

Factorial versus multi-arm multi-stage designs for clinical trials with multiple treatments

Thomas Jaki*† and Despina Vasileiou

When several treatments are available for evaluation in a clinical trial, different design options are available. We compare multi-arm multi-stage with factorial designs, and in particular, we will consider a 2×2 factorial design, where groups of patients will either take treatments A, B, both or neither. We investigate the performance and characteristics of both types of designs under different scenarios and compare them using both theory and simulations. For the factorial designs, we construct appropriate test statistics to test the hypothesis of no treatment effect against the control group with overall control of the type I error. We study the effect of the choice of the allocation ratios on the critical value and sample size requirements for a target power. We also study how the possibility of an interaction between the two treatments A and B affects type I and type II errors when testing for significance of each of the treatment effects. We present both simulation results and a case study on an osteoarthritis clinical trial. We discover that in an optimal factorial design in terms of minimising the associated critical value, the corresponding allocation ratios differ substantially to those of a balanced design. We also find evidence of potentially big losses in power in factorial designs for moderate deviations from the study design assumptions and little gain compared with multi-arm multi-stage designs when the assumptions hold. © 2016 The Authors. *Statistics in Medicine* Published by John Wiley & Sons Ltd.

Keywords: factorial design; multi-arm multi-stage designs (MAMS); familywise error rate

1. Introduction

Despite the increased understanding of many diseases unfolding in recent years and the increased spending in research and development, the US Food and Drug Administration identified a slowdown in the approval of innovative medical therapies [1] and called for methods that achieved reliable results more quickly. It is well known that clinical trials are expensive [2, 3], and that only a small proportion (about 10%) of new drugs in phase I trials reaches the market [4]. Moreover, phase II clinical trials designed to assess the treatment's therapeutic capacity and safety to warrant further testing in a phase III trial do not necessarily succeed in identifying potentially effective treatments, because on average only about 50% of the large phase III confirmatory trials are successful [5]. Studying two or more treatments simultaneously within one trial is more efficient compared with the traditional separate treatment evaluation, as both the sample size and the duration of the trial will be smaller. Combination therapies are also of interest especially when monotherapies fail to prove their effectiveness or when agents are known to have different mechanisms of efficacy. Examples of such therapies include medications for the treatment of hypercholesterolemia where there is a class of cholesterol-lowering agents that inhibit the intestinal absorption of cholesterol and another class (statins) which inhibit cholesterol biosynthesis [6].

Couper *et al.* [7] explicitly discuss the use of factorial designs in order to investigate interactions in combination treatment studies. A factorial experiment is one that two or more experimental factors are

Medical and Pharmaceutical Statistics Research Unit, Department of Mathematics and Statistics, Lancaster University, Lancaster, U.K.

*Correspondence to: Thomas Jaki, Medical and Pharmaceutical Statistics Research Unit, Department of Mathematics and Statistics, Lancaster University, Lancaster, U.K.

†E-mail: t.jaki@lancaster.ac.uk

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

studied simultaneously [8]. Specifically a 2×2 factorial model is one that has two factors and two levels for each factor. The core assumption for the accurate estimation of the main treatment effects is the absence of an interaction between the two treatments which directly implies that the combination of both treatments has a truly additive effect.

As an alternative to factorial designs a multi-arm multi-stage (MAMS) trial design (e.g. [9, 10]) could be used for testing a number of new agents and their combinations at the same time. A multi-arm (MA) design is a clinical trial design that allows the simultaneous assessment of a number of experimental treatments, which can be either different treatments or combinations of treatments, against a single control arm [11, Chapter 16]. MAMS designs are a class of the MA designs which utilise the group sequential methods [12–14] thus allowing multiple looks on the data.

The objective of this paper is to compare 2×2 factorial designs, where groups of patients will either take treatment A, B, both treatments or none, to MA designs where the three arms (A, B and A with B) are tested against control. The issues that need to be considered when planning a factorial trial have for example been discussed in the literature by Montgomery *et al.* [15] and those arising when applying MAMS methodology have been addressed in [16, 17]. Our aim is to find the situations when the use of one design is preferable to the other in terms of the power or the sample size requirement.

2. Example

To illustrate the difference between a 2×2 factorial design and an MA design, we use a study on the clinical effectiveness of manual or exercise physiotherapy or both, in addition to usual care for patients with osteoarthritis (OA) of the hip or knee [18]. Manual therapy is intended to modify the quality and range of the target joint and improve musculoskeletal function and pain. Exercise therapy is used for muscle strengthening, stretching and neuromuscular control and has been shown to be effective in the increase of physical function and pain reduction. However, there is little evidence regarding the effectiveness of manual treatment and the long-term effectiveness of exercise therapy. The study has been designed as a 2×2 factorial randomised controlled trial, where 206 people are equally allocated to receive one of the following interventions: usual care, manual physiotherapy, exercise physiotherapy or manual and exercise physiotherapy, with each of the physiotherapies administered in nine treatment sessions each lasting 50 min. The primary outcome was change in the Western Ontario and McMaster osteoarthritis index (WOMAC) at 1-year follow-up. The primary endpoint was assumed to be normally distributed, and the trial sample size was chosen to detect a difference in WOMAC points of $\Delta = 28$ for each main effect assuming a standard deviation, $\sigma = 50$. Using a two-sided type I error of 5% ($\alpha = 0.05$) – it was found that a total of 180 participants are necessary to detect a main effects difference of a comparison in the margins with 95% power, that is comparing the presence against the absence of manual therapy or exercise therapy, a within the table comparison of all of the active interventions versus control with 75% – that is when comparing manual, exercise or the combination therapy to control and an interaction between the interventions manual therapy and exercise therapy with 46% power. The study sample size allowed for 20% attrition and the study protocol planned for a total of 224 participants, but 11 months into the trial due to higher than anticipated retention rates sample size recalculation was performed allowing for 10% attrition, which reduced the required sample size to 200 and recruitment stopped at 206 patients. General linear regression adjusted for baseline WOMAC score, stratification variable of knee or hip OA and some pre-specified potential confounding factors at baseline was used for the primary analysis. The results showed a significant difference in manual therapy versus no manual therapy and a non-significant difference in exercise versus no exercise. However, a significant strong antagonistic interaction was found in this study questioning the appropriateness of a 2×2 factorial design.

3. Methods

In this section, we describe the methods used to compare factorial and MA designs in clinical trials with normally distributed endpoints in terms of sample size and power. We focus on superiority clinical trials where the question of interest lies in determining the efficacy of treatments A, B and their combination A and B together, denoted AB. Let Y be the variable that measures the endpoint of interest, and let n denote the sample size. Subscripts A, B, AB and 0 correspond to patient on treatment A, B, AB and control, respectively. Assuming that the response variables are normally distributed, we model $Y_j \sim \mathcal{N}(\mu_j, \sigma^2)$ with $j = A, B, AB, 0$, where μ_j is the mean effect of the response to treatment or control. Correspondingly n_j is the number of patients allocated to a treatment or control. The one sided global null hypothesis

family (H_0) to be tested is

$$H_0 = \{H_{0A} : \mu_A \leq \mu_0, H_{0B} : \mu_B \leq \mu_0, H_{0AB} : \mu_{AB} \leq \mu_0\} \quad (1)$$

We want to control the probability of rejecting at least one true null hypothesis (the familywise error rate) at level α for both designs. Therefore, we are interested in finding a critical value k such that

$$P_{H_0}(Z_A \leq k \cap Z_B \leq k \cap Z_{AB} \leq k) = 1 - \alpha. \quad (2)$$

We also intent to compare the number of participants that need to be included in a study so that a pre-specified power level is achieved. The probability of correctly rejecting the global null hypothesis when the alternative is true should be large as this ensures a high powered test. Thus, the probability of type II error when the alternative hypothesis is true is defined by $\beta = P_{H_1}(Z_A \leq k \cap Z_B \leq k \cap Z_{AB} \leq k)$, with the design family of alternatives being $H_1 : \{H_{1A} : \mu_A > \mu_0, \text{ or } H_{1B} : \mu_B > \mu_0, \text{ or } H_{1AB} : \mu_{AB} > \mu_0\}$. This formula either helps define the relevant sample size that ensures control over the type I and type II errors with the use of numerical integration or gives the power $1 - \beta$ for a given sample size. Finally, this design alternative is specifying the union of the events $\mu_{1j} > \mu_0$, for $j = A, B, AB$ occurs, that is that any of the treatments A, B, AB have a bigger effect compared with control.

In the subsequent sections, we give the estimates of the treatment effects in a factorial design and describe our method of finding the distribution of the statistics relevant to the hypothesis testing which relate to the power of the test. Equivalently, the estimates of treatment effects in an MAMS design, as well as the statistics and distributions for the hypothesis tests and the resulting power of the tests are also provided.

3.1. Factorial design treatment effect estimation

By introducing indicator variables I_A and I_B , specifying whether treatment A or B, respectively has been administered, we can express the response variable, Y as the following linear model $Y_i = \beta_0 + \beta_1 I_{Ai} + \beta_2 I_{Bi} + \beta_3 I_{Ai} I_{Bi} + \varepsilon_i$, with the total number of participants being $n_0 + n_A + n_B + n_{AB}$, where and $i = 1, 2, \dots, n_0 + n_A + n_B + n_{AB}$. Subscript i corresponds to patient i , and Y, I_A, I_B denote the vectors of the random variable values for all study participants. Additionally, $\varepsilon_i \sim \mathcal{N}(0, \sigma^2)$ and $\text{Cov}(\varepsilon_i, \varepsilon_j^*) = 0$ for $i \neq j$.

The underlying true effect of a 2×2 factorial model in terms of the linear model parametrisation is illustrated by Table I. Therefore, the hypothesis testing regarding the treatment effects μ_j described previously may be conducted through testing linear combinations of the coefficients $\beta = (\beta_0, \beta_1, \beta_2, \beta_3)$ of the linear model.

