

ARTICLE TYPE

Unlocking Insights: Enhanced Analysis of Covariance in General Factorial Designs through Multiple Contrast Tests under Variance Heteroscedasticity

Matthias Becher^{*1} | Ludwig A. Hothorn² | Frank Konietschke¹

¹Institut für Biometrie und klinische Epidemiologie, Charité - Universitätsmedizin Berlin, Berlin, Germany

²Im Grund 12, D-31867 Lauenau, Germany (retired from Leibniz University Hannover)

Correspondence

*Corresponding author Email: matthias.becher@charite.de

Present Address

Matthias Becher, Institute of Biometry and Clinical Epidemiology, Charitéplatz 1, 10117 Berlin, Germany.

Summary

A common goal in clinical trials is to conduct tests on estimated treatment effects adjusted for covariates such as age or sex. Analysis of Covariance (ANCOVA) is often used in these scenarios to test the global null hypothesis of no treatment effect using an F -test. However, in several samples, the F -test does not provide any information about individual null hypotheses and has strict assumptions such as variance homoscedasticity. We extend the method proposed by Konietschke et al.¹ to a multiple contrast test procedure (MCTP), which allows us to test arbitrary linear hypotheses and provides information about the global- as well as the individual null hypotheses. Further, we can calculate compatible simultaneous confidence intervals for the individual effects. We derive a small sample size approximation of the distribution of the test statistic via a multivariate t -distribution. As an alternative, we introduce a Wild-bootstrap method. Extensive simulations show that our methods are applicable even when sample sizes are small. Their application is further illustrated within a real data example.

KEYWORDS:

ANCOVA, MCTP, Box-type approximation, bootstrap, experimental designs

1 | INTRODUCTION

A common goal in clinical and pre-clinical research is to estimate and test the impact of factor levels and their combinations on an outcome. Even in complete randomized designs, the outcome is often not just dependent on the factors of interest, but also on other variables such as age, gender, weight, etc. which are called covariates. They can distort the factor effects and thus bias the results and especially the treatment effect estimates. Therefore, the chosen model needs to adjust the treatment effect estimates for their impact, which is usually carried out in a so-called *Analysis of Covariance* (ANCOVA) model. In statistical practice, especially in early and pre-clinical stages, researchers are often confronted with other data situations than (multivariate) normality and homogeneous variances; rather skewed data and especially variance heteroscedasticity are no rarity. The traditional F -test (which assumes homogeneity and normal errors) is thus often not applicable in these situations and therefore there is a need for statistical methods that neither postulate a specific data distribution nor homogeneous variances.

Different methods have been proposed to address the problem of variance heteroscedasticity in the general ANOVA designs without covariates^{2,3}. More recently Konietschke et al.¹ introduced a method for the general ANCOVA allowing for group-wise variance heteroscedasticity. In their paper they propose new test statistics, using methods of moments to get unbiased estimators

of the variance components. They approximate the distribution of the test statistics with Box-type approximation methods. One downside of these types of methods is that they can only be used to test the global null hypothesis H_0 of no effect difference between any groups. Finding significant differences in effects between specific groups of interest and computing corresponding simultaneous confidence intervals (SCIs) requires three steps. First, the global null hypothesis is tested. Then, multiple comparisons are tested to evaluate the individual hypotheses. Lastly, the SCIs are computed. There are, however, some problems with this approach. The global null hypothesis may be rejected, but none of the individual and vice versa. In addition, even if an individual hypothesis is rejected the corresponding SCI may include the null, i.e. the value of no treatment effect.

To address these problems multiple contrast test procedures (MCTPs) were developed. These test procedures along with compatible SCIs were derived by Mukerjee et al.⁴ and Bretz et al.⁵. Within the MCTP framework, multiple individual null hypotheses are tested with adequate test statistics. Since they are not necessarily independent, critical- and p-values are computed from their joint distribution taking their correlation into account. The global null hypothesis is rejected if any of the individual ones are rejected. This way the results of the individual null hypothesis tests are consonant and coherent with the result of the global null hypothesis test. MCTPs are highly flexible and can be used to test a variety of contrasts including all-pairs, many-to-one, trends, etc. As a result, they can be adapted depending on the research question.

Hothorn et al.⁶ have proposed a method to apply the MCTP framework to general parametric models such as linear regression or ANCOVA. However, their method relies on distributional assumptions, at least asymptotically. In small sample sizes, the tests often behave liberal and over-reject the null hypotheses.

In this article, we aim to close that gap and approximate the distributions of the tests for small sample sizes. In addition, this approximation works under group-wise variance heteroscedasticity. Finally, we also introduce a wild-bootstrap-based method to estimate the distribution of the test statistic, when complete variance heteroscedasticity is present. The remainder of the paper is organized as follows: In Section 2, we introduce the statistical model as well as the effect and variance estimators. In Section 3 we extend the introduced methods to the MCTP framework and present our approximations of the distribution of the test statistic. We will evaluate our new methods using a simulation study in Section 4. This is followed in Section 5 by a data example using data from a toxicological animal study. Finally, we will discuss our findings in Section 6.

2 | BACKGROUND

2.1 | Statistical model

We consider a one-way layout involving a groups of n_i independent subjects each that have been (will be) observed under the respective treatments as well as under M covariate conditions. Of major interest is estimating and testing treatment effect(s) adjusted for the impact of the covariates. Here and throughout, the assumption(s) on the variances of the data plays a major role in the inferential framework. The covariates and/or the treatments may induce heteroscedasticity group-wise (often in pre-clinical trials) or even subject-wise on an individual level (e.g., in observational trials). The common ANCOVA model assumes homogeneity and therefore does not constitute a suitable framework. Here and throughout, we consider a general ANCOVA model

$$Y_{ij} = b_i + \sum_{l=1}^M p_l M_{ij}^{(l)} + \epsilon_{ij}, \quad i = 1, \dots, a; \quad j = 1, \dots, n_i,$$

where b_i denotes the fixed and unknown effect of treatment i , and p_l the regression parameter of the l^{th} covariate. Two-way and higher-way layouts in means of general factorial designs are obtained by sub-indexing the index i . No less important, ϵ_{ij} denotes the error term with $E(\epsilon_{ij}) = 0$ and either group-wise variance $Var(\epsilon_{ij}) = \sigma_i^2 > 0$ for $i = 1, \dots, a$, or subject-wise variances $\sigma_{ij}^2 > 0$, for $i = 1, \dots, a; j = 1, \dots, n_i$, respectively. For ease of reading, we tackle both heteroscedastic models simultaneously and not separately. The test statistics proposed in the coming sections use variance estimators, which are consistent in either model, and therefore, the risk of confusing them is minimal. The total number of subjects is denoted by $N = \sum_{i=1}^a n_i$. In matrix notation, the model can be written as

$$\mathbf{Y} = \mathbf{X}\mathbf{b} + \mathbf{M}\mathbf{p} + \boldsymbol{\epsilon},$$

where \mathbf{Y} denotes the $N \times 1$ response vector, $\mathbf{X} = \bigoplus_{i=1}^a \mathbf{1}_{n_i}$ the design matrix, $\mathbf{b} = (b_1, \dots, b_a)'$ the vector of treatment effects, \mathbf{M} the $N \times l$ matrix of the covariates, $\mathbf{p} = (p_1, \dots, p_l)'$ the vector of regression parameters, and $\boldsymbol{\epsilon}$ the vector of the independent

error terms with

$$E(\epsilon) = \mathbf{0} \text{ and } Var(\epsilon) = \begin{cases} \Sigma_G = \bigoplus_{i=1}^a \sigma_i^2 \mathbf{I}_{n_i}, & \text{(group-wise), or} \\ \Sigma_I = diag(\sigma_{11}^2, \dots, \sigma_{an_a}^2), & \text{(subject-wise).} \end{cases} \quad (1)$$

The latter variance structure refers to *complete* heteroscedasticity in which the random variables from each subject may have different variances and thus may come from different populations. Of major interest is making statistical inferences in estimating the unknown model parameters \mathbf{b} , \mathbf{p} , Σ_G and Σ_I , as well as testing the null hypotheses

$$H_0^{(b)} : \mathbf{C}\mathbf{b} = \mathbf{0} \text{ and } H_1^{(b)} : \mathbf{C}\mathbf{b} \neq \mathbf{0}$$

formulated in terms of the treatment effects \mathbf{b} . Here, \mathbf{C} denotes the hypothesis or contrast matrix, which we will discuss in more detail in the next subsection.

2.2 | Hypotheses and contrast matrices

In studies involving more than two randomized treatment levels of the factor, research questions are manifold and rarely restricted to testing the global null hypothesis $H_0^{(b)} : b_1 = \dots = b_a$ solely. If $H_0^{(b)}$ gets rejected by an appropriate test procedure (e.g., F -test), any of the a groups differ. The arising question is “Which ones?”, which then can only be answered using omnibus test procedures (post-hoc) and multiplicity adjustments. Hence, testing q multiple local null hypotheses (alternatives) $H_0^{(\ell)} : \mathbf{c}'_{\ell}\mathbf{b} = 0$ with a user-defined $q \times a$ contrast matrix \mathbf{C} is of major interest. Here, $\mathbf{c}'_{\ell} = (c_{\ell 1}, \dots, c_{\ell a})$ denotes the ℓ th row vector of \mathbf{C} , $\ell = 1, \dots, q$. For instance, many-to-one comparisons (so-called Dunnett-type contrasts⁷) are performed using

$$H_0 : \begin{cases} b_1 = b_2 \\ b_1 = b_3 \\ \vdots \\ b_1 = b_a \end{cases} \Leftrightarrow H_0 : \mathbf{C}\mathbf{b} = \begin{pmatrix} -1 & 1 & 0 & 0 & \dots & 0 \\ -1 & 0 & 1 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ -1 & 0 & 0 & 0 & \dots & 1 \end{pmatrix} \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_a \end{pmatrix} = \mathbf{0},$$

whereas all-pairwise comparisons (Tukey-type^{8,9}) or comparisons to the mean (Grand-mean type⁵) are performed using

$$H_0 : \begin{cases} b_1 = b_2 \\ b_1 = b_3 \\ \vdots \\ b_1 = b_a \\ b_2 = b_3 \\ \vdots \\ b_{a-1} = b_a \end{cases} \Leftrightarrow H_0 : \mathbf{C}\mathbf{b} = \begin{pmatrix} -1 & 1 & 0 & 0 & \dots & 0 \\ -1 & 0 & 1 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ -1 & 0 & 0 & 0 & \dots & 1 \\ 0 & -1 & 1 & 0 & \dots & 0 \\ 0 & -1 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & -1 & 1 \end{pmatrix} \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_a \end{pmatrix} = \mathbf{0}, \text{ or}$$

$$H_0 : \begin{cases} b_1 = \bar{b} \\ b_2 = \bar{b} \\ \vdots \\ b_a = \bar{b} \end{cases} \Leftrightarrow H_0 : \mathbf{C}\mathbf{b} = \begin{pmatrix} 1 - 1/a & -1/a & -1/a & -1/a & \dots & -1/a \\ -1/a & 1 - 1/a & -1/a & -1/a & \dots & -1/a \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ -1/a & -1/a & -1/a & -1/a & \dots & 1 - 1/a \end{pmatrix} \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_a \end{pmatrix} = \mathbf{0}, \bar{b} = \frac{1}{a} \sum_{i=1}^a b_i,$$

respectively. Which contrast to use depends on the research question of interest and cannot be recommended in a general way. Bretz et al.⁵ provide an overview of different contrasts, which are also implemented in the R function *contrMat* within the *multcomp*¹⁰ package. Note that the Grand-mean type contrast matrix is the centering projection matrix $\mathbf{P}_a = \mathbf{I}_a - \mathbf{J}_a$ well known from linear model theory.

