

Efficient Adaptive Designs for Clinical Trials of Interventions for COVID-19

Nigel Stallard, Lisa Hampson, Norbert Benda, Werner Brannath, Thomas Burnett, Tim Friede, Peter K. Kimani, Franz Koenig, Johannes Krisam, Pavel Mozgunov, Martin Posch, James Wason, Gernot Wassmer, John Whitehead, S. Faye Williamson, Sarah Zohar & Thomas Jaki

To cite this article: Nigel Stallard, Lisa Hampson, Norbert Benda, Werner Brannath, Thomas Burnett, Tim Friede, Peter K. Kimani, Franz Koenig, Johannes Krisam, Pavel Mozgunov, Martin Posch, James Wason, Gernot Wassmer, John Whitehead, S. Faye Williamson, Sarah Zohar & Thomas Jaki (2020) Efficient Adaptive Designs for Clinical Trials of Interventions for COVID-19, *Statistics in Biopharmaceutical Research*, 12:4, 483-497, DOI: [10.1080/19466315.2020.1790415](https://doi.org/10.1080/19466315.2020.1790415)

To link to this article: <https://doi.org/10.1080/19466315.2020.1790415>



© 2020 The Author(s). Published with license by Taylor & Francis Group, LLC.



Published online: 29 Jul 2020.



[Submit your article to this journal](#)



Article views: 3249



[View related articles](#)



[View Crossmark data](#)



Citing articles: 5 [View citing articles](#)

Efficient Adaptive Designs for Clinical Trials of Interventions for COVID-19

Nigel Stallard^{*a}, Lisa Hampson^{*b}, Norbert Benda^c, Werner Brannath^d, Thomas Burnett^e, Tim Friede^f, Peter K. Kimani^a, Franz Koenig^g, Johannes Krisam^h, Pavel Mozgunov^e, Martin Posch^g, James Wason^{ij}, Gernot Wassmer^k, John Whitehead^e, S. Faye Williamson^e, Sarah Zohar^a, and Thomas Jaki^{*e,j}

^aStatistics and Epidemiology, Division of Health Sciences, Warwick Medical School, University of Warwick, Coventry, UK; ^bAdvanced Methodology and Data Science, Novartis Pharma AG, Basel, Switzerland; ^cThe Federal Institute for Drugs and Medical Devices (BfArM), Bonn, Germany; ^dInstitute for Statistics, University of Bremen, Bremen, Germany; ^eDepartment of Mathematics and Statistics, Lancaster University, Lancaster, UK; ^fDepartment of Medical Statistics, University Medical Center Göttingen, Göttingen, Germany; ^gSection for Medical Statistics, CeMSIIS, Medical University of Vienna, Vienna, Austria; ^hInstitute of Medical Biometry and Informatics, University of Heidelberg, Heidelberg, Germany; ⁱPopulation Health Sciences Institute, Newcastle University, Newcastle upon Tyne, UK; ^jMRC Biostatistics Unit, University of Cambridge, Cambridge, UK; ^kRPACT GbR, Sereetz, Germany; ^lINSERM, Centre de Recherche des Cordeliers, Sorbonne Université, Université de Paris, Paris, France

ABSTRACT

The COVID-19 pandemic has led to an unprecedented response in terms of clinical research activity. An important part of this research has been focused on randomized controlled clinical trials to evaluate potential therapies for COVID-19. The results from this research need to be obtained as rapidly as possible. This presents a number of challenges associated with considerable uncertainty over the natural history of the disease and the number and characteristics of patients affected, and the emergence of new potential therapies. These challenges make adaptive designs for clinical trials a particularly attractive option. Such designs allow a trial to be modified on the basis of interim analysis data or stopped as soon as sufficiently strong evidence has been observed to answer the research question, without compromising the trial's scientific validity or integrity. In this article, we describe some of the adaptive design approaches that are available and discuss particular issues and challenges associated with their use in the pandemic setting. Our discussion is illustrated by details of four ongoing COVID-19 trials that have used adaptive designs.

ARTICLE HISTORY

Received May 2020
Accepted June 2020

KEYWORDS

Adaptive trial; Group sequential design; Multi-arm multi-stage; Pandemic research; Platform trial; SARS-CoV-2

1. Background and the Value and Need for Adaptive Designs

1.1. Introduction

Well conducted high quality randomized controlled clinical trials are the gold standard when evaluating the potential of a novel intervention, a standard that should not be compromised during a pandemic. They are an essential component of an outbreak response (National Academies of Sciences, Engineering, and Medicine 2017). Ideally, a rapidly conducted trial will provide definitive evidence about an intervention, allowing its immediate wide-spread deployment in the field. A vaccine that reduces risk of infection has the potential to end an outbreak. Alternatively, a treatment that improves outcomes in those who are ill may reduce transmissions by encouraging infected individuals to visit clinics (to receive treatment), thereby reducing community spread.

1.2. Context of Trials in a Pandemic

Studies must start quickly to track with the epidemic curve and enroll enough cases. This poses a particular challenge as trials therefore need to be initiated and started before the natural history of the disease is fully understood (Baucher and Fontanarosa

2020). In the West African Ebola outbreak, clinical trials were devised, funded, and initiated in record time. However, by the time most studies opened, the incidence of new cases was dropping and most studies failed to complete enrolment (Thielman et al. 2016). A similar experience is now repeated in COVID-19 where early trials in China (e.g., Wang et al. 2020) were stopped prior to reaching the preplanned sample sizes due to dwindling patient numbers and a large number of separate trials have been initiated investigating the same treatments. Such experience has led to some feeling that treatment trials conducted during one outbreak are about finding therapies to use during subsequent episodes of the disease, although this is certainly not always the case.

In the context of a novel infectious disease (such as Ebola or COVID-19), limited experience with treatment is common and heterogeneity of patient characteristics makes randomization an important tool for establishing intervention efficacy (Dodd et al. 2016, 2019; Proschan, Dodd, and Price 2016). Additionally, careful considerations about the endpoints used in the study are required (Dodd et al. 2020). Despite these challenges, there exist recent examples of trials that immediately changed practice such as the rVSV Ebola RING vaccine trial (Henao-Restrepo et al. 2017) and PALM in Ebola (Mulangu et al. 2019).

CONTACT Lisa Hampson  lisa.hampson@novartis.com  Advanced Methodology and Data Science, Novartis Pharma AG, CH-4002 Basel, Switzerland.

*These authors contributed equally.

© 2020 The Author(s). Published with license by Taylor and Francis Group, LLC

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

1.3. Aims and Remit of This Article

The recent experience in COVID-19 is that poorly conducted studies together with a certain degree of desperation mixed with endorsement of treatments by national and international leaders have meant that treatments have become widely used despite limited evidence of either benefit or safety.

This article seeks to make the case that scientific rigor is equally important during a pandemic outbreak as it usually is (London and Kimmelman 2020). We highlight the benefits provided by the flexibility of adaptive designs. Following the recent CONSORT extension for adaptive designs (Dimairo et al. 2020), we define an adaptive design as “A clinical trial design that offers preplanned opportunities to use accumulating trial data to modify aspects of an ongoing trial while preserving the validity and integrity of that trial.” We provide a guide to the literature on statistical methods for such designs, briefly describing the main approaches and giving key references. We seek to remind readers that one must ensure that possible adaptations are considered in the final trial analysis, as some critics will otherwise raise questions about the scientific integrity of the study, and definitive trials are meant to resolve, not stir, controversy. We provide examples of COVID-19 trials that use different adaptive features and finally discuss some of the more challenging aspects of implementing adaptive designs during a major disease outbreak.

2. Examples of COVID-19 Trials With Adaptive Designs

Adaptive designs have been used (and were planned to be used) during epidemic outbreaks in the past. During the 2014 Ebola virus disease outbreak in West Africa, for example, the PREVAIL II trial (Dodd et al. 2016) used a Bayesian sequential design with non-informative priors (termed “barely Bayesian design”), while the RAPIDE development platform (Cooper et al. 2015) used frequentist group-sequential approaches.

To illustrate some of the choices to be made when implementing an adaptive design, and the issues to be considered when making these choices, this section presents four different trials in COVID-19. These range in scope considerably, including an early (phase I/II) dose-finding platform, a confirmatory (phase III) trial, and two trials of COVID-19 patients embedded within previously-running trials. In all cases, an adaptive design has been employed, the aim of which is to obtain a meaningful result from the trial as quickly as possible without any loss of scientific integrity. Our aim in including these examples is not to give a comprehensive review of COVID-19 research, or even a full description of each of these four trials, details of which can be found via the trial registration numbers given, but to highlight the nature and purpose of the adaptive design approach used in each case. More detailed descriptions of the different methodological approaches used for the adaptive designs are then given in Section 3.

2.1. A Randomized Phase I/II Platform Trial Using a Bayesian Sequential Phase II Design: The AGILE Platform (EudraCT 2020-001860-27)

The first example is of a phase I/II dose-finding platform. Before investigating a new or existing compound in COVID-19

patients in a large efficacy trial, it is essential to establish safe doses and preliminary signs of activity. The AGILE (Accelerating COVID-19 Drug Development, EudraCT 2020-001860-27) trial is a randomized seamless phase I/II trial platform in which multiple different candidate treatments, potentially in different populations, can be evaluated. An illustration of the design is given in Figure 1 with the full protocol given by Griffiths et al. (2020). Each candidate compound entering the platform undergoes dose-escalation initially to establish its safety profile. Safe doses then progress to a Bayesian adaptive group-sequential (GS, Gsponer et al. 2014) phase in which the efficacy of the doses is established. For each compound, the trial is seamless in that the efficacy data observed during the dose-escalation phase of the trial are used in the efficacy phase, and similarly safety information from the efficacy phase contributes to the safety model.

The safety of the doses is defined in terms of the risk of a dose-limiting toxicity (DLT) in the first 7 days, with dose-escalation decisions based on a randomized Bayesian model-based design (Mozgunov, Jaki, and Paoletti 2019). Cohorts of patients are randomized between the control arm (standard of care, SoC) and the highest active dose that is safe according to the safety model. Patients within each cohort are randomized 2:1 in favor of the active dose. The inclusion of control in the dose-escalation phase is motivated by the potential overlap between DLTs and symptoms associated with COVID-19 and the need to avoid labeling potential treatments as unsafe due to misclassifying nontreatment related toxicities. Every dose that is found to be sufficiently safe can proceed to the efficacy evaluation phase so that it is possible to have several parallel efficacy evaluations of different doses of the same compound.

The efficacy of a given dose of a compound is established via a randomized Bayesian GS trial with a time-to-event primary endpoint. Patients are again randomized 2:1 in favor of the active dose being studied. The endpoint depends on the population under study. For mild and moderate cases, time to negative viral titers in nose and/or throat swab within 29 days is used while time to a 2-point improvement in the WHO clinical severity score (WHO Working Group on the Clinical Characteristics of COVID-19 infection 9-point ordinal scale) within 29 days is used for patients with severe disease. After the efficacy data are observed, the dose can be (i) dropped for futility; (ii) progressed to a definitive trial; (iii) be evaluated in an additional cohort of patients, again randomized 2:1. Stopping is based on predefined stopping boundaries that have been derived under a Bayesian GS design such that the Type I error for the evaluation of each dose is controlled to be below 10%. The GS trial gains further efficiency by allowing up to 20 concurrent control observations from other compounds to be shared. The final analysis for efficacy for each arm is based on a Bayesian Cox-model with an uninformative prior for the treatment effect.