Most factorial designs, despite the possibility that there may be an interaction between the two single treatments, make the assumption that there is no interaction present between the treatments, that is $\beta_3 = 0$. Thus, in effect, the nested linear model in Equation (3) is used.

$$Y_i = \beta_0 + \beta_1 I_{Ai} + \beta_2 I_{Bi} + \varepsilon_i \quad (3)$$

Denoting the design matrix of the linear model in Equation (3) by $\mathbf{X} = (\mathbf{1}, \mathbf{I}_A, \mathbf{I}_B)$, then the maximum likelihood estimator for the vector of parameters can be found as $\hat{\beta} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{Y}$; and the variance covariance matrix of the estimates is $\text{Cov}(\hat{\beta}) = \sigma^2 (\mathbf{X}^\top \mathbf{X})^{-1}$ [19] so that $\hat{\beta} \sim \mathcal{N}(\beta, \sigma^2 (\mathbf{X}^\top \mathbf{X})^{-1})$. The treatment effects of interest are a linear combination of these parameters (as shown in Table I). The distribution of a linear combination \mathbf{c} of the coefficients β is $\mathbf{c}^\top \hat{\beta} \sim \mathcal{N}(\mathbf{c}^\top \beta, \sigma^2 \mathbf{c}^\top (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{c})$, with $(\mathbf{X}^\top \mathbf{X})^{-1}$ given by Equation (A.1) in the Appendix.

Table I. The mean response for all treatment combinations in a 2×2 factorial experiment.

		B	
		Treatments	Presence
A	Presence	$\beta_0 + \beta_1 + \beta_2 + \beta_3$	$\beta_0 + \beta_1$
	Absence	$\beta_0 + \beta_2$	β_0

Assuming that $\beta_3 = 0$ then yields that the treatment effects for A, B and AB are β_1 , β_2 and $\beta_1 + \beta_2$, respectively and correspondingly $\mathbf{c}_A^\top = (0, 1, 0)$, $\mathbf{c}_B^\top = (0, 0, 1)$ and $\mathbf{c}_{AB}^\top = (0, 1, 1)$.

3.1.1. Joint comparison for both sole treatments and their combination to placebo. The relevant statistics for the hypothesis tests in Equation (1) are based on the formula of Equation (4).

$$Z_{\mathbf{c}_j^\top \boldsymbol{\beta}} = \frac{\mathbf{c}_j^\top \hat{\boldsymbol{\beta}}}{\sigma \sqrt{\mathbf{c}_j^\top (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{c}_j}} \quad (4)$$

The statistic to test for an effect of treatment A for $j = A, j^* = B$ and vice versa is given by $Z_A = \frac{\frac{n_0 n_A}{n_0 + n_A} (\bar{Y}_A - \bar{Y}_0) + \frac{n_A n_B}{n_B + n_A} (\bar{Y}_{AB} - \bar{Y}_B)}{\sigma \sqrt{\frac{n_0 n_A}{n_0 + n_A} + \frac{n_A n_B}{n_B + n_A}}}$ and $Z_B = \frac{\frac{n_0 n_B}{n_0 + n_B} (\bar{Y}_B - \bar{Y}_0) + \frac{n_B n_A}{n_A + n_B} (\bar{Y}_{AB} - \bar{Y}_A)}{\sigma \sqrt{\frac{n_0 n_B}{n_0 + n_B} + \frac{n_A n_B}{n_A + n_B}}}$ with \bar{Y}_j , $j = A, B, AB$ or 0 denoting the sample means of all observations in the j th treatment group. These are the statistics typically used in factorial experiments (e.g. [8]). Notice that the aforementioned statistics include information from the sole treatment arms and also from the combination arm as, under the assumption of additivity, one can extract the information about each individual arm.

To test the effect of the combination treatment, the standard test statistic $Z_{AB} = \frac{\bar{Y}_{AB} - \bar{Y}_0}{\sigma \sqrt{\frac{n_0 n_{AB}}{n_0 + n_{AB}}}}$ is used, instead of that derived using Equation (4) and \mathbf{c}_{AB} which relies heavily on the additivity of the sole treatment effects and the absence of any interaction between them. Information about the combination treatment in the full model comes only from patients on that arm.

To simplify notation and recognising that at the planning stage unequal numbers in the single treatment group are unlikely, we assume that the allocation ratios for the single treatment groups relative to the control group is r and the allocation ratio in the combination treatment group relative to control is q . Under these assumptions, $n_A = n_B = rn_0$ and $n_{AB} = qn_0$, which result to the simplified z-statistics shown by Equation (5).

$$\begin{aligned} Z_A &= \sqrt{n_0} \frac{r(r+q)(\bar{Y}_A - \bar{Y}_0) + qr(1+r)(\bar{Y}_{AB} - \bar{Y}_B)}{\sigma \sqrt{(1+r)(r+q)(r^2 + 2rq + r^2q)}} \\ Z_B &= \sqrt{n_0} \frac{r(r+q)(\bar{Y}_B - \bar{Y}_0) + qr(1+r)(\bar{Y}_{AB} - \bar{Y}_A)}{\sigma \sqrt{(1+r)(r+q)(r^2 + 2rq + r^2q)}} \\ Z_{AB} &= \sqrt{n_0} \frac{\bar{Y}_{AB} - \bar{Y}_0}{\sigma \sqrt{\frac{1+q}{q}}} \end{aligned} \quad (5)$$

3.1.2. Designing a factorial study with familywise error rate control. In a 2×2 factorial design, usually three different hypothesis are tested, and hence, it is of interest to avoid an inflation of the overall type-I error. More specifically, we wish to ensure that the familywise error rate (FWER) defined as $P(\text{reject at least one } H_{0j} \text{ incorrectly})$ is controlled at a pre-specified level α . To determine the FWER, the joint distribution of the test statistics is necessary. Because the test statistics are marginally normal, it can be shown that they jointly follow a trivariate normal distribution (Appendix B). We can use this result to find the parameters of this normal distribution under the null hypothesis and determine an overall critical value for all comparisons that controls the FWER at a pre-specified level alpha.

In particular, because $\text{Cov}(\bar{Y}_A, \bar{Y}_B) = 0$, $\text{Cov}(\bar{Y}_A, \bar{Y}_{AB}) = 0$ and $\text{Cov}(\bar{Y}_B, \bar{Y}_{AB}) = 0$, then $\text{Var}(Z_A) = \text{Var}(Z_B) = \text{Var}(Z_{AB}) = 1$. Further, because of the symmetry of the statistics Z_A, Z_B, Z_{AB} , it is easily shown that $\text{Cov}(Z_A, Z_{AB}) = \text{Cov}(Z_B, Z_{AB})$. Under the null hypothesis, the joint distribution of the statistics is a trivariate normal density with mean zero and variance covariance matrix V , which is shown in Equation (6), and the special case of a balanced design ($r = q = 1$) is given in Appendix B, Equation (B.4).

$$V = \begin{pmatrix} 1 & \frac{r(r+q)^2 + qr(1+r)^2 - 2q(1+r)(r+q)}{(1+r)(r+q)(r+2q+rq)} & \frac{\sqrt{rq}(1+2r+q)}{\sqrt{(1+r)(1+q)(r+q)(r+2q+rq)}} \\ \frac{r(r+q)^2 + qr(1+r)^2 - 2q(1+r)(r+q)}{(1+r)(r+q)(r+2q+rq)} & 1 & \frac{\sqrt{rq}(1+2r+q)}{\sqrt{(1+r)(1+q)(r+q)(r+2q+rq)}} \\ \frac{\sqrt{rq}(1+2r+q)}{\sqrt{(1+r)(1+q)(r+q)(r+2q+rq)}} & \frac{\sqrt{rq}(1+2r+q)}{\sqrt{(1+r)(1+q)(r+q)(r+2q+rq)}} & 1 \end{pmatrix} \quad (6)$$

To determine the critical value k to ensure FWER control, one can therefore numerically search for the value of k that satisfies Equation (2) using the distributional results given here.

Under the alternative hypothesis, the joint distribution of the statistics changes, and the trivariate normal is no longer centred at 0; however, the variance-covariance matrix of the distribution remains the same. The mean of the distribution under the alternative is given by Equation (7).

$$E \begin{pmatrix} Z_A \\ Z_B \\ Z_{AB} \end{pmatrix} = \begin{pmatrix} \frac{\sqrt{rn_0}((r+q)(\mu_A - \mu_0) + q(1+r)(\mu_{AB} - \mu_B))}{\sigma\sqrt{(1+r)(r+q)(r+2q+qr)}} \\ \frac{\sqrt{rn_0}((r+q)(\mu_B - \mu_0) + q(1+r)(\mu_{AB} - \mu_A))}{\sigma\sqrt{(1+r)(r+q)(r+2q+qr)}} \\ \frac{\sqrt{qn_0}(\mu_{AB} - \mu_0)}{\sigma\sqrt{1+q}} \end{pmatrix} \quad (7)$$

3.2. Multi-arm multi-stage designs

Multi-arm multi-stage designs evaluate the effect of several treatments in one trial by testing more than one hypothesis simultaneously for times up to the maximum number of stages. The special case of such a trial with three active treatments and the additional restriction of the third arm being assigned to the combination treatment is the closest equivalent to a 2×2 factorial design, where the null hypothesis is in the form of Equation (1). In essence MA trials treat each arm independently when estimating treatment effects rather than learning something about the individual treatments through the combination arm as in a factorial design. It still remains possible to make observations on the relationship between the individual treatments by reviewing their joint effect on the combination arm. Further to that, to enable a comparison between them, it is necessary to control the FWER at the same level, α , which is used for the factorial design hypothesis testing.

3.2.1. Single stage multi-arm design. Effectively the test statistics in an MA design with one stage can be viewed as the result of using formula (4) for the complete model which includes an interaction term. Specifically, the statistic for the comparison of the j th active treatment to control in one-stage design is defined as in Equation (8).

$$Z'_j = \frac{\bar{Y}_j - \bar{Y}_0}{\sigma \sqrt{\frac{n_0 + n_j}{n_0 n_j}}}, \text{ with } j = A, B, AB. \quad (8)$$

The joint distribution of the statistics under the null hypothesis is a trivariate normal density with 0 mean and variance covariance matrix $W = \begin{pmatrix} 1 & \frac{r}{1+r} & \frac{\sqrt{rq}}{\sqrt{(1+r)(1+q)}} \\ \frac{r}{1+r} & 1 & \frac{\sqrt{rq}}{\sqrt{(1+r)(1+q)}} \\ \frac{\sqrt{rq}}{\sqrt{(1+r)(1+q)}} & \frac{\sqrt{rq}}{\sqrt{(1+r)(1+q)}} & 1 \end{pmatrix}$, where $n_A = n_B = rn_0$ and $n_{AB} = qn_0$.