2.3 | Effect and variance estimators

To test the null hypotheses formulated above, the unknown model parameters \mathbf{b} and \mathbf{p} as well as the variances of their estimators must be estimated from the data. We use least squares for the former and obtain the well-known generalized least squares estimators (GLS)

$$\hat{\mathbf{b}}_c = (\mathbf{X}'\Sigma_c^{-1}\mathbf{X})^{-1} \mathbf{X}'\Sigma_c^{-1} (\mathbf{Y} - \mathbf{M}\hat{\mathbf{p}}_c) \text{ and } \hat{\mathbf{p}}_c = (\mathbf{M}'\mathbf{Q}\Sigma_c^{-1}\mathbf{M})^{-1} \mathbf{M}'\mathbf{Q}\Sigma_c^{-1}\mathbf{Y}, \quad c \in \{G, I\},$$

with the projection matrices $\mathbf{P} = \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$, $\mathbf{Q} = \mathbf{I}_N - \mathbf{P}$, and covariance matrix $\mathbf{\Sigma}_c$ as defined in (1). As usual, the ordinary least squares estimator (OLS) is obtained with $\mathbf{\Sigma}_c = \mathbf{I}_N$. Since $\mathbf{\Sigma}_G$ and $\mathbf{\Sigma}_I$ are unknown, they must be replaced by consistent estimators. For the ease of representation, we define the generating matrices

$$\mathbf{A}_c = (\mathbf{M}'\mathbf{Q}\mathbf{\Sigma}_c^{-1}\mathbf{M})^{-1}\mathbf{M}'\mathbf{Q}\mathbf{\Sigma}_c^{-1} \text{ and } \mathbf{D}_c = (\mathbf{X}'\mathbf{\Sigma}_c^{-1}\mathbf{X})^{-1}\mathbf{X}'\mathbf{\Sigma}_c^{-1} - (\mathbf{X}'\mathbf{\Sigma}_c^{-1}\mathbf{X})^{-1}\mathbf{X}'\mathbf{\Sigma}_c^{-1}\mathbf{M}\mathbf{A}_c, \quad c \in \{G, I\}, \quad (2)$$

with the use of which the estimators $\hat{\mathbf{b}}_c$ and $\hat{\mathbf{p}}_c$ as well as their variances can be easily written as $\hat{\mathbf{b}}_c = \mathbf{D}_c\mathbf{Y}$, $\hat{\mathbf{p}}_c = \mathbf{A}_c\mathbf{Y}$, and

$$\mathbf{\Psi}_c = \text{Cov}(\hat{\mathbf{b}}_c) = \mathbf{D}_c\mathbf{\Sigma}_c\mathbf{D}_c' \text{ and } \mathbf{\Xi}_c = \text{Cov}(\hat{\mathbf{p}}_c) = \mathbf{A}_c\mathbf{\Sigma}_c\mathbf{A}_c', \quad c \in \{G, I\}, \quad (3)$$

respectively. This representation is charming because the covariance matrices $\mathbf{\Psi}_c$ and $\mathbf{\Xi}_c$ can be written as matrix products (*Sandwich*)¹¹ with $\mathbf{\Sigma}_G$ and $\mathbf{\Sigma}_I$ being the only unknown components. An unbiased estimator^{1,12} of $\mathbf{\Sigma}_G$ is given by

$$\hat{\mathbf{\Sigma}}_G = \bigoplus_{i=1}^a \hat{\sigma}_i^2 \mathbf{1}_{n_i}, \text{ where } \hat{\sigma}_i^2 = \frac{\mathbf{Y}_i' \mathbf{Q}_i \mathbf{Y}_i}{n_i - 1 - \text{rank}(\mathbf{M}_i)} \quad (4)$$

is an unbiased estimator of σ_i^2 . Here, \mathbf{Y}_i is the $n_i \times 1$ vector of the response in group i , $\mathbf{Q}_i = \mathbf{I}_{n_i} - \mathbf{B}_i(\mathbf{B}_i'\mathbf{B}_i)^{-1}\mathbf{B}_i'$ is a projection matrix, $\mathbf{B}_i = (\mathbf{X}_i, \mathbf{M}_i)$, $\mathbf{X}_i = \mathbf{1}_{n_i}$, and \mathbf{M}_i is the group specific matrix of covariates, respectively. Similarly, a heteroscedasticity consistent estimator^{13,14} of $\mathbf{\Sigma}_I$ is given by

$$\hat{\mathbf{\Sigma}}_I = \bigoplus_{i=1}^a \bigoplus_{j=1}^{n_i} \hat{\epsilon}_{ij}^2. \quad (5)$$

Finally, we obtain consistent estimators $\hat{\mathbf{\Psi}}_c$ and $\hat{\mathbf{\Xi}}_c$ of $\mathbf{\Psi}_c$ and $\mathbf{\Xi}_c$ by replacing $\mathbf{\Sigma}_G$ and $\mathbf{\Sigma}_I$ with their estimators $\hat{\mathbf{\Sigma}}_G$ and $\hat{\mathbf{\Sigma}}_I$ in (3), respectively. For making inferences, information about the (asymptotic) joint distribution of the estimators is the only remaining missing information at this stage. It can be shown that for increasing (large) sample sizes, i.e., if $N \rightarrow \infty$ such that $N/n_i \rightarrow \lambda_i < \infty$, and under a few mild assumptions on the variances, the estimators follow multivariate normal distributions, i.e.,

$$\sqrt{N}(\hat{\mathbf{b}}_c - \mathbf{b}) \rightsquigarrow N(\mathbf{0}, N\mathbf{\Psi}_c), \text{ and } \sqrt{N}(\hat{\mathbf{p}}_c - \mathbf{p}) \rightsquigarrow N(\mathbf{0}, N\mathbf{\Xi}_c),$$

where $\mathbf{\Psi}_c$ and $\mathbf{\Xi}_c$ are as given in (3), respectively. Next, different test statistics for testing the null hypotheses motivated above in Section 2.2 will be introduced.

3 | MULTIPLE CONTRAST TEST PROCEDURES

Since the pioneering works of Dunnett and Tukey^{7,8}, and many others for making multiple comparisons, multiple contrast tests are well established tools in (bio)-statistical sciences. They have been developed for testing multiple null hypotheses formulated in different effects such as mean differences (or ratios thereof)^{5,15,16,17,18}, proportions¹⁹ or purely nonparametric relative effects^{20,21,22,23}. In the following, we will revisit the general terminology for testing null hypotheses in ANCOVA models in line with Hothorn et al.⁶ and propose small sample size approximations of the distributions of the test statistics afterward. To test the individual null hypothesis $H_0^{(\ell)} : \mathbf{c}'_\ell \mathbf{b} = 0$, consider the test statistic

$$T_{\ell,c} = \frac{\mathbf{c}'_\ell \hat{\mathbf{b}}_c}{\sqrt{\mathbf{c}'_\ell \hat{\mathbf{\Psi}}_c \mathbf{c}_\ell}}, \quad c \in \{G, I\}, \quad \ell = 1, \dots, q. \quad (6)$$

Even though each test statistic $T_{\ell,c}$ follows a $N(0, 1)$ distribution for large samples sizes, its finite and small sample size distribution is yet unknown. Furthermore, the test statistics $T_{\ell,c}$ and $T_{\ell',c}$, $\ell \neq \ell'$, are not necessarily independent. Ignoring their dependency by using the Bonferroni inequality within the inferential framework would result in a loss of power. We therefore collect them in the vector $\mathbf{T}_c = (T_{1,c}, \dots, T_{q,c})'$ and use its joint distribution for the computation of critical- and p-values. Here, \mathbf{T}_c follows, for large sample sizes, if $N \rightarrow \infty$ such that $\frac{N}{n_i} \leq N_0 < \infty$ (and further mild assumptions on variances being finite), a multivariate normal distribution with expectation $\mathbf{0}$ and correlation matrix

$$\mathbf{R}_c = \text{diag}(\mathbf{C}\mathbf{\Psi}_c\mathbf{C}')^{-1/2} (\mathbf{C}\mathbf{\Psi}_c\mathbf{C}') \text{diag}(\mathbf{C}\mathbf{\Psi}_c\mathbf{C}')^{-1/2}. \quad (7)$$

For large sample sizes, the null hypothesis $H_0^{(\ell)} : \mathbf{c}'_\ell \mathbf{b} = 0$ will be rejected at level $\alpha \in (0, 1)$, if $|T_{\ell,c}| \geq z_{1-\alpha}(\mathbf{R}_c)$. A compatible $(1 - \alpha)$ simultaneous confidence interval for the treatment effect $\delta_\ell = \mathbf{c}'_\ell \mathbf{b}$ is obtained from

$$CI_\ell = \left[\mathbf{c}'_\ell \hat{\mathbf{b}}_c \pm z_{1-\alpha}(\mathbf{R}_c) \sqrt{\mathbf{c}'_\ell \hat{\Psi}_c \mathbf{c}_\ell} \right], \ell = 1, \dots, q, \quad (8)$$

where $z_{1-\alpha}(\mathbf{R}_c)$ denotes the two-sided $(1 - \alpha)$ -equicoordinate quantile of the $N(\mathbf{0}, \mathbf{R}_c)$ distribution⁶. For large sample sizes, the global null hypothesis $H_0 : \mathbf{C}\mathbf{b} = \mathbf{0}$ will be rejected at level α , if

$$T_{0,c} = \max \{ |T_{1,c}|, \dots, |T_{q,c}| \} \geq z_{1-\alpha}(\mathbf{R}_c). \quad (9)$$

For small sample sizes, however, the MCTP does not control the type-1 error rate, behaves liberal and over-rejects the null hypothesis, resulting in false positive conclusions. We therefore develop an approximation of the joint distribution of \mathbf{T}_c for small sample sizes in the next subsections. We hereby differ between group-wise and complete heteroscedasticity and develop MCTPs for each model separately.

