

Online error rate control for platform trials

David S. Robertson¹  | James M. S. Wason²  | Franz König³  |
 Martin Posch³  | Thomas Jaki^{1,4} 

¹MRC Biostatistics Unit, School of Clinical Medicine, University of Cambridge, Cambridge, UK

²Population Health Sciences Institute, Faculty of Medical Sciences, Newcastle University, Newcastle upon Tyne, UK

³Section of Medical Statistics, Medical University of Vienna, Vienna, Austria

⁴Faculty of Informatics and Data Science, University of Regensburg, Regensburg, Germany

Correspondence

David S. Robertson, MRC Biostatistics Unit, School of Clinical Medicine, University of Cambridge, Cambridge, UK.
 Email: david.robertson@mrc-bsu.cam.ac.uk

Funding information

Biometrika Trust; EU-PEARL, Grant/Award Number: 853966; Medical Research Council, Grant/Award Number: MC_UU_00002/14; National Institute for Health Research, Grant/Award Number: BRC1215-20014

Platform trials evaluate multiple experimental treatments under a single master protocol, where new treatment arms are added to the trial over time. Given the multiple treatment comparisons, there is the potential for inflation of the overall type I error rate, which is complicated by the fact that the hypotheses are tested at different times and are not necessarily pre-specified. Online error rate control methodology provides a possible solution to the problem of multiplicity for platform trials where a relatively large number of hypotheses are expected to be tested over time. In the online multiple hypothesis testing framework, hypotheses are tested one-by-one over time, where at each time-step an analyst decides whether to reject the current null hypothesis without knowledge of future tests but based solely on past decisions. Methodology has recently been developed for online control of the false discovery rate as well as the familywise error rate (FWER). In this article, we describe how to apply online error rate control to the platform trial setting, present extensive simulation results, and give some recommendations for the use of this new methodology in practice. We show that the algorithms for online error rate control can have a substantially lower FWER than uncorrected testing, while still achieving noticeable gains in power when compared with the use of a Bonferroni correction. We also illustrate how online error rate control would have impacted a currently ongoing platform trial.

KEY WORDS

multiple testing, online hypothesis testing, platform trial, type I error rate

1 | INTRODUCTION

There is a strong need to conduct high-quality evaluations of new interventions aimed at improving the health of patients. The highest quality of evidence comes from randomized controlled trials (RCTs). However, the cost of RCTs is high and increasing, leading to the very high costs of bringing new drugs to market¹ and evaluation of other types of intervention.² This has led to focus on methods that can improve the operational and statistical efficiency of conducting clinical trials.

One important class of efficient approaches are platform trials.³⁻⁵ Platform trials set up an infrastructure that allows evaluation of multiple intervention arms under a single master protocol. Platform trials may use an adaptive design to drop non-promising intervention arms early and can add new interventions as they become available for evaluation. Some

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

© 2023 The Authors. *Statistics in Medicine* published by John Wiley & Sons Ltd.

interventions can also be evaluated only within particular patient subgroups if warranted. These properties all can lead to increased efficiency.^{3,6,7} However, they also introduce challenges, both operational⁸ and statistical.⁹

One statistical challenge, given the multiple treatment comparisons, is control of type I error rates. Although the need for formal control of type I error rates in trials of distinct treatments is controversial,¹⁰⁻¹⁶ it may be required in some regulatory settings and desirable in other situations where the treatments are related in some way. In a platform trial, controlling type I error rates is complicated by the fact that the hypotheses are tested at different times and are not all pre-specified.

Online error rate control methodology provides a possible solution to the problem of multiplicity for platform trials where multiple treatments are tested over time.^{17,18} In the online multiple hypothesis testing framework, hypotheses are tested one-by-one over time, where at each time-step a decision is made whether to reject the current null hypothesis without knowledge of future hypothesis tests but based solely on past decisions. Methodology has recently been developed for online control of the false discovery rate (FDR)¹⁹⁻²⁴ as well as the familywise error rate (FWER),²⁵ so that the relevant error rate is controlled at all times throughout the trial.

In order to apply online error rate control methodology to platform trials and evaluate its performance, there are some general features specific to this setting that need to be explored. First, one key consideration is the total number of hypothesis tests that is reasonable to envisage for a platform trial. Thus far, most of the literature on online testing has tended to focus on applications with a very large (≥ 1000) number of hypotheses. Second, many existing algorithms^{19-22,25} for online FDR or FWER control assume independence between the P -values (or equivalently, the test statistics). However, in the platform trial setting there is a positive dependence induced by shared control arm information, and so potential inflation of the relevant type I error rate needs to be considered. Lastly, in the clinical trial setting the usual error rate considered (at least for confirmatory trials) is the FWER.^{26,27} Hence, it would be useful to see how the FWER is impacted when using methods that control the FDR.

Our aim in this article is to demonstrate the advantages and disadvantages of using online error rate control methodology in the platform trial setting through an extensive simulation study looking at varying numbers of treatment arms, patterns of expected treatment responses and arm entry times, as well as assumptions about the upper bound on the total number of experimental treatments to be tested.

In Section 2, we introduce the different algorithms for online error rate control that we will explore, and then in Section 3 we describe the idealized platform trial set-up that we use for the simulations and the different simulation scenarios. Section 4 provides the simulation results, focusing on testing treatments one-by-one (Section 4.1) as well as in batches (Section 4.2). In Section 5, we present a case study based on the STAMPEDE platform trial²⁸ of therapies for prostate cancer. We conclude with recommendations and discussion in Section 6.

2 | ALGORITHMS FOR ONLINE ERROR RATE CONTROL

We now briefly introduce the online multiple hypothesis testing algorithms that we consider, classified into those for fully online testing (ie, where each hypothesis is tested one after another) and those for online batched testing (ie, where groups of hypotheses are available to be tested at the same time). An illustration of the difference is given in Figure 1. All of the algorithms require a pre-specified significance level α , as well as a sequence of non-negative numbers γ_i that sum to 1. We provide references giving further details of the algorithms for the interested reader. The algorithms described below are also available to use via the `onlineFDR` R package.²⁹

2.1 | Fully online testing

The algorithms for fully online testing output an adjusted testing level α_i for testing the null hypothesis H_i , that is, H_i is rejected if the associated P -value p_i satisfies $p_i \leq \alpha_i$.

- **LOND:** One of the first online testing algorithms proposed, LOND (levels based on number of discoveries) provably controls the FDR for independent¹⁹ and positively dependent P -values.²³ In LOND, the α_i are calculated by simply multiplying the sequence $\alpha\gamma_i$ by the number of discoveries (ie, the number of rejected hypotheses) that have been made thus far.
- **LORD:** Another early online testing algorithm, LORD (levels based on recent discovery) provably controls the FDR for independent P -values.^{19,20} LORD takes advantage of so-called “alpha-investing,” in which hypothesis tests costs

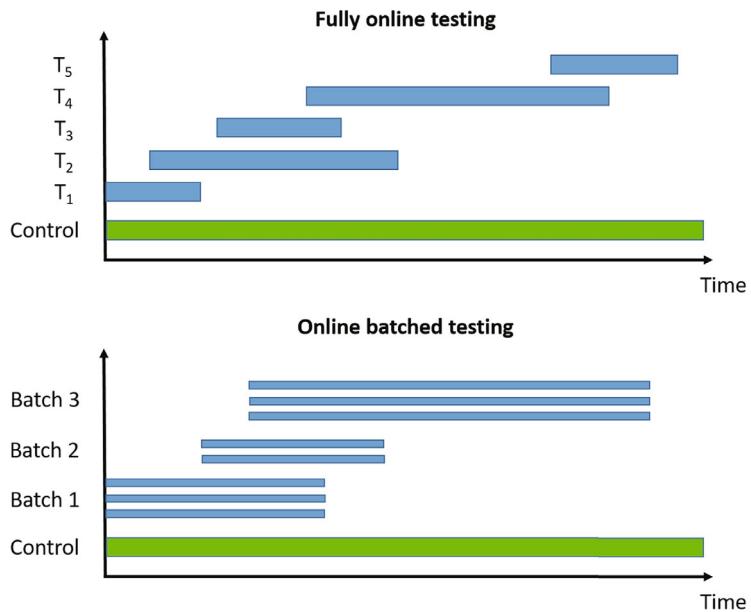


FIGURE 1 Illustration of fully online testing and online batched testing for a platform trial testing experimental treatments T_i against a common control. For the online testing procedures, a hypothesis test can be conducted whenever new results become available (ie, whichever treatment arm is the i th one to be analysed is tested at level α_i). For offline methods (eg, the Benjamini-Hochberg procedure), testing can only occur when all the trial data are available.

some amount from the error budget (or “alpha-wealth”), but a discovery earns some of the error budget back. The α_i in LORD depend not only on how many discoveries have been made but also on the timing of these discoveries. The more recent discoveries there are, the higher the alpha-wealth will be.

- **SAFFRON:** Proposed as an improvement to LORD, the SAFFRON (serial estimate of the alpha fraction that is futilely rationed on true null hypotheses) algorithm provably controls the FDR for independent P -values.²¹ Intuitively, SAFFRON focuses on the stronger signals in an experiment (ie, the smaller P -values). By never rejecting weaker signals (ie, larger P -values), SAFFRON preserves alpha-wealth. When a substantial fraction of null-hypotheses are false, SAFFRON will often be more powerful than LORD.
- **ADDIS:** Proposed as an additional improvement to both LORD and SAFFRON, the ADDIS (ADaptive algorithm that DIScards conservative nulls) algorithm provably controls the FDR for independent P -values.²⁵ Intuitively, ADDIS builds on the SAFFRON algorithm by potentially investing alpha-wealth more effectively through the discarding of the weakest signals (ie, the largest P -values) in a principled way. This procedure can gain appreciable power if the null P -values are conservative, that is, stochastically larger than the uniform distribution.
- **ADDIS-spending:** Unlike all of the algorithms mentioned above, ADDIS-spending provably controls the FWER (in the strong sense) for independent P -values.²² ADDIS-spending shares similarities with ADDIS in setting thresholds based on the size of the P -values for hypotheses to be selected for testing.

2.2 | Online batched testing

The algorithms for online batched testing use well-known offline procedures for FDR control (ie, procedures that require all of the P -values to be known before testing) for each batch, in such a way that the FDR is controlled across all of the batches over time. Three related algorithms were proposed by Zrnic et al:²⁴

- **BatchBH:** Provably controls the FDR when the P -values are independent within and across batches. BatchBH runs the Benjamini-Hochberg (BH) procedure at level α_i on each batch, where the values of α_i depend on the number of previous discoveries.
- **BatchPRDS:** A modification of BatchBH, which provably controls the FDR when the P -values in one batch are positively dependent, and independent across batches.

- *BatchStBH*: Provably controls the FDR when the P -values are independent within and across batches. BatchStBH runs the Storey Benjamini-Hochberg (StBH) procedure at level α_i on each batch, where the values of α_i depend on the number of previous discoveries. The StBH procedure can potentially make more rejections than BH by adapting to the (estimated) number of non-nulls.

3 | SIMULATION STUDY

We now describe the platform trial set-up (Section 3.1) used in our simulation study, as well as the simulation scenarios (Section 3.2) that we explore.

3.1 | Platform trial set-up

Consider an idealized platform trial that eventually tests a total of K experimental treatments T_1, \dots, T_K against a common control T_0 (see Figure 13 for a graphical example). Let μ_i denote the expected response for treatment T_i , and define the treatment effect $\theta_i = \mu_i - \mu_0$ for $i \in \{1, \dots, K\}$. The null hypotheses of interest are given by $H_{0i} : \theta_i \leq 0$ with corresponding alternatives $H_{1i} : \theta_i > 0$, which we test at error level α (for either FWER or FDR control). Let n_i denote the pre-specified sample size for each experimental treatment, and for simplicity assume that an equal number of patients are eventually allocated to each experimental arm (ie, $n_1 = \dots = n_K = n$). We use the simplifying assumption that an experimental treatment continues in the trial until n patients have been allocated to that arm and their outcomes observed. The experimental treatment is then formally tested for effectiveness by comparing with *concurrent* controls (ie, only the outcomes from patients allocated to the control group during the time that the experimental treatment was active in the trial). We return to the issue of conducting interim analyses in Section 6.

We assume that the observation X_{ij} from patient j on treatment T_i is distributed (at least asymptotically) as $X_{ij} \sim N(\mu_i, \sigma^2)$, where σ is known. The mean response for patients on experimental treatment T_i ($i > 0$) is then given by $\bar{X}_i = \frac{1}{n} \sum_{j=1}^n X_{ij}$, with corresponding concurrent control mean $\bar{X}_0^{(i)} = \frac{1}{N_0(i)} \sum_{j \in \mathcal{T}_i} X_{0j}$, where $N_0(i)$ and \mathcal{T}_i denote the number and index set, respectively, of the control observations that are concurrent with T_i . The comparisons between the experimental treatments and the control are based on normally-distributed test-statistics

$$Z_i = \frac{\bar{X}_i - \bar{X}_0^{(i)}}{\sigma \sqrt{\frac{1}{N_0(i)} + \frac{1}{n}}} \sim N\left(\frac{\theta_i}{\sigma \sqrt{\frac{1}{N_0(i)} + \frac{1}{n}}}, 1\right),$$

where $\bar{X}_i - \bar{X}_0$ is the observed difference in means. The corresponding one-sided P -value for treatment T_i is given by $p_i = 1 - \Phi(Z)$, where Φ is the standard normal cdf.

