


# False discovery rate control with e-values

Ruodu Wang<sup>1</sup>

Aaditya Ramdas<sup>2</sup>



Users of these tests speak of the  
5 per cent. point [p-value of 5%]  
in much the same way as I should  
speak of the  $K = 10^{-1/2}$  point  
[e-value of  $10^{1/2}$ ], and of the 1  
per cent. point [p-value of 1%]  
as I should speak of the  
 $K = 10^{-1}$  point [e-value of 10].

## Project “Hypothesis testing with e-values”

Working Paper #9

November 21, 2020

Project web site:  
<http://alrw.net/e>

<sup>1</sup>Department of Statistics and Actuarial Science, University of Waterloo. E-mail: [wang@uwaterloo.ca](mailto:wang@uwaterloo.ca).

<sup>2</sup>Departments of Statistics and Machine Learning, Carnegie Mellon University. E-mail: [aramdas@cmu.edu](mailto:aramdas@cmu.edu).

# Abstract

E-values have gained attention as potential alternatives to p-values as measures of uncertainty, significance and evidence. In brief, e-values are realized by random variables with expectation at most one under the null; examples include betting scores, (point null) Bayes factors, likelihood ratios and stopped supermartingales. We design a natural analog of the Benjamini-Hochberg (BH) procedure for false discovery rate (FDR) control that utilizes e-values, called the e-BH procedure, and compare it with the standard procedure for p-values. One of our central results is that, unlike the usual BH procedure, the e-BH procedure controls the FDR at the desired level—with no correction—for any dependence structure between the e-values. We illustrate that the new procedure is convenient in various settings of complicated dependence, structured and post-selection hypotheses, and multi-armed bandit problems. Moreover, the BH procedure is a special case of the e-BH procedure through calibration between p-values and e-values. Overall, the e-BH procedure is a novel, powerful and general tool for multiple testing under dependence, that is complementary to the BH procedure, each being an appropriate choice in different applications.

**Keywords:** multiple testing, FDR, p-values, betting scores, supermartingales

## Contents

<b>1 Introduction</b>	<b>1</b>
<b>2 When might one prefer e-values over p-values?</b>	<b>2</b>
<b>3 Recap: the BH procedure</b>	<b>5</b>
<b>4 The e-BH procedure</b>	<b>6</b>
<b>5 An FDR guarantee via compliance</b>	<b>8</b>
<b>6 A technical, but interpretable, lemma</b>	<b>11</b>
<b>7 FDR of e-BH for arbitrary e-values</b>	<b>12</b>
<b>8 FDR of e-BH for PRDS e-values</b>	<b>14</b>
<b>9 An optimality result</b>	<b>16</b>
<b>10 Applying the e-BH procedure to p-values</b>	<b>17</b>
<b>11 Simulations: multiple testing for ordered hypotheses using multi-armed bandits</b>	<b>19</b>
<b>12 Conclusion</b>	<b>21</b>
<b>References</b>	<b>22</b>
<b>A Simulations for correlated z-tests</b>	<b>24</b>
<b>B Omitted proofs from the paper</b>	<b>25</b>
<b>C Using e-BH with multiple decreasing transforms</b>	<b>28</b>

# Contents

## 1 Introduction

We study procedures for controlling the false discovery rate (FDR) as in Benjamini and Hochberg [3]. We will encounter p-values in this paper, but focus more on e-values and e-tests as in Vovk and Wang [25]. We use “e-value” as an abstract umbrella term which encompasses betting scores, likelihood ratios, and stopped supermartingales, which appear in the recent literature, e.g., Shafer [20], Grünwald et al. [11], Wasserman et al. [27] and Howard et al. [12, 13]. (Since Bayes factors [16] for point nulls are also e-values, one may hope that e-values appeal to adherents of frequentist, Bayesian and game-theoretic foundations of probability, but this paper abstains from further philosophical discussion, treating e-values as a useful technical tool that arises naturally in many situations.)

Throughout, let  $H_1, \dots, H_K$  be  $K$  hypotheses, and write  $\mathcal{K} = \{1, \dots, K\}$ . Let the true (unknown) data-generating joint probability measure be denoted by  $\mathbb{P}$ . For each  $k \in \mathcal{K}$ , it is useful to think of hypothesis  $H_k$  as implicitly defining a set of joint probability measures, and  $H_k$  is called a true null if  $\mathbb{P} \in H_k$ .

Following Vovk and Wang [25], a *p-variable*  $P$  is a random variable that satisfies  $\mathbb{P}(P \leq \alpha) \leq \alpha$  (often with equality) for all  $\alpha \in (0, 1)$ . In other words, a p-variable is stochastically larger than  $U[0, 1]$  (values of  $P > 1$  can be treated as 1). An *e-variable*  $E$  is a  $[0, \infty]$ -valued random variable satisfying  $\mathbb{E}[E] \leq 1$ .

Let  $\mathcal{N} \subseteq \mathcal{K}$  be the set of indices of null cases, which is unknown to the decision maker, and  $K_0$  be the number of null cases, thus the cardinality of  $\mathcal{N}$ . The ratio  $K_0/K$  may be close to 1, meaning that the signals are sparse.

For reasons to use e-values over p-values, see Shafer [20], Vovk and Wang [25] and Grünwald et al. [11]; however, we summarize our own perspectives later in Section 2. We consider two settings of testing multiple hypotheses, and sometimes convert between them:

1. For each  $k \in \mathcal{K}$ ,  $H_k$  is associated with p-value  $p_k$ , which is a realization of a random variable  $P_k$ . If  $k \in \mathcal{N}$ , then  $P_K$  is a p-variable. A p-testing procedure  $\mathcal{D} : [0, 1]^K \rightarrow 2^{\mathcal{K}}$  gives the indices of rejected hypotheses based on observed p-values.
2. For each  $k \in \mathcal{K}$ ,  $H_k$  is associated with e-value  $e_k$ , which is a realization of a random variable  $E_k$ . If  $k \in \mathcal{N}$ , then  $E_K$  is an e-variable. An e-testing procedure  $\mathcal{G} : [0, \infty]^K \rightarrow 2^{\mathcal{K}}$  gives the indices of rejected hypotheses based on observed e-values.

We tacitly require that all testing procedures are Borel functions. The terms “p-values/e-values” refer to both the random variables and their realized values; these should be clear from the context.

The rejected hypotheses by a p- or e-testing procedure  $\mathcal{D}$  are called discoveries. We write  $F_{\mathcal{D}} := |\mathcal{D} \cap \mathcal{N}|$  as the number of null cases that are rejected (i.e., false discoveries), and  $R_{\mathcal{D}} := |\mathcal{D}|$  as the total number of discoveries. The value of interest is  $F_{\mathcal{D}}/R_{\mathcal{D}}$ , called the false discovery proportion (FDP), which is the ratio of the number of false discoveries to that of all claimed discoveries, with the convention  $0/0 = 0$  (i.e., FDP is 0 if there is no discovery). Since both quantities  $F_{\mathcal{D}}$  and  $R_{\mathcal{D}}$  are random, Benjamini and Hochberg [3] proposed to control FDR, which is the expected value of FDP, that is,  $\text{FDR}_{\mathcal{D}} := \mathbb{E}[F_{\mathcal{D}}/R_{\mathcal{D}}]$ , where the expected value is taken under the true probability.

For those unfamiliar with e-values, we offer a couple of quick remarks. First, an e-variable  $E$  can be converted to a p-variable  $P = 1/E$ ; the validity of  $P$  can be checked using Markov’s inequality. However, since Markov’s inequality is not tight, such a translation is not airtight. A p-variable  $P$  can also be converted to an e-variable, but this is more complex; one simple example is to set  $E = P^{-1/2} - 1$  (and more generally  $f(P)$  is an e-variable if  $\int_0^1 f(u) du \leq 1$ ). Importantly, small p-values correspond to large e-values, so we will reject e-values above some threshold.

## A brief summary of our contributions

For  $k \in \mathcal{K}$ , let  $e_{[k]}$  be the  $k$ -th order statistic of  $e_1, \dots, e_K$ , from *the largest to the smallest*. We design a simple e-value analog of the standard BH procedure, which will be called the *base e-BH procedure*. For  $\alpha > 0$ , define the e-testing procedure  $\mathcal{G}(\alpha) : [0, \infty]^K \rightarrow 2^{\mathcal{K}}$  which rejects hypotheses with the largest  $k_e^*$  e-values, where

$$k_e^* = \max \left\{ k \in \mathcal{K} : \frac{ke_{[k]}}{K} \geq \frac{1}{\alpha} \right\}, \quad (1)$$

with the convention  $\max(\emptyset) = 0$ , and accepts the rest.

The astute reader will note that since  $1/e_k$  is a valid p-value, the base e-BH procedure is simply the BH procedure applied to the corresponding p-values. Hence, at first glance, apparently nothing was gained with the e-value viewpoint. However, one of our central results is the surprising property that

the base e-BH procedure controls FDR at level  $\alpha$  even under unknown *arbitrary dependence* between the e-values.

It is well known that such a statement is not true for general arbitrarily dependent p-values. Moreover, the *full e-BH procedure* involves a pre-screening of e-values, allowing us to “boost” them up by a factor before feeding them to the base e-BH procedure; see Section 4 for details.

We also derive several important, but secondary, results. These include an analysis of compliant e-testing procedures, the fact that known results on FDR control for BH and BY [5] procedures can be derived as special cases of the e-BH procedure, and the optimality of e-BH procedure amongst all FDR controlling procedures acting on arbitrarily dependent e-values.

Finally, the paper contains several other tertiary contributions, including a mathematical study of how conservative the e-BH procedure is under arbitrary dependence, independence and PRDS (positive regression dependence on a subset). We provide a few interesting motivating examples along the way: identifying skillful traders (finance) and identifying promising slot machines by adaptive sampling (multi-armed bandits). Simulation results on the e-BH procedure are produced in Section 11 and Appendix A. Proofs of most results are put in Appendix B.

Before delving into the above, it seems critical to first answer a motivating question: why use e-values in the first place?

## 2 When might one prefer e-values over p-values?

Given that the Benjamini-Hochberg procedure using p-values is so widely established and accepted in the sciences, one must first answer the question of why one should care about FDR control (or multiple testing) with e-values. This is a philosophical, subjective, question and the opinions of the authors (and readers) are complementary to the contributions of this paper. For example, in his upcoming discussion paper, Shafer [20] eloquently provides various arguments for the use of “betting scores”, which are e-values with a game-theoretic interpretation, as measures of evidence and uncertainty in the scientific enterprise while Grünwald et al. [11] find them to be a middle ground between Bayesian and frequentist (and financial, when combined with the earlier view) interpretations of evidence. Other authors have also already weighed in on the subject, but rather than deferring to them, we wish to offer our own perspectives in light of recent research.

Below, we list some situations when one may potentially prefer to use e-values over p-values for non-philosophical, purely statistical, reasons. To avoid confusion, we note that *the following discussion relates to the construction of a single e-value or a single p-values*—in other words, it applies equally to single and multiple hypothesis testing.

1. **High-dimensional asymptotics.** One of the most classical ways to compute p-values is to use the asymptotic distribution of the likelihood ratio test statistic, as given by Wilk’s theorem.

However, the correctness of Wilk’s theorem is typically justified when the dimensionality  $d$  of the data remains fixed, and the sample size  $n$  tends to infinity. There are several results on the high-dimensional asymptotics of likelihood ratios, but the resulting p-value relies on the practitioner making assumptions on the relative scalings of  $n$  and  $d$ ; see Jiang and Yang [15] for example. In contrast, the likelihood ratio for a point null hypothesis is a valid e-value in finite samples, meaning that its expectation equals one under the null regardless of  $d$  or  $n$ .

2. **Irregular (composite) models.** There are many composite null hypothesis testing problems for we may know of no direct way to construct a valid p-value even under low-dimensional asymptotics—this could happen because the model is singular or irregular and Wilk’s theorem fails to hold [7] and in such cases the validity of the bootstrap is also typically unknown. Recently, the *split* likelihood ratio statistic [27] was developed and shown to yield an e-value under no assumptions on the composite null, or on  $d$  and  $n$ . Examples of new settings in which one can now construct e-values under no regularity assumptions include mixtures (e.g., testing if data comes from a mixture of  $\leq k$  versus  $> k$  components), shape-constraints (e.g., log-concavity), dependence structures (e.g., multivariate total positivity), and several latent variable models.
3. **Sequential inference.** The correctness of Wald’s sequential likelihood ratio test is based on noting that the likelihood ratio process  $(L_t)$  is a nonnegative martingale with initial value one. Thus, at any stopping time  $\tau$ , the optional stopping theorem implies that  $L_\tau$  is an e-value since its expectation is at most one. Moving beyond parametric settings, sequential inference is often enabled by designing nonparametric supermartingales [12] that immediately yield e-values at the stopping time. These have been called test martingales in prior work [21]. In summary, e-values arise very naturally in sequential inference as stopped supermartingales.
4. **Accumulation of information and evidence.** Given a set of data, if the ultimate goal of a scientist is to either reject (with level  $\alpha$ ) or not reject a given hypothesis *in a single run*, then the most powerful method is to reject a p-value  $p$  no larger than  $\alpha$ . This testing procedure can also be achieved by a simple e-value  $e := \mathbb{1}_{\{p \leq \alpha\}}/\alpha$ , using a threshold  $1/\alpha$  (called an all-or-nothing bet by Shafer [20]); in fact, no other e-values would be more powerful in general. The situation becomes quite different when the hypothesis (if seen as promising by the first experiment) will be tested with future evidence and possibly by other scientists: the all-or-nothing e-value carries little information for the next studies, whereas a generic e-value can provide a continuum of evidence strength. Even if the e-value is moderate (e.g., 3), it is useful and can be easily merged into the next experiment [25]. On the other hand, a p-value of 0.1 is difficult to use for future studies that may depend on data from the current one, and may be discarded (or worse, “p-hacked”) due to the all-or-nothing nature of the p-test. Hence, for a *dynamic flow* of experiments, common in modern sciences, it is beneficial to report e-values (see Shafer [20] and Grünwald et al. [11]).
5. **Robustness to misspecification.** In genetics, it is not uncommon to encounter p-values with astronomically small values (like  $10^{-20}$ ), or sometimes point masses near the value one, even though the sample sizes may not intuitively support such extreme evidence. This is often (but not always) reflective of utilizing a model that is not perfectly specified. The validity of p-values is quite sensitive to model misspecification, because they utilize the entire (hypothesized) distribution of the test statistic. In contrast, e-values can be constructed without over-reliance on fine-grained tail information, and so they are typically more robust than p-values to misspecification (but less powerful under perfect specification). For example, instead of assuming that the data  $X$  is Gaussian to build a p-value, we may instead assume that it is symmetric about the origin (under no further moment constraints), in which case  $\exp(\lambda X - \lambda^2 X^2/2)$  is a valid e-value for any  $\lambda \in \mathbb{R}$ . The proof, by Bahadur and Eaton, was communicated by Efron [8].
6. **Robustness to dependence.** To construct a single valid p-value, it is often assumed (for convenience of deriving limiting distributions) that the underlying observations are independent,

and indeed their validity is often hurt if this assumption is violated. However, we can often construct an e-value quite easily in settings where the observations are dependent, and this is because we are requiring less of an e-value (just bounded in expectation) than of a p-value (knowledge of its whole distribution). As one example, suppose we observe non-negative data  $X_1, \dots, X_n$ , and we wish to test  $H_0 : \mathbb{E}[X_i] \leq \mu$  for all  $i$ . Then,  $(X_1 + \dots + X_n)/(n\mu)$  is an e-value for any dependence structure of  $X_1, \dots, X_n$ , and does not require making any distributional assumption. In addition, combination of multiple p-values relies heavily on dependence assumptions [24], unlike that of e-values [25]. In the field of risk management, analysis of risks with unknown or complicated dependence has recently been an active topic; see e.g., [9].