3.1 | A small size approximation of the distribution of \mathbf{T}_G

Group-wise heteroscedasticity is a common assumption in pre-clinical and clinical trials. Motivated by the Satterthwaite-Welch t -test and its generalization to MCTPs¹⁵, we aim to approximate the distribution of \mathbf{T}_G by a central multivariate $t(\nu, \mathbf{R}_G, \mathbf{0})$ -distribution with ν degrees of freedom and a correlation matrix \mathbf{R}_G . The correlation matrix is derived as shown in (7) based on the group-wise variance estimator we introduced in (4). A straight forward computation of the degree of freedom ν is, however, not possible in the statistical model considered here: First, the multivariate t -distribution is, by definition, homoscedastic. Second, the variances of the contrasts $\mathbf{c}'_\ell \hat{\mathbf{b}}_G$ differ in their degrees of heteroscedasticity. We therefore compute the degree of freedom of the distribution of each test statistic $T_{\ell,G}$ separately using a Box-type approximation, and select one of the q candidates to approximate the joint distribution conservatively in a second step. Later, we compare the qualities of the different selections in extensive simulation studies in Section 4. To begin with the former, we approximate the distribution of the estimated variance of $\mathbf{c}'_\ell \hat{\mathbf{b}}_G$ by a scaled $\chi^2_{f_\ell}$ distribution,

$$\mathbf{c}'_\ell \hat{\Psi}_c \mathbf{c}_\ell \approx g_\ell \chi^2_{f_\ell},$$

such that the first two moments coincide and obtain

$$\nu_\ell = \frac{\text{diag}(\mathbf{c}'_\ell \hat{\Psi}_c \mathbf{c}_\ell)^2}{(\mathbf{k}'_\ell \mathbf{X})^2 \mathbf{D}_\delta^{-2} \mathbf{\Omega}^{-1}}, \ell = 1, \dots, q, \quad (10)$$

where, $\mathbf{k}_\ell = (\mathbf{c}'_\ell \mathbf{D})^2$ denotes the ℓ th row of the matrix $\mathbf{K} = (\mathbf{C}\mathbf{D})^2$, $\mathbf{D}_\delta = \text{diag}(\hat{\sigma}_1^2, \dots, \hat{\sigma}_q^2)$ denote the diagonal matrix of the variance estimators $\hat{\sigma}_i^2$ as given in (4) and $\mathbf{\Omega}^{-1} = \text{diag}\left(\frac{1}{(n_1-1-\text{rank}(\mathbf{M}_1))}, \dots, \frac{1}{(n_q-1-\text{rank}(\mathbf{M}_q))}\right)$ denote the diagonal matrix of their numerators (degrees of freedom) of each variance estimator separately. As potential candidates for the degree of freedom of the joint distribution, we choose the minimum, maximum, or mean (rounded to the nearest integer) of the candidate values ν_ℓ . The larger the degree of freedom the closer will be the joint distribution to the asymptotic $N(\mathbf{0}, \mathbf{R}_G)$ distribution, or, in other words, the larger ν , the more liberal will be the approximation. The individual null hypothesis $H_0^{(\ell)} : \mathbf{c}'_\ell \mathbf{b} = 0$ will be rejected at multiple α level of significance, if $|T_\ell| \geq t_{1-\alpha}(\nu, \mathbf{R}_G)$. Compatible $(1 - \alpha)$ -simultaneous confidence intervals are given by

$$CI_\ell = \left[\mathbf{c}'_\ell \mathbf{b} \pm t_{1-\alpha}(\nu, \mathbf{R}_G) \sqrt{\mathbf{c}'_\ell \hat{\Psi}_c \mathbf{c}_\ell} \right], \ell = 1, \dots, q. \quad (11)$$

Finally, the global null hypothesis $H_0 : \mathbf{C}\mathbf{b} = \mathbf{0}$ will be rejected at a significance level α , if

$$T_0 \geq t_{1-\alpha}(\nu, \mathbf{R}_G)$$

with $t_{1-\alpha}(\nu, \mathbf{R})$ the $(1 - \alpha)$ -equicoordinate quantile of the multivariate $t(\nu, \mathbf{R}_G, \mathbf{0})$ distribution. Note that the correlation matrix \mathbf{R}_G is unknown in applications and must be estimated from the data. We recommend to replace \mathbf{R}_G with the consistent estimator

$$\hat{\mathbf{R}}_G = \text{diag}(\mathbf{C}\hat{\Psi}_G\mathbf{C}')^{-1/2} (\mathbf{C}\hat{\Psi}_G\mathbf{C}') \text{diag}(\mathbf{C}\hat{\Psi}_G\mathbf{C}')^{-1/2}$$

in (8), (9), (11), and (12) above, respectively.

3.2 | A small size approximation of the distribution of T_I

In the previous section, we introduced an approximation of the distribution of the maximum test statistic $T_{0,G}$ (T_G) for small sample sizes under group-wise variance heteroscedasticity assumption, also known as the *Behrens-Fisher* situation. In statistical practice, e.g., in observational studies, however, the assumption may be too strict and unrealistic. We therefore also investigate a multiple contrast test that is valid under complete heteroscedasticity. In this situation, approximating the distribution of the test statistic by a multivariate t -distribution results—if it is even possible—in very cumbersome computations of its degree of freedom. We, therefore, propose a Wild-Bootstrap resampling-based approximation as proposed by^{1,14}. Let W_{11}, \dots, W_{an_a} be independent and identically distributed random weights with $E(W_{11}) = 0$ and $Var(W_{11}) = 1$. We use random signs (Rademacher weights) throughout, i.e., $P(W_{11} = 1) = P(W_{11} = -1) = 1/2$, independently of the original data. The resampling variables are now computed by multiplying the random weights with the (scaled) residuals

$$Y_{ij}^* = \hat{\epsilon}_{ij} W_{ij} (1 - p_{ij})^{-1/2}.$$

Here, p_{ij} denotes the diagonal elements of the matrix $\mathbf{P} = \mathbf{B}\mathbf{P}_B$, where $\mathbf{B} = (\mathbf{X}, \mathbf{M})$ and $\mathbf{P}_B = (\mathbf{B}'\mathbf{B})^{-1}\mathbf{B}'$ denotes a projection matrix, respectively. To approximate the distribution of $T_{0,I}$ for small sample sizes, consider the conditional (i.e., given the data \mathbf{Y}) distribution of the maximum test statistic

$$T_{0,I}^* = \max \left(|T_{1,I}^*|, \dots, |T_{q,I}^*| \right), \text{ where } T_{\ell,I}^* = \frac{\mathbf{c}'_{\ell} \hat{\mathbf{b}}_I^*}{\sqrt{\mathbf{c}'_{\ell} \hat{\Psi}_I^* \mathbf{c}_{\ell}}}, \ell = 1, \dots, q. \quad (12)$$

In comparison with the test statistic $T_{\ell,I}$ given in (6), $T_{\ell,I}^*$ is computed with the resampling variables Y_{ij}^* instead of Y_{ij} , i.e., $\hat{\mathbf{b}}_I^*$ denotes the OLS estimator as given in (2), and $\hat{\Psi}_I^*$ the upper left block matrix of the matrix $\mathbf{P}_B \hat{\Sigma}_I^* \mathbf{P}_B'$ with $\hat{\Sigma}_I^*$ denoting the estimator as given in (5) computed with Y_{ij}^* instead of Y_{ij} , respectively. The individual null hypothesis $H_0^{(\ell)} : \mathbf{c}_{\ell}' \mathbf{b} = 0$ will be rejected at level $\alpha \in (0, 1)$, if

$$|T_{\ell,I}| \geq T_{1-\alpha}^*,$$

where $T_{1-\alpha}^*$ denotes the $(1 - \alpha) \cdot 100\%$ quantile of the distribution of $T_{0,I}^*$ as given in (12). Compatible simultaneous confidence intervals for the treatment effects $\delta_{\ell} = \mathbf{c}_{\ell}' \mathbf{b}$ are obtained from

$$CI_{\ell}^* = \left[\mathbf{c}_{\ell}' \hat{\mathbf{b}} \pm T_{1-\alpha}^* \sqrt{\mathbf{c}_{\ell}' \hat{\Psi}_I^* \mathbf{c}_{\ell}} \right].$$

Finally, the global null hypothesis $H_0 : \mathbf{C}\mathbf{b} = \mathbf{0}$, will be rejected at level α , if

$$T_{0,I} \geq T_{1-\alpha}^*. \quad (13)$$

The numerical computation of the critical values $T_{1-\alpha}^*$ (or p-values) is as follows:

1. Given the data, compute the test statistic $T_{0,I}$ in (9).
2. Generate N random weights W_{ij} and compute the statistic $T_{0,I}^*$. Save the value in $T_{0,I,1}$.
3. Repeat the previous step a n_{boot} times, e.g., $n_{boot} = 10,000$, and obtain the values $T_{0,I,1}, \dots, T_{0,I,n_{boot}}$.
4. Estimate the critical value $T_{1-\alpha}^*$ by the empirical $(1 - \alpha) \times 100\%$ -quantile of the values $T_{0,I,1}, \dots, T_{0,I,n_{boot}}$.

Note that the bootstrap approach is constructed in such a way that the correlation matrix of the multiple test statistics $T_{\ell,I}$ must not be estimated from the data. The information about the correlation is implicitly involved in the (resampling) distribution of $T_{0,I}^*$ (i.e., in $T_{1-\alpha}^*$). The approximation quality will be investigated in extensive simulation studies in the next section.

4 | SIMULATION

4.1 | Setup

All of the methods investigated are approximate. We therefore examine their behavior (type-I error rate control and power to detect selected alternatives) in small samples in an extensive simulation study; each with $nsim = nboot = 10K$ simulation and bootstrap runs. As a benchmark we use the general linear hypothesis multiple comparisons for parametric models (glht) proposed

by Hothorn et al.⁶, which we mentioned in chapter 1. We compare both the bootstrap method from 3.2 and the approximation with the t-distribution as described in 3.1 to the benchmark. For the latter, we consider three possibilities (min, max, mean) to choose from the candidates for the degrees of freedom of the joint distribution (10).