We assume for simplicity that experimental treatment T_i enters the trial at time t_i , and remains in the trial until time $(t_i + r)$. The platform trial starts at time 0 and finishes once all K experimental treatments have been tested. Hence the temporal structure of the platform trial can be described by the value of r and the vector $\mathbf{t} = (t_1, \dots, t_K)$. Without loss of generality, we impose the restriction that $t_1 \leq \dots \leq t_K$. We assume that n/r patients are allocated to each arm (including the control) per unit time, where r is chosen so that n/r is an integer.

3.2 | Simulation scenarios

In our simulation study, we consider $K \in \{5, 10, 15, 20\}$, and set $\alpha = 0.025$, $n = 50$, $\mu_0 = 0$, $\sigma = 1$ and $r = 10$. Table 1 gives an overview of the scenarios, algorithms and performance metrics used in the simulation study, with further details given below.

We consider the following scenarios for the treatment means μ_i , $i \in \{1, \dots, K\}$:

1. Global null: $\mu_1 = \dots = \mu_K = 0$

TABLE 1 Overview of simulation scenarios and performance metrics.

Treatment means	Arm entry times	Algorithms	Performance metrics
Global null	All arms at once	Bonferroni, ADDIS-spending	FWER and FDR
Fixed means	Batches of size 5	ADDIS, LORD, SAFFRON, BH	Disjunctive power
Staircase scenario	Staggered starts One-by-one	BatchBH, BatchPRDS, BatchStBH	Sensitivity

2. Fixed means: m effective treatments which are either all tested at the start or the end of the platform trial, or are tested in a random order. We set $\mu_i = 0.5$ for $i \in I$ and $\mu_i = 0$ otherwise, where

$$I = \{i : 1 \leq i \leq m\} \text{ (early).}$$

$$I = \{i : K - m + 1 \leq i \leq K\} \text{ (late).}$$

I is a set of m values drawn randomly from $\{1, \dots, K\}$ without replacement (random) for $m \in \{1, K/5 + 1, 2K/5 + 1\}$.

3. Staircase scenario:

$$\mu_i = (i - \lceil K/2 \rceil)/K \text{ (increasing).}$$

$$\mu_i = (\lceil K/2 \rceil - i + 1)/K \text{ (decreasing).}$$

μ_i is drawn randomly from the set $\{(i - \lceil K/2 \rceil)/K : 1 \leq i \leq K\}$ without replacement (random).

We also consider the following patterns of arm entry times t :

1. All arms at once (offline testing): $t_i = 0$.
2. Batches of size b : $t_i = r(j-1)$ for $i \in \{b(j-1)+1, \dots, bj\}$ and $j \in \{1, \dots, K/b\}$, where $b = 5$.
3. Staggered starts: $t_i = r(i-1)/s$ for $s \in \{2, 5\}$.
4. One-by-one (fully online): $t_i = r(i-1)$.

Finally, we explore different assumptions about the upper bound, N_{bound} , on the total number of experimental treatments to be tested. We consider $N_{\text{bound}} \in \{K, 2K, 5K\}$, that is, considering the impact of more conservative choices of N_{bound} . Note that N_{bound} should be chosen before the start of the platform trial. We return to the issue of what happens if $N_{\text{bound}} < K$ in Section 6.

For each of the simulation scenarios described above, we compare the following algorithms for FWER and FDR control:

FWER: Bonferroni and ADDIS-spending.

FDR: ADDIS, SAFFRON, and LOND. For batched entry times, we consider the BatchBH, BatchPRDS, and BatchStBH procedures. We also use the standard Benjamini-Hochberg (BH) procedure as a comparator.

We give more details about the exact implementation of these algorithms in Section A of the Supporting web materials. Note that Bonferroni tests each hypothesis at level α/N_{bound} (as opposed to α/K), and the BH procedure is an offline procedure that requires all hypotheses to be tested at once (and hence could not be used in practice for a platform trial).

For each simulation scenario, we report the following performance measures:

Type I error rates: FWER and FDR.

Power metrics: Disjunctive power (probability of rejecting at least one non-null hypothesis) and sensitivity (proportion of non-null hypotheses that are rejected).

The use of disjunctive power in our comparisons reflects how it is the power analog of FWER. However, as pointed out by a reviewer, arguably the most relevant definition of “power” when making multiple comparisons is the power to detect a given treatment effect and hence sensitivity may be a more relevant metric to use. We note that for simplicity we have kept type I error and power metrics separate in the results that follow. Alternatively, it could be useful to consider a metric that simultaneously incorporates both power and type I error rates. However, to combine both metrics in this way would require the specification of the relative costs of making type I and type II errors, which would vary across different stakeholders.

4 | SIMULATION RESULTS

To start with, we consider the type I error rate of the different algorithms, in particular checking whether the FDR is controlled. This is of interest because apart from LOND, all of the online testing algorithms only provably control the FDR under independence of the P -values. Figure B1 in Section B.1 of the Supporting web materials shows the FDR under the global null (and hence the FDR = FWER), for varying patterns of arm entry times and $N_{\text{bound}} = K$. Note that the global null maximizes the FWER for multi-arm multi-stage (MAMS) designs.³⁰ Using uncorrected testing leads to a highly inflated FDR/FWER above the nominal 2.5%, which can be as high as 40% when $K = 20$. In contrast, the online testing algorithms have a FDR/FWER which is equal to or below the nominal level. The one exception is SAFFRON, which has a slight inflation of the FDR/FWER to around 3%.

4.1 | Fully online setting

Continuing with the FDR and now focusing on the fully online setting, Figure B2 in Section B.2 of the Supporting web materials shows the worst and best case FDR (corresponding to all the effective treatments appearing early and late, respectively) for fixed means and different numbers of effective treatments, with $N_{\text{bound}} = K$. Again we see how uncorrected testing can lead to a highly inflated FDR, although this inflation reduces as the number of effective treatments increases. This time, all the online testing procedures control the FDR at or below the nominal 2.5%, but they can be noticeably conservative. This is the case for LORD, as well as ADDIS-spending for $(K/5 + 1)$ and $(2K/5 + 1)$ effective treatments. Note though that the Bonferroni correction is also highly conservative in these settings.

We now move on to examining the FWER for the different procedures. Of course, the FDR-controlling algorithms would not be expected to also control the FWER in general, but it is of interest to see how large any inflation in the FWER can be. In general, the magnitude of the inflation depends critically on the effect sizes chosen for the alternative hypotheses. Figure 2 shows the FWER under the early and late scenarios, as well as the FWER under the random ordering of the effective treatments, all with $N_{\text{bound}} = K$. When there is only 1 effective treatment the FWER is controlled by all the online testing algorithms, but (apart from LOND) they are even more conservative than Bonferroni. A possible explanation for this observation is that the test statistics are independent in the fully online setting, and the Bonferroni correction is not as highly conservative under independence. Under the global null (eg.) the probability of at least one error when using Bonferroni is $1 - (1 - \alpha/K)^K$ and this is close to α for small α .

When there are $(K/5 + 1)$ effective treatments, the FWER is inflated above 2.5% for all the online testing procedures (apart from ADDIS-spending) when $K > 5$ in the early scenario and $K > 15$ in the random scenario. The LOND algorithm has the largest inflation in the early scenario, which is similar to the inflation seen using the BH procedure. Finally, when there are $(2K/5 + 1)$ effective treatments, there is FWER inflation for all the online algorithms (except for ADDIS-spending) when $K > 5$ in the early scenario and $K > 10$ in the random scenario. This time, SAFFRON has the highest inflation, which can be higher than BH in the early scenario, although still noticeably lower than uncorrected testing. Note that in the random setting, for both $(K/5 + 1)$ and $(2K/5 + 1)$ effective treatments, when $K = 15$ the FWER of the online testing procedures (apart from SAFFRON) is controlled at 5%.

The other side of the story is the power of the trial, and to start with we focus on sensitivity as a measure of the proportion of the truly effective treatments that are declared efficacious. Figure 3 shows the sensitivities for the early, late, and random scenarios, again with $N_{\text{bound}} = K$. For ease of reference, we also include plots showing the sensitivity and FWER together in Figures B3–B5 in Section B.2 of the supporting web materials. With only 1 effective treatment, the sensitivity of all of the online algorithms apart from LOND is substantially lower than using Bonferroni. Even in the best case (early scenario), only SAFFRON and ADDIS-spending have a higher sensitivity than Bonferroni when $K \geq 10$. When there are $(K/5 + 1)$ effective treatments, again the sensitivities of all of the online algorithms apart from LOND are lower than Bonferroni in the late and random scenarios. LOND maintains a sensitivity roughly halfway that between the Bonferroni and BH procedures. In the best case (early scenario), SAFFRON has a higher sensitivity than Bonferroni for $K > 5$ (and even BH for $K \geq 15$), while the other online algorithms also achieve a higher sensitivity for $K > 10$. Finally, with $(2K/5 + 1)$ effective treatments, we again see that LOND maintains a sensitivity roughly halfway that between the Bonferroni and BH procedures. Under the random scenario, SAFFRON, ADDIS and LORD have a higher sensitivity than Bonferroni for $K > 5$, $K > 10$, and $K \geq 15$, respectively. In the worst case (late scenario), the equivalent values of K

