

Multiple testing with anytime-valid Monte-Carlo p-values

Lasse Fischer and Aaditya Ramdas
University of Bremen and Carnegie Mellon University
fischer1@uni-bremen.de, aramdas@cmu.edu

April 25, 2024

Abstract

In contemporary problems involving genetic or neuroimaging data, thousands of hypotheses need to be tested. Due to their high power, and finite sample guarantees on type-1 error under weak assumptions, Monte-Carlo permutation tests are often considered as gold standard for these settings. However, the enormous computational effort required for (thousands of) permutation tests is a major burden. Recently, Fischer and Ramdas [12] constructed a permutation test for a single hypothesis in which the permutations are drawn sequentially one-by-one and the testing process can be stopped at any point without inflating the type I error. They showed that the number of permutations can be substantially reduced (under null and alternative) while the power remains similar. We show how their approach can be modified to make it suitable for a broad class of multiple testing procedures. In particular, we discuss its use with the Benjamini-Hochberg procedure and illustrate the application on a large dataset.

Contents

1	Introduction	2
2	Recap: Sequential Monte-Carlo testing by betting	3
3	Multiple testing with anytime-valid permutation tests by α-dependent p-value calibration	6
4	The Benjamini-Hochberg procedure with anytime-valid permutation p-values	8
4.1	Stopping early for rejection	9
4.2	Stopping early for futility with independent p-values	9
4.3	Stopping early for futility with PRDS p-values	10
5	Simulated experiments	11
5.1	Power and average number of permutations	11
5.2	Early reporting of decisions	11
6	Real data analysis	13
7	Conclusion	15
A	Sequential permutation based multiple testing procedures	18
B	Omitted proofs	19
C	Detailed algorithms	20

1 Introduction

In a classical Monte-Carlo permutation test, we observe some test statistic Y_0 , generate B additional test statistics Y_1, \dots, Y_B and calculate the permutation p-value by

$$\mathbf{p}_B^{\text{perm}} = \frac{1 + \sum_{t=1}^B \mathbb{1}\{Y_t \geq Y_0\}}{1 + B}. \quad (1)$$

In order to be able to reject the null hypothesis at level α , we would need to generate at least $B \geq 1/\alpha - 1$ permuted datasets with corresponding test statistics. This can already lead to considerably computational effort when the data generation process is demanding. However, it is even much higher when M hypotheses are tested instead of a single one. First, we need to perform a permutation test for each hypothesis, leading to a minimum number of $M(1/\alpha - 1)$ permutations in total. Second, when multiple testing corrections are performed, each hypothesis is tested at a lower individual level than the overall level α . For example, if the Bonferroni correction is used, each hypothesis is tested at level α/M leading to the requirement of at least $M(M/\alpha - 1) \approx M^2/\alpha$ permutations. Even if we use more powerful multiple testing procedures such as the Benjamini-Hochberg (BH) procedure [5], this lower bound for the total number of permutations usually remains. The problem is that B need to be chosen in a manner that protects against the worst case in which all p-values are large except for one that is then essentially tested at level α/M . In this paper, we consider applying multiple testing procedures to anytime-valid permutation p-values which allows to stop resampling early and reject (or accept) a hypothesis as soon as the corresponding p-value is rejected (or accepted) by the multiple testing procedure. In this way, we can adapt the number of permutations to the number of rejections and potentially save a lot of permutations.

Despite our terminology and notation focusing on *permutation* tests, everything in this paper identically applies to other Monte Carlo methods such as *conditional randomization tests* as well, which are also common in the causal inference literature, and have been applied to genetic settings under the so-called model-X assumption [9, 4, 3].

There are many existing approaches considering sequential Monte-Carlo tests for multiple testing — in particular for the BH procedure [28, 20, 41, 33, 2, 13, 14, 15, 32]. All these works can be categorized into the following three approaches, from which only the first and third one are useful for saving computational resources.

1. The first idea is to use early stopping rules that ensure obtaining the same decisions as applying the multiple testing procedure to the permutation p-value $\mathbf{p}_B^{\text{perm}}$ for a fixed $B \in \mathbb{N}$ with high probability [28, 20, 41].
2. Another line of work provides bounds for the probability of obtaining a different decision than the limiting permutation p-value [13, 14, 15]. Here, the main goal is not to reduce the number of permutations but to minimize the resampling risk [10].
3. As a third approach, Sandve et al. [33], Pounds et al. [32] and Bancroft et al. [2] consider applying multiple testing procedures to the sequential permutation p-value by Besag and Clifford [7].

The problem with the third approach is that the Besag-Clifford p-value only allows to stop early for accepting the null hypothesis because the evidence is weak, and it cannot be stopped early for rejection. For this reason, the previous works either did not stop for rejection [2, 32] or use heuristics that do not provide provable error control [33]. To the best of our knowledge, the only existing method allowing to adapt the number of permutations to the number of rejections while guaranteeing FDR control, is the AMT algorithm by Zhang et al. [41]. They fix the number of permutations $B \in \mathbb{N}$ in advance and consider the generation process of the test statistics as a multi-armed bandit (MAB) problem, and their guarantee falls into the first category above.

We will follow a strategy similar to the third path, however, instead of the Besag-Clifford p-value, we consider the anytime-valid permutation p-values recently constructed by Fischer and Ramdas [12].

These reduce the number of permutations under both the null and the alternative, without compromising power or type I error. With this, our approach allows to adapt the number of permutations to the number of rejections while providing provable error control. If the used multiple testing procedure provides the desired error control under arbitrary dependence of the p-values, our approach works with all anytime-valid permutation p-values and allows to stop fully data-adaptively. If a specific dependence structure is required, some care regarding the stopping time and used anytime-valid permutation p-value needs to be taken. For brevity, we restrict to the BH procedure in this case and prove that FDR control is guaranteed, if the limiting permutation p-values are independent, or if the limiting permutation p-values are PRDS and we use the anytime-valid BC method. In addition, we argue heuristically why we also do not expect the FDR to be inflated, if the limiting permutation p-values are PRDS and we use other strategies than the anytime-valid BC method. Consequently, in contrast to the AMT algorithm [41], our method can be applied with a large variety of p-value based multiple testing procedures, allows to stop for rejection data-adaptively at any time and does not require a prespecified maximum number of permutations. Furthermore, our method outperformed the AMT algorithm significantly in all considered experiments (see Sections 5 and 6).

In Section 2, we begin with a recap of the betting approach for sequential permutation tests by Fischer and Ramdas [12]. Afterward, we extend this methodology to multiple testing by using α -dependent e-to-p calibration (Section 3). The proposed technique is straightforward for all p-value based multiple testing procedures that do not require assumptions on the dependence structure. In Section 4, we discuss what needs to be considered when applying multiple testing procedures for which the p-values must be independent or PRDS with the example of the Benjamini-Hochberg procedure. Finally, we demonstrate the application of the BH procedure with sequential permutation p-values to simulated data and real neuroimaging data in Sections 5 and 6, respectively.

2 Recap: Sequential Monte-Carlo testing by betting

In this section, we recap the method by [12] on Monte Carlo testing of a single hypothesis via betting, since this is a key module in the multiple testing methods that we return to in the following sections.

Suppose we observed some data X_0 and have the ability to generate additional data X_1, X_2, \dots , e.g., by permuting the treatment labels of X_0 in a treatment vs. control trial. Let $Y_i = T(X_i)$ for some test statistic T . If the same generating mechanism is used independently for each X_i , $i \geq 1$, then the test statistics Y_1, Y_2, \dots are always exchangeable conditional on Y_0 , which means that the joint distribution of the sequence does not change for any finite permutation of the test statistics. We consider testing the null hypothesis

$$H_0 : Y_0, Y_1, \dots \text{ are exchangeable} \quad (2)$$

against the alternative

$$H_1 : Y_1, Y_2, \dots \text{ are exchangeable conditional on } Y_0, \text{ and } Y_0 \text{ is stochastically larger than } Y_1, Y_2, \dots \quad (3)$$

Note that these hypotheses implicitly assume that we sample the permutations with replacement, since there would be an upper bound for the number of permutations otherwise. However, this is the usual state of affairs in practice, since remembering which permutations have already been drawn is computationally intensive and if the data size is large, not all permutations can be drawn anyway [36]. Thus, the standard approach to this testing problem uses Monte Carlo sampling, that is the permutation p-value (1). In general, one would want to choose B as large as possible, since the power increases with B . However, the computational effort can be too high for large B . In order to reduce the computational cost in permutation testing, Besag and Clifford [7] introduced a sequential permutation

p-value:

$$\mathbf{p}_{\gamma(h,B)}^{\text{BC}} = \begin{cases} h/\gamma(h,B), & K_{\gamma(h,B)} = h \\ (K_{\gamma(h,B)} + 1)/(B + 1), & \text{otherwise,} \end{cases} \quad (4)$$

where h is some predefined parameter, $K_B := \sum_{t=1}^B 1\{Y_t \geq Y_0\}$ is the random number of “losses” after B permutations and $\gamma(h, B) = \min(\inf\{t \in \mathbb{N} : K_t = h\}, B)$ is a stopping time. In case of $B = \infty$, we just write $\gamma(h)$. The Besag-Clifford (BC) p-value only allows to stop at the particular time $\gamma(h, B)$ and not at any other stopping time. Hence, it allows to stop early when h losses occurred before all B permutations were sampled and therefore only saves computations under the null hypothesis, but not when the evidence against the null is strong.

To solve this issue, Fischer and Ramdas [12] recently introduced an anytime-valid permutation test based on a testing by betting technique [19, 34]. This makes it possible to look at the data at any time point and then decide whether to stop and make a decision or to continue sampling based on the evidence gathered so far. They showed that this approach is fundamental for the construction of (sequential) permutation tests, meaning that any rank-based permutation test can be obtained by betting. Furthermore, they were able to reduce the number of permutations compared to the classical permutation p-value and the BC method substantially, while achieving a similar power.