First, we examine the ability of the methods to accurately control the nominal type-1 error rate $\alpha = 5\%$. In the second step, we conduct further simulation studies to estimate the power of the methods to detect selected alternatives. We would like to investigate whether the methods are applicable in a range of possible settings. Therefore, we generate data using different parameter combinations to emulate various scenarios. Data has been generated from

$$Y_{ik} = b_i + \mathbf{M}'_{ik} \mathbf{p} + \epsilon_{ik}, i = 1, \dots, a; k = 1, \dots, n_i,$$

$$\epsilon_{ik} = \sigma_{ik} * \frac{Z_{ik} - E(Z_{ik})}{\sqrt{Var(Z_{ik})}}$$

Here, Z_{ik} are random values to base the error terms drawn from one of the following distributions: $N(0, 1)$, t_5 , χ^2_{12} , or $Exp(1)$ to evaluate the methods in symmetric and tailed distributional scenarios with different numbers of groups ($a \in \{3, 4, 5\}$) and sample sizes. Throughout, we use $m=4$ covariates drawn from $N(7, 1)$ and set $b_i = 7$ (under the null hypothesis). The covariate effects are given by $\mathbf{p} = (0.2, 1, 1.5, 2)'$. In the case of variance homoscedasticity, we set $\sigma_1 = \dots = \sigma_4 = \sigma_5 = 1$; otherwise in case of group-wise heteroscedasticity we choose $\sigma_1 \in \{2, 4, 6\}$; $\sigma_2 = 1.5$; $\sigma_3 = 1$; $\sigma_4 = 0.5$; $\sigma_5 = 0.75$, and for complete variance heteroscedasticity $\sigma_{ik} \sim U(0.5, 4)$, respectively.

We examine balanced and unbalanced designs with sample sizes $\mathbf{n} = (8, 8, 8, 8, 8) + i$, with $i \in \{0, 2, 4, \dots, 18\}$ in the balanced design and $\mathbf{n} = (8, 10, 13, 17, 20) + i$ or $\mathbf{n} = (20, 17, 13, 10, 8) + i$ for negative pairing (NP) or positive pairing (PP) respectively, with $i \in \{0, 2, 4, \dots, 18\}$. We investigate the impact of different contrasts ("Dunnett", "Tukey", and "GrandMean") on their type-I error rate control.

The implementation and simulation were done in R²⁴ version 4.2.2. Computation has been performed on the HPC for Research cluster of the Berlin Institute of Health. The code for the simulation and the data example can be found in the following repository on the GitLab of the Charité (<https://git-ext.charite.de/matthias.becher/mctp-ancova>).

4.2 | Results

Due to the many different settings, we will only present a few selected scenarios and discuss general findings in this section. The full results are provided in the supplementary material.

Table 1 Overview of the different parameter constellations and statistical models being simulated.

Setting Number	Number of Groups	Balanced/Unbalanced	Variance Structure
1	3	Unbalanced NP	Group-wise heteroscedasticity
2	5	Balanced	Complete heteroscedasticity
3	3	Balanced	Homoscedasticity
4	4	Unbalanced	Homoscedasticity
5	3	Unbalanced	Complete heteroscedasticity

The different parameter constellations and models simulated are summarized in Table 1. For each setting, we show the results for all different distributions and contrasts by sample size. The results are shown in Figure 1.

There are a few general observations we can make regardless of the setting. Firstly, we can see that glht is quite liberal in all settings and only converges towards the nominal type-1 error rate slowly with increasing sample sizes. Our methods outperform the competitors in all settings. Another observation we can make is that the bootstrap method appears to be slightly less stable than non-bootstrap methods. However, for very small sample sizes the bootstrap method performs best overall.

In setting 1 there are a few important points. As mentioned, we can see that glht is quite liberal, especially for small sample sizes, and converges to the nominal type-1 error rate with an increasing number of observations. In comparison, our methods

generally work well even when sample sizes are small, even though mean and max tend to be liberal with very few observations. We can also see that the methods perform best with symmetrical distributions, while skewed data, especially the exponential, poses a bigger challenge. This is expected, but overall the performance is still quite good.

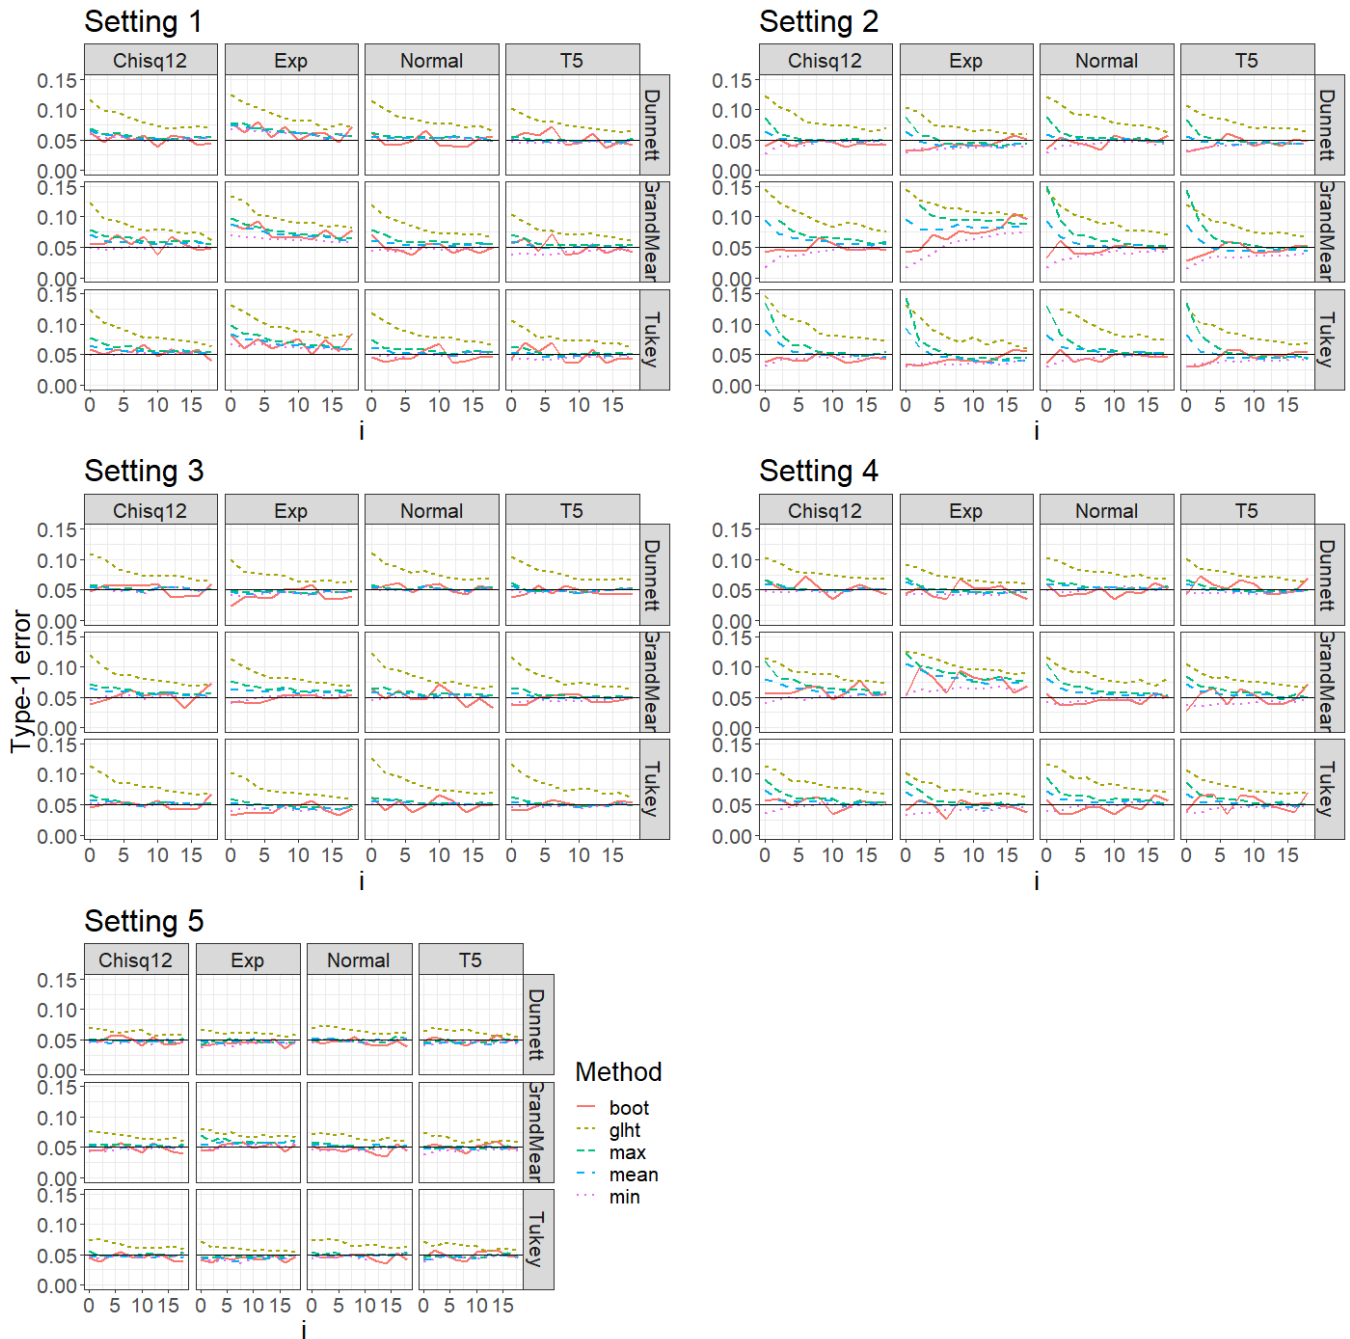


Figure 1 Type-I error simulation results of the competing methods in the five settings given in Table 1.

In Setting 2 in which contrasts of $a = 5$ groups are tested, we can see that the methods behave more liberally than in situations with fewer groups when sample sizes are small. The bootstrap method generally outperforms the other methods for small sample sizes. As before, glht is liberal in all scenarios and never manages to reach the 5% even with a higher sample size. In this setting,

the multivariate t-approximation methods also tend to be liberal in small samples. While the mean and max methods are liberal, min is overly conservative. However, the type-I errors of the methods converge quickly towards the nominal 5% level and show promising results when a moderate number of samples are available. The one exception to the overall good performance in Setting 2 is the case with the exponential distribution and the GrandMean contrast. There the methods are converging towards around 8%. They tend to be accurate with large sample sizes ($n = 100$). That being said, this specific setting is a somewhat extreme scenario and the results are expected. The performance overall in Setting 2, especially from the bootstrap method is very good.