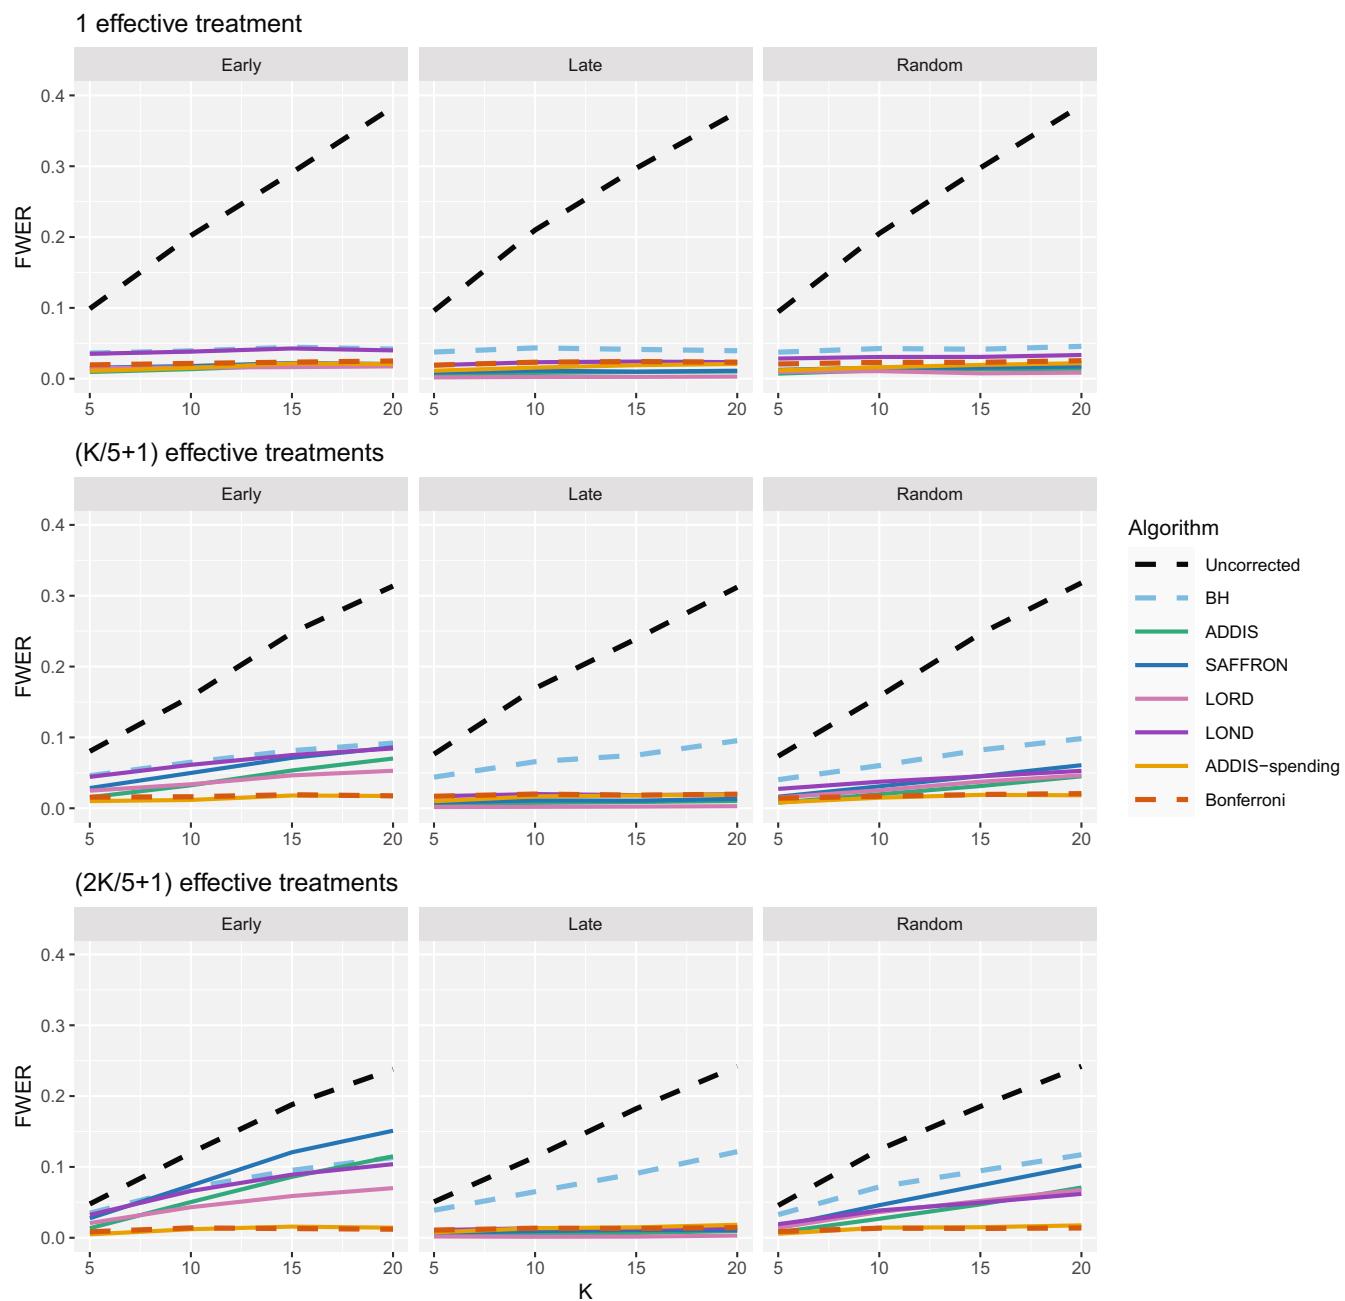


FIGURE 2 FWER for fixed means and different numbers of effective treatments, with $N_{\text{bound}} = K$.

are $K \geq 10$, $K > 10$, and $K > 15$, respectively. Finally, in the best case (early scenario), all the online testing algorithms outperform Bonferroni for $K \geq 10$.

Thus far we have used an upper bound $N_{\text{bound}} = K$, that is, exactly equal to the number of treatments that are tested. In practice, this is unlikely to be known prior to the platform trial starting, and a more conservative (ie, higher) value of N_{bound} would be used. Figure 4 shows the effect of increasing the value of N_{bound} on the sensitivity of the algorithms, where there are $(2K/5 + 1)$ effective treatments. The key takeaway here is that as N_{bound} increases, the sensitivity of Bonferroni, LOND and LORD noticeably decreases, with Bonferroni having the largest relative decrease, whereas the sensitivity of all the other algorithms remains virtually unchanged. This means that there is a greater advantage in using online algorithms (in terms of sensitivity) compared with Bonferroni as N_{bound} increases. For example, in the extreme case where $N_{\text{bound}} = 5K$, then under a random ordering of treatment means, all of the online algorithms outperform Bonferroni, with the relative advantage increasing with K . The same is seen for the more realistic scenario of $N_{\text{bound}} = 2K$, apart from LORD for $K < 10$.

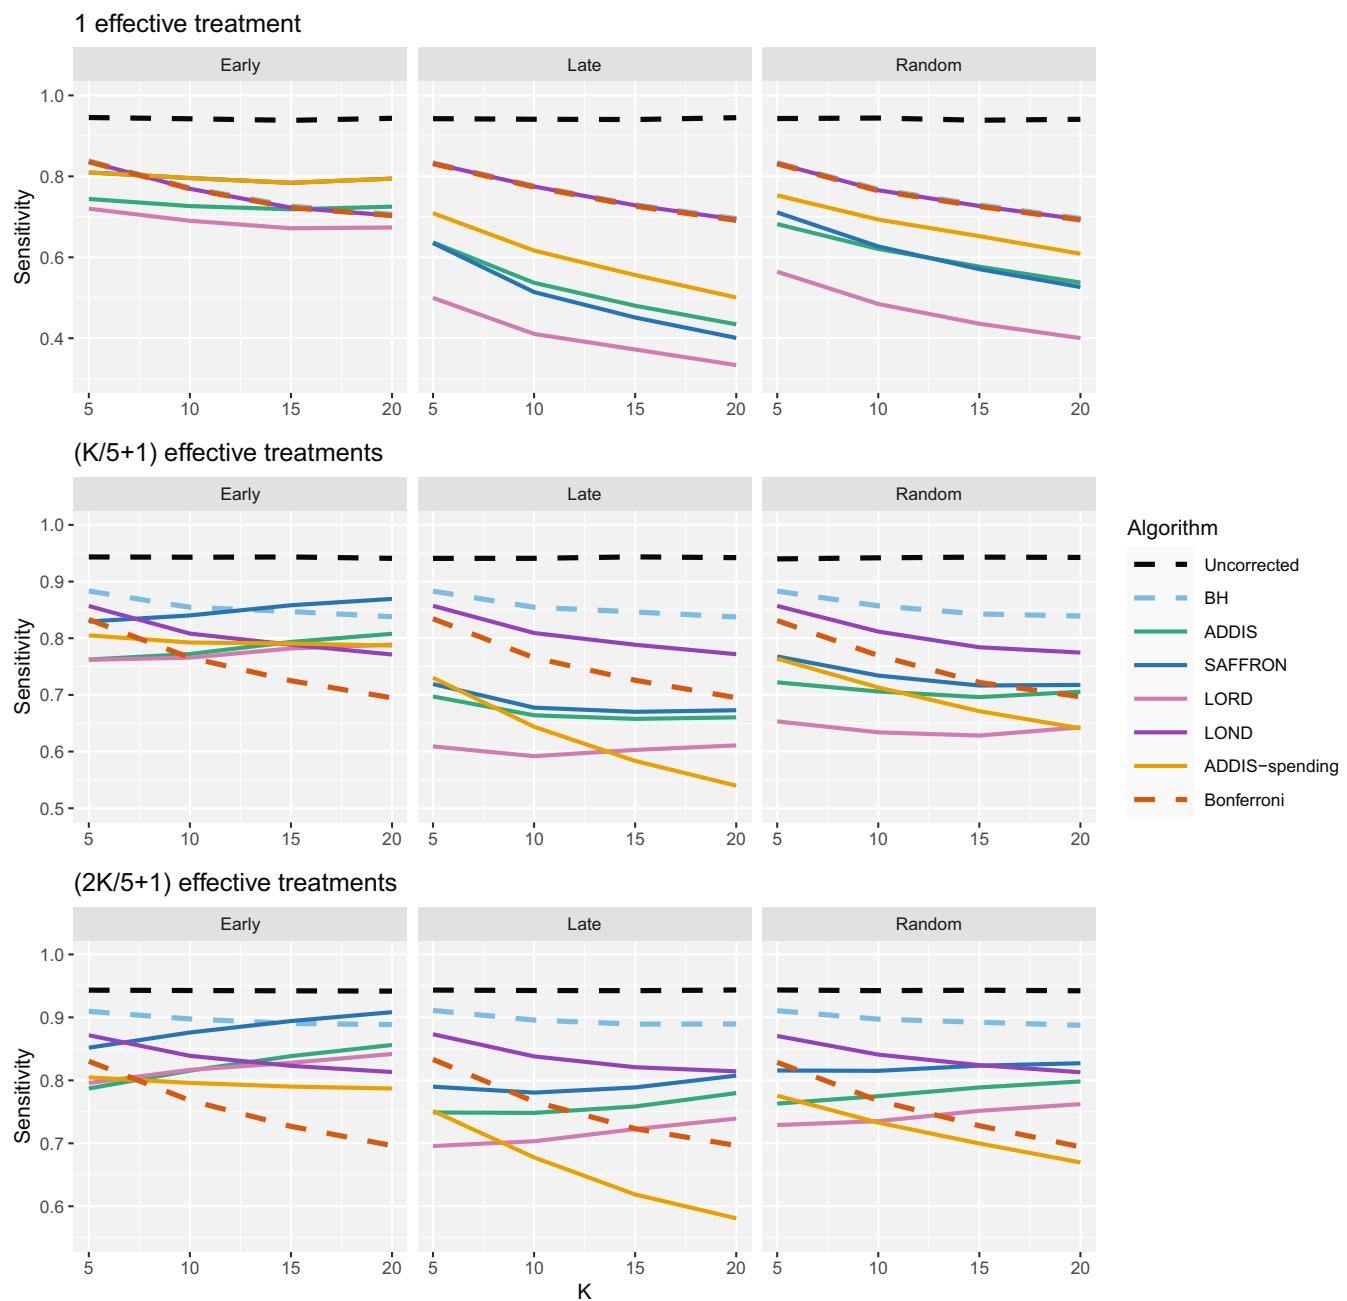


FIGURE 3 Sensitivity for fixed means and different numbers of effective treatments, with $N_{\text{bound}} = K$. Note that the y-axis scales differ as the number of effective treatments vary for readability.

In terms of the effect of N_{bound} on the FWER, Figure 5 shows the trade-off between sensitivity and the FWER under a random ordering of treatment means (the results for the early and late scenarios are given in Figures B6 and B7 in Section B.2 of the Supporting web materials). We see that as N_{bound} increases, the FWER of Bonferroni, LOND and LORD decreases, whereas the FWER of the other algorithms remain virtually unchanged. Hence, when $N_{\text{bound}} = 2K$, LOND controls the FWER at 2.5% while still having substantial gains in power over Bonferroni. Meanwhile, the power gains of SAFFRON (and to a lesser extent, ADDIS) come at the cost of a large inflation of the FWER. However, we note that in order to be conservative it may still be useful to consider the FWER when $N_{\text{bound}} = K$ since it represents a sort of “worst-case” scenario where the assumed upper bound on the total number of treatments is actually reached.

Figure 6 shows the disjunctive power for 1 and $(K/5 + 1)$ effective treatments, with $N_{\text{bound}} = 2K$ (results for $(2K/5 + 1)$ effective treatments are not shown since most methods have a disjunctive power very close to 1). Using this