2.2. A Multi-Arm Multi-Stage (MAMS) Trial: The RECOVERY-Respiratory Support Trial (ISRCTN16912075)

The second example is the RECOVERY-Respiratory Support trial (ISRCTN16912075). This is a multi-arm multi-stage

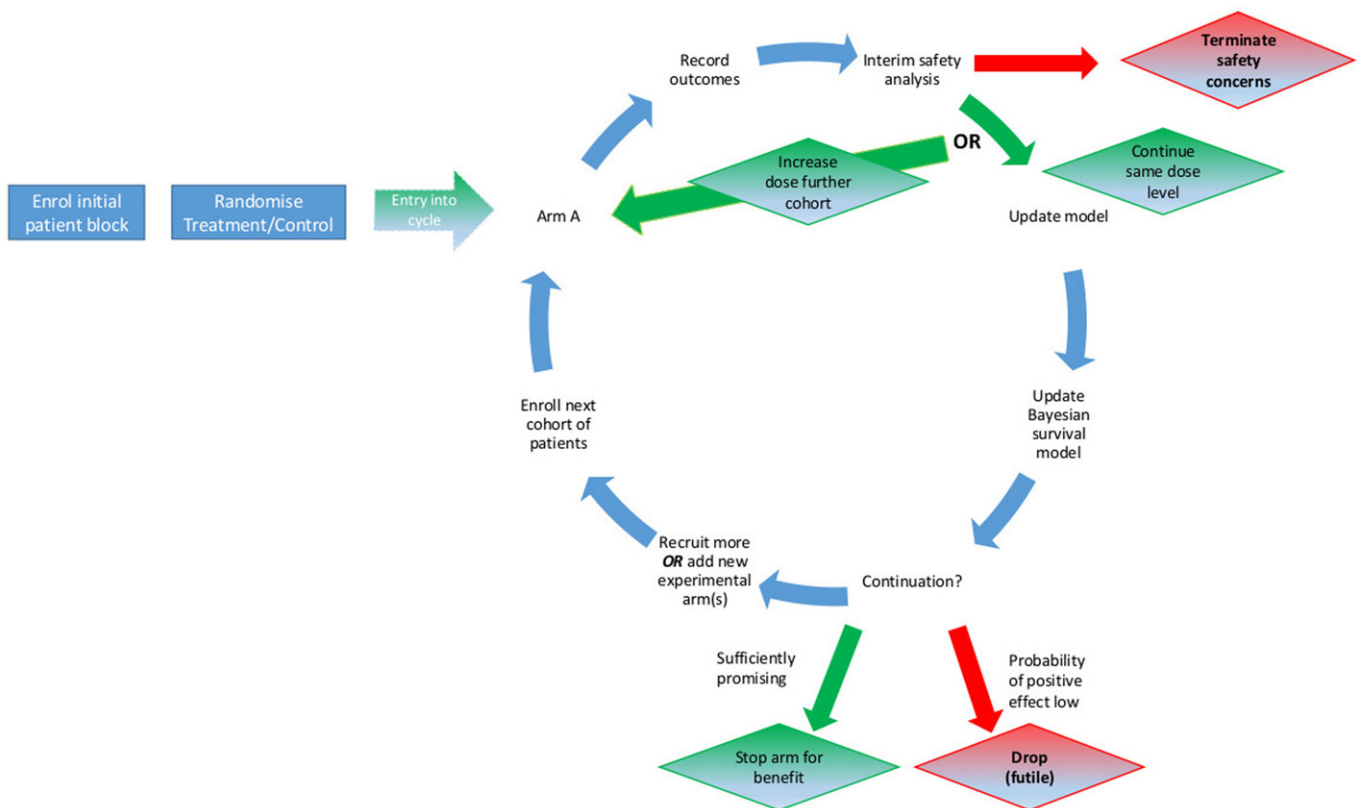


Figure 1. Illustration of AGILE design.

(MAMS) randomized clinical trial to compare three non-invasive ventilation methods of standard care (oxygen given via face masks or nasal tubes), high flow nasal oxygen (HFNO), and continuous positive airway pressure (CPAP) for patients with confirmed or suspected COVID-19 who require oxygen ($\text{FiO}_2 \geq 0.4$ and $\text{SpO}_2 \leq 94\%$). The primary endpoint is the requirement of intubation or death within 30 days.

The trial team acknowledged that not all centers would be able to randomize patients between all three arms, for example, due to the lack of availability of one of the interventions. To maximize recruitment to the trial it was decided to include centers provided they could randomize patients between standard care and at least one of HFNO and CPAP.

The inclusion of centers randomizing between different interventions gives the potential for bias, for example, due to differences between centers randomizing to all arms and those randomizing only to HFNO and standard care. To avoid this bias, analyses will only compare HFNO patients with standard care patients from centers that included HFNO patients, and only compare CPAP patients only with standard care patients from centers that included CPAP patients. The sample size will account for this.

The trial started in April 2020 with a target recruitment of 4002 patients from 40 centers. This would give 90% power to detect a reduction in the proportion of patients requiring intubation or dying from 15% to 10%. To achieve this sample size, the trial is expected to be completed within 18 months. Given the urgent need for effective treatments for COVID-19 patients and the potential long duration of the trial, the investigators wanted to stop the study early if either CPAP or HFNO was shown to be more effective than standard care, or to drop an

arm or stop the study completely if one or both experimental arms was not sufficiently promising. This ability was achieved through the planning of interim analyses on a monthly basis. A formal GS stopping rule for efficacy of either experimental arm over standard care was proposed based on the alpha-spending function approach (Kim and Demets 1987), as described in Section 3.3.

The stopping rule is applied separately for each pairwise comparison with standard care. A conventional one-sided overall Type I error rate of 0.025 was used, and the final analysis will take account of the GS design using the approach proposed by J. Whitehead (1997). No adjustment is planned for the multiple comparisons with standard care. No futility boundary was specified, though the trial data monitoring committee (DMC), who will monitor the data monthly, will be able to recommend that an experimental arm be dropped or the trial be stopped for futility at any interim analysis.

2.3. Embedding a Trial Within a Trial: The CAPE-Covid and the CAPE-Cod (Community-Acquired Pneumonia: Evaluation of Corticosteroids) Studies (NCT02517489)

The third example shows how an existing trial was modified to include an embedded trial of COVID-19 patients, with this embedded trial given a GS design. The CAPE-Cod trial (NCT02517489), assessing the efficacy of hydrocortisone to improve 28 day survival in intensive care units (ICU) for patients suffering severe community-acquired pneumonia (CAP), was active and recruiting patients at the beginning of the COVID-19 pandemic. As SARS-CoV-2 pneumonia was not an exclusion criterion of the CAPE-Cod trial, centers started to




Figure 2. REMAP-CAP trial structure for COVID-19 patients (adapted from a webinar by Berry Consultants, <https://www.youtube.com/watch?v=mhPsj1j3hik>).

include COVID-19 infected patients in the study. As the clinical characteristics of the two indications differ, trial stakeholders have decided to study the COVID-19 patients by embedding a specific study considering the COVID-19 indication only, with the CAPE-Cod study interrupted temporarily to account for patients' heterogeneity in disease characteristics and evolution.

A GS design using the alpha-spending approach (Kim and Demets 1987) was used for the embedded COVID-19 study. This was chosen for its simplicity in terms of training the local methodological team, availability of software and medical community acceptance of study results. A conservative stopping boundary was selected for efficacy with an aggressive boundary for futility; that is stopping for superiority only if the evidence is particularly strong while stopping early if the experimental treatment is not sufficiently promising. This will be adjusted for in the final analysis. A blinded sample size re-estimation was included in the study protocol because of uncertainty about the mortality rate for placebo patients with COVID-19, with estimates ranging from 26% to 73% (Grasselli et al. 2020; Ruan et al. 2020).

2.4. Evaluating COVID-19 Treatments in an Ongoing Adaptive Platform Trial: The REMAP-CAP Trial (NCT0273570)

The final example is of a trial that was designed to be able to adapt to an acute pandemic need. Motivated by the 2009 swine flu (H1N1) pandemic, the Randomized, Embedded, Multifactorial Adaptive Platform trial for Community-Acquired Pneumonia (REMAP-CAP, NCT0273570) seeks to: (i) evaluate multiple interventions to improve outcomes of patients admitted to an ICU with severe CAP, and (ii) provide a platform that can respond rapidly in the event of a respiratory pandemic.

The primary endpoint for COVID-19 patients entering REMAP-CAP is a composite of in-hospital death and the number of ICU free days over 21 days, thus forming a 23-point ordinal scale, that is, -1 (death), 0, 1, ..., 21. Adult patients with suspected or proven COVID-19 are enrolled into the COVID-19 stratum of REMAP-CAP, on an open-label basis, and classified into severe (in ICU) or moderate (hospitalized but not in ICU) disease states.

Within each disease state, there are multiple domains, each of which is a set of mutually exclusive interventions. An example is the antiviral domain comprising four interventions: (1) no antiviral; (2) lopinavir/ritonavir; (3) hydroxychloroquine; (4) lopinavir/ritonavir and hydroxychloroquine combination. The trial is adaptive so that new domains or interventions can be added at any time. Sites do not necessarily need to include all of the available domains or interventions.

Each patient is randomized to a regimen—one intervention from every domain—so that they receive multiple interventions simultaneously; this forms the multifactorial component of the trial design (Figure 2). By randomizing patients to multiple interventions, only a few receive no active treatment. Response-adaptive randomization is used to allocate patients to regimens so that more patients receive the most promising regimen(s) as the trial progresses.

Interim analyses occur frequently and are currently expected to take place weekly, although this will fluctuate with the rate of enrolment. They are also used to detect superiority, inferiority, or equivalence of interventions. Depending on regional differences or treatment availability, for example, an intervention may not necessarily be dropped if it is needed in the pandemic. All inferences in this trial are based on a Bayesian statistical model, which estimates the posterior probability of the primary endpoint for each regimen and takes into account the variation in outcomes by region, stratum, disease state, age group, and time since the start of the trial. The full statistical analysis plan is available at <https://www.remapcap.org/protocol-documents>.

3. Statistical Principles of Adaptive Design

3.1. Motivation for Using Adaptive Design Methods

Traditionally, clinical trials testing one or more null hypotheses have followed nonadaptive designs where all details of the design are specified ahead of time. In recent decades, however, there has been a rise of adaptive designs for clinical trials in all phases of drug development (Pallmann et al. 2018) and, in response, regulatory authorities have issued several guidelines for adaptive designs (European Medicines Agency 2007; Food and Drug Administration 2019).

Suppose we want to perform a clinical trial to test whether an experimental treatment is superior to control (placebo or current SoC). Measuring the benefit of the novel treatment over control by θ , we want to test the one-sided null hypothesis $H_0: \theta \leq 0$ versus the alternative $H_1: \theta > 0$. To claim control of the Type I error rate for the trial, the probability of incorrectly rejecting H_0 must not exceed a prespecified level α . The need to maintain strict control of the Type I error rate is well established, particularly in the context of late-phase studies (European Medicines Agency 1998) where a typical choice for α is 0.025. In many situations, however, clinical trials have several objectives. For example, trialists may want to compare several experimental treatments or active doses with control, or evaluate the effect of a novel treatment in several patient subgroups. In this setting, a trial begins testing m null hypotheses $H_{0i}: \theta_i \leq 0$, for $i = 1, \dots, m$, and often is designed to maintain strong control of the family-wise error rate (FWER) at some level α , which is achieved if the probability of rejecting one or more true null hypotheses does not exceed α . The requirement to control the FWER is similar to that of controlling the Type I error rate in that it restricts the probability of making a false claim of efficacy.