Thus, there exist many settings when one can, and should, use e-values to quantify evidence against a null. (However, there of course remain innumerable situations in which p-values are perfectly reasonable choices.)

To reflect some of the above points, in a concrete example below, we will see that a complicated dependence structure appears both within experiment and across experiments, and one can easily construct useful e-values in such situations, with almost no model assumptions on the test statistics.

**Example 2.1.** Suppose there are  $K$  traders (or machines), and a researcher is interested in knowing which ones are skillful (or useful). This is a classic problem in finance; see e.g., Barras et al. [2] in the context of detecting mutual fund skills. For  $k = 1, \dots, K$ , the null hypothesis  $H_k$  is that trader  $k$  is not skillful, meaning that they make no profit on average (without loss of generality we can assume the market risk-free return rate is 0). The nonnegative random variables  $X_{k,1}, \dots, X_{k,n}$  are the monthly realized performance (i.e., the ratio of payoff to investment;  $X_{k,j} > 1$  presents a profit and  $X_{k,j} < 1$  means a loss) of agent  $k$  from month 1 to month  $n$ . The no-skill null hypothesis is  $\mathbb{E}[X_{k,j} | \mathcal{F}_{j-1}] \leq 1$  for  $j = 1, \dots, n$ , where the  $\sigma$ -field  $\mathcal{F}_t$  represents the available market information up to time  $t \in \{0, \dots, n\}$ , and we naturally assume that  $(X_{k,t})_t$  is adapted to  $(\mathcal{F}_t)_t$ .

Since the agents are changing investment strategies over time and all strategies depend on the financial market evolution, there is complicated serial dependence within  $(X_{k,1}, \dots, X_{k,n})$  for single  $k$ , as well as cross dependence among agents  $k = 1, \dots, K$ . Because of the complicated serial dependence and the lack of distributional assumptions of the performance data, it is difficult to obtain useful p-values for these agents. Nevertheless, we can easily obtain useful e-values: for instance,  $E_k = \prod_{j=1}^n X_{k,j}$  is a valid e-value, as well as any mixture of U-statistics of  $X_{k,1}, \dots, X_{k,n}$ , including the mean and the product ( $X_{k,1}, \dots, X_{k,n}$  are called sequential e-values by [25], who also propose methods to merge them into one e-value). Moreover, the obtained e-values  $E_1, \dots, E_K$  are dependent in a complicated way. Even if these e-values are not very large, they can be useful for other studies on these traders. We return to a similar example in the simulations (Section 11).

As a side note, other sophisticated e-values can also be constructed for this problem using exponential self-normalized processes [12]. For example, define  $\psi(\lambda) := -\lambda - \log(1 - \lambda)$ . Then,  $E_k = \prod_{j=1}^n \exp(\lambda(X_{k,j} - 1) - \psi(\lambda)(X_{k,j} - 1)^2)$  for any  $\lambda \in [0, 1)$ . Note that  $\lim_{\lambda \rightarrow 0} \frac{\psi(\lambda)}{\lambda^2/2} = 1$ , so this e-value captures subGaussian-like right tails of  $\sum_{j=1}^n X_{k,j} - 1$  relative to its empirical variance  $\sum_{j=1}^n (X_{k,j} - 1)^2$ . To avoid picking a fixed  $\lambda$ , one can instead simply use a mixture over  $\lambda$ , for example using a Gamma distribution peaking near 0 [13]. One can also pick  $\lambda_j$  predictably at each step. These observations form the basis of all time-uniform Empirical-Bernstein inequalities [13, 28, 29].

To summarize, there are several reasons to work with e-values: they arise naturally in sequential settings, we know how to construct e-values in settings where we do not know how to construct p-values, and e-values can be more robust to misspecification or uncertain asymptotics in high-dimensional settings. Of course, there are also many reasons *not* to work with e-values: in particular, p-values will often yield more “powerful” (single/multiple) tests when the underlying modeling assumptions are true. Thus, e-values and p-values can also be viewed as picking different points on the validity-power curve: if one takes more risks with validity (p-values), one can gain more power, but if one questions one or more modeling assumptions, e-values may provide a safer choice.

*Remark 2.2.* Recall that we can always convert an e-value  $e$  into a p-value  $p = 1/e$ . One can also transform a p-value  $p$  into an e-value  $e = f(p)$  using “calibrators”; recalling the definition by Shafer et al. [21], a calibrator is a decreasing function  $f : [0, 1] \rightarrow [0, \infty)$  such that

$$\int_0^1 f(u) \, du = 1. \quad (2)$$

If the only known way to construct an e-value is by calibrating a bonafide p-value, then e-BH will typically not be more powerful than BH. Nevertheless, in some of our examples, we will assume that we can construct p-values, and calibrate them to e-values and study the resulting e-BH procedure, in order to understand the difference between the two approaches. In contrast, if the only known way to construct a p-value is by inverting an e-value, then working with e-values directly is beneficial because the e-BH procedure is robust to arbitrary dependence without corrections, unlike the BH procedure.

### 3 Recap: the BH procedure

We briefly review the Benjamini-Hochberg (BH) and the Benjamini-Yekutieli (BY) procedures (Benjamini and Hochberg [3], Benjamini and Yekutieli [5]) for FDR control, as the benchmark for our discussions.

To discuss the dependence structure among p-values and e-values, we rely on the notion of *positive regression dependence on a subset (PRDS)* of Benjamini and Yekutieli [5], flipped when imposed on e-values. A set  $A \subseteq \mathbb{R}^K$  is said to be *decreasing* (resp. *increasing*) if  $\mathbf{x} \in A$  implies  $\mathbf{y} \in A$  for all  $\mathbf{y} \leq \mathbf{x}$  (resp. all  $\mathbf{y} \geq \mathbf{x}$ ). In this paper, all terms “increasing” and “decreasing” are in the non-strict sense.

1. A random vector  $\mathbf{P}$  of p-values is PRDS on  $\mathcal{N}$  (or for simplicity, satisfies PRDS) if for any null index  $k \in \mathcal{N}$  and increasing set  $A \subseteq \mathbb{R}^K$ , the function  $x \mapsto \mathbb{P}(\mathbf{P} \in A \mid P_k \leq x)$  is increasing on  $[0, 1]$ .
2. A random vector  $\mathbf{E}$  of e-values satisfies PRDS if for any null index  $k \in \mathcal{N}$  and decreasing set  $A \subseteq \mathbb{R}^K$ , the function  $x \mapsto \mathbb{P}(\mathbf{E} \in A \mid E_k \geq x)$  is decreasing on  $[0, \infty)$ .

We chose to use the version of PRDS in Barber and Ramdas [1] which is weaker than the original one used in Benjamini and Yekutieli [5], but the difference is relatively minor: “ $P_k \leq x$ ” in the former is replaced by “ $P_k = x$ ” in the latter.

If the null p-values (e-values) are mutually independent and independent of the non-null p-values (e-values), then PRDS holds; as such, PRDS is a generalization of independence. Moreover, increasing individual transforms do not affect the PRDS property. Further, the PRDS property is preserved when moving from p-values to e-values using calibrators, or vice versa by inversion. We record this fact below.

**Fact 3.1.** The PRDS property of p-values and that of e-values are equivalent in the following sense: If  $\mathbf{P}$  is a vector of PRDS p-values, then  $\mathbf{E}$  is a vector of PRDS e-values, where  $E_k = f_k(P_k)$  for any calibrator  $f_k$  from (2). Similarly, if  $\mathbf{E}$  is a vector of PRDS e-values, then  $\mathbf{P}$  is a vector of PRDS p-values, where  $P_k = 1/E_k$ . As an important example, PRDS holds if p-values/e-values are built on positively correlated (jointly) Gaussian test statistics  $X_1, \dots, X_K$  (themselves PRDS). For instance, this includes p-values obtained from Neyman-Pearson tests, i.e.,  $P_k = 1 - \Phi(X_k)$  where  $\Phi$  is the standard Gaussian cdf, and e-values obtained from mixture likelihood ratios,  $E_k = \int \exp(\delta X_k - \delta^2/2) \, d\nu_k(\delta)$ , where  $\nu_k$  is an arbitrary mixing distribution; see Example 7.3 for instance. The PRDS property of these p-values and e-values directly follows from the invariance of PRDS under monotone transforms.

Following common practice in the literature, we will study FDR in both the case of PRDS and that of arbitrary dependence. The input of the Benjamini-Hochberg (BH) procedure includes three ingredients:



- (a)  $K$  realized p-values  $p_1, \dots, p_K$  associated to  $H_1, \dots, H_K$ , respectively;
- (b) an FDR level  $\alpha \in (0, 1)$ ;
- (c) (optional) dependence information or assumption on p-values, such as independence, PRDS or no information.

For  $k \in \mathcal{K}$ , let  $p_{(k)}$  be the  $k$ -th order statistics of  $p_1, \dots, p_K$ , from the smallest to the largest. The (base) BH procedure  $\mathcal{D}(\alpha)$  rejects all hypotheses with the smallest  $k^*$  p-values, where

$$k^* = \max \left\{ k \in \mathcal{K} : \frac{K p_{(k)}}{k} \leq \alpha \right\}, \quad (3)$$

with the convention  $\max(\emptyset) = 0$ , and accepts the rest. We summarize some known results of Benjamini and Hochberg [3] and Benjamini and Yekutieli [5] on the procedure  $\mathcal{D}(\alpha)$  below. Throughout, we write

$$\ell_K := \sum_{k=1}^K \frac{1}{k} \approx \log K.$$

**Theorem 3.2.** *For arbitrary p-values and  $\alpha \in (0, 1)$ , the base BH procedure  $\mathcal{D}(\alpha)$  satisfies*

$$\mathbb{E} \left[ \frac{F_{\mathcal{D}(\alpha)}}{R_{\mathcal{D}(\alpha)}} \right] \leq \frac{\ell_K K_0}{K} \alpha \leq \ell_K \alpha.$$

Moreover, if the p-values satisfy PRDS, then

$$\mathbb{E} \left[ \frac{F_{\mathcal{D}(\alpha)}}{R_{\mathcal{D}(\alpha)}} \right] \leq \frac{K_0}{K} \alpha \leq \alpha,$$

and the first inequality is an equality if the null p-values are iid uniform on  $[0, 1]$ .

By Theorem 3.2, if the p-values are PRDS, then the base BH procedure has an FDR of at most  $\alpha$ . If the p-values are arbitrarily dependent, then we need to replace  $\alpha$  in the base BH procedure by  $\alpha/\ell_K$ , resulting in the Benjamini-Yekutieli (BY) procedure  $\mathcal{D}'(\alpha) := \mathcal{D}(\alpha/\ell_K)$ , which rejects all hypotheses with the smallest  $k^*$  p-values, where

$$k^* = \max \left\{ k \in \mathcal{K} : \frac{\ell_K K p_{(k)}}{k} \leq \alpha \right\}. \quad (4)$$

The extra factor  $\ell_K$  in the BY procedure reflects the fact that it is harder to justify discoveries without the PRDS assumption, which includes independence as a special case. It is well known that this factor cannot be improved in general.

## 4 The e-BH procedure

In this section, we design an analog of the BH procedure for e-values, which we call the *e-BH procedure*. Similar to the BH procedure, the input of the e-BH procedure includes three ingredients:

- (a)  $K$  realized e-values  $e_1, \dots, e_K$  associated to  $H_1, \dots, H_K$ , respectively;
- (b) an FDR level  $\alpha \in (0, 1)$ ;
- (c) (optional) distributional information or assumption on e-values.



While (a) and (b) are the same as those of the BH procedure, (c) is somewhat different: in addition to dependence information, e-BH can also accommodate information of marginal distributions, since e-values have more freedom than p-values in terms of their distributions (the former are only constrained by their expectations under the null).

The e-BH procedure can be described in two simple steps.

1. (Optional) Boost the *raw e-values* in (a) using information in (c).
2. Apply the *base e-BH procedure* to the boosted e-values and level  $\alpha$ .

We will first describe the base e-BH procedure in step 2 and then explain how to boost e-values in step 1. If there is no available information in (c), we can simply skip step 1 above, and directly apply the base e-BH procedure to the raw e-values, which always has a valid FDR control of  $\alpha$ .

### The base e-BH procedure

Let  $e'_1, \dots, e'_K$  be the boosted e-values obtained from step 1 of the procedure. Keep in mind that they can be chosen as identical to the raw e-values  $e_1, \dots, e_K$  if there is no information in (c) or one opts to skip the boosting in step 1.

The base e-BH procedure is applying the BH procedure to  $1/e'_1, \dots, 1/e'_K$ . To be precise, for  $k \in \mathcal{K}$ , let  $e'_{[k]}$  be the  $k$ -th order statistic of  $e'_1, \dots, e'_K$ , sorted from the largest to the smallest so that  $e'_{[1]}$  is the largest boosted e-value. The base e-BH procedure  $\mathcal{G}(\alpha) : [0, \infty]^K \rightarrow 2^{\mathcal{K}}$  then rejects hypotheses with the largest  $k_e^*$  (boosted) e-values, where

$$k_e^* = \max \left\{ k \in \mathcal{K} : \frac{ke'_{[k]}}{K} \geq \frac{1}{\alpha} \right\}, \quad (5)$$

with the convention  $\max(\emptyset) = 0$ , and accepts the rest.

We note that, analogous to the BH procedure, there are many equivalent ways of describing the e-BH procedure, including using an estimate of the false discovery proportion as pioneered by Storey [22], or the self-consistency viewpoint of Blanchard and Roquain [6].

### Boosting e-values

To boost e-values in step 1, there are two types of information we can use: marginal distributional information and joint dependence information. Below we first describe our standard choice of boosting, and then the more general choices which can be freely chosen by the user (Remark 4.2). Although we generally recommend the standard choice, the flexibility of choosing boosting functions allows us to include the p-based methods (like BH and BY) as special cases.