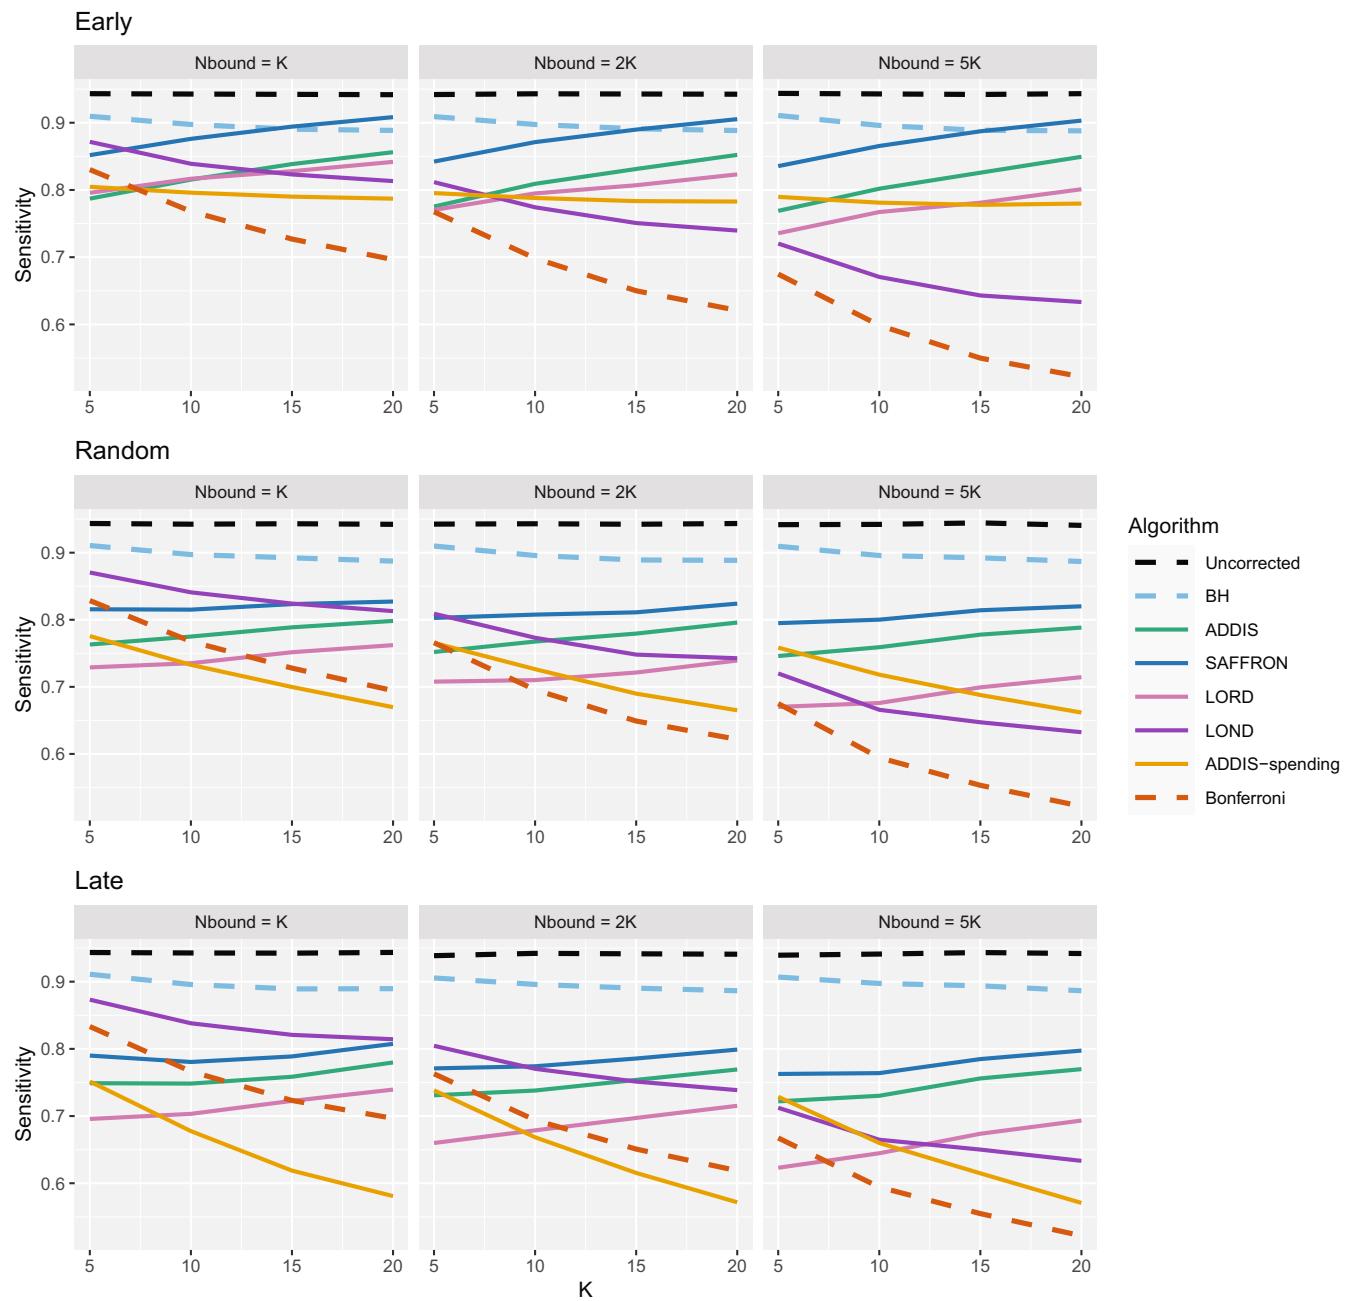


FIGURE 4 Sensitivity for fixed means and $(2K/5 + 1)$ effective treatments, with varying N_{bound} .

power metric, the online algorithms all have substantially lower power than Bonferroni in the random and late scenarios, apart from LOND (and ADDIS-spending for randomly ordered means). Only in the early scenario do we see advantages for ADDIS-spending and (to a lesser extend) ADDIS, particularly when there is only 1 effective treatment.

Finally, we consider the impact of conservative nulls (ie, the staircase scenario) in the fully online setting. A-priori we would expect ADDIS and ADDIS-spending to perform better here. Due to the treatment means used, the sensitivity of all algorithms is low (<50%) and hence we only focus on the disjunctive power (the results for sensitivity are available in Section B.2 of the Supporting web materials, Figure B8). Figure 7 shows the results for the staircase scenario, where $N_{\text{bound}} = 2K$ for the disjunctive power plots and $N_{\text{bound}} = K$ for the FWER plots. In this setting and for this power metric, ADDIS-spending performs particularly well with comparable disjunctive power to Bonferroni in the worst case (ascending) and noticeable increases in power under the random and descending staircase scenarios. Note that the FWER of all the procedures (except uncorrected testing) control the FWER below the nominal 2.5% level.

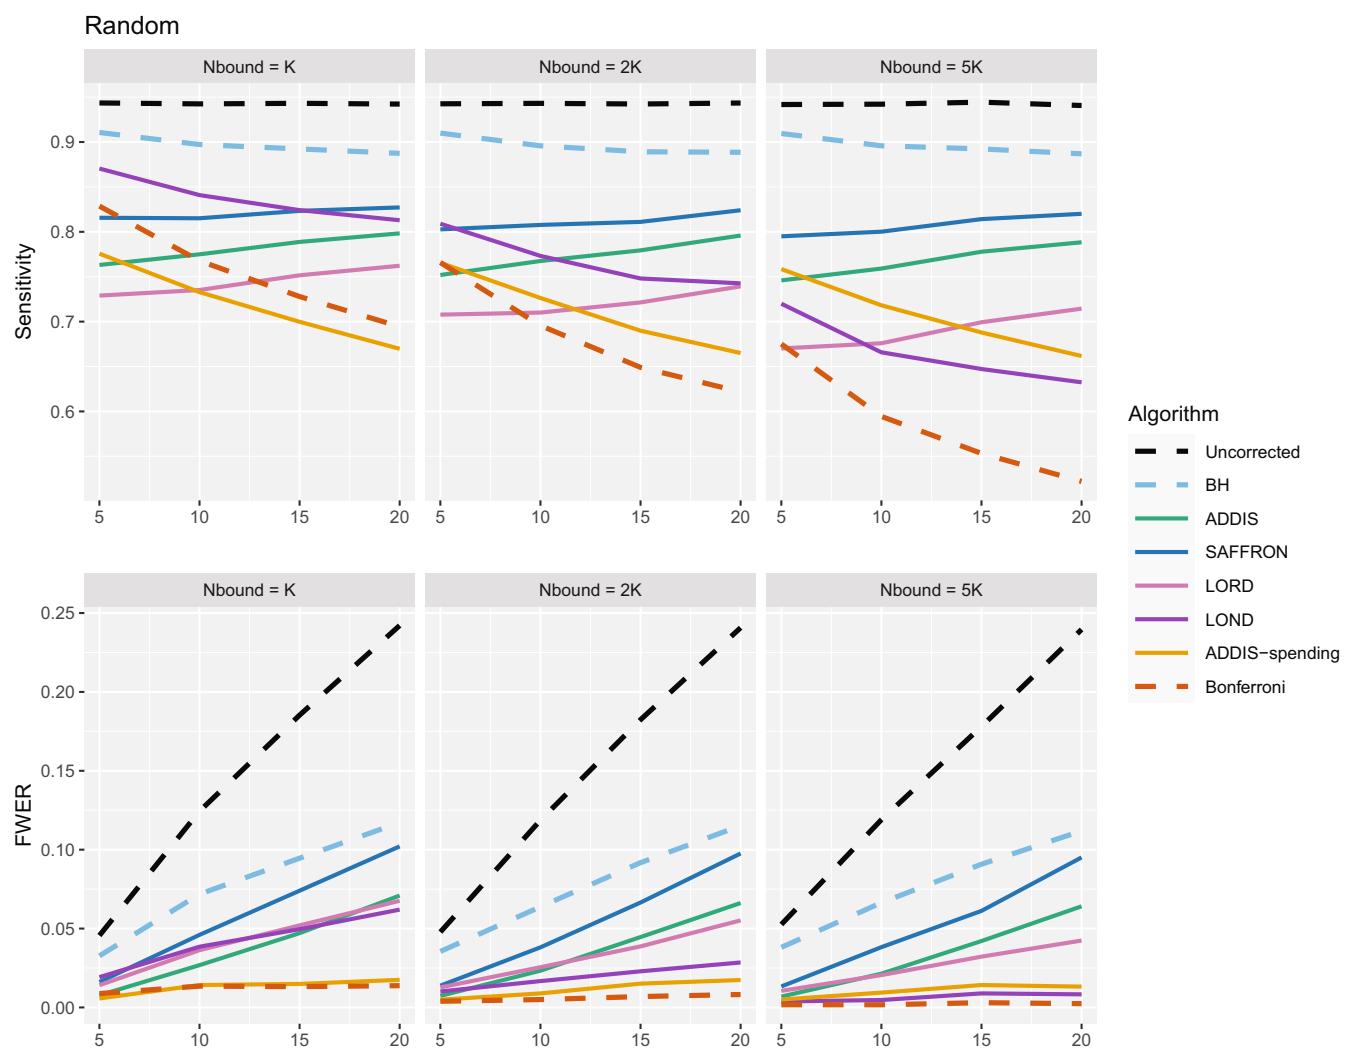


FIGURE 5 Sensitivity with corresponding FWER for random fixed means and $(2K/5 + 1)$ effective treatments, with varying N_{bound} .

In summary, we have seen that the performance of the fully online testing algorithms strongly depends on the number of effective treatments, the order in which they are tested in the trial, and the assumed upper bound N_{bound} on the total number of treatment arms. As would be expected, the power of the online algorithms increases when there is a relatively large number of effective treatments which appear early on in the trial. However, as is the case for offline FDR controlling procedures, the online FDR algorithms do not control the FWER, although the FWER is still noticeably lower than with uncorrected testing.

Overall, LOND is a robust choice across the scenarios considered, since it provably always make at least as many rejections as Bonferroni, which can sometimes translate into substantial sensitivity gains, while maintaining a reasonable FWER for all but the most extreme cases. For example, under (arguably) the most plausible setting of a random ordering of the truly effective treatment, when $N_{\text{bound}} = 2K$ even when $N_{\text{bound}} = 2K$ the FWER on LOND is below the nominal 2.5%. Even when $N_{\text{bound}} = K$, the FWER is still below 5% for $K \leq 15$. In the staircase scenario though, ADDIS-spending performs surprisingly well in terms of disjunctive power, with a noticeably higher power than LOND and Bonferroni in the random and descending scenarios (while still maintaining FWER control). We return to our general recommendations in the discussion (Section 6).

4.2 | Online batched algorithms

We now turn our attention to the batched setting (with a batch size of 5) in order to assess how the online batched algorithms (BatchBH, BatchPRDS and BatchStBH) compare with the fully online algorithms. Starting with the FDR,

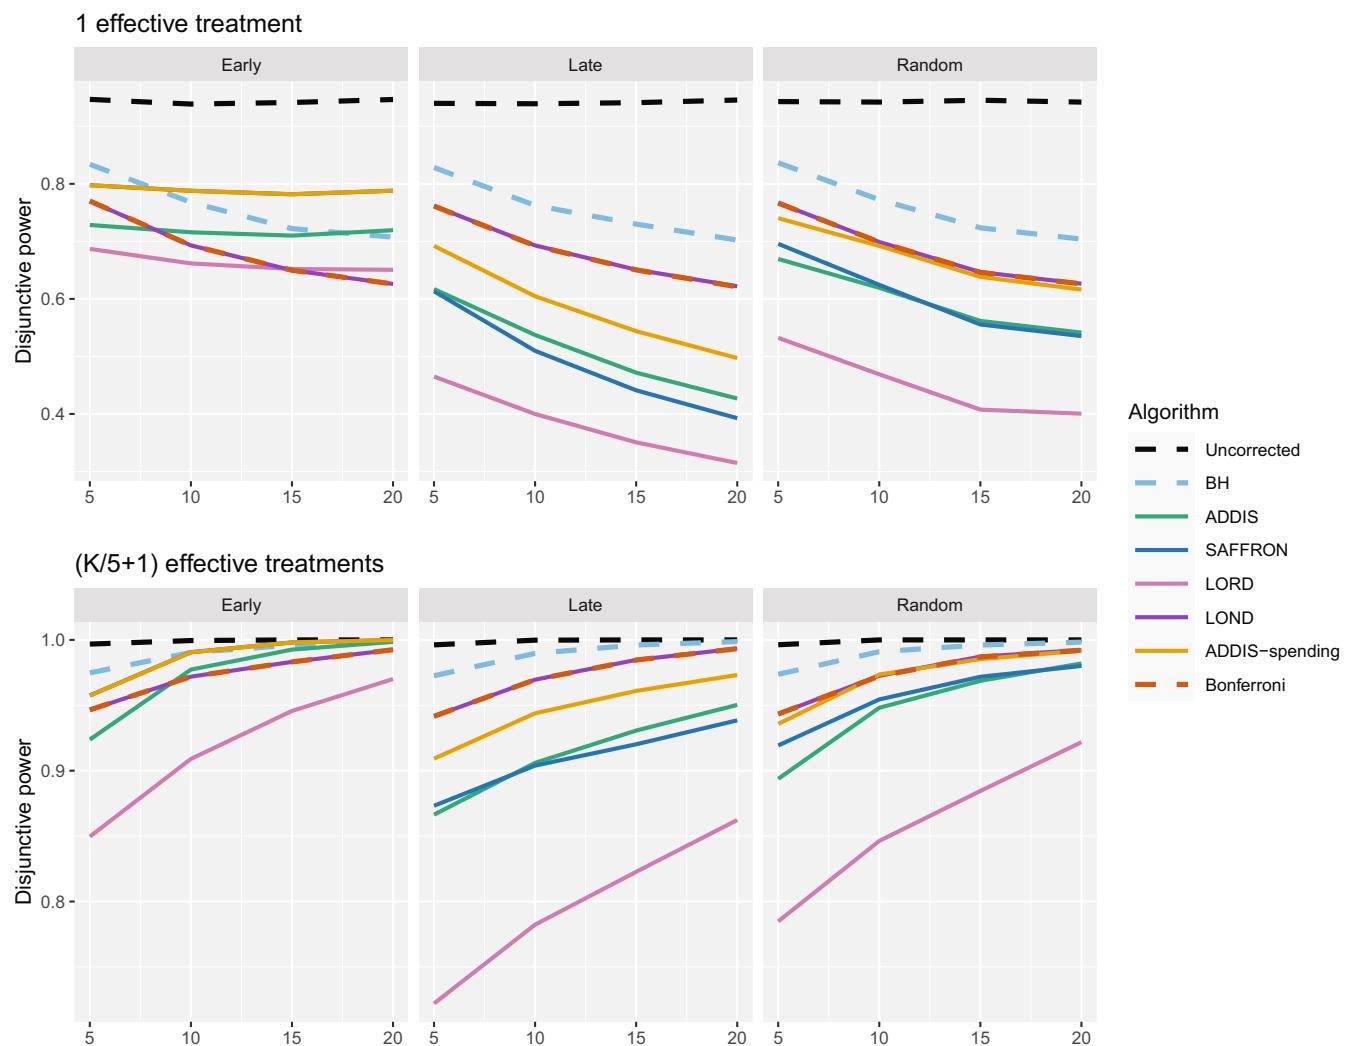


FIGURE 6 Disjunctive power for fixed means and different numbers of effective treatments, with $N_{\text{bound}} = 2K$.