Why should one use an adaptive design in the first place? When planning a clinical trial there might be (substantial) uncertainty about important design parameters, for example, which effect sizes and value of the outcome variance are realistic and in which study population. If the study population is set too broadly, or the assumed variance is too small, the study will be underpowered. To address such uncertainties, it is tempting to learn from accumulating data and adapt if necessary. Adaptations of interest for COVID-19 trials may include early stopping for futility or efficacy at an interim analysis, sample size reassessment, dose-group selection, change of allocation ratio or modification of the study population in a data dependent way. Emerging knowledge of the natural history of COVID-19 can also make the choice of a primary endpoint for a trial difficult (Dodd et al. 2020). Adaptive methods are available to enable a change in endpoint (Kieser, Bauer, and Lehman 1999), though it is important to ensure that the results of a trial will be sufficiently convincing to the scientific community when making such a change.

MAMS designs (e.g., Stallard and Friede 2008; Magirr, Jaki, and Whitehead 2012; Wason et al. 2016) are of great interest for COVID-19 trials. They increase efficiency by comparing multiple experimental treatments against a shared control group (Parmar, Carpenter, and Sydes 2014), while the adaptive design allows early stopping of nonpromising (or highly efficacious) arms using prespecified decision rules (Wason and Jaki 2012).

Platform trial designs (Meyer et al. 2020) allow new experimental treatments to be added as the trial progresses and often start as, or become, MAMS trials. Platform trials provide notable operational efficiency (Schiavone et al. 2019) since evaluation of a new treatment within an existing trial will typically be much quicker than setting up a new trial. This flexible approach enabled the REMAP-CAP trial described in Section 2.4 to rapidly be adapted to study COVID-19 patients. Platform trials are cited as one potential approach for evaluating drugs intended to treat or prevent COVID-19 in a recent guidance (Food and Drug Administration 2020) and several other platform trials, such as the SOLIDARITY trial (ISRCTN83971151) and the RECOVERY trial (ISRCTN50189673, www.recoverytrial.net), are currently running.

For a more detailed review of adaptive clinical trial designs, we refer the interested reader to Bauer et al. (2016) and Pallmann et al. (2018), and the references therein. For a technical description of the methods for confirmatory clinical trials, we refer to the book of Wassmer and Brannath (2016). For an overview of Bayesian adaptive methods, which have attracted much attention in the context of exploratory trials, see Berry et al. (2010).

For all adaptive design approaches, data dependent changes may introduce bias and, particularly for late stage trials, appropriate statistical methods must be used to preserve the Type I error rate. More details on methods to ensure error rate control is maintained are given in the next sections, with estimation methods discussed in Section 3.5.

3.2. Type I Error Rate Control in Adaptive Designs

All regulatory guidelines emphasize that, for confirmatory clinical trials following an adaptive design, control of the Type I error rate is paramount. This requirement remains for trials of COVID-19 treatments (Food and Drug Administration 2020). Elsässer et al. (2014) reviewed scientific advice given by the European Medicines Agency (EMA) on adaptive design proposals and their outcomes (Collignon et al. 2018). The authors showed that if adaptive designs are properly implemented, then such innovative designs are well accepted by health authorities.

Strict control of the Type I error rate can be achieved by using prespecified adaptation rules such as early stopping boundaries in GS designs (Jennison and Turnbull 1999) or blinded sample size reassessment rules when estimating nuisance parameters at an interim analysis (Friede and Kieser 2006). However, prespecified designs do not permit unplanned modifications. To allow for further flexibility, fully adaptive designs have been suggested which allow adaptations not completely specified in the adaptation rule while still controlling the Type I error rate. Many designs of this type use p -value combination tests (Bauer and Köhne 1994) to allow for more flexible decision making, while still preserving the integrity of the design. Alternatively, designs can follow the conditional error principle (Proschan and Hunsberger 1995; Müller and Schäfer 2004) to control the Type I error rate. The conditional error principle enables one to introduce flexibility to any type of predefined design, even if no interim analysis was originally planned (Müller and Schäfer 2004); the key is that the significance level of the adapted trial

must not exceed the level of the initial design conditional on the data already observed. If data are combined across stages of the adaptive trial using the inverse normal function proposed by Lehmacher and Wassmer (1999) and Cui, Hung, and Wang (1999), and monitored against GS boundaries, the test decision should coincide with the study outcome had no unplanned adaptation been performed.

If more than one null hypothesis is to be tested, for example, in a multi-arm study, several extensions (Bauer and Kieser 1999; Hommel 2001) to the designs discussed above have been proposed which use adaptive designs within a closed testing framework (Marcus, Peritz, and Gabriel 1976). A closed testing procedure requires level α tests of each individual null hypothesis to be specified, as well as level α tests of all possible intersections of null hypotheses; if all tests involving a specific hypothesis are significant at level α , then that null hypothesis may be rejected globally at level α . Tests maintaining strong control of the FWER must (whether explicitly defined or not) be a closed testing procedure. For adaptive trials, for each null hypothesis, whether individual or an intersection, an adaptive test is prespecified.

The advantage of the fully adaptive design is that both internal and external information can be used to adapt the trial without fully specifying the adaptation rule. This is crucial in the pandemic situation as it means emerging data from parallel trials can be easily incorporated into the adaptation decision rule of an ongoing adaptive trial without compromising its integrity.

The methods discussed above focus on ensuring Type I error rate control in a frequentist setting. It is also possible to incorporate Bayesian methodology into the decision making process for adaptations without compromising error rate control, often using simulation methods to select critical values and other design parameters (Schmidli, Bretz, and Racine-Poon 2007; Brannath et al. 2009; Berry et al. 2011; Yu, Ramakrishnan, and Meinzer 2019; Stallard et al. 2020). Such an approach is illustrated by the examples in Sections 2.1 and 2.4.

3.3. GS Stopping Rules in Confirmatory Adaptive Design Trials

Incorporating a GS stopping rule into a confirmatory clinical trial allows for greater flexibility as compared to a standard trial design without any interim data looks. If an extraordinarily large benefit of the experimental treatment over control is observed at an interim analysis, it is desirable to have the option to stop the trial for efficacy with early rejection of the null hypothesis. If such an option is incorporated, a multiple testing problem emerges which needs to be handled by adjusting the local significance levels for the interim and final analysis to maintain the nominal Type I error rate. In GS trials with one or more interim looks, this can be done by using adjusted rejection regions tailored to prespecified interim analysis times (O'Brien and Fleming 1979; J. Whitehead 1997). In practice, many GS trials are run allowing flexibility with regards to the timing and number of interim analyses, using the alpha-spending function approach (Kim and Demets 1987) or flexible boundary methods (J. Whitehead 1997, 2011). The two examples described above in

Sections 2.2 and 2.3 that are based on frequentist analyses both use the alpha-spending function approach to ensure that Type I error rates are controlled. GS stopping rules for efficacy can also be incorporated in adaptive designs using a combination test approach (Bauer and Köhne 1994; Lehmacher and Wassmer 1999) or the conditional error principle (Proschan and Hunsberger 1995; Müller and Schäfer 2004).

Alternatively, one might also consider stopping a confirmatory trial early for futility if data are consistent with clinically irrelevant or harmful treatment effects. Allowing stopping for futility can be particularly useful in confirmatory trials of COVID-19 treatments which, due to high levels of unmet medical need, may follow accelerated development programs, so that limited phase II data are available before the confirmatory study begins (Food and Drug Administration 2020). Stopping for futility also means that resources can be quickly diverted to test more promising interventions. If recruitment into a COVID-19 trial is rapid, many patients will already have been enrolled by the time the futility interim analysis is performed, and the reductions in expected sample size possible by testing group sequentially will be smaller than those typically cited (Hampson and Jennison 2013). To mitigate this issue, accrual could be paused in the run-up to a futility interim analysis (Food and Drug Administration 2020) although in practice restarting recruitment is very challenging. A futility rule can either be binding in the sense that it needs to be strictly adhered to at the interim analysis, or nonbinding, meaning that the rule may also not be followed. While binding futility rules allow for a relaxation of the local significance levels, nonbinding rules can be regarded as more advantageous from a practical perspective due to their increased flexibility. However, the Type II error rate increases with the use of a nonbinding futility rule, requiring an increased sample size to achieve the target power. The COVID-19 trial examples given above include both trials with binding (Section 2.3) and nonbinding (Section 2.2) futility stopping rules.

3.4. Multi-Arm Multi-Stage (MAMS) and Platform Trials

MAMS designs allow several experimental treatments to be assessed in a single trial with treatments dropped from the trial as soon as interim analysis data suggest they are ineffective. The design requires specification of stopping boundaries prior to the trial starting: these usually consist of futility stopping boundaries and, less commonly, efficacy stopping boundaries. At an interim analysis, each experimental arm is compared against control using all patients with outcome information up to that point: experimental arms with test statistics below the futility boundary have further recruitment stopped and experimental arms with test statistics above the efficacy boundary can be concluded to be effective (and the respective null hypothesis rejected). Setting stopping boundaries is not trivial, but can be done using analytical formulas (Royston, Parmar, and Qian 2003; Magirr, Jaki, and Whitehead 2012) or simulation in the frequentist (Wason and Jaki 2012) or Bayesian (Yu, Ramakrishnan, and Meinzer 2019) settings. Software exists such as the R package MAMS (Jaki, Pallmann, and Magirr 2019) and the Stata module nstage (Bratton, Choodari-Oskooei, and Royston 2015).

The stopping and selection rules are preplanned, meaning that the trial must be conducted according to the prespecified rules. A frequentist MAMS design is being used in the COVID-19 RECOVERY-Respiratory Support trial described in Section 2.2. The design is efficient in that two experimental treatments are compared with a shared control group with the chance to stop as soon as sufficiently positive results have been observed or to drop an arm or stop the trial completely if results are not sufficiently promising.

The MAMS trial methodology can be generalized to a flexible confirmatory adaptive approach (Magirr, Stallard, and Jaki 2014) which is usually based on the combination testing principle (Bauer and Köhne 1994). This principle can be applied for testing multiple hypotheses by use of the closed testing principle (Bauer and Köhne 1994; Bretz et al. 2009; Bauer et al. 2016; Wassmer and Brannath 2016). An application is the multi-arm setting where, for example, treatment arms can be selected and sample sizes for the selected treatment arms can be recalculated based on the observed response (Bauer and Kieser 1999; Koenig et al. 2008; Magirr, Stallard, and Jaki 2014; Wassmer and Brannath 2016). For these designs, the adaptation rules, need not be prespecified. Consequently, this approach allows for testing in MAMS trials controlling the FWER in a strong sense. R packages are available for confirmatory adaptive designs in the multi-arm setting, such as *asd* (Friede, Stallard, and Parsons 2020) and *rpact* (rpact 2020), with MAMS (Jaki, Pallmann, and Magirr 2019) also having an extension for dealing with unexpected design modifications.

Some MAMS trials use adaptive randomization (AR) to guide allocation to better performing arms. In a multi-arm setting, if the shared control group is kept separate from the AR procedure, this maintains the power at a high level (Wason and Trippa 2014; Williamson and Villar 2020). Frequentist and Bayesian approaches to AR can be used, although in practice it is more common to use a Bayesian approach, as in the REMAP-CAP trial described in Section 2.4. AR has been criticized for being susceptible to temporal trends in the trial (Thall, Fox, and Wathen 2015; Proschan and Evans 2020), which may be possible in COVID-19 settings due to mutation of the virus. However, if the allocation to the shared control group is maintained, there is little impact on the statistical properties of the trial (Villar, Bowden, and Wason 2018).

Another type of MAMS design is the drop-the-losers design (Sampson and Sill 2005; Wason et al. 2017), which sets a fixed number of arms to progress at each interim analysis. An advantage of this is that the design has a fixed sample size (in contrast to a random sample size for the approach described above). The main disadvantage is that it may lead to dropping of promising arms if there are more than the design is permitted to take forward. Approaches that combine the drop-the-losers approach with prespecified stopping boundaries have been proposed (Friede and Stallard 2008) to allow early stopping for futility and efficacy.