Setting 3 is probably the "easiest" setting. As expected all of our methods perform well across the board, although the bootstrap method is a bit less stable. Nevertheless, there are no major deviations and it manages to control the type-1 error rate well overall. In contrast, the competitor again is liberal for small sample sizes and only slowly converges towards the 5% level of significance. Settings 4 and 5 again show that among the 3 parameters we are varying here, the number of groups has the strongest impact on the performances of the methods. Both settings are unbalanced and Setting 5 has complete variance heteroscedasticity compared to homoscedastic variance in Setting 4, which should be more challenging. However, with 3 groups in Setting 5 being involved, all of our methods control the nominal type-1 error rate in all scenarios for almost all sample sizes. In contrast, with 4 groups in setting 4, the performance suffers especially when only a few samples are available. Similar to Setting 2 the bootstrap method performs best at very small sample sizes.

Overall, some general observations can be made across all the results. First, the choice of the research question and as a result the choice of the contrast has a sizeable effect on the performance of the methods. Since the GrandMean contrast induces a negative correlation among the test statistics, it is more demanding, particularly when combined with a skewed distribution like the exponential. This leads to the next observation. In general, symmetric distributions are handled a bit better. However, even for skewed distributions the performance of our methods, and in particular the bootstrap method, is still good, even with only a few or a moderate number of samples. As one would expect the performance gets worse with an increasing number of groups being compared. Somewhat counter-intuitively the performance is mostly better in the unbalanced design. This is likely due to the simulation setup where the unbalanced design has an overall larger sample size. Consequently, the total number of samples is more important than some imbalance. We also investigated settings where N was constant between balanced and unbalanced scenarios (Figure A1 in the appendix), but there was no major difference between the two. Results were slightly better for positive pairing compared to negative pairing, especially for small sample sizes. This is expected as negative pairing combines the problems of small sample size and large variance, whereas in positive pairing one is offset by the other. For a comparison we refer to Figure B2 in the appendix. Interestingly, group-wise heteroscedasticity seems to be a bit more challenging for the bootstrap method, compared to complete heteroscedasticity, though it still generally performs well even in those scenarios.

4.3 | Power simulations

In the previous section, we investigated the ability of our methods to control the type-1 error rate. In practice, we are also interested in the statistical power of the tests. To that end, we conduct an extensive power simulation study based on the simulation setup described earlier. Of course, there is a wide range of possible alternatives to test. We investigate two possible alternatives:

$$\text{Alternative 1: } \mathbf{b}' = (71'_{a-1} - \delta, 71'_{a-1}) \text{ and Alternative 2: } \mathbf{b}' = (71'_{a-2} - \delta, 71'_{a-2} + \delta, 71'_{a-2}).$$

Here $\delta \in \{0, 0.2, \dots, 2\}$ is the effect we want to detect. It is the same across all settings. We only estimate the power for the smallest sample sizes $\mathbf{n} = (8, 8, 8, 8, 8)$ and $\mathbf{n} = (8, 10, 13, 17, 20)$ in the balanced and unbalanced designs respectively. As we have seen the competitor is quite liberal in virtually all scenarios at that sample size. We therefore exclude it from the power analysis.

We present the results for the same settings we showed for the type-1 error simulation. The results for Alternative 1 are shown in Figure 2. We can see that the most important impact on the power is the variance. In Settings 3 and 4 with homoscedastic variances, the power rises much faster with an increasing effect δ compared to the group-wise heteroscedasticity (Setting 1) which in turn has a higher power compared to complete variance heteroscedasticity (settings 2 and 5).

Overall, the methods perform comparably across the 5 settings. The min-method tends to have slightly lower power, especially in the settings with more than 3 groups (Settings 2 and 4). This is in line with the previous results where it was the most conservative.

As before, positive pairing leads to better results compared to negative pairing, which can be seen in Figure B3 in the appendix. Additionally, in the case of group-wise heteroscedasticity the variance of the first group σ_1 had a noticeable difference on the

power, with a larger σ_1 and therefore a larger difference in variance between the groups, leading to lower power.

In the situations considered here, the choice of contrast seems to have a significant impact on the power of the methods; this is different in other alternatives. Furthermore, since all error distributions are standardized, they seem to impact the power only on a minor level. This is in contrast to the results we have seen in the type-1 error simulation.

The results for the same settings with alternative 2 are shown in Figure C4 in the appendix. Overall, the results for Alternative 2

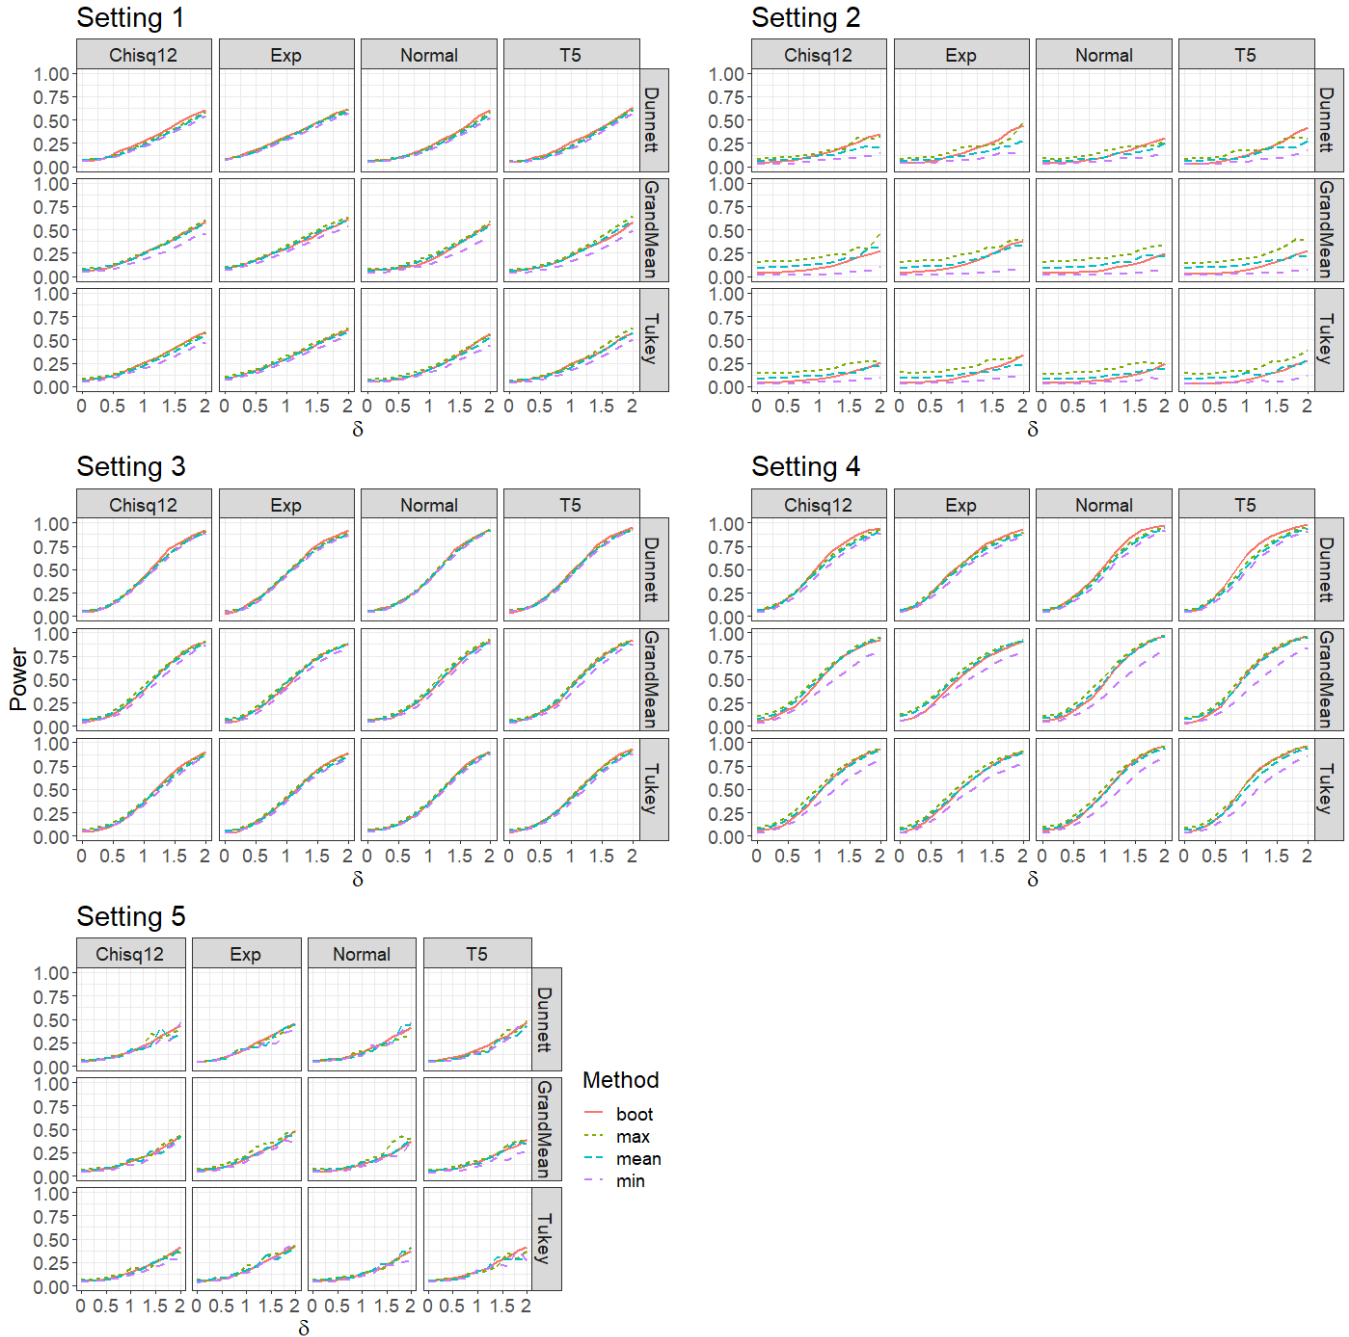


Figure 2 Power analysis results for the 5 settings with alternative 1

are similar to Alternative 1. The power generally is higher which we would expect, as there are two groups with effects compared

to one. The most important factor influencing the power is still the variance structure. We can also observe again, that the min-method tends to have the lowest power of our proposed methods.