Their approach works as follows. At each step $t = 1, 2, \dots$, the statistician bets on the outcome of the relative rank $R_t = \sum_{i=0}^t 1\{Y_i \geq Y_t\}$ by specifying a non-negative $t + 1$ -dimensional betting vector B_t . If $R_t = r$, the wealth of the statistician gets multiplied by $B_t(r)$. Starting with an initial wealth of $W_0 = 1$, after t rounds of gambling the wealth of the statistician is given by

$$W_t = \prod_{s=1}^t B_s(R_s).$$

The wealth process $(W_t)_{t \in \mathbb{N}}$ is a test martingale with respect to the filtration generated by the relative ranks $(R_t)_{t \in \mathbb{N}}$, if

$$\sum_{r=1}^{t+1} B_t(r) = t + 1 \quad (t \in \mathbb{N}).$$

By the optional stopping theorem, this implies that $(W_t)_{t \in \mathbb{N}}$ is an anytime-valid e-value, meaning $\mathbb{E}_{H_0}[W_\tau] \leq 1$ for any stopping time τ . In addition, by Ville’s inequality $(\mathbf{p}_t)_{t \in \mathbb{N}}$, where

$$\mathbf{p}_t := \frac{1}{\max_{s=1, \dots, t} W_s}, \quad (5)$$

is an anytime-valid p-value, meaning $\mathbb{P}_{H_0}(\mathbf{p}_\tau \leq \alpha) \leq \alpha$ for all $\alpha \in [0, 1]$. In particular, this shows that a hypothesis could be rejected as soon as the wealth is larger or equal than $1/\alpha$. The concrete structure of the wealth depends on the betting strategy used.

An anytime-valid generalization of the Besag-Clifford method. The wealth of the *aggressive strategy* after t permutations is given by

$$W_t^{\text{agg}} = (t + 1)1\{Y_1 < Y_0, \dots, Y_t < Y_0\}. \quad (6)$$

This is the most aggressive strategy, as the wealth hits zero as soon as one of the generated test statistic is larger or equal than the observed one. However, it also requires the least possible number of permutations. At the stopping time $\gamma(1, B)$, the p-value of the aggressive strategy equals the BC p-value with parameter $h = 1$. Since the aggressive strategy is also valid at any stopping time $\tau < \gamma(1, B)$, it can be seen as a generalization of the BC method for $h = 1$. Fischer and Ramdas [12]

showed that the BC method for arbitrary h can also be generalized by their strategy, which results in the p-value

$$\mathbf{p}_t^{\text{avBC}} = \frac{h}{t + h - K_t} \quad (7)$$

at time $t \leq \gamma(h)$, if no maximum number of permutations is prespecified. For $t > \gamma(h)$, we just have $\mathbf{p}_t^{\text{avBC}} = \mathbf{p}_{\gamma(h)}^{\text{avBC}}$. This anytime-valid BC method might seem very simple, since $\mathbf{p}_t^{\text{avBC}}$ is just the classical BC p-value (with $B = \infty$) in the case that only losses occur after step t . However, we have not seen it being explored in the literature before and it seems particularly useful in the multiple testing setting. Note that for a fixed level α , the anytime-valid BC method is equivalent to the classical permutation p-value for $B = \lceil h/\alpha \rceil - 1$, although it might stop earlier [35]. Hence, if the significance level an individual hypothesis is tested at is not predefined, as it is the case in multiple testing, the anytime-valid BC method naturally adapts the number of permutations to this data-dependent level. This simple idea of stopping early because the decision will remain the same if we continue sampling was previously explored for the classical permutation p-value, but not found to be very useful [21, 18]. However, it is much more efficient for the BC method. First, because it automatically adapts to the individual significance level. Second, since usually $h \ll B$, it is much less conservative to assume that all permutations are losses until hitting h losses than B permutations. For example, suppose $h = 10$ and $B = 999$. If $\alpha = 0.01$, the anytime-valid generalizations of the BC method and the permutation p-value are equivalent. They can both stop for accepting the null hypothesis as soon as the number of losses equals 10 or stop for rejecting the null if the number of losses after $990 + k$ permutations is k , $k \in \{0, 1, \dots, 9\}$. Now suppose α increases to 0.05. Suddenly, the anytime-valid BC method allows to stop for rejection after $190 + k$ permutations, if k losses are observed, while at least 950 permutations would need to be sampled such that a rejection can be obtained by the permutation p-value.

The binomial mixture strategy. A more sophisticated betting strategy that particularly leads to desirable asymptotic behavior is the *binomial mixture strategy* [12]. For a predefined parameter $c \in [0, 1]$, the wealth of the *binomial mixture strategy* is given by

$$\bar{W}_t^c = (1 - \text{Bin}(K_t; t + 1, c))/c, \quad (8)$$

where $\text{Bin}(K_t; t + 1, c)$ is the CDF of a binomial distribution with size parameter $t + 1$ and probability c . Fischer and Ramdas [12] showed that

$$\bar{W}_t^c | \{\mathbf{p}_{\text{lim}} = p^{\text{lim}}\} \xrightarrow{a.s.} \begin{cases} 1/c, & \text{if } p^{\text{lim}} \in [0, c) \\ 0, & \text{if } p^{\text{lim}} \in (c, 1] \end{cases} \quad (9)$$

for $t \rightarrow \infty$, where \mathbf{p}_{lim} is the limiting permutation p-value $\mathbf{p}_{\text{lim}} = \lim_{B \rightarrow \infty} \mathbf{p}_B^{\text{perm}}$. Therefore, if we choose $c < \alpha$ for some predefined significance level α , the binomial mixture strategy rejects almost surely after a finite number of permutations if $\mathbf{p}_{\text{lim}} < c$. Hence, we can make the loss compared to the limiting permutation p-value arbitrarily small by choosing $c < \alpha$ arbitrarily close to α . In addition to this desirable asymptotic behavior, an advantage of the binomial mixture strategy compared to the Besag-Clifford method is that it automatically adapts to the difficulty of the decision. If the evidence against the null is strong, the wealth will go fast to $1/c$, and if the evidence for the null hypothesis is strong, it will go fast to 0. Consequently, we obtain a fast decision if the data are unambiguous, but if the decision is close, meaning the limiting permutation p-value is close to α , the binomial mixture strategy will take its time and we can always opt to continue sampling.

In multiple testing scenarios the level of an individual hypothesis might not be predefined and depends on the rejection or non-rejection of the other hypotheses. In the following section, we will show how we can still use the binomial mixture strategy and all other α -dependent betting strategies in multiple testing procedures.

3 Multiple testing with anytime-valid permutation tests by α -dependent p-value calibration

Suppose we have M null hypotheses H_0^1, \dots, H_0^M of the form (2) with corresponding alternatives H_1^1, \dots, H_1^M as in (3). We denote by Y_0^i, Y_1^i, \dots the sequence of test statistics for H_0^i and by K_B^i the random number of losses for hypothesis H_0^i after B permutations. Furthermore, let $I_0 \subseteq \{1, \dots, M\}$ be the index set of true hypotheses, $R \subseteq \{1, \dots, M\}$ the set of rejected hypotheses and $V = I_0 \cap R$ the set of falsely rejected hypotheses. Two of the most common error rates considered in multiple testing are the familywise error rate (FWER) and the false discovery rate (FDR). The FWER is defined as probability of rejecting any true null hypothesis

$$\text{FWER} := \mathbb{P}(|V| > 0). \quad (10)$$

The FDR is the expected proportion of falsely rejected hypotheses among all rejections

$$\text{FDR} := \mathbb{E} \left[\frac{|V|}{|R| \vee 1} \right]. \quad (11)$$

The aim is to control either the FWER or FDR at some prespecified level $\alpha \in (0, 1)$. Although we focus on FWER and FDR in this paper for concreteness, our approach also works with all other common error metrics (FDP tail probabilities, k-FWER, etc.). Control of the FWER implies control of the FDR [5] and is therefore more conservative. Hence, FDR controlling procedures often lead to more rejections and are more appropriate in large-scale exploratory hypothesis testing. In non-exploratory validation studies, FWER control is usually the norm.

Many multiple testing methods take p-values $\mathbf{p}^1, \dots, \mathbf{p}^M$ as inputs and reject H_0^i , if $\mathbf{p}^i \leq \alpha_i$ for some potentially data-dependent individual significance level $\alpha_i \in [0, 1]$. Therefore, we do not know in advance at which individual significance level the hypothesis H_0^i will be tested at. However, most betting strategies derived by Fischer and Ramdas [12], including the binomial mixture strategy defined in Section 2, require the significance level as input and thus it need to be fixed in advance. In the following, we show how this can be circumvented by calculating an anytime-valid e-value for each level $\alpha \in [0, 1]$ and then calibrating these into an anytime-valid p-value by taking the smallest level α at which the corresponding e-value can reject the null hypothesis.