Let  $K/\mathcal{K} := \{K/k : k \in \mathcal{K}\}$ , and define a truncation function  $T : [0, \infty] \rightarrow [0, K]$  by letting  $T(x)$  be the largest number in  $K/\mathcal{K} \cup \{0\}$  that is no larger than  $x$ . In other words,

$$T(x) = \frac{K}{\lceil K/x \rceil} \mathbf{1}_{\{x \geq 1\}} \quad \text{with } T(\infty) = K. \quad (6)$$

Note that  $T$  truncates  $x$  to take only values in  $K/\mathcal{K}$  (or zero). The truncation function  $T$  will be frequently used in this paper. The standard boosting method is now described below. For each  $k \in \mathcal{K}$ , take a *boosting factor*  $b_k \geq 1$  (the larger, the better) such that, depending on dependence information,

$$\max_{x \in K/\mathcal{K}} x \mathbb{P}(\alpha b_k E_k \geq x) \leq \alpha \quad \text{if e-values are PRDS;} \quad (7)$$

$$\mathbb{E}[T(\alpha b_k E_k)] \leq \alpha \quad \text{otherwise,} \quad (8)$$

where the expectation and the probability are computed under the null distribution of  $E_k$ , which is sometimes known. In case of a composite null, (7) and (8) need to hold for all probability measures

in  $H_k$ , so an additional supremum over  $H_k$  must be appended to the left hand side. By Markov's inequality,  $b_k$  in (7) is larger or equal to that in (8), leading to a stronger boosting under PRDS. Since  $\mathbb{E}[E_k] \leq 1$  for a null e-value  $E_k$ , (7) and (8) always hold for  $b_k = 1$ , making  $b_k = 1$  always a safe choice even if there is no additional (marginal or joint) distributional information. There is no boosting, PRDS or not, if no marginal information on the distributions is available. In either case, define the boosted e-values

$$e'_k = b_k e_k$$

for  $k \in \mathcal{K}$ . On first reading, the reader may skip the next two remarks.

*Remark 4.1.* Note that the left-hand sides of (7) and (8) are increasing in  $b_k$ , and hence a practical value of  $b_k$  can be obtained by simply trying a few choices of  $b_k \geq 1$ . It would be ideal to find the largest  $b_k$  such that (7) or (8) becomes an equality. This is possible in some cases but it may not be possible if the null distribution of  $E_k$  is not continuous or not completely specified; see Examples 7.2, 7.3, 8.3 and 8.4, where we either get a precise value or an analytical approximation of the best  $b_k$ . Under an extra assumption, we have a simple formula (Proposition 8.2) for the best  $b_k$  in case e-values are PRDS.

*Remark 4.2.* Instead of using the standard method in (7) and (8) described above, the user can instead choose increasing functions  $\phi_1, \dots, \phi_K : [0, \infty] \rightarrow [0, \infty]$  such that

$$\max_{x \in K/\mathcal{K}} x \mathbb{P}(\alpha \phi_k(E_k) \geq x) \leq \alpha \quad \text{if e-values are PRDS;} \quad (9)$$

$$\mathbb{E}[T(\alpha \phi_k(E_k))] \leq \alpha \quad \text{otherwise,} \quad (10)$$

and define the boosted e-values  $e'_k = \phi_k(e_k)$  for  $k \in \mathcal{K}$ . The choice  $\phi_k : x \mapsto b_k x$  corresponds to the standard method explained earlier. Note that if e-values are PRDS, then so are the boosted e-values by any choice of boosting, since PRDS is invariant under monotone transforms.

## 5 An FDR guarantee via compliance

We are now ready to state our main result.

**Theorem 5.1.** *The e-BH procedure has FDR at most  $K_0\alpha/K$ , under any dependence structure of the e-values. In particular, the base e-BH procedure applied to raw e-values has FDR at most  $K_0\alpha/K$ .*

The full proof of Theorem 5.1 follows from Theorems 7.1 and 8.1 in the next few sections, which requires delicate technical treatment. Here, we provide a simple proof for the second statement of Theorem 5.1 that the base e-BH procedure has the desired FDR guarantee. This simple argument illustrates the advantages of working with e-values, and it applies to any ‘‘complaint’’ e-testing procedures including the base e-BH procedure. Following Su [23] who defined compliant p-testing procedures, an e-testing procedure  $\mathcal{G}$  is said to be *compliant at level  $\alpha \in (0, 1)$*  if, denoting by  $R_{\mathcal{G}}$  the number of rejections, every rejected e-value  $e_k$  satisfies

$$e_k \geq \frac{K}{\alpha R_{\mathcal{G}}}.$$

Using  $1/p_k$  in place of  $e_k$  above, we recover the definition from Su [23]. Clearly, the base e-BH procedure is compliant because of (5); moreover, the base e-BH procedure dominates all other compliant e-testing procedures.

**Proposition 5.2.** *Any compliant e-testing procedure at level  $\alpha$  has FDR at most  $\alpha K_0/K$  for arbitrary configurations of e-values.*

*Proof.* Let  $\mathcal{G}$  be a compliant e-testing procedure. Let  $\mathbf{E} = (E_1, \dots, E_K)$  be an arbitrary vector of e-variables fed to the testing procedure  $\mathcal{G}$ . The FDP of  $\mathcal{G}$  satisfies

$$\frac{F_{\mathcal{G}}}{R_{\mathcal{G}}} = \frac{|\mathcal{G}(\mathbf{E}) \cap \mathcal{N}|}{R_{\mathcal{G}} \vee 1} = \sum_{k \in \mathcal{N}} \frac{\mathbb{1}_{\{k \in \mathcal{G}(\mathbf{E})\}}}{R_{\mathcal{G}} \vee 1} \leq \sum_{k \in \mathcal{N}} \frac{\mathbb{1}_{\{k \in \mathcal{G}(\mathbf{E})\}} \alpha E_k}{K} \leq \sum_{k \in \mathcal{N}} \frac{\alpha E_k}{K}, \quad (11)$$

where the first inequality is due to compliance. As  $\mathbb{E}[E_k] \leq 1$  for  $k \in \mathcal{N}$ , we have

$$\mathbb{E} \left[ \frac{F_{\mathcal{G}}}{R_{\mathcal{G}}} \right] \leq \sum_{k \in \mathcal{N}} \mathbb{E} \left[ \frac{\alpha E_k}{K} \right] \leq \frac{\alpha K_0}{K},$$

thus the desired FDR guarantee.  $\square$

In contrast, general compliant p-testing procedures do not have the nice property in Proposition 5.2 even if p-values are independent. Under a condition slightly weaker than PRDS (see [23]), a compliant p-testing procedure at level  $\alpha$  has a weaker FDR guarantee  $\alpha(1 + \log(1/\alpha)) > \alpha$ . Hence, for an FDR guarantee of  $\alpha$ , the p-testing procedure needs to be compliant at level  $\alpha' < \alpha$  satisfying

$$\alpha' \left( 1 + \log \frac{1}{\alpha'} \right) = \alpha. \quad (12)$$

### The (compliant) structured, post-selection and grouped e-BH procedures

Without further structure, compliant p- or e-testing procedures are dominated by the base BH or e-BH procedure, and hence we do not directly apply them for a stand-alone data set of p-values or e-values. Nevertheless, they become useful in multi-armed bandit problems; see Jamieson and Jain [14] and our Section 11. Compliant procedures are also useful in structured settings. Suppose we would like to insist that the rejected set cannot be an arbitrary subset of  $\mathcal{K}$ , but must be one amongst a given class of sets  $\mathcal{S} \subseteq 2^{\mathcal{K}}$ . For example, if the hypotheses are structured as a graph, then we may choose  $\mathcal{S}$  to be all connected subgraphs. The appropriate modification of the e-BH procedure (“structured e-BH procedure”) would try to

find any set  $S \in \mathcal{S}$  such that every e-value in  $S$  is larger than  $\frac{K}{\alpha|S|}$ .

Importantly, solving for the *largest*  $S$  may not be computationally feasible, so one may wish to use an approximation algorithm (like a greedy heuristic that grows  $S$  by starting with the largest e-value and adding a few elements at a time to this set) to find *some*  $S \in \mathcal{S}$  such that every e-value in  $S$  is larger than  $\frac{K}{\alpha|S|}$ . Such a procedure would immediately satisfy an FDR guarantee, for any  $\mathcal{S}$  and dependence structure, by virtue of being compliant. In contrast, the “structured BH procedure” in Ramdas et al. [19, Section 8] requires the BY reshaping correction in order to maintain validity in general.

Another application of compliance is in post-processing e-values that have been screened or filtered based on the data. Formally, suppose the original set of hypotheses  $\mathcal{K}$  has been shrunk to a data-dependent subset  $S \subset \mathcal{K}$  in some arbitrary way based on the e-values. This could be based on any subjective choices of the user (based on side-information, prior scientific knowledge, whatever). The question is: how do we now find a subset of  $S$  that guarantees FDR control? The “post-selection e-BH procedure” is simple:

run the e-BH procedure on  $S$  at the amended level  $\tilde{\alpha} := \alpha|S|/K$ .

Then every rejected e-value will be larger than  $\frac{|S|}{\tilde{\alpha} R_{\mathcal{G}}}$ , which equals  $\frac{K}{\alpha R_{\mathcal{G}}}$ . Thus, this procedure controls FDR under arbitrary dependence for any initial selection procedure. This is once more in contrast to the “post-selection BH procedure” in Ramdas et al. [19, Section 8], where conditions on both the dependence and the selection are needed for FDR control, devoid of which one must use the BY reshaping correction.

Combining the above two ideas (imposition of structural constraints and screening/filtering) in any order immediately lends a great deal of flexibility to the user to design new procedures. For example, it immediately yields a “focused e-BH” procedure, an extension of the recent focused BH algorithm by Katsevich et al. [17], whose details are omitted for brevity. In contrast to [17, Theorem 1], our procedure controls FDR under no assumptions whatsoever on the dependence, structure and filtering.

As a final extension, imagine that we are given a partitioning of the hypotheses into  $G$  prespecified groups. It is easy to define group-level e-values by merging the individual level e-values (e.g., averaging under arbitrary dependence, or multiplying under independence). This immediately yields a group-level e-BH procedure that can be used as an initial filtering step before running the post-selection e-BH procedure. Finally, an e-filter extension of the p-filter algorithm [19] is also possible.

### A few clarifying remarks on the e-BH procedure

At this point, we have seen that the base e-BH procedure has the desirable FDR control. Let us make a few remarks on the (full) e-BH procedure to further explain its features and subtleties.

1. Regardless of any available information, one can always opt to skip the boosting step, and directly apply the base e-BH method. This gives a valid FDR of  $\alpha$  under arbitrary setup of the e-values and arbitrary dependence, although it may cause a loss of detection power, as the boosting factor  $b_k$  may be substantial (see Examples 7.2, 7.3, 8.3 and 8.4); a boosting factor between 1.3 and 10 is common in stylized settings.
2. The base e-BH procedure is precisely the BH procedure applied to the p-values  $1/e_1, \dots, 1/e_K$  (truncated to one if needed; such a truncation does not affect (5)). Therefore, if the e-values are independent (or PRDS), then by Theorem 3.2, the FDR of the base e-BH procedure is at most  $\alpha$ . This result is not of primary interest, but is worthy of note. The main advantage of e-BH is that, unlike the BH procedure, the above FDR guarantee holds under arbitrary dependence, where a conversion to p-values would be suboptimal due to the extra  $\approx \log K$  correction factor. As a consequence, for p-values derived from e-values (for example via universal inference [27]), the BH procedure does not need any correction even under arbitrary dependence.
3. The boosting in step 1 of the BH procedure can easily incorporate partial information on distributions of the e-values. If we know a set of null possible distributions of  $E_k$ , then it suffices to calculate a boosting factor under each distribution, and take their infimum. Moreover, if only the null distributions of some e-values are known, we can simply choose  $b_k = 1$  for the ones whose distributions we do not know. Note that the post-selection and structured e-BH procedures also enjoy FDR control after boosting under arbitrary dependence.
4. Since  $b_k$  may be different for each  $k$ , the boosted e-values may not have the same order as the raw e-values. Hence, the full e-BH procedure may reject a hypothesis with a smaller e-value while accepting one with a larger e-value. This feature is intentional. For instance, an observation of  $e_k = 1.999$  from a uniform distribution on  $[0, 2]$  carries stronger evidence against the null hypothesis, compared to an observation of  $e_k = 2$  from an exponential distribution with mean 1, or from an unknown distribution.
5. In case both p-values and e-values are available for the hypotheses, we typically do not expect the e-BH procedure to beat the BH procedure. This is because e-values typically contain less information than p-values (the flip side of needing less assumptions/structure/knowledge/modeling in order to be constructed), and hence a simplistic comparison with the BH procedure is not particularly insightful. Nevertheless, if the distributions of null e-values are fully known, then the e-BH procedure performs comparably to the BH and the BY procedures; see Section 10 for their connection and Section 11 and Appendix A for simulation results.

6. The base BH procedure has FDR control  $\alpha$  under PRDS, and in the absence of such assumption, a penalty ( $\ell_K$  in the BY procedure, or a milder penalty if some other information is available, e.g., [10]) needs to be applied. The e-BH procedure is the opposite: it starts with FDR control  $\alpha$  with no assumption, and a boosting can be applied to improve power if further information on the distributions of e-values and dependence is available. Similarly, a penalty in (12) needs to be applied for compliant p-testing procedures, which are dominated by BH, but not for compliant e-testing procedures.

In practice, it may be desirable to assign weights based on prior knowledge of each hypothesis. Here, we describe an analog of the weighted BH procedure [4]. For arbitrary constant weights  $w_1, \dots, w_K$  summing up to  $K$ , a similar boosting scheme is obtained by each replacing  $\alpha$  with  $w_k \alpha$  in (7) and (8). We can alternatively skip the boosting in step 1 and directly use  $w_1 E_1, \dots, w_K E_K$  as the input e-values fed to the base e-BH procedure. In either case, the full e-BH procedure still has the valid FDR control of  $\alpha$ , under arbitrary dependence of the e-values as before. This claim is also justified by Theorems 7.1 and 8.1 (see the explanation after Theorem 7.1).

At this point, most (but not all) messages for the practitioner have been delivered and the casual reader may be warned that the paper intentionally takes a rather theoretical turn, and analyzes the e-BH procedure more carefully, comparing repeatedly to the BH procedure or to the case when e-values are simply calibrated p-values. Proofs of all results beyond this section are put in Appendix B.

## 6 A technical, but interpretable, lemma

While we have provided a direct and simple proof for FDR control of the base e-BH that does not require any correction under arbitrary dependence, the reader may be curious to understand better why exactly BH has to pay an extra  $\approx \log K$  factor, but e-BH, on the contrary, allows for boosting. We provide one technical answer here in the form of a lemma, similar to Ramdas et al. [19, Lemma 1] in the setting of p-values. Lemma 6.1 leads to FDR control of the full e-BH procedure applied to boosted e-values.

Recall that for  $k \in \mathcal{K}$ , a random vector  $\mathbf{X} = (X_1, \dots, X_K)$  is PRDS on  $X_k$  if for any decreasing set  $A \subseteq \mathbb{R}^K$ , the function  $x \mapsto \mathbb{P}(\mathbf{X} \in A \mid X_k \geq x)$  is decreasing on  $[0, \infty)$ .