Magirr, Jaki, and Whitehead (2012), Wason et al. (2017), and Robertson and Wason (2019) describe how FWER control can be achieved in a MAMS design. However, there has been debate about which type of multiplicity adjustment, for example, pairwise or family-wise, is needed for a multi-armed design (e.g., Wason, Stecher, and Mander 2014; Stallard et al. 2019; Collignon

et al. 2020). Generally, the decision will be determined by the relative importance of making a Type I error and a Type II error, as correcting for multiple testing reduces the power. In a platform setting, it is more difficult to formally control the FWER unless it is known in advance how many treatments will be tested. New methods for online control of error rates may provide a way of doing this if it is required (Robertson and Wason 2018).

Most methodology work in MAMS has focused on settings with a shared control group. In the case of no obvious control group, we would refer the reader to the Magaret et al. (2016) and the motivating example in Whitehead, Desai, and Jaki (2020).

When there is a shared control group, the relative allocation between control and experimental treatments can affect the power. A natural choice is to randomize patients equally between each available treatment. One can gain a small amount of power by increasing the allocation to the control group (Wassmer 2011; Wason et al. 2016), although in a MAMS study the optimal allocation depends on the likelihood of dropping arms only. Setting the allocation ratio of new arms in a platform study can also be optimized (Bennett and Mander 2020).

3.5. Estimation After an Adaptive Design Trial

Point estimates and confidence intervals (CIs) for the treatment effects assessed in a confirmatory clinical trial play a key role in deciding whether or not to recommend adoption of the new interventions in clinical practice. Therefore, the estimators used in confirmatory adaptive clinical trials need to have good accuracy. In some settings conventional estimation approaches may be used provided the bias can be shown to be small. In other settings naive point estimators that assume a two-arm single-stage design can be substantially biased for an adaptive trial because of the possibility of stopping early for futility or overwhelming effect (Whitehead 1986; Koopmeiners, Feng, and Pepe 2012) and/or selecting the most promising experimental treatments at an interim analysis (Bauer et al. 2010; Carreras and Brannath 2013; Kimani, Todd, and Stallard 2013). For two-stage trials, conditional on the trial continuing to stage 2, the selection rule and the treatments continuing to stage 2, several uniformly minimum variance conditional unbiased estimators (UMVCUEs) exist (Bowden and Glimm 2008; Kimani, Todd, and Stallard 2013; Robertson, Prevost, and Bowden 2016; Stallard and Kimani 2018). The UMVCUEs can have large mean squared errors (MSEs) and are computed if the trial continues to stage 2. Alternative shrinkage estimators that are based on empirical Bayes approaches tend to have smaller MSEs but they are not necessarily unbiased (Carreras and Brannath 2013; Brückner, Titman, and Jaki 2017). Estimators that also reduce bias when a trial stops in stage 1 exist (e.g., Stallard and Todd 2005; Carreras and Brannath 2013) as well as generalizations to multiple stages (Whitehead, Desai, and Jaki 2020).

Naive CIs can have incorrect coverage. When an adaptive design selects the most effective experimental treatment in stage 1, among the methods that adjust for adaptation (Sampson and Sill 2005; Stallard and Todd 2005; Wu, Wang, and Yang 2010; Magirr et al. 2013), the Sampson and Sill (2005) method performs best in terms of coverage while the approach of Wu,

Wang, and Yang (2010) is the most straightforward for computing CIs (Kimani, Todd, and Stallard 2014). Methods for computing CIs with other selection rules are available (Neal et al. 2011; Magirr et al. 2013). One of these methods allows continuing with multiple experimental treatments in stage 2 but sometimes the computed CIs are non-informative (Magirr et al. 2013).

Most of the methods above assume normally distributed outcomes with known variance. Extensions to the case of unknown variance (Robertson and Glimm 2019) exist and, using asymptotic approximations, they can be extended to other outcomes. Point estimation with time-to-event outcomes has been considered separately as these outcomes pose specific challenges (Brückner, Titman, and Jaki 2017). One of the UMVCUEs is a generalization of the others enabling MAMS trials (Stallard and Kimani 2018).

4. Specific Statistical Issues in Adaptive Designs in COVID-19 Intervention Trials

4.1. Characterizing the Dose–Response Relationship

When studying a compound in a completely new disease population, such as COVID-19, it is essential to identify a therapeutic window, that is, a range of doses which simultaneously control the risk of adverse events while having a clinically meaningful efficacy effect. Importantly, the same requirement applies to repurposed medicines as the dose selected for the original indication is not guaranteed to be optimal for the COVID-19 population. Characterization of the dose–response relationship, however, poses several statistical challenges.

Both the toxicity and efficacy signals should be attributable to the compound rather than to the baseline toxicity and efficacy rates. Therefore, the dose–response relationship should be established accounting for outcomes on control, by quantifying differences between the active doses and control (SoC or placebo) for both toxicity and efficacy endpoints. Furthermore, to gain efficiency by using data across different doses of the compound, a model-based approach that stipulates a parametric relationship for the underlying differences should be used. This, in turn, poses a challenge in terms of how to select the dose–response model. To overcome this, an MCP-Mod approach which considers several candidate models (Bretz, Pinheiro, and Branson 2005; Pinheiro et al. 2014), or a model-averaging approach can be used. To establish the dose-response relationship, randomization of patients between doses is needed to allow for efficient dose-response estimation. Consequently, the safety of the dosing range must be established before the dose-response relationship can be reliably characterized. A two-stage procedure adjusting allocation probabilities to each dose could also be used (Mielke and Dragalin 2017).

To ensure the most efficient use of the data, the most informative yet clinically meaningful endpoint should be used. While it is common to dichotomize continuous or time-to-recovery endpoints to simplify the analysis, it is inevitable that this will result in the loss of some statistical information; specifically, this strategy has been found to result in noticeable losses when characterizing the dose–response relationship (Mozgunov, Jaki, and Paoletti 2020). While dose–response models for contin-

uous and binary endpoints are more common, there are also parametric and semiparametric models available for time-to-event data with implementations available in standard statistical software.

In the AGILE trial described in Section 2.1 the risk of toxicity (yes/no) is modeled and excess risk of toxicity of the investigated treatment above control is used for dose-selection.

4.2. Choice of Primary Endpoint

Choosing the primary endpoint is often a challenge in clinical trials, but a novel disease with lots of uncertainty makes this choice even more challenging. Whitehead and Horby (2017) approached this challenge by recognizing that outcomes such as recovered versus not recovered are almost always of interest. In recent work discussing different choices of endpoints for COVID-19, Dodd et al. (2020) also state that mortality is a desirable endpoint for COVID-19 that might be difficult to use in practice due to the large number of patients required. Other endpoints are therefore often used, but as knowledge about the disease increases, alternative endpoints might emerge that are deemed to be more suitable. In such a situation it is natural to want to adapt the primary endpoint in the light of this new information. Such an approach is possible within an adaptive design (Kieser, Bauer, and Lehmacher 1999), but challenges may arise in convincing end-users of the trial results that such a change has not been made with ill intent. A recent example of a trial that did change the primary endpoint is the ACTT trial (Beigel et al. 2020) that initially used the WHO ordinal scale at a fixed time point but later changed, on the basis of information external to the trial, to time-to-recovery.

4.3. Stopping a Trial Early Due to Recruitment Difficulties

The incidence of COVID-19 will vary over the course of the pandemic. A trial starting recruitment during the “post-peak” phase of the pandemic, as disease activity levels trail off, may struggle to reach the target information level before it ends. If a GS test is stopped early due to feasibility issues, steps can be taken to draw valid inferences about the target treatment effect. As mentioned in Section 3.3, many GS trials are monitored according to the alpha-spending function approach (Kim and Demets 1987), and accommodating early termination in this framework is relatively straightforward. At the final analysis, we choose the critical value of the test so as to “spend” all the remaining Type I error rate. Choosing the final critical value in this way preserves the Type I error rate of the trial, although there will be a drop in power relative to the originally planned design. CIs and p -values based on the stagewise ordering of the test’s outcome space can be calculated conditioning on the observed sequence of information levels using the methods described in Section 3.5. A similar idea can be used to reach a final hypothesis decision in a GS trial which does not follow an alpha-spending design: one finds the critical value for the final analysis so that the trial has Type I error rate α given the observed information levels and the test boundaries used at previous interim analyses. One additional caveat that applies to all GS trials stopped early for lack of feasibility: while the

decision of when to truncate the GS test can (and should) be informed by trends in recruitment rates, the timing of the final analysis should not be influenced by previous treatment effect estimates (Jennison and Turnbull 1999). If, for example, we bring forward the final analysis because we were close to crossing a boundary at the last interim look, test statistics will not follow the anticipated canonical joint distribution used to calculate stopping boundaries (Jennison and Turnbull 1997), leading to deviations from the nominal Type I error rate.

If the pandemic ends earlier in some geographic regions than others, the sponsor of a global study can compensate by switching to recruit more patients at sites where the disease is still active. While this strategy would preserve the Type II error rate of the trial, one would still see perturbations in the power of region-specific subgroup analyses. In addition, the impact of this shift in recruitment on the interpretation of the trial should be carefully considered, particularly if the target estimand (ICH 2020) remains as the effect of treatment in a global patient population (which would still be pertinent if all regions remained at risk of subsequent waves of infection). To align with this target estimand, the analyst could perform a random effects meta-analysis of region-specific treatment effect estimates, where the mean of the random-effects distribution would be taken as the overall treatment effect of interest.

4.4. Heterogeneity of the Patient Population

The prognosis of COVID-19 is highly variable, ranging from nonsymptomatic courses of the disease in the majority of patients, to severe clinical courses requiring hospitalization, intensive care and leading to death in many cases. A number of prognostic factors have been identified, including age, sex, chronic lung disease, diabetes, hypertension, and other cardiovascular comorbidities (Grasselli et al. 2020; Zhou et al. 2020). In addition, the time span and course of the disease up to the time of study inclusion is prognostic for the outcome. For the design of clinical trials this implies that the inclusion criteria have a strong impact on the distribution of the primary endpoint and the corresponding sample size planning. Furthermore, in the analysis of the clinical trial, the precision of treatment effect estimates can be substantially improved by adjusting for important covariates in the analysis. Also, stratified randomization according to the risk factors can improve the efficiency of the design.

An important question is to what extent the prognostic factors are also predictive for the efficacy of the treatment. For example, for anti-viral drugs it is usually assumed that efficacy is largest when treatment starts early after infection. However, comorbidities and demographic variables can also play a role in the mechanism of action or affect the safety profile. If the treatment effect varies across subgroups that are defined by baseline characteristics, treatment effect estimates are only meaningful for a well-defined patient population. In addition, subgroup analyses will be critical to better understand the treatment effect. Challenges are the typically small sample sizes in subgroups and the multiplicity problem connected with multiple analyses. Adaptive enrichment designs can be a useful tool in this context (see, e.g., Friede, Stallard, and Parsons 2020

for a recent overview). Based on interim results, they allow one to increase the sample size in specific subgroups or to restrict randomization to these populations, while controlling frequentist error rates (Ondra et al. 2016, 2019; Chiu et al. 2018). However, pre-specification of subgroups can be difficult, especially for quantitative predictive factors (e.g., time from the onset of the disease), where thresholds need to be prespecified to define the corresponding sub-populations. If there is not enough prior information to select a single threshold, subgroup analyses with multiple thresholds can be implemented if an appropriate adjustment for multiple testing is foreseen (Graf et al. 2020).