Figure 8 shows the worst and best case FDR (corresponding to the early and late scenarios, respectively) of the online batched algorithms when $N_{\text{bound}} = K$. We see that while BatchBH and BatchPRDS control the FDR below the nominal 2.5% in all the scenarios, BatchStBH has an inflated FDR when there is 1 effective treatment which approaches 5% in the worst case when $K = 20$. This inflation of the FDR of BatchStBH is theoretically justified, since BatchStBH only controls the FDR under independence of the P -values both within and across batches.

Figures 9–11 compares the sensitivity and FWER for fully online and online batched algorithms, with $N_{\text{bound}} = 2K$ for the sensitivity and $N_{\text{bound}} = K$ for the FWER. Starting first with 1 effective treatment (Figure 9), there is no advantage in terms of sensitivity when using the online batched algorithms compared to LOND under the late and random scenarios, or ADDIS-spending and SAFFRON in the early scenario. In all scenarios, BatchStBH has a noticeably inflated FWER, particularly under the early scenario. In contrast, BatchBH and BatchPRDS are reasonably close to the nominal 2.5% level.

When there are $(K/5 + 1)$ effective treatments (Figure 10), there is a slight sensitivity advantage in using BatchBH and BatchStBH over LOND in the random scenario for $K > 10$, with similar or lower FWER. In the early scenario, BatchBH and BatchPRDS have similar sensitivity to SAFFRON but substantially lower FWER (of approximately 2.5%). Meanwhile, BatchStBH has a higher sensitivity than SAFFRON for $K > 10$, but at the cost of an even more inflated FWER. Finally, in the late scenario, BatchBH and BatchPRDS have a slightly lower sensitivity than SAFFRON, whereas BatchStBH has a substantially higher sensitivity but at the cost of a noticeable inflation of the FWER.

When there are $(2K/5 + 1)$ effective treatments (Figure 11), BatchPRDS has similar sensitivity to LOND but with a substantially lower FWER that is below the 2.5% level. BatchBH and BatchStBH have similar sensitivity to SAFFRON and

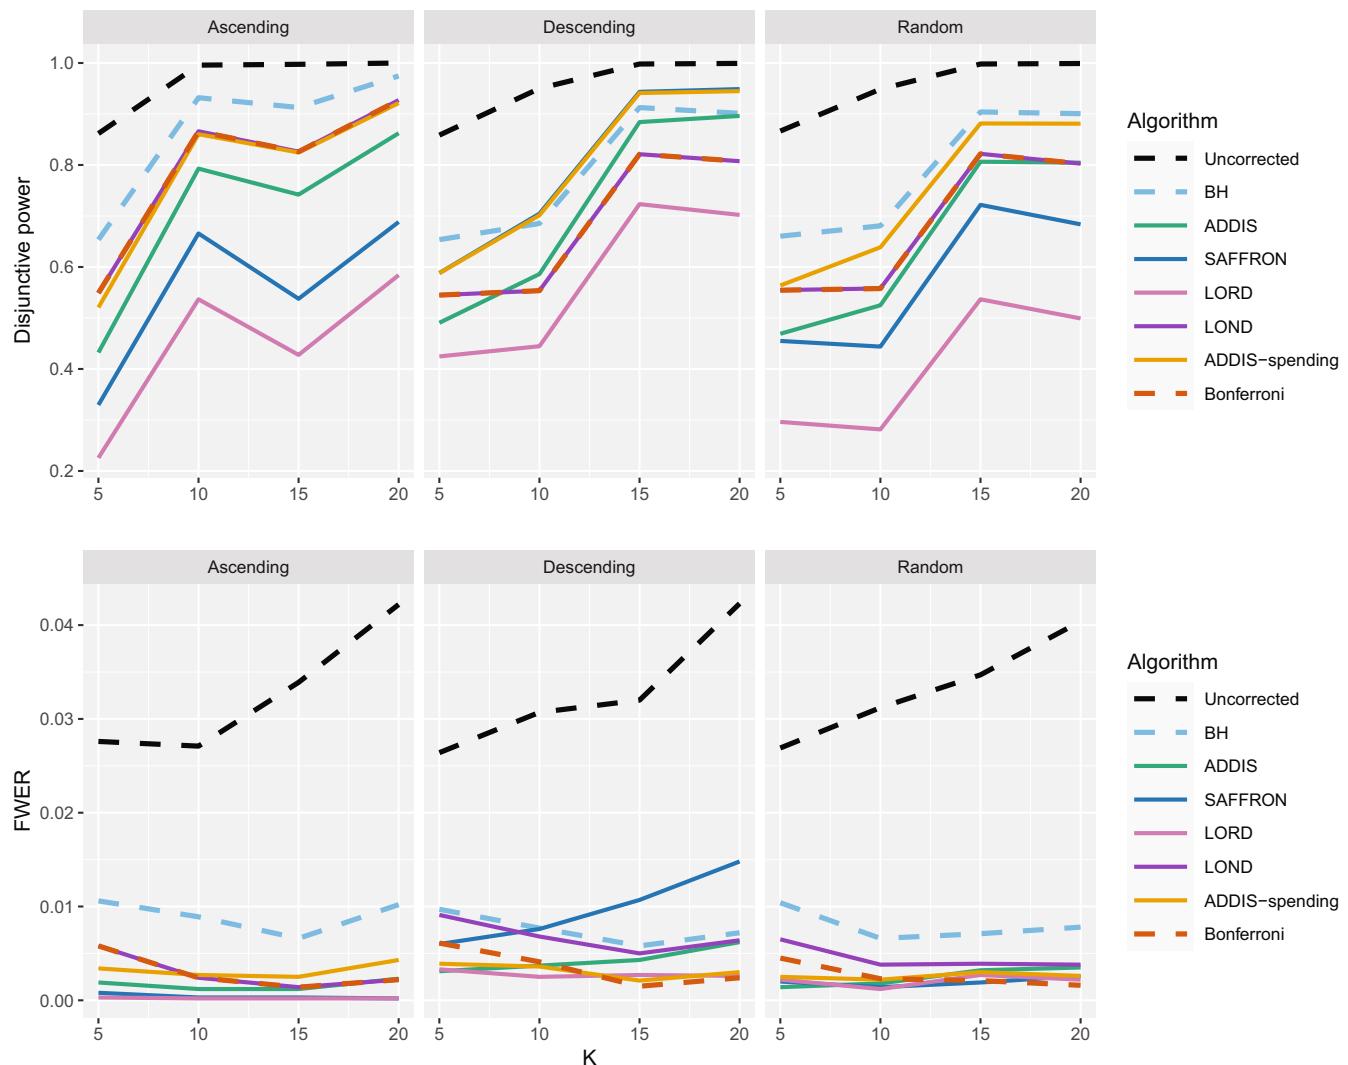


FIGURE 7 Disjunctive power and FWER for the staircase scenarios. Here $N_{\text{bound}} = 2K$ for the disjunctive power plots and $N_{\text{bound}} = K$ for the FWER plots.

ADDIS respectively when $K \geq 10$, but BatchStBH has substantially lower FWER (approximately 2.5%) while BatchBH has similar FWER to ADDIS. In the early scenario, BatchBH and BatchPRDS have a roughly similar sensitivity to SAFFRON, but with a substantially lower FWER (which is at or below 2.5% for BatchPRDS). BatchStBH has a higher sensitivity than SAFFRON (in fact approaching the level of uncorrected testing) but with a similarly highly inflated FWER. Lastly, in the late scenario BatchPRDS has a lower sensitivity than LOND for $K \geq 10$ although still noticeably higher than Bonferroni. BatchBH has a slightly higher sensitivity than LOND with a similar FWER. Meanwhile BatchStBH has a substantially higher sensitivity than LOND (or any of the other fully online algorithms) but again at the cost of a noticeably higher FWER.

In terms of disjunctive power, Figure B9 in Section B.3 of the Supporting web materials shows that all of the online batched algorithms have fairly similar power to ADDIS-spending across the different scenarios. However, as already seen this comes at the cost of a (potentially highly) inflated FWER for BatchStBH and to a lesser extent BatchBH.

Finally, Figure 12 gives the results for the staircase scenario with conservative nulls, in terms of disjunctive power and FWER. We again see that the online batched algorithms tend to have similar disjunctive power to ADDIS-spending, although BatchStBH has slightly higher values in the ascending and descending scenarios and all the online batched algorithms have slightly lower values in the random scenarios. As for the FWER, all the online batched algorithms control the FWER below the nominal 2.5%, except for BatchStBH under the descending scenario.

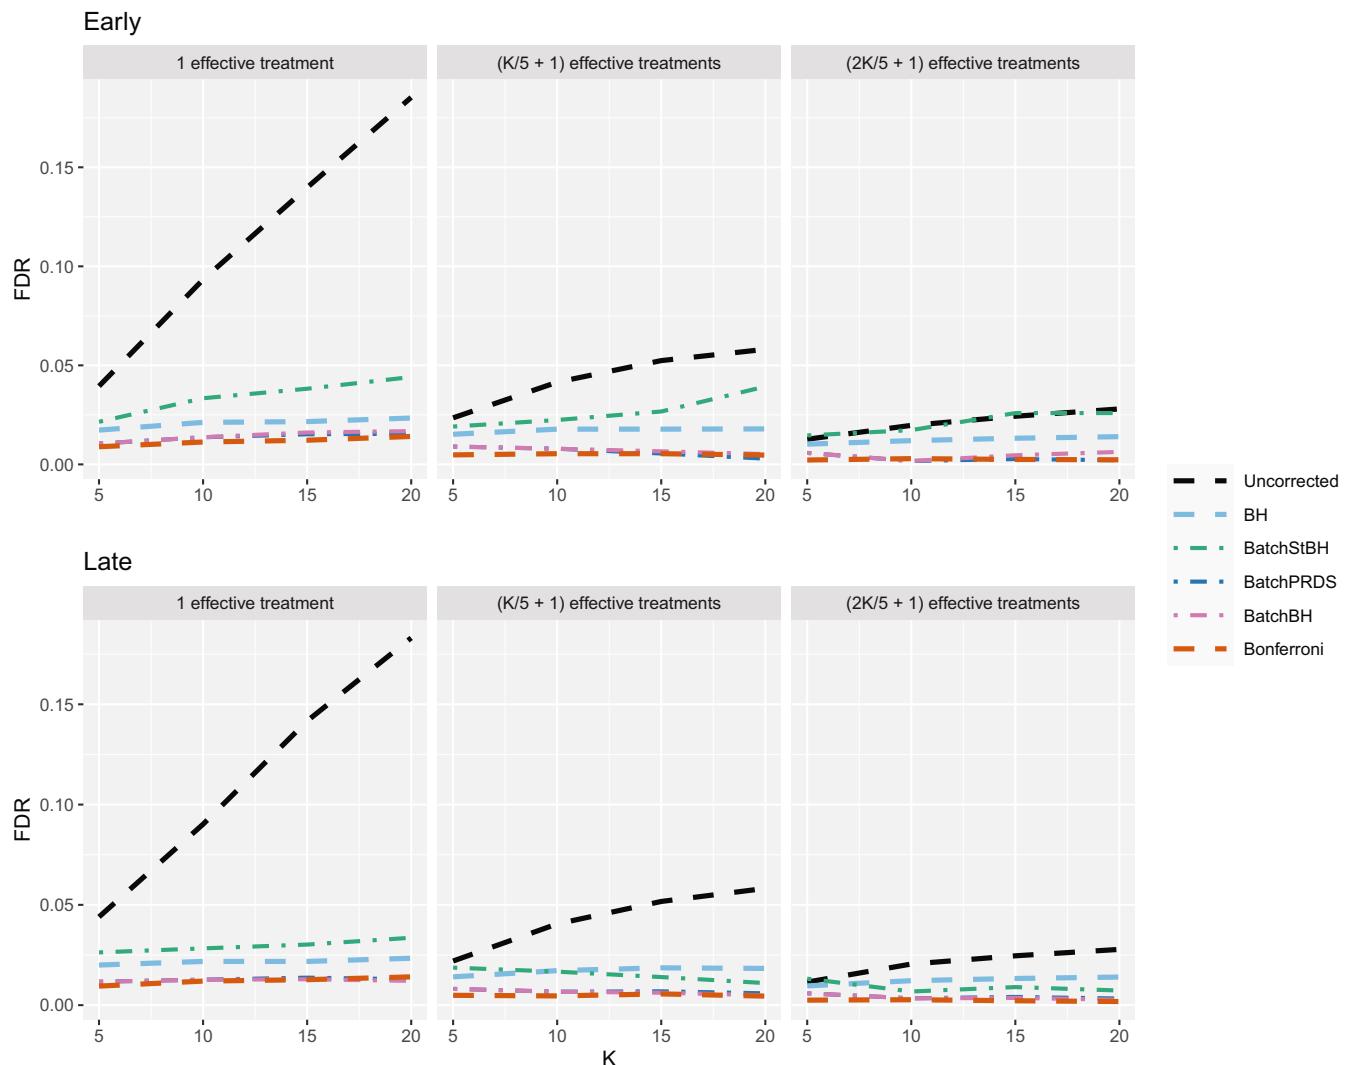


FIGURE 8 FDR for the online batched algorithms, with $N_{\text{bound}} = K$.