In order to prepare for the lemma that follows, note that for any positive constant  $c$ , we have that

$$\mathbb{E}[c \mathbb{1}_{\{X_k \geq c\}}] \leq \mathbb{E}[X_k], \quad (13)$$

which is simply a restatement of Markov's inequality. The following lemma addresses the situation where  $c$  is a random data-dependent quantity  $f(\mathbf{X})$ , as in the computation of FDR of the e-BH procedure.

**Lemma 6.1.** *Take an arbitrary random vector  $\mathbf{X} = (X_1, \dots, X_K) : \Omega \rightarrow [0, \infty]^K$  and fix  $k \in \mathcal{K}$ . Let  $f : [0, \infty]^K \rightarrow [0, \infty)$  be a measurable function with range  $I_f$ , and  $X'_k := \sup\{x \in I_f \cup \{0\} : x \leq X_k\} \leq X_k$ .*

(i) *If  $X_k$  is independent of  $\mathbf{X}^{-k} := (X_j)_{j \neq k}$ , then*

$$\mathbb{E}[f(\mathbf{X}) \mathbb{1}_{\{X_k \geq f(\mathbf{X})\}} \mid \mathbf{X}^{-k}] \leq \mathbb{E}[X'_k] \leq \mathbb{E}[X_k].$$

(ii) *If  $f$  is decreasing and  $\mathbf{X}$  satisfies PRDS on  $X_k$ , then*

$$\mathbb{E}[f(\mathbf{X}) \mathbb{1}_{\{X_k \geq f(\mathbf{X})\}}] \leq \sup_{x \geq 0} x \mathbb{P}(X'_k \geq x) \leq \mathbb{E}[X_k].$$

(iii) *For any dependence structure, it always holds that*

$$\mathbb{E}[f(\mathbf{X}) \mathbb{1}_{\{X_k \geq f(\mathbf{X})\}}] \leq \mathbb{E}[X'_k] \leq \mathbb{E}[X_k].$$

In particular, if  $X_k$  is a null e-value, all expectations above are bounded by 1.

Later we will apply Lemma 6.1 to  $\mathbf{X} = \mathbf{E}'$  and a specific choice of  $f$  to show Theorems 7.1 and 8.1.

Comparing Lemma 6.1 to Ramdas et al. [19, Lemma 1] for the p-value setting, we observe that the first two statements have a very direct parallel. However the third statements are where the difference appears; the p-value statement requires a correction while the above e-value statement does not. To briefly elaborate, recall (13) and correspondingly note that a null p-variable  $P_k$  satisfies for any constant  $c \in (0, 1]$ ,

$$\mathbb{E} \left[ \frac{\mathbb{1}_{\{P_k \leq c\}}}{c} \right] \leq 1. \quad (14)$$

While Lemma 6.1 showed that (13) easily generalized from a constant  $c$  to a data-dependent  $f(X)$ , Ramdas et al. [19, Lemma 1] showed that a corresponding statement for (14) is true for p-variables only under a PRDS assumption, but otherwise an extra  $\ell_K$  factor is paid under arbitrary dependence.

We anticipate this lemma to aid with the design and proof of FDR procedures with e-values in other contexts—just as the corresponding p-value lemma already has [18]—and we shall see how it is applied to e-BH next.

## 7 FDR of e-BH for arbitrary e-values

We analyze some properties of the base e-BH procedure. Below, the e-values we feed into the base e-BH procedure are the boosted e-values from step 1 of the BH procedure. We will omit the term “boosted” while keeping in mind that the base e-BH can also be directly applied to the raw e-values. Let  $(E'_1, \dots, E'_K)$  be the random vector of (boosted) e-values and  $(e'_1, \dots, e'_K)$  be its realized value. The only property we will use on  $E'_1, \dots, E'_K$  is that they are nonnegative.

For  $t \in [0, \infty)$ , let  $f(t)$  be the number of null cases with an e-value  $e'_k$  larger than or equal to  $t$ , and  $R(t)$  be the total number of e-values larger than or equal to  $t$ , but it is set to 1 if there is no such e-value. Define the quantity

$$Q(t) = tR(t)/K,$$

and let

$$t_\alpha = \inf\{t \in [0, \infty) : Q(t) \geq 1/\alpha\}.$$

Clearly  $t_\alpha \in [1/\alpha, K/\alpha]$  since  $Q(t) \leq t$  and  $R(t) \geq 1$ . Since  $Q$  only has downside jumps and  $Q(0) = 0$ , we know  $Q(t_\alpha) = 1/\alpha$ , and thus

$$t_\alpha R(t_\alpha) = \frac{K}{\alpha}. \quad (15)$$

From (15),  $\alpha t_\alpha$  takes values in the set  $K/\mathcal{K}$ . We will use the notation  $t_\alpha$  for both the random variable, which is a decreasing function of  $(E'_1, \dots, E'_K)$  by (15), and its realization, which is the same function applied to  $(e'_1, \dots, e'_K)$ .

If  $e'_{[k]} < t_\alpha$ , then by definition of  $Q$ , we have

$$\frac{ke'_{[k]}}{K} \leq \frac{R(e'_{[k]})e'_{[k]}}{K} = Q(e'_{[k]}) < 1/\alpha.$$

Thus, each  $H_k$  is rejected by  $\mathcal{G}(\alpha)$  if and only if  $e'_k \geq t_\alpha$ . Therefore, we obtain that the FDR equals

$$\mathbb{E} \left[ \frac{F_{\mathcal{G}(\alpha)}}{R_{\mathcal{G}(\alpha)}} \right] = \mathbb{E} \left[ \frac{f(t_\alpha)}{R(t_\alpha)} \right] = \mathbb{E} \left[ \frac{t_\alpha f(t_\alpha)}{KQ(t_\alpha)} \right] = \frac{\alpha}{K} \mathbb{E} [t_\alpha f(t_\alpha)]. \quad (16)$$

Next, we give a more precise analysis of the FDR of the base e-BH procedure under arbitrary dependence. Similarly to the observation made by Vovk and Wang [25], arbitrarily dependent e-values are nicer to work with than independent e-values, in sharp contrast to the case of p-values.

**Theorem 7.1.** Applied to arbitrary non-negative random variables  $E'_1, \dots, E'_K$  and  $\alpha \in (0, 1)$ , the base e-BH procedure  $\mathcal{G}(\alpha)$  satisfies

$$\mathbb{E} \left[ \frac{F_{\mathcal{G}(\alpha)}}{R_{\mathcal{G}(\alpha)}} \right] \leq \frac{K_0}{K} y_\alpha,$$

where

$$y_\alpha = \frac{1}{K_0} \sum_{k \in \mathcal{N}} \mathbb{E} [T(\alpha E'_k)]. \quad (17)$$

In particular, if  $E'_1, \dots, E'_K$  are the raw e-values or the boosted e-values via (8) or (10), then  $y_\alpha \leq \alpha$ .

Let  $(E_1, \dots, E_K)$  be the vector of raw e-values. Theorem 7.1 shows that for any choice of  $b_1 E_1, \dots, b_K E_K$  satisfying (8) or  $\phi_1(E_1), \dots, \phi_K(E_K)$  satisfying (10), the FDR of the e-BH procedure is at most  $\alpha$  under arbitrary dependence. Moreover, the FDR control of the weighted e-BH follows from

$$\frac{K_0}{K} y_\alpha = \frac{1}{K} \sum_{k \in \mathcal{N}} \mathbb{E} [T(\alpha E'_k)] \leq \frac{1}{K} \sum_{k \in \mathcal{N}} w_k \alpha \leq \frac{1}{K} \sum_{k \in \mathcal{K}} w_k \alpha = \alpha.$$

The value of  $y_\alpha$  depends on the distribution of the boosted e-values as well as the number  $K$  of hypotheses. In case the null distributions of the e-values are known, we would ideally set  $y_\alpha$  to  $\alpha$  (or close to  $\alpha$ ) by properly choosing  $b_k$  in (8) or  $\phi_k$  in (10).

In case  $\mathbb{E}[T(\alpha b_k E_k)]$  is not easy to compute, it might be convenient to use a weaker bound

$$\bar{y}_{\alpha,k}(b_k) := \mathbb{E}[\alpha b_k E_k \mathbf{1}_{\{\alpha b_k E_k \geq 1\}}] \geq \mathbb{E}[T(\alpha b_k E_k)], \quad (18)$$

and set  $\bar{y}_{\alpha,k}(b_k) \leq \alpha$  (ideally an equality) by choosing  $b_k \geq 1$ . Such a choice of  $b_k$  always guarantees the FDR of the e-BH procedure to be at most  $\alpha$  by Theorem 7.1, since

$$\frac{1}{K_0} \sum_{k \in \mathcal{N}} \mathbb{E}[T(\alpha b_k E_k)] \leq \frac{1}{K_0} \sum_{k \in \mathcal{N}} \bar{y}_{\alpha,k}(b_k) \leq \alpha.$$

In addition to being easier to compute, an advantage of  $\bar{y}_{\alpha,k}(b_k)$  is that it depends purely on the distribution of  $E_k$  and not on  $K$ , and hence the boosted e-value  $b_k E_k$  is ready to use for other experiments.

**Example 7.2.** We illustrate  $y_\alpha$  and  $\bar{y}_{\alpha,k}(b_k)$  with a popular class of calibrators in Shafer [20] and Vovk and Wang [25]. Take  $\lambda \in (0, 1)$  and assume the raw e-values are given by, for  $k \in \mathcal{K}$ ,

$$E_k = \lambda P_k^{\lambda-1}, \quad (19)$$

where  $P_k$  is a uniform random variable on  $[0, 1]$  if  $k \in \mathcal{N}$ . We will consider the boosted e-values  $bE_1, \dots, bE_K$  where  $b$  is a common boosting factor since the null e-values are identically distributed. In this case,

$$y_\alpha = \mathbb{E}[T(\alpha b E_1)] = \int_0^1 T(\alpha b \lambda u^{\lambda-1}) du.$$

A formula for  $\bar{y}_{\alpha,k}(b)$  is simple:

$$\bar{y}_{\alpha,k}(b) = \alpha b \int_0^{(\alpha b \lambda)^{1/(1-\lambda)}} \lambda u^{\lambda-1} du = (\lambda^\lambda \alpha b)^{1/(1-\lambda)}.$$

For instance, if  $\lambda = 1/2$ , then  $\bar{y}_{\alpha,k}(b) = (\alpha b)^2/2$ . Setting  $\bar{y}_{\alpha,k}(b) = \alpha$  yields  $b = (2/\alpha)^{1/2}$ . In this example, all e-values are boosted by a multiplier of  $(2/\alpha)^{1/2}$ , which is substantial; e.g.,  $b \approx 6.32$  if  $\alpha = 0.05$ . Finally, note that the same boosting factor  $b$  above is also valid whenever the p-values are not exactly uniform but instead stochastically larger than uniform.



**Example 7.3.** We consider e-values obtained from the likelihood ratio between two normal distributions with different mean and variance 1 as used in the numerical experiment of Vovk and Wang [25]. Take  $\delta > 0$  which represents the difference in the alternative and the null means, and assume the raw e-values are given by, for  $k \in \mathcal{K}$ ,

$$E_k = e^{\delta X_k - \delta^2/2} \quad (20)$$

where  $X_k$  is a standard normal random variable if  $k \in \mathcal{N}$ . Note that each null e-value is a log-normal random variable with parameter  $(-\delta^2/2, \delta)$ . For the boosted e-values  $bE_1, \dots, bE_K$ , We have

$$\bar{y}_{\alpha,k}(b) = \alpha b \mathbb{E}[E_k \mathbf{1}_{\{E_k \geq 1/(\alpha b)\}}] = \alpha b \Phi \left( \frac{\delta}{2} + \frac{\log(\alpha b)}{\delta} \right),$$

where  $\Phi$  is the standard Gaussian cdf. Setting  $\bar{y}_{\alpha,k}(b) = \alpha$  yields the equation

$$b \Phi \left( \frac{\delta}{2} + \frac{\log(\alpha b)}{\delta} \right) = 1,$$

which can be easily solved numerically. For instance, if  $\delta = 3$  and  $\alpha = 0.05$ , then  $b \approx 1.37$  by solving  $\bar{y}_{\alpha,k}(b) = \alpha$ , and if  $\delta = 4$  and  $\alpha = 0.05$ , then  $b \approx 1.11$ . These choices of  $b$  are all conservative since  $\bar{y}_{\alpha,k}(b)$  is a conservative bound for  $y_\alpha$  (but  $b$  is only slightly conservative if  $K$  is large).

For raw e-values  $(E_1, \dots, E_K)$  with unspecified distributions, Theorem 7.1 gives an upper bound  $\alpha K_0/K$  on the FDR of the base e-BH procedure. This upper bound is usually quite loose since  $y_\alpha$  in (17) is typically much smaller than  $\alpha$ ; see e.g., Example 7.2. This upper bound cannot be improved in general without any additional information.

**Example 7.4** (Sharpness of the upper FDR bound in Theorem 7.1). Consider the following setup. Let the  $K_0$  null raw e-values be given by  $E_k = K/(K_0\alpha) \mathbf{1}_A$ ,  $k \in \mathcal{N}$ , where  $A$  is an event with  $\mathbb{P}(A) = \alpha K_0/K$ . Moreover, we set all other e-values to 0, which means that they cannot be rejected. Hence, the false discovery proportion is 1 as soon as there is any rejection. It follows that

$$\text{FDR}_{\mathcal{G}(\alpha)} = \mathbb{P}(R_{\mathcal{G}(\alpha)} > 0) = \mathbb{P}(A) = \alpha K_0/K,$$

which is the upper bound provided by Theorem 7.1.

## 8 FDR of e-BH for PRDS e-values

In case p-values are independent or positively dependent, the BH procedure has a lower FDR guarantee than the one obtained with arbitrary dependence.

We investigate a similar matter for e-values, aiming for a bound better than that of Theorem 7.1. This is possible via the result of Lemma 6.1. To summarize, if e-values are PRDS, the base e-BH procedure has a smaller FDR guarantee, and, consequently, the e-BH procedure allows a more powerful boosting in its step 1.

**Theorem 8.1.** *Suppose that the raw e-values are PRDS. Applied to arbitrary boosted e-values  $(E'_1, \dots, E'_K)$  and  $\alpha \in (0, 1)$ , the base e-BH procedure  $\mathcal{G}(\alpha)$  satisfies*

$$\mathbb{E} \left[ \frac{F_{\mathcal{G}(\alpha)}}{R_{\mathcal{G}(\alpha)}} \right] \leq \frac{K_0}{K} z_\alpha. \quad (21)$$

where

$$z_\alpha = \frac{1}{K_0} \sum_{k \in \mathcal{N}} \max_{x \in \mathcal{K}/\mathcal{K}} x \mathbb{P}(\alpha E'_k \geq x). \quad (22)$$

In particular, if  $E'_1, \dots, E'_K$  are the raw e-values or the boosted e-values via (7) or (9), then  $z_\alpha \leq \alpha$ .