For each null hypothesis H_0^i we define a strategy we would use if we test at level $\alpha' \in (0, 1)$. Let $W_t^{i, \alpha'}$ be the wealth of the strategy for hypothesis H_i and level α' after t permutations. We suppose that for all $\alpha'_1 < \alpha'_2$,

$$\{W_t^{i, \alpha'_1} \geq 1/\alpha'_1\} \subseteq \{W_t^{i, \alpha'_2} \geq 1/\alpha'_2\}. \quad (12)$$

This means that if our strategy for level α'_1 rejects at level α'_1 , our strategy for some larger level $\alpha'_2 > \alpha'_1$, needs to reject at level α'_2 as well. In particular, this is satisfied if we use for each $\alpha' \in (0, 1)$ the binomial mixture strategy with parameter $c = b\alpha'$ for some constant $b \in (0, 1)$. To see this, note that

$$\bar{W}_t^c \geq 1/\alpha' \iff \frac{1 - \text{Bin}(K_t; t+1, b\alpha')}{b} \geq 1$$

and $(1 - \text{Bin}(k; t+1, b\alpha'))/b$ is monotone increasing in α' . In the following, we therefore also use the parameter b instead of c to parameterize the binomial mixture strategy in an α -independent way.

We define the p-value for the i -th hypothesis at step $t \in \mathbb{N}$ as

$$\mathbf{p}_t^i = \inf\{\alpha \in (0, 1) \mid \exists s \in \{1, \dots, t\} : W_s^{i, \alpha} \geq 1/\alpha\}. \quad (13)$$

For example, the anytime-valid version of the BC method (7) is the result of such an α -dependent calibration, however, Fischer and Ramdas [12] just used this implicitly and did not write this approach down in its general form. Note that if we use the same strategy for each α , as it is for example the case with the aggressive strategy, then (12) is trivially satisfied and the p-value in (13) reduces to the one in (5).

Proposition 3.1. *If (12) holds, then $(\mathbf{p}_t^i)_{t \in \mathbb{N}}$ in (13) is an anytime-valid p-value.*

Proof. Due to Ville's inequality and Condition (12), for every stopping time τ_i we have

$$\mathbb{P}_{H_0^i}(\mathbf{p}_{\tau_i}^i \leq \alpha) = \mathbb{P}_{H_0^i}(\exists t \in \{1, \dots, \tau_i\} : W_t^{i, \alpha} \geq 1/\alpha) \leq \alpha.$$

□

It is important to note that $(\mathbf{p}_t^i)_{t \in \mathbb{N}}$ is not only valid at any stopping time τ_i , but also at any random time T_i which is not required to be adapted to the filtration generated by the sequential ranks for H_0^i . This means, we can use any (data-dependent) information to stop the testing process without violating the validity of the stopped p-value. In particular, we can stop and reject H_0^i as soon as \mathbf{p}_t^i is less or equal than α_i for some $t \geq 1$, or equivalently, if $W_t^{i, \alpha_i} \geq 1/\alpha_i$, where α_i is the (possibly) data-dependent threshold obtained by a multiple testing procedure. Nevertheless, since random and stopping times are equivalent for anytime-valid p-values [27], we will stick to the term stopping time in the following.

Consequently, we can choose any betting strategy from the very general class provided by (12), stop at any time we want and then apply the multiple testing procedure to the stopped p-values. This is the first multiple testing method for sequential permutation p-values providing such flexibility. We are mainly interested in monotone multiple testing procedures [37]. That means, if $d : [0, 1]^M \rightarrow \{0, 1\}^M$ is the multiple testing procedure mapping the p-values to the decisions (1 = reject; 0 = accept), then d should be coordinatewise nonincreasing. Hence, if some of the p-values become smaller, all previous rejections remain and additional rejections might be obtained. A useful property of our anytime-valid p-value $(\mathbf{p}_t^i)_{t \in \mathbb{N}}$ is that it is nonincreasing in time (see (13)). Together with the monotonicity of the multiple testing procedure, this implies that as soon as \mathbf{p}_t^i can be rejected by the multiple testing procedure, we can stop the sampling process for H_0^i and report a rejection, since the decision would not change if we continue sampling. Our general method is summarized in Algorithm 1. This also shows that we do not need to evaluate $W_t^{i, \alpha}$ for each $\alpha \in [0, 1]$ in order to calculate \mathbf{p}_t^i , but to compare \mathbf{p}_t^i with α_i we can just check whether $W_t^{i, \alpha_i} \geq 1/\alpha_i$, saving a lot of computational effort. It should be noted that in its most general case, error control is only provided if the multiple testing procedure works under arbitrary dependence. This is captured in the following theorem, whose proof follows immediately from the explanations above.

Theorem 3.2. *If d is a monotone multiple testing procedure with FWER or FDR control under arbitrary dependence of the p-values, then Algorithm 1 controls the FWER or FDR respectively at level α .*

For FWER control, the class of monotone multiple testing procedures under arbitrary dependence includes all Bonferroni-based procedures such as the Bonferroni-Holm [25], the sequentially rejective graphical multiple testing procedure [8] and the fixed sequence method [30]. For FDR control, the Benjamini-Yekutieli (BY) procedure [6] is most common. In order to control the desired error rate under arbitrary dependence, the procedures usually need to protect against the worst case distribution. Therefore, improvements can be derived under additional assumptions. A typical assumption is positive regression dependence on a subset (PRDS), under which Hochberg's [24] and Hommel's [26] procedure provide uniform improvements of Holm's procedure and the famous Benjamini-Hochberg (BH) procedure uniformly improves BY. However, if PRDS is required, caution with respect to the choice of betting strategy and stopping time needs to be taken. We address this in the next section by the example of the BH procedure.

Algorithm 1 General sequential multiple testing with α -dependent p-value calibration

Input: Overall significance level α , wealth strategies $W_t^{i, \alpha'}$ that satisfy (12), monotone p-value based multiple testing procedure d , data-dependent stopping rules \mathcal{S}^i and sequences of test statistics $Y_0^i, Y_1^i, \dots, i \in \{1, \dots, M\}$.

Output: Stopping times τ_1, \dots, τ_M and rejection set R .

```
1:  $A = \{1, \dots, M\}$ 
2:  $R = \emptyset$ 
3: for  $t = 1, 2, \dots$  do
4:   for  $i \in A$  do
5:     Check whether  $W_t^{i, \alpha_i} \geq 1/\alpha_i$ , where  $\alpha_i$  is the potentially data dependent significance level
6:     obtained by multiple testing procedure  $d$ .
7:     if  $W_t^{i, \alpha_i} \geq 1/\alpha_i$  then
8:        $R = R \cup \{i\}$ 
9:        $A = A \setminus \{i\}$ 
10:       $\tau_i = t$ 
11:    end if
12:    if  $\mathcal{S}^i(\text{data}) = \text{stop}$  then
13:       $A = A \setminus \{i\}$ 
14:       $\tau_i = t$ 
15:    end if
16:  end for
17:  if  $A = \emptyset$  then
18:    return  $\tau_1, \dots, \tau_M, R$ 
19:  end if
20: end for
```

4 The Benjamini-Hochberg procedure with anytime-valid permutation p-values

The Benjamini-Hochberg (BH) procedure [5] rejects all hypotheses H_0^i with p-value $\mathbf{p}^i \leq m^* \alpha / M$, where

$$m^* = \max \left\{ m \in \{1, \dots, M\} : \sum_{i=1}^M 1\{\mathbf{p}^i \leq m\alpha/M\} \geq m \right\} \quad (14)$$

with the convention $\max(\emptyset) = 0$. In the following we write m_t^* for the BH threshold at time $t \in \mathbb{N}$ when the p-values in (14) are replaced by anytime-valid p-values at time t . The BH procedure controls the FDR when the p-values are positive regression dependent on a subset (PRDS). In order to define PRDS we need the notion of an increasing set. A set $D \subseteq \mathbb{R}^M$ is increasing, if $z \in D$ implies $y \in D$ for all $y \geq z$.

Property 1 (PRDS [6, 11]). *A random vector of p-values \mathbf{p} is weekly PRDS on I_0 if for any null index $i \in I_0$ and increasing set $D \subseteq \mathbb{R}^M$, the function $x \mapsto \mathbb{P}(\mathbf{p} \in D \mid \mathbf{p}^i \leq x)$ is nondecreasing in x for any $x \in \mathbb{R}$. We call \mathbf{p} strongly PRDS on I_0 , if “ $\mathbf{p}^i \leq x$ ” is replaced with “ $\mathbf{p}^i = x$ ”.*

Strong PRDS implies weak PRDS, but for controlling the FDR with the BH procedure, weak PRDS on I_0 is sufficient [11]. When we just write PRDS in the remainder of this paper, we always mean weak PRDS. For example, the PRDS condition holds when the null p-values are independent from each other and the non-nulls. However, it also holds when there is some kind of positive dependence, which is an appropriate assumption in many trials. To guarantee FDR control by applying the BH procedure to our anytime-valid p-values, we need to ensure that our stopped p-values are PRDS. In the following we discuss when this is a reasonable assumption.

4.1 Stopping early for rejection

When PRDS is required, one needs to be careful with choosing the stopping time, as it could possibly induce some kind of negative dependence. For example, suppose we stop sampling for hypothesis H_0^i as soon as $\mathbf{p}_t^i \leq m_t^* \alpha / M$. The smaller the other p-values, the sooner we could stop, which potentially induces some negative dependence between the p-values, even if the data used for calculating the different p-values is independent. However, the following proposition shows that stopping early for rejection with the BH procedure cannot inflate the FDR, which is very important as it allows to adapt the number of permutations to the number of rejections.

Theorem 4.1. *Let τ'_i be a stopping time for each $i \in \{1, \dots, M\}$ such that the stopped anytime-valid p-values $\mathbf{p}_{\tau'_1}^1, \dots, \mathbf{p}_{\tau'_M}^M$ are PRDS and define*

$$\tau_i = \inf\{t \geq 1 : \tau'_i = t \vee \mathbf{p}_t^i \leq m_t^* \alpha / M\}. \quad (15)$$

Then the BH procedure applied on $\mathbf{p}_{\tau_1}^1, \dots, \mathbf{p}_{\tau_M}^M$ rejects the same hypotheses as if applied on $\mathbf{p}_{\tau'_1}^1, \dots, \mathbf{p}_{\tau'_M}^M$ and therefore controls the FDR.

