be equal to the observed outcomes; i. e., the observed outcomes will be the same across all realizations of \mathbf{W} under the Sharp Null. Thus, under H_0 , the value of any test statistic $t(Y(\mathbf{W}), \mathbf{W}, \mathbf{X})$ can be computed for any particular realization of the treatment assignment \mathbf{W} . A common choice of test statistic is $t(Y(\mathbf{W}), \mathbf{W}, \mathbf{X}) = \hat{\tau}_{sd}$. Our framework can incorporate any test statistic that differentiates between treatment and control response; for now we will focus on the test statistic $\hat{\tau}_{sd}$, and later we will discuss model-adjusted test statistics. See [27, Chapter 2] for further discussion on choices of test statistics for randomization tests.

To test Fisher's Sharp Null, one compares the observed value of the test statistic, t^{obs} , to the randomization distribution of the test statistic under the Sharp Null. Importantly, the randomization distribution of the test statistic depends on the set of treatment assignments that one considers possible within the randomization test.

We follow the notation of [16, Chapter 4] and define \mathbb{W} as the set of treatment assignments with positive probability within a given randomization test. Given any test statistic $t(Y(\mathbf{W}), \mathbf{W}, \mathbf{X})$, the two-sided randomization test p -value for Fisher's Sharp Null is

$$P(|t(Y(\mathbf{W}), \mathbf{W}, \mathbf{X})| \geq |t^{obs}|) = \sum_{\mathbf{w} \in \mathbb{W}} \mathbb{I}(|t(Y(\mathbf{w}), \mathbf{w}, \mathbf{X})| \geq |t^{obs}|) P(\mathbf{W} = \mathbf{w}) \quad (5)$$

In other words, the p -value (5) is the probability that a test statistic larger than the observed one would have occurred under the Sharp Null, given the assignment mechanism $P(\mathbf{W})$.

Typically, the set \mathbb{W} is too large to feasibly compute (5). Instead, (5) can be approximated by randomly sampling $\mathbf{w}^{(1)}, \dots, \mathbf{w}^{(M)}$ from $P(\mathbf{W})$; then, the randomization-test p -value (5) is approximated by

$$P(|t(Y(\mathbf{W}), \mathbf{W}, \mathbf{X})| \geq |t^{obs}|) \approx \frac{\sum_{m=1}^M \mathbb{I}(|t(Y(\mathbf{w}^{(m)}), \mathbf{w}^{(m)}, \mathbf{X})| \geq |t^{obs}|)}{M} \quad (6)$$

Thus, testing Fisher's Sharp Null is a three-step procedure [8]:

1. Specify the distribution $P(\mathbf{W})$ (and, consequentially, \mathbb{W}) to be used within the randomization test.
2. Choose a test statistic $t(Y(\mathbf{W}), \mathbf{W}, \mathbf{X})$.
3. Compute or approximate the p -value (5).

In the remainder of this section we will discuss two randomization tests: one that does not condition on covariate balance and one that does. The only difference between the two tests is the first step in the procedure above, i. e., the choice of the assignment mechanism $P(\mathbf{W})$.

2.2 Unconditional randomization tests

The most common randomization test in the literature utilizes the same assignment mechanism used to design the experiment, the completely randomized assignment mechanism defined in (1). A completely randomized assignment mechanism assumes that $\mathbf{W} = \{\mathbf{w} : \sum_{i=1}^N w_i = N_T\}$, i. e., it only considers assignments where N_T units are assigned to treatment. Hennessy et al. [10] call randomization tests that assume a completely randomized assignment mechanism “unconditional randomization tests” because they do not condition on forms of covariate balance. Once $P(\mathbf{W})$ and a test statistic are specified, the randomization test follows the three-step procedure from Section 2.1. This test is also called a permutation test because random samples from $P(\mathbf{W})$ can be obtained by randomly permuting the observed treatment assignment \mathbf{W}^{obs} .

Instead of using $P(\mathbf{W})$ in the randomization test procedure, Hennessy et al. [10] proposed using an assignment mechanism that conditions on covariate balance.

2.3 Conditional randomization tests

Because the number of treated units is prespecified as part of the design of a completely randomized experiment, the unconditional randomization test in Section 2.2 follows the typical recommendation to “analyze as you randomize.” However, many have recommended conditioning on the observed number of treated units even when the number of treated units was not specified by design [9, 31, 20, 28]. The goal of conditional inference in general (and conditional randomization tests specifically) is to focus inference on experiments that are most relevant to the data at hand by conditioning on pertinent statistics such as the number of treated units or forms of covariate balance. As we show through simulation in Section 4, conditional randomization tests can have the benefit of being valid conditional on the data as well as being valid unconditionally, whereas unconditional randomization tests are only valid unconditionally.

To formalize this idea of conditioning on pertinent statistics, define a criterion that is a function of the treatment assignment and pre-treatment covariates:

$$\phi(\mathbf{W}, \mathbf{X}) = \begin{cases} 1 & \text{if } \mathbf{W} \text{ is an acceptable treatment assignment} \\ 0 & \text{if } \mathbf{W} \text{ is not an acceptable treatment assignment.} \end{cases} \quad (7)$$

This notation mimics that of [21], who use $\phi(\mathbf{W}, \mathbf{X})$ to define treatment assignments that are desirable for an experimental design, and that of [3], who were the first to introduce such notation for randomization tests. The unconditional randomization test in Section 2.2 inherently defines $\phi(\mathbf{W}, \mathbf{X}) = 1$ if $\sum_{i=1}^N W_i = N_T$ and 0 otherwise. In general, conditional randomization tests involve sampling from the conditional distribution $P(\mathbf{W}|\phi(\mathbf{W}, \mathbf{X}) = 1)$ rather than the unconditional distribution $P(\mathbf{W})$ in Section 2.2.

Hennessy et al. [10] focus on $\phi(\mathbf{W}, \mathbf{X})$ that indicate some specified degree of categorical covariate balance. Assume there are covariate strata $s = 1, \dots, S$ specified by the researcher such that each unit belongs to only one stratum, and define $c_i = s$ if the i th unit belongs to the s th stratum. The strata may be defined using all of the covariates or some subset of them. Then, Hennessy et al. [10] define the criterion $\phi(\mathbf{W}, \mathbf{X})$ as¹

$$\phi_s(\mathbf{W}, \mathbf{X}) = \begin{cases} 1 & \text{if } \sum_{i:c_i=s} W_i = N_{T,s}, \text{ for } s = 1, \dots, S \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

In other words, each stratum is treated as a completely randomized experiment. Hennessy et al. [10] assume that the conditional distribution $P(\mathbf{W}|\phi_s(\mathbf{W}, \mathbf{X}) = 1)$ is uniform, i. e.,

$$P(\mathbf{W}|\phi_s(\mathbf{W}, \mathbf{X}) = 1) = \begin{cases} \left(\prod_{s=1}^S \binom{N_s}{N_{T,s}} \right)^{-1} & \text{if } \sum_{i:c_i=s} W_i = N_{T,s}, \text{ for } s = 1, \dots, S \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

¹ Hennessy et al. [10] use slightly different notation, instead defining a balance function $B(\mathbf{W}, \mathbf{X})$ and condition on the balance function being equal to some prespecified b . The more general notation that uses $\phi(\mathbf{W}, \mathbf{X})$ will become helpful in our discussion of continuous covariate balance.

Random samples from $P(\mathbf{W}|\phi_s(\mathbf{W}, \mathbf{X}) = 1)$ can be obtained by randomly permuting the observed treatment assignment \mathbf{W}^{obs} within the covariate strata $s = 1, \dots, S$. Once a test statistic is specified, the conditional randomization test follows the three-step procedure in Section 2.1, but using $P(\mathbf{W}|\phi_s(\mathbf{W}, \mathbf{X}) = 1)$ instead of $P(\mathbf{W})$.

Hennessy et al. [10] showed via simulation that this conditional randomization test using the test statistic $\hat{\tau}_{sd}$ is more powerful than the unconditional randomization test in Section 2.2 using $\hat{\tau}_{sd}$. Furthermore, they found that this conditional randomization test using $\hat{\tau}_{sd}$ is comparable to the unconditional randomization test using the post-stratification test statistic

$$\hat{\tau}_{ps} = \sum_{s=1}^S \frac{N_s}{N} \hat{\tau}_{sd}(s), \quad (10)$$

where $\hat{\tau}_{sd}(s)$ is the estimator $\hat{\tau}_{sd}$ within stratum s [20].

Note that the set of possible treatment assignments \mathbf{W} must be large enough to perform a powerful randomization test. For example, if $|\mathbf{W}| < 20$, then it is impossible to obtain a randomization test p -value less than 0.05. It may be surprising that conditional randomization tests can be more powerful than unconditional randomization tests, because the former utilizes fewer treatment assignments than the latter. However, these fewer treatment assignments are more relevant to the observed treatment assignment in terms of covariate balance, which leads to more powerful inference, as discussed by works such as [24] and [10].

When the criterion $\phi(\mathbf{W}, \mathbf{X})$ is defined as in (8), $|\mathbf{W}| = \prod_{s=1}^S \binom{N_s}{N_{T,s}}$, which is typically large. Furthermore, assuming that $P(\mathbf{W}|\phi(\mathbf{W}, \mathbf{X}) = 1)$ is uniform, random samples from this distribution can be obtained directly, and thus implementation of the conditional randomization test is straightforward. However, this approach is less straightforward when \mathbf{X} contains non-categorical covariates, because \mathbf{X} is no longer composed of strata where there are treatment and control units in each stratum. One option is to coarsen \mathbf{X} into strata and then use the conditional randomization test of [10]. Instead of throwing away information via coarsening, we propose a criterion $\phi(\mathbf{W}, \mathbf{X})$ that incorporates covariate balance for non-categorical covariates. We define $\phi(\mathbf{W}, \mathbf{X})$ such that $|\mathbf{W}|$ is large enough while still sufficiently conditioning on covariate balance. Furthermore, as we discuss below, random samples from $P(\mathbf{W}|\phi(\mathbf{W}, \mathbf{X}) = 1)$ will no longer be equivalent to random permutations of \mathbf{W}^{obs} ; thus, we develop an algorithm to obtain random samples from $P(\mathbf{W}|\phi(\mathbf{W}, \mathbf{X}) = 1)$.

3 A conditional randomization test for the case of non-categorical covariates

The conditional randomization test discussed in Section 2.3 is equivalent to a permutation test within S strata. This is analogous to analyzing a completely randomized experiment as if it were a blocked randomized experiment. We follow this intuition by proposing a conditional randomization test that is analogous to analyzing a completely randomized experiment as if it were a rerandomized experiment, where the rerandomization scheme incorporates a general form of covariate balance.

Rerandomization involves randomly allocating units to treatment and control until a certain level of pre-specified covariate balance is achieved. Thus, rerandomization requires specifying a metric for covariate balance. We first consider an omnibus measure of covariate balance and the corresponding conditional randomization test. We then extend this conditional randomization test to flexibly incorporate multiple measures of covariate balance, rather than a single omnibus measure, which we find yields more powerful randomization tests.

3.1 Conditional randomization test using an omnibus measure of covariate balance

The most common covariate balance metric used in the rerandomization literature is the Mahalanobis distance [19], which is defined as

$$M \equiv (\bar{\mathbf{X}}_T - \bar{\mathbf{X}}_C)^T [\text{cov}(\bar{\mathbf{X}}_T - \bar{\mathbf{X}}_C)]^{-1} (\bar{\mathbf{X}}_T - \bar{\mathbf{X}}_C) \quad (11)$$

$$= \frac{N_T N_C}{N} (\bar{\mathbf{X}}_T - \bar{\mathbf{X}}_C)^T [\text{cov}(\mathbf{X})]^{-1} (\bar{\mathbf{X}}_T - \bar{\mathbf{X}}_C) \quad (12)$$

where $\bar{\mathbf{X}}_T$ and $\bar{\mathbf{X}}_C$ are p -dimensional vectors of the covariate means in the treatment and control groups, respectively, and $\text{cov}(\mathbf{X})$ is the sample covariance matrix of \mathbf{X} , which is fixed across randomizations. The derivation for the equality in (12) can be found in [21]. Note that $\bar{\mathbf{X}}_T - \bar{\mathbf{X}}_C = \frac{\mathbf{X}^T \mathbf{W}}{\sum_{i=1}^N W_i} - \frac{\mathbf{X}^T (\mathbf{1} - \mathbf{W})}{\sum_{i=1}^N 1 - W_i}$, and so M is stochastic through \mathbf{W} .

We focus on using the Mahalanobis distance for our conditional randomization test because of its widespread use in measuring covariate balance for non-categorical covariates. Note that the Mahalanobis distance is an omnibus measure for balance among the individual covariates as well as their interactions (see, e. g., [30]). Following [10], we define a criterion $\phi(\mathbf{W}, \mathbf{X})$ such that:

1. It is asymmetric in treatment and control.²
2. It conditions on the covariate balance being similar to the observed balance for a particular randomization.

To fulfill these two desires, we consider the following criterion for our conditional randomization test:

$$\phi_{b_L, b_U}(\mathbf{W}, \mathbf{X}) = \begin{cases} 1 & \text{if } b_L \leq M^{obs} \leq b_U \text{ and } \text{sign}(\bar{\mathbf{X}}_{Tj} - \bar{\mathbf{X}}_{Cj}) = \text{sign}(\bar{\mathbf{X}}_{Tj}^{obs} - \bar{\mathbf{X}}_{Cj}^{obs}) \quad \forall j = 1, \dots, p \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

The equality of signs for all covariate mean differences addresses the first item above—in particular, it recognizes whether the treatment or control group has higher covariate values—while the bounds (b_L, b_U) address the second item.

The criterion (13) only considers randomizations that correspond to covariate balance similar to the observed M . Restricting M to be within the bounds (b_L, b_U) is analogous to stratifying the Mahalanobis distance and restricting M to be in the same stratum as the observed M . Now we outline two procedures for selecting (b_L, b_U) for our conditional randomization test.

3.1.1 How to choose the bounds (b_L, b_U)

To gain some intuition for how to choose the bounds, note that the interval (b_L, b_U) should be narrow enough around the observed M such that the corresponding \mathbf{W} sufficiently conditions on the observed covariate balance, but also the interval should be wide enough such that a powerful randomization test can still be performed. For example, consider the most narrow interval possible, when $b_L = b_U = M^{obs}$. In this case, there may be only a single randomization such that $M = M^{obs}$ (i. e., $|\mathbf{W}| = 1$) and thus our conditional randomization test completely loses its power, even though it is fully conditioning on the observed covariate balance.