Using this distribution, the critical value k can be obtained in the same manner as for the factorial design. Note that this test is a Dunnett test [20].

3.2.2. Two-stage multi-arm design. Finally, we also compare the factorial design with an MA design with two stages, where we assume that $r = q$, that is the allocation ratio is the same between all active treatments and control. Extensions to more stages and other allocation ratios are possible [21, 22], but for simplicity, we focus on two-stage designs only. A two-stage design is a sequential design where by one is allowed to examine the data at a specific time point or after a defined number of patients have been followed up, based on a stopping rule derived from repeated significance tests. Group sequential designs allow for early stopping of the trial, either because of efficacy or futility, whilst still fully controlling the pre-specified type I error [11]. At the interim analysis, the test statistics are compared against pre-determined boundaries. If at least one test statistic exceeds the upper boundary (u), the null hypothesis can be rejected and the study stopped. If the study can not be stopped for efficacy, any treatment whose test statistic falls below the lower bound (l) will be dropped from the remainder of the study. If all treatments are dropped, the study is stopped. Note that, although a similar strategy could be conceived for a factorial

design, we do not consider such a design as dropping an arm in a factorial design impacts on the arms one is still interested in.

In this multi-stage design, the set of previously stated null hypotheses in Equation (1) is potentially tested twice. Because of making multiple comparisons, we need to control the familywise error rate, α , and we use the multiple testing procedure for multiple stages described by [21] which is the multi-stage extension of the Dunnett test used for the single stage design [20]. In this type of multi-stage clinical trial designs fixing the type I error and power, is not sufficient for their full specification. The probability of rejecting the null hypothesis depends on the stopping boundaries at each stage which need to be specified. Lower boundaries are used to stop a treatment whose test statistic falls below the threshold, and upper boundaries are used to stop the trial when any test statistic exceeds this boundary, as a treatment that is superior to control is found. In our two-stage design, we use the O'Brien–Fleming boundary shape [23] for u_1, u_2 and a fixed lower boundary at 0, that is $l_1 = 0$. The actual type I error equation in this two-stage design is specified as

$$\alpha = 1 - \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \left[\Phi(t_2) + \Phi_2 \left(u_1 \sqrt{2} + t_2, u_2 \sqrt{2} + \frac{t_1 + t_2}{\sqrt{2}}, (\sqrt{2})^{-1} \right) \right. \\ \left. - \Phi_2 \left(t_2, u_2 \sqrt{2} + \frac{t_1 + t_2}{\sqrt{2}}, (\sqrt{2})^{-1} \right) \right]^3 d\Phi(t_1) d\Phi(t_2)$$

where Φ denotes the standard normal distribution function and $\Phi_2(a, b, (\sqrt{2})^{-1})$ denotes the result of the integration of a bivariate standard normal density with covariance $(\sqrt{2})^{-1}$ over region $[a, b]$. To ensure control of the type II error, we need to be able to reject the null hypothesis if the mean treatment response is large, and for that purpose, we employ the least favourable configuration (LFC) to specify the power of the design, which is defined as the following probability $P(\text{Reject } H_{0AB} | \mu_{AB} = \Delta, \mu_A = \delta_0, \mu_B = \delta_0)$. The specific equation for our implemented two-stage design given the effect sizes Δ, δ_0 once more involves two-dimensional integration and is shown here.

$$1 - \beta = \int_{-\infty}^{\infty} \Phi \left(u_1 \sqrt{2} + t + \frac{\sqrt{n_0}}{\sigma} \Delta \right) \left[\Phi \left(t + \frac{\sqrt{n_0}}{\sigma} (\Delta - \delta_0) \right) \right]^2 d\Phi(t) \\ + \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \left[\Phi \left(t_2 - \frac{n_0}{\sigma} \delta_0 \right) + \Phi_2 \left(u_1 \sqrt{2} + t_2 - \frac{n_0}{\sigma} \delta_0, t_1 + (\Delta - \delta_0) \frac{\sqrt{2n_0}}{\sigma}, (\sqrt{2})^{-1} \right) \right. \\ \left. - \Phi_2 \left(t_2 - \frac{\sqrt{n_0}}{\sigma} \delta_0, t_1 + (\Delta - \delta_0) \frac{\sqrt{2n_0}}{\sigma}, (\sqrt{2})^{-1} \right) \right]^2 \\ \left[\Phi \left(2u_1 + t_2 \sqrt{2} - t_1 - \frac{\sqrt{2n_0}}{\sigma} \Delta \right) - \Phi \left(t_2 \sqrt{2} - t_1 - \frac{\sqrt{2n_0}}{\sigma} \Delta \right) \right] \Phi(t_1 \sqrt{2} - t_2 - 2u_2) d\Phi(t_1) d\Phi(t_2)$$

The aforesaid equations are provided here for reproducibility of our results. Detailed derivations of these equations for a general number of treatments and stages can be found in [21]. We have used the R package MAMS [24] to obtain the full specification of the design.

4. Results

In this section, we firstly explore some of the design features of factorial designs and subsequently compare these designs to MA and MAMS designs. We begin by looking into how different allocation ratios affect properties of the factorial design. Similar evaluations of MAMS designs can be found in [17]. All results in this section are based on the analytic formulae in the previous section and have been verified with 10 000 fold simulations.

A study of covariances between the test statistics of Equations (5) presented in matrix (6) demonstrates that the correlation between Z_A, Z_{AB} is always positive and that Z_A, Z_B are uncorrelated when $q = r^2$. We can show that by studying the correlation as a function of the allocation ratio r , specifically when $q = r^2$, we have $\text{Cov}(Z_A, Z_{AB}) = f(r) = \sqrt{\frac{r}{1+r^2}}$, with $r > 0$. The first derivative of this is $\frac{df(r)}{dr} = \frac{1-r^2}{2\sqrt{r(1+r^2)^{3/2}}}$ and

is used to find the extremes of $f(r)$. Setting the first derivative to 0, we find that $r = 1$ which maximises the function $f(r)$ and $f(1) = 1/\sqrt{2}$. Therefore, the information obtained about the treatment effects by statistics Z_A, Z_B is maximised, and the additional information from statistic Z_{AB} is minimised as they are correlated to the highest level, making this setup very appealing when the design assumptions are met.

4.1. The different treatment allocations combination effect on the critical value

The joint distribution of the test statistics in Equations (5) under the null hypothesis can be used to find the critical value that corresponds to the probability of committing a type I error, which we set to 5% here. In the case of a balanced design, the critical value is found to be $k = 2.028$. We investigated further how differences in the allocation ratios r, q affect the choice of the critical value. Figure 1 and Table II show the critical values for the different allocation ratio combinations of the sole experimental treatments and combination treatment. We find that the critical value varies substantially for the different allocation ratio combinations and reveals the extent to which changes in the patient recruitment ratios for each of the treatments will make it easier or more difficult to reject the global null hypothesis compared with the standard set by the balanced design.

Whilst retaining the type I error rate fixed, we can use the information in Figure 1 to find the optimal choice of q for a given r in terms of minimising the critical value, which by implication ensures that the required sample size will be small. It is worth noting that whilst r ranges, the optimal value for q is largely unaffected. It seems that for any value of the allocation ratio r to the single treatments, the optimal choice for the allocation ratio of the combination treatment is less than 1. More specifically for the most interesting range of $r \in [0.5, 2.5]$, the optimal choice for q is 0.8, when $r = 2.5$ yielding the

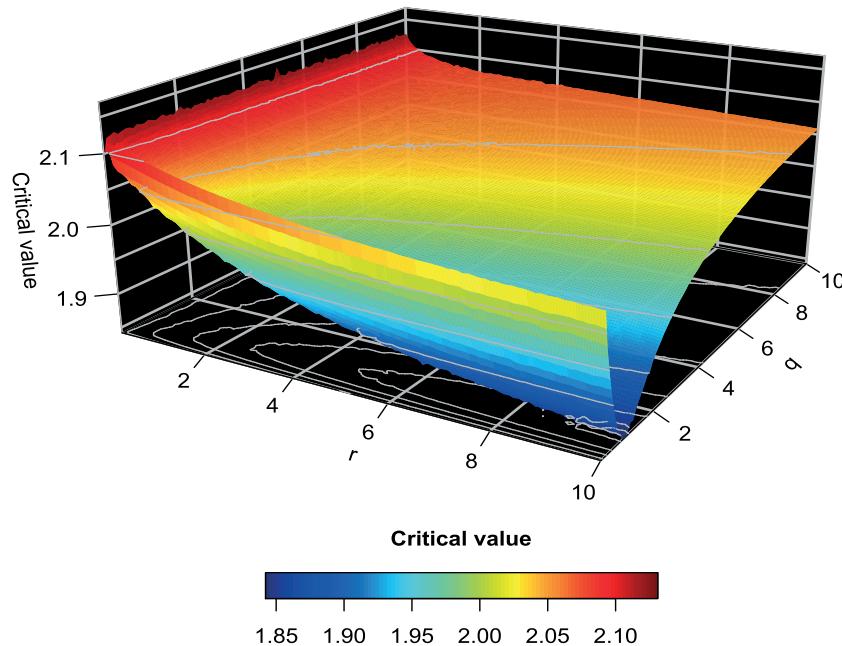


Figure 1. The critical values of a factorial design with varying allocation ratio when $\alpha = 0.05$. A r, q grid size of 0.01/0.1, respectively is used.

Table II. Critical value demonstration for specific allocation combination ratios r, q .

r	q				
	0.1	0.5	1	1.5	2
0.1	2.09	2.10	2.11	2.11	2.11
0.5	2.07	2.05	2.07	2.08	2.08
1	2.06	2.02	2.03	2.04	2.05
2	2.04	1.97	1.98	2.00	2.02

optimum critical value of 1.954. In other words, two and a half times as many patients should be allocated to treatments A and B compared with the control arm, whilst only 80% of the number on control should be devoted to the combination arm. When (numerically) searching the optimal choice of an allocation ratio when $r = q$, we find $r = q = 1.7$ to be optimal which corresponds to critical value 2.017. Note that this optimal allocation ratio is in the opposite direction with the one determined for the Dunnett test for which \sqrt{K} , with K being the number of active treatments, has been shown to be optimal [25].