Proof. Note that \mathbf{p}_t^i , $i \in \{1, \dots, M\}$, is nonincreasing and m_t^* is nondecreasing in t . Hence, if $\mathbf{p}_t^i \leq m_t^* \alpha / M$ for some $t \in \mathbb{N}$, then $\mathbf{p}_s^i \leq m_s^* \alpha / M$ for all $s \geq t$. \square

In the following subsections, we discuss how to choose the stopping times τ'_1, \dots, τ'_M in Theorem 4.1 such that the p-values $\mathbf{p}_{\tau'_1}^1, \dots, \mathbf{p}_{\tau'_M}^M$ are PRDS.

4.2 Stopping early for futility with independent p-values

We differentiate between two different stops; stop for rejection and stop for futility. By stop for futility, we mean that we stop early because it is unlikely that we would reject the hypothesis if we continue sampling. Thus, provided that the wealth of the used strategy is nonincreasing in the number of losses, one could write a stop for futility quite generally as

$$\tau'_i = \inf\{t \geq 1 : K_t^i > z_t^i\}, \quad (16)$$

where the parameters z_t^i can either be constant or dependent on the number of losses of the other hypotheses based on the multiple testing procedure used. For example, the stopping time of the Besag-Clifford method $\gamma_i(h, B)$ is obtained by $z_1^i = \dots = z_{B-1}^i = h - 1$ and $z_B^i = -1$. A stop for futility in combination with a stop for rejection includes most reasonable stopping times. Theorem 4.1 shows that stopping for rejection is never a problem. Therefore, it remains to show that stopping for futility does not violate the PRDS condition as well. One obvious sufficient condition is if the data for different hypotheses is independent and the stopping time τ'_i only depends on the data for the corresponding hypothesis H_0^i , which is, for example, the case if the z_t^i in (16) are constants. This is captured in the following proposition.

Proposition 4.2. *Suppose we generate the permuted samples independently for all hypotheses and the vector of limiting p-values \mathbf{p}_{lim} is independent on I_0 . Furthermore, for each $i \in \{1, \dots, M\}$, let the stopping time τ'_i be given by $\tau'_i = \inf\{t \in \mathbb{N} : K_t^i > z_t^i\}$, where z_t^i are constant parameters. Then the stopped p-values $\mathbf{p}_{\tau'} = (\mathbf{p}_{\tau'_1}^1, \dots, \mathbf{p}_{\tau'_M}^M)$ are independent on I_0 . In particular, applying BH to $\mathbf{p}_{\tau} = (\mathbf{p}_{\tau_1}^1, \dots, \mathbf{p}_{\tau_M}^M)$, where τ_i is defined as in (15), controls the FDR.*

Proof. Due to De Finetti's theorem K_t^i , $i \in \{1, \dots, M\}$ and $t \geq 1$, follows a mixture binomial distribution with size parameter t and random probability $\mathbf{p}_{\text{lim}}^i$. Since \mathbf{p}_{lim} is independent on I_0 , $\mathbf{p}_{\tau'}$ is independent on I_0 as well. \square

4.3 Stopping early for futility with PRDS p-values

If there is some positive dependence between the data for the different hypotheses, proving PRDS for the anytime-valid p-values becomes more difficult. The reason for this is that \mathbf{p}_t^i is not only a function of the losses at step t , but might depend on the losses of all previous steps K_1^i, \dots, K_t^i . This can lead to situations in which a larger p-value \mathbf{p}_t^i yields more evidence against a large $\mathbf{p}_{\text{lim}}^i$, which makes it hard to use that $\mathbf{p}_{\text{lim}}^i$ is PRDS for proving that \mathbf{p}_t^i is PRDS. To illustrate this, consider a simple example in which for some $x < x^*$: $\mathbf{p}_3^i \leq x$ is equivalent to $K_1^i \leq 0 \cup K_2^i \leq 0 \cup K_3^i \leq 0$, which reduces to $K_1^i \leq 0$, and $\mathbf{p}_3^i \leq x^*$ is equivalent to $K_1^i \leq 0 \cup K_2^i \leq 0 \cup K_3^i \leq 1$, which reduces to $K_1^i \leq 0 \cup K_3^i \leq 1$. Recall that K_t^i follows a binomial distribution with size parameter t and random probability $\mathbf{p}_{\text{lim}}^i$. With this, under the assumption that $\mathbf{p}_{\text{lim}}^i$ follows a uniform distribution, one can check that $\mathbb{P}(\mathbf{p}_{\text{lim}}^i > 0.9 \mid K_1^i \leq 0) = 0.01 > 0.0091 = \mathbb{P}(\mathbf{p}_{\text{lim}}^i > 0.9 \mid K_1^i \leq 0 \cup K_3^i \leq 1)$. To understand why this is the case, note that $K_1^i \leq 0$ yields the following set of possible loss sequences $\{(0, 1, 1), (0, 0, 1), (0, 1, 0), (0, 0, 0)\}$, while $K_1^i \leq 0 \cup K_3^i \leq 1$ only adds $(0, 0, 1)$ to this set. For this reason, $K_1^i \leq 0 \cup K_3^i \leq 1$ provides more evidence against small values for $\mathbf{p}_{\text{lim}}^i$, but also more evidence against very large values of $\mathbf{p}_{\text{lim}}^i$. Therefore, $\mathbf{p}_{\text{lim}}^i$ is not PRDS on \mathbf{p}_t^i in general.

Nevertheless, we expect the effect from the above phenomenon, that larger bounds for the number of losses can lead to a lower probability for very large values of the limiting permutation p-value, to be very minor and not inflating the FDR. If we choose a betting strategy with nonincreasing wealth in the number of losses and the z_t^i in (16) are constant, then the stopped anytime-valid p-values $\mathbf{p}_{\tau_i}^i$, $i \in \{1, \dots, M\}$, are coordinatewise nondecreasing and nonrandom functions of the number of losses $(K_t^i)_{t \in \mathbb{N}}$. Due to De Finetti's theorem, \mathbf{p}_{lim} being PRDS implies that (K_t^1, \dots, K_t^M) is PRDS for any $t \in \mathbb{N}$. Hence, in this case $\mathbf{p}_{\tau_1}^1, \dots, \mathbf{p}_{\tau_M}^M$ potentially have some kind of positive association as well and therefore we argue that it is reasonable to replace the classical permutation p-values with our anytime-valid ones in the BH procedure. Even if the z_t^i depend in a nonincreasing way on the losses of the other hypotheses, we do not see any reason why this should hurt the PRDS of the stopped p-values. Since we can additionally stop for rejection (Theorem 4.1), this provides a large class of possible betting strategies and stopping times. Fischer and Ramdas [12] proposed to stop for futility if the wealth drops below α . This can easily be transferred to a condition of the form (16) with constant parameters z_t^i . In the multiple testing case the level an individual hypothesis is tested at is not fixed in advance. For this reason, we propose to stop sampling for H_0^i at time t , if $W_t^{i, \alpha_t^{\max}} < \alpha_t^{\max}$, where $\alpha_t^{\max} := \alpha(|A_t| + m_t^*)/M$ and A_t is the index set of hypotheses for which the testing process did not stop before step t . Hence, α_t^{\max} can be interpreted as the maximum level a hypothesis will be tested at by the BH procedure according to the information up to step t . When using the binomial mixture strategy, the combined stopping time of this stop for futility with a stop for rejection is almost surely finite if $\mathbf{p}_{\text{lim}}^i$ is continuously distributed (see (9)). The detailed algorithm of the entire BH procedure using the binomial mixture strategy is illustrated in Appendix C. Of course, this is only one possible stop for futility and there are many other reasonable choices. For example, in situations where the sampling process takes several weeks or months it would also be possible to choose the stopping time interactively, meaning to revisit the data at some point and decide for which hypotheses to continue sampling based on the study interests and the evidence gathered so far.

We believe that for most applications this heuristic argumentation should be reasonable enough to replace the classical permutation p-values with our sequential ones in the BH procedure. Note that the PRDS of the classical permutation p-values cannot usually be checked in practice either and is just assumed if there is no explicit evidence against this. However, if this heuristic argumentation for PRDS of the stopped p-values is not allowed, one could still use the anytime-valid version of the BC method (7) with its incorporated stop for futility $\gamma(h)$ as shown by the next proposition.

Proposition 4.3. *Suppose we generate the permuted samples independently for all hypotheses and the vector of limiting p-values \mathbf{p}_{lim} is strongly PRDS on I_0 . Then the anytime-valid BC p-values $\mathbf{p}_{\gamma(h)}^{\text{avBC}} = (\mathbf{p}_{\gamma_1(h)}^{\text{avBC}, 1}, \dots, \mathbf{p}_{\gamma_M(h)}^{\text{avBC}, M})$, where $\gamma_i(h) = \inf\{t \geq 1 : K_t^i = h\}$, are weakly PRDS on I_0 . In particular, applying BH to $\mathbf{p}_{\tau}^{\text{avBC}} = (\mathbf{p}_{\tau_1}^{\text{avBC}, 1}, \dots, \mathbf{p}_{\tau_M}^{\text{avBC}, M})$, where τ_i is defined as in (15) for $\tau_i' = \gamma_i(h)$,*

controls the FDR.

The proof can be found in the appendix.

5 Simulated experiments

In this section, we aim to characterize the behavior of applying the BH procedure with the proposed sequential permutation p-values using simulated data. We consider an independent Gaussian mean multiple testing problem. That means, each hypothesis is given by $H_0^i : \mathbb{E}[Y_0^i] = 0$, $i \in \{1, \dots, M\}$, where Y_0^i follows a standard normal distribution under the null hypothesis and a shifted standard normal distribution with mean μ_A under the alternative. The probability for a hypothesis being false was set to $\pi_A \in [0, 1]$ and the generated test statistics Y_1^i, \dots, Y_B^i were sampled from $N(0, 1)$. Note that this matches the setting described in Section 2.