5 | DATA EXAMPLE

To illustrate the need for the methods and their applications, we consider a subproject of a toxicological study on pyridine (number C55301B) from the US National Toxicology Program (https://manticore.niehs.nih.gov/cebssearch/test_article/110-86-1; assessed November 2020). In this animal study, the researchers investigated the effect of pyridine on various clinical chemistry parameters. As the outcome of interest, we chose blood urea nitrate (BUN) measured in mg/dl. The study included $N = 120$ rats (60 male and 60 female) randomly assigned to one of six treatment groups, six dose levels of pyridine (0, 50, 100, 250, 500, 1000). BUN was measured twice, once at baseline and once after 90 days. BUN after 90 days is the outcome of interest. As covariates, we include BUN at baseline as well as the change in body weight during the trial. This trial can then be analyzed as a two-way factorial design with the factors dose (6 levels), gender (2 levels), and 2 covariates. Figure 3 shows a boxplot of the outcomes for each combination of the factor levels. We can see a difference in the outcome between male and female rats but no obvious relationship between dose levels and the outcome. Based on the boxplot no distributional assumptions can be made. In addition, the empirical variances between the different factor combinations are quite different, so we cannot assume variance homogeneity.




Figure 3 BUN at day 90 stratified by factors sex and dose

The trial can be statistically modeled using a general ANCOVA design

$$Y_{ijk} = b_{ij} + \sum_{l=1}^2 p_l M_{ijk}^{(l)} + \epsilon_{ijk}, i = 1, 2; j = 1, \dots, 6; k = 1, \dots, n_i$$

where Y_{ijk} is BUN on day 90, b_{ij} are the factor-wise treatment effects, $M_{ijk}^{(1)}$ and $M_{ijk}^{(2)}$ are BUN at baseline and change in body weight with p_1 and p_2 the corresponding coefficients. Their OLS estimates are $\hat{p}_1 = -0.036$ and $\hat{p}_2 = -0.009$, respectively, i.e., none of them significantly impacts the response.

We now use our method to test for group-wise differences for either the individual factors or the combination of the two. Testing the null hypotheses of no gender effect, no dose effect and no interaction between dose and gender with global testing procedures, e.g., with the ones proposed by¹, yields global significant dose and gender effects at 5% level of significance. Which dose levels differ from each other is unknown but of paramount importance for the researcher. To illustrate the benefit of analyzing using MCTPs, we use the same contrasts as being used by the ANCOVA and show the results for testing the global null hypotheses in Table 2. We see that all methods detect a significant effect on the individual factors. The data do not provide the evidence to reject the null hypothesis of no interaction effect. Since the sample sizes are very small ($n \leq 10$ in every cell), we tested the interaction effect with the bootstrap method only. As already mentioned, one of the advantages of the MCTP framework is

Table 2 Global test results for the individual and combined effects for each our different methods

Method	Effect	Test statistic	p-value
min	Dose	3.525	0.015 *
	Sex	4.591	0.000 *
mean	Dose	3.525	0.009 *
	Sex	4.591	0.000 *
max	Dose	3.525	0.008 *
	Sex	4.591	0.000 *
boot	Dose	3.544	0.013 *
	Sex	4.291	0.000 *
	Dose \times Sex	5.231	0.139

the ability to also test the individual hypotheses. Based on the global results we know that there is a significant effect of the dose levels. However, we do not know which levels deviate from the average. Using MCTP we can look at the results for the individual hypotheses. Table 3 shows the results for individual dose levels calculated using the bootstrap method. This allows us to see not only that there is a difference between dose levels, but also which levels differ from the overall mean. In this case, dose levels 250 and 1000 both differ significantly from the overall mean at 5% level of significance.

Table 3 Individual results for dose level estimated by the bootstrap method

	Contrast	Effect	CI_Lower	CI_Upper	Test statistic	p-value
1	0	0.27	-1.20	1.74	0.55	0.99
2	50	-1.07	-2.59	0.46	2.07	0.28
3	100	-0.07	-1.75	1.62	0.11	1.00
4	250	-1.70	-3.11	-0.28	3.54	0.01 *
5	500	0.20	-1.58	1.99	0.34	1.00
6	1000	2.77	0.42	5.12	3.49	0.01 *

6 | DISCUSSION

ANCOVA is one of the most widely used methods in clinical and non-clinical research to estimate and test null hypotheses concerning adjusted treatment effects. However, common ANCOVA models have several important limitations. Firstly, as parametric models, they rely on strict distributional assumptions such as normality and variance homoscedasticity. Secondly, the commonly used ANCOVA F -test is limited to assessing the global null hypothesis of no difference in effects between groups only.

Konietschke et al.¹ addressed the first issue by proposing an alternative approach for the analysis of covariance in general factorial designs using Box-type approximation methods, which are also valid under variance heteroscedasticity.

In this paper, we extended their approach to the MCTP framework, thereby addressing the second shortcoming. This extension allows us to test not only the global null hypothesis but also individual null hypotheses and construct consistent simultaneous confidence intervals. To achieve this, we derived a small sample size approximation for the distribution of the test statistics via a multivariate t -distribution, which does not necessitate variance homoscedasticity but is also applicable for group-wise variance heteroscedasticity. Additionally, we proposed an alternative method using a wild-bootstrap approach to empirically approximate the distribution of the test statistic through resampling. This method further relaxes the model assumptions and is also valid in the presence of complete variance heteroscedasticity in the data.

To evaluate the ability of our methods to control the type-1 error rate, we conducted a simulation with various settings. As a point of comparison, we utilized an approach by Hothorn et al.⁶, which is based on general parametric models.

Our simulations demonstrated that our methods outperformed the competitor in all scenarios, particularly in small sample sizes. In complex settings with numerous groups and very few samples, only the bootstrap method outperformed the multivariate t -approximations. However, with increasing sample sizes, the type-I error rates of the methods quickly converged towards the nominal level, demonstrating their applicability even with a moderate number of samples.

Furthermore, we conducted a power analysis, which revealed that the degree of variance heteroscedasticity primarily impacts the power of the procedures. The procedures attained the highest/lowest power under homoscedasticity/complete heteroscedasticity, respectively. Except a few scenarios in which the min-method exhibited significantly lower power compared to the other methods, all methods displayed similar power.

Consequently, the general recommendation would be to employ the bootstrap method when faced with few samples or challenging scenarios such as skewed data or a large number of groups. The downside of the bootstrap method is its computational cost compared to the other methods. However, if a moderate or large number of observations ($n_i > 15$) is available per group, then the approximate methods also consistently perform well. All of the methods discussed are mean-based and thus, only applicable in metric models. MCTP for nonparametric ANCOVA models, e.g., probabilistic index models^{25,26}, will be part of future research.

ACKNOWLEDGMENTS

The work was supported by the Deutsche Forschungsgemeinschaft grant number DFG KO 4680/4-1.

Financial disclosure

None reported.

Conflict of interest

The authors declare no potential conflict of interests.

References

1. Konietschke F, Cao C, Gunawardana A, Zimmermann G. Analysis of covariance under variance heteroscedasticity in general factorial designs. *Statistics in Medicine* 2021; 40: 4732–4749.

2. Pauly M, Brunner E, Konietzschke F. Asymptotic permutation tests in general factorial designs. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 2015; 77(2): 461–473.
3. Brunner E, Dette H, Munk A. Box-type approximations in nonparametric factorial designs. *Journal of the American Statistical Association* 1997; 92: 1494–1502.
4. Mukerjee H, Robertson T, Wright F. Comparison of several treatments with a control using multiple contrasts. *Journal of the American Statistical Association* 1987; 82: 902–910.
5. Bretz F, Genz A, Hothorn L. On the numerical availability of multiple comparison procedures. *Biometrical Journal* 2001; 43: 645–656.
6. Hothorn T, Bretz F, Westfall P. Simultaneous inference in general parametric models. *Biometrical Journal* 2008; 50(3): 346–363.
7. Dunnett CW. A multiple comparison procedure for comparing several treatments with a control. *Journal of the American Statistical Association* 1955; 50(272): 1096–1121.
8. Tukey JW. Comparing individual means in the analysis of variance. *Biometrics* 1949: 99–114.
9. Tukey JW. The problem of multiple comparisons. *Multiple comparisons* 1953.
10. Hothorn T, Bretz F, Westfall P. Simultaneous Inference in General Parametric Models. *Biometrical Journal* 2008; 50(3): 346–363.
11. Zeileis A. Object-oriented computation of sandwich estimators. *Journal of Statistical Software* 2006; 16: 1–16.
12. Cao C, Pauly M, Konietzschke F. The Behrens-Fisher problem with covariates and baseline adjustments. *Metrika* 2020; 83: 197–215.
13. White H. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica: Journal of the Econometric Society* 1980: 817–838.
14. Zimmermann G, Pauly M, Bathke A. Small-sample performance and underlying assumptions of a bootstrap-based inference method for a general analysis of covariance model with possibly heteroskedastic and nonnormal errors. *Statistical Methods in Medical Research* 2019; 28: 3808–3821.
15. Hasler M, Hothorn LA. Multiple contrast tests in the presence of heteroscedasticity. *Biometrical Journal: Journal of Mathematical Methods in Biosciences* 2008; 50(5): 793–800.
16. Hasler M. Multiple contrast tests for multiple endpoints in the presence of heteroscedasticity. *The International Journal of Biostatistics* 2014; 10(1): 17–28.
17. Dilba G, Bretz F, Guiard V, Hothorn LA. Simultaneous confidence intervals for ratios with applications to the comparison of several treatments with a control. *Methods of Information in Medicine* 2004; 43(05): 465–469.
18. Dilba G, Schaarschmidt F, Hothorn LA. Inferences for ratios of normal means. *R News* 2007; 7(1): 20–23.
19. Schaarschmidt F, Sill M, Hothorn LA. Approximate simultaneous confidence intervals for multiple contrasts of binomial proportions. *Biometrical Journal: Journal of Mathematical Methods in Biosciences* 2008; 50(5): 782–792.
20. Konietzschke F, Hothorn LA, Brunner E. Rank-based multiple test procedures and simultaneous confidence intervals. 2012.
21. Friedrich S, Konietzschke F, Pauly M. A wild bootstrap approach for nonparametric repeated measurements. *Computational Statistics & Data Analysis* 2017; 113: 38–52.
22. Rubarth K, Pauly M, Konietzschke F. Ranking procedures for repeated measures designs with missing data: Estimation, testing and asymptotic theory. *Statistical Methods in Medical Research* 2022; 31(1): 105–118.

23. Gunawardana A, Konietzschke F. Nonparametric multiple contrast tests for general multivariate factorial designs. *Journal of Multivariate Analysis* 2019; 173: 165–180.
24. R Core Team . *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing; Vienna, Austria: 2021.
25. De Neve J, Thas O. A regression framework for rank tests based on the probabilistic index model. *Journal of the American Statistical Association* 2015; 110(511): 1276–1283.
26. Thas O, Neve JD, Clement L, Ottoy JP. Probabilistic index models. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 2012; 74(4): 623–671.