Theorem 8.1 shows that, under the assumption of PRDS, using the boosted e-values  $b_1 E_1, \dots, b_K E_K$  satisfying (7) or  $\phi_1(E_1), \dots, \phi_K(E_K)$  satisfying (9), the FDR of the e-BH procedure is at most  $\alpha$ . For  $b_k \geq 1$ , let

$$z_{\alpha,k}(b_k) = \max_{x \in K/\mathcal{K}} x \mathbb{P}(\alpha b_k E_k \geq x). \quad (23)$$

In step 1 of the e-BH procedure, we need to choose a boosting factor  $b_k \geq 1$  such that  $z_{\alpha,k}(b_k) \leq \alpha$  (ideally, an equality). Under an extra condition, a suitable choice of  $b_k$  admits a simple formula.

**Proposition 8.2.** *Suppose that  $E_k$  is a continuously distributed null e-value. Let  $q_{1-\alpha}(E_k)$  be the left  $(1-\alpha)$ -quantile of  $E_k$ . If*

$$t \mapsto t \mathbb{P}(E_k \geq t) \text{ is decreasing on } [q_{1-\alpha}(E_k), \infty), \quad (24)$$

then  $b_k := (\alpha q_{1-\alpha}(E_k))^{-1}$  is the largest boosting factor which satisfies (7).

Condition (24) is not uncommon as  $\alpha$  is typically small, and the value  $t \mathbb{P}(E_k \geq t)$  goes to 0 as  $t \rightarrow \infty$  since  $E_k$  has a finite mean. Condition (24) is satisfied by, for instance, the e-values in Example 7.2 as well as their distributional mixtures. For a continuously distributed  $E_k$ , condition (24) can be equivalently expressed as

$$\beta \mapsto \beta q_{1-\beta}(E_k) \text{ is decreasing on } [1-\alpha, 1).$$

Similar to  $\bar{y}_{\alpha,k}$  in (18), there is a conservative version of  $z_{\alpha,k}$  which does not depend on  $K$ , given by

$$\bar{z}_{\alpha,k}(b_k) = \sup_{x \geq 1} x \mathbb{P}(\alpha b_k E_k \geq x). \quad (25)$$

In particular, if (24) holds, then  $\bar{z}_{\alpha,k}(b_k) = z_{\alpha,k}(b_k)$  for the choice  $b_k$  in Proposition 8.2.

**Example 8.3.** We look at the e-values in Example 7.2. For  $k \in \mathcal{N}$ , we have

$$t \mathbb{P}(E_k \geq t) = t(\lambda/t)^{1/(1-\lambda)} = (\lambda t^{-\lambda})^{1/(1-\lambda)},$$

which is a decreasing function in  $t$ , and thus (24) holds. In this case, for  $b \geq 1$ ,

$$z_{\alpha,k}(b) = \mathbb{P}(\alpha b E_k \geq 1) = (\lambda \alpha b)^{1/(1-\lambda)},$$

and by Proposition 8.2, the best choice of  $b_k$  is

$$b = (\alpha q_{1-\alpha}(E_k))^{-1} = (\lambda \alpha^\lambda)^{-1}.$$

Since  $\bar{y}_{\alpha,k}(b) = (\lambda^\lambda \alpha b)^{1/(1-\lambda)}$  in Example 7.2, we get  $z_{\alpha,k}(b)/\bar{y}_{\alpha,k}(b) = \lambda$ . Hence, the FDR is improved by a factor of roughly  $\lambda$  from Theorem 7.1 (arbitrary dependence) to Theorem 8.1 (PRDS), which could be substantial if  $\lambda$  is small. For instance, if  $\lambda = 1/2$ , then  $b = 2\alpha^{-1/2}$ . For  $\alpha = 0.05$ , we have  $b \approx 8.94$ , which should be compared with  $b \approx 6.32$  in Example 7.2 under arbitrary dependence.

**Example 8.4.** For the setup of e-values in Example 7.3, (24) does not always hold. Nevertheless, for  $b \geq 1$ ,  $z_{\alpha,k}(b)$  and  $\bar{z}_{\alpha,k}(b)$  in (23) and (25) have simple formulas

$$z_{\alpha,k}(b) = \max_{x \in K/\mathcal{K}} x \Phi \left( \frac{\log(\alpha b)}{\delta} - \frac{\log x}{\delta} - \frac{\delta}{2} \right),$$

and

$$\bar{z}_{\alpha,k}(b) = \max_{x \geq 1} x \Phi \left( \frac{\log(\alpha b)}{\delta} - \frac{\log x}{\delta} - \frac{\delta}{2} \right).$$

Without specifying  $K$ , we use  $\bar{z}_{\alpha,k}(b)$ , and the value  $b$  for  $\bar{z}_{\alpha,k}(b) = \alpha$  can be easily computed numerically. For instance, if  $\delta = 3$  and  $\alpha = 0.05$ , then  $b \approx 7.88$  (compared with  $b \approx 1.37$  in Example 7.3), and if  $\delta = 4$  and  $\alpha = 0.05$ , then  $b \approx 10.31$  (compared with  $b \approx 1.11$  in Example 7.3).

## 9 An optimality result

There is a simple way of generating other e-testing procedures similar to the base e-BH procedure by transforming the e-values. In this section, working under the assumption that no distributional information on e-values is available, we focus on a procedure resulted from applying a common transform to all raw e-values. We obtain an optimality of the base e-BH procedure among all such transforms.

Take a strictly increasing and continuous function  $\phi : [0, \infty] \rightarrow [0, \infty]$  with  $\phi(\infty) = \infty$  and  $\phi(0) < 1$ . We shall call  $\phi$  an *increasing transform*. Define

$$\tau_\phi = \inf\{\tau \geq 0 : \phi(\tau)R(\tau) \geq K\}.$$

Since  $R$  only has downside jumps, and  $\phi(0)R(0) < K$ , we have  $\phi(\tau_\phi)R(\tau_\phi) = K$ , and  $\phi(\tau_\phi) \in K/\mathcal{K}$ .

We design an e-testing procedure  $\mathcal{G}(\phi)$  by rejecting all hypotheses with an e-value larger than or equal to  $\tau_\phi$ . Alternatively, we can describe  $\mathcal{G}(\phi)$  in the following way, more similar to the base e-BH procedure: we reject  $k_{e,\phi}^*$  hypotheses with the largest e-values, where

$$k_{e,\phi}^* = \max \left\{ k \in \mathcal{K} : \frac{k\phi(e_{[k]})}{K} \geq 1 \right\}. \quad (26)$$

It is clear that the choice of  $\phi : t \mapsto \alpha t$  corresponds to  $\tau_\phi = t_\alpha$  in Section 4, which yields the base e-BH procedure  $\mathcal{G}(\alpha)$ .

The FDR of  $\mathcal{G}(\phi)$  can be calculated as

$$\mathbb{E} \left[ \frac{F_{\mathcal{G}(\phi)}}{R_{\mathcal{G}(\phi)}} \right] = \mathbb{E} \left[ \frac{f(\tau_\phi)}{R(\tau_\phi)} \right] = \mathbb{E} \left[ \frac{\phi(\tau_\phi)f(\tau_\phi)}{K} \right] = \frac{1}{K} \sum_{k \in \mathcal{N}} \mathbb{E}[\phi(\tau_\phi) \mathbf{1}_{\{E_k \geq \tau_\phi\}}]. \quad (27)$$

Hence, using the fact that  $\phi(\tau_\phi) \in K/\mathcal{K}$  and Lemma 6.1, we obtain

$$\mathbb{E} \left[ \frac{F_{\mathcal{G}(\phi)}}{R_{\mathcal{G}(\phi)}} \right] \leq \frac{K_0}{K} y_\phi, \quad (28)$$

where

$$y_\phi := \frac{1}{K_0} \sum_{k \in \mathcal{N}} \mathbb{E}[T(\phi(E_k))].$$

The FDR control (28) also follows from Theorem 7.1 by choosing the boosted e-values as  $\phi(E_k)/\alpha$ ,  $k \in \mathcal{K}$ .

The special choice  $\phi : t \mapsto \alpha t$ , corresponding to the e-BH procedure, always guarantees  $y_\phi = y_\alpha \leq \alpha$ , where  $y_\alpha$  is defined in (17). For a general choice of  $\phi$ , the FDR of  $\mathcal{G}(\phi)$  requires the knowledge of  $y_\phi$ , or at least an upper bound.

In the next result, we show that, if we require an FDR guarantee for arbitrary e-values, then the base e-BH procedure is optimal among the class  $\mathcal{G}(\phi)$ . This is reassuring for us to use the base e-BH procedure as a canonical candidate for handling arbitrarily dependent e-values without further information.

**Theorem 9.1.** *Fix  $\alpha \in (0, 1)$  and  $K$ . For any increasing transform  $\phi$ , if  $\mathcal{G}(\phi)$  satisfies*

$$\mathbb{E} \left[ \frac{F_{\mathcal{G}(\phi)}}{R_{\mathcal{G}(\phi)}} \right] \leq \alpha$$

*for arbitrary configurations of e-values, then  $\mathcal{G}(\phi) \subseteq \mathcal{G}(\alpha)$ .*

## 10 Applying the e-BH procedure to p-values

It would be interesting to comparing the BY procedure

$$\mathcal{D}'(\alpha) := \mathcal{D}(\alpha/\ell_K)$$

with the e-BH procedure  $\mathcal{G}(\alpha)$ , since both require no assumptions on the dependence structure. Although the two procedures are comparable on the dependence assumption, we remark that e-values generally use less information than p-values, as p-values require full specification of the distributions of the test statistics, whereas e-values only require the information of a known mean. Note that information on the boosting factors  $b_k$ ,  $k \in \mathcal{K}$  is contained in the null distribution of the test statistics.

To properly compare the p- and e-based procedures, we need to calibrate between p-values and e-values using calibrators. With such a calibration, the e-BH procedure gives rise to a general class of p-testing procedures with FDR guarantee for arbitrary p-values, similar to  $\mathcal{G}(\phi)$  in Section 9, and this class includes the base BH and the BY procedures as special cases. On a related note, calibration to e-values serves as a crucial intermediate step in combining p-values under arbitrary dependence for testing a global null [26, Section 5].

Let  $\psi : [0, 1] \rightarrow [0, \infty]$  be a strictly decreasing and continuous function with  $\psi(0) = \infty$ . We shall call  $\psi$  a *decreasing transform*. Define

$$\tau_\psi = \sup\{\tau \in [0, 1] : \psi(\tau)R(\tau) \geq K\}, \quad (29)$$

where  $R(t)$  is the number of p-values no larger than  $t \in [0, 1]$  (and set to 1 if there is no such p-value). We design a p-testing procedure  $\mathcal{D}(\psi)$  by rejecting all hypotheses with a p-value no larger than  $\tau_\psi$ . Alternatively,  $\mathcal{D}(\psi)$  can be described as rejecting the  $k_\psi^*$  hypotheses with the smallest p-values, where

$$k_\psi^* = \max\left\{k \in \mathcal{K} : \psi(p_{(k)}) \geq \frac{K}{k}\right\},$$

with the convention  $\max(\emptyset) = 0$ . A smaller transform function leads to less power of the testing procedure. A similar procedure using different transforms on individual p-values is analyzed in Appendix C.

An important example of  $\mathcal{D}(\psi)$  is to choose  $\psi : p \rightarrow \alpha/p$  for some  $\alpha \in (0, 1)$ . Note that

$$\psi(p_{(k)}) \geq \frac{K}{k} \iff \frac{Kp_{(k)}}{k} \leq \alpha.$$

In this case,  $\mathcal{D}(\psi)$  is precisely the BH procedure  $\mathcal{D}(\alpha)$ .

The p-testing procedure  $\mathcal{D}(\psi)$  is equal to the step-up procedure of Benjamini and Yekutieli [5] which rejects  $k^* := \max\{k \in \mathcal{K} : p_{(k)} \leq \alpha_k\}$  hypotheses with the smallest p-values, where  $\alpha_k = \psi^{-1}(K/k)$ ,  $k \in \mathcal{K}$ . From here, it is clear that only the values of  $\psi^{-1}$  on  $K/\mathcal{K}$  matter for  $\mathcal{D}(\psi)$ . Our main purpose here is to compare the p- and e-testing procedures, not to propose new step-up p-testing procedures.

The objects  $\phi(p_1), \dots, \phi(p_K)$  can be treated as boosted e-values. Hence, using Theorems 7.1 and 8.1, we can easily calculate the FDR of  $\mathcal{D}(\psi)$ .

**Proposition 10.1.** *For arbitrary p-values and a decreasing transform  $\psi$ , the testing procedure  $\mathcal{D}(\psi)$  satisfies*

$$\mathbb{E} \left[ \frac{F_{\mathcal{D}(\psi)}}{R_{\mathcal{D}(\psi)}} \right] \leq \frac{K_0}{K} y_\psi,$$

where

$$y_\psi := \psi^{-1}(1) + \sum_{j=1}^{K-1} \frac{K}{j(j+1)} \psi^{-1}(K/j). \quad (30)$$

If the  $p$ -values are PRDS, then

$$\mathbb{E} \left[ \frac{F_{\mathcal{D}(\psi)}}{R_{\mathcal{D}(\psi)}} \right] \leq \frac{K_0}{K} z_\psi,$$

where

$$z_\psi := \max_{t \in K/\mathcal{K}} t\psi^{-1}(t). \quad (31)$$

If the decreasing transform  $\psi$  satisfies

$$t \mapsto t\psi^{-1}(t) \text{ is decreasing on } [1, \infty), \quad (32)$$

similarly to condition (24), we have  $z_\psi = \psi^{-1}(1)$ . We can replace  $[1, \infty)$  by  $[1, K]$  or  $K/\mathcal{K}$  in (32); the current condition (32) is slightly stronger but it does not depend on  $K$ .

For the specific choice  $\psi : p \rightarrow \alpha/p$  which gives the BH procedure, we have  $\psi^{-1} = \psi$ ,

$$y_\psi = \alpha + \sum_{j=1}^{K-1} \frac{\alpha}{j+1} = \alpha \ell_K,$$

and

$$z_\psi = \max_{t \in K/\mathcal{K}} t\psi^{-1}(t) = \alpha.$$

By Proposition 10.1,  $\mathcal{D}(\psi) = \mathcal{D}(\alpha)$  has FDR guarantee  $K_0 \ell_K \alpha / K$  for arbitrarily dependent  $p$ -values, and FDR guarantee  $K_0 \alpha / K$  for PRDS  $p$ -values. This gives an analytical proof of the FDR guarantee of the BH and BY procedures in Theorem 3.2.

The inequality  $\alpha \mathbb{E}[t_\alpha f(t_\alpha)] \leq K_0 z_\alpha$  in the proof of Theorem 8.1 is generally not an equality. Hence, the FDR bound provided by Proposition 10.1 may not be sharp in the case of PRDS  $p$ -values. Nevertheless, in the next proposition we shall see that this bound is almost sharp under some extra conditions. From there, we obtain a weak optimality result on the base BH procedure.