We will consider two ways to pick (b_L, b_U) , presented as Procedures 1 and 2 below. Procedure 1 selects the bounds unconditionally of M^{obs} , while Procedure 2 does the same conditional on M^{obs} . In Section 3.3 we

² In particular, we would like the criterion to be able to distinguish between assignments \mathbf{W} where treated units have higher covariate values and \mathbf{W} where control units have higher covariate values. As discussed in [10], this can be useful information to condition on during a randomization test. In contrast, the Mahalanobis distance is symmetric in treatment and control. Asymmetry allows us to condition on the direction of the mean imbalance between treatment and control, in addition to the degree of covariate overlap between the two groups (as measured by the Mahalanobis distance).

establish that Procedure 1 yields a valid randomization test, and we also discuss the extent to which Procedure 2 yields a valid randomization test.

Procedure 1 for Selecting (b_L, b_U) : Bin the Mahalanobis Distance

1. Approximate the *sign-constrained* randomization distribution of the Mahalanobis distance by generating randomizations $\mathbf{w}^{(1)}, \dots, \mathbf{w}^{(D)}$ such that $\text{sign}(\bar{\mathbf{X}}_{T,j} - \bar{\mathbf{X}}_{C,j}) = \text{sign}(\bar{\mathbf{X}}_{T,j}^{obs} - \bar{\mathbf{X}}_{C,j}^{obs}) \forall j = 1, \dots, p$, and computing the corresponding $M^{(1)}, \dots, M^{(D)}$.
2. Before observing M^{obs} , bin the aforementioned randomization distribution into C categories. Denote the cutoff points for these C bins as m_1, \dots, m_{C+1} , where $m_1 \leq \dots \leq m_{C+1}$ and $m_1 = 0$ and $m_{C+1} = \infty$.
3. After observing M^{obs} , set $b_L = m_c$ and $b_U = m_{c+1}$ for the $c \in \{1, \dots, C\}$ such that $b_L \leq M^{obs} \leq b_U$.

Procedure 2 for Selecting (b_L, b_U) : Build a Neighborhood around M^{obs}

1. Approximate the *sign-constrained* randomization distribution of the Mahalanobis distance by generating randomizations $\mathbf{w}^{(1)}, \dots, \mathbf{w}^{(D)}$ such that $\text{sign}(\bar{\mathbf{X}}_{T,j} - \bar{\mathbf{X}}_{C,j}) = \text{sign}(\bar{\mathbf{X}}_{T,j}^{obs} - \bar{\mathbf{X}}_{C,j}^{obs}) \forall j = 1, \dots, p$, and computing the corresponding $M^{(1)}, \dots, M^{(D)}$.
2. Specify an acceptance probability $p_a \in (0, 1]$ that denotes the proportion of the aforementioned randomization distribution to be included in (b_L, b_U) .
3. After observing M^{obs} , let \mathcal{M}_L be the set of $\frac{Dp_a}{2}$ Mahalanobis distances that are immediately below M^{obs} , and let \mathcal{M}_U be the set of $\frac{Dp_a}{2}$ Mahalanobis distances that are immediately above M^{obs} . Then, set $b_L = \min \mathcal{M}_L$ and $b_U = \max \mathcal{M}_U$.
 - If there are fewer than $\frac{Dp_a}{2}$ Mahalanobis distances immediately below M^{obs} , set \mathcal{M}_L as the set of all Mahalanobis distances below M^{obs} , and set \mathcal{M}_U as the set of Mahalanobis distances immediately above M^{obs} such that $|\mathcal{M}_L| + |\mathcal{M}_U| = Dp_a$.
 - If there are fewer than $\frac{Dp_a}{2}$ Mahalanobis distances immediately above M^{obs} , set \mathcal{M}_U as the set of all Mahalanobis distances above M^{obs} , and set \mathcal{M}_L as the set of Mahalanobis distances immediately below M^{obs} such that $|\mathcal{M}_L| + |\mathcal{M}_U| = Dp_a$.

Procedure 1 categorizes the Mahalanobis distance and then sets (b_L, b_U) according to the category that M^{obs} falls into. Procedure 2 sets (b_L, b_U) according to the Mahalanobis distances that are immediately around M^{obs} , such that Dp_a of the Mahalanobis distances $M^{(1)}, \dots, M^{(D)}$ are contained in (b_L, b_U) , with M^{obs} being the median of (b_L, b_U) (except for the two corner cases noted in the final step of Procedure 2). Furthermore, one can use rejection sampling to generate the randomizations in Step 1 of Procedures 1 and 2: generate a complete randomization $\mathbf{w} \sim P(\mathbf{W})$, where $P(\mathbf{W})$ is defined in (1), and only keep \mathbf{w} if the sign constraint is fulfilled by \mathbf{w} . In the simulation study discussed in Section 4, we focus on Procedure 2, because it ensures that the hypothetical randomizations $\mathbf{w}^{(1)}, \dots, \mathbf{w}^{(D)}$ used during the conditional randomization test are the randomizations most similar to the observed one in terms of covariate balance.

3.1.2 Rejection-sampling approach for performing the conditional randomization test

The conditional randomization test proceeds according to the three-step procedure in Section 2.1 after b_L and b_U are specified and the criterion (13) is defined. While we assume that $P(\mathbf{W} | \phi_{b_L, b_U}(\mathbf{W}, \mathbf{X}) = 1)$ is uniformly distributed, random samples from this conditional distribution no longer correspond to random permutations of \mathbf{W}^{obs} as in the unconditional randomization test in Section 2.2 or the conditional randomization test in Section 2.3. Similar to how the randomizations in Step 1 of Procedures 1 and 2 can be generated, we propose a simple rejection-sampling algorithm to generate a random draw from $P(\mathbf{W} | \phi_{b_L, b_U}(\mathbf{W}, \mathbf{X}) = 1)$:

1. Generate a random draw \mathbf{w} from $P(\mathbf{W})$ defined in (1).
2. Accept \mathbf{w} if $\phi_{b_L, b_U}(\mathbf{w}, \mathbf{X}) = 1$; otherwise, repeat Step 1.

Note that, as p_a gets smaller, it will be more computationally intensive to generate random samples from $P(\mathbf{W}|\phi_{b_L, b_U}(\mathbf{W}, \mathbf{X}) = 1)$, but it corresponds to more precisely conditioning on the observed covariate balance. If generating random samples from $P(\mathbf{W}|\phi_{b_L, b_U}(\mathbf{W}, \mathbf{X}) = 1)$ via rejection-sampling is computationally intensive, one can use an alternative approach proposed by [3], which uses importance-sampling to approximate randomization test p -values at a lower computational cost than rejection-sampling.

In Section 4 we show via simulation that this conditional randomization test is more powerful than the standard unconditional randomization test, because the former conditions on a measure of covariate balance. However, the criterion (13) uses an omnibus measure of covariate balance, which may not sufficiently condition on the observed randomization if the number of covariates p is large. We now extend this procedure to more precisely condition on the observed covariate balance for a given randomization by incorporating multiple measures of covariate balance. We show in Section 4 that this extension results in a further gain in statistical power.

3.2 Conditional randomization test using multiple measures of covariate balance

Consider $t = 1, \dots, T$ tiers (or sets) of covariates that are of interest as specified by the researcher. Let $\mathbf{X}^{(t)} \equiv (\mathbf{X}_1^{(t)}, \dots, \mathbf{X}_{k_t}^{(t)})$ denote the covariates in tier t , where each covariate only appears in one of the T tiers. Then, define

$$M^{(t)} \equiv \frac{N_T N_C}{N} (\bar{\mathbf{X}}_T^{(t)} - \bar{\mathbf{X}}_C^{(t)})^T [\text{cov}(\mathbf{X}^{(t)})]^{-1} (\bar{\mathbf{X}}_T^{(t)} - \bar{\mathbf{X}}_C^{(t)}) \quad (14)$$

as the Mahalanobis distance for the covariates in tier t . This setup of dividing covariates into tiers is similar to [22], who developed a rerandomization framework that forces each $M^{(t)}$ to be sufficiently small by design. Note that the setup in Section 3.1 corresponds to $T = 1$ tiers.

Our proposed conditional randomization test follows a procedure similar to that in Section 3.1, but within each tier t . Define the criterion

$$\phi^{(t)}(\mathbf{W}, \mathbf{X}) = \begin{cases} 1 & \text{if } b_{L_t} \leq M^{(t)} \leq b_{U_t} \text{ and } \text{sign}(\bar{\mathbf{X}}_{j,T} - \bar{\mathbf{X}}_{j,C}) = \text{sign}(\bar{\mathbf{X}}_{j,T}^{obs} - \bar{\mathbf{X}}_{j,C}^{obs}) \quad \forall j = 1, \dots, k_t \\ 0 & \text{otherwise,} \end{cases} \quad (15)$$

for some lower and upper bounds b_{L_t} and b_{U_t} for each tier t . Then, define the overall criterion

$$\phi_T(\mathbf{W}, \mathbf{X}) = \prod_{t=1}^T \phi^{(t)}(\mathbf{W}, \mathbf{X}) \quad (16)$$

The bounds (b_{L_t}, b_{U_t}) are chosen separately for each tier using the procedure discussed in Section 3.1.1. This requires choosing an acceptance probability p_{a_t} for each tier. Because a smaller p_{a_t} corresponds to more stringent conditional inference, tiers with covariates that are believed to be most relevant to the outcomes should be assigned smaller p_{a_t} . However, recall that smaller p_{a_t} corresponds to more computational time required to obtain draws from $P(\mathbf{W}|\phi_T(\mathbf{W}, \mathbf{X}) = 1)$ via our rejection-sampling algorithm discussed in Section 3.1.2.

3.3 The validity of conditional randomization tests

A test is valid if $P(p \leq \alpha | H_0) \leq \alpha$, where H_0 is the Sharp Null Hypothesis and p is the calculated p -value. In our context, p is a function of the observed assignment and the testing procedure, and the probability is taken over the true assignment mechanism with the potential outcomes held fixed. For our conditional tests, the p -value is calculated as the probability of observing a test statistic more extreme than the observed one across randomizations \mathbf{w} such that $\phi_T(\mathbf{w}, \mathbf{X}) = 1$, for a specific $\phi_T(\mathbf{w}, \mathbf{X})$ determined by the observed assignment and covariates. Thus, the validity of our conditional randomization test depends on the criterion $\phi_T(\mathbf{W}, \mathbf{X})$, which—as shown in (15)—is defined by the bounds in each tier and the covariate sign constraints. In

Section 3.1.1, Procedure 1 defines the bounds before randomization, whereas Procedure 2 defines the bounds based on \mathbf{W}^{obs} after randomization. This latter case induces complications to establishing validity that we believe have not been previously discussed in the literature. In what follows, we discuss why exact validity may not necessarily hold for the conditional randomization test that uses Procedure 2, and establish validity for the test that uses Procedure 1.

Define \mathcal{B} as the set of possible bounds and \mathcal{S} as the set of possible covariate signs across all randomizations, and define $\mathbb{W}_{b,s}$ as the set of all randomizations that lead to particular bounds $\mathbf{b} \in \mathcal{B}$ and signs $\mathbf{s} \in \mathcal{S}$. The collection of $\mathbb{W}_{b,s}$ partition \mathbb{W} into non-overlapping sets.³ The overall probability of our conditional randomization test falsely rejecting the null can then be decomposed as

$$P(p \leq \alpha | H_0) = \sum_{b \in \mathcal{B}} \sum_{s \in \mathcal{S}} P(p \leq \alpha | H_0, \mathbf{W} \in \mathbb{W}_{b,s}) P(\mathbf{W} \in \mathbb{W}_{b,s}), \quad (17)$$

Given the above, a sufficient condition for establishing validity is that $P(p \leq \alpha | H_0, \mathbf{W} \in \mathbb{W}_{b,s}) \leq \alpha$ for all $\mathbf{b} \in \mathcal{B}$ and $\mathbf{s} \in \mathcal{S}$.

A given b and s pair specify a specific conditioning function $\phi_T(\mathbf{W}, \mathbf{X})$. Let $\mathbb{W}_\phi \equiv \{\mathbf{W} : \phi_T(\mathbf{W}, \mathbf{X}) = 1\}$ be the set of randomizations satisfying a given function $\phi_T(\cdot)$. Then, our calculated p -value, conditioned on an observed randomization, consequent $\phi_T(\cdot)$, and outcome will be

$$p \equiv \sum_{\mathbf{w} \in \mathbb{W}_\phi} \mathbb{I}(|t(Y(\mathbf{w}), \mathbf{w}, \mathbf{X})| \geq |t^{obs}|) P(\mathbf{W} = \mathbf{w} | \mathbf{W} \in \mathbb{W}_\phi), \quad (18)$$

where $t^{obs} \equiv t(\mathbf{y}^{obs}, \mathbf{W}^{obs}, \mathbf{X})$.

Under the null, \mathbf{y}^{obs} and \mathbf{X} are both invariant to random assignment, making our test statistic solely a function of \mathbf{W} . Under the null, then, let U_ϕ be a random variable whose distribution is that of $|t(\mathbf{y}^{obs}, \mathbf{w}, \mathbf{X})|$, where \mathbf{w} is uniformly distributed across the elements of \mathbb{W}_ϕ , and let $U_{b,s}$ be analogously defined for $\mathbb{W}_{b,s}$. (Note that $\mathbf{W}^{obs} \in \mathbb{W}_{b,s}$, because the realized \mathbf{b} and \mathbf{s} are specified by \mathbf{W}^{obs} .) Now consider our conditional probability $P(p \leq \alpha | H_0, \mathbf{W} \in \mathbb{W}_{b,s})$ for some specific b and s . Given this conditioning, our original test statistic is distributed as $U_{b,s}$. Regardless of the observed value of our test statistic, we have that our reference distribution will be U_ϕ , for our given $\phi_T(\cdot)$. Thus, our p -value, conditioned on our original randomization giving us our given b, s pair will then be the upper tail of our reference distribution, calculated as

$$1 - F_{U_\phi}(U_{b,s}), \quad (19)$$

where $F_{U_\phi}(\cdot)$ is the cumulative distribution function of U_ϕ . Here U_ϕ is a function of b and s , given the potential outcomes and covariates.