How to cite this article: Becher M., Hothorn L., and Konietzschke F. (2024), Unlocking Insights: Enhanced Analysis of Covariance in General Factorial Designs through Multiple Contrast Tests under Variance Heteroscedasticity, 2024.

APPENDIX

A BALANCED VS UNBALANCED WITH CONSTANT N

To compare balanced and unbalanced designs with the same N for both we also ran the following setting in addition to the ones discussed in 4. In the balanced case $n_j = 10+i$ for all groups. In the unbalanced case $\mathbf{n} = (8, 10, 12)+i$ or $\mathbf{n} = (8, 9, 10, 11, 12)+i$ for 3 and 5 groups respectively with negative pairing and equivalently for positive pairing. Figure A1 shows the results.

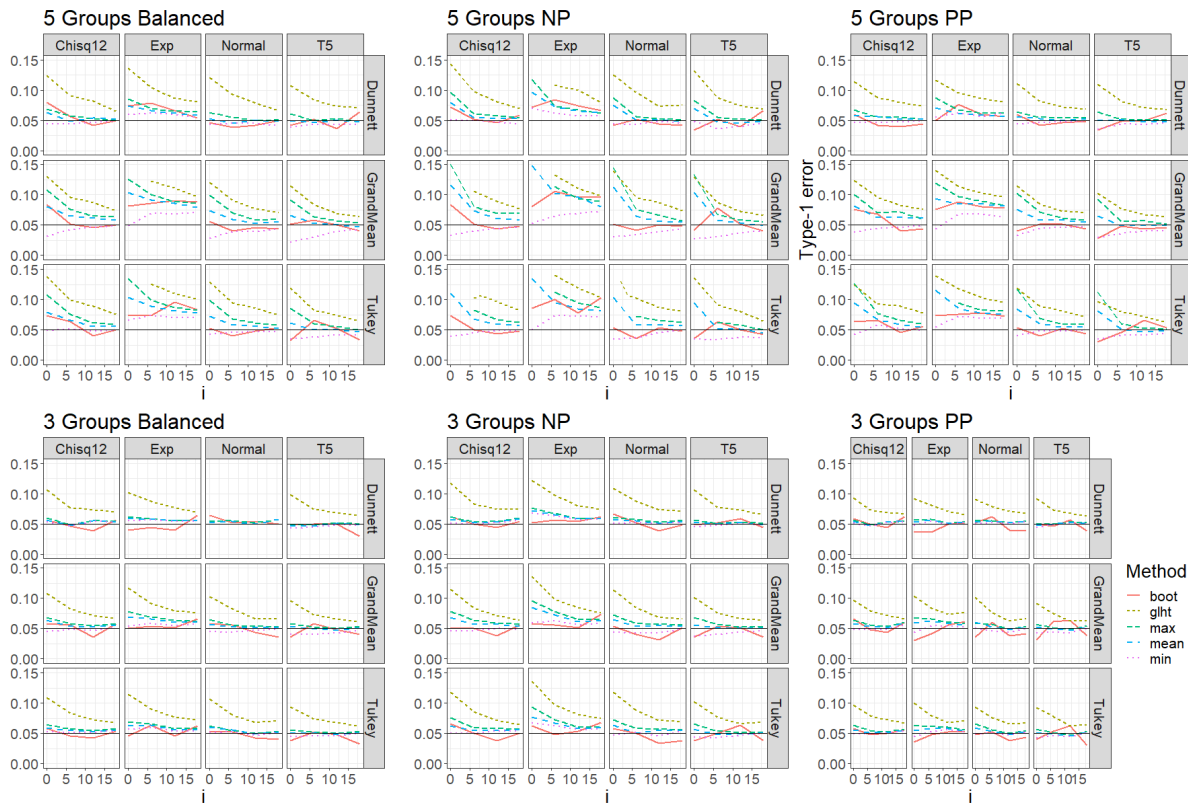


Figure A1 Type-I error simulation balanced vs unbalanced with constant N

B NEGATIVE PAIRING VS POSITIVE PAIRING

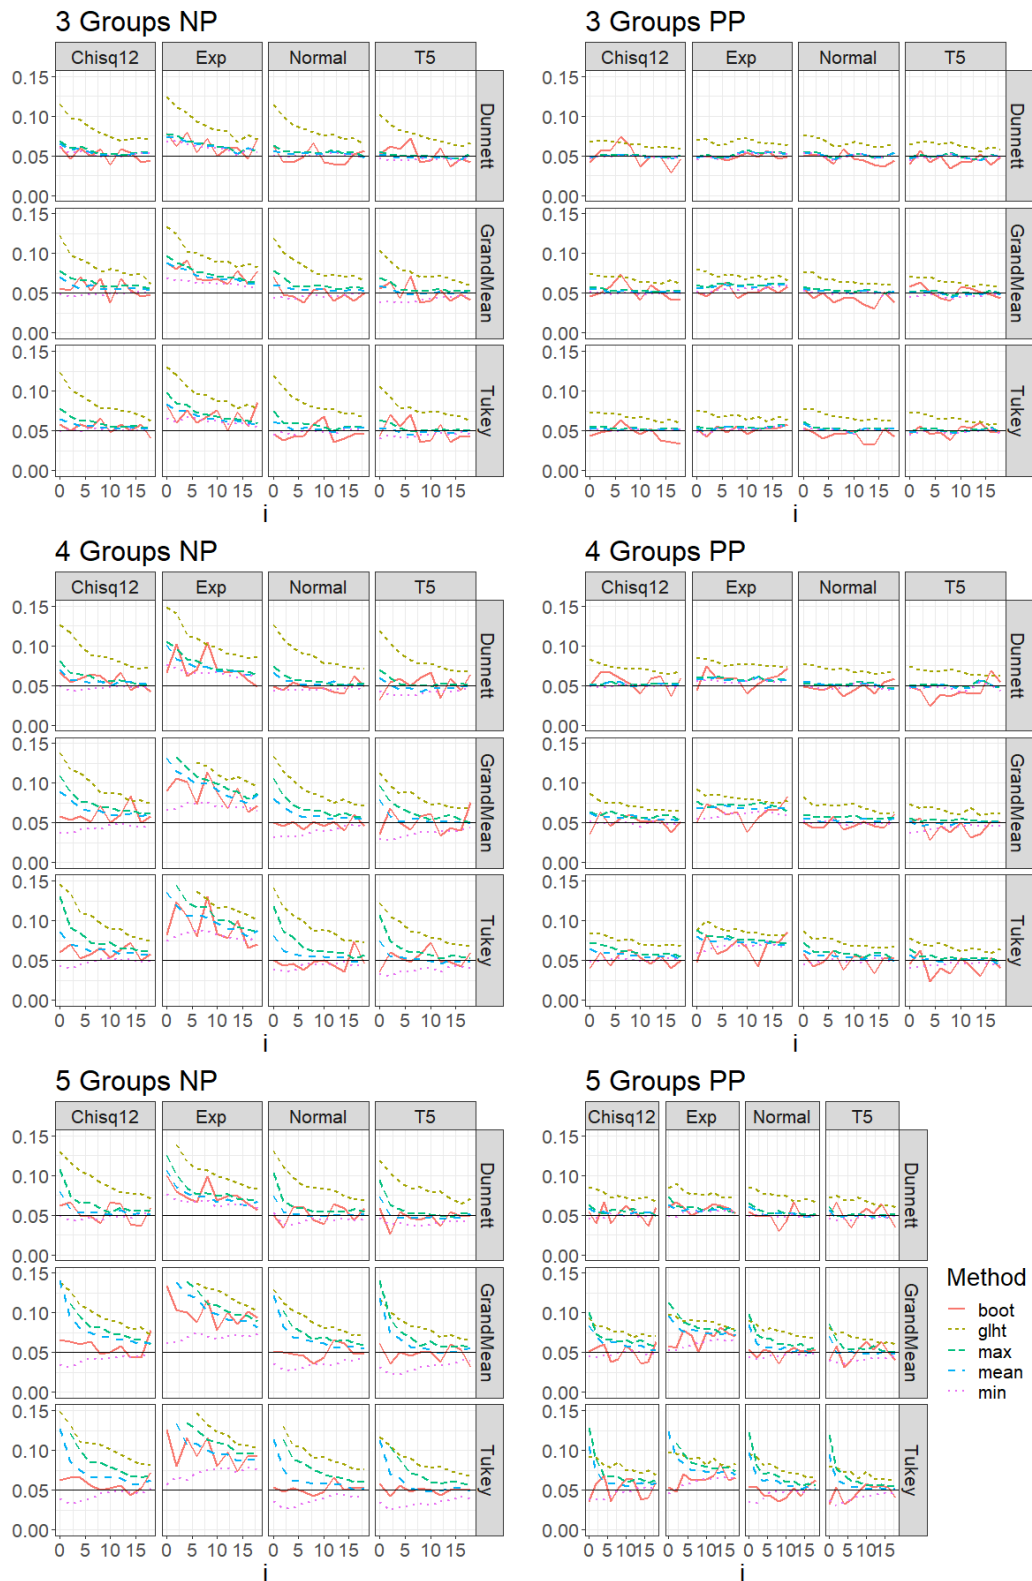


Figure B2 Type-I error simulation results negative pairing vs positive pairing

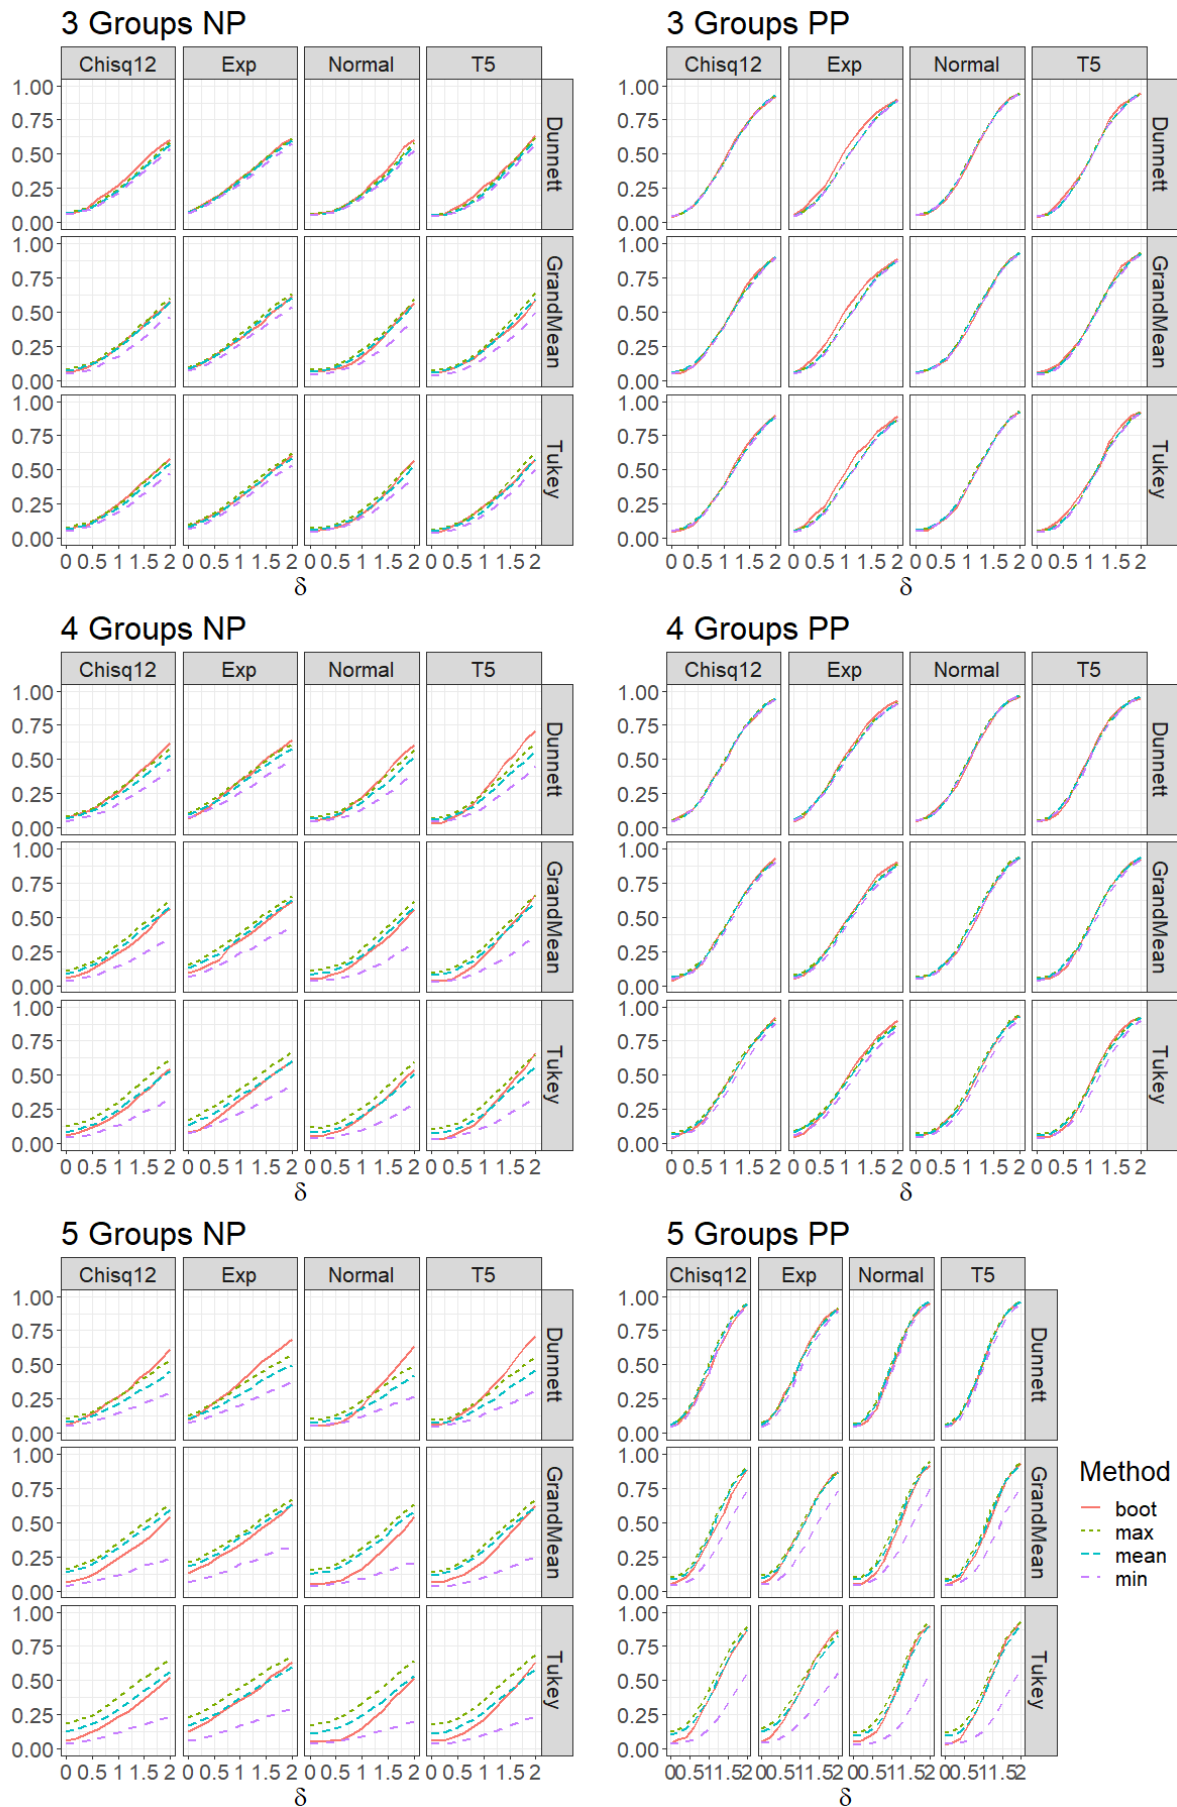


Figure B3 Power simulation results negative pairing vs positive pairing

C POWERANALYSIS FOR ALTERNATIVE 2

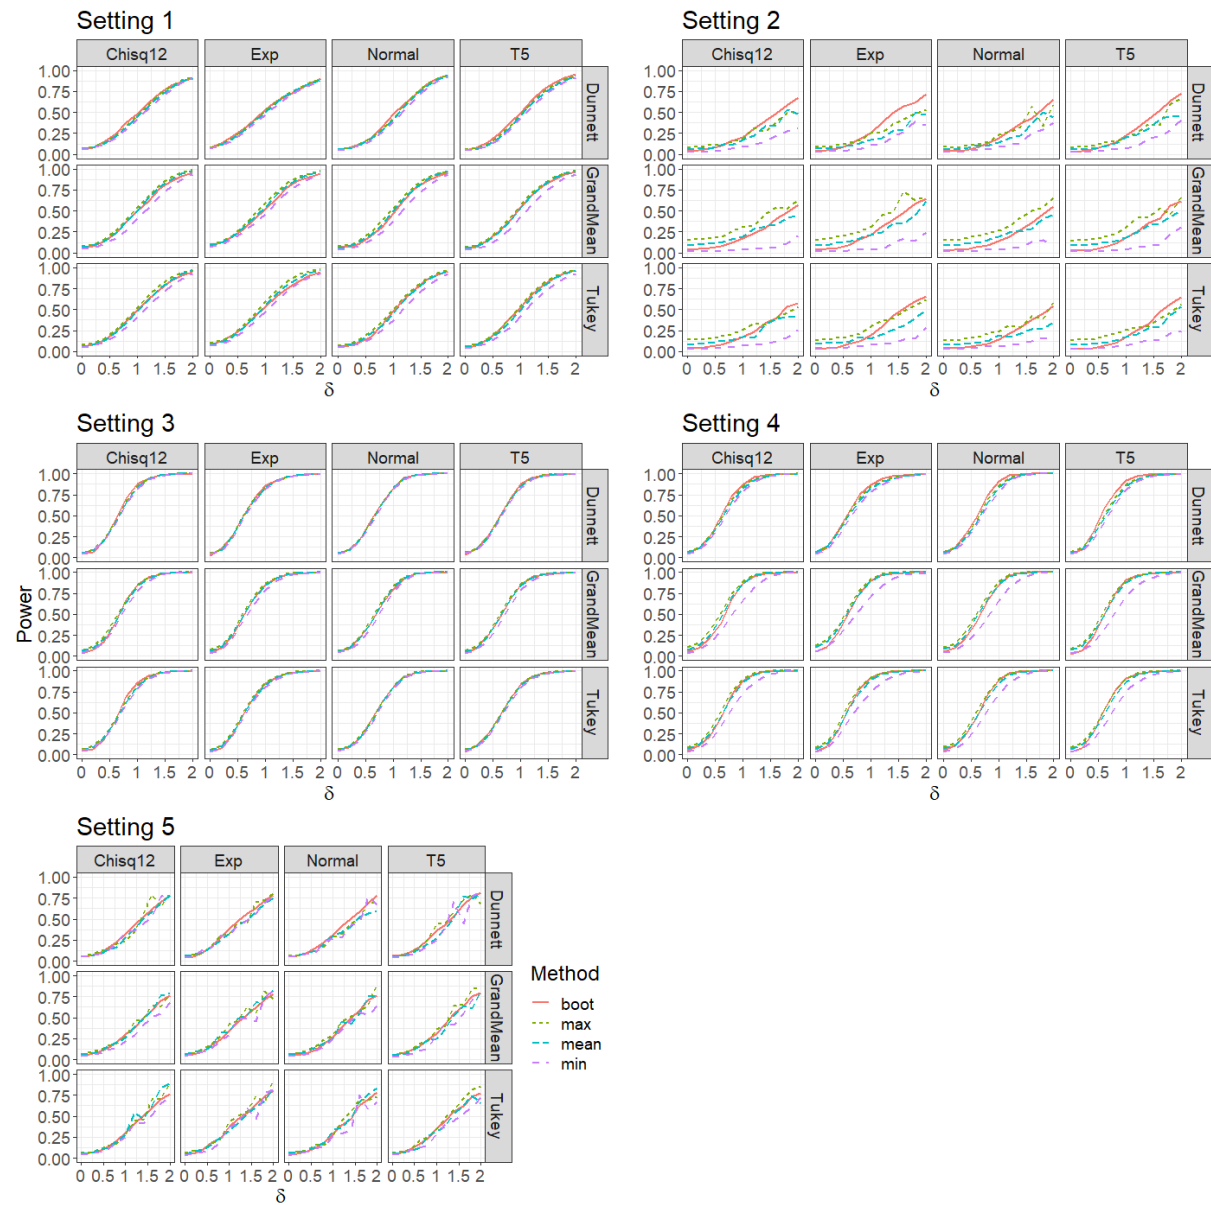


Figure C4 Power simulation results for alternative 2