4.2. Sample size when varying allocation ratio combinations

We now focus on the effect of varying allocation ratio combinations on the sample sizes when the alternative hypothesis scenarios are in the first instance consistent with the additivity assumptions of the factorial design and in the second when they are inconsistent with it, with the sole treatments having either a synergistic or an antagonistic effect. Throughout, we wish to control the FWER at 5% and seek a power of 90% to detect at least one superior treatment.

4.2.1. Allocation ratio effect in alternative scenarios consistent with the factorial design assumptions.

We consider finding the optimal sample size for any design configuration in a specific range of possible allocation ratios r, q , and we denote an interesting effect by Δ whilst a positive yet uninteresting effect is δ_0 . We deliberately consider a wide range of allocation ratios to ensure that the optimal allocation is contained in the display and also show that extreme allocation ratios do result in very large sample sizes. Figure 2 and Table III show the total sample size for different allocation ratios when one of the sole treatment arms has the interesting effect whilst the other sole treatment has an uninteresting effect ($\mu_A - \mu_0 = \Delta, \mu_B - \mu_0 = \delta_0, \mu_{AB} - \mu_0 = \Delta + \delta_0, \mu_0 = 0$) in panel (a) whilst in panel (b) both sole treatments have an uninteresting effect ($\mu_A - \mu_0 = \mu_B - \mu_0 = \delta_0, \mu_{AB} - \mu_0 = 2\delta_0$). Under the first configuration, the optimal total sample size when we assume that $\Delta = 0.5$ and $\delta_0 = 0.1$ in the case of a balanced design is that 160 patients need to be included in the study to achieve power of at least 0.9,

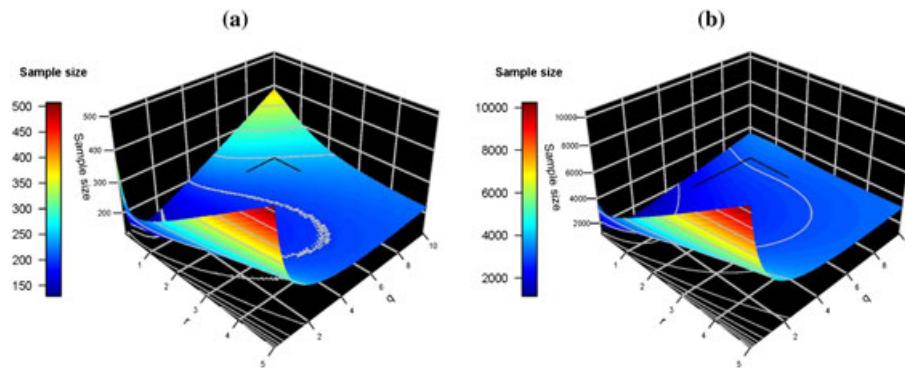


Figure 2. Plots (a) and (b) show the optimal total sample size that ensures power at least 0.9 for each combination of r, q when the alternative hypothesis scenarios are consistent with the factorial design assumptions. Specifically, in plot (a), any of the sole experimental treatment groups may have an interesting effect, with either $\mu_A - \mu_0 = \Delta$ and $\mu_B - \mu_0 = \delta_0$ or correspondingly $\mu_A - \mu_0 = \delta_0$ and $\mu_B - \mu_0 = \Delta$ whilst finally $\mu_{AB} - \mu_0 = \Delta + \delta_0$, and in (b) both of the sole experimental treatment groups have an uninteresting effect, with $\mu_A - \mu_0 = \mu_B - \mu_0 = \delta_0$ whilst $\mu_{AB} - \mu_0 = 2\delta_0$.

Table III. Total sample size for specific allocation combination ratios r, q as per Figure 2(a).

r	q				
	0.1	0.5	1	1.5	2
0.01	380	145	130	135	145
0.1	308	149	132	139	148
0.5	219	160	144	147	156
1	233	175	160	157	160
2	294	218	190	179	175

that is $P((Z_A, Z_B, Z_{AB}) \geq (k, k, k) | (Z_A, Z_B, Z_{AB}) \sim MVN((\Delta, \delta_0, \Delta + \delta_0), V) \geq 0.9$. The smallest total sample size of 129 patients is achieved when $r = 0.01$ and $q = 0.9$. Clearly, such a small allocation ratio would not be useful in practice, but it does show that, under the assumptions of a factorial design, one might as well allocate no patients to the single arms as their treatment effect can be estimated from the combination arm anyway. We see that in this scenario changes in the allocation ratios do not have a big effect on the required sample size to achieve a target power. We only see notable increases in the sample size when either both q and r are very small or q is small whilst r is big and vice versa. Additionally, if we look at the sample size requirements when δ_0 decreases, or is set to 0, we notice that the sample size requirements are a bit higher. For example, the number of patients to be recruited for the control group in the case of a balanced design is $n_0 = 43$ which means that a total sample size of 172 patients need to be included in the study to achieve power of at least 0.9. In general, the range where the required sample size is big when both q, r and q alone are small whilst r increased is a bit broader.

Looking into the alternative hypothesis when both single experimental treatments have a weak effect and the combination treatment has the addition of both of those effects (Figure 2(b) and Table IV), we find that the optimal sample size when we assumed that $\delta_0 = 0.1$ in the case of a balanced design is $n_0 = 502$ which means that a total sample size of 2008 patients need to be included in the study to achieve power of at least 0.9. We also find the configurations which result to the smallest sample size. The minimum occurs when $r = 0.01$ and $q = 1$, which gives a total sample size of 1150 patients to be included in the study, which is about half of that required by a balanced design. In fact, the smallest sample size occurs when r is at the border of the grid we are considering. In general, we notice that the smallest sample sizes occur when r, q are moderately small and that the differences between the allocation combinations in terms of sample size are substantial.

Table IV. Total sample size for specific allocation combination ratios r, q as per Figure 2(b).

r	q				
	0.1	0.5	1	1.5	2
0.01	3455	1295	1150	1198	1291
0.1	3492	1404	1224	1254	1341
0.5	3297	1894	1572	1544	1589
1	3866	2475	2008	1898	1894
2	5404	3603	2872	2591	2475

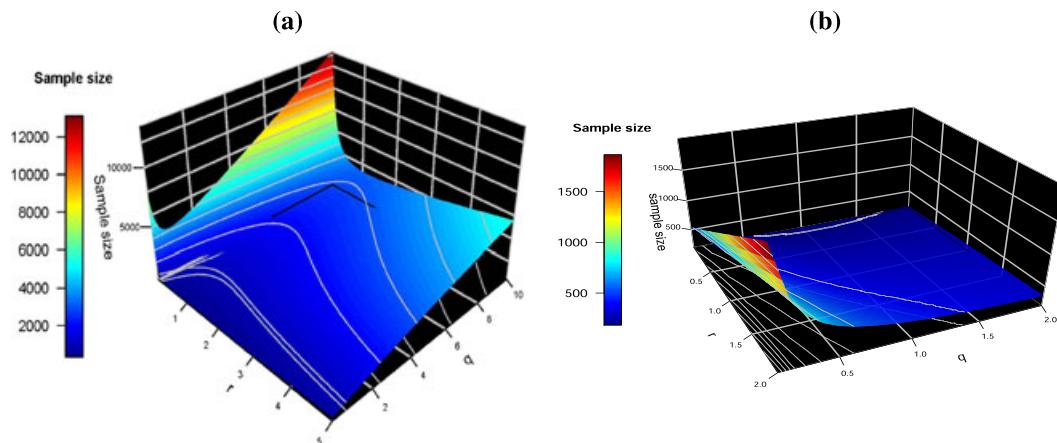


Figure 3. Plots (a) and (b) show the optimal total sample size that ensures power at least 0.9 for each combination of r, q in consistent with the alternative hypothesis scenarios where in (a) any sole treatment group may have an interesting effect, either $\mu_A = \Delta$ or $\mu_B = \Delta$ and (b) the combination treatment $\mu_{AB} = \Delta$ has the interesting effect, in a least favourable configuration set up where the remaining treatment groups have an uninteresting effect δ_0 and for the controls $\mu_0 = 0$.

4.2.2. Allocation ratio effect when alternative hypotheses are inconsistent with the factorial design assumptions. The minimum total trial sample size requirement for the most interesting range of different combinations of allocation ratios r, q , when any of the three treatments have an interesting effect ensuring that $P(\text{Reject } H_0 | H_1 \text{ is true}) \geq 1 - \beta$ and $H_{1j}^{(LFC)} : \mu_j = \Delta \text{ & } \mu_{j^*} = \delta_0 \text{ & } \mu_0 = 0$ for $j, j^* \in (A, B, AB)$, an LFC scenario is shown in Figure 3. Specifically, Figure 3 presents two scenarios for the different combinations of allocation ratios r, q and the optimal choice of sample size necessary to detect an interesting effect $\Delta = 0.5$, whilst assuming an uninteresting effect $\delta_0 = 0.1$ and standard deviation $\sigma = 1$, for a pre-specified target power $1 - \beta = 0.9$ when the assumption of additivity of the sole treatments effects to produce the combination treatment effect is not satisfied. The two alternative hypothesis scenarios are (a) $H_{1A} : \mu_A - \mu_0 = \Delta \text{ & } \mu_B - \mu_0 = \mu_{AB} - \mu_0 = \delta_0$ when the sole treatments interact antagonistically and (b) $H_{1AB} : \mu_{AB} - \mu_0 = \Delta, \mu_A - \mu_0 = \mu_B - \mu_0 = \delta_0$, when the sole treatments interact synergistically.

In Figure 3(a) and Table V, we notice that the smallest trial sizes with 90% power occur when r is around one and q is small. The optimal sample size in the case of a balanced design is when $n_0 = 176$ which means that a total sample size of 704 patients need to be included in the study to achieve at least 90% power. We also find what is the configuration which results in the smallest sample size. The minimum occurs when $r = 0.81$ and $q = 0.1$, which gives a total sample size of 326 patients to be included in the study which strikingly is less than half of the patients that need to be included when the design is balanced. Once again, a very extreme allocation ratio of $q = 0.1$ is found to be best. Because the additivity is violated, information on the combination arm can no longer readily be used to extract information on the single arms, and hence, fewer patients are allocated to this arm. We have found that the highest sample sizes correspond to the cases when the allocation ratios r, q result to a high negative correlation between the sole treatment statistics and a low correlation between the one of the sole treatment statistics and the combination treatment statistic. Conversely, the smallest sample sizes occur when the correlation between the z-statistics is high. Finally, it seems that even when the assumptions of the factorial design are not satisfied, and it becomes necessary to increase the sample size compared with when they are satisfied, the allocation ratio choices have an important bearing to the choice of sample size through their control of the correlation between the statistics.