Looking at the results for online batched testing as a whole, we see that the BatchPRDS algorithm is competitive when compared to LOND (and hence also Bonferroni) by having similar or higher power, in terms of sensitivity and disjunctive power, while still controlling the FWER at 2.5% across all the scenarios considered.

Finally, an alternative strategy to try to control the FDR (or FWER) in the batch setting is to apply the BH or Holm procedure to each batch. However, a repeated application of the BH (Holm) procedure each at level α will not in general control the FDR (FWER) at level α across batches. For example, in the extreme case of all batches being of size 1, then such a strategy using the Holm procedure would be identical to using uncorrected testing at level α . Nonetheless, we have evaluated the use of a repeated BH and a repeated Holm procedure for the batch setting, with the results found in Section B.3 of the Supporting information. There can be a highly inflated FDR and FWER for this strategy but also a substantial increase in sensitivity in some scenarios.

5 | CASE STUDY: STAMPEDE TRIAL

The STAMPEDE (Systemic Therapy for Advancing or Metastatic Prostate Cancer) trial is a flagship platform trial for research into the effect of systemic therapies for prostate cancer²⁸ on overall survival. The trial opened to accrual in October 2005 with 6 arms (A-F), including the control arm A, which was standard-of-care (SOC) hormone therapy. Figure 13 shows a schematic of the treatment comparisons that have already been reported from STAMPEDE. Two additional arms

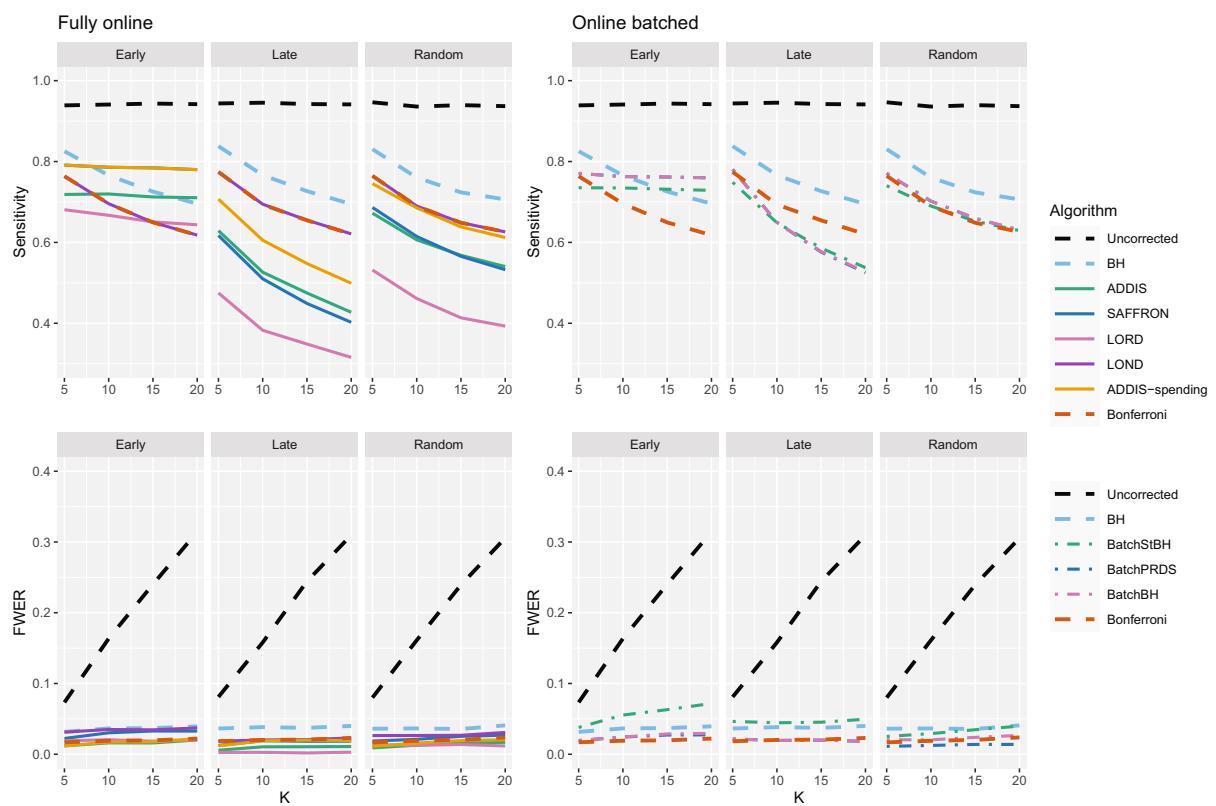


FIGURE 9 Comparison of sensitivity and FWER for fully online and online batched algorithms, with 1 effective treatment. Here $N_{\text{bound}} = 2K$ for the sensitivity and $N_{\text{bound}} = K$ for the FWER.

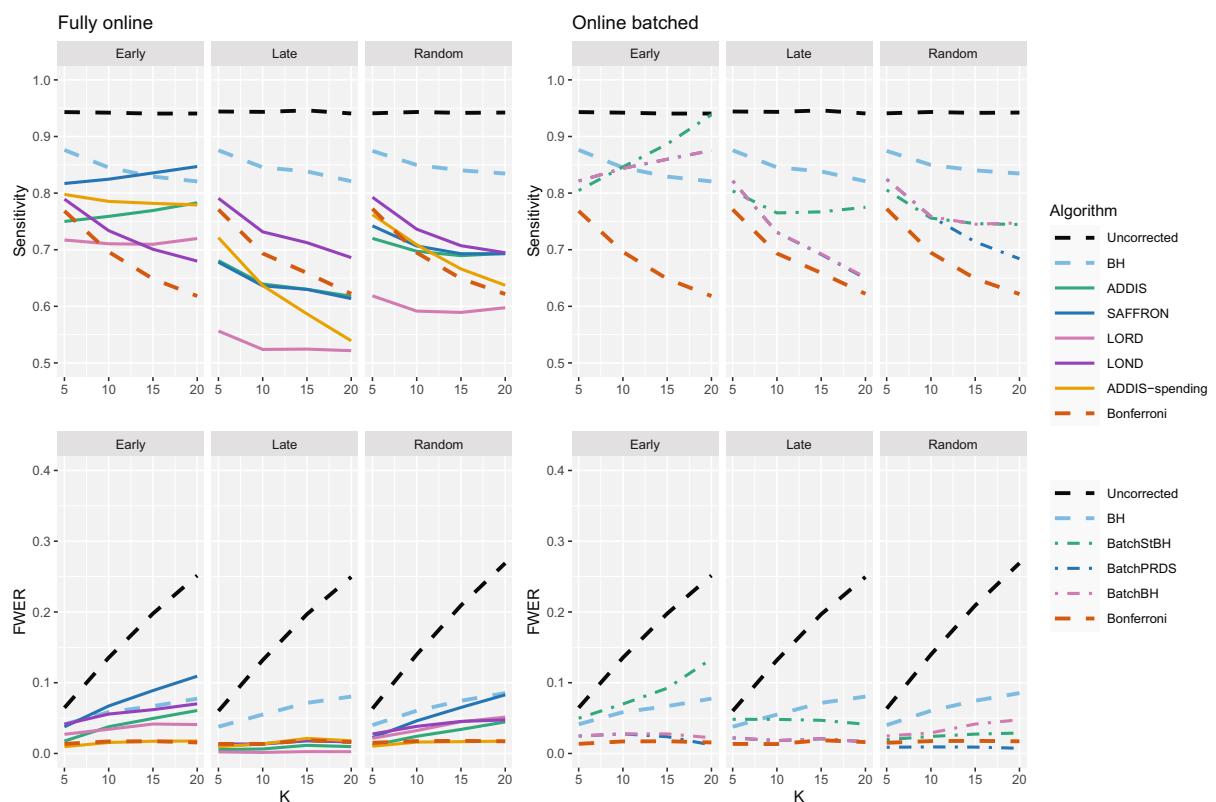


FIGURE 10 Comparison of sensitivity for fully online and online batched algorithms, with $(K/5 + 1)$ effective treatments. Here $N_{\text{bound}} = 2K$ for the sensitivity and $N_{\text{bound}} = K$ for the FWER.

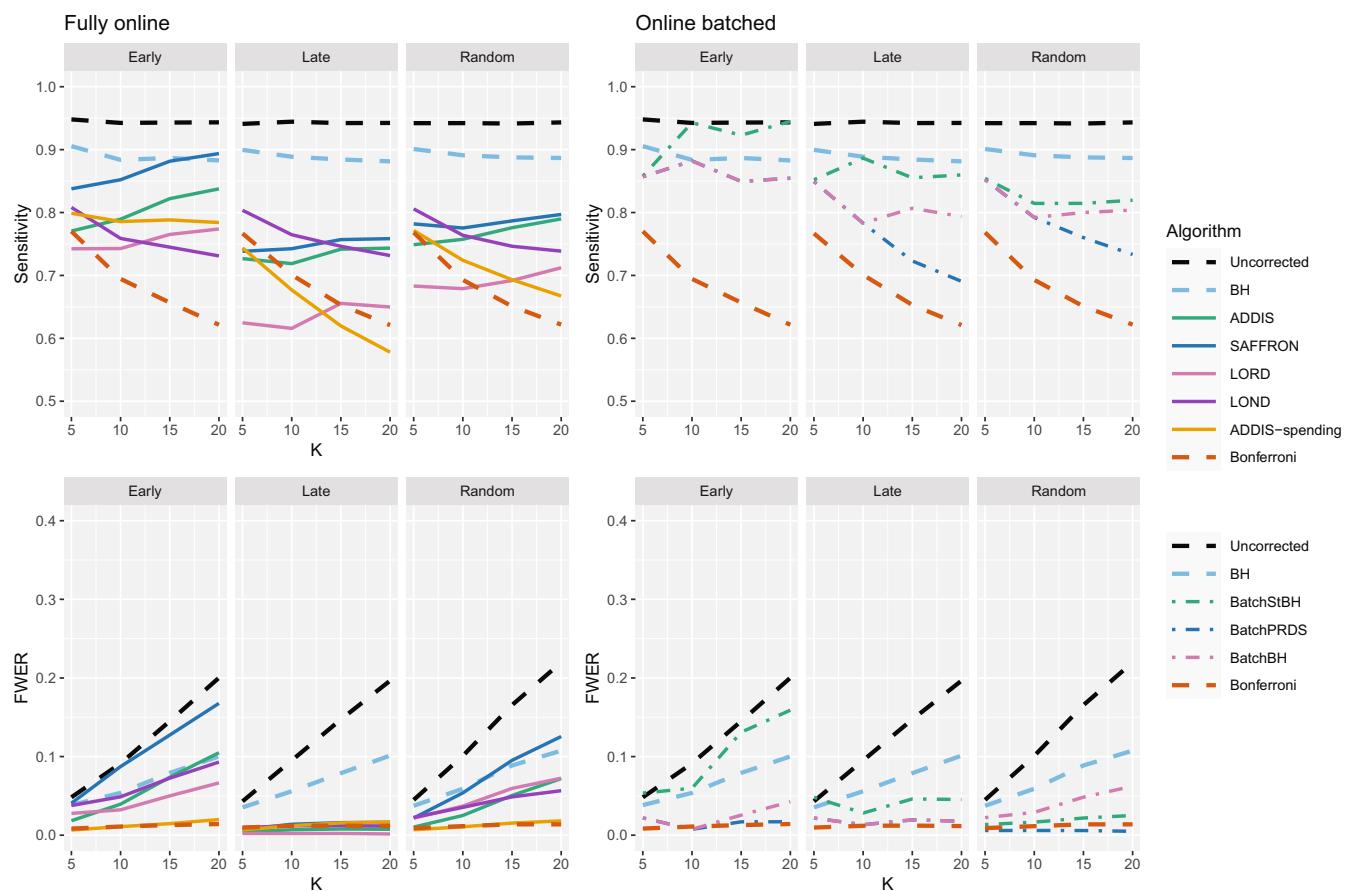


FIGURE 11 Comparison of sensitivity for fully online and online batched algorithms, with $(2K/5 + 1)$ effective treatments. Here $N_{\text{bound}} = 2K$ for the sensitivity and $N_{\text{bound}} = K$ for the FWER.