Typically, validity of a randomization test is proven by arguing that p -values of the form (19) are uniformly distributed by applying the probability integral transform (for an example of this method of proof, see [10, Section 2]). When Procedure 1 is used to select the bounds, $\mathbb{W}_\phi = \mathbb{W}_{b,s}$; i. e., all of the assignments used in the conditional randomization test are the same assignments that would lead to the realized \mathbf{b} and \mathbf{s} . Therefore, $U_{b,s}$ and U_ϕ have the same distribution under Procedure 1, and validity immediately follows from (19). However, because Procedure 2 specifies the bounds as a neighborhood around M^{obs} , the conditional randomization test under Procedure 2 uses randomizations that may not have led to the realized \mathbf{b} and \mathbf{s} . As a result, \mathbb{W}_ϕ and $\mathbb{W}_{b,s}$ will differ, and $U_{b,s}$ and U_ϕ will not necessarily have the same distribution. Consequently, our conditional randomization test that uses Procedure 2 for selecting the bounds is not necessarily valid. Nonetheless, in Section 4 and the Appendix, we find that our conditional randomization test using Procedure 2 is empirically valid under a wide variety of scenarios. This in part stems from the centering of our reference distributions around the test statistics; by contrast, if we had always selected distributions less extreme than the observed, we could induce invalidity. We leave investigating when validity formally holds when randomization test p -values are of the form (19) for two differing distributions as a promising line for future research.

³ Recall that \mathbb{W} is the set of treatment assignments with positive probability; see Section 2.

4 Simulation study: Conditional and unconditional performance of conditional and unconditional randomization tests

We now conduct a simulation study to explore the statistical power of the unconditional randomization test from Section 2.2, our conditional randomization tests from Sections 3.1 and 3.2, and another conditional randomization test inspired by Coarsened Exact Matching (CEM). CEM was designed for observational studies to find a subset of treatment and control units that match exactly on a coarsened covariate space [13, 14]. Even though CEM was developed for observational studies and not randomization tests, we include it in our comparison because—as we noted at the end of Section 2—coarsening \mathbf{X} into strata is one option for performing a conditional randomization test in the face of continuous covariates. Thus, it is the most natural test to compare to our conditional randomization test.

In what follows, we find that our conditional randomization test using $\hat{\tau}_{sd}$ is more powerful than the unconditional randomization test using $\hat{\tau}_{sd}$ as well as the CEM-based tests. Furthermore, we find that our test is comparable to an unconditional randomization test using a regression-based test statistic. Finally, we find that the conditional randomization tests and an unconditional randomization test using a regression-based test statistic are all valid both unconditionally and conditional on the data, whereas the unconditional randomization test that uses an unadjusted test statistic is only valid unconditionally.

4.1 Simulation procedure

Consider $N = 100$ units whose potential outcomes are generated according to the following model:

$$\begin{aligned} Y_i(0)|X_i &= \beta(0.1X_{i1} + 0.2X_{i2} + 0.3X_{i3} + 0.4X_{i4}) + \epsilon_i, \quad i = 1, \dots, 100 \\ Y_i(1) &= Y_i(0) + \tau \end{aligned} \quad (20)$$

where X_{i1} , X_{i2} , X_{i3} , X_{i4} , and ϵ_i are independently and randomly sampled from a $N(0, 1)$ distribution. The parameters β and τ take on values $\beta \in \{0, 1.5, 3\}$ and $\tau \in \{0, 0.1, \dots, 1\}$ across simulations. As β increases, the covariates become more associated with the outcome; as τ increases, the treatment effect increases and thus should be easier to detect.

Once the potential outcomes are generated, units are randomized to treatment and control such that $N_T = 50$ units receive treatment and $N_C = 50$ units receive control; in other words, units are assigned according to the completely randomized assignment mechanism (1). This is repeated such that 1,000 randomizations are produced using the same fixed potential outcomes. In the Appendix we also consider an unbalanced design where an unequal number of units are assigned to treatment and control; however, the results for that scenario are largely the same as the results presented here, where $N_T = N_C = 50$.

For each randomization, five separate randomization tests were performed:

1. **Unconditional Randomization Test:** The procedure described in Section 2.2, using the test statistic $\hat{\tau}_{sd}$ given in (3).
2. **Conditional Randomization Test:** The procedure described in Section 3.2 using the criterion (16), which requires specifying the number of covariate tiers T and acceptance probability p_a in Procedure 2 for selecting the bounds within each tier. We consider number of tiers $T \in \{1, 2, 4\}$ and acceptance probabilities $p_a \in \{0.1, 0.25, 0.5\}$. The $T = 1$ case corresponds to the procedure described in Section 3.1.⁴ For each tier, we choose (b_L, b_U) by setting all tier-level acceptance probabilities p_{a_t} to be equal, where the overall acceptance probability is $p_a = \prod_{t=1}^T p_{a_t}$.⁵ We use the test statistic $\hat{\tau}_{sd}$.

⁴ For $T = 2$, the first two covariates are in one tier while the last two are in another tier. For $T = 4$, all covariates are in their own tier.

⁵ Note that this equality holds only because the covariates in each tier are independent. Thus, $p_{a_t} = (p_a)^{1/T}$ for all tiers $t = 1, \dots, T$.

3. **Unconditional Randomization (with model-adjusted test statistic):** The procedure described in Section 2.2, using the test statistic $\hat{\tau}_{int}$, which is defined as the estimated coefficient for W_i from the linear regression of Y_i on W_i , \mathbf{x}_i , and $W_i(\mathbf{x}_i - \bar{\mathbf{X}})$. This test statistic was discussed in [18], but within the context of Neymanian inference rather than randomization tests.
4. **Coarsened Exact Matching (Prespecified Groups):** Each covariate is coarsened into G groups according to the quantiles of $N(0, 1)$, thus coarsening the \mathbb{R}^4 covariate space into G^4 strata. Then, to perform a randomization test using $\hat{\tau}_{sd}$ as a test statistic, \mathbf{W}^{obs} is permuted many times within each stratum. We consider number of groups $G \in \{2, 3, 4\}$.
5. **Coarsened Exact Matching (Automatic Groups):** The same as the previous test, but the G groups are chosen automatically by the R function `cem`.

Our motivation for including the third randomization test in our comparison is that Hennessy et al. [10] found that their conditional randomization test using $\hat{\tau}_{sd}$ is comparable to the unconditional randomization test using $\hat{\tau}_{ps}$ defined in (10), and that $\hat{\tau}_{ps}$ is equivalent to $\hat{\tau}_{int}$ when covariates are categorical [18]. We also considered our conditional randomization test using $\hat{\tau}_{int}$ instead of $\hat{\tau}_{sd}$, and found that the power results for that test are essentially the same as those for the unconditional randomization test using $\hat{\tau}_{int}$; we relegate those results to the Appendix.

Meanwhile, the last two procedures utilize CEM. These conditional tests are identical to the test of Hennessy et al. [10] using the assignment mechanism (9), where the strata are chosen via CEM. In the CEM (Prespecified Groups) procedure, the strata are specified according to the quantiles of the known distributions of the covariates. Meanwhile, in the CEM (Automatic Groups) procedure, the strata are automatically specified according to Sturges' rule, which uses the range of the covariates and is the default option in the `cem` R package [12]. Details about this procedure and other automated procedures in the context of CEM are discussed in [14].

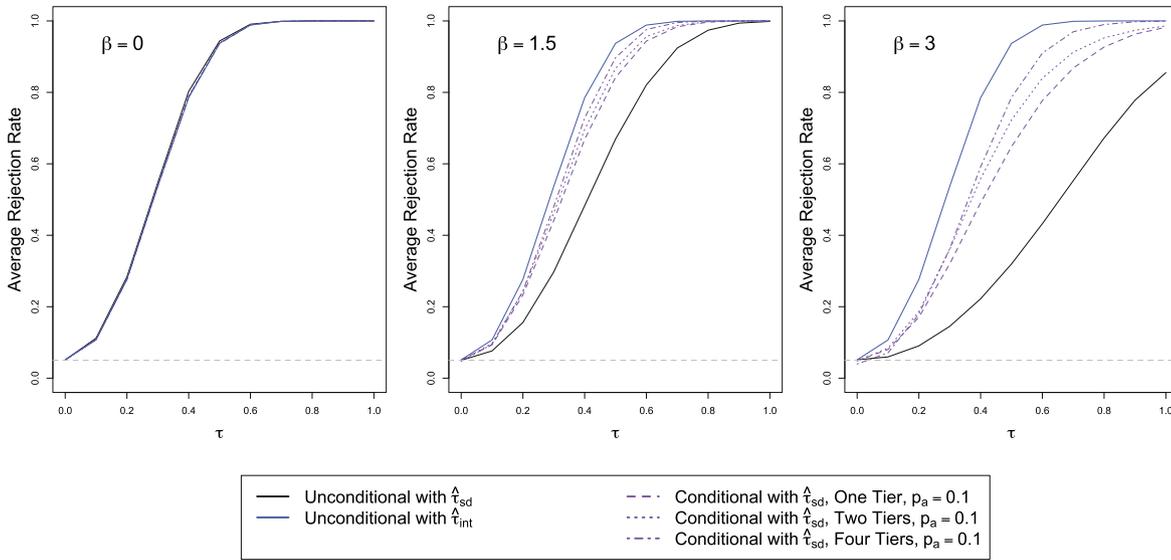
4.2 Simulation results: Unconditional performance

We first assess statistical power, which corresponds to how often each randomization test rejected Fisher's Sharp Null across the 1,000 complete randomizations when $\tau > 0$. The average rejection rates for the unconditional randomization tests using $\hat{\tau}_{sd}$ and $\hat{\tau}_{int}$ as well as our conditional randomization test are presented in Figure 1 for various values of β and τ . Figure 1a displays results for a fixed acceptance probability $p_a = 0.1$ and different numbers of tiers, while Figure 1b displays results for a fixed number of $T = 4$ tiers and different acceptance probabilities.

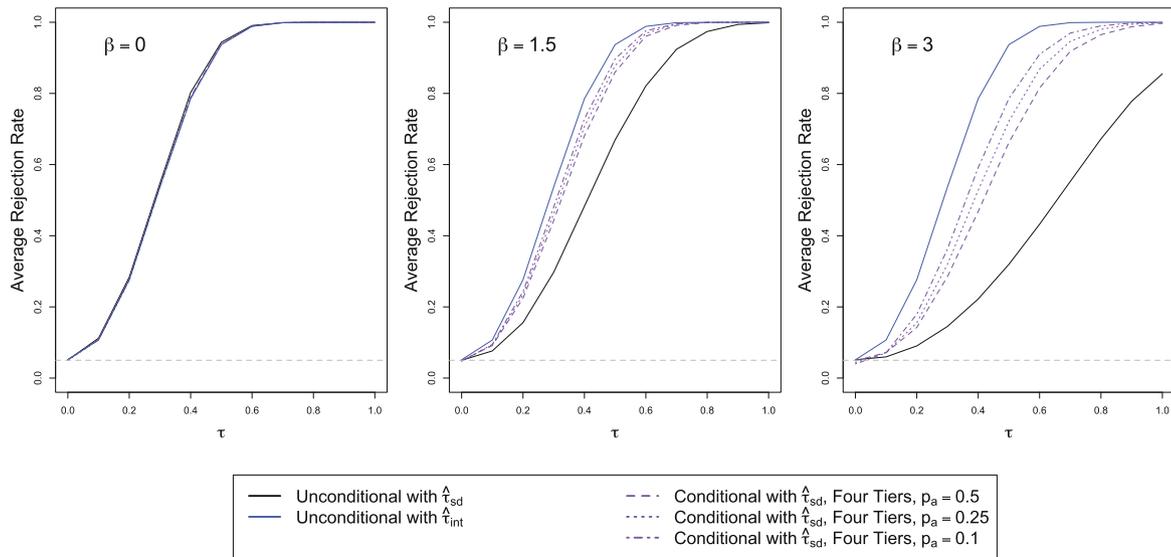
Several conclusions can be made from Figure 1. First, when $\beta = 0$ (i. e., when the covariates are not associated with the outcome), all of the randomization tests are essentially equivalent. When the covariates are associated with the outcome, our conditional randomization test is more powerful than the unconditional randomization test that uses $\hat{\tau}_{sd}$. Furthermore, the power of our conditional randomization test increases as the acceptance probability p_a decreases and/or the number of tiers increases; this is expected: lower p_a and higher T corresponds to more stringent conditioning.

Figure 1a suggests that practitioners can increase power by increasing the number of tiers without any additional computational cost (i. e., without decreasing the acceptance probability). Furthermore, Figure 1b suggests that the *additional* gain in power decreases as p_a decreases, which echoes the observation made by Li et al. [17] in the rerandomization literature that the marginal benefit to decreasing p_a decreases as p_a decreases. Analogous figures for the $T = 1$ and $T = 2$ cases are in the Appendix; by comparing those figures with Figure 1b, it can be seen that the additional gain in power from decreasing p_a increases as T increases. This observation emphasizes the benefits of conditioning on multiple measures of covariate balance rather than a single omnibus measure. Further discussion on this point is in the Appendix.

Meanwhile, Figure 1 also shows that the unconditional randomization test using $\hat{\tau}_{int}$ was more powerful than all of the conditional and unconditional randomization tests using $\hat{\tau}_{sd}$. However, as p_a gets smaller and



(a) For the conditional randomization test, different tiers and a fixed $p_a = 0.1$ acceptance probability.

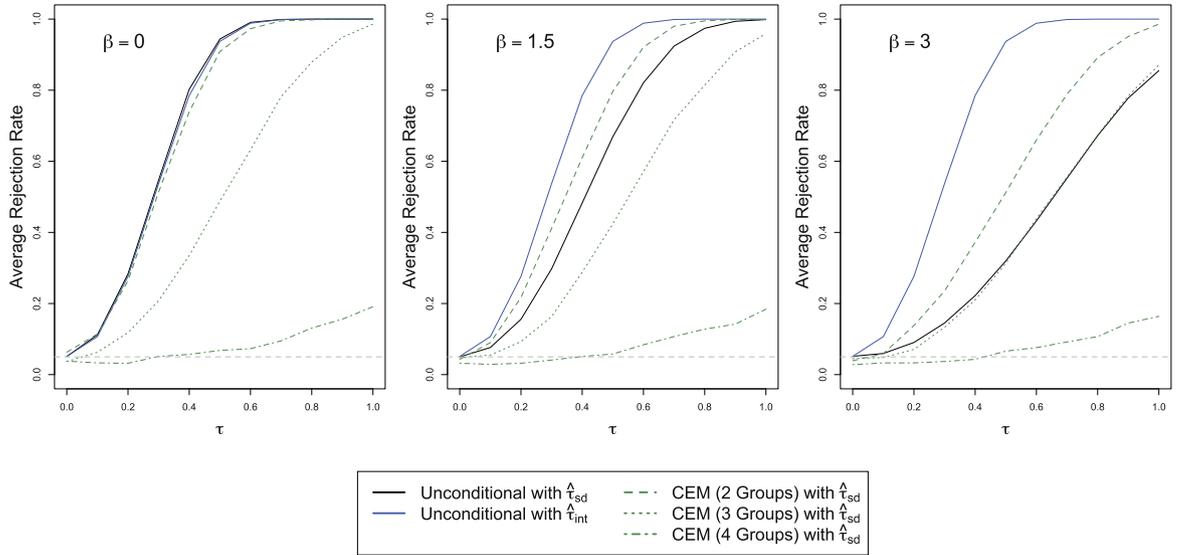


(b) For the conditional randomization test, different acceptance probabilities and a fixed $T = 4$ tiers.