4.5. Dynamically Changing Standard of Care

SoC is likely to evolve quickly in the early stages of a pandemic, particularly of a previously unknown disease such as COVID-19, as clinicians' understanding of the disease and experience treating it increases rapidly. Of course, many trials in an acute pandemic setting recruit quickly and measure short-term endpoints; see, for example, the primary endpoints of the RECOVERY and REMAP-CAP trials described in Section 2. This means that SoC would be expected to remain relatively stable over the course of an individual study. However, more substantial changes may be seen over the course of a trial measuring the long-term outcomes of patients with less severe disease at baseline. Alternatively, we may see changes in SoC across the course of several successive studies which we intend to combine via meta-analysis.

For trials of add-on treatments (comparing SoC + novel therapy vs. SoC alone), block-randomization or pre-specification of the allocation probabilities in a parallel group study will be sufficient to ensure control of the Type I error rate at the nominal level. If changes in SoC result in a time-varying treatment effect, the study investigators should consider what the target estimand is, that is, an aggregate treatment effect averaging across the course of the study, or the effect of treatment compared with the most recent version of SoC, and whether the target effect can be estimated.

Changes in SoC can present a particular challenge in a platform trial adding in new experimental treatments over time. In this setting, there is debate about whether one should use non-concurrent control patients in the evaluation of experimental treatments. If a new treatment is added in part-way through the trial, concurrent control patients are those who are recruited from that time; nonconcurrent controls would be those patients recruited prior to the addition. If it is likely that there is strong potential for a temporal trend in the data, as may be the case in COVID-19 trials, it is likely best to use only concurrent controls as the primary analysis.

4.6. Leveraging Information Across Trials Using Meta-Analytic Approaches

Comparison-specific treatment effects may accumulate over the course of a pandemic as randomized controlled trials are replicated by different research groups, and as platform trials with common treatment arms report pairwise comparisons. While there may be substantial heterogeneity among the control

groups of successive trials in a pandemic due to the evolving SoC (Dodd et al. 2019), two-group comparative effects may remain more stable. Combining evidence across studies in a meta-analysis will increase the statistical information available for estimating a target treatment effect. If we prospectively intend to update the meta-analysis as each new study is published, the cumulative results should be monitored using a GS test to avoid over-interpretation of random highs and lows in treatment effect estimates, and inflation of the Type I and Type II error rates. Different approaches have been proposed for monitoring a cumulative meta-analysis (Simmonds et al. 2017). Sequential meta-analysis approaches use GS boundaries to monitor standardized Z statistics obtained from random-effects meta-analyses. For example, boundaries could be from a triangular GS test (A. Whitehead 1997) or a restricted procedure (Higgins, Whitehead, and Simmonds 2011), or an alpha-spending test spending Type I error as a function of the estimated proportion of the target information level accumulated thus far. The DerSimonian and Laird (1986) estimate of the between-study heterogeneity may be unstable in the early stages of a cumulative meta-analysis when few studies are available, so that estimated information levels may decrease between successive interim analyses (Higgins, Whitehead, and Simmonds 2011). With (very) few studies also Bayesian approaches with weakly informative priors for the between-study heterogeneity were found to be more robust than standard methods (Friede et al. 2017). It may be possible to reduce between-study heterogeneity and increase borrowing across studies by leveraging data on baseline covariates, either in a meta-regression approach if only aggregate data are available (Neuenschwander and Schmidli 2020), or in models for patient responses if individual patient data are available.

5. Practical Issues in Conducting Adaptive Designs

As highlighted in Section 1, time is essential when conducting clinical trials during a disease outbreak. When conducting an adaptive trial, one particular aspect, namely decision making about undertaking any of the (preplanned) adaptations, needs to be particularly considered in addition. This time-sensitivity means that the decision processes (which often involve review of unblinded data) are streamlined and decision making is clear before undertaking any interim data looks to avoid potential delays in implementing decisions. The DMC of the trial will often have a particularly crucial role to play in these decision processes.

5.1. Data Monitoring Committees for an Adaptive Trial

A DMC, also known as data safety monitoring board (DSMB) or data safety monitoring committee (DSMC), is “a committee created by a sponsor to provide independent review of accumulating safety and efficacy data” (Herson 2009). Therefore, the DMC is “responsible for [...] protecting trial integrity and patient safety” (Herson 2009). Although all trials require safety monitoring, this does not mean that a DMC needs to be established for all trials. However, safety monitoring by a DMC is recommended “for any controlled trial of any size that will compare rates of mortality or major morbidity” or more

specifically for a “population at elevated risk of death or other serious outcomes, even when the study objective addresses a lesser endpoint” (Food and Drug Administration 2006), which is likely to be the case for any COVID-19 trial. Compared to other studies, COVID-19 trials might be shorter with faster recruitment (given the pandemic situation) and shorter length of follow-up (e.g., restricted to length of stay in hospital). For safety monitoring by a DMC this, in turn, requires shorter review intervals, which might lead to weekly safety reviews in some instances. This has obvious consequences for logistics in terms of data management, statistical analyses, as well as the required availability and commitment of the DMC members. Given these demands it might be practical to set up committees to monitor not only an individual trial but to oversee several trials on a programme level. This would not be uncommon for trials investigating the same compound, but in COVID-19 might be applied also to settings where different compounds are being investigated by the same sponsor. This would also enable the DMC to consider data emerging from other ongoing trials in their decision making, allowing them to integrate rapidly increasing and changing information. This type of oversight committee tends to be larger than the standard DMC which typically comprises two physicians and one statistician, since a wider range of expertise might be required and several regions might need to be represented. In the case of adaptive designs or complex platform trials, the DMC might be charged with extra responsibilities beyond safety monitoring. For example, when performing a formal interim analysis, the DMC might make recommendations regarding treatment or subgroup selection, as well as futility stopping or sample size re-estimation. Since COVID-19 trials often use quickly observed endpoints, such adaptations might still be meaningful, although the trials might be quite short overall given rapid recruitment.

5.2. Regulatory Considerations

Basic regulatory requirements on adaptive designs to be used in confirmatory study are laid down by the Food and Drug Administration and EMA (European Medicines Agency 2007; Food and Drug Administration 2019). In general, regulatory concerns relate to the validity and robustness of the conclusions to be made with respect to a statement on the drug’s efficacy and a thorough benefit risk assessment. In statistical terms, proper Type I error control and unbiased effect estimates (or at least minimization of bias) are of major concern for the primary efficacy endpoint(s), whereas the potential lack of mature data in trials with interim decision relate to the totality of the data package. In addition to proper statistical methods to control Type I error and bias in estimation, issues with the integrity of the trial due to change in trial conduct following an interim evaluation (and leading to inconsistency between stages) play a major role in the assessment of the trial’s validity.

Referring to the specific pandemic setting, the request for rapid decisions, investigation of multiple treatments, and potentially changing controls could complicate the multiplicity issues but the assessment of the trial’s validity remains an overriding regulatory principle. In case the high medical need for effective drugs in a pandemic situation leads to a different appraisal of

Type I error rate control to make new treatments available as fast as possible. It remains paramount to understand the trial's operating characteristics to ensure a fully informed regulatory decision making.

In the specific setting of the COVID-19 pandemic with a number of highly relevant risk factors, high heterogeneity with respect to the severity of the disease and the onset of treatment, a good characterization of the population that would profit from treatment is crucial. Obviously, an extensive search for the optimal subgroup and subsequent confirmation may not be compatible with the need for rapid development. Regulatory principles of subgroup selection are given in the EMA guideline on the investigation of subgroups in confirmatory clinical trials (European Medicines Agency 2019), to be applied in case a positive effect has been established in the overall population but plausible heterogeneity with respect to efficacy and severe side effects asks for a refinement of the population.

6. Conclusions and Discussion

The examples provided in Section 2 give an indication of the range of adaptive design approaches available, and show how these methods have been embraced to advance research even within the challenging setting of a pandemic. In all cases, the objectives of the investigators were to choose an approach that would yield scientifically rigorous results in as rapid a timeframe as possible. The adaptive design framework provides sufficient flexibility to enable trials of COVID-19 patients to be embedded in ongoing trials, while MAMS and platform trial approaches enable rapid evaluation of a number of potential treatments, perhaps repurposed from other disease areas. Bayesian or frequentist methods are available, and in each case can, provide Type I error rate control, at least approximately, required in a confirmatory clinical trial. It is also important that trials are adequately powered, as control of the Type II error rate is particularly relevant in the COVID-19 setting in which there are no existing treatment options.

Although there can be no doubt that clinical research in the setting of a pandemic raises a number of unique challenges, the need for the evaluation of potential treatments in scientifically rigorous randomized controlled clinical trials is as great as in any other area of clinical research.

Whilst it is always desirable to conduct clinical trials of new medical interventions in a timely fashion, faced with a pandemic for which no vaccines or treatments are available, this need is particularly acute, both because of the urgent need for treatments and because of the limited time-window in which large-scale research may be possible in a particular location, or even globally.

Although trials that are unable to fully enroll in time to inform a current outbreak may provide crucial data for future outbreaks or may be analyzed in the context of a meta-analysis, the need for clinical trials that can yield a conclusion, either positive or negative, as quickly as possible, makes adaptive design methods particularly attractive, especially if they are flexible enough to allow modifications to a trial in light of emerging clinical knowledge and results from other ongoing or completed clinical trials. It should be possible to suspend recruitment to

a trial during times when the disease is under control and to resume if it re-emerges (Dean et al. 2020). The use of adaptive platform designs, in particular for clinical trials in COVID-19, has been proposed by a number of authors (Baucher and Fontanarosa 2020) and, as illustrated by the examples given in Section 2, a range of sequential and adaptive design methods are being employed in ongoing clinical trials. Although not a focus of this article, adaptive design methods also provide valuable tools to allow for modification of trials in other diseases necessitated by the COVID-19 pandemic (Kunz et al. 2020).

To yield timely results, it is clear that treatment trials should start as soon as possible after an epidemic has begun. The time limiting factor should be identification of candidate treatments, and not the development of a statistical design. In the current COVID-19 pandemic, although adaptive designs were used for some trials, as indicated by the examples above, others which could have benefited from such an approach will have been planned with a fixed sample design because of the pressure to start trials rapidly. Hence, research on trial designs that are most suitable for the purposes of a pandemic should become a research topic that goes beyond the actual case to be better prepared for future pandemics.

A suitable design, or small suite of designs, should be available on the shelf ready for use at any time. Debates about trial design, including whether "best available care" can vary amongst centers and whether Bayesian or frequentist approaches are appropriate should be settled between pandemics, and not reopened when they strike. It might be supposed that such a generic approach is impossible and that the design cannot be created without details of the natural history of the disease or the specification of candidate treatments. However, in almost all cases the primary patient response will be status, a certain number of D days after randomization, classified as dead/severely ill/moderately ill/recovered. The magnitude of differences between treatments can be conveniently assessed in terms of odds-ratios, and an odds-ratio of 2 is almost always clinically important. A potential generic design has been suggested by Whitehead and Horby (2017) that might be a starting point in the search for a suitable vehicle. Perhaps the COVID-19 experience will motivate the clinical trials community to devise and agree upon such a generic design once the current pandemic subsides. It should accommodate the evaluation of multiple treatments compared with best available care, with treatments entering the study as they are devised and are approved for experimental use, and leaving as they are demonstrated to be inferior to others. The design should be adaptive and capable of being used across many sites, preferably internationally.