5.1 Power and average number of permutations

First, we evaluate the power and average number of permutations per hypothesis using the aggressive strategy (6), the anytime-valid generalization of the BC method with $h = 10$ (7), the binomial mixture strategy (8) with $b = 0.9$, the AMT algorithm by Zhang et al. [41] with $\delta = 0.01$ and the classic permutation p-value (1). The early stops for our sequential methods were chosen as described in Section 4, but all methods were stopped after a maximum number of $B = 10\,000$ permutations. In addition to applying the classical permutation p-value for $B = 10\,000$, we also evaluated it for $B = 200$ as a reference. It should be noted that the AMT algorithm was applied at an overall level of $\alpha - \delta$ such that FDR control at level α is provided [41].

The results in Figure 1 were obtained by averaging over 10 independently simulated trials. Similarly as in [41], we set the standard values of the simulation parameters to $\pi_A = 0.4$, $\mu_A = 2.5$, $\alpha = 0.1$ and $M = 1000$, while one of them was varied in each of the plots. It can be seen that the anytime-valid BC method and binomial mixture strategy lead to a similar power as the classical permutation p-value for $B = 10\,000$ permutations, while being able to reduce the number of permutations by orders of magnitude. The anytime-valid BC method performs slightly better than the binomial mixture one, however, the binomial mixture strategy provides more flexibility with respect to the stopping time and one could, for example, continue sampling for the hypotheses where the data looks promising but did not lead to a rejection after the first 10 000 permutations. In addition, it has advantages with respect to early reporting of rejections as described in the next section. When the number of permutations of the classical permutation p-value is reduced to 200, the power reduces substantially, particularly if the proportion of false hypotheses, strength of the alternative or significance level is small. Since the anytime-valid BC and binomial mixture method also need approximately 200 permutations on average, this shows that the performance of these sequential methods cannot be accomplished by the permutation p-value with a fixed number of permutations. The AMT algorithm was outperformed by the anytime-valid BC method and the binomial mixture strategy in terms of power and number of permutations in all considered scenarios. The use of the aggressive strategy can be reasonable when the main goal is to reduce the number of permutations, while a power loss is acceptable. It should be noted that the behavior of the methods does not change much with the number of hypotheses, since the other parameters remain fixed, which implies that m_t^*/M , and thus the significance level of BH procedure, remain approximately constant. Lastly, we would like to highlight that the results for all these different constellations of the simulation parameters were obtained with the same hyperparameters for the sequential methods, which thus seem to be universally applicable choices.

5.2 Early reporting of decisions

In the previous subsection, we have shown that a lot of permutations can be saved by our sequential strategies, while the power remains similar to the classical permutation p-value. However, reducing the

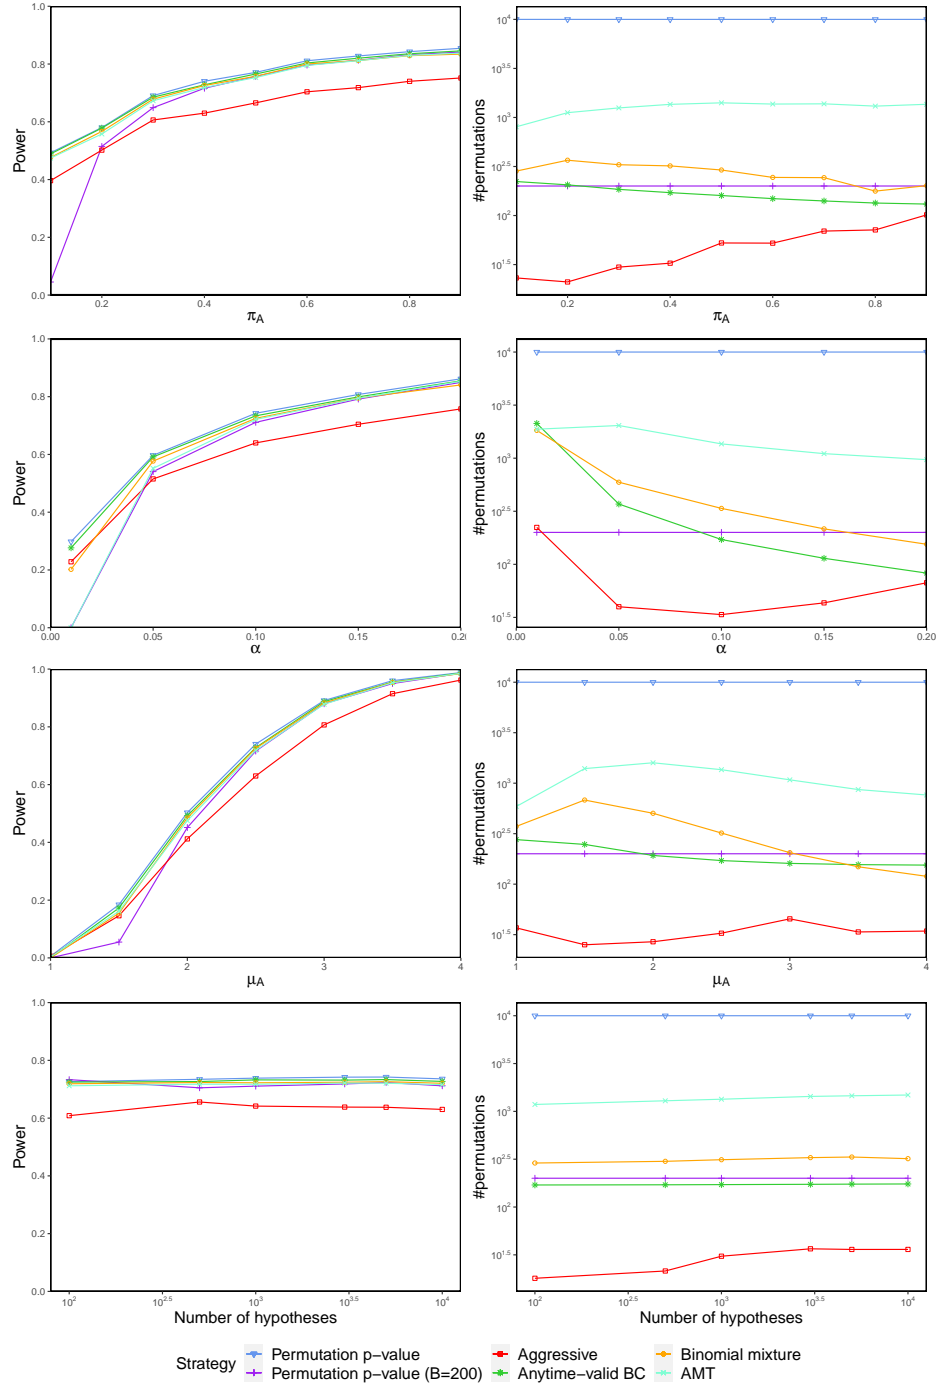


Figure 1: Power and average number of permutations per hypothesis for varying simulation parameters obtained by applying the BH procedure to different (sequential) permutation strategies. The anytime-valid BC method and the binomial mixture strategy lead to marginally less power than the classical permutation p-value, while reducing the number of permutations by orders of magnitude.

total number of permutations is not the only way of increasing efficiency with sequential permutation tests. We can also report already made decisions, and particularly rejections, before the entire process has stopped. This might not be important if the entire procedure only takes some hours or one day to run. However, in large-scale trials generating all required permutations and performing the tests can take up to several months or even longer. This can slow down the research process, as the results are crucial for writing a scientific paper or identifying follow up work. Therefore, it would be very helpful to be able to already report the unambiguous decisions at an earlier stage of the process. Note that this is not possible with the AMT algorithm by Zhang et al. [41], since no decisions can be obtained until the entire process has finished.

In Figure 2 we show the distribution of the rejection times in an arbitrary simulation run using the anytime-valid BC and the binomial mixture strategy in the standard setting described in Section 5.1. In this case, the binomial mixture strategy rejected 305 and the anytime-valid BC method 304 hypotheses, while the former needed a mean number of 459 and the latter of 328 permutations to obtain a rejection. However, the distribution of the stopping times looks very different. More than 50% of the rejections made by the binomial mixture strategy were obtained after 147 permutations and more than 75% after less than 180 permutations, while the rest is distributed quite broadly up to 10 000 permutations. In contrast, all rejections by the anytime-valid BC procedure were made at time 328 such that the entire process was stopped at that time. Basically, this is because for a fixed level α the rejections made by the anytime-valid BC procedure are not obtained in a sequential manner, but we can just reject all hypotheses with $K_{t_\alpha} \leq h - 1$, where $t_\alpha = \lceil h/\alpha \rceil - 1$ (see Section 2). Hence, with the BH procedure we sample until there is a $m^* \in \{1, \dots, M\}$ such that $K_{t_{\alpha m^*/M}} \leq h - 1$ for at least m^* hypotheses. Since $K_{t_{\alpha m^*/M}} \leq h - 1$ for all hypotheses that have not been stopped for futility yet, we can stop the entire process and reject all the remaining hypotheses. However, if the sampling process is stopped before reaching the time $t_{\alpha m^*/M}$ such that $K_{t_{\alpha m^*/M}} \leq h - 1$ for m^* hypotheses (in our example time 328), the anytime-valid BC method would not reject any hypothesis, while the binomial mixture strategy could have already achieved a lot of rejections. Consequently, while the anytime-valid BC method needed less permutations in total, the binomial mixture strategy would be more useful if early reporting of the rejections is desired, since the majority of decisions is obtained much faster. In practice, the binomial mixture strategy could be used to start interpreting the results while keeping to generate further permutations in order to increase the total number of discoveries. Indeed, if the data for some of the undecided hypotheses looks promising, sampling could be continued after the first 10 000 permutations to possibly make even more rejections. In this case, one could also consider increasing the precision of the binomial mixture strategy by choosing the parameter b closer to 1, since a larger average number of permutations might be acceptable when the majority of decisions is obtained fast.