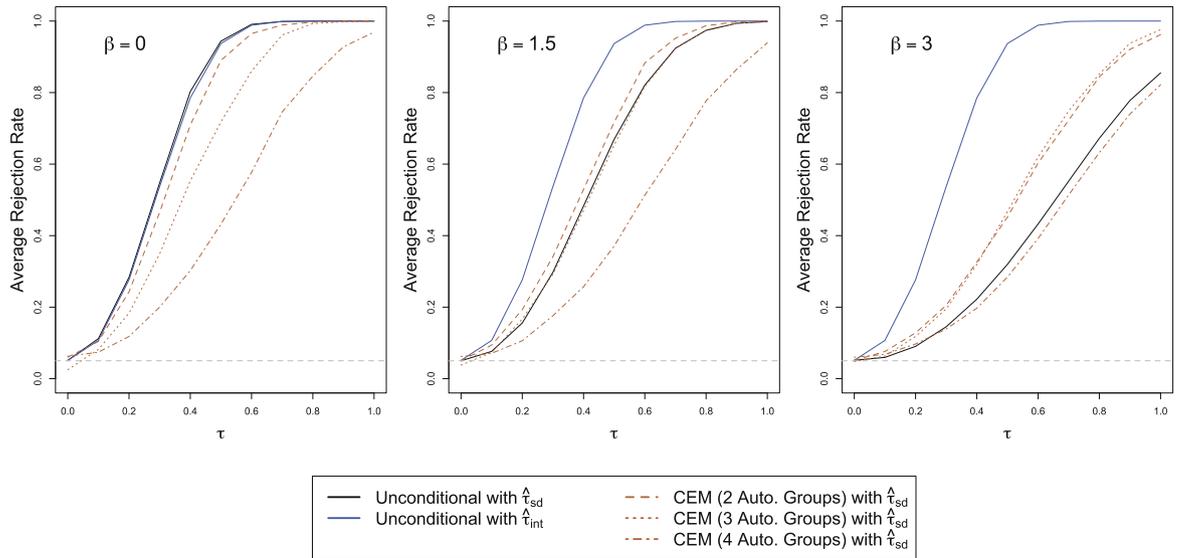
Figure 1: Average rejection rate (power) of Fisher’s Sharp Null for the unconditional randomization test using $\hat{\tau}_{sd}$ and $\hat{\tau}_{int}$, as well as our conditional randomization test using $\hat{\tau}_{sd}$.

T gets larger—i. e., as conditioning becomes more stringent—the performance of our conditional randomization test appears to approach that of the unconditional randomization test that uses $\hat{\tau}_{int}$. This reinforces the claim made by Li et al. [17] that—in a Neymanian inference context— $\hat{\tau}_{int}$ under complete randomization is equivalent to $\hat{\tau}_{sd}$ under very stringent rerandomization. However, Li et al. [17] made this claim about the rerandomization scheme that uses an omnibus measure of covariate balance; our findings suggest that this claim should be qualified to state that the equivalence between $\hat{\tau}_{int}$ under complete randomization and $\hat{\tau}_{sd}$ under rerandomization holds when the rerandomization scheme incorporates separate measures of balance for each covariate used in $\hat{\tau}_{int}$, rather than a single omnibus measure.

Here, $\hat{\tau}_{int}$ is correctly specified because the potential outcomes are generated from a linear model, and one may wonder how the unconditional randomization test using $\hat{\tau}_{int}$ performs when this model is misspecified.



(a) CEM (Prespecified Groups).



(b) CEM (Automatic Groups).

Figure 2: Average rejection rate (power) of Fisher’s Sharp Null for the CEM (Prespecified Groups) and CEM (Automatic Groups) procedures for various number of groups for each covariate. Also shown are results for the unconditional tests using $\hat{\tau}_{sd}$ and $\hat{\tau}_{int}$, which are the same results from Figure 1.

We consider this in the Appendix and obtain findings very similar to those presented here. In particular, for the simulation settings considered, we find that it is still beneficial to use the unconditional randomization test with $\hat{\tau}_{int}$ or our conditional randomization test with $\hat{\tau}_{sd}$ in the misspecified case as long as the functions of the covariates used in the regression to construct $\hat{\tau}_{int}$ are correlated with the response; when they are not correlated, these tests are essentially equivalent. In the Appendix we also explore a variety of additional simulation scenarios—when the covariates have positive and negative effects on the potential outcomes, when there are heterogeneous treatment effects, and when the covariates are not normally distributed—and we again find results that are very similar to the results presented here. This suggests that these results hold under a wide variety of scenarios.

Now we assess the performance of the conditional randomization tests that use CEM. Figure 2 shows the average rejection rate of Fisher’s Sharp Null for the CEM-based randomization tests. To anchor our comparison, Figure 2 also includes the results for the unconditional randomization tests using $\hat{\tau}_{sd}$ and $\hat{\tau}_{int}$ (i. e., the same results presented in Figure 1). When the covariate space for each covariate is coarsened into $G = 2$ groups, these conditional randomization tests are more powerful than the unconditional randomization test using $\hat{\tau}_{sd}$ when $\beta > 0$, although they are not as powerful as our conditional randomization test or the unconditional randomization test using $\hat{\tau}_{int}$. When the number of groups for each covariate is increased, the power of the conditional randomization tests tend to decrease, especially for the CEM procedure that specifies strata according to the quantiles of the known covariate distributions. At first this finding may be surprising, because more groups should correspond to more stringent conditioning and thus possibly higher power. However, as the number of groups increases, there are fewer strata with both treatment and control units, and thus more units are discarded and there are fewer possible randomizations used during the randomization test. For example, for the CEM (Prespecified Groups) procedure, when there were $G = 2$ groups, on average 4 of the 100 units were discarded across the 1,000 randomizations; when $G = 3$, on average 54 of the 100 units were discarded; and when $G = 4$, on average 88 of the 100 units were discarded. In the most extreme case, if we let the number of groups go to infinity—i. e., not coarsen the continuous covariate space at all—there would not be any treatment and control units with the same covariate values, and thus all units would be discarded.

Meanwhile, there is not a clear winner between the CEM (Prespecified Groups) and CEM (Automatic Groups) procedures, although the CEM (Automatic Groups) procedure is not as severely underpowered for the $G = 4$ case as the CEM (Prespecified Groups) procedure. In their development of CEM, Iacus et al. [12, 13, 14] recommend researchers use context-specific knowledge for specifying strata rather than automated procedures, but it is unclear if this should be the recommendation when using CEM for conditional randomization tests. Indeed, the CEM (Prespecified Groups) procedure uses the most context-specific knowledge possible (the actual data-generating process for the covariates), but it does not necessarily perform as well as the automated procedure.

In summary, these findings suggest that it is beneficial to condition on forms of covariate balance that account for continuous covariates, rather than condition on a coarsened version of the continuous covariate space. Furthermore, our methodology allows researchers to condition on the data at hand in a way that increases the power of randomization tests, while coarsening the covariate space may lead to a lack of possible treatment assignments to perform a powerful randomization test.

4.3 Simulation results: Conditional performance

We next examine the performance of the five tests across randomizations that are particularly balanced or imbalanced. First, we generated the potential outcomes using model (20) with $\tau = 0$ (which corresponds to no treatment effect) and $\beta = 3$ (which corresponds to a strong association between the covariates and potential outcomes). Then, we generated 10,000 randomizations and divided these randomizations into 10 groups according to quantiles of the Mahalanobis distance. Thus, the first group consists of the 1,000 best randomizations according to the Mahalanobis distance, while the tenth group consists of the 1,000 worst randomizations. Now we consider whether the five randomization tests are valid (i. e., reject Fisher’s Sharp Null when it is true 5% of the time) for randomizations conditional on a particular level of covariate balance. Conditional validity assesses to what extent these tests are valid across randomizations that are similar to the observed randomization.

Figure 3 displays the average rejection rate of each randomization test for each of the 10 quantile groups of the Mahalanobis distance. For the CEM-based tests, we display results for $G = 2$ groups, because this resulted in the most power in Section 4.2. The conditional performance for higher groups are similar. Our conditional randomization test that uses $\hat{\tau}_{sd}$ and the unconditional randomization test that uses $\hat{\tau}_{int}$ both exhibit average rejection rates close to the 5% level across all quantile groups, which suggests that both tests are conditionally valid across randomizations of any particular balance level. The story is quite different for the unconditional randomization test that uses $\hat{\tau}_{sd}$: for low levels of covariate imbalance, the average rejection

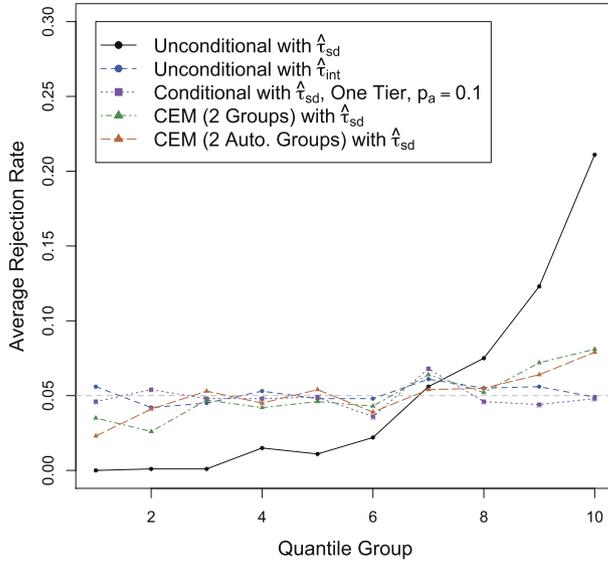


Figure 3: The rejection rate of the five randomization tests when Fisher’s Sharp Null Hypothesis is true (i. e., $\tau = 0$) and $\beta = 3$. Rejection rates are shown within each quantile group of the Mahalanobis distance, such that each quantile group corresponds to 1,000 randomizations.

rate is below the 5% level, while for high levels of covariate imbalance the average rejection rate is notably above the 5% level. These rejection rates average out to 5%—as can be seen in Figure 1—and thus the unconditional randomization test that uses $\hat{\tau}_{sd}$ is unconditionally valid, but—as can be seen in Figure 3—it is not conditionally valid conditional on a particular balance level. In particular, the false rejection rate for the unconditional randomization test that uses $\hat{\tau}_{sd}$ appears to be monotonically increasing in covariate imbalance, which is intuitive given that treatment effects will be increasingly confounded with covariate effects as covariate imbalance increases. Meanwhile, the false rejection rate for the CEM-based tests also appears to be monotonically increasing in covariate imbalance according to the Mahalanobis distance, but to a much less severe degree. This is likely because these tests condition on balance for a coarsened version of the covariate space instead of balance for the continuous covariate space as measured by the Mahalanobis distance. In short, they are conditionally valid for the coarsened covariate space but not the continuous covariate space.

In summary, statistically powerful randomization tests can be constructed by conditioning on covariate balance through the assignment mechanism or by using a model-adjusted test statistic; either option will result in a more powerful test than an unconditional randomization test that uses an unadjusted test statistic. We also find that our conditional randomization test using unadjusted test statistics or unconditional randomization tests using model-adjusted test statistics appear to be approximately equivalent, both across complete randomizations as well as across randomizations of a particular balance level. Furthermore, we find that our conditional randomization test that directly conditions on group-level balance for continuous covariates is more powerful than other conditional randomization tests that condition on a coarsened version of the covariate space. Finally, it is particularly important to condition on group-level covariate balance or use a model-adjusted test statistic to ensure validity across randomizations of a particular balance level, because covariate imbalances can break the conditional validity of unconditional randomization tests that use unadjusted test statistics.

5 Discussion and conclusion

Hennessy et al. [10] outlined a conditional randomization test that conditions on the covariate balance observed after an experiment has been conducted, and showed that these tests are more powerful than standard unconditional randomization tests and comparable to randomization tests that use model-adjusted estimators, such as the post-stratified estimator in [20]. However, Hennessy et al. [10] focused on the case when there

are only categorical covariates. Here we proposed a methodology for conducting a randomization test that conditions on a form of covariate balance that allows for non-categorical covariates.

Through simulation, we found that our conditional randomization test is more powerful than unconditional randomization tests that use unadjusted test statistics as well as other conditional randomization tests inspired by the observational study literature, and that it is approximately equivalent to an unconditional randomization test that uses a regression-based test statistic. We also found that the conditional randomization tests and the unconditional randomization tests that use adjusted test statistics appear valid conditional on the observed covariate balance; the more traditional unconditional randomization tests that use unadjusted test statistics, however, are clearly not.

The above findings hold under a variety of data-generating scenarios, such as ones with treatment effect heterogeneity or model misspecification. Most of the literature has focused on increasing the power of randomization tests through the choice of the test statistic; to our knowledge, we are the first to do the same through the choice of the assignment mechanism for the general case when non-categorical covariates are present. Furthermore, we found evidence that these two avenues for constructing randomization tests are approximately equivalent in terms of statistical power. Thus, our methodology can achieve the power of model-adjustment while preserving the transparency of an unadjusted treatment effect estimate, thereby taking advantage of the benefits of both adjusted and unadjusted estimators as discussed by Lin [18]. Relatedly, we also discussed how this finding suggests connections between regression-based estimators after complete randomization and unadjusted estimators after rerandomization, which refines observations previously made by Li et al. [17].