Even if suitable trial designs are available, the planning and conduct of clinical trials remains difficult. This is true for any trial being conducted during a pandemic, when healthcare resources are overstretched and the possibility of collecting data from face-to-face assessments may be limited (Anker et al. 2020). However, in trials focusing on the pandemic itself, there are additional challenges in that the number of patients likely to be affected, as well as the duration of the pandemic, is unknown and may vary considerably across the globe. In addition, pressures on healthcare systems mean that real-time reporting of trial data is often difficult, reducing the information

available at planned interim analyses. Given the importance of solid evidence on the effectiveness of treatments, however, we believe that every effort should be made to ensure that logistical challenges in data collection and transfer do not undermine otherwise high quality trials. For rigorous and quick progress in evaluating treatments, there is a need for scientific collaboration, and the necessary political cooperation to enable and facilitate this. A multi-center trial can recruit patients as quickly as possible by recruiting internationally, or from different countries at different times as the pandemic progresses. In a setting where scientific knowledge is rapidly changing and research studies are being set up quickly, data sharing among global scientific communities is key in allowing meta-analysis methods, such as those described in Section 4.5, to be used. Consequently, results from a number of studies can be combined to provide robust evidence of treatment effectiveness, thus preventing wastage of resources arising from the conduct of multiple similar trials, each of which are individually too small to yield definitive results.

Acknowledgments

We are grateful to an associate editor and three anonymous referees for their helpful comments that have improved the article.

Disclosure Statement

Martin Posch reports grants from Novartis Pharma AG, grants from Medicconomics GmbH, grants from Merck KGaA, grants from Abbott Laboratories, grants from Almirall, outside the submitted work. He is also a member of the IMI Project Consortium EU-PEARL.

Funding

This article was initiated by the Adaptive Designs and Multiple Testing Procedures Joint Working Group of the Austro-Swiss (ROeS) and the German (IBS-DR) Regions of the International Biometric Society, chaired by Rene Schmidt and Lisa Hampson, and with support from Werner Brannath and Andreas Faldum. J. Wason received funding from UK Medical Research Council (MC_UU_00002/6 and MR/N028171/1). T. Jaki received funding from UK Medical Research Council (MC_UU_00002/14). This report is independent research arising in part from Prof Jaki's Senior Research Fellowship (NIHR-SRF-2015-08-001) supported by the National Institute for Health Research. 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 and Social Care (DHCS), or the Federal Institute for Drugs and Medical Devices (BfArM). Martin Posch and Franz Koenig (Medical University of Vienna) are members of the EU Patient-centric clinical trial platform (EU-PEARL). EU-PEARL has received funding from the Innovative Medicines Initiative 2 Joint Undertaking under grant agreement no. 853966. This Joint Undertaking receives support from the European Union's Horizon 2020 Research and Innovation Programme and EFPIA.

References

Anker, S. D., Butler, J., Khan, M. S., Abraham, W. T., Bauersachs, J., Bocchi, E., Bozkurt, B., Braunwald, E., Chopra, V. K., Cleland, J. G., Ezekowitz, J., Filippatos, G., Friede, T., Hernandez, A. F., Lam, C. S. P., Lindenfeld, J., McMurray, J. J. V., Mehra, M., Metra, M., Packer, M., Pieske, B., Pocock, S. J., Ponikowski, P., Rosano, G. M. C., Teerlink, J. R., Tsutsui, H., Van Veldhuisen, D. J., Verma, S., Voors, A. A., Wittes, J., Zannad, F., Zhang, J., Sefderovic, P., and Coats, A. J. S. (2020), "Conducting Clinical Trials in Heart Failure During (and After) the COVID-19 Pandemic: An Expert

Consensus Position Paper From the Heart Failure Association (HFA) of the European Society of Cardiology (ESC)," *European Heart Journal*, 41, 2109–2117. [493]

Baucher, H., and Fontanarosa, P. B. (2020), "Randomized Clinical Trials and COVID-19: Managing Expectations," *Journal of the American Medical Association*, DOI: 10.1001/jama.2020.8115. [483,493]

Bauer, P., Bretz, F., Dragalin, V., König, F., and Wassmer, G. (2016), "Twenty-Five Years of Confirmatory Adaptive Designs: Opportunities and Pitfalls," *Statistics in Medicine*, 35, 325–347. [487,489]

Bauer, P., and Kieser, M. (1999), "Combining Different Phases in the Development of Medical Treatments Within a Single Trial," *Statistics in Medicine*, 18, 1833–1848. [488,489]

Bauer, P., Koenig, F., Brannath, W., and Posch, M. (2010), "Selection and Bias—Two Hostile Brothers," *Statistics in Medicine*, 29, 1–13. [489]

Bauer, P., and Köhne, K. (1994), "Evaluation of Experiments With Adaptive Interim Analyses," *Biometrics*, 50, 1029–1041, Corrigenda (1996) 52, 380. [487,488,489]

Beigel, J. H., Tomashek, K. M., Dodd, L. E., Mehta, A. K., Zingman, B. S., Kalil, A. C., Hohmann, E., Chu, H. Y., Luetkemeyer, A., Kline, S., de Castilla, D. L., Finberg, R. W., Dierberg, K., Tapson, V., Hsieh, L., Patterson, T. F., Paredes, R., Sweeney, D. A., Short, W. R., Touloumi, G., Lye, D. C., Ohmagari, N., Oh, M., Ruiz-Palacios, G. M., Benfield, T., Fätkenheuer, G., Kortepeter, M. G., Atmar, R. L., Creech, C. B., Lundgren, J., Babiker, A. G., Pett, S., Neaton, J. D., Burgess, T. H., Bonnett, T., Green, M., Makowski, M., Osinusi, A., Nayak, S., and Lane, H. C. (2020), "Remdesivir for the Treatment of Covid-19—Preliminary Report," *New England Journal of Medicine*, DOI: 10.1056/NEJMoa2007764. [490]

Bennett, M., and Mander, A. P. (2020), "Designs for Adding a Treatment Arm to an Ongoing Clinical Trial," *Trials*, 21, 251, DOI: 10.1186/s13063-020-4073-1. [489]

Berry, S. M., Carlin, B. P., Lee, J. J., and Müller, P. M. (2010), *Bayesian Adaptive Methods for Clinical Trials*, Boca Raton, FL: CRC Press. [487]

——— (2011), *Bayesian Adaptive Methods for Clinical Trials*, Boca Raton, FL: CRC Press. [488]

Bowden, J., and Glimm, E. (2008), "Unbiased Estimation of Selected Treatment Means in Two-Stage Trials," *Biometrical Journal*, 50, 515–527. [489]

Brannath, W., Zuber, E., Branson, M., Bretz, F., Gallo, P., Posch, M., and Racine-Poon, A. (2009), "Confirmatory Adaptive Designs With Bayesian Decision Tools for a Targeted Therapy in Oncology," *Statistics in Medicine*, 28, 1445–1463. [488]

Bratton, D. J., Choodari-Oskooei, B., and Royston, P. (2015), "A Menu-Driven Facility for Sample-Size Calculation in Multiarm, Multistage Randomized Controlled Trials With Time-to-Event Outcomes: Update," *Stata Journal*, 15, 350–368, DOI: 10.1177/1536867x1501500202. [488]

Bretz, F., Koenig, F., Brannath, W., Glimm, E., and Posch, M. (2009), "Adaptive Designs for Confirmatory Clinical Trials," *Statistics in Medicine*, 28, 1181–1217, DOI: 10.1002/sim.3538. [489]

Bretz, F., Pinheiro, J. C., and Branson, M. (2005), "Combining Multiple Comparisons and Modeling Techniques in Dose-Response Studies," *Biometrics*, 61, 738–748. [490]

Brückner, M., Titman, A., and Jaki, T. (2017), "Estimation in Multi-Arm Two-Stage Trials With Treatment Selection and Time-to-Event Endpoint," *Statistics in Medicine*, 36, 3137–3153. [489,490]

Carreras, M., and Brannath, W. (2013), "Shrinkage Estimation in Two-Stage Adaptive Designs With Mid-Trial Treatment Selection," *Statistics in Medicine*, 32, 1677–1690. [489]

Chiu, Y. D., Koenig, F., Posch, M., and Jaki, T. (2018), "Design and Estimation in Clinical Trials With Subpopulation Selection," *Statistics in Medicine*, 37, 4335–4352. [491]

Collignon, O., Gartner, C., Haidich, A. B., Hemmings, R. J., Hofner, B., Pétavy, F., Posch, M., Rantell, K., Roes, K., and Schiel, A. (2020), "Current Statistical Considerations and Regulatory Perspectives on the Planning of Confirmatory Basket, Umbrella and Platform Trials," *Clinical Pharmacology & Therapeutics*. [489]

Collignon, O., Koenig, F., Koch, A., Hemmings, R. J., Pétavy, F., Saint-Raymond, A., Papaluca-Amati, M., and Posch, M. (2018), "Adaptive Designs in Clinical Trials: From Scientific Advice to Marketing Authorisation to the European Medicine Agency," *Trials*, 19, 642. [487]