**Proposition 10.2.** Fix  $\alpha \in (0, 1)$  and  $K$ . For any decreasing transform  $\psi$ , if  $\mathcal{D}(\psi)$  satisfies

$$\mathbb{E} \left[ \frac{F_{\mathcal{D}(\psi)}}{R_{\mathcal{D}(\psi)}} \right] \leq \alpha$$

for arbitrary configurations of PRDS  $p$ -values, then  $\psi^{-1}(1) \leq \alpha$ . Moreover, if  $\psi$  satisfies (32), then  $\mathcal{D}(\psi) \subseteq \mathcal{D}(\alpha)$ .

Comparing the optimality of the e-BH procedure in Theorem 9.1 with the optimality of the BH procedure in Proposition 10.2, we note two differences: Theorem 9.1 is stated for arbitrary  $e$ -values whereas Proposition 10.2 is stated for PRDS  $p$ -values; Theorem 9.1 imposes no assumption on  $\phi$  whereas Proposition 10.2 requires  $\psi$  to satisfy (32).

Different from the case of PRDS  $p$ -values,  $\mathcal{D}'(\alpha)$  and  $\mathcal{D}(\psi)$  may not be strictly comparable for arbitrarily dependent  $p$ -values. We discuss this issue in Example 10.3 below. The general message is that, in contrast to the case of PRDS  $p$ -values as shown in Proposition 10.2, the BY procedure is not necessarily always the best for arbitrarily dependent  $p$ -values.

**Example 10.3.** We consider the calibrators in (19) by choosing  $\psi : p \mapsto \beta \lambda p^{\lambda-1}$  for some  $\beta > 0$ . For simplicity, we take  $\lambda = 1/2$ . As we see from Example 7.2,  $\mathcal{D}(\psi)$  has an FDR guarantee of  $y_\psi \leq \beta^2/2$ . To compare with  $\mathcal{D}'(\alpha)$ , we choose  $\beta = (2\alpha)^{1/2}$ , so that both procedures have FDR guarantee of  $\alpha$ . Let  $k^*$  be the number of hypotheses rejected by  $\mathcal{D}'(\alpha)$ . Note that a sufficient condition for  $k_\psi^* \geq k^*$  is  $\psi(p_{k^*}) \geq K/k^*$ . If we set

$$\gamma := \frac{k^* \alpha}{K \ell_K} \approx p_{k^*},$$

then, approximately, the above condition is

$$\psi(\gamma) \geq \frac{K}{k^*} \iff \frac{k^*}{K} \geq \frac{2}{\ell_K}.$$

Note that  $\ell_K \sim \log K$ , and  $k^*/K$  is the proportion of rejection among all hypotheses. Hence,  $\mathcal{D}(\psi)$  is more powerful than  $\mathcal{D}'(\alpha)$  roughly when the proportion of rejections exceeds  $2/\log K$ . This conclusion is independent of  $\alpha$ .

## 11 Simulations: multiple testing for ordered hypotheses using multi-armed bandits

In this section, we conduct simulation studies for a setting that merges multiple testing with a simple ordered variant of the multi-armed bandit problem. Appendix A contains additional simulation results for a more classical non-sequential multiple testing setup involving z-tests. While it is more traditional to present the latter application in a paper, we find the former a more conceptually interesting case study.

Consider a multi-armed bandit with  $K$  arms, where pulling arm  $k$  produces an iid sample  $(X_1^k, X_2^k, \dots)$  from a non-negative random reward  $X^k$ . The  $K$  null hypotheses are  $\mathbb{E}[X^k] \leq 1$ ,  $k = 1, \dots, K$ , which are nonparametric and similar to Example 2.1. For interpretation's sake, think of  $X_j^k$  as follows: each arm is an investment and  $X_j^k$  is the return of the  $k$ -th investment on its  $j$ -th trial. Specifically, imagine that each investment  $k$  starts with unit wealth and the wealth can be multiplied by  $X_j^k$  when the arm is pulled the  $j$ -th time. Our high-level aim is thus to find profitable investments as quickly as possible; although we do not write down a formal objective, this is not necessary to convey our main points. (One may equally think of testing if the  $k$ -th drug has an effect or not.)

Consider a simple setting in which we have used side information or prior knowledge to arrange the arms in a prespecified order and we must collect data on arm  $k$  before visiting arm  $k + 1$ ; thus previous arms cannot be revisited to obtain more samples. Each arm can be pulled at most  $n$  times (the budget) before moving on to the next one.

We describe a natural algorithm to tackle this problem (an algorithm consists of two parts: deciding when to stop pulling an arm and move to the next one, and deciding which nulls to reject at the end). Our algorithm will use a multiple testing procedure  $\mathcal{D}$  such as BH or e-BH for both steps: we move onto the next arm if we have exhausted the budget for the current arm, or our multiple testing procedure can reject the null with the available information. To formalize this, we take an e-testing procedure as an example. (For a p-testing procedure, e-values are replaced by p-values.) For  $k = 1, \dots, K$ :

1. At the end of dealing with arm  $k - 1$ , we can summarize the available information as a vector of e-values  $\mathbf{e}_{k-1} := (e_1, \dots, e_{k-1}, 1, \dots, 1)$ .
2. Denote by  $e_{k,j}$  the e-value obtained after pulling the  $k$ -th arm  $j$  times. We pull arm  $k$  and stop pulling after  $T_k$  trials if either
  - (a)  $|\mathcal{D}(e_1, \dots, e_{k-1}, e_{k,T_k}, 1, \dots, 1)| > |\mathcal{D}(\mathbf{e}_{k-1})|$ , i.e., there is at least one new discovery, or
  - (b) arm  $k$  has been pulled  $T_k = n$  times, i.e., the budget is exhausted.
3. Set  $e_k := e_{k,T_k}$  and move on to arm  $k + 1$ .

After all  $K$  arms are finished, we apply the procedure  $\mathcal{D}$  to  $\mathbf{e}_K := (e_1, \dots, e_K)$  to determine the rejected hypotheses.

Since we are in a nonparametric setting, the natural choices of e-values and p-values are based on nonnegative supermartingales. More precisely, the e-value  $e_{k,j}$  and the p-value  $p_{k,j}$  are realized by, respectively,

$$E_{k,j} := \prod_{i=1}^j X_i^k \quad \text{and} \quad P_{k,j} := \left( \max_{i=1, \dots, j} E_{k,i} \right)^{-1},$$

for each  $k = 1, \dots, K$  and  $j = 1, \dots, n$ . Also, define  $E_{k,0} = P_{k,0} = 1$  for simplicity. It is not hard to see that if  $H_k$  is true then  $(E_{k,j})_{j=0}^n$  is a nonnegative supermartingale. The final e-variables  $E_k$  and p-variables  $P_k$  are obtained by

$$E_k = E_{k,T_k} \quad \text{and} \quad P_k = P_{k,T_k}, \quad k = 1, \dots, K.$$

The fact that the above stopped processes yield valid e- and p-values is by the optional stopping theorem and Ville's inequality [21, 12] respectively. The above algorithm is a natural baseline in the sense that, for  $\mathcal{D}$  being BH, e-BH or BY, one cannot obtain more rejections than this algorithm, even if all arms are pulled  $n$  times. The algorithm may be seen as an ordered version of that of Jamieson and Jain [14], but a detailed discussion is not necessary here.

It is important to note that, even assuming data across arms are mutually independent, the produced e-values (or p-values) are dependent in a complicated way; a larger previous realization leads to a smaller threshold for the next stopping rule. Indeed, we have  $T_k < n$  only if Step 2(a) was invoked, meaning that  $e_{k,T_k} \geq \frac{K}{|\mathcal{D}(\mathbf{e}_k)|\alpha}$ . Hence, some negative dependence exists and the BH procedure does not control FDR theoretically. There are, however, two p-testing procedures that have valid FDR: the BY procedure  $\mathcal{D}(\alpha/\ell_K)$  which is designed for arbitrary dependence, and the *compliant BH* (cBH) procedure  $\mathcal{D}(\alpha')$ , where  $\alpha'$  is given in (12); i.e., it satisfies  $\alpha'(1 - \log(\alpha')) = \alpha$  (e.g.,  $\alpha' = 0.0087$  for  $\alpha = 0.05$ ). Both the BY and the cBH procedure give a FDR guarantee of  $\alpha$  in our experiments, although cBH relies on the independence of data across arms (BY and e-BH do not). Therefore, we will compare the four procedures e-BH, BH, BY and cBH in our numerical experiments, while keeping in mind that BH does not have a valid FDR in theory, although by design it has more rejections than the other three procedures. We summarize the allowed conditions for a valid FDR guarantee for these four methods in Table 1.

Table 1: Conditions for the validity of the testing algorithm. To allow for dependence among data across arms, we only require for  $k \in \mathcal{N}$ ,  $\mathbb{E}[X_j^k | \mathcal{G}] \leq 1$  where  $\mathcal{G}$  is the available information before we observe  $X_j^k$  (weaker than independence used in our experiments); see Example 2.1.

	arbitrarily dependent data across arms	arbitrarily dependent stopping rules $T_k$	FDR guarantee in our experiments
e-BH	YES	YES	valid at level $\alpha K_0/K$
BH	NO	NO	not valid
BY	YES	YES	valid at level $\alpha K_0/K$
cBH	NO	YES	valid at level $\alpha K_0/K$

Below we describe the data generating process used in our experiments (note that the design and the validity of the above procedures do not depend on any knowledge of the data generating process). As mentioned above, arms are naturally ordered such that more promising arms come first. More specifically, arm  $k$  is non-null with probability  $\theta(K - k + 1)/(K + 1)$  where  $\theta \in [0, 1]$  is a parameter. The expected number of non-nulls in this setting is  $\theta/2$ . Further, let  $s_k$  be the strength of signal in each non-null hypothesis, which follows an iid exponential distribution with mean  $\mu$ . With this setting, some non-nulls may have a very weak signal and they are almost impossible to distinguish from a null, whereas some other non-nulls with a strong signal may be rejected quickly with a few pulls. Conditional on  $s_k$ , the data  $X_1^k, \dots, X_n^k$  for arm  $k$  are iid following a log-normal distribution

$$X^k = \exp(Z^k + s_k \mathbb{1}_{\{k \in \mathcal{K} \setminus \mathcal{N}\}} - 1/2)$$

where  $Z^1, \dots, Z^K$  are iid standard normal. Note that the data are independent across arms, and hence the cBH procedure is valid (see Table 1).

In the experiments, we set  $\alpha = 0.05$  and  $\theta = 1/2$  (on average, 1/4 of all hypotheses are non-null). Results for some values of other parameters are summarized in Table 2.



Table 2: Simulation results in the multi-armed bandit setting, where  $R$  is the number of rejected hypothesis,  $B$  is the proportion of unused budget (number of unused pulls divided by  $nK$ ), and TD is the number of true discoveries (rejected non-null hypotheses). Each number is computed over an average of 500 trials. The default values of parameters are  $K = 500$ ,  $n = 50$  and  $\mu = 1$ , and each panel may have one parameter value different from the default.

(a) Default					(b) $K = 2000$				
	$R$	$B\%$	TD	FDP%		$R$	$B\%$	TD	FDP%
e-BH	74.4	11.42	73.2	1.58	e-BH	297.6	11.39	293.2	1.48
BH	77.0	11.44	75.3	2.13	BH	307.8	11.41	301.4	2.07
BY	70.6	10.06	70.4	0.31	BY	281.2	9.95	280.4	0.26
cBH	71.1	10.16	70.8	0.36	cBH	284.5	10.15	283.5	0.36

(c) $n = 10$					(d) $n = 100$				
	$R$	$B\%$	TD	FDP%		$R$	$B\%$	TD	FDP%
e-BH	47.7	3.99	47.3	0.83	e-BH	79.1	13.48	77.9	1.50
BH	49.3	4.01	48.7	1.06	BH	81.3	13.50	79.5	2.13
BY	38.4	2.77	38.4	0.08	BY	76.4	12.36	76.1	0.35
cBH	39.2	2.85	39.2	0.11	cBH	76.7	12.44	76.4	0.41

(e) $\mu = 0.5$					(f) $\mu = 2$				
	$R$	$B\%$	TD	FDP%		$R$	$B\%$	TD	FDP%
e-BH	43.5	5.77	42.9	1.54	e-BH	97.4	16.46	95.9	1.54
BH	46.3	5.80	45.3	2.13	BH	99.3	16.47	97.2	2.07
BY	39.6	4.66	39.5	0.27	BY	94.3	15.23	94.1	0.29
cBH	40.1	4.74	40.0	0.35	cBH	94.6	15.32	94.3	0.35

In all settings of Table 2, we see that the e-BH procedure outperforms both BY and cBH by providing more discoveries and using a smaller number of pulls. The BY and cBH procedures are both quite conservative and they perform similarly, with cBH being better for large  $K$  as expected. The numbers of true discoveries of e-BH are slightly lower than those of BH, but, as we explained above, BH does not have a theoretical guarantee of FDR (its realized FDP is not larger than  $\alpha$  in our experiments, but this may be due to the fact that the null p-values are not uniform).

For many other settings of parameters and distributions of  $s_k$  that we tried, the relative performance of the four methods is always qualitatively similar.

## 12 Conclusion

We have introduced the e-BH procedure to achieve FDR control with e-values. Some highlights of the e-BH procedure are summarized below:

1. It works for arbitrarily dependent e-values;
2. it requires no information on the configuration of the input e-values, and also works for weighted e-values;
3. it allows for boosting e-values if partial distributional information is available on some e-values;
4. it gives rise to a class of p-testing procedures which include both the BH procedure and the BY procedure as special cases (along with proofs of FDR control);

5. it is optimal among the class of e-testing procedures  $\mathcal{G}(\phi)$  (Theorem 9.1) in the setting of arbitrary e-values.

Although the e-BH procedure is primarily designed for cases where e-values are directly available (not calibrated from bonafide p-values), even the p-value calibrated e-BH procedure may outperform the BY procedure in some settings (see the simulation results in Appendix A). Based on our results, the overall recommendation on using e-BH is very simple, assuming that one is not willing to risk violating FDR control:

1. If e-values are directly available and p-values are calibrated from e-values (e.g., those obtained from stopped test supermartingales or universal inference), then one should use the e-BH procedure.
2. If bonafide p-values are available but arbitrarily dependent, then one can use the e-BH procedure with calibrated e-values, the BY procedure being a special case; these methods generally do not dominate each other.
3. If bonafide p-values are available and known to be PRDS, then the base BH procedure should be used (it is a special case of the calibrated e-BH, but an optimal one in the sense of Proposition 10.2).