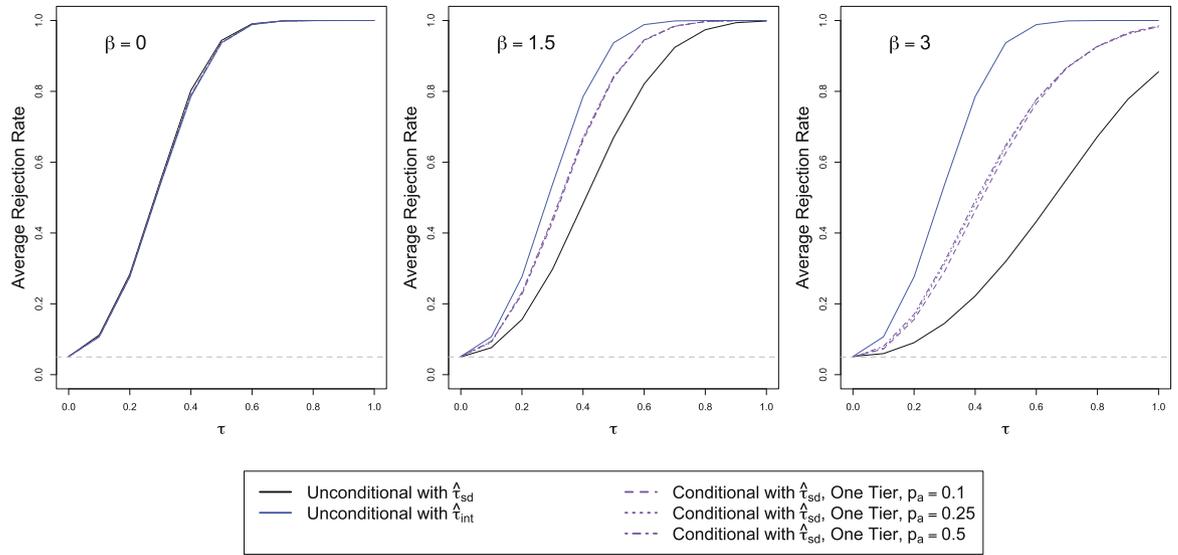
We focused on randomization tests for randomized experiments, but we believe that this work has implications beyond tests and experiments. Randomization tests can be inverted to yield confidence intervals for treatment effects [27, 16], and thus our method can go beyond testing the presence of a treatment effect. Some have criticized such randomization-based confidence intervals because they commonly make the assumption of a constant treatment effect for all units. However, recent works have suggested how to incorporate treatment effect heterogeneity in randomization tests (e. g., [6, 4]), and our work adds to this literature by suggesting how forms of covariate balance can be incorporated in randomization tests as well. An interesting line of future work would be to combine our conditional randomization test with these works to conduct randomization-based inference that incorporates both treatment effect heterogeneity and covariate balance.

Furthermore, most work on randomization tests for observational studies has focused on cases where only categorical covariates are present [24, 25, 26, 27]. Our work suggests a way to conduct randomization-based inference for observational studies when non-categorical covariates are present. However, because the assignment mechanism in an observational study is unknown, researchers need to determine when certain assignment mechanisms can be assumed within an observational study before conducting randomization-based inference. See [2] for a framework for how to conduct conditional randomization-based inference in this context.

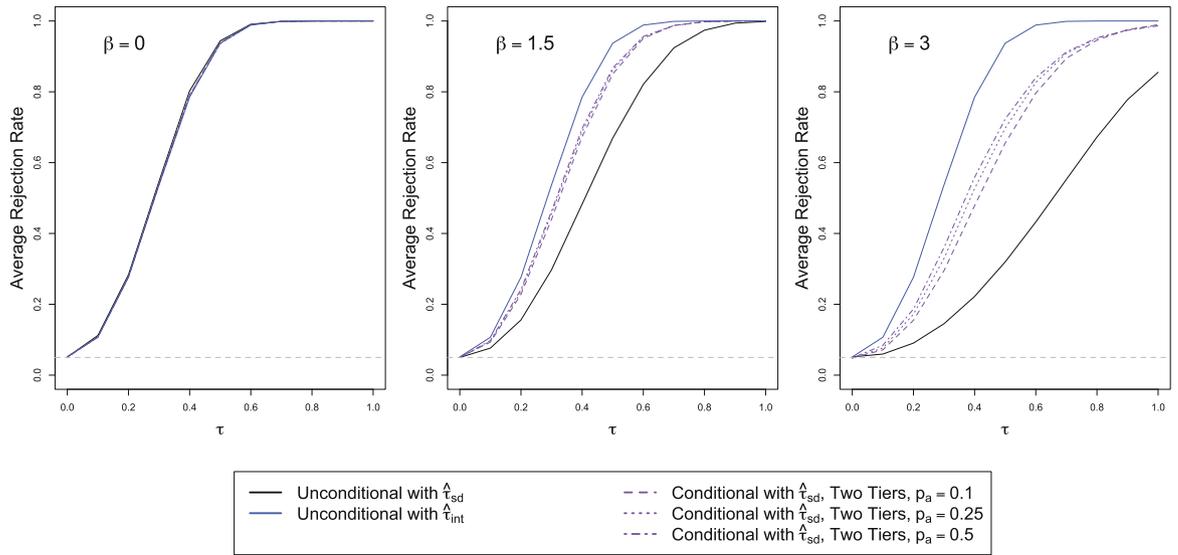
6 Appendix: Additional simulation results

Here we present further power results of randomization tests similar to those presented in Section 4. All of the following sections and figures discuss the average rejection rate of Fisher's Sharp Null for various randomization tests. In Section 6.1, we consider the same setup discussed in Section 4 and present results for our conditional randomization test for various acceptance probabilities and one or two tiers (instead of four tiers), as well as results for our conditional randomization test using the regression-adjusted test statistic $\hat{\tau}_{int}$ (instead of $\hat{\tau}_{sd}$). Then, in Sections 6.3 and 6.4 we consider other data-generating processes not explored in Section 4, including:

1. when some covariate effects are positive and some are negative,
2. when there is treatment effect heterogeneity,
3. when there are non-normal covariates,



(a) Power results of our conditional randomization test using one tier.



(b) Power results of our conditional randomization test using two tiers.

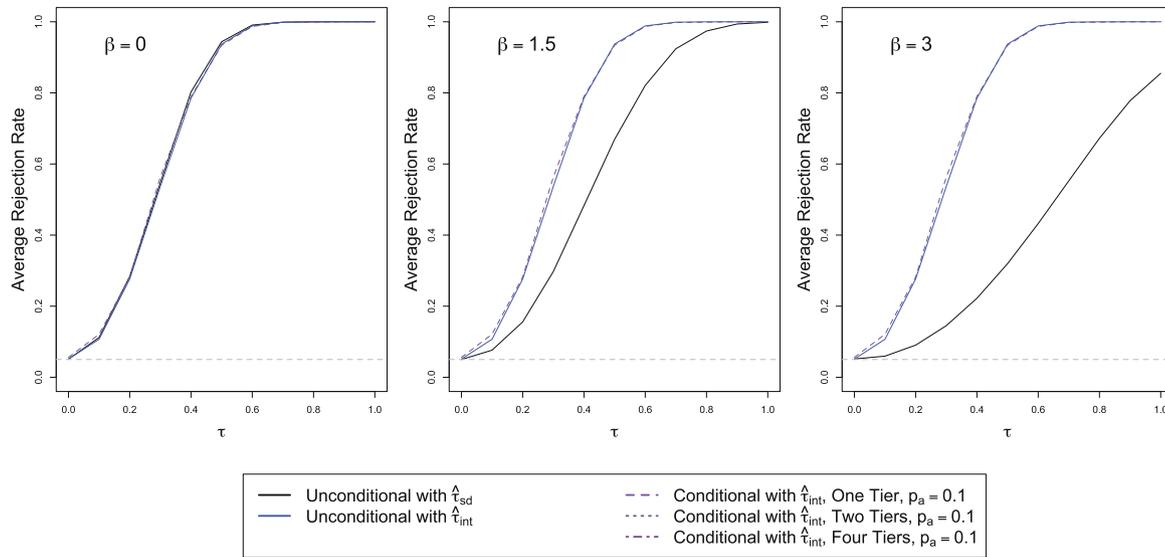
Figure 4: The rejection rate of the same tests discussed in Figure 1b, but for one or two tiers instead of four.

4. when the linear regression used in $\hat{\tau}_{int}$ is misspecified.

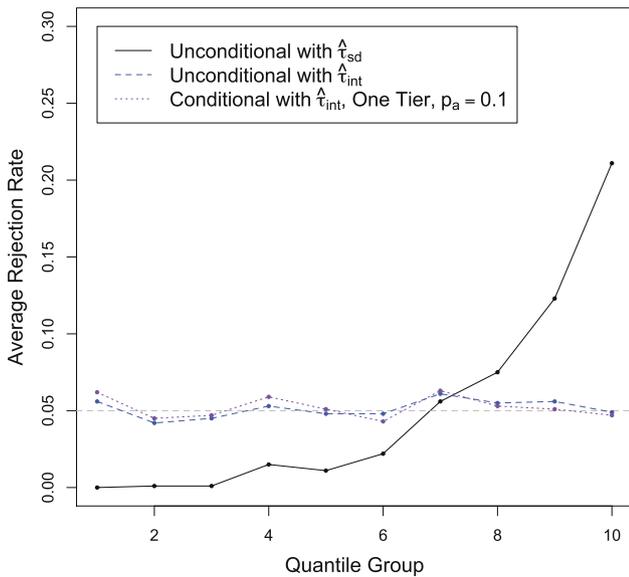
The results for the first three are quite similar to the results presented in Section 4, and so we discuss them together in Section 6.3. We discuss results for the misspecified case in Section 6.4.

6.1 Simulation results for one and two tiers and for conditional randomization using $\hat{\tau}_{int}$

Consider the same simulation setup as Section 4, where the potential outcomes for $N = 100$ units are generated using the model (20). In Section 4.2, we examined the power of our conditional randomization test



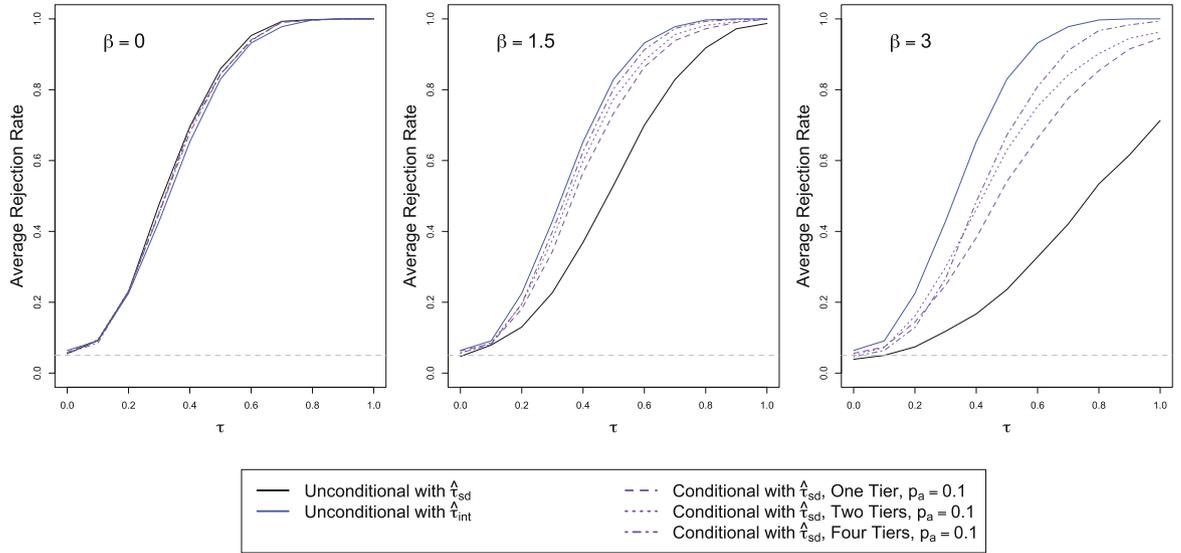
(a) Conditional randomization tests using $\hat{\tau}_{int}$ for various tiers and a fixed acceptance probability.



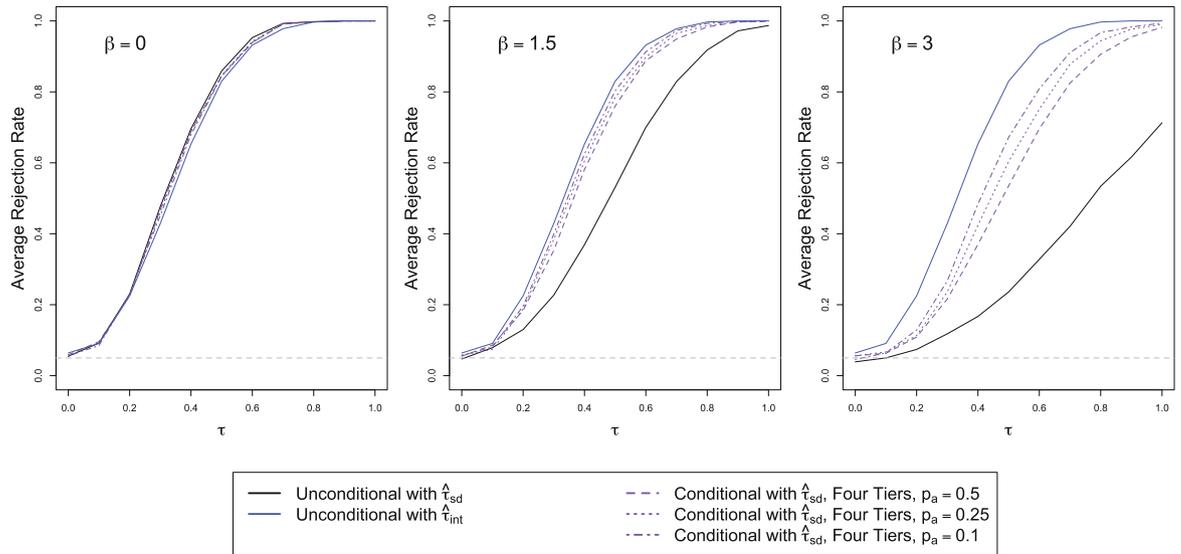
(b) Average rejection rate of each randomization test when Fisher’s Sharp Null Hypothesis is true. Rejection rates are shown within each quantile group of the Mahalanobis distance, such that each quantile group corresponds to 1,000 randomizations. Data were generated using (20) with $\tau = 0$ and $\beta = 3$, as in Section 4.3.

Figure 5: The unconditional and conditional performance of our conditional randomization test using $\hat{\tau}_{int}$. Figure 5a is analogous to Figure 1a; Figure 5b is analogous to Figure 3.

for various acceptance probabilities for a fixed number of four tiers. Figure 4 shows the same results for one and two tiers, respectively. In other words, Figure 4 is analogous to Figure 1b, but for one or two tiers instead of four. The results are quite similar to those presented in Figure 1b: the power of our conditional randomization test increases as the acceptance probability decreases. Furthermore, by comparing Figures 1b and 4, one can see that the additional benefit of decreasing the acceptance probability increases with the number of tiers. This emphasizes the benefit of conditioning on multiple measures of balance, rather than just a single measure.