In Figure 3(b) and Table VI, we find that in the case of a balanced factorial design, the required sample size for the control group which achieves this power is $n_0 = 81$, with a total sample size of 324, where as under the previously discussed alternative where only the sole experimental treatment had an interesting effect, the required sample size was $n_0 = 176$. The configuration that gives, in this setting, the total minimum sample size of 199 occurs when $r = 0.1$ and $q = 1$. We also notice when studying the effect of the different allocation ratios on the sample size that only when q is very small, the sample size requirement is big and indeed bigger than that of the aforementioned LFC scenario. This becomes more prominent as r increases. In all other allocation ratio combinations, the sample size requirement to

Table V. Total sample size for specific allocation combination ratios r, q as per Figure 3(a).

r	q				
	0.1	0.5	1	1.5	2
0.1	1000	1547	1944	2299	2650
0.5	353	585	804	1000	1180
1	329	508	704	877	1035
2	388	561	766	962	1145

Table VI. Total sample size for specific allocation combination ratios r, q as per Figure 3(b).

r	q				
	0.1	0.5	1	1.5	2
0.1	606	228	196	203	215
0.5	890	315	252	245	253
1	1256	431	324	297	290
2	1993	664	462	387	347

achieve the target power is small and in fact quite a lot smaller when compared with the requirement discussed in the aforementioned paragraph. Lastly, we note that in a factorial design where the main focus is on the sole treatments and testing if those have a significant effect, such an alternative hypothesis scenario would result to a high probability of rejecting the null hypothesis for both sole treatments and simulation study results have shown that is the case for the vast majority of allocation combinations r, q , apart from the cases when r and q are very small.

In conclusion, we found that for the majority of allocation ratio combinations, the power under the first type of hypothesis when the single experimental treatments have the interesting effect is less than under the second hypothesis where the combination treatment has the interesting effect. There are some allocation ratio combinations when q is small that give equal power under both types of alternative hypotheses scenarios. A setup where one experimental treatment has a strong effect whilst the combination of that treatment with another having no effect implies an antagonistic effect between the treatments so that the combination has a dissimilar effect to that of the working single treatment. It therefore makes it very difficult for the other experimental treatment to demonstrate its true effect, because of the way the effect is evaluated. Additionally, we note that in the case of the combination treatment having a strong effect in the LFC that amplifies the effect of the single experimental treatments further compared with the case that $\delta_0 = 0$ thus necessitating the inclusion of more people in the study when that is the case.

4.3. Effect of interaction

One of the fundamental assumptions used in factorial designs to find the critical value and required sample size is the additivity of the treatment effects. To explore the impact of deviating from this assumption, we investigate a scenario where the data for each of the single experimental treatments is consistent with the null hypothesis, but the data for the combination of the treatments are drawn from a distribution which includes varying degrees of interaction between the treatments implying either an antagonist or synergistic effect or no relation between them. The strength of the interaction, β_3 , ranges from -1 (strong antagonistic effect) to 1 (strong synergistic effect).

Figure 4 shows the empirical probability of rejecting the four null hypotheses (global null hypothesis and the three single treatment to control comparison) based on 10 000 simulations for a balanced factorial design with $\beta_0 = \beta_1 = \beta_2 = 0$ and different levels of interaction β_3 . An overall type I error of $\alpha = 0.05$ is used and $n_0 = n_A = n_B = n_{AB} = 50$ is chosen to give power of $1 - \beta = 0.9$ for $\Delta = 0.5$, $\delta_0 = 0.1$ and standard deviation $\sigma = 1$. The four plots relate to the global null hypothesis given in Equation (1) (top left) as well as the three individual null hypotheses comparing one arm against control. We can clearly observe a notable inflation of the type I error of the single treatment arm comparison in the cases when the treatment interact in a synergistic manner ($\beta_3 > 0$). It is also worth pointing out that despite the observations on the single arms being consistent with the null hypothesis, the chance to reject the hypothesis related to the single treatment arms is also increasing. This is because the observations on the combination arm are also contributing to the test statistics for the single treatment arms.

We, moreover, investigated the effect on the power of the hypothesis testing when the experimental treatments have an effect when individually administered but interact antagonistically when combined in a balanced design. We assumed that treatment A has a strong effect $\Delta = 0.5$ and treatment B has a weaker one $\delta_0 = 0.25$, and by using a similar simulation set-up where the rest of the design parameters are common, we find a similar pattern to Figure 4 with the curves shifted further to the left regarding the rejection probability. We found a loss of power when the treatments have in combination an antagonistic effect and a level of power close to the nominal value otherwise. Further results on the effect of the power of hypothesis testing for a broader variety of alternative scenarios are presented forwards in Section 4.4 in relation to a case study.

Finally, we have looked into what happens to the power and type one error if say two of the treatments have an interesting effect. Using simulation, we devised two scenarios were in the first one both of the experimental treatments have the interesting effect, Δ , but the combination treatment does not, having an uninteresting effect δ_0 , thus suggesting that the treatments interact antagonistically with one another, and a second one where only one of the experimental treatments as well as the combination of treatments exhibit the interesting effect, thus implying no interference between the treatments. In the first case where the effect between the experimental treatments is strongly antagonistic, the reduction in power is apparent. Specifically, we notice that the power in the hypothesis testing decreases rapidly as the allocation ratio q in the combination treatment group increases. In the second case where only one of the experimental treatments as well as the combination of treatments exhibit an interesting effect, we see that the estimated power is increased.

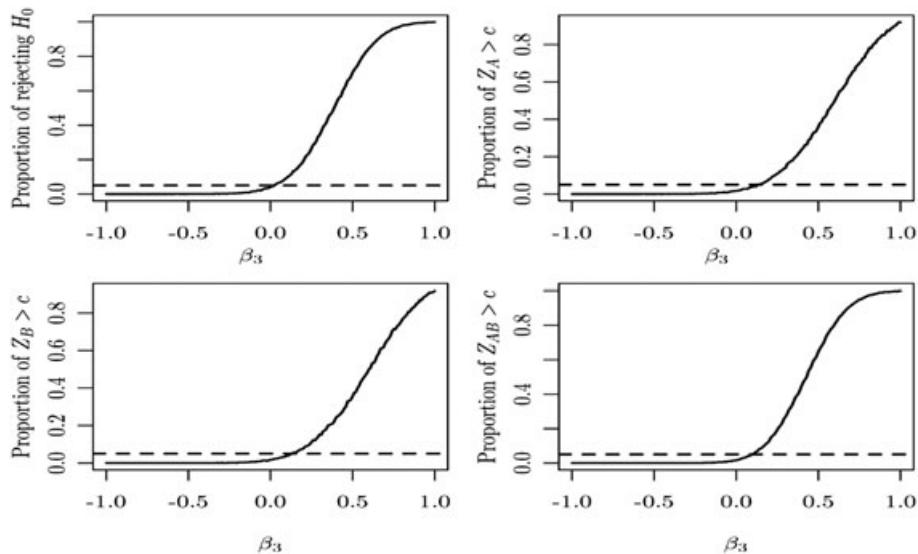


Figure 4. The Monte Carlo estimate of the probability of rejecting the null hypothesis based on 10 000 simulated samples when the sole experimental treatments are assumed not to have an effect (solid line) and the type I error reference line at 0.05 (dashed line). The top left graph shows the probability to reject the global null hypothesis whilst the remaining graphs provide the probabilities for the individual null hypotheses.

4.4. Case study

For our case study, we use the clinical trial described in Section 2 where we show that in a balanced factorial design, recruiting 45 people per treatment group ensures power of 95%. In this section, we investigate the properties of a factorial, an MA and an MAMS design, for different alternative hypothesis scenarios for this design. Using the protocol information on the clinically relevant treatment effect ($\Delta = 28$) and the effect standard deviation $\sigma = 50$ of the designed clinical trial on the use of physiotherapy on OA [26], we have also assumed that the threshold of a small uninteresting effect is one quarter of the interesting effect ($\delta_0 = 7$). The power for all hypothetical alternative scenarios is calculated by the trivariate normal with the corresponding mean and variance-covariance matrix for both the factorial and MA design. The Monte Carlo estimate of the expected value based on 10 000-fold simulations is used for a two-stage MA design with an O'Brien–Fleming efficacy boundary and a futility bound of 0. The required boundaries are specified such that the pre-specified power level $1 - \beta = 0.95$ is ensured under the alternative hypothesis, by the computation of the probability to reject the null hypothesis given that the alternative is true, and that the FWER is set to $\alpha = 0.05$.

4.4.1. Direct power comparison between a factorial, a multi-arm and a multi-arm multi-stage design. We investigate the probability of rejecting the null hypotheses for a factorial, an MA and an MAMS design under five different types of alternatives in Figure 5. Scenario (0) $\mu_0 = \mu_A = \mu_B = \mu_{AB} = 0$ is investigating the global null hypothesis where none of the treatments has an effect. Scenario (i) uses $\mu_A - \mu_0 = \Delta$, $\mu_B - \mu_0 = \mu_{AB} - \mu_0 = \delta_0$ which is consistent with a LFC setting where one of the single experimental treatments has an interesting effect, whereas the other and the combination treatment have an uninteresting effect. Scenario (ii) investigates another case that is consistent with the LFC ($\mu_A - \mu_0 = \mu_B - \mu_0 = \delta_0$, $\mu_{AB} - \mu_0 = \Delta$) whilst scenario (iii) $\mu_A - \mu_0 = \Delta$, $\mu_B - \mu_0 = \delta_0$, $\mu_{AB} - \mu_0 = \Delta + \delta_0$ is consistent with the factorial design assumptions. In the final scenario, we are considering (iv) $\mu_A - \mu_0 = \delta_0$, $\mu_B - \mu_0 = \delta_0$, $\mu_{AB} - \mu_0 = 2\delta_0$ both single experimental treatments have an uninteresting effect, and the strongest effect is in the combination treatment which may still be uninteresting and is consistent with the factorial design assumptions.