6 Real data analysis

Henderson et al. [23] evaluated the neural responses in the higher visual cortex to natural scene images using fMRI data from the Natural Scenes Dataset [1]. For each considered voxel in each of the eight participants, resulting in a total number of more than 150 000 voxels, they modeled the voxel response by texture statistics [31] in a ridge regression, where the regularization parameter was obtained by cross-validation. To evaluate model accuracy, they calculated the coefficient of determination R^2 on a held-out validation set and assessed significance using a permutation test. For this, they permuted the image labels in the training and validation data and performed the entire fitting and evaluation process for each of these permuted datasets. The final decisions were obtained by applying the BH procedure to the classical permutation p-values at level $\alpha = 0.01$.

Due to the large number of hypotheses and refitting step for each permutation, this problem requires an enormous computational effort. Henderson et al. [23] drew 1000 permutations for each hypothesis and made the resulting R^2 publicly available at <https://osf.io/8fsnx/>, which we also use in our analysis. Note that 1000 permutations is very low for testing such a large number of hypotheses

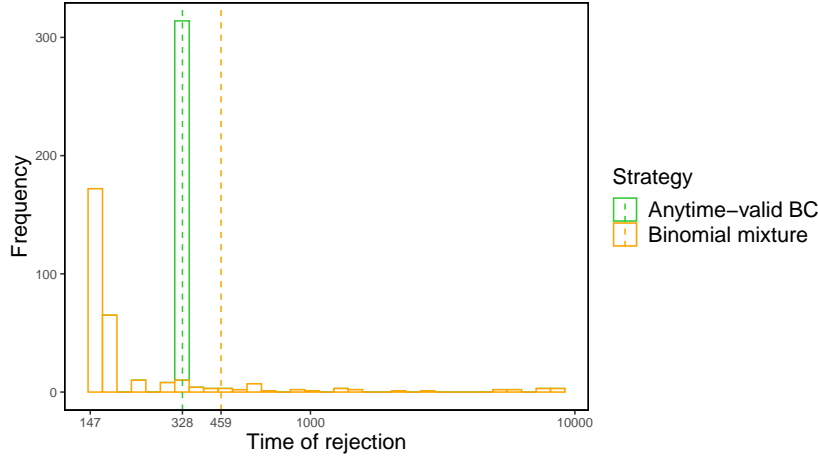


Figure 2: Distribution of the rejection times obtained by the anytime-valid BC and the binomial mixture strategy in a standard simulation run according to Section 5.1. All discoveries by the anytime-valid BC method were made at time 328, while the binomial mixture strategy made more than 75% of the rejections after less than 180 permutations but needed more time on average.

at level $\alpha = 0.01$, but it worked out since there were many very small p-values. The proportion of rejections and average number of permutations per hypothesis obtained by the different strategies are illustrated in Table 1. We applied the anytime-valid BC method and the binomial mixture strategy with parameters $h = 3$ and $b = 0.6$, respectively, to keep the number of permutations low and stopped at the latest when all 1000 permutations were drawn. For the AMT algorithm [41] we changed δ to 0.001 due to the lower α in this case.

Table 1: Proportion of rejected hypotheses and average number of permutations per hypothesis obtained by different methods applied on fMRI data.

Method	Proportion of rejections (%)	Number of permutations
Permutation p-value	62.8	1000
Anytime-valid BC	62.7	411
Aggressive strategy	61.9	138
Binomial mixture strategy	62.2	245
AMT algorithm [41]	62.3	687

The results show that, although Henderson et al. [23] have already chosen a rather low number of permutations, our sequential methods were able to further reduce it significantly while rejecting a similar amount of hypotheses. Also note that the binomial mixture strategy could have continued sampling for hypotheses where the data looks promising after generating the first 1000 permutations to possibly obtain further rejections. Moreover, Henderson et al. [23] were lucky that their risky approach of sampling just 1000 permutations worked out and the brain regions showed the expected response. If the proportion of low p-values was smaller, e.g., if less than 10% of the p-values were smaller than α , then no hypothesis could have been rejected no matter how strong the evidence was. In contrast, our sequential strategies with parameters $h = 3$ and $b = 0.6$ would still have reasonable power, but might have sampled more than 1000 permutations. Note that in this case the AMT algorithm (for a maximum number of $B = 1000$ permutations) would be powerless as well.

7 Conclusion

There are several advantages of using the proposed sequential permutation p-values in large-scale multiple testing problems:

1. They automatically adapt the number of permutations to the number of rejections and thus the proportion of false hypotheses and strengths of the alternatives. For this reason, there is no need to adapt the parameter choice to the unknown data generating process. We found $h = 10$ for the anytime-valid BC method and $b = 0.9$ for the binomial mixture strategy to be universally applicable. Depending on whether the main focus lies in increasing precision or reducing the number of permutations, h and b could also be chosen larger or smaller, respectively.
2. Due to their ability to adapt to the actual significance level obtained by the multiple testing procedure and stopping early as soon as a decision can be made, they reduce the required number of permutations by orders of magnitude. In particular, simulations showed that there is no fixed number of permutations such that the classical permutation p-value performs similar as the sequential permutation p-values. The anytime-valid version of the Besag-Clifford method performed best in terms of power and required number of permutations.
3. They allow to report decisions early such that one can start with writing a paper or identifying follow up work before the entire testing process is finished. In particular, the binomial mixture strategy was found to be useful for such a proceeding.
4. During computing times of several months or more, the initial goals and ideas might change based on external information or the data observed so far. The sequential permutation tests allow to adapt the stopping time and even the betting strategy interactively based on the data. It should be noted that the betting strategy can only be changed for the upcoming permutations and not for the permutations that have already been observed.

In this paper, we showed how anytime-valid permutation tests can be used with p-value based multiple testing procedures. Another type of multiple testing procedures does not rely on individual p-values and explicitly uses the permutation tests to adapt to the unknown dependence structure of the data. This includes the famous MaxT approach for FWER control by Westfall & Young [40], which is particularly powerful when there is a large positive correlation between the test statistics. In Appendix A, we derive sequential versions of such permutation based multiple testing procedures for FWER and simultaneous FDP control. These permutation based multiple testing procedures are based on the closure principle [29] and therefore only consist of tests at level α , which reduces the need for sequential permutation tests. Still, we believe that these can be useful in certain applications.

Acknowledgments

The authors thank Leila Wehbe and Timothy Barry for stimulating practical discussions. LF acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project number 281474342/GRK2224/2. AR was funded by NSF grant DMS-2310718.

References

- [1] Emily J Allen, Ghislain St-Yves, Yihan Wu, Jesse L Breedlove, Jacob S Prince, Logan T Dowdle, Matthias Nau, Brad Caron, Franco Pestilli, Ian Charest, et al. A massive 7T fMRI dataset to bridge cognitive neuroscience and artificial intelligence. *Nature Neuroscience*, 25(1):116–126, 2022.
- [2] Tim Bancroft, Chuanlong Du, and Dan Nettleton. Estimation of false discovery rate using sequential permutation p-values. *Biometrics*, 69(1):1–7, 2013.

- [3] Timothy Barry, Xuran Wang, John A Morris, Kathryn Roeder, and Eugene Katsevich. Sceptre improves calibration and sensitivity in single-cell crispr screen analysis. *Genome biology*, 22:1–19, 2021.
- [4] Stephen Bates, Matteo Sesia, Chiara Sabatti, and Emmanuel Candès. Causal inference in genetic trio studies. *Proceedings of the National Academy of Sciences*, 117(39):24117–24126, 2020.
- [5] Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodology)*, 57(1):289–300, 1995.
- [6] Yoav Benjamini and Daniel Yekutieli. The control of the false discovery rate in multiple testing under dependency. *Annals of Statistics*, pages 1165–1188, 2001.
- [7] Julian Besag and Peter Clifford. Sequential Monte Carlo p-values. *Biometrika*, 78(2):301–304, 1991.
- [8] Frank Bretz, Willi Maurer, Werner Brannath, and Martin Posch. A graphical approach to sequentially rejective multiple test procedures. *Statistics in medicine*, 28(4):586–604, 2009.
- [9] Emmanuel Candès, Yingying Fan, Lucas Janson, and Jinchi Lv. Panning for gold: ‘model-X’ knockoffs for high dimensional controlled variable selection. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 80(3):551–577, 2018.
- [10] Michael P Fay and Dean A Follmann. Designing Monte Carlo implementations of permutation or bootstrap hypothesis tests. *The American Statistician*, 56(1):63–70, 2002.
- [11] Helmut Finner, Thorsten Dickhaus, and Markus Roters. On the false discovery rate and an asymptotically optimal rejection curve. 2009.
- [12] Lasse Fischer and Aaditya Ramdas. Sequential permutation testing by betting. *arXiv preprint arXiv:2401.07365*, 2024.
- [13] Axel Gandy and Georg Hahn. MMCTest—a safe algorithm for implementing multiple Monte Carlo tests. *Scandinavian Journal of Statistics*, 41(4):1083–1101, 2014.
- [14] Axel Gandy and Georg Hahn. A framework for Monte Carlo based multiple testing. *Scandinavian Journal of Statistics*, 43(4):1046–1063, 2016.
- [15] Axel Gandy and Georg Hahn. QuickMMCTest: quick multiple Monte Carlo testing. *Statistics and Computing*, 27:823–832, 2017.
- [16] Christopher R Genovese and Larry Wasserman. Exceedance control of the false discovery proportion. *Journal of the American Statistical Association*, 101(476):1408–1417, 2006.
- [17] Jelle J Goeman and Aldo Solari. Multiple testing for exploratory research. *Statistical Science*, 26(4):584–597, 2011.
- [18] Phillip Good. *Permutation tests: a practical guide to resampling methods for testing hypotheses*. Springer Science & Business Media, 2013.
- [19] Peter Grünwald, Rianne de Heide, and Wouter M Koolen. Safe testing. *Journal of the Royal Statistical Society Series B: Statistical Methodology (with discussion)*, 2024.
- [20] Wenge Guo and Shyamal Peddada. Adaptive choice of the number of bootstrap samples in large scale multiple testing. *Statistical Applications in Genetics and Molecular Biology*, 7(1), 2008.