Finally, we remark that there are two more complicated situations where the e-BH procedure can be applied:

1. Natural e-values are available for some hypotheses, whereas natural p-values are available for the others. Depending on proportion, one may calibrate them all to p-values or all to e-values.
2. A hypothesis has both an e-value and a p-value (or even multiple p-values and e-values), which may be dependent and obtained from different experiments. In such a situation, the power comparison of e- and p-based procedures depends on the quality of those e-values and p-values. One may choose to combine them into a single e-value or p-value for each hypothesis (e.g., using the methods in [25]).

A more detailed theoretical and empirical comparison of the last two “hybrid” situations is left for future work.

## Acknowledgements

We thank Vladimir Vovk for many helpful discussions and his kind help on some numerical experiments of the e-BH procedure. AR acknowledges NSF DMS CAREER grant 1945266. RW acknowledges financial support from NSERC grants (RGPIN-2018-03823, RGPAS-2018-522590).

## References

- [1] Rina Foygel Barber and Aaditya Ramdas. The p-filter: multilayer false discovery rate control for grouped hypotheses. *Journal of the Royal Statistical Society Series B*, 79(4):1247–1268, 2017.
- [2] Laurent Barras, Olivier Scaillet, and Russ Wermers. False discoveries in mutual fund performance: Measuring luck in estimated alphas. *Journal of Finance*, 65(1):179–216, 2010.
- [3] Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society B*, 57:289–300, 1995.
- [4] Yoav Benjamini and Yosef Hochberg. Multiple hypotheses testing with weights. *Scandinavian Journal of Statistics*, 24(3):407–418, 1997.

- [5] Yoav Benjamini and Daniel Yekutieli. The control of the false discovery rate in multiple testing under dependency. *Annals of Statistics*, 29:1165–1188, 2001.
- [6] Gilles Blanchard and Etienne Roquain. Two simple sufficient conditions for FDR control. *Electronic Journal of Statistics*, 2:963–992, 2008.
- [7] Mathias Drton. Likelihood ratio tests and singularities. *The Annals of Statistics*, 37(2):979–1012, 2009.
- [8] Bradley Efron. Student’s t-test under symmetry conditions. *Journal of the American Statistical Association*, 64(328):1278–1302, 1969.
- [9] Paul Embrechts, Bin Wang, and Ruodu Wang. Aggregation-robustness and model uncertainty of regulatory risk measures. *Finance and Stochastics*, 19(4):763–790, 2015.
- [10] William Fithian and Lihua Lei. Conditional calibration for false discovery rate control under dependence. Technical Report [arXiv:2007.10438 \[math.ST\]](#), [arXiv.org](#) e-Print archive, July 2020.
- [11] Peter Grünwald, Rianne de Heide, and Wouter M. Koolen. Safe testing. Technical Report [arXiv:1906.07801 \[math.ST\]](#), [arXiv.org](#) e-Print archive, June 2020.
- [12] Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform Chernoff bounds via nonnegative supermartingales. *Probab. Surveys*, 17:257–317, 2020.
- [13] Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics (to appear)*, 2021.
- [14] Kevin G Jamieson and Lalit Jain. A bandit approach to sequential experimental design with false discovery control. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 3660–3670, 2018.
- [15] Tiefeng Jiang and Fan Yang. Central limit theorems for classical likelihood ratio tests for high-dimensional normal distributions. *The Annals of Statistics*, 41(4):2029–2074, 2013.
- [16] Robert E Kass and Adrian E Raftery. Bayes factors. *Journal of the american statistical association*, 90(430):773–795, 1995.
- [17] Eugene Katsevich, Chiara Sabatti, and Marina Bogomolov. Controlling fdr while highlighting selected discoveries. *arXiv preprint arXiv:1809.01792*, 2018.
- [18] Aaditya Ramdas, Jianbo Chen, Martin J Wainwright, and Michael I Jordan. A sequential algorithm for false discovery rate control on directed acyclic graphs. *Biometrika*, 106(1):69–86, 2019.
- [19] Aaditya K. Ramdas, Rina F. Barber, Martin J. Wainwright, and Michael I. Jordan. A unified treatment of multiple testing with prior knowledge using the p-filter. *Annals of Statistics*, 47:2790–2821, 2019.
- [20] Glenn Shafer. The language of betting as a strategy for statistical and scientific communication. Technical Report [arXiv:1903.06991 \[math.ST\]](#), [arXiv.org](#) e-Print archive, March 2019. To appear as discussion paper in the *Journal of the Royal Statistical Society Series A*.
- [21] Glenn Shafer, Alexander Shen, Nikolai Vereshchagin, and Vladimir Vovk. Test martingales, bayes factors and p-values. *Statistical Science*, 26(1):84–101, 2011.
- [22] John Storey. A direct approach to false discovery rates. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(3):479–498, 2002.

- [23] Weijie Su. The fdr-linking theorem. Technical Report [arXiv:1812.08965 \[math.ST\]](#), [arXiv.org](#) e-Print archive, December 2018.
- [24] Vladimir Vovk and Ruodu Wang. Combining p-values via averaging. *Biometrika*, 2020. doi: 10.1093/biomet/asaa027.
- [25] Vladimir Vovk and Ruodu Wang. E-values: Calibration, combination, and applications. Technical Report [arXiv:1912.06116 \[math.ST\]](#), [arXiv.org](#) e-Print archive, September 2020. To appear in the *Annals of Statistics*.
- [26] Vladimir Vovk, Bin Wang, and Ruodu Wang. Admissible ways of merging p-values under arbitrary dependence. Technical Report [arXiv:2007.14208 \[math.ST\]](#), [arXiv.org](#) e-Print archive, July 2020.
- [27] Larry Wasserman, Aaditya Ramdas, and Sivaraman Balakrishnan. Universal inference. *Proceedings of the National Academy of Sciences*, 117(29):16880–16890, 2020.
- [28] Ian Waudby-Smith and Aaditya Ramdas. Confidence sequences for sampling without replacement. *Neural Information Processing Systems*, 2020.
- [29] Ian Waudby-Smith and Aaditya Ramdas. Variance-adaptive confidence sequences by betting. *arXiv preprint arXiv:2010.09686*, 2020.

## A Simulations for correlated z-tests

We provide further simulation results in a classic setting, where test statistics  $X_k$ ,  $k \in \mathcal{K}$  are generated from correlated z-tests. The null hypotheses are  $N(0, 1)$  and the alternatives are  $N(\delta, 1)$ , where we take  $\delta := -3$  throughout the section. We generate  $K - K_0$  observations from the alternative distribution  $N(\delta, 1)$  and then  $K_0$  observations from the null distribution  $N(0, 1)$ .

We take the p-values as

$$P(x) := \Phi(x), \tag{33}$$

where  $x \in \{X_k : k \in \mathcal{K}\}$ ; these are the p-values found using the most powerful test given by the Neyman–Pearson lemma. The raw e-values are, following Vovk and Wang [25], the likelihood ratios

$$E(x) := \frac{\exp(-(x - \delta)^2/2)}{\exp(-x^2/2)} = \exp(\delta x - \delta^2/2) \tag{34}$$

of the alternative to the null density.

We report the numbers of discoveries for the (base) BH, BY, e-BH (boosted), and base e-BH procedures, in several different settings in Table 3. In particular, we consider two settings of negative correlation (thus, PRDS fails to hold). We set the target FDR level  $\alpha \in \{10\%, 5\%, 2\%\}$ . All numbers in Table 3 are produced on the average from 1,000 trials.

From the simulation results in Table 3, we can see that the BH procedure has slightly more discoveries than the e-BH procedure for PRDS e-values via (7); recall that we do not expect e-BH to outperform BH when PRDS and precise (standard uniform) p-values are available (cf. Proposition 10.2). The e-BH procedure for arbitrarily dependent (AD) e-values performs slightly better than the BY procedure in most settings, and this advantage is more pronounced in the case of a large experiment (Table 3b, left), as the BY correction  $\ell_K$  is penalized by a large  $K$ . Only in the case of sparse signal and small  $\alpha$  (Table 3b, last two columns), BY performs better than e-BH. This is consistent with Example 10.3 (with different e-values from our experiments), where we see that BY outperforms e-BH if the proportion of rejections is very small.

These simulation results suggest that, even in case p-values are available, the e-BH procedure (with boosted e-values) has competitive performance in most situations, especially when PRDS fails to hold (thus the base BH does not have a correct FDR in theory).

Table 3: Simulation results for correlated z-tests, where  $\rho_{i,j}$  is the correlation between two test statistics  $X_i$  and  $X_j$  for  $i \neq j$ . Each cell gives the number of rejections and, in parentheses, the realized FDP (in %). Each number is computed over an average of 1,000 trials.

(a) Independent and positively correlated tests,  $K = 1000$ ,  $K_0 = 800$

	$\rho_{ij} = 0$			$\rho_{ij} = 0.5$		
	$\alpha = 10\%$	$\alpha = 5\%$	$\alpha = 2\%$	$\alpha = 10\%$	$\alpha = 5\%$	$\alpha = 2\%$
BH	177.3 (8.01)	148.7 (4.07)	115.0 (1.63)	180.0 (7.00)	144.8 (3.64)	109.8 (1.50)
e-BH PRDS	171.8 (7.07)	147.6 (3.95)	114.6 (1.62)	170.2 (5.71)	142.5 (3.35)	108.0 (1.50)
BY	101.1 (1.10)	78.8 (0.57)	53.2 (0.22)	96.6 (1.03)	76.7 (0.50)	55.0 (0.20)
e-BH AD	109.4 (1.41)	85.4 (0.68)	54.6 (0.24)	103.1 (1.32)	81.4 (0.70)	56.6 (0.28)
base e-BH	97.5 (1.00)	70.6 (0.43)	36.9 (0.11)	91.9 (0.97)	69.1 (0.45)	43.6 (0.16)

(b) Independent tests with large number of hypotheses

	$K = 20,000$ , $K_0 = 10,000$			$K = 20,000$ , $K_0 = 19,000$		
	$\alpha = 10\%$	$\alpha = 5\%$	$\alpha = 2\%$	$\alpha = 10\%$	$\alpha = 5\%$	$\alpha = 2\%$
BH	9567 (5.00)	8564 (2.49)	7164 (1.00)	681.3 (9.58)	520.2 (4.79)	357.7 (1.93)
e-BH PRDS	9092 (3.60)	8330 (2.13)	7124 (0.98)	681.3 (9.58)	509.3 (4.54)	312.1 (1.40)
BY	5956 (0.48)	4818 (0.24)	3417 (0.10)	254.1 (0.89)	177.6 (0.46)	103.1 (0.19)
e-BH AD	6811 (0.80)	5809 (0.44)	4384 (0.18)	271.0 (1.02)	159.5 (0.39)	51.4 (0.07)
base e-BH	6426 (0.64)	5234 (0.31)	3509 (0.10)	224.8 (0.69)	109.2 (0.21)	16.4 (0.01)

(c) Negatively correlated tests,  $K = 1000$ ,  $K_0 = 800$ .

	$\rho_{ij} = -1/(K-1)$			$\rho_{ij} = -0.5\mathbb{1}_{\{ i-j =1\}}$		
	$\alpha = 10\%$	$\alpha = 5\%$	$\alpha = 2\%$	$\alpha = 10\%$	$\alpha = 5\%$	$\alpha = 2\%$
BH	177.7 (8.14)	149.0 (4.09)	115.2 (1.61)	177.2 (8.10)	148.8 (4.00)	115.3 (1.62)
e-BH PRDS	172.0 (7.13)	147.9 (3.98)	114.9 (1.59)	171.5 (7.13)	147.7 (3.89)	114.9 (1.61)
BY	101.2 (1.08)	78.8 (0.52)	53.3 (0.20)	101.3 (1.11)	78.8 (0.56)	53.2 (0.22)
e-BH AD	109.7 (1.38)	85.5 (0.65)	54.6 (0.22)	109.8 (1.40)	85.6 (0.69)	54.6 (0.24)
base e-BH	97.8 (0.98)	70.7 (0.40)	37.2 (0.11)	97.6 (0.99)	70.7 (0.41)	36.7 (0.12)

## B Omitted proofs from the paper

*Proof of Lemma 6.1.* Note that  $f(\mathbf{X}) \leq X'_k$  on the event  $\{X_k \geq f(\mathbf{X})\}$ .

(i) If  $X_k$  is independent of  $\mathbf{X}^{-k}$ , then

$$\mathbb{E} [f(\mathbf{X})\mathbb{1}_{\{X_k \geq f(\mathbf{X})\}} \mid \mathbf{X}^{-k}] \leq \mathbb{E} [X'_k\mathbb{1}_{\{X_k \geq f(\mathbf{X})\}} \mid \mathbf{X}^{-k}] \leq \mathbb{E} [X'_k] \leq \mathbb{E}[X_k].$$

(ii) Let  $g : x \mapsto \mathbb{P}(X'_k \geq x)$  and  $P_k := g(X'_k)$ . Clearly  $P_k$  is a null p-value. For each  $y \geq 0$ , since  $\mathbf{X}$  is PRDS on  $X_k$  and  $g \circ f$  is an increasing function, the function  $x \mapsto \mathbb{P}(g \circ f(\mathbf{X}) \leq y \mid X_k \geq x)$  is decreasing on  $[0, \infty)$ . Noting that  $P_k$  is a decreasing function of  $X_k$ , the function  $t \mapsto \mathbb{P}(g \circ f(\mathbf{X}) \leq y \mid P_k \leq t)$  is increasing on  $[0, 1]$ . Using Lemma 1 of Ramdas et al. [19], we have

$$\mathbb{E} \left[ \frac{\mathbb{1}_{\{P_k \leq g \circ f(\mathbf{X})\}}}{g \circ f(\mathbf{X})} \right] \leq 1.$$

It follows that

$$\begin{aligned}
\mathbb{E}[f(\mathbf{X})\mathbf{1}_{\{E_k \geq f(\mathbf{X})\}}] &= \mathbb{E}[f(\mathbf{X})\mathbf{1}_{\{E'_k \geq f(\mathbf{X})\}}] \\
&\leq \mathbb{E}[f(\mathbf{X})\mathbf{1}_{\{P_k \leq g \circ f(\mathbf{X})\}}] \\
&= \mathbb{E}\left[f(\mathbf{X})\mathbb{P}(E'_k \geq f(\mathbf{X})) \frac{\mathbf{1}_{\{P_k \leq g \circ f(\mathbf{X})\}}}{g \circ f(\mathbf{X})}\right] \\
&\leq \sup_{x \geq 0} x \mathbb{P}(X'_k \geq x) \mathbb{E}\left[\frac{\mathbf{1}_{\{P_k \leq g \circ f(\mathbf{X})\}}}{g \circ f(\mathbf{X})}\right] \\
&\leq \sup_{x \geq 0} x \mathbb{P}(X'_k \geq x).
\end{aligned}$$

(iii) It follows directly that

$$\mathbb{E}[f(\mathbf{X})\mathbf{1}_{\{X_k \geq f(\mathbf{X})\}}] \leq \mathbb{E}[X'_k \mathbf{1}_{\{X_k \geq f(\mathbf{X})\}}] \leq \mathbb{E}[X'_k] \leq \mathbb{E}[X_k].$$