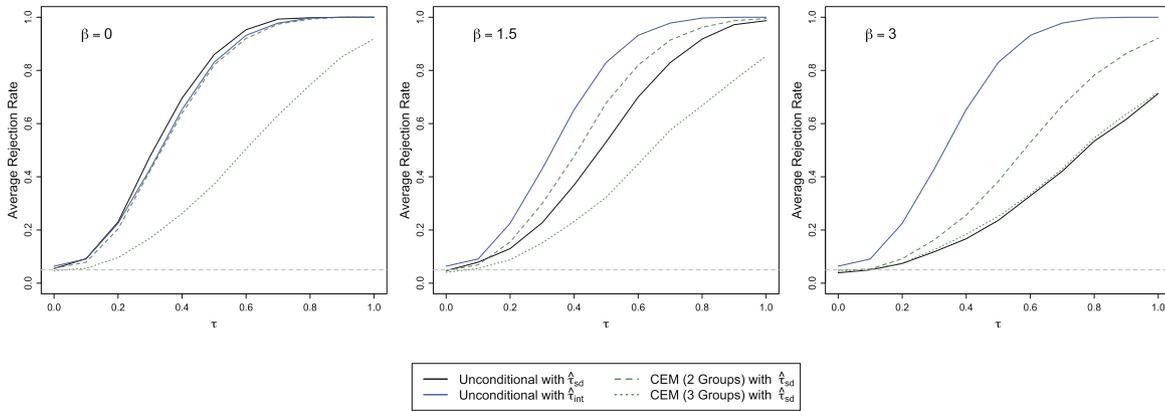
(a) For the conditional randomization test, different tiers and a fixed $p_a = 0.1$ acceptance probability.



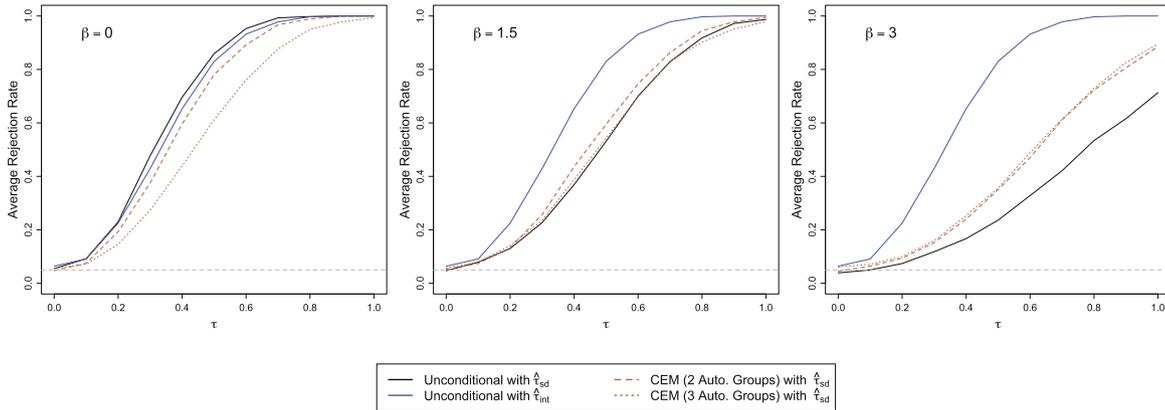
(b) For the conditional randomization test, different acceptance probabilities and a fixed $T = 4$ tiers.

Figure 6: The rejection rate of the same tests discussed in Figure 1, but for an unbalanced design where $N_T = 25$ and $N_C = 75$ instead of a balanced design where $N_T = N_C = 50$.

Furthermore, in Section 4 we focused on our conditional randomization test using the simple mean-difference test statistic $\hat{\tau}_{sd}$. Figure 5 presents the unconditional and conditional performance of our conditional randomization test using the regression-adjusted test statistic $\hat{\tau}_{int}$. In other words, Figures 5a and 5b are the same as Figures 1 and 3, respectively, except we use $\hat{\tau}_{int}$ instead of $\hat{\tau}_{sd}$ for our conditional randomization test. We find that the power results for our conditional randomization test using $\hat{\tau}_{int}$ are essentially the same as those using $\hat{\tau}_{sd}$, and thus there does not appear to be an additional benefit of using a conditional randomization distribution for the randomization test if a model-adjusted test statistic is used (or vice versa).



(a) CEM (Prespecified Groups).



(b) CEM (Automatic Groups).

Figure 7: Power results for the CEM (Prespecified Groups) and CEM (Automatic Groups) procedures for various number of groups for each covariate under the unbalanced design scenario where $N_T = 25$ and $N_C = 75$. Also shown are results for the unconditional tests using $\hat{\tau}_{sd}$ and $\hat{\tau}_{int}$, which are the same results from Figure 6.

6.2 Simulation results for unbalanced designs

Consider the same simulation setup as Section 4, where the potential outcomes for $N = 100$ units are generated using the model (20). In Section 4, we considered balanced designs, where an equal number of units are assigned to treatment and control (i. e., $N_T = N_C = 50$). Here we consider an unbalanced design, where $N_T = 25$ and $N_C = 75$. Otherwise, the simulation setup discussed here is identical to the one discussed in Section 4. The results for this unbalanced design scenario are essentially identical to the results discussed in Section 4.

Figure 6 shows the power results of (1) the unconditional randomization test using $\hat{\tau}_{sd}$; (2) the unconditional randomization test using $\hat{\tau}_{int}$; and (3) our conditional randomization test using $\hat{\tau}_{sd}$. In other words, Figure 6 is analogous to Figure 1, except the results are for an unbalanced design where $N_T = 25$ and $N_C = 75$ instead of a balanced design where $N_T = N_C = 50$. The power of all three tests are slightly lower for this case as compared to their power for the balanced design, but otherwise the results from Figure 6 are identical to those from Figure 1: Our conditional randomization test is more powerful than the unconditional randomization test using $\hat{\tau}_{sd}$, and the results of our conditional randomization test approach those of the unconditional randomization test using $\hat{\tau}_{int}$ when the number of tiers increases or the acceptance probability p_a decreases.

Meanwhile, Figure 7 shows the power results of the CEM-based randomization tests discussed in Section 4. In other words, Figure 7 is analogous to Figure 2, except the results are for the unbalanced design instead of the balanced design. For this unbalanced design scenario, we were only able to obtain results for $G = 2$ and $G = 3$ groups for the CEM-based tests, the reason being that there were less treated units in this

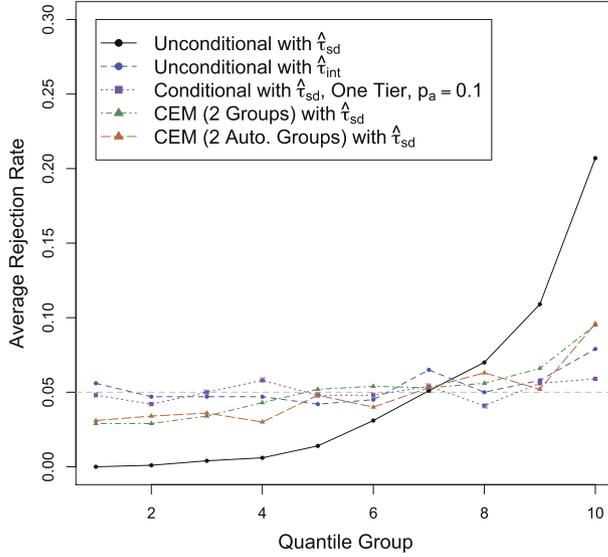


Figure 8: The rejection rate of the five randomization tests when Fisher’s Sharp Null Hypothesis is true (i. e., $\tau = 0$) and $\beta = 3$ for the unbalanced design scenario. Rejection rates are shown within each quantile group of the Mahalanobis distance, such that each quantile group corresponds to 1,000 randomizations.

unbalanced design, and thus less opportunities for CEM to find matches across many strata. This problem is the same as the issue that the CEM-based tests discard more and more units as the number of groups (or coarsened strata) G increases, as discussed in Section 4. This again emphasizes the benefit of conditioning on forms of covariate balance that account for continuous covariates, instead of conditioning on a coarsened version of the covariate space. Otherwise, the results from Figure 7 are identical to those from Figure 2: These CEM-based tests tend to be more powerful than the unconditional randomization test that uses $\hat{\tau}_{sd}$ but not as powerful as our conditional randomization test, and their power tends to decrease as G increases.

Similar to Section 4.3, we also examined the conditional performance of these randomization tests for this unbalanced design scenario. After the potential outcomes were generated from (20) for $\tau = 0$ and $\beta = 3$, we simulated 10,000 randomizations (where $N_T = 25$ and $N_C = 75$) and computed the Mahalanobis distance for each randomization. Then, we divided these randomizations into 10 groups according to the 10 quantiles of the 10,000 Mahalanobis distances. Figure 8 shows the rejection rate of each of the five randomization tests for each quantile group of the Mahalanobis distance. In other words, Figure 8 is analogous to Figure 3, except for the unbalanced design instead of the balanced design. The results are again largely the same as those presented in Section 4.3: The unconditional randomization test using $\hat{\tau}_{int}$ and the conditional randomization test using $\hat{\tau}_{sd}$ are conditionally valid across quantile groups, while the unconditional randomization test using $\hat{\tau}_{sd}$ is not conditionally valid and its rejection rate is monotonically increasing in covariate imbalance. Meanwhile, similar to Section 4.3, the false rejection rate for the CEM-based tests also appears to be monotonically increasing in covariate imbalance according to the Mahalanobis distance, but to a much less severe degree, suggesting that these tests are approximately conditionally valid.

6.3 Simulation results for alternative data-generating linear models

In Section 4, the potential outcomes were generated using the linear model (20) where all the covariates had positive effects on the outcomes, were unrelated to the treatment effect, and were normally distributed. Here we consider alternative linear models for the potential outcomes and compare power results for the unconditional randomization tests using $\hat{\tau}_{sd}$ and $\hat{\tau}_{int}$ as well as our conditional randomization test using $\hat{\tau}_{sd}$ for these alternative models. We examine the performance of the randomization tests under each of the following models:

– Positive/Negative Covariate Effects

$$\begin{aligned}
 Y_i(0)|X_i &= \beta(-0.1X_{i1} + 0.2X_{i2} + 0.3X_{i3} - 0.4X_{i4}) + \epsilon_i, \quad i = 1, \dots, 100 \\
 Y_i(1) &= Y_i(0) + \tau
 \end{aligned} \tag{21}$$

where $(X_{i1}, X_{i2}, X_{i3}, X_{i4}, \epsilon_i) \stackrel{iid}{\sim} \mathcal{N}_5(0, I_5)$.

– **Heterogeneous Treatment Effects**

$$\begin{aligned} Y_i(0)|X_i &= \beta(0.1X_{i1} + 0.2X_{i2} + 0.3X_{i3} + 0.4X_{i4}) + \epsilon_i, \quad i = 1, \dots, 100 \\ Y_i(1) &= Y_i(0) + \tau + \sigma_\tau Y_i(0) \end{aligned} \quad (22)$$

where $(X_{i1}, X_{i2}, X_{i3}, X_{i4}, \epsilon_i) \stackrel{iid}{\sim} \mathcal{N}_5(0, I_5)$. Following [6], we set $\sigma_\tau = 0.5$ to induce strong treatment effect heterogeneity.

– **Different Covariate Distributions**

$$\begin{aligned} Y_i(0)|X_i &= \beta(0.1X_{i1} + 0.2X_{i2} + 0.3X_{i3} + 0.4X_{i4}) + \epsilon_i, \quad i = 1, \dots, 100 \\ Y_i(1) &= Y_i(0) + \tau \end{aligned} \quad (23)$$

where $X_{i1} \sim N(0, 1)$, $X_{i2} \sim N(X_{i1}, 1)$, $X_{i3} \sim \text{Pois}(5)$, $X_{i4} \sim \text{Bern}(0.2)$, and $\epsilon_i \sim N(0, 1)$.

Similar to Section 4, the parameters β and τ take on values $\beta \in \{0, 1.5, 3\}$ and $\tau \in \{0, 0.1, \dots, 1\}$ across simulations for the above models.

Figure 9 shows the power results of the randomization tests when the potential outcomes were generated from the above models. Figure 9 is analogous to Figure 1, except the potential outcomes were generated from models (21), (22), or (23) instead of model (20) used in Section 4. The results are largely the same: The conditional randomization test is more powerful than the unconditional randomization test that uses the unadjusted test statistic $\hat{\tau}_{sd}$; furthermore, as the number of tiers increases, the conditional randomization test approaches the unconditional randomization test that uses the regression-adjusted test statistic.

Similar to Section 4.3, we also examined the conditional performance of the randomization tests when the potential outcomes were generated from the above models. After the potential outcomes were generated for $\tau = 0$ and $\beta = 3$ for each of the three models, we simulated 10,000 randomizations and computed the Mahalanobis distance for each randomization. Then, we divided these randomizations into 10 groups according to the 10 quantiles of the 10,000 Mahalanobis distances. Figure 10 shows the rejection rate of each randomization test for each quantile group for each of the three potential outcome models. Figure 10 is analogous to Figure 3, except the potential outcomes were generated from models (21), (22), or (23) instead of model (20). The results are again largely the same as those presented in Section 4.3: The unconditional randomization test using $\hat{\tau}_{int}$ and the conditional randomization test using $\hat{\tau}_{sd}$ are conditionally valid across quantile groups, while the unconditional randomization test using $\hat{\tau}_{sd}$ is not conditionally valid and its rejection rate is monotonically increasing in covariate imbalance. In short, Figures 9 and 10 suggest that the results found in Section 4 hold across many data-generating processes.