Figure 5 demonstrates that under the factorial design the probability of rejecting a hypothesis where the treatment does not have an interesting effect is unreasonably high. We notice that plot (a) of Figure 5 demonstrates how all designs perform the same way under the global null hypothesis (0) and in the case where the treatments have uninteresting effects (iv). At the same time, we can see that the factorial design is highly powered under scenario (iii) which is consistent with the factorial design assumptions, but we can see that it is greatly underpowered in scenario (i), which is a LFC setup. Specifically in this

scenario, we observe an antagonistic effect between the experimental treatments and an interaction effect which would imply that $\beta_3 = -\Delta$ in the full linear model setting described in Section 3.1. Finally, we notice that in general, the MA design is more empowered compared with the MAMS design throughout, without however noticing big differences in general. The biggest difference occurs when the factorial assumptions are met, and we assume a strong treatment effect alternative, scenario (iii). One can also observe that under scenario (ii), the factorial design has the largest overall power. This is, however, due to a large chance of rejecting the individual treatment hypothesis despite them not having an interesting

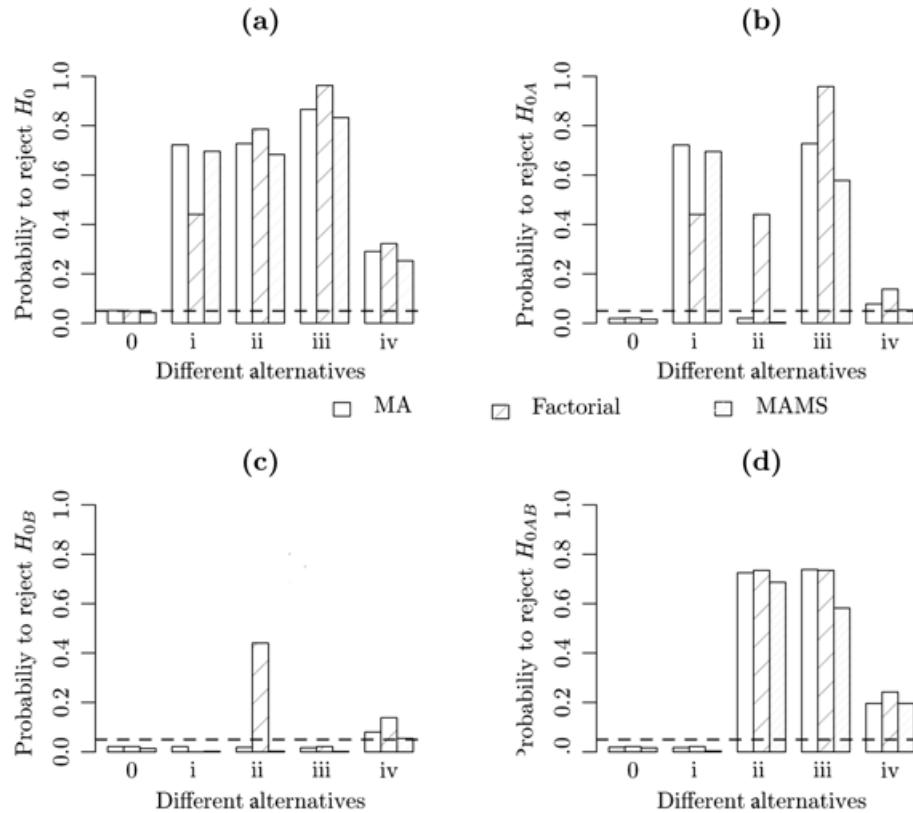


Figure 5. In each of the four plots the probability of rejecting the null hypothesis is depicted for multi-arm, factorial and multi-arm multi-stage designs for all five alternative hypothesis scenarios when (a) the global null hypothesis is rejected (b) the hypothesis for experimental treatment A is rejected (c) the hypothesis for experimental treatment B is rejected (d) the hypothesis for the combination treatment is rejected. We depict with the dashed line the type I error reference level at 0.05.

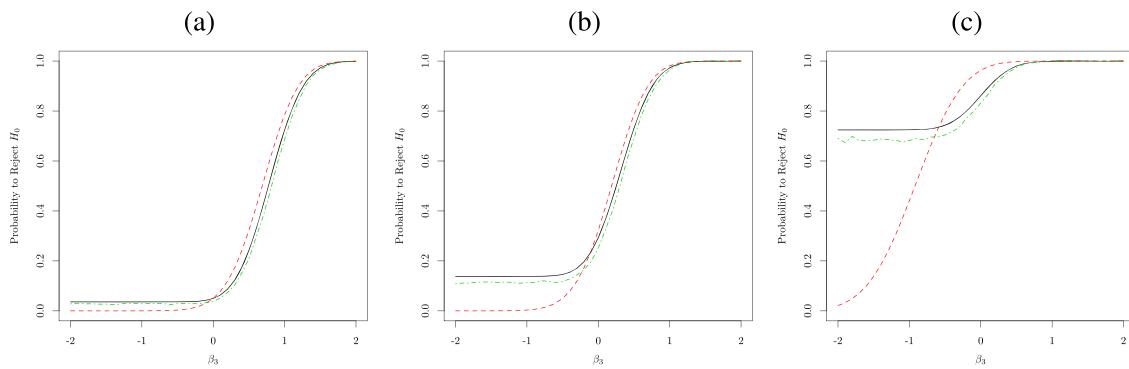


Figure 6. The probability of rejecting the global null hypothesis is depicted for a multi-arm, a factorial and an multi-arm multi-stage design with a black, a red and a green line respectively, for a varying degree of interaction between the treatments in each of the plots, for a different sole treatment level effect pair. We have assumed in (a) $\mu_A - \mu_0 = \mu_B - \mu_0 = 0$, (b) $\mu_A - \mu_0 = \mu_B - \mu_0 = 7$ and (c) $\mu_A - \mu_0 = 0$ & $\mu_B - \mu_0 = 28$.

effect. Their effect has been increased by borrowing information from the combination arm that truly has an interesting effect in this setting.

To investigate the impact of the interaction further, Figure 6 considers different combinations for the mean effect on the single treatment arms ((a) where the sole experimental treatment effects are assumed to be 0, that is $\mu_A - \mu_0 = \mu_B - \mu_0 = 0$, (b) where $\mu_A - \mu_0 = \mu_B - \mu_0 = 7$ and (c) where $\mu_A - \mu_0 = 0$ & $\mu_B - \mu_0 = 28$, whilst the interaction ranges from -2 to 2 . We find that in the case of no effect in the single treatments, the MA and MAMS design maintain the type I error level for antagonistic effects whilst the factorial design is conservative whereas for a synergistic effect the factorial design has a slightly increased power over both the MA and MAMS design. For all other scenarios, the power of both the MA and MAMS designs is larger (sometimes by a large margin) than for the factorial design when there is an antagonistic effect. For values of β_3 close to zero a slightly larger power for the factorial design is observed, whilst large synergistic effects lead to no difference between the methods. Moreover, we notice a small persistent difference in power between the MA and MAMS design for most values of β_3 , especially for strong antagonistic effects, with the exception of large synergistic effects, where there is no difference.

4.4.2. Design differences between an MAMS, a multi-arm and a factorial trial. Multi-arm multi-stage designs will in general be leading to a larger maximum sample size but smaller expected sample size compared with MA ones. The expected sample size is the expected number of patients required to detect a pre-specified treatment effect and accounts for early stopping at the first stage of the analysis. Because an optimal design has the lowest expected sample size for a given treatment effect, the expected sample size can be used as a measure of the efficiency of the design. In practice, it has been obtained through the evaluation of the expectation of 100 000-fold simulations of such designs. We do not consider a multi-stage for a factorial design as it is not possible to stop any treatment early either for efficacy or futility, because all treatments contribute to the computation of the sole treatments' effect. Our previous analysis in Section 4.4.1 revealed a similarity in terms of the relative merits between an MAMS versus factorial and MA versus factorial designs regarding the type I error and power, despite the small deficiency in power of the MAMS design when compared with the MA one. Therefore, the only remaining feature of the MAMS design that is of interest is the expected sample size.

We use the protocol information of our case study trial to see the differences in the total sample size that different choices of clinical trial designs result to, for varying degrees of interaction between the sole experimental treatments. We compare a balanced factorial design, an MA design and an MAMS design with two stages, assuming parameters $\sigma = 50$, interesting treatment effect $\Delta = 28$, uninteresting treatment effect $\delta_0 = 7$ and $\alpha = 0.05$, $1 - \beta = 0.9$ for all of them. We have chosen an MA two-stage

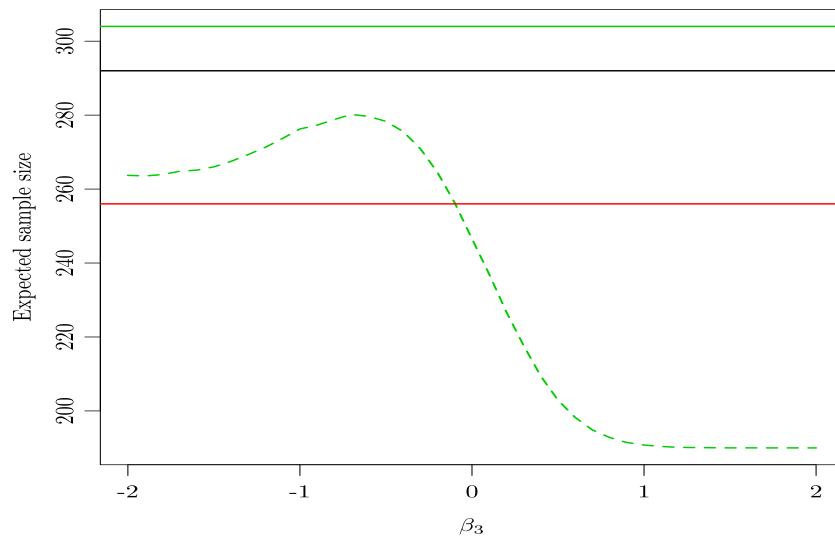


Figure 7. The total trial expected sample size estimate for a two-stage multi-arm design (green dotted line) as well as its maximum sample size (solid green line), the sample size of a balanced factorial design (red solid line) and the sample size of a multi-arm design (solid black line) for varying levels of interaction between the sole experimental treatments when $\beta_3 \in [-2, 2]$ for target power $1 - \beta = 0.9$.

design with a 0 futility boundary, an O'Brien–Fleming efficacy boundary [23] and equal sample sizes for all treatments and control for each stage. Thus by implication, the interim analysis is conducted at the half-way point of the process, and a treatment is dropped if it performs worse than the control, whereas the trial stops early if any treatment's performance exceeds the efficacy boundary, and the null hypothesis is rejected with the conclusion that this treatment is superior to control. The resulting sample size per treatment at the first stage is 38 and cumulatively at the second is 76. The critical value for the upper bound in the first stage is 2.932 and 2.073 for the second. The critical value for the lower bound in the first stage was set to 0 and for the second it is 2.073.