- Cooper, B. S., Boni, M. F., Pan-ngum, W., Day, N. P., Horby, P. W., Olliaro, P., Lang, T., White, N. J., White, L. J., and Whitehead, J. (2015), "Evaluating Clinical Trial Designs for Investigational Treatments of Ebola Virus Disease," *PLoS Medicine*, 12. [484]
- Cui, L., Hung, H. M. J., and Wang, S.-J. (1999), "Modification of Sample Size in Group Sequential Clinical Trials," *Biometrics*, 55, 853–857. [488]
- Dean, N. E., Gsell, P. S., Brookmeyer, R., Crawford, F. W., Donnelly, C. A., Ellenberg, S. S., Fleming, T. R., Halloran, M. E., Horby, P., Jaki, T., Krause, P. R., Longini, I. M., Mulangu, S., Muyembe-Tamfum, J.-J., Nason, M. C., Smith, P. G., Wang, R., Henao-Restrepo, A. M., and De Gruttola, V. (2020), "Creating a Framework for Conducting Randomized Clinical Trials During Disease Outbreaks," *New England Journal of Medicine*, DOI: 10.1056/NEJMs1905390. [493]
- DerSimonian, R., and Laird, N. (1986), "Meta-Analysis in Clinical Trials," *Controlled Clinical Trials*, 7, 177–188. [492]
- Dimairo, M., Pallmann, P., Wason, J., Todd, S., Jaki, T., Julious, S. A., Mander, A. P., Weir, C. J., Koenig, F., Walton, M. K., Nicholl, J. P., Coates, E., Biggs, K., Hamasaki, T., Proschan, M. A., Scott, J. A., Ando, Y., Hind, D., and Altman, D. G. (2020), "The Adaptive Designs CONSORT Extension (ACE) Statement: A Checklist With Explanation and Elaboration Guideline for Reporting Randomized Trials That Use an Adaptive Design," *BMJ*, 369, m115. [484]
- Dodd, L. E., Follmann, D., Proschan, M., Wang, J., Malvy, D., van Griensven, J., Ciglenecki, I., Horby, P. W., Ansumana, R., Jiang, J. F., and Davey, R. T. (2019), "A Meta-Analysis of Clinical Studies Conducted During the West Africa Ebola Virus Disease Outbreak Confirms the Need for Randomized Control Groups," *Science Translational Medicine*, 11, eaaw1049. [483,492]
- Dodd, L. E., Follmann, D., Wang, J., Koenig, F., Korn, L. L., Schoergenhofer, C., Proschan, M., Hunsberger, S., Bonnett, T., Makowski, M., Belhadi, D., Wang, Y., Cao, B., Mentré, F., and Jaki, T. (2020), "Endpoints for Randomized Controlled Clinical Trials for COVID-19 Treatments," *Clinical Trials*, DOI: 10.1177/1740774520939938. [483,487,490]
- Dodd, L. E., Proschan, M. A., Neuhaus, J., Koopmeiners, J. S., Neaton, J., Beigel, J. D., Barrett, K., Lane, H. C., and Davey, R. T., Jr. (2016), "Design of a Randomized Controlled Trial for Ebola Virus Disease Medical Countermeasures: PREVAIL II, the Ebola MCM Study," *The Journal of Infectious Diseases*, 213, 1906–1913. [483,484]
- Elsässer, A., Regnstrom, J., Vetter, T., Koenig, F., Hemmings, R. J., Greco, M., Papaluca-Amati, M., and Posch, M. (2014), "Adaptive Clinical Trial Designs for European Marketing Authorization: A Survey of Scientific Advice Letters From the European Medicines Agency," *Trials*, 15, 383. [487]
- European Medicines Agency (1998), "ICH E9: Statistical Principles for Clinical Trials." [487]
- (2007), "Reflection Paper on Methodological Issues in Confirmatory Clinical Trials Planned With an Adaptive Design." [487,492]
- (2019), "Guideline on the Investigation of Subgroups in Confirmatory Clinical Trials." [493]
- Food and Drug Administration (2006), "Guidance for Clinical Trial Sponsors: Establishment and Operation of Clinical Trial Data Monitoring." [492]
- (2019), "Guidance for industry: Adaptive Design Clinical Trials for Drugs and Biologics." [487,492]
- (2020), "COVID-19: Developing Drugs and Biologic Products for Treatment or Prevention. Guidance for Industry." [487,488]
- Friede, T., and Kieser, M. (2006), "Sample Size Recalculation in Internal Pilot Study Designs: A Review," *Biometrical Journal*, 48, 537–555. [487]
- Friede, T., Röver, C., Wandel, S., and Neuenschwander, B. (2017), "Meta-analysis of Few Small Studies in Orphan Diseases," *Research Synthesis Methods*, 8, 79–91. [492]
- Friede, T., and Stallard, N. (2008), "A Comparison of Methods for Adaptive Treatment Selection," *Biometrical Journal*, 50, 767–781, DOI: 10.1002/bimj.200710453. [489]
- Friede, T., Stallard, N., and Parsons, N. (2020), "Adaptive Seamless Clinical Trials Using Early Outcomes for Treatment or Subgroup Selection: Methods, Simulation Model and Their Implementation in R," *Biometrical Journal*, DOI: 10.1002/bimj.201900020. [489,491]
- Graf, A. C., Magirr, D., Dmitrienko, A., and Posch, M. (2020), "Optimized Multiple Testing Procedures for Nested Sub-Populations Based on a Continuous Biomarker," *Statistical Methods in Medical Research*, DOI: 10.1177/0962280220913071. [491]
- Grasselli, G., Zangrillo, A., Zanella, A., Antonelli, M., Cabrini, L., Castelli, A., Cereda, D., Coluccello, A., Foti, G., Fumagalli, R., Iotti, G., Latronico, N., Lorini, L., Merler, S., Natalini, G., Piatti, A., Ranieri, M. V., Scandroglio, A. M., Storti, E., Cecconi, M., and Pesenti, A. (2020), "Baseline Characteristics and Outcomes of 1591 Patients Infected With SARS-CoV-2 Admitted to ICUs of the Lombardy Region, Italy," *Journal of the American Medical Association*, 323, 1574–1581. [486,491]
- Griffiths, G., Fitzgerald, R., Jaki, T., Corkhill, A., Marwood, E., Reynolds, H., Stanton, L., Ewings, S., Condie, S., Wrixon, E., Norton, A., Radford, M., Yeats, S., Robertson, J., Darby-Dowman, R., Walker, L., and Khoo, S. (2020), "AGILE-ACCORD: A Randomized, Multicentre, Seamless, Adaptive Phase I/II Platform Study to Determine the Optimal Dose, Safety and Efficacy of Multiple Candidate Agents for the Treatment of COVID-19: A Structured Summary of a Study Protocol for a Randomised Platform Trial," *Trials*, 21, 1–3. [484]
- Gsponer, T., Gerber, F., Bornkamp, B., Ohlssen, D., Vandemeulebroecke, M., and Schmidli, H. (2014), "A Practical Guide to Bayesian Group Sequential Designs," *Pharmaceutical Statistics*, 13, 71–80. [484]
- Hampson, L. V., and Jennison, C. (2013), "Group Sequential Tests for Delayed Responses" (with discussion), *Journal of the Royal Statistical Society, Series B*, 75, 3–54. [488]
- Henao-Restrepo, A. M., Camacho, A., Longini, I. M., Watson, C. H., Edmunds, W. J., Egger, M., Carroll, M. W., Dean, N. E., Diatta, I., Doumbia, M., and Draguez, B. (2017), "Efficacy and Effectiveness of an rVSV-Vectored Vaccine in Preventing Ebola Virus Disease: Final Results From the Guinea Ring Vaccination, Open-Label, Cluster-Randomized Trial (Ebola Ça Suffit!)," *The Lancet*, 389, 505–518. [483]
- Herson, J. (2009), *Data and Safety Monitoring Committees in Clinical Trials*, Boca Raton, FL: Chapman & Hall/CRC. [492]
- Higgins, J. P., Whitehead, A., and Simmonds, M. (2011), "Sequential Methods for Random-Effects Meta-Analysis," *Statistics in Medicine*, 30, 903–921. [492]
- Hommel, G. (2001), "Adaptive Modifications of Hypotheses After an Interim Analysis," *Biometrical Journal*, 43, 581–589. [488]
- ICH (2020), "Addendum on Estimands and Sensitivity Analyses in Clinical Trials to the Guideline on Statistical Principles for Clinical Trials: E9(R1)." [491]
- Jaki, T., Pallmann, P., and Magirr, D. (2019), "The R Package MAMS for Designing Multi-Arm Multi-Stage Clinical Trials," *Journal of Statistical Software*, 88, 1–25, DOI: 10.18637/jss.v088.i04. [488,489]
- Jennison, C., and Turnbull, B. W. (1997), "Group-Sequential Analysis Incorporating Covariate Information," *Journal of the American Statistical Association*, 92, 1330–1341. [491]
- (1999), *Group Sequential Methods With Applications to Clinical Trials*, Boca Raton, FL: CRC Press. [487,491]
- Kieser, M., Bauer, P., and Lehmacher, W. (1999), "Inference on Multiple Endpoints in Clinical Trials With Adaptive Interim Analyses," *Biometrical Journal*, 41, 261–277. [487,490]
- Kim, K., and Demets, D. L. (1987), "Design and Analysis of Group Sequential Tests Based on the Type I Error Spending Rate Function," *Biometrika*, 74, 149–154. [485,486,488,490]
- Kimani, P. K., Todd, S., and Stallard, N. (2013), "Conditionally Unbiased Estimation in Phase II/III Clinical Trials With Early Stopping for Futility," *Statistics in Medicine*, 32, 2893–2910. [489]
- (2014), "A Comparison of Methods for Constructing Confidence Intervals After Phase II/III Clinical Trials," *Biometrical Journal*, 56, 107–128. [490]
- Koenig, F., Brannath, W., Bretz, F., and Posch, M. (2008), "Adaptive Dunnett Tests for Treatment Selection," *Statistics in Medicine*, 27, 1612–1625. [489]
- Koopmeiners, J. S., Feng, Z., and Pepe, M. S. (2012), "Conditional Estimation After a Two-Stage Diagnostic Biomarker Study That Allows Early Termination for Futility," *Statistics in Medicine*, 31, 420–435. [489]
- Kunz, C. U., Jorgens, S., Bretz, F., Stallard, N., Van Lancker, K., Xi, D., Zohar, S., Gerlinger, C., and Friede, T. (2020), "Clinical Trials Impacted by the COVID-19 Pandemic: Adaptive Designs to the Rescue?" arXiv no. 2005.13979. [493]