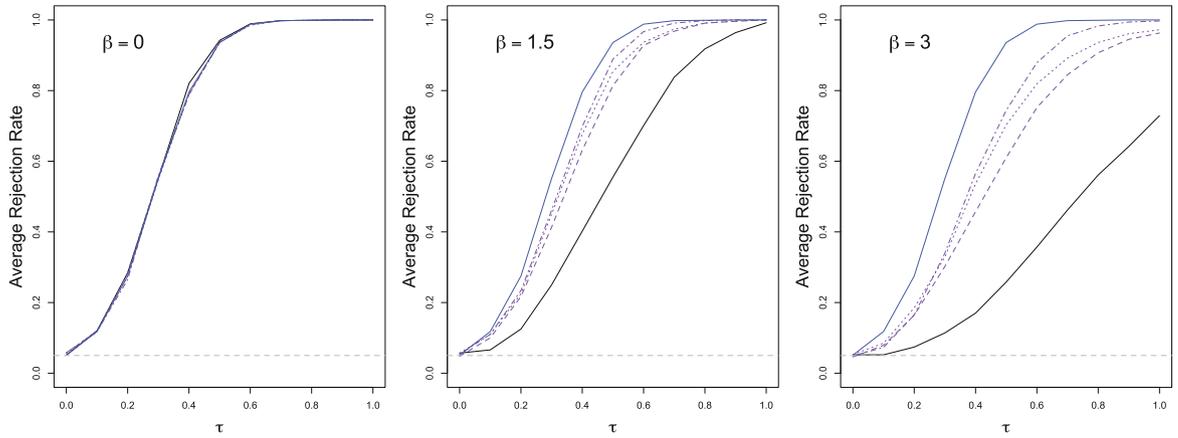
6.4 Simulation results for misspecified linear models

In the simulation study discussed in Section 4, the potential outcomes were generated from the linear model (20). We considered using the test statistic $\hat{\tau}_{int}$, which is defined as the estimated coefficient for W_i from the linear regression of Y_i on W_i , \mathbf{x}_i , and $W_i(\mathbf{x}_i - \bar{\mathbf{X}})$. Thus, $\hat{\tau}_{int}$ is a correctly specified model in the simulation setup presented in Section 4. We now consider cases when $\hat{\tau}_{int}$ is still defined as in Section 4 but the potential outcomes are generated from a nonlinear model, making the model $\hat{\tau}_{int}$ assumes misspecified. Consider $N = 100$ units whose potential outcomes are generated from one of the following models:

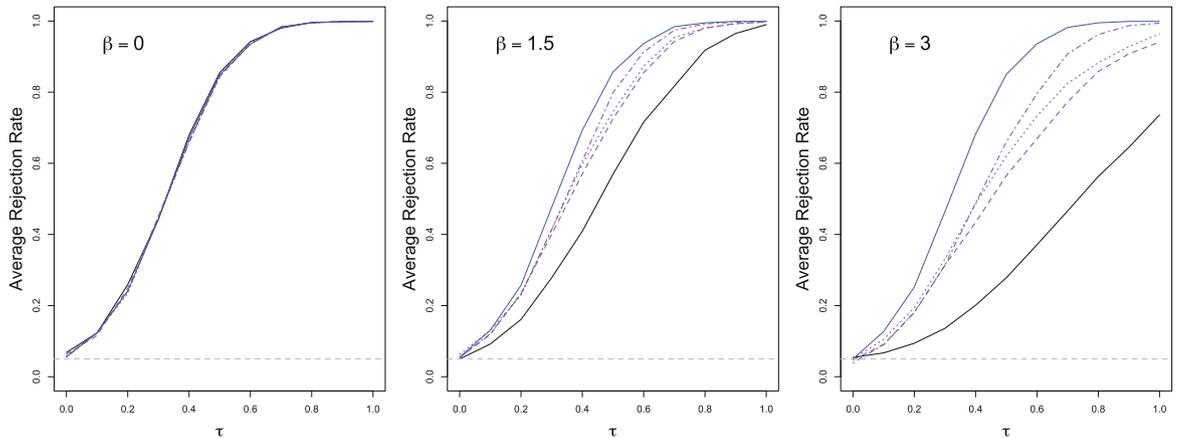
– **Model with Moderate Correlation**

$$\begin{aligned} Y_i(0)|X_i &= \beta(0.1X_{i1}^2 + 0.2X_{i2} + 0.3X_{i3}^2 + 0.4X_{i4}) + \epsilon_i, \quad i = 1, \dots, 100 \\ Y_i(1) &= Y_i(0) + \tau \end{aligned} \quad (24)$$

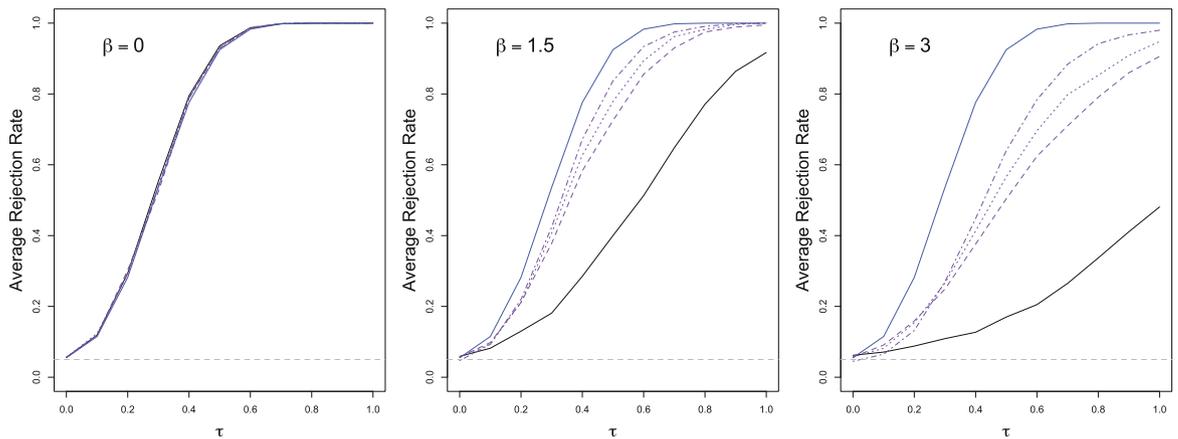
where $(X_{i1}, X_{i2}, X_{i3}, X_{i4}, \epsilon_i) \stackrel{iid}{\sim} \mathcal{N}_5(0, I_5)$.



(a) Potential outcomes generated from the Positive/Negative Covariate Effects model (21).



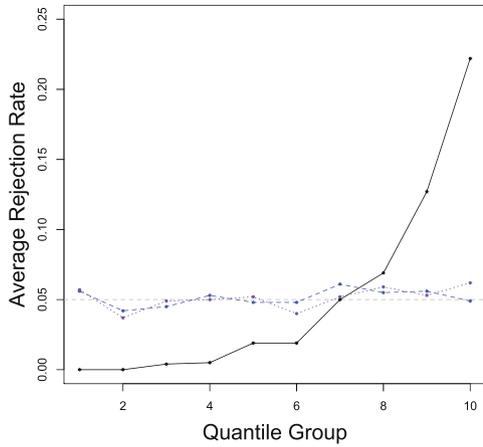
(b) Potential outcomes generated from the Heterogeneous Treatment Effects model (22).



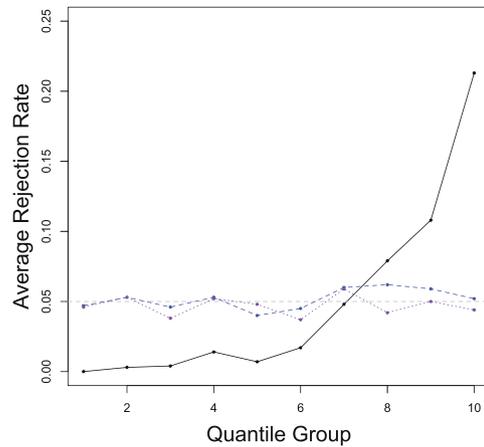
(c) Potential outcomes generated from the Different Covariate Distributions model (23).



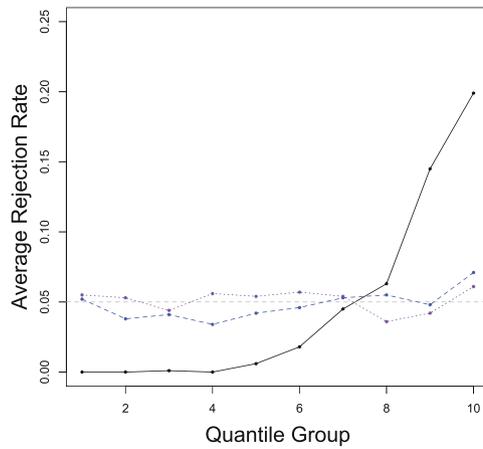
Figure 9: Average rejection rate for the unconditional randomization tests using $\hat{\tau}_{sd}$ and $\hat{\tau}_{int}$ as well as our conditional randomization test using $\hat{\tau}_{sd}$ for various tiers and a fixed acceptance probability, where the potential outcomes were generated from the Positive/Negative Covariate Effects model (21), Heterogeneous Treatment Effects model (22), or Different Covariate Distributions model (23).



(a) Positive/Negative Covariate Effects model (21).



(b) Heterogeneous Treatment Effects model (22).



(c) Different Covariate Distributions model (23).

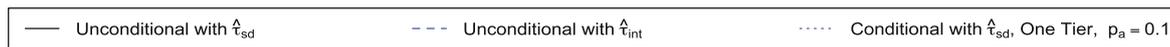


Figure 10: The rejection rate of the randomization tests within each quantile group of the Mahalanobis distance when the potential outcomes were generated from the Positive/Negative Covariate Effects model (21), Heterogeneous Treatment Effects model (22), or Different Covariate Distributions model (23).

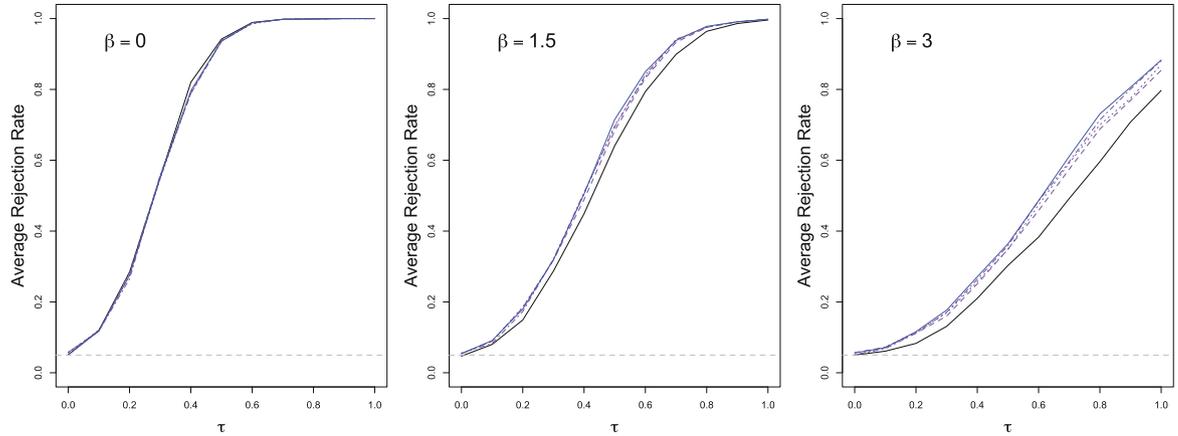
– **Model with No Correlation**

$$\begin{aligned}
 Y_i(0)|X_i &= \beta \left(0.1\sqrt{|X_{i1}|} + 0.2X_{i2}^2 + 0.3\sqrt{|X_{i3}|} + 0.4X_{i4}^2 \right) + \epsilon_i, \quad i = 1, \dots, 100 \\
 Y_i(1) &= Y_i(0) + \tau
 \end{aligned}
 \tag{25}$$

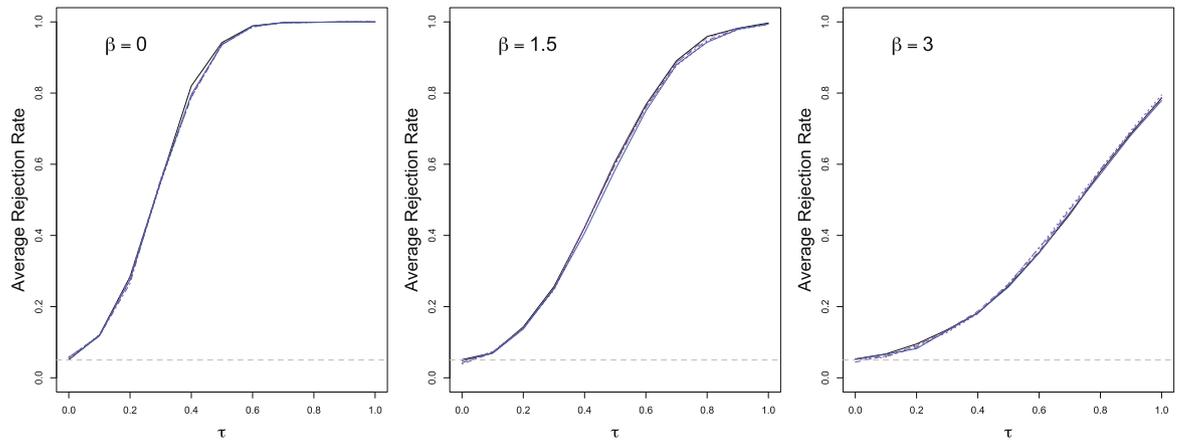
where $(X_{i1}, X_{i2}, X_{i3}, X_{i4}, \epsilon_i) \stackrel{iid}{\sim} \mathcal{N}_5(0, I_5)$.

Similar to Section 4, the parameters β and τ take on values $\beta \in \{0, 1.5, 3\}$ and $\tau \in \{0, 0.1, \dots, 1\}$ across simulations for the above models.

In the first model, there is a moderate correlation between the raw covariates and the potential outcomes: For the specific set of potential outcomes generated from (24) with $\beta = 3$ for the simulation, the empirical R^2 between $Y(0)$ and $(\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \mathbf{X}_4)$ was 0.33. Meanwhile, in the second model, there is no correlation between the raw covariates and the potential outcomes: For the specific set of potential outcomes generated from (25) with $\beta = 3$ for the simulation, the empirical R^2 was only 0.075. These cases differ from the case discussed in



(a) Potential outcomes generated from the Moderate Correlation model (24).



(b) Potential outcomes generated from the No Correlation model (25).



Figure 11: Average rejection rate of the unconditional randomization tests using $\hat{\tau}_{sd}$ and $\hat{\tau}_{int}$ as well as our conditional randomization test when the potential outcomes were generated from the Moderate Correlation model (24) or the No Correlation model (25).

Section 4, where the empirical R^2 was 0.82 and thus there was a strong correlation between the raw covariates and the potential outcomes.

Figure 11 shows the power results of the randomization tests when the potential outcomes were generated from the above models. The results for the Moderate Correlation case are similar to those presented in Section 4: The conditional randomization test is more powerful than the unconditional randomization test that uses $\hat{\tau}_{sd}$; furthermore, as the number of tiers increases, the conditional randomization test approaches the unconditional randomization test that uses $\hat{\tau}_{int}$. Meanwhile, for the No Correlation case, the power of all the tests appear to be essentially equivalent. These results suggest that there is a benefit of using our conditional randomization test or the unconditional randomization test with a regression-adjusted test statistic if there is a correlation between the covariates and the potential outcomes.