Figure 7 interestingly shows that there is no difference in the expected sample size of an MAMS trial and a factorial one when there is no interaction between the treatments. The sample size requirement of an MAMS study increases when the treatments interact antagonistically and decreases when they interact synergistically compared with that of a factorial design. Finally, the maximum sample size of an MAMS study is larger than the sample size required by a simpler MA trial.

5. Discussion

The apparent need for efficient clinical trial designs that quickly identify potentially effective treatments has led to the use of designs that simultaneously study many treatments. In this paper, we compared MAMS designs to a factorial with FWER control, and investigated their properties, in terms of their effectiveness in finding good working treatments through their power characteristics under different hypothetical scenarios. We found that the adoption of a factorial design which simply assumes no interaction between the sole treatments may even for modest antagonistic interaction between the sole treatments reduce the power of a hypothesis testing substantially when compared with an MA design. At the same time, we noticed an increase in the type one error rate in the cases of a moderate synergistic interaction between the sole treatments. Furthermore, we have become aware of the role that the choices of allocation ratios to each treatment have on the choice of critical value and the sample size required to achieve a certain target power in the cases both when the factorial assumptions are met or not. The level to which the necessary sample size changes for different allocation ratio combinations is greatly increased in alternative hypothesis scenarios with some degree of interaction between the sole treatments. Finally, through our case study analysis, we discovered that in the scenarios when the factorial assumptions hold, the gain of using this design over an MA in terms of power is small, whereas the losses in power when the assumptions are not satisfied are substantial. In addition, when we include a two-stage MA design in a comparison of the necessary sample size for a nominal level of target power amongst factorial and MA designs, we find that even in the absence of any interaction the expected trial sample size is the same as that of the factorial, a design that according to previous results in the presence of antagonistic interaction is greatly underpowered.

Based on these results, it is apparent that factorial designs should only really be considered instead of an MA design when the researchers are very sure that the assumption of additivity is met. As soon as there is some doubt, MAMS designs provide a robust alternative that loses little power compared with a factorial design but can gain drastically in other situations. MAMS designs also seem to be preferable to MA designs in such situations as despite their small deficiency in power, they are expected to require a much smaller sample size.

Appendix A: Nested and full linear model design matrix relevant quantities

Equation (A.1) shows the inverse of $\mathbf{X}^\top \mathbf{X}$, where $\det(\mathbf{X}^\top \mathbf{X})$ is its determinant and $\det(\mathbf{X}^\top \mathbf{X}) = n_0 n_A (n_B + n_{AB}) + n_B n_{AB} (n_0 + n_A)$ of the nested model in Equation (3).

$$(\mathbf{X}^\top \mathbf{X})^{-1} = \frac{1}{\det(\mathbf{X}^\top \mathbf{X})} \begin{pmatrix} n_A n_B + n_A n_{AB} + n_B n_{AB} & -n_A (n_B + n_{AB}) & -n_B (n_A + n_{AB}) \\ -n_A (n_B + n_{AB}) & (n_0 + n_A)(n_B + n_{AB}) & n_B n_A - n_0 n_{AB} \\ -n_B (n_A + n_{AB}) & n_B n_A - n_0 n_{AB} & (n_A + n_{AB})(n_0 + n_B) \end{pmatrix} \quad (\text{A.1})$$

Furthermore, using the information on the number of people in each treatment group, we calculate the inverse of $X^\top X$ (Equation (A.2)) of the full model presented first in Section 3.1.

$$(X^\top X)^{-1} = \begin{pmatrix} \frac{1}{n_0} & -\frac{1}{n_0} & -\frac{1}{n_0} & \frac{1}{n_0} \\ -\frac{1}{n_0} & \frac{n_0+n_A}{n_0 n_A} & \frac{1}{n_0} & -\frac{n_0+n_A}{n_0 n_A} \\ -\frac{1}{n_0} & \frac{1}{n_0} & \frac{n_0+n_B}{n_0 n_B} & -\frac{n_0+n_B}{n_0 n_B} \\ \frac{1}{n_0} & -\frac{n_0+n_A}{n_0 n_A} & -\frac{n_0+n_B}{n_0 n_B} & \frac{n_B n_{AB}(n_A+n_{AB})+n_A n_{AB}(n_0+n_B)}{n_0 n_A n_B n_{AB}} \end{pmatrix} \quad (\text{A.2})$$

Appendix B: Derivation of joint distribution of test statistics

We show that the joint distribution of $\mathbf{Z} = (Z_A, Z_B, Z_{AB})$ is a trivariate normal by expressing these dependent random variables in terms of linear combinations of the independent variables $\bar{Y}_j \sim \mathcal{N}(\mu_j, \frac{\sigma^2}{n_j})$ for $j = A, B, AB, 0$, with $n_A = n_B = rn_0$ and $n_{AB} = qn_0$, which precisely due to their independence property, are known to be jointly distributed as a multivariate normal. Specifically, for $\bar{\mathbf{Y}} = (\bar{Y}_0, \bar{Y}_A, \bar{Y}_B, \bar{Y}_{AB})$ we

have, $\bar{\mathbf{Y}} \sim \mathcal{N}_4(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, where $\boldsymbol{\mu} = (\mu_0, \mu_A, \mu_B, \mu_{AB})$ and $\boldsymbol{\Sigma} = \begin{pmatrix} \frac{\sigma^2}{n_0} & 0 & 0 & 0 \\ 0 & \frac{\sigma^2}{n_A} & 0 & 0 \\ 0 & 0 & \frac{\sigma^2}{n_B} & 0 \\ 0 & 0 & 0 & \frac{\sigma^2}{n_{AB}} \end{pmatrix}$. Thus,

$$\begin{aligned} Z_A &= d_1 \bar{Y}_A - d_2 \bar{Y}_B + d_2 \bar{Y}_{AB} - d_1 \bar{Y}_0 \\ Z_B &= -d_2 \bar{Y}_A + d_1 \bar{Y}_B + d_2 \bar{Y}_{AB} - d_1 \bar{Y}_0 \\ Z_{AB} &= d_3 \bar{Y}_{AB} - d_3 \bar{Y}_0 \end{aligned}$$

where $d_1 = \frac{(r+q)\sqrt{rn_0}}{\sigma\sqrt{(1+r)(r+q)(r+2q+rq)}}$, $d_2 = \frac{q(1+r)\sqrt{rn_0}}{\sigma\sqrt{(1+r)(r+q)(r+2q+rq)}}$ and $d_3 = \frac{q\sqrt{n_0}}{\sigma\sqrt{1+q}}$. We can express $\mathbf{Z} = \mathbf{D}\bar{\mathbf{Y}}$, using matrix $\mathbf{D} = \begin{pmatrix} -d_1 & d_1 & -d_2 & d_2 \\ -d_1 & -d_2 & d_1 & d_2 \\ -d_3 & 0 & 0 & d_3 \end{pmatrix}$ which defines the linear combination of vector $\bar{\mathbf{Y}}$. Utilising the basic property of the multivariate normal distribution (MVN) whereby any linear combination of $\bar{\mathbf{Y}}$ is also MVN, we see that $\mathbf{Z} \sim \mathcal{N}_3(\mathbf{D}\boldsymbol{\mu}, \mathbf{D}\boldsymbol{\Sigma}\mathbf{D}^\top)$. Let $\boldsymbol{\mu}_Z = \mathbf{D}\boldsymbol{\mu}$, therefore for each element of $\boldsymbol{\mu}_Z = (\mu_{Z_A}, \mu_{Z_B}, \mu_{Z_{AB}})$, we have:

$$\begin{aligned} \mu_{Z_A} &= \frac{\sqrt{rn_0}((r+q)(\mu_A - \mu_0) + q(1+r)(\mu_{AB} - \mu_B))}{\sigma\sqrt{(1+r)(r+q)(r+2q+rq)}} \\ \mu_{Z_B} &= \frac{\sqrt{rn_0}((r+q)(\mu_B - \mu_0) + q(1+r)(\mu_{AB} - \mu_A))}{\sigma\sqrt{(1+r)(r+q)(r+2q+rq)}} \\ \mu_{Z_{AB}} &= \frac{\sqrt{qn_0}(\mu_{AB} - \mu_0)}{\sigma\sqrt{1+q}} \end{aligned}$$

Equation (B.1) shows the general form of a variance–covariance matrix $\mathbf{D}\boldsymbol{\Sigma}\mathbf{D}^\top = V$. We find that $\text{Var}(Z_A) = \text{Var}(Z_B) = \text{Var}(Z_{AB}) = 1$ and $\text{Cov}(Z_A, Z_{AB}) = \text{Cov}(Z_B, Z_{AB})$, and in the following equations, we present some of the work in deriving $\text{Cov}(Z_A, Z_B)$, $\text{Cov}(Z_A, Z_{AB})$ under the null hypothesis where $\text{Cov}(Z_A, Z_B) = E(Z_A Z_B) - E(Z_A)E(Z_B) = E(Z_A Z_B)$ and $\text{Cov}(Z_A, Z_{AB}) = E(Z_A Z_{AB}) - E(Z_A)E(Z_{AB}) = E(Z_A Z_{AB})$.