- [21] Alexander Hapfelmeier, Roman Hornung, and Bernhard Haller. Efficient permutation testing of variable importance measures by the example of random forests. *Computational Statistics & Data Analysis*, 181:107689, 2023.
- [22] Jesse Hemerik and Jelle J Goeman. False discovery proportion estimation by permutations: confidence for significance analysis of microarrays. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 80(1):137–155, 2018.
- [23] Margaret M Henderson, Michael J Tarr, and Leila Wehbe. A texture statistics encoding model reveals hierarchical feature selectivity across human visual cortex. *Journal of Neuroscience*, 43(22):4144–4161, 2023.
- [24] Yosef Hochberg. A sharper Bonferroni procedure for multiple tests of significance. *Biometrika*, 75(4):800–802, 1988.
- [25] Sture Holm. A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, pages 65–70, 1979.
- [26] Gerhard Hommel. A stagewise rejective multiple test procedure based on a modified Bonferroni test. *Biometrika*, 75(2):383–386, 1988.
- [27] Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *Annals of Statistics*, 2021.
- [28] Hui Jiang and Julia Salzman. Statistical properties of an early stopping rule for resampling-based multiple testing. *Biometrika*, 99(4):973–980, 2012.
- [29] Ruth Marcus, Peritz Eric, and K Ruben Gabriel. On closed testing procedures with special reference to ordered analysis of variance. *Biometrika*, 63(3):655–660, 1976.
- [30] Willi Maurer, Ludwig Hothorn, and Walter Lehmacher. Multiple comparisons in drug clinical trials and preclinical assays: a-priori ordered hypotheses. In Vollmar Joachim, editor, *Biometrie in der Chemisch-Pharmazeutischen Industrie*, pages 3–18. Fischer Verlag, Stuttgart, 1995.
- [31] Javier Portilla and Eero P Simoncelli. A parametric texture model based on joint statistics of complex wavelet coefficients. *International Journal of Computer Vision*, 40:49–70, 2000.
- [32] Stan Pounds, Xueyuan Cao, Cheng Cheng, Jun J Yang, Dario Campana, Ching-Hon Pui, William Evans, and Mary Relling. Integrated analysis of pharmacologic, clinical and snp microarray data using projection onto the most interesting statistical evidence with adaptive permutation testing. *International Journal of Data Mining and Bioinformatics*, 5(2):143–157, 2011.
- [33] Geir Kjetil Sandve, Egil Ferkingstad, and Ståle Nygård. Sequential Monte Carlo multiple testing. *Bioinformatics*, 27(23):3235–3241, 2011.
- [34] Glenn Shafer. Testing by betting: a strategy for statistical and scientific communication. *Journal of the Royal Statistical Society Series A: Statistics in Society (with discussion)*, 184(2):407–431, 2021.
- [35] Ivair Silva, Renato Assunção, and Marcelo Costa. Power of the sequential Monte Carlo test. *Sequential Analysis*, 28(2):163–174, 2009.
- [36] Philip B Stark and Kellie Ottoboni. Random sampling: Practice makes imperfect. *arXiv preprint arXiv:1810.10985*, 2018.
- [37] Ajit C Tamhane and Lingyun Liu. On weighted Hochberg procedures. *Biometrika*, 95(2):279–294, 2008.

- [38] Virginia Goss Tusher, Robert Tibshirani, and Gilbert Chu. Significance analysis of microarrays applied to the ionizing radiation response. *Proceedings of the National Academy of Sciences*, 98(9):5116–5121, 2001.
- [39] Anna Vesely, Livio Finos, and Jelle J Goeman. Permutation-based true discovery guarantee by sum tests. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 85(3):664–683, 2023.
- [40] Peter H Westfall and S Stanley Young. *Resampling-based multiple testing: examples and methods for p-value adjustment*, volume 279. John Wiley & Sons, 1993.
- [41] Martin Zhang, James Zou, and David Tse. Adaptive Monte Carlo multiple testing via multi-armed bandits. In *International Conference on Machine Learning*, pages 7512–7522. PMLR, 2019.

A Sequential permutation based multiple testing procedures

We consider the same setting as in Section 3. However, instead of defining the p-values as in (13), we construct a level- α permutation test ϕ_I for each intersection hypothesis $H_0^I = \bigcap_{i \in I} H_0^i$, $I \subseteq \{1, \dots, M\}$, and then use the closure principle to obtain decisions for individual hypotheses. There are two versions of the closure principle. The initial version was proposed by Marcus et al. [29] who showed that the rejection set defined by

$$R = \{i \in \{1, \dots, M\} : \phi_I = 1 \text{ for all } I \text{ with } i \in I\} \quad (17)$$

controls the FWER. However, Goeman and Solari [17] showed that the closure principle can also be used for simultaneous true discovery control. In particular, they proved that

$$\mathbb{P}(\mathbf{d}(S) \leq |S \cap I_1| \text{ for all } S \subseteq \{1, \dots, M\}) \geq 1 - \alpha,$$

where $\mathbf{d}(S) := \min\{|S \setminus I| : I \subseteq \mathbb{N}, \phi_I = 0\}$ and $I_1 = \{1, \dots, M\} \setminus I_0$. In an earlier paper, Genovese and Wasserman [16] proposed an equivalent approach without using the closure principle.

Most permutation based multiple testing procedures are based on choosing some combination function C_I for each $I \subseteq M$, defining the intersection tests ϕ_I by the classical permutation p-value applied on $C_I((Y_0^i)_{i \in I}), C_I((Y_1^i)_{i \in I}), \dots, C_I((Y_B^i)_{i \in I})$ and then applying the closure principle with the intersection tests ϕ_I . To make sure that the ϕ_I are indeed level- α tests such that this yields a valid procedure, the following assumption is usually made.

Assumption 1. *The vectors of test statistics corresponding to true hypotheses $(Y_0^i)_{i \in I_0}, (Y_1^i)_{i \in I_0}, \dots$, are jointly exchangeable.*

This general approach encompasses many existing permutation based multiple testing methods. For example, the MaxT approach for FWER control [40] is obtained by $C_I((Y_j^i)_{i \in I}) = \max_{i \in I} Y_j^i$. Vesely et al. [39] focus on procedures providing true discovery control with $C_I((Y_j^i)_{i \in I}) = \sum_{i \in I} Y_j^i$. Furthermore, the method by Hemerik and Goeman [22], which uniformly improves the popular significance analysis of microarrays (SAM) procedure [38], is obtained by $C_I((Y_j^i)_{i \in I}) = \#\{i \in I : Y_j^i \in D^i\}$ for some prespecified sets D^i .

It is straightforward to derive sequential versions of these multiple testing methods by replacing the classical permutation tests ϕ_I by the anytime-valid ones introduced in Section 2. We summarized this general approach in Algorithm 2.

For a large number of hypotheses M this general approach is computationally infeasible. For this reason, short cuts have been proposed for specific choices of the combination function C_I [40, 22, 39]. In case of the MaxT approach, the entire closed test can be performed with a maximum number of M intersection tests [40]. It might be the case that using these short cuts only works for specific betting

Algorithm 2 Sequential permutation based multiple testing with the closure principle

Input: Combination functions C_I , $I \subseteq \mathbb{N}$, and sequences of test statistics $Y_0^i, Y_1^i, \dots, i \in \{1, \dots, M\}$.

Optional output: Rejection set R for FWER control or function \mathbf{d} for simultaneous FDP control.

- 1: **for** $I \subseteq \{1, \dots, M\}$ **do**
 - 2: Test H_0^I by applying a level- α anytime-valid permutation test ϕ_I to $C_I((Y_0^i)_{i \in I}), C_I((Y_1^i)_{i \in I}), \dots$.
 - 3: **end for**
 - 4: $R = \{i \in \{1, \dots, M\} : \phi_I = 1 \text{ for all } I \text{ with } i \in I\}$
 - 5: $\mathbf{d}(S) = \min\{|S \setminus I| : I \subseteq \{1, \dots, M\}, \phi_I = 0\}$ for all $S \subseteq \{1, \dots, M\}$
 - 6: **return** R, \mathbf{d}
-

strategies. For example, the MaxT short cut can be used with all betting strategies with nonincreasing wealth for an increasing number of losses. The detailed procedure for applying the MaxT approach with the binomial mixture strategy is provided in Algorithm 3.