(G and H) were added to the trial in 2011 and 2013, respectively. Three arms (B, C, and E) reported main analyses in 2015,³¹ with two additional arms (D and F) reporting in 2015/2016.³² Arm G reported main analyses in 2017³³ and Arm H reported results in 2018.³⁴

Table 2 shows the reported P -values (unadjusted for multiplicity) when comparing arms B to H with the SOC (arm A), as given in References 31-34. The dashed lines denote the 4 batches in the trial.

We now apply the offline and online testing algorithms to the observed P -values given above in Table 2, keeping the alphabetical ordering of P -values within the three batches. As a sensitivity analysis, we show what happens if the ordering of the P -values in the first batch changes in Table C2 in Section C of the Supporting web materials. We set the upper bound on the number of treatments $N_{\text{bound}} = 20$, that is, twice as many arms that have already entered the STAMPEDE trial as of the end of 2021.

Table 3 shows which of the hypotheses corresponding to each trial arm can be rejected at level $\alpha \in \{0.025, 0.05, 0.1\}$, as well as the current significance level α_8 that would be used to test the next treatment arm after the 7 already evaluated in the trial. We consider larger values of α since this has been a suggestion for multi-arm trials³⁵ as a compromise between not correcting for multiplicity at all (as would be the case when running a series completely independent two-arm trials) and strict FWER control at 2.5%.

When $\alpha = 0.025$, uncorrected testing rejects hypotheses C, E, and G, but only BH and BatchStBH reject more than one hypothesis (C and G). All the other online algorithms as well as Bonferroni only reject hypothesis G, except for ADDIS and LORD which reject no hypotheses. When $\alpha = 0.05$ the rejections made by the various algorithms remain the same, except now SAFFRON, BatchBH and BatchPRDS reject hypotheses C and G, while BatchStBH rejects hypotheses C, E, and G. Finally, when $\alpha = 0.1$ we see that BH, SAFFRON, BatchBH, BatchPRDS, and BatchStBH all make the same rejections as uncorrected testing. In terms of the current adjusted testing level α_8 , we see that ADDIS-spending, ADDIS, and LORD have a smaller value than even Bonferroni. SAFFRON, LOND, BatchBH,

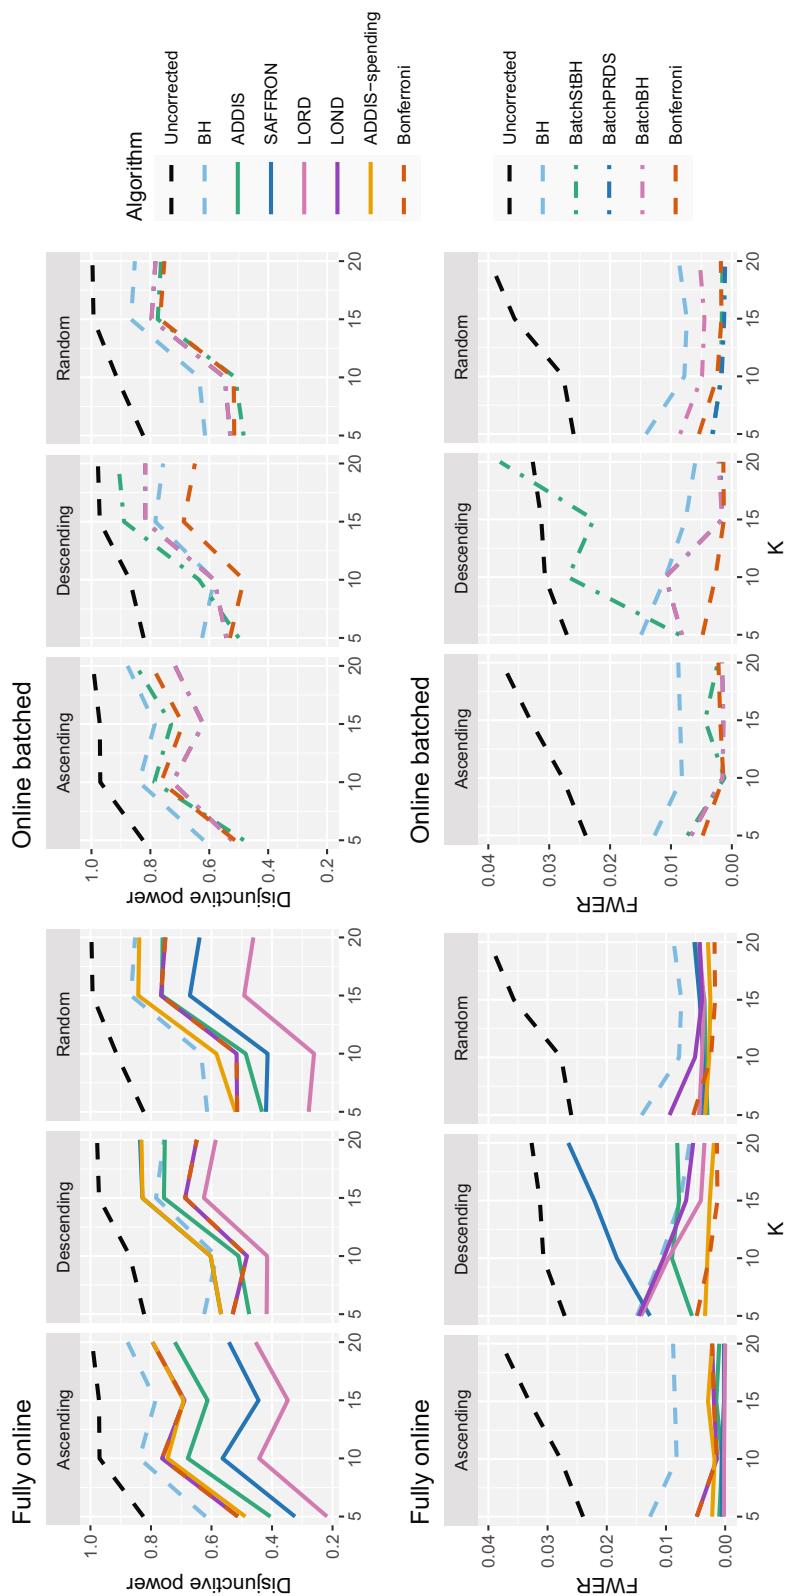


FIGURE 12 Comparison of the FWER ($N_{\text{bound}} = K$) and disjunctive power ($N_{\text{bound}} = 2K$) for fully online and online batched algorithms under the staircase scenario.

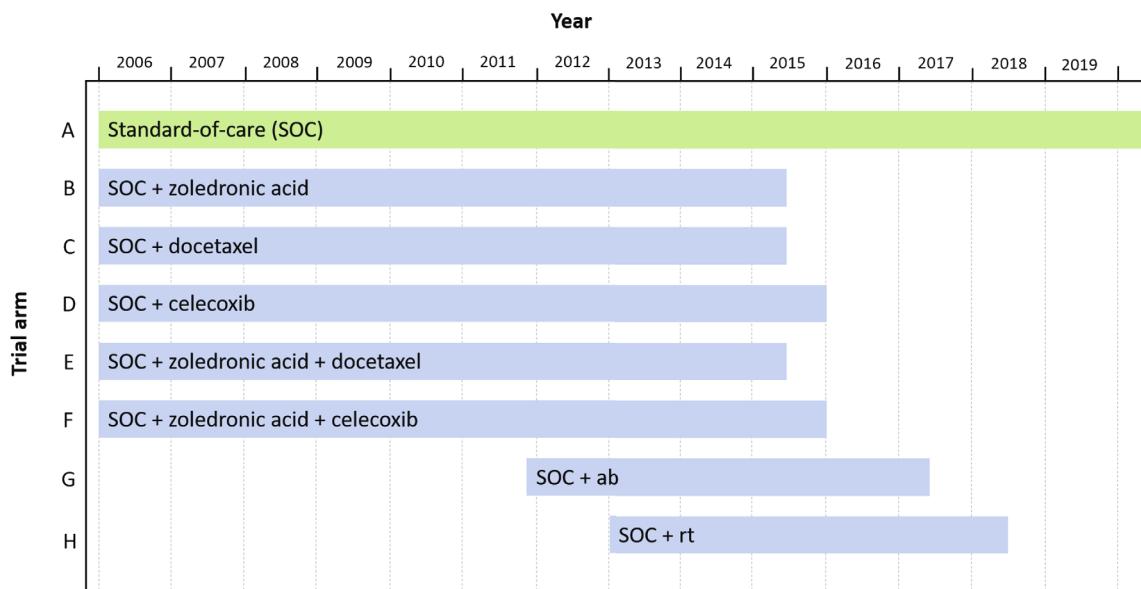


FIGURE 13 Schematic of the STAMPEDE trial. ab, abiraterone; rt, radiotherapy.

TABLE 2 Reported results for the STAMPEDE trial.

Trial arm	P-value
B: SOC + zoledronic acid	0.450
C: SOC + docetaxel	0.006
E: SOC + zoledronic acid + docetaxel	0.022
D: SOC + celecoxib	0.847
F: SOC + zoledronic acid + celecoxib	0.130
G: SOC + abiraterone	0.001
H: SOC + radiotherapy	0.266

Abbreviation: SOC, standard-of-care.

TABLE 3 Rejections and current significance level α_8 of different algorithms using the results of the STAMPEDE trial, with the ordering as in Table 2.

Algorithm	Hypotheses rejected			α_8		
	$\alpha = 0.025$	$\alpha = 0.05$	$\alpha = 0.1$	$\alpha = 0.025$	$\alpha = 0.05$	$\alpha = 0.1$
Uncorrected	C, E, G	C, E, G	C, E, G	0.0250	0.0500	0.1000
Bonferroni	G	G	G	0.0013	0.0025	0.0050
ADDIS-spending	G	G	G	0.0005	0.0011	0.0021
BH	C, G	C, G	C, E, G	–	–	–
ADDIS	–	G	G	0.0003	0.0016	0.0031
SAFFRON	G	C, G	C, E, G	0.0041	0.0165	0.0412
LORD	–	–	–	0.0001	0.0002	0.0003
LOND	G	G	G	0.0025	0.0050	0.0100
BatchBH	G	C, G	C, E, G	0.0019	0.0057	0.0151
BatchPRDS	G	C, G	C, E, G	0.0019	0.0057	0.0151
BatchStBH	C, G	C, E, G	C, E, G	0.0381	0.1015	0.1238

and BatchPRDS have a larger value of α_8 than Bonferroni, but only BatchStBH has a higher value than uncorrected testing.

6 | DISCUSSION

In this article, we have shown how online multiple hypothesis testing can be applied to the platform trial setting to achieve overall control of type I errors. In many of the simulation scenarios, there were noticeable gains in sensitivity compared with using Bonferroni, although this has to be considered carefully with respect to the potential inflation of the FWER and how acceptable that may be to stakeholders (particularly from a regulatory viewpoint). However, in all cases the FWER was lower than uncorrected testing, and so from that perspective, any of the online algorithms would offer an improvement.

We have focused on trial settings that ultimately test $N \leq 20$ hypotheses, which is a relatively small number of hypotheses compared with most previous simulation results in the online testing literature. Our results show that the online testing framework offers a smaller advantage in terms of power properties in this small N setting compared to when a very large number (≥ 1000) of hypotheses are tested. However, in some clinical trial settings (eg, testing for interaction with genomic data) there can be a large number of hypotheses and a desire for FDR control (as opposed to FWER control), and so online error rate control may be especially applicable.

As noted in Section 4, the relative performance of the online algorithms compared with Bonferroni and uncorrected testing crucially depends on the number of effective treatments, the order in which they are tested in the trial, as well as the assumed upper bound N_{bound} . Hence, the “best” online algorithm to use may be quite different depending on the trial context and goals. In general, across the simulation scenarios for the fully online setting, we see that LOND seems to offer quite a good compromise by guaranteeing at least as many rejections as Bonferroni and often seeing noticeable increases in sensitivity, while still maintaining (depending on the trial context) reasonable FWER levels. In the online batched setting however, BatchPRDS seems to be preferable to LOND as it has a similar or higher sensitivity and disjunctive power, while maintaining FWER control below the nominal level.

As seen in the case study results, in practice the online testing algorithms may not be “competitive” in terms of the number of rejections when compared with uncorrected testing unless a more relaxed α is used. However, the potential inflation on top of that relaxation for many of the online algorithms also needs to be taken into account. In that sense, BatchPRDS is particularly appealing due to the lack of further FWER inflation.