- Lehmacher, W., and Wassmer, G. (1999), "Adaptive Sample Size Calculations in Group Sequential Trials," *Biometrics*, 55, 1286–1290. [488]
- London, A. J., and Kimmelman, J. (2020), "Against Pandemic Research Exceptionalism," *Science*, 368, 476–477. [484]
- Magaret, A., Angus, D. C., Adhikari, N. K. J., Banura, P., Kissoon, N., Lawler, J. V., and Jacob, S. T. (2016), "Design of a Multi-Arm Randomized Clinical Trial With No Control Arm," *Contemporary Clinical Trials*, 46, 12–17, DOI: 10.1016/j.cct.2015.11.003. [489]
- Magirr, D., Jaki, T., Posch, M., and Klinglmueller, F. (2013), "Simultaneous Confidence Intervals That Are Compatible With Closed Testing in Adaptive Designs," *Biometrika*, 100, 985–996. [489,490]
- Magirr, D., Jaki, T., and Whitehead, J. (2012), "A Generalized Dunnett Test for Multi-Arm Multi-Stage Clinical Studies With Treatment Selection," *Biometrika*, 99, 494–501. [487,488,489]
- Magirr, D., Stallard, N., and Jaki, T. (2014), "Flexible Sequential Designs for Multi-Arm Clinical Trials," *Statistics in Medicine*, 33, 3269–3279. [489]
- Marcus, R., Peritz, E., and Gabriel, K. R. (1976), "On Closed Testing Procedures With Special Reference to Ordered Analysis of Variance," *Biometrika*, 63:655–660. [488]
- Meyer, E. L., Mesenbrink, P., Dunger-Baldauf, C., Fülle, H.-J., Glimm, E., Li, Y., Posch, M., and König, F. (2020), "The Evolution of Master Protocol Clinical Trial Designs—A Systematic Literature Review," *Clinical Therapeutics* (to appear). [487]
- Mielke, T., and Dragalin, V. (2017), "Two-Stage Designs in Dose Finding," in *Handbook of Methods for Designing, Monitoring, and Analyzing Dose-Finding Trials*, eds. J. O'Quigley, A. Iasonos, and B. Bornkamp, Boca Raton, FL: CRC Press, Taylor and Francis Group, pp. 247–265. [490]
- Mozgunov, P., Jaki, T., and Paoletti, X. (2019), "Randomized Dose-Escalation Designs for Drug Combination Cancer Trials With Immunotherapy," *Journal of Biopharmaceutical Statistics*, 29, 359–377. [484]
- (2020), "Using a Dose-Finding Benchmark to Quantify the Loss Incurred by Dichotomisation in Phase II Dose-Ranging Studies," *Biometrical Journal*, DOI: 10.1002/bimj.201900332. [490]
- Mulangu, S., Dodd, L. E., Davey Jrm, R. T., Tshiani Mbaya, O., Proschan, M., Mukadi, D., Lusakibanza Manzo, M., Nzolo, D., Tshomba Oloma, A., Ibanda, A., and Ali, R. (2019), "A Randomized, Controlled Trial of Ebola Virus Disease Therapeutics," *New England Journal of Medicine*, 381, 2293–2303. [483]
- Müller, H. H., and Schäfer, H. (2004), "A General Statistical Principle for Changing a Design Any Time During the Course of a Trial," *Statistics in Medicine*, 23, 2497–2508. [487,488]
- National Academies of Sciences, Engineering, and Medicine (2017), *Integrating Clinical Research Into Epidemic Response: The Ebola Experience*, Washington, DC: The National Academies Press, DOI: 10.17226/24739. [483]
- Neal, D., Casella, G., Yang, M. C., and Wu, S. S. (2011), "Interval Estimation in Two-Stage Drop-the-Losers Clinical Trials With Flexible Treatment Selection," *Statistics in Medicine*, 30, 2804–2814. [490]
- Neuenschwander, B., and Schmidli, H. (2020), "Use of Historical Data," in *Bayesian Methods in Pharmaceutical Research*, eds. E. Lesaffre, G. Baio, and B. Boulanger, Boca Raton, FL: Chapman & Hall/CRC. [492]
- O'Brien, P. C., and Fleming, T. R. (1979), "A Multiple Testing Procedure for Clinical Trials," *Biometrics*, 35, 549–556. [488]
- Ondra, T., Dmitrienko, A., Friede, T., Graf, A., Miller, F., Stallard, N., and Posch, M. (2016), "Methods for Identification and Confirmation of Targeted Subgroups in Clinical Trials: A Systematic Review," *Journal of Biopharmaceutical Statistics*, 26, 99–119. [491]
- Ondra, T., Jobjörnsson, S., Beckman, R. A., Burman, C. F., König, F., Stallard, N., and Posch, M. (2019), "Optimized Adaptive Enrichment Designs," *Statistical Methods in Medical Research*, 28, 2096–2111. [491]
- Pallmann, P., Bedding, A. W., Choodari-Oskoei, B., Dimairo, M., Flight, L., Hampson, L. V., Holmes, J., Mander, A. P., Odondi, L., Sydes, M. R., Villar, S. S., Wason, J. M. S., Weir, C. J., Wheeler, G. M., Yap, C., and Jaki, T. (2018), "Adaptive Designs in Clinical Trials: Why Use Them, and How to Run and Report Them," *BMC Medicine*, 16, 29. [487]
- Parmar, M. K. B., Carpenter, J., and Sydes, M. R. (2014), "More Multiarm Randomized Trials of Superiority Are Needed," *The Lancet*, 384, 283–284, DOI: 10.1016/S0140-6736(14)61122-3. [487]
- Pinheiro, J., Bornkamp, B., Glimm, E., and Bretz, F. (2014), "Model-Based Dose Finding Under Model Uncertainty Using General Parametric Models," *Statistics in Medicine*, 33, 1646–1661. [490]
- Proschan, M. A., Dodd, L. E., and Price, D. (2016), "Statistical Considerations for a Trial of Ebola Virus Disease Therapeutics," *Clinical Trials*, 13, 39–48. [483]
- Proschan, M. A., and Evans, S. (2020), "The Temptation of Response-Adaptive Randomization," *Clinical Infectious Diseases*, DOI: 10.1093/cid/ciaa334. [489]
- Proschan, M. A., and Hunsberger, S. A. (1995), "Designed Extension of Studies Based on Conditional Power," *Biometrics*, 51, 1315–1324. [487,488]
- Robertson, D. S., and Glimm, E. (2019), "Conditionally Unbiased Estimation in the Normal Setting With Unknown Variances," *Communications in Statistics—Theory and Methods*, 48, 616–627. [490]
- Robertson, D. S., Prevost, A. T., and Bowden, J. (2016), "Unbiased Estimation in Seamless Phase II/III Trials With Unequal Treatment Effect Variances and Hypothesis-Driven Selection Rules," *Statistics in Medicine*, 35, 3907–3922. [489]
- Robertson, D. S., and Wason, J. M. S. (2018), "Online Control of the False Discovery Rate in Biomedical Research," arXiv no. 809.07292. [489]
- (2019), "Familywise Error Control in Multi-Armed Response-Adaptive Trials," *Biometrics*, 75, 885–894, DOI: 10.1111/biom.13042. [489]
- Royston, P., Parmar, M. K. B., and Qian, W. (2003), "Novel Designs for Multi-Arm Clinical Trials With Survival Outcomes With an Application in Ovarian Cancer," *Statistics in Medicine*, 22, 2239–2256, DOI: 10.1002/sim.1430. [488]
- rpact (2020), "rpact: Confirmatory Adaptive Clinical Trial Design and Analysis | RPACT," available at <https://www.rpact.org/>. [489]
- Ruan, Q., Yang, K., Wang, W., Jiang, L., and Song, J. (2020), "Clinical Predictors of Mortality Due to COVID-19 Based on an Analysis of Data of 150 Patients From Wuhan, China," *Intensive Care Medicine*. [486]
- Sampson, A. R., and Sill, M. W. (2005), "Drop-the-Losers Design: Normal Case," *Biometrical Journal*, 47, 257–268. [489]
- Schiavone, F., Bathia, R., Letchemanan, K., Masters, L., Amos, C., Bara, A., Brown, L., Gilson, C., Pugh, C., Atako, N., Hudson, F., Parmar, M., Langley, R., Kaplan, R. S., Parker, C., Attard, G., Clarke, N. W., Gillissen, S., James, N. D., Maughan, T., and Sydes, M. R. (2019), "This Is a Platform Alteration: A Trial Management Perspective on the Operational Aspects of Adaptive and Platform and Umbrella Protocols," *Trials*, 20, 264, DOI: 10.1186/s13063-019-3216-8. [487]
- Schmidli, H., Bretz, F., and Racine-Poon, A. (2007), "Bayesian Predictive Power for Interim Adaptation in Seamless Phase II/III Trials Where the Endpoint Is Survival up to Some Specified Timepoint," *Statistics in Medicine*, 26, 4925–4938. [488]
- Simmonds, M., Salanti, G., McKenzie, J., and Elliot, J. (2017), "Living Systematic Reviews: 3. Statistical Methods for Updating Meta-Analysis," *Journal of Clinical Epidemiology*, 91, 38–46. [492]
- Stallard, N., and Friede, T. (2008), "A Group-Sequential Design for Clinical Trials With Treatment Selection," *Statistics in Medicine*, 27, 6209–6227. [487]
- Stallard, N., and Kimani, P. K. (2018), "Uniformly Minimum Variance Conditionally Unbiased Estimation in Multi-Arm Multi-Stage Clinical Trials," *Biometrika*, 105, 495–501. [489,490]
- Stallard, N., and Todd, S. (2005), "Point Estimates and Confidence Regions for Sequential Trials Involving Selection," *Journal of Statistical Planning and Inference*, 135, 402–419. [489]
- Stallard, N., Todd, S., Parashar, D., Kimani, P., and Renfro, L. (2019), "On the Need to Adjust for Multiplicity in Confirmatory Clinical Trials With Master Protocols," *Annals of Oncology*, 30, 506. [489]
- Stallard, N., Todd, S., Ryan, E. G., and Gates, S. (2020), "Comparison of Bayesian and Frequentist Group-Sequential Clinical Trial Designs," *BMC Medical Research Methodology*, 20, 4. [488]
- Thall, P., Fox, P., and Wathen, J. (2015), "Statistical Controversies in Clinical Research: Scientific and Ethical Problems With Adaptive Randomization in Comparative Clinical Trials," *Annals of Oncology*, 26, 1621–1628. [489]
- Thielman, N. M., Cunningham, C. K., Woods, C., Petzold, E., Spreng, M., and Russell, J. (2016), "Ebola Clinical Trials: Five Lessons Learned and a Way Forward," *Clinical Trials*, 13, 83–86. [483]

- Villar, S. S., Bowden, J., and Wason, J. (2018), "Response-Adaptive Designs for Binary Responses: How to Offer Patient Benefit While Being Robust to Time Trends?," *Pharmaceutical Statistics*, 17, 182–197, DOI: 10.1002/pst.1845. [489]
- Wang, Y., Zhang, D., Du, G., Du, R., Zhao, J., Jin, Y., Fu, S., Gao, L., Cheng, Z., Lu, Q., Hu, Y., Luo, G., Wang, K., Lu, Y., Li, H., Wang, S., Ruan, S., Yang, C., Mei, C., Wang, Y., Ding, D., Wu, F., Tang, X., Ye, X., Ye, Y., Liu, B., Yang, J., Yin, W., Wang, A., Fan, G., Zhou, F., Liu, Z., Gu, X., Xu, J., Shang, L., Zhang, L., Cao, L., Guo, T., Wan, Y., Qin, H., Jiang, Y., Jaki, T., Hayden, F. G., Horby, P. W., Cao, B., and Wang, C. (2020), "Remdesivir in Adults With Severe COVID-19: A Randomized, Double-Blind, Placebo-Controlled, Multicentre Trial," *The Lancet*, 395, 1569–1578. [483]
- Wason, J. M. S., and Jaki, T. (2012), "Optimal Design of Multi-Arm Multi-Stage Trials," *Statistics in Medicine*, 31, 4269–4279. [487,488]
- Wason, J., Magirr, D., Law, M., and Jaki, T. (2016), "Some Recommendations for Multi-Arm Multi-Stage Trials," *Statistical Methods in Medical Research*, 25, 716–727. [487,489]
- Wason, J., Stallard, N., Bowden, J., and Jennison, C. (2017), "A Multi-Stage Drop-the-Losers Design for Multi-Arm Clinical Trials," *Statistical Methods in Medical Research*, 26, 508–524, DOI: 10.1177/0962280214550759. [489]
- Wason, J. M. S., Stecher, L., and Mander, A. P. (2014), "Correcting for Multiple-Testing in Multi-Arm Trials: Is It Necessary and Is It Done?," *Trials*, 15, 364, DOI: 10.1186/1745-6215-15-364. [489]
- Wason, J. M. S., and Trippa, L. (2014), "A Comparison of Bayesian Adaptive Randomization and Multi-Stage Designs for Multi-Arm Clinical Trials," *Statistics in Medicine*, 33, 2206–2221. [489]
- Wassmer, G. (2011), "On Sample Size Determination in Multi-Armed Confirmatory Adaptive Designs," *Journal of Biopharmaceutical Statistics*, 21, 802–817, DOI: 10.1080/10543406.2011.551336. [489]
- Wassmer, G., and Brannath, W. (2016), *Group Sequential and Confirmatory Adaptive Designs in Clinical Trials*, Heidelberg: Springer. [487,489]
- Whitehead, A. (1997), "A Prospectively Planned Cumulative Meta-Analysis Applied to a Series of Concurrent Clinical Trials," *Statistics in Medicine*, 16, 2901–2913. [492]
- Whitehead, J. (1997), *The Design and Analysis of Sequential Clinical Trials*, Chichester: Wiley. [485,488]
- (1986), "On the Bias of Maximum Likelihood Estimation Following a Sequential Test," *Biometrika*, 73, 573–581. [489]
- (2011), "Group Sequential Trials Revisited: Simple Implementation Using SAS," *Statistical Methods in Medical Research* 20, 636–656, Corrigenda (2017) 26, 2481. [488]
- Whitehead, J., Desai, Y., and Jaki, T. (2020), "Estimation of Treatment Effects Following a Sequential Trial of Multiple Treatments," *Statistics in Medicine*, 39, 1593–1609. [489]
- Whitehead, J., and Horby, P. (2017), "GOST: A Generic Ordinal Sequential Trial Design for a Treatment Trial in an Emerging Pandemic," *PLoS Neglected Tropical Diseases*, 11, e0005439, Corrigenda (2018), 12(4), e0006414. [490,493]
- Williamson, S. F., and Villar, S. S. (2020), "A Response-Adaptive Randomization Procedure for Multi-Armed Clinical Trials With Normally Distributed Outcomes," *Biometrics*, 76, 197–209. [489]
- Wu, S. S., Wang, W., and Yang, M. C. K. (2010), "Interval Estimation for Drop-the-Losers Design," *Biometrika*, 97, 405–418. [489,490]
- Yu, Z., Ramakrishnan, V., and Meinzer, C. (2019), "Simulation Optimization for Bayesian Multi-Arm Multi-Stage Clinical Trial With Binary Endpoints," *Journal of Biopharmaceutical Statistics*, 29, 306–317. [488]
- Zhou, F., Yu, T., Du, R., Fan, G., Liu, Y., Liu, Z., Xiang, J., Wang, Y., Song, D. B., Gu, X., Guan, D. L., Wei, Y., Li, H., Wu, X., Xu, J., Tu, S., Zhang, Y., Chen, H., and Cao, B. (2020), "Clinical Course and Risk Factors for Mortality of Adult Inpatients With COVID-19 in Wuhan, China: A Retrospective Cohort Study," *The Lancet*, 395, 1054–1062. [491]