If  $X_k$  is a null e-value, then  $\mathbb{E}[X_k] \leq 1$  by definition.  $\square$

*Proof of Theorem 7.1.* We will apply Lemma 6.1 to  $\mathbf{X} = \mathbf{E}'$  and the function  $f$  given by  $f(\mathbf{X}) = t_\alpha$ . Note that  $I_f = \alpha^{-1}K/\mathcal{K}$ , and in this case,  $X'_k = T(\alpha X_k)/\alpha$ . By Lemma 6.1 (iii) and (16), we have

$$\mathbb{E}\left[\frac{F_{\mathcal{G}(\alpha)}}{R_{\mathcal{G}(\alpha)}}\right] = \frac{\alpha}{K} \mathbb{E}\left[\sum_{k \in \mathcal{N}} t_\alpha \mathbf{1}_{\{E'_k \geq t_\alpha\}}\right] \leq \frac{1}{K} \mathbb{E}\left[\sum_{k \in \mathcal{N}} X'_k\right] = \frac{K_0}{K} y_\alpha,$$

thus showing the first inequality. If  $E'_1, \dots, E'_K$  are the raw e-values or the boosted e-values obtained by (8) or (10), then  $y_\alpha \leq \alpha$  by construction.  $\square$

*Proof of Theorem 8.1.* Write  $\mathbf{E}' = (E'_1, \dots, E'_K)$ . For each  $k \in \mathcal{N}$ , denote by

$$z_{\alpha,k} := \max_{x \in K/\mathcal{K}} x \mathbb{P}(\alpha E'_k \geq x) = \max_{x \geq 1} x \mathbb{P}(T(\alpha E'_k) \geq x) \leq \mathbb{E}[T(\alpha E'_k)].$$

Note that  $t_\alpha = f(\mathbf{E}')$  for a decreasing function  $f$  with range  $I_f = \alpha^{-1}K/\mathcal{K}$ . Using the notation in Lemma 6.1 with  $\mathbf{X} = \mathbf{E}'$ , we have

$$X'_k = \sup\{x \in I_f \cup \{0\} : x \leq E'_k\} = \frac{T(\alpha E'_k)}{\alpha}.$$

Lemma 6.1 (ii) gives

$$\mathbb{E}[t_\alpha \mathbf{1}_{\{E'_k \geq t_\alpha\}}] \leq \sup_{x \geq 0} x \mathbb{P}(X'_k \geq x) = \max_{x \geq 1} \frac{x}{\alpha} \mathbb{P}(T(\alpha E'_k) \geq x) = \frac{z_{\alpha,k}}{\alpha}.$$

Hence, by (16),

$$\mathbb{E}\left[\frac{F_{\mathcal{G}(\alpha)}}{R_{\mathcal{G}(\alpha)}}\right] = \frac{\alpha}{K} \mathbb{E}[t_\alpha f(t_\alpha)] = \frac{\alpha}{K} \sum_{k \in \mathcal{N}} \mathbb{E}[t_\alpha \mathbf{1}_{\{E'_k \geq t_\alpha\}}] \leq \frac{1}{K} \sum_{k \in \mathcal{N}} z_{\alpha,k}.$$

Note that  $z_\alpha \leq y_\alpha$  by Markov's inequality. Hence, if  $E'_1, \dots, E'_K$  are raw e-values, then  $z_\alpha \leq y_\alpha \leq \alpha$ . If  $E'_1, \dots, E'_K$  are boosted e-values, then  $z_\alpha \leq \alpha$  by construction.  $\square$

*Proof of Proposition 8.2.* Since  $b_k = (\alpha q_{1-\alpha}(E_k))^{-1}$ , (24) implies

$$t \mapsto t \mathbb{P}(b_k E_k \geq t) \text{ is decreasing on } [1/\alpha, \infty).$$

Therefore,  $z_{\alpha,k}(b_k)$  in (23) satisfies  $z_{\alpha,k}(b_k) = \mathbb{P}(\alpha b_k E_k \geq 1)$ . Direct calculation gives

$$z_{\alpha,k}(b_k) = \mathbb{P}(\alpha E_k \geq 1/b_k) = \mathbb{P}(\alpha E_k \geq \alpha q_{1-\alpha}(E_k)) = \alpha.$$

Since  $E_k$  is continuously distributed,  $b_k$  cannot be enlarged without violating (7).  $\square$

*Proof of Theorem 9.1.* First, from (26), only the values of  $\phi^{-1}(K/k)$  for  $k \in \mathcal{K}$  affects the testing procedure  $\mathcal{G}(\phi)$ . Note that  $\mathcal{G}(\alpha)$  uses the linear transform  $t \mapsto \alpha t$ . Therefore, if  $\phi^{-1}(x) \geq x/\alpha$  for all  $x \in K/\mathcal{K}$ , then  $\mathcal{G}(\phi)$  is dominated by the e-BH procedure  $\mathcal{G}(\alpha)$ , i.e.,  $\mathcal{G}(\phi) \subseteq \mathcal{G}(\alpha)$ .

Next, suppose for the purpose of contradiction that  $\phi^{-1}(K/k_0) < K/(k_0\alpha)$  for some  $k_0 \in \mathcal{K}$ . Write

$$m = \phi^{-1}(K/k_0) \quad \text{and} \quad m' = \max\{m, K/k_0\}.$$

Note that  $m \leq m' < K/(k_0\alpha)$ . Consider the following setting of e-values: Let  $\mathcal{N} = \mathcal{K}$ , i.e., all e-values are null. For  $k \in \mathcal{K}$ , set  $E_k = m' \mathbf{1}_{A_k}$  for some set  $A_k$  with  $\mathbb{P}(A_k) = 1/m'$ . We design the sets  $A_k$ ,  $k \in \mathcal{K}$ , in the following way: for each  $\omega \in \Omega$ ,

$$|\{k \in \mathcal{K} : \omega \in A_k\}| = k_0,$$

i.e., exactly  $k_0$  events out of  $A_1, \dots, A_K$  occur together (the existence of such an arrangement is easy to verify). Let  $A = \bigcup_{k \in \mathcal{K}} A_k$ , and we have  $\mathbb{P}(A) = K/(k_0 m') \in (\alpha, 1]$ .

If event  $A$  happens, then exactly  $k_0$  of  $E_1, \dots, E_K$  take the value  $m'$ , and the rest take the value 0. Hence, for  $\tau > 0$ , we have  $R(\tau) = \max(k_0 \mathbf{1}_{\{\tau \leq m'\}} \mathbf{1}_A, 1)$ . If event  $A$  happens, then  $\tau_\phi \leq m'$  since  $\phi(m')R(m') \geq \phi(m)k_0 = K$ , and  $\tau_\phi \geq m$  since  $\phi(\tau_\phi) \geq K/k_0$ . Therefore,  $\tau_\phi \in [m, m']$  if  $A$  happens. It follows that, for each  $k$ ,

$$\mathbb{E}[\phi(\tau_\phi) \mathbf{1}_{\{E_k \geq \tau_\phi\}}] \geq \mathbb{E}[\phi(m) \mathbf{1}_{\{E_k \geq m'\}} \mathbf{1}_{A_k}] = \mathbb{E}[\phi(m) \mathbf{1}_{A_k}] = \frac{K}{k_0 m'} > \alpha.$$

Using the FDR control of  $\mathcal{G}(\phi)$  and (27), this leads to

$$\alpha \geq \mathbb{E} \left[ \frac{F_{\mathcal{G}(\phi)}}{R_{\mathcal{G}(\phi)}} \right] = \frac{1}{K} \sum_{k \in \mathcal{K}} \mathbb{E}[\phi(\tau_\phi) \mathbf{1}_{\{E_k \geq \tau_\phi\}}] > \alpha,$$

a contradiction. Hence,  $\mathcal{G}(\phi)$  is dominated by  $\mathcal{G}(\alpha)$ .  $\square$

*Proof of Proposition 10.1.* For  $k \in \mathcal{N}$ , let  $\alpha = \mathbb{E}[T(\psi(P))]$  where  $P$  is uniform on  $[0, 1]$ . If  $\psi^{-1}(1) = 0$ , then  $\mathbb{P}(\psi(P_k) \geq 1) = 0$  for  $k \in \mathcal{N}$ , and thus  $\mathcal{D}(\psi)$  will never reject any hypotheses with a non-zero p-value, leading to a zero FDR for both PRDS and arbitrarily dependent p-values. Next, assume  $\psi^{-1}(1) > 0$ , which implies  $\alpha > 0$ .

Define

$$E_k = \frac{T(\psi(P_k))}{\alpha}, \quad k \in \mathcal{K}, \quad (35)$$

which is a null e-value for  $k \in \mathcal{N}$ . Note that

$$\psi(p_{(k)}) \geq \frac{K}{k} \iff e_{[k]} \geq \frac{K}{k\alpha}.$$

Hence,  $\mathcal{G}(\alpha)$  applied to  $(e_1, \dots, e_K)$  is equivalent to  $\mathcal{D}(\psi)$  applied to  $(p_1, \dots, p_K)$ . Note that

$$\begin{aligned} y_\alpha &= \frac{1}{K_0} \sum_{k \in \mathcal{N}} \mathbb{E}[T(\psi(P_k))] \\ &\leq \sum_{j=1}^K \frac{K}{j} \left( \psi^{-1} \left( \frac{K}{j} \right) - \psi^{-1} \left( \frac{K}{j-1} \right) \right) = y_\psi. \end{aligned}$$

and

$$z_\alpha = \max_{x \in K/\mathcal{K}} x \mathbb{P}(\psi(P_k) \geq x) \leq \max_{x \in K/\mathcal{K}} x \psi^{-1}(x) = z_\psi.$$

Using Theorems 7.1 and 8.1, the testing procedure  $\mathcal{D}(\psi)$  satisfies

$$\mathbb{E} \left[ \frac{F_{\mathcal{D}(\psi)}}{R_{\mathcal{D}(\psi)}} \right] = \mathbb{E} \left[ \frac{F_{\mathcal{G}(\alpha)}}{R_{\mathcal{G}(\alpha)}} \right] \leq \frac{K_0}{K} y_\alpha \leq \frac{K_0}{K} y_\psi,$$

and  $y_\psi$  can be replaced by  $z_\psi$  if the p-values are PRDS.  $\square$



*Proof of Proposition 10.2.* Consider the setup where all p-values are null, and they are identical following a uniform distribution on  $[0, 1]$ . Note that identical p-values satisfy PRDS. In this case,  $F_{\mathcal{D}(\psi)}/R_{\mathcal{D}(\psi)} = 1$  as soon as there is any discovery. If  $\psi^{-1}(1) > \alpha$ , then

$$\mathbb{P}(\psi(p_{(K)}) \geq 1) = \mathbb{P}(p_{(K)} \leq \psi^{-1}(1)) = \min\{\psi^{-1}(1), 1\} > \alpha.$$

Hence, the probability of having a false discovery is more than  $\alpha$ , violating the assumption that  $\mathbb{E}[F_{\mathcal{D}(\psi)}/R_{\mathcal{D}(\psi)}] \leq \alpha$ . Therefore,  $\psi^{-1}(1) \leq \alpha$ .

If  $\psi$  satisfies (32) and  $\psi^{-1}(1) \leq \alpha$ , then we have  $t\psi^{-1}(t) \leq \alpha$  for all  $t \geq 1$ . Hence,  $\psi(p) \leq \alpha/p$  for  $p \in (0, \alpha]$ . As a consequence,  $\mathcal{D}(\psi) \subseteq \mathcal{D}(\alpha)$ .  $\square$

## C Using e-BH with multiple decreasing transforms

We can generalize the p-testing procedure  $\mathcal{D}(\psi)$  in Section 10 to the case of multiple decreasing transforms, which we briefly describe here.

Let  $\psi_1, \dots, \psi_K$  be decreasing transforms, and write  $r_k := \psi_k(p_k)$ ,  $k \in \mathcal{K}$ . Design a p-testing procedure  $\mathcal{D}(\psi_1, \dots, \psi_K)$  by rejecting the  $k^*$  hypotheses with the largest  $r_k$ , where

$$k^* = \max \left\{ k \in \mathcal{K} : r_{[k]} \geq \frac{K}{k} \right\},$$

with the convention  $\max(\emptyset) = 0$ . Here, we are not rejecting hypotheses with the smallest p-values (a hypothesis with a smaller p-value may not be rejected before one with a larger p-value), but rather those with the largest values of  $r_k$ ; thus the procedure is no longer a step-up procedure in Section 10. Recall that for a decreasing transform  $\psi$ ,  $y_\psi$  is defined by (30), and  $z_\psi$  is defined by (31). The following result is a stronger version of Proposition 10.1.

**Proposition C.1.** *For arbitrary p-values and decreasing transforms  $\psi_1, \dots, \psi_K$ , the testing procedure  $\mathcal{D}(\psi_1, \dots, \psi_K)$  satisfies*

$$\mathbb{E} \left[ \frac{F_{\mathcal{D}(\psi_1, \dots, \psi_K)}}{R_{\mathcal{D}(\psi_1, \dots, \psi_K)}} \right] \leq \frac{1}{K} \sum_{k \in \mathcal{N}} y_{\psi_k}.$$

*If the p-values are PRDS, then*

$$\mathbb{E} \left[ \frac{F_{\mathcal{D}(\psi_1, \dots, \psi_K)}}{R_{\mathcal{D}(\psi_1, \dots, \psi_K)}} \right] \leq \frac{1}{K} \sum_{k \in \mathcal{N}} z_{\psi_k}.$$

*Proof.* The proof is similar to that of Proposition 10.1. Define

$$E_k = \frac{T(\psi_k(P_k))}{\beta}, \quad k \in \mathcal{K},$$

where  $\beta = \max_{k \in \mathcal{N}} \mathbb{E}[T(\psi_k(P_k))]$ . Clearly,  $E_k$  is an e-value for  $k \in \mathcal{N}$ . Note that

$$r_{[k]} \geq \frac{K}{k} \iff e_{[k]} \geq \frac{K}{k\beta}.$$

Therefore,  $\mathcal{G}(\beta)$  applied to  $(e_1, \dots, e_K)$  is equal to  $\mathcal{D}(\psi_1, \dots, \psi_K)$  applied to  $(p_1, \dots, p_K)$ . Using Theorems 7.1 and 8.1,  $\mathcal{D}(\psi_1, \dots, \psi_K)$  satisfies

$$\mathbb{E} \left[ \frac{F_{\mathcal{D}(\psi_1, \dots, \psi_K)}}{R_{\mathcal{D}(\psi_1, \dots, \psi_K)}} \right] = \mathbb{E} \left[ \frac{F_{\mathcal{G}(\beta)}}{R_{\mathcal{G}(\beta)}} \right] \leq \frac{K_0}{K} y_\beta = \frac{1}{K} \sum_{k \in \mathcal{N}} y_{\psi_k},$$

and  $y_{\psi_k}$  can be replaced by  $z_{\psi_k}$  if the p-values are PRDS.  $\square$