We have started with the simplifying assumption that each experimental treatment is tested exactly once, as soon as the outcomes from a pre-specified number of patients have been observed. A useful extension would be to allow multiple looks (stages) for each treatment arm, with early stopping for futility and efficacy. Indeed, the STAMPEDE trial included interim analyses that we did not take into account in our analysis in Section 5. A complicating factor is that to directly apply the online testing algorithms, we have only assigned the adjusted testing level α_i to hypothesis H_i at the point of testing. To allow repeated testing of hypothesis H_i and early stopping, we would require α_i to be determined in advance. For further discussion and proposals for this setting, we refer the interested reader to Zrnic et al,²³ who proposed a framework for asynchronous online testing. As well, recent work by Zehetmayer et al³⁶ shows how to use LOND specifically for group-sequential platform trials.

Another assumption that we have made is that $N_{\text{bound}} \geq K$, that is, that we never test more treatment arms than originally planned for. In the case where $N_{\text{bound}} < K$, then a simple method of accommodating this change has been proposed by Robertson and Wason,^{17(p. 10)} but the power properties have not been explored in the context of platform trials. We could also make further comparisons of online testing with existing methodologies for controlling the FWER when adding treatment arms to a platform trial.^{37,38} Finally, one limitation of our investigation is that in practice, future comparisons depend on results of past trials, which in turn depend on which online testing algorithm is being used. We did not account for this in our simulations. Taking this into consideration would be highly complex, but may be a useful avenue for future research.

Alternative type I error rates to the FWER and FDR have been proposed in the literature. One example is the Error Rate per Family (ERpF) proposed by Tukey.³⁹ The ERpF is defined as the number of type I errors divided by the number of families of hypotheses. As pointed out by an anonymous reviewer, while the FWER may be appropriate for multiplicity within a study, the ERpF is meant for multiplicity control across studies. Hence, if one can view a platform trial as a series of distinct studies (eg, corresponding to a change to the standard-of-care), the ERpF may be a more appropriate metric to control. The ERpF can be controlled exactly across studies by using what is called an additive multiplicity correction.⁴⁰

As future research, it would be interesting to compare online control of the FWER or FDR with ERpF control. It would also be interesting to explore whether recently proposed methods of FDR control using e-values⁴¹ could be applied to the platform trial setting.

As mentioned in the introduction, the need for multiplicity adjustments for platform trials is controversial, and there are some platform trial contexts where applying such adjustments may be difficult and/or undesirable. For example, platform trials often involve multiple pharmaceutical sponsors and a natural question to ask is why a sponsor would participate in a platform trial that uses *any* multiplicity adjustment, given that they can conduct their own two-arm trial with *no* multiplicity adjustment. To help address this issue, one partial solution is to adjust for multiplicity (using online testing or otherwise) at a level α that is higher than the usual 2.5% (one-sided). As shown in the case study in Section 5, this can result in adjusted testing levels above 2.5%. More generally, the efficiencies that a platform trial bring (eg, in terms of time and cost of recruiting control patients) may still mean that a sponsor would be willing to join a trial even if there is a price to be paid in terms of multiplicity adjustment. Meanwhile, in the context of emerging infectious diseases, there is a real need to quickly find a treatment that works and it is unclear how the use of multiplicity adjustment may affect that time. A fuller discussion around multiplicity adjustment (and online testing more specifically) for platform trials would be useful, but is out of scope of this article.

ACKNOWLEDGEMENTS

This research was supported by the NIHR Cambridge Biomedical Research Centre (BRC1215-20014). 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). David S. Robertson and Thomas Jaki received funding from the UK Medical Research Council (MC_UU_00002/14). David S. Robertson also received funding from the Biometrika Trust. For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) licence to any author accepted manuscript version arising. Franz König and Martin Posch were supported by EU-PEARL. EU-PEARL (EU Patient-cEtric clinicAl tRial pLatforms) project has received funding from the Innovative Medicines Initiative (IMI) 2 Joint Undertaking (JU) under Grant agreement no. 853966. This joint undertaking receives support from the European Union's Horizon 2020 research and innovation programme and EFPIA and Children's Tumor Foundation, Global Alliance for TB Drug Development non-profit organization, Springworks Therapeutics Inc. This publication reflects the authors' views. Neither IMI nor the European Union, EFPIA, or any Associated Partners are responsible for any use that may be made of the information contained herein.

DATA AVAILABILITY STATEMENT

All of the data that support the findings of this study are available within the article itself.

ORCID

David S. Robertson  <https://orcid.org/0000-0001-6207-0416>
James M. S. Wason  <https://orcid.org/0000-0002-4691-126X>
Franz König  <https://orcid.org/0000-0002-6893-3304>
Martin Posch  <https://orcid.org/0000-0001-8499-8573>
Thomas Jaki  <https://orcid.org/0000-0002-1096-188X>

REFERENCES

1. DiMasi JA, Grabowski HG, Hansen RW. Innovation in the pharmaceutical industry: new estimates of R&D costs. *J Health Econ.* 2016;47:20-33.
2. Bentley C, Cressman S, van der Hoek K, Arts K, Dancey J, Peacock S. Conducting clinical trials—costs, impacts, and the value of clinical trials networks: a scoping review. *Clin Trials.* 2019;16(2):183-193.
3. Saville BR, Berry SM. Efficiencies of platform clinical trials: a vision of the future. *Clin Trials.* 2016;13(3):358-366.
4. Park JJ, Harari O, Dron L, Lester RT, Thorlund K, Mills EJ. An overview of platform trials with a checklist for clinical readers. *J Clin Epidemiol.* 2020;125:1-8.
5. Meyer EL, Mesenbrink P, Dunger-Baldauf C, et al. The evolution of master protocol clinical trial designs: a systematic literature review. *Clin Ther.* 2020;42(7):1330-1360.
6. Parmar MK, Carpenter J, Sydes MR. More multiarm randomised trials of superiority are needed. *Lancet.* 2014;384(9940):283-284.

7. Pallmann P, Bedding AW, Choodari-Oskooei B, et al. Adaptive designs in clinical trials: why use them, and how to run and report them. *BMC Med.* 2018;16(1):1-15.
8. Hague D, Townsend S, Masters L, et al. Changing platforms without stopping the train: experiences of data management and data management systems when adapting platform protocols by adding and closing comparisons. *Trials.* 2019;20(1):1-16.
9. Lee KM, Brown LC, Jaki T, Stallard N, Wason J. Statistical consideration when adding new arms to ongoing clinical trials: the potentials and the caveats. *Trials.* 2021;22(1):1-10.
10. Wason JM, Stecher L, Mander AP. Correcting for multiple-testing in multi-arm trials: is it necessary and is it done? *Trials.* 2014; 15(1):1-7.
11. Howard DR, Brown JM, Todd S, Gregory WM. Recommendations on multiple testing adjustment in multi-arm trials with a shared control group. *Stat Methods Med Res.* 2018;27(5):1513-1530.
12. Collignon O, Gartner C, Haidich AB, et al. Current statistical considerations and regulatory perspectives on the planning of confirmatory basket, umbrella, and platform trials. *Clin Pharmacol Ther.* 2020;107(5):1059-1067.
13. Collignon O, Burman CF, Posch M, Schiel A. Collaborative platform trials to fight COVID-19: methodological and regulatory considerations for a better societal outcome. *Clin Pharmacol Ther.* 2021;110:311-320.
14. Wason JM, Robertson DS. Controlling type I error rates in multi-arm clinical trials: a case for the false discovery rate. *Pharm Stat.* 2021;20(1):109-116.
15. Parker RA, Weir CJ. Non-adjustment for multiple testing in multi-arm trials of distinct treatments: rationale and justification. *Clin Trials.* 2020;17(5):562-566.
16. Bretz F, Koenig F. Commentary on Parker and Weir. *Clin Trials.* 2020;17(5):567-569.
17. Robertson DS, Wason JMS. Online control of the false discovery rate in biomedical research. arXiv preprint arXiv:1809.07292 [stat.ME], September 19, 2018.
18. Robertson DS, Wason J, Ramdas A. Online multiple hypothesis testing for reproducible research. arXiv preprint arXiv:220811418, 2022.
19. Javanmard A, Montanari A. Online rules for control of false discovery rate and false discovery exceedance. *Ann Stat.* 2018;46(2): 526-554.
20. Ramdas A, Yang F, Wainwright MJ, Jordan MI. Online control of the false discovery rate with decaying memory. Proceedings of the Advances in Neural Information Processing Systems; 2017:5650-5659. <http://papers.nips.cc/paper/7148-online-control-of-the-false-discovery-rate-with-decaying-memory.pdf>
21. Ramdas A, Zrnic T, Wainwright M, Jordan M. SAFFRON: an adaptive algorithm for online control of the false discovery rate. Proceedings of the 35th International Conference on Machine Learning; 2018:4286-4294. <http://proceedings.mlr.press/v80/ramdas18a.html>
22. Tian J, Ramdas A. ADDIS: an adaptive discarding algorithm for online FDR control with conservative nulls. Proceedings of the Advances in Neural Information Processing Systems; 2019:32.
23. Zrnic T, Ramdas A, Jordan MI. Asynchronous online testing of multiple hypotheses. *J Mach Learn Res.* 2021;22:1-39.
24. Zrnic T, Jiang D, Ramdas A, Jordan M. The power of batching in multiple hypothesis testing. Proceedings of the International Conference on Artificial Intelligence and Statistics PMLR; 2020:3806-3815.
25. Tian J, Ramdas A. Online control of the familywise error rate. *Stat Methods Med Res.* 2021;30(4):976-993.
26. European Medicines Agency. Guideline on multiplicity issues in clinical trials; 2016. https://www.ema.europa.eu/en/documents/scientific-guideline/draft-guideline-multiplicity-issues-clinical-trials_en.pdf
27. U.S. Food and Drug Administration. Multiple endpoints in clinical trials guidance for industry; 2017. <https://www.fda.gov/media/102657/download>
28. James ND, Sydes MR, Clarke NW, et al. STAMPEDE: systemic therapy for advancing or metastatic prostate cancer—a multi-arm multi-stage randomised controlled trial. *Clin Oncol.* 2008;20(8):577-581.
29. Robertson DS, Wildenhain J, Javanmard A, Karp NA. onlineFDR: an R package to control the false discovery rate for growing data repositories. *Bioinformatics.* 2019;35(20):4196-4199.
30. Magirr D, Jaki T, Whitehead J. A generalized Dunnett test for multi-arm multi-stage clinical studies with treatment selection. *Biometrika.* 2012;99(2):494-501.
31. James ND, Sydes MR, Clarke NW, et al. Addition of docetaxel, zoledronic acid, or both to first-line long-term hormone therapy in prostate cancer (STAMPEDE): survival results from an adaptive, multiarm, multistage, platform randomised controlled trial. *Lancet.* 2016;387(10024):1163-1177.
32. Mason MD, Clarke NW, James ND, et al. Adding celecoxib with or without Zoledronic acid for hormone-Naïve prostate cancer: long-term survival results from an adaptive, multiarm, multistage, platform, randomized controlled trial. *J Clin Oncol.* 2017;35(14): 1530-1541.
33. James ND, de Bono JS, Spears MR, et al. Abiraterone for prostate cancer not previously treated with hormone therapy. *N Engl J Med.* 2017;377(4):338-351. doi:[10.1056/NEJMoa1702900](https://doi.org/10.1056/NEJMoa1702900)
34. Parker CC, James ND, Brawley CD, et al. Radiotherapy to the primary tumour for newly diagnosed, metastatic prostate cancer (STAMPEDE): a randomised controlled phase 3 trial. *Lancet.* 2018;392(10162):2353-2366.
35. Wason J, Magirr D, Law M, Jaki T. Some recommendations for multi-arm multi-stage trials. *Stat Methods Med Res.* 2016;25(2): 716-727.
36. Zehetmayer S, Posch M, Koenig F. Online control of the false discovery rate in group-sequential platform trials. *Stat Methods Med Res.* 2022;31(12):2470-2485.

37. Burnett T, König F, Jaki T. Adding experimental treatment arms to multi-arm multi-stage platform trials in progress. arXiv preprint arXiv: 2007.04951 [stat.AP], July 9, 2020.
38. Choodari-Oskooei B, Bratton DJ, Gannon MR, Meade AM, Sydes MR, Parmar MK. Adding new experimental arms to randomised clinical trials: impact on error rates. *Clin Trials*. 2020;17:273-284.
39. Tukey JW. The problem of multiple comparisons [Unpublished manuscript]. *The Collected Works of John W Tukey VIII. Multiple Comparisons: 1948-1983*. New York, NY: Chapman and Hall; 1994:1-300.
40. Ding Y, Li YG, Liu Y, Ruberg SJ, Hsu JC. Confident inference for SNP effects on treatment efficacy. *Ann Appl Stat*. 2018;12(3): 1727-1748.
41. Wang R, Ramdas A. False discovery rate control with e-values. *J R Stat Soc Series B Stat Methodol*. 2020;84(3):822-852.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Robertson DS, Wason JMS, König F, Posch M, Jaki T. Online error rate control for platform trials. *Statistics in Medicine*. 2023;42(14):2475-2495. doi: 10.1002/sim.9733