$$V = \begin{pmatrix} \text{Var}(Z_A) & \text{Cov}(Z_A, Z_B) & \text{Cov}(Z_A, Z_{AB}) \\ \text{Cov}(Z_B, Z_A) & \text{Var}(Z_B) & \text{Cov}(Z_B, Z_{AB}) \\ \text{Cov}(Z_{AB}, Z_A) & \text{Cov}(Z_B, Z_{AB}) & \text{Var}(Z_{AB}) \end{pmatrix} \quad (\text{B.1})$$

Equations (B.2) and (B.3) explicitly show the computations. We also find, using that $V = D\boldsymbol{\Sigma}D^\top$, the formulae remain the same under any alternative hypothesis scenario.

$$\begin{aligned}
 \text{Cov}(Z_A, Z_B) &= E\left(\frac{rn_0[(r+q)(\bar{Y}_A - \bar{Y}_0) + q(1+r)(\bar{Y}_{AB} - \bar{Y}_B)][(r+q)(\bar{Y}_B - \bar{Y}_0) + q(1+r)(\bar{Y}_{AB} - \bar{Y}_A)]}{\sigma^2(1+r)(r+q)(r^2 + 2rq + r^2q)}\right) \\
 &= n_0\left(\frac{r(r+q)^2E[(\bar{Y}_A - \bar{Y}_0)(\bar{Y}_B - \bar{Y}_0)]}{\sigma^2(1+r)(r+q)(r+2q+rq)} + \frac{q^2r(1+r)^2E[(\bar{Y}_{AB} - \bar{Y}_B)(\bar{Y}_{AB} - \bar{Y}_A)]}{\sigma^2(1+r)(r+q)(r+2q+rq)}\right. \\
 &\quad \left.+ \frac{qr(r+q)(1+r)\{E[(\bar{Y}_A - \bar{Y}_0)(\bar{Y}_{AB} - \bar{Y}_A)] + E[(\bar{Y}_{AB} - \bar{Y}_B)(\bar{Y}_B - \bar{Y}_0)]\}}{\sigma^2(1+r)(r+q)(r+2q+rq)}\right) \\
 &= \frac{n_0[r(r+q)^2(E(\bar{Y}_0^2) - \mu^2) + q^2r(1+r)^2(E(\bar{Y}_{AB}^2) - \mu^2) + qr(r+q)(1+r)(2\mu^2 - E(\bar{Y}_A^2) - E(\bar{Y}_B^2))]}{\sigma^2(1+r)(r+q)(r+2q+rq)} \\
 &= \frac{r(r+q)^2 + qr(1+r)^2 - 2q(1+r)(r+q)}{(1+r)(r+q)(r+2q+rq)} \tag{B.2}
 \end{aligned}$$

$$\begin{aligned}
 \text{Cov}(Z_A, Z_{AB}) &= \frac{\sqrt{qn_0}E[(r(r+q)(\bar{Y}_A - \bar{Y}_0) + qr(1+r)(\bar{Y}_{AB} - \bar{Y}_B))(\bar{Y}_{AB} - \bar{Y}_0)]}{\sigma^2\sqrt{(1+r)(1+q)(r+q)(r^2 + 2rq + r^2q)}} \\
 &= \frac{\sqrt{qn_0}(E[r(r+q)(\bar{Y}_A - \bar{Y}_0)(\bar{Y}_{AB} - \bar{Y}_0)] + E[qr(1+r)(\bar{Y}_{AB} - \bar{Y}_B)(\bar{Y}_{AB} - \bar{Y}_0)])}{\sigma^2\sqrt{(1+r)(1+q)(r+q)(r^2 + 2rq + r^2q)}} \\
 &= \frac{\sqrt{qn_0}(r(r+q)\text{Var}(\bar{Y}_0) + qr(1+r)\text{Var}(\bar{Y}_{AB}))}{\sigma^2\sqrt{(1+r)(1+q)(r+q)(r^2 + 2rq + r^2q)}} \\
 &= \frac{\sqrt{qn_0}\left(r(r+q)\frac{\sigma^2}{n_0} + qr(1+r)\frac{\sigma^2}{qn_0}\right)}{\sigma^2\sqrt{(1+r)(1+q)(r+q)(r^2 + 2rq + r^2q)}} \\
 &= \frac{\sqrt{rq}(1+2r+q)}{\sqrt{(1+r)(1+q)(r+q)(r+2q+rq)}} \tag{B.3}
 \end{aligned}$$

Equation (B.4) shows the variance–covariance matrix, a balanced factorial design.

$$V = \begin{pmatrix} 1 & 0 & \frac{1}{\sqrt{2}} \\ 0 & 1 & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & 1 \end{pmatrix} \tag{B.4}$$

Acknowledgements

This work is independent research arising in part from Dr Jaki's Senior Research Fellowship (NIHR-SRF-2015-08-001) supported by the National Institute for Health Research. Funding for this work was also provided by the Medical Research Council (MR/J004979/1). The views expressed in this publication are those of the authors and not necessarily those of the NHS, the National Institute for Health Research or the Department of Health. Both authors have made equal contributions to this manuscript.

References

1. Food and DA. *Innovation or Stagnation: Challenge and Opportunity on the Critical Path to New Medical Products*. Food and Drug Administration: Washington, DC, 2004.
2. DiMasi JA, Hansen RW, Grabowski HG. The price of innovation: new estimates of drug development costs. *Journal of Health Economics* 2003; **22**(2):151–185.
3. Paul SM, Mytelka DS, Dunwiddie CT, Persinger CC, Munos BH, Lindborg SR, Schacht AL. How to improve R&D productivity: the pharmaceutical industry's grand challenge. *Nature Reviews Drug discovery* 2010; **9**(3):203–214.

4. Hay M, Thomas DW, Craighead JL, Economides C, Rosenthal J. Clinical development success rates for investigational drugs. *Nature Biotechnology* 2014; **32**(1):40–51.
5. Arrowsmith J. A decade of change. *Nature Reviews Drug Discovery* 2012; **11**(1):17–18.
6. Bays HE, Ose L, Fraser N, Tribble DL, Quinto K, Reyes R, Johnson-Levonas AO, Sapre A, Donahue SR, Group ES. *et al.* A multicenter, randomized, double-blind, placebo-controlled, factorial design study to evaluate the lipid-altering efficacy and safety profile of the ezetimibe/simvastatin tablet compared with ezetimibe and simvastatin monotherapy in patients with primary hypercholesterolemia. *Clinical Therapeutics* 2004; **26**(11):1758–1773.
7. Couper DJ, Hosking JD, Cisler RA, Gastfriend DR, Kivlahan DR. Factorial designs in clinical trials: options for combination treatment studies. *Journal of Studies on Alcohol, Supplement* 2005; (15):24–32.
8. Fleiss JL. *The Design and Analysis of Clinical Experiments*. John Wiley & Sons: New York, 1986.
9. Parmar MKB, Barthel FM-S, Sydes M, Langley R, Kaplan R, Eisenhauer E, Brady M, James N, Bookman MA, Swart A-M. *et al.* Speeding up the evaluation of new agents in cancer. *Journal of the National Cancer Institute* 2008; **100**(17):1204–1214.
10. Jaki T. Designing multi-arm multi-stage clinical studies. In *Developments in Statistical Evaluation of Clinical Trials*. Springer: Berlin, 2014; 51–69.
11. Jennison C, Turnbull BW. *Group Sequential Methods with Applications to Clinical Trials*. CRC Press: Boca Raton, 1999.
12. Royston P, Parmar MKB, Qian W. Novel designs for multi-arm clinical trials with survival outcomes with an application in ovarian cancer. *Statistics in Medicine* 2003; **22**(14):2239–2256.
13. Royston P, Barthel FM, Parmar MK, Choodari-Oskooei B, Isham V. Designs for clinical trials with time-to-event outcomes based on stopping guidelines for lack of benefit. *Trials* 2011; **12**(1):81.
14. Magirr D, Stallard N, Jaki T. Flexible sequential designs for multi-arm clinical trials. *Statistics in Medicine* 2014; **33**(19):3269–3279.
15. Montgomery AA, Peters TJ, Little P. Design, analysis and presentation of factorial randomised controlled trials. *BMC Medical Research Methodology* 2003; **3**(1):26.
16. Wason J, Jaki T. Optimal design of multi-arm multi-stage trials. *Statistics in medicine* 2012; **31**(30):4269–4279.
17. Wason J, Magirr D, Law M, Jaki T. Some recommendations for multi-arm multi-stage trials. *Statistical Methods in Medical Research* 2016; **25**(2):716–727.
18. Abbott JH, Robertson MC, Chapple C, Pinto D, Wright AA, de la Barra SL, Baxter GD, Theis J-C, Campbell AJ, MOA Trial Team. *et al.* Manual therapy, exercise therapy, or both, in addition to usual care, for osteoarthritis of the hip or knee: a randomized controlled trial. 1: clinical effectiveness. *Osteoarthritis and Cartilage* 2013; **21**(4):525–534.
19. Graybill FA. *An Introduction to Linear Statistical Models*, Vol. 1. McGraw-Hill: New York, 1961.
20. Dunnett CW. Selection of the best treatment in comparison to a control with an application to a medical trial. *Design of Experiments: Ranking and Selection* 1984:47–66.
21. Magirr D, Jaki T, Whitehead J. A generalized dunnett test for multi-arm multi-stage clinical studies with treatment selection. *Biometrika* 2012; **99**(2):494–501.
22. Jaki T, Magirr D. Considerations on covariates and endpoints in multi-arm multi-stage clinical trials selecting all promising treatments. *Statistics in Medicine* 2013; **32**(7):1150–1163.
23. O'Brien PC, Fleming TR. A multiple testing procedure for clinical trials. *Biometrics* 1979; **35**(3):549–556.
24. Jaki T, Magirr D, Pallmann P. *MAMS: designing multi-arm multi-stage studies*, 2015. <http://CRAN.R-project.org/package=MAMS> [R package version 0.7].
25. Hu F, Rosenberger WF. *The theory of response-adaptive randomization in clinical trials*, Vol. 525. John Wiley & Sons: Hoboken, New Jersey, 2006.
26. Abbott JH, Robertson MC, McKenzie JE, Baxter GD, Theis J-C, Campbell AJ, et al. Exercise therapy, manual therapy, or both, for osteoarthritis of the hip or knee: a factorial randomised controlled trial protocol. *Trials* 2009; **10**(11):6215–10.