Similar to Section 4.3, we also examined the conditional performance of the randomization tests when the potential outcomes were generated from the Moderate Correlation and No Correlation models. Figure 12 shows the rejection rate of each randomization test for each quantile group for each potential outcome model,

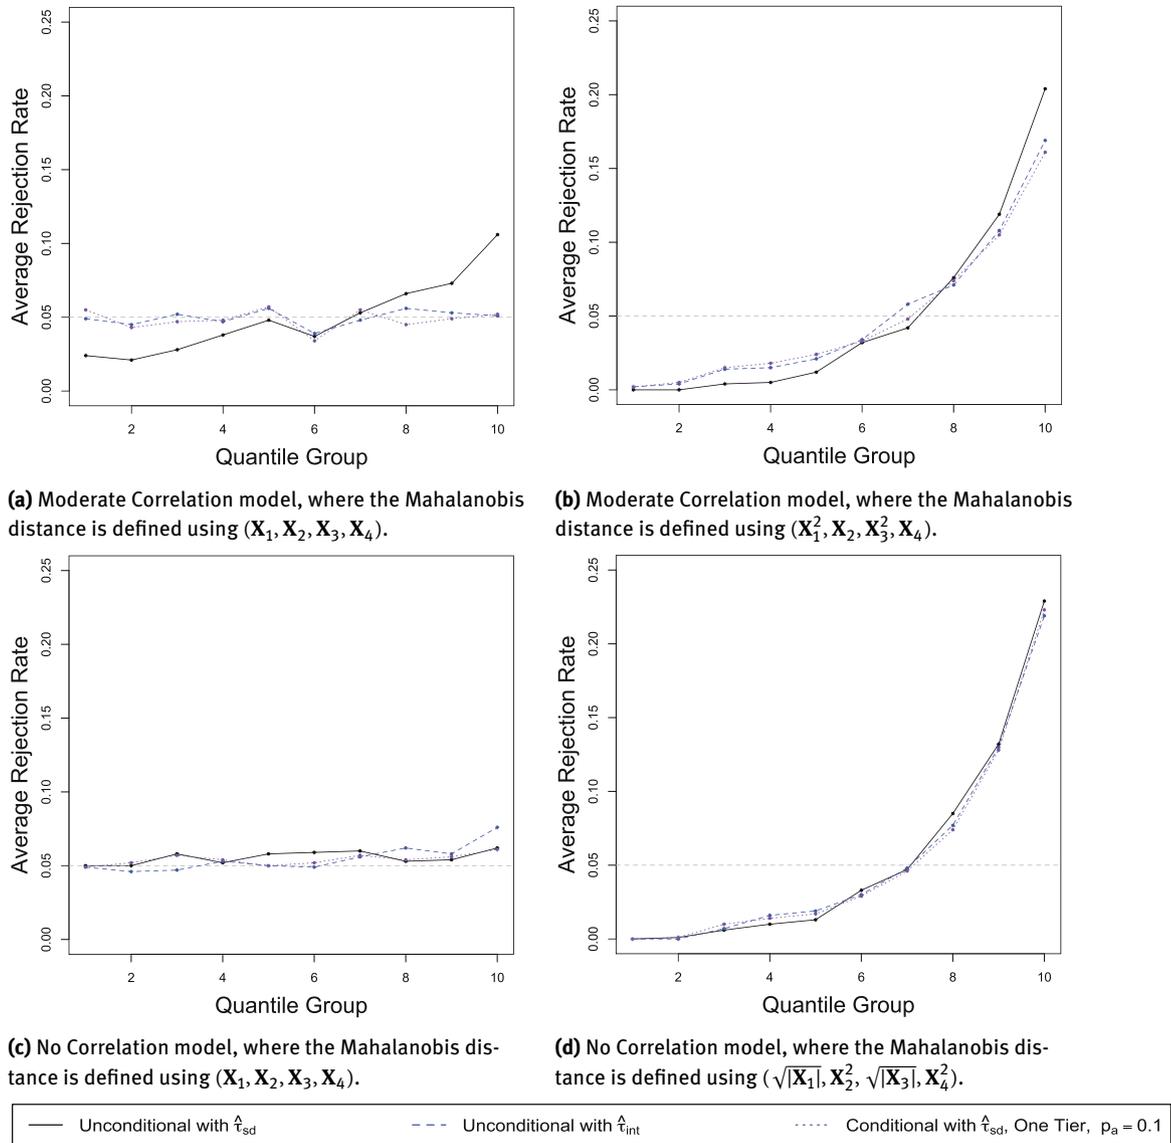


Figure 12: The rejection rate of the randomization tests within each quantile group of the Mahalanobis distance when the potential outcomes were generated from the Moderate Correlation model (24) or the No Correlation model (25). In Figures 12a and 12c, the Mahalanobis distance is defined using the raw covariates $(\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \mathbf{X}_4)$; in Figures 12b and 12d, the Mahalanobis distance is defined using the functions of the covariates that are linearly related with the potential outcomes for each model $((\mathbf{X}_1^2, \mathbf{X}_2, \mathbf{X}_3^2, \mathbf{X}_4)$ and $(\sqrt{|\mathbf{X}_1|}, \mathbf{X}_2^2, \sqrt{|\mathbf{X}_3|}, \mathbf{X}_4^2)$, respectively).

where we followed the same quantile-binning procedure as Section 4.3. In particular, in the left-hand plots of Figure 12, the Mahalanobis distance is defined using the raw covariates $(\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \mathbf{X}_4)$, whereas in the right-hand plots it is defined using the functions of the covariates that are linearly related to the potential outcomes, i. e., $(\mathbf{X}_1^2, \mathbf{X}_2, \mathbf{X}_3^2, \mathbf{X}_4)$ and $(\sqrt{|\mathbf{X}_1|}, \mathbf{X}_2^2, \sqrt{|\mathbf{X}_3|}, \mathbf{X}_4^2)$ for the Moderate Correlation and No Correlation models, respectively.

When the Mahalanobis distance is defined using $(\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \mathbf{X}_4)$, the results are similar to those presented in Section 4.3: The unconditional randomization test using $\hat{\tau}_{int}$ and the conditional randomization test using $\hat{\tau}_{sd}$ are conditionally valid across quantile groups, while the rejection rate of the unconditional randomization test using $\hat{\tau}_{sd}$ increases with covariate imbalance. For the No Correlation model, even the unconditional randomization test using $\hat{\tau}_{sd}$ appears to be conditionally valid across quantile groups; this is because the

covariates are not correlated with the outcome, and thus the treatment effect is not confounded by covariate imbalances in $(\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \mathbf{X}_4)$.

However, when the Mahalanobis distance is defined using the functions of the covariates that are linearly related to the potential outcomes, the rejection rate of all the randomization tests are monotonically increasing in the covariate imbalance defined by this Mahalanobis distance. This is because the treatment effect is *confounded* by covariate imbalances in $(\mathbf{X}_1^2, \mathbf{X}_2, \mathbf{X}_3^2, \mathbf{X}_4)$ and $(\sqrt{|\mathbf{X}_1|}, \mathbf{X}_2^2, \sqrt{|\mathbf{X}_3|}, \mathbf{X}_4^2)$ for the Moderate Correlation and No Correlation models, respectively. Because none of the randomization tests incorporate these functions of the covariates, we see this monotonic behavior in the rejection rate for all randomization tests, as shown in Figures 12b and 12d. In other words, similar to how the unconditional randomization test using $\hat{\tau}_{sd}$ does not adjust for linear imbalances in the covariates and thus exhibited this monotonic behavior in Section 4, the conditional randomization test using $\hat{\tau}_{sd}$ and the unconditional randomization test using $\hat{\tau}_{int}$ similarly do not fully account for imbalances in $(\mathbf{X}_1^2, \mathbf{X}_2, \mathbf{X}_3^2, \mathbf{X}_4)$ or $(\sqrt{|\mathbf{X}_1|}, \mathbf{X}_2^2, \sqrt{|\mathbf{X}_3|}, \mathbf{X}_4^2)$, and thus we again see the monotonic behavior in Figures 12b and 12d. The conditional randomization test using $\hat{\tau}_{sd}$ and the unconditional randomization test using $\hat{\tau}_{int}$ are only accounting for imbalances in $(\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \mathbf{X}_4)$. This also suggests why, in Figure 12b (when the covariates are moderately correlated with the outcome), the monotonicity of the rejection rate for these two tests is less pronounced than that of the unconditional randomization test using $\hat{\tau}_{sd}$, whereas in Figure 12d (when the covariates are not correlated with the outcome), the behavior of the rejection rate for all the randomization tests is essentially the same.

In summary, when the Mahalanobis distance (or test statistic $\hat{\tau}_{int}$) is defined using functions of the covariates that are moderately correlated with the potential outcomes, then it is still beneficial to use our conditional randomization test (or the unconditional randomization test using $\hat{\tau}_{int}$) over the unconditional randomization test using $\hat{\tau}_{sd}$. Furthermore, the equivalence of the unconditional randomization test using $\hat{\tau}_{int}$ and our conditional randomization test appears to still hold when the regression used to construct $\hat{\tau}_{int}$ is misspecified. Finally, the unconditional randomization test using $\hat{\tau}_{int}$ and our conditional randomization test appear to be valid across various degrees of imbalance in functions of the covariates used to define $\hat{\tau}_{int}$ or the Mahalanobis distance. However, this does not guarantee that these tests will be conditionally valid across covariate imbalances that are not captured by $\hat{\tau}_{int}$ or the Mahalanobis distance but nonetheless confound treatment effect estimates. Regardless, both the unconditional and conditional performance of our conditional randomization test and the unconditional randomization test using $\hat{\tau}_{int}$ appear to be preferable to those of the unconditional randomization test using $\hat{\tau}_{sd}$ if covariates are correlated with outcomes, and otherwise they appear to be equivalent.

Acknowledgment: We would like to thank two anonymous reviewers and the Associate Editor, Peter Aronow, for their insightful comments that led to notable improvements in this work.

Funding: This research was supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 1144152. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

1. Aronow PM, Middleton JA. A class of unbiased estimators of the average treatment effect in randomized experiments. *J Causal Inference*. 2013;1(1):135–54.
2. Branson Z. Is my matched dataset as-if randomized, more, or less? Unifying the design and analysis of observational studies. 2018. arXiv preprint. arXiv:1804.08760.
3. Branson Z, Bind M-A. Randomization-based inference for Bernoulli trial experiments and implications for observational studies. *Stat Methods Med Res*. 2018:1–21.
4. Caughy D, Dafoe A, Miratrix L. Beyond the sharp null: Permutation tests actually test heterogeneous effects. In: Summer Meeting of the Society for Political Methodology, Rice University, July. vol. 22. 2016.

5. Ding P. A paradox from randomization-based causal inference. *Stat Sci.* 2017;32(3):331–45.
6. Ding P, Feller A, Miratrix L. Randomization inference for treatment effect variation. *J R Stat Soc, Ser B, Stat Methodol.* 2016;78(3):655–71.
7. Freedman DA. On regression adjustments to experimental data. *Adv Appl Math.* 2008;40(2):180–93.
8. Good P. *Permutation Tests: A Practical Guide to Resampling Methods for Testing Hypotheses.* Berlin: Springer; 2013.
9. Hansen BB, Bowers J. Covariate balance in simple, stratified and clustered comparative studies. *Stat Sci* 2008;219–36.
10. Hennessy J, Dasgupta T, Miratrix L, Pattanayak C, Sarkar P. A conditional randomization test to account for covariate imbalance in randomized experiments. *J Causal Inference.* 2016;4(1):61–80.
11. Hernández AV, Steyerberg EW, Habbema JDF. Covariate adjustment in randomized controlled trials with dichotomous outcomes increases statistical power and reduces sample size requirements. *J Clin Epidemiol.* 2004;57(5):454–60.
12. Iacus S, King G, Porro G, et al. CEM: software for coarsened exact matching. *J Stat Softw.* 2009;30(13):1–27.
13. Iacus SM, King G, Porro G. Multivariate matching methods that are monotonic imbalance bounding. *J Am Stat Assoc.* 2011;106(493):345–61.
14. Iacus SM, King G, Porro G. Causal inference without balance checking: Coarsened exact matching. *Polit Anal.* 2012;20(1):1–24.
15. Imai K, King G, Stuart EA. Misunderstandings between experimentalists and observationalists about causal inference. *J R Stat Soc, Ser A, Stat Soc.* 2008;171(2):481–502.
16. Imbens GW, Rubin DB. *Causal Inference in Statistics, Social, and Biomedical Sciences.* Cambridge University Press; 2015.
17. Li X, Ding P, Rubin DB. Asymptotic theory of rerandomization in treatment–control experiments. *Proc Natl Acad Sci.* 2018;115(37):9157–62.
18. Lin W. Agnostic notes on regression adjustments to experimental data: Reexamining freedman’s critique. *Ann Appl Stat.* 2013;7(1):295–318.
19. Mahalanobis PC. On the generalized distance in statistics. *Proc Natl Inst Sci Calcutta.* 1936;2:49–55.
20. Miratrix LW, Sekhon JS, Yu B. Adjusting treatment effect estimates by post-stratification in randomized experiments. *J R Stat Soc, Ser B, Stat Methodol.* 2013;75(2):369–96.
21. Morgan KL, Rubin DB. Rerandomization to improve covariate balance in experiments. *The Annals of Statistics.* 2012. 1263–1282.
22. Morgan KL, Rubin DB. Rerandomization to balance tiers of covariates. *J Am Stat Assoc.* 2015;110(512):1412–21.
23. Raz J. Testing for no effect when estimating a smooth function by nonparametric regression: A randomization approach. *J Am Stat Assoc.* 1990;85(409):132–8.
24. Rosenbaum PR. Conditional permutation tests and the propensity score in observational studies. *J Am Stat Assoc.* 1984;79(387):565–74.
25. Rosenbaum PR. Permutation tests for matched pairs with adjustments for covariates. *Appl Stat.* 1988;401–11.
26. Rosenbaum PR. Covariance adjustment in randomized experiments and observational studies. *Stat Sci.* 2002;17(3):286–327.
27. Rosenbaum PR. *Observational Studies.* Springer; 2002.
28. Rosenberger WF, Lachin JM. *Randomization in Clinical Trials: Theory and Practice.* John Wiley & Sons; 2015.
29. Rubin DB. Estimating causal effects of treatments in randomized and nonrandomized studies. *J Educ Psychol.* 1974;66(5):688.
30. Stuart EA. Matching methods for causal inference: A review and a look forward. *Stat Sci.* 2010;25(1):1.
31. Zheng L, Zelen M. Multi-center clinical trials: Randomization and ancillary statistics. *Ann Appl Stat.* 2008;582–600.