B Omitted proofs

Lemma B.1. *Let Y, Z and X be real valued random variables on $(\Omega, \mathcal{A}, \mathcal{P})$. If Y is weakly PRDS on X, Z strongly PRDS on Y , and Z and X independent conditional on Y , then Z is weakly PRDS on X .*

Proof. Let D be an increasing set and $x^* \geq x$, then

$$\begin{aligned} \mathbb{P}(Z \in D \mid X \leq x) &= \mathbb{E}_{Y \mid X \leq x}[\mathbb{P}(Z \in D \mid Y, X \leq x)] \\ &= \mathbb{E}_{Y \mid X \leq x}[\mathbb{P}(Z \in D \mid Y)] \\ &= \int_{\Omega} \mathbb{P}(Z \in D \mid Y = y) dF_{Y \mid X \leq x}(y) \\ &\leq \int_{\Omega} \mathbb{P}(Z \in D \mid Y = y) dF_{Y \mid X \leq x^*}(y) \\ &= \mathbb{P}(Z \in D \mid X \leq x^*), \end{aligned}$$

where the inequality follows from the fact that $P(Z \in D \mid Y = y)$ is increasing in y and $F_{Y \mid X \leq x^*}(y) \leq F_{Y \mid X \leq x}(y)$ for all y . \square

Proof of Proposition 4.3. Let $x, x^* \in [0, 1]$ with $x^* \geq x$, $D \subseteq [0, 1]^M$ be an increasing set and $i \in I_0$ be arbitrary but fixed. In the following, we just write \mathbf{p} and \mathbf{p}^i instead of \mathbf{p}^{avBC} and $\mathbf{p}^{\text{avBC}, i}$, respectively. We want to show that

$$\mathbb{P}(\mathbf{p}_{\gamma(h)}^{-i} \in D \mid \mathbf{p}_{\gamma_i(h)}^i \leq x) \leq \mathbb{P}(\mathbf{p}_{\gamma(h)}^{-i} \in D \mid \mathbf{p}_{\gamma_i(h)}^i \leq x^*),$$

where $\mathbf{p}_{\gamma(h)}^{-i} = (\mathbf{p}_{\gamma_1(h)}^1, \dots, \mathbf{p}_{\gamma_{i-1}(h)}^{i-1}, \mathbf{p}_{\gamma_{i+1}(h)}^{i+1}, \dots, \mathbf{p}_{\gamma_M(h)}^M)$. The proof mainly consists of showing the following two claims.

1. $\mathbf{p}_{\text{lim}}^i$ is weakly PRDS on $\mathbf{p}_{\gamma_i(h)}^i$.
2. $\mathbf{p}_{\gamma(h)}^{-i}$ is strongly PRDS on $\mathbf{p}_{\text{lim}}^i$.

Since we assumed that $\mathbf{p}_{\text{lim}}^i$ is strongly PRDS on I_0 , the first claim and Lemma B.1 imply that $\mathbf{p}_{\text{lim}}^i$ is weakly PRDS on $\mathbf{p}_{\gamma_i(h)}^i$. Together with the second claim and the fact that $\mathbf{p}_{\gamma(h)}^{-i}$ is independent of $\mathbf{p}_{\gamma_i(h)}^i$ conditional on $\mathbf{p}_{\text{lim}}^i$ (since we sample independently for all hypotheses), the final proposition follows by Lemma B.1.

We start with proving the first claim. Note that $\mathbf{p}_{\gamma_i(h)}^i \leq x$ iff $K_{t_x} \leq h - 1$, where $t_x = \lceil h/x \rceil - 1$. With this and the fact that $K_{t_x}^i | \mathbf{p}_{\text{lim}}^i = p$ follows a binomial distribution with size parameter t_x and probability p , we obtain

$$\begin{aligned} \mathbb{P}(\mathbf{p}_{\text{lim}}^i \leq p^* \mid \mathbf{p}_{\gamma_i(h)}^i \leq x) &= \frac{\int_0^{p^*} \mathbb{P}(\mathbf{p}_{\gamma_i(h)}^i \leq x \mid \mathbf{p}_{\text{lim}}^i = p) dp}{\int_0^1 \mathbb{P}(\mathbf{p}_{\gamma_i(h)}^i \leq x \mid \mathbf{p}_{\text{lim}}^i = p) dp} \\ &= \frac{\sum_{k=0}^{h-1} \binom{t_x}{k} \int_0^{p^*} p^k (1-p)^{t_x-k} dp}{\sum_{k=0}^{h-1} \binom{t_x}{k} \int_0^1 p^k (1-p)^{t_x-k} dp} \\ &= \frac{\sum_{k=0}^{h-1} (1 - \text{Bin}(k; t_x + 1, p^*))}{h}, \end{aligned}$$

where $\text{Bin}(k; t_x + 1, p^*)$ is the CDF of a binomial distribution with size parameter $t_x + 1$ and probability p^* . Since t_x is decreasing in x , $\mathbb{P}(\mathbf{p}_{\text{lim}}^i \leq p^* \mid \mathbf{p}_{\gamma_i(h)}^i \leq x)$ is decreasing in x as well.

For the second claim, note that $\mathbf{p}_{\gamma_j(h)}^j$ is independent of $\mathbf{p}_{\text{lim}}^{-j}$ and $\mathbf{p}_{\gamma(h)}^{-j}$ conditional on $\mathbf{p}_{\text{lim}}^j$ for all $j \in \{1, \dots, M\}$. Therefore, it only remains to show that $\mathbf{p}_{\gamma_j(h)}^j$ is strongly PRDS on $\mathbf{p}_{\text{lim}}^j$. Since $K_{t_u}^j | \mathbf{p}_{\text{lim}}^j = p$ follows a binomial distribution with probability p , it is immediately implied by the fact that the CDF of a binomial distribution is decreasing in its probability parameter. \square

C Detailed algorithms

Algorithm 3 Sequential MaxT approach with the binomial mixture strategy

Input: Significance level $\alpha \in (0, 1)$, parameter $c \in (0, \alpha)$ for the binomial mixture strategy with uniform prior u_c , ordered test statistics $Y_0^1 \geq \dots \geq Y_0^M$ and sequences of generated test statistics $Y_1^1, Y_2^1, \dots, Y_1^2, Y_2^2, \dots, \dots, Y_1^M, Y_2^M, \dots$.

Output: Stopping times τ_1, \dots, τ_M and index set of rejections $R \subseteq \{1, \dots, M\}$.

```
1:  $\tau^i = \infty$  for all  $i \in \{1, \dots, M\}$ 
2:  $k^i = 0$  for all  $i \in \{1, \dots, M\}$ 
3:  $R = \emptyset$ 
4:  $A = \{1, \dots, M\}$ 
5: for  $t = 1, 2, \dots$  do
6:   for  $i \in A$  do
7:     if  $\max\{Y_t^j : j \in \{i, \dots, M\}\} \geq \max\{Y_0^j : j \in \{i, \dots, M\}\}$  then
8:        $k^i = k^i + 1$ 
9:       if  $(1 - \text{Bin}(k^i; t + 1, c))/c < \alpha$  then
10:         $\tau_i = \min(\tau_i, t)$ 
11:         $\vdots$ 
12:         $\tau_M = \min(\tau_M, t)$ 
13:         $A = A \setminus \{i, \dots, M\}$ 
14:         $R = R \setminus \{i, \dots, M\}$ 
15:      end if
16:    else if  $(1 - \text{Bin}(k^i; t + 1, c))/c \geq 1/\alpha$  then
17:       $\tau_i = t$ 
18:       $R = R \cup \{i\}$ 
19:       $A = A \setminus \{i\}$ 
20:    end if
21:  end for
22:  if  $A = \emptyset$  then
23:    return  $\tau, R$ 
24:  end if
25: end for
```

Algorithm 4 Sequential BH procedure based on the binomial mixture strategy

Input: Significance level $\alpha \in (0, 1)$, parameter $b \in (0, 1)$ for the binomial mixture strategy and sequences of generated test statistics $Y_0^1, Y_1^1, \dots, Y_0^2, Y_1^2, \dots, \dots, Y_0^M, Y_1^M, \dots$.

Output: Stopping times τ_1, \dots, τ_M and index set of rejections $R \subseteq \{1, \dots, M\}$.

```
1:  $m^* = 0$ 
2:  $r_{i,j} = 0$  for all  $i, j \in \{1, \dots, M\}$ 
3:  $k^i = 0$  for all  $i \in \{1, \dots, M\}$ 
4:  $\text{crit}_i = 0$  for all  $i \in \{1, \dots, M\}$ 
5:  $A = \{1, \dots, M\}$ 
6: for  $t = 1, 2, \dots$  do
7:   for  $j = 1, \dots, M$  do
8:      $\text{crit}_j = (\text{Bin})^{-1}(1 - b; t + 1, b\alpha j/M) - 1$ 
9:   end for
10:  for  $i \in A$  do
11:    if  $Y_t^i \geq Y_0^i$  then
12:       $k^i = k^i + 1$ 
13:    end if
14:    for  $j = 1, \dots, M$  do
15:      if  $k^i \leq \text{crit}_j$  then
16:         $r_{i,j} = 1$ 
17:      end if
18:    end for
19:  end for
20:   $m^* = \max\{m \in \{1, \dots, M\} : \sum_{i=1}^m r_{i,m} \geq m\}$ 
21:  for  $i \in A$  do
22:    if  $r_{i,m^*} = 1$  then
23:       $R = R \cup \{i\}$ 
24:       $\tau_i = t$ 
25:       $A = A \setminus \{i\}$ 
26:    end if
27:    if  $1 - \text{Bin}(k^i; t + 1, b\alpha(|A| + m^*)/M) < b[\alpha(|A| + m^*)/M]^2$  then
28:       $\tau_i = t$ 
29:       $A = A \setminus \{i\}$ 
30:    end if
31:  end for
32:  if  $A = \emptyset$  then
33:    return  $\tau_1, \dots, \tau_M, R$ 
34:  end if
35: end for
```
