

An online generalization of the (e-)Benjamini-Hochberg procedure

Lasse Fischer*

Ziyu Xu†

Aaditya Ramdas‡

July 7, 2025

Abstract

In online multiple testing, the hypotheses arrive one by one, and at each time we must immediately reject or accept the current hypothesis solely based on the data and hypotheses observed so far. Many online procedures have been proposed, but none of them are generalizations of the Benjamini-Hochberg (BH) procedure based on p-values, or of the e-BH procedure that uses e-values. In this paper, we consider a relaxed problem setup that allows the current hypothesis to be rejected at *any later step*. We show that this relaxation allows us to define — what we justify extensively to be — the natural and appropriate online extension of the BH and e-BH procedures. We show that the FDR guarantees for BH (resp. e-BH) and online BH (resp. online e-BH) are identical under positive, negative or arbitrary dependence, at fixed and stopping times. Further, the online BH (resp. online e-BH) rule recovers the BH (resp. e-BH) rule as a special case when the number of hypotheses is known to be fixed. Of independent interest, our proof techniques also allow us to prove that numerous existing online procedures, which were known to control the FDR at fixed times, also control the FDR at stopping times.

*University of Bremen, Germany. Email: fischer1@uni-bremen.de

†Carnegie Mellon University, USA. Email: xzy@cmu.edu

‡Carnegie Mellon University, USA. Email: aramdas@cmu.edu

Contents

1	Introduction	3
1.1	Online ARC (accept-to-reject) procedures	3
1.2	Problem Setup	4
1.3	Our contributions	4
1.4	The TOAD algorithm	5
2	The online e-BH procedure	6
2.1	SupFDR control under arbitrary dependence	6
2.2	Online e-BH uniformly improves e-LOND	7
2.3	Boosting the online e-BH procedure	8
2.4	Online e-BH with p-values	9
3	The online BH procedure	11
3.1	OnlineFDR control under positive dependence	11
3.2	SupFDR control under positive dependence	12
3.3	SupFDR control under negative dependence	13
3.4	SupFDR K control under arbitrary dependence	14
4	Related literature	14
5	Conclusion	15
A	Details about boosting with online e-BH	19
A.1	Handling the infinite sum in the truncation function used for boosting	19
A.2	Boosting under local dependence	20
A.3	Boosting with PRDS e-values	22
B	The e-TOAD procedure	23
C	Online BH versus LORD	24
D	The online Storey-BH procedure	25
E	Positive and negative dependence	27
F	Simulation setup	28
G	Omitted proofs and derivations	28
G.1	Online e-BH with p-values	28
G.2	SupFDR control of online BR and r-LOND	29
G.3	Weighted FDR-Linking theorem	29
G.4	Relationship between SupFDR and StopFDR	31
G.5	Sharpness of SupFDR and StopFDR control of online BH	31

1 Introduction

Suppose we observe a stream of (null) hypotheses $(H_t)_{t \in \mathbb{N}}$ that comes either with a sequence of p-values $(P_t)_{t \in \mathbb{N}}$ or with a sequence of e-values $(E_t)_{t \in \mathbb{N}}$. Let \mathbb{P} denote the true but unknown probability distribution and \mathbb{E} the corresponding expected value, while \mathbb{P}_{H_0} and \mathbb{E}_{H_0} denote the probability and expectation under the null hypothesis (that is, according to any distribution that satisfies H_0). A p-value P_t for H_t is a random variable with values in $[0, 1]$ such that $\mathbb{P}_{H_t}(P_t \leq x) \leq x$ for all $x \in [0, 1]$. In contrast, an e-value E_t is a nonnegative random variable with $\mathbb{E}_{H_t}[E_t] \leq 1$. Previous works [14, 17, 11, 34] defined an online multiple testing procedure as a sequence of test decisions $(r_t)_{t \in \mathbb{N}}$, where each r_t is only allowed to depend on the p-values (or e-values) P_1, \dots, P_t . In this paper we relax this setup, which allows to define more powerful procedures that are valid under more general stopping rules. This allows us to propose (what we justify to be) online generalizations of the Benjamini-Hochberg (BH) and e-BH procedures (recapped later).

1.1 Online ARC (accept-to-reject) procedures

An *online ARC procedure* is an online multiple testing method that can change earlier accept decisions (or non-rejections) into a rejection at a later point in time. The motivation is that early in the procedure, there have been no discoveries yet, and so the level at which hypotheses are tested at must be quite stringent leading to fewer discoveries initially. However, as the number of discoveries builds up, we may realize that we could have tested earlier hypotheses at more relaxed levels, leading to rejections at earlier times.

In more detail, an online ARC procedure goes through all p-values P_1, P_2, \dots (or e-values) and at each time $t \in \mathbb{N}$ it decides based on P_1, \dots, P_t which of the hypotheses H_1, \dots, H_t are to be rejected with the restriction that as soon as a hypothesis is rejected, this decision cannot be overturned in the future and thus the rejection remains during the entire testing process. However, in contrast to a fully online multiple testing procedure, an online ARC procedure allows to turn an accepted hypothesis into a rejected one based on information obtained in the future. This is why we call this online with *accept-to-reject changes* (online ARC) procedure. Formally one can define an online ARC procedure as follows.

Definition 1. An online ARC procedure is a nested sequence of rejection sets $R_1 \subseteq R_2 \subseteq R_3 \subseteq \dots$, where each $R_t \subseteq [t] := \{1, \dots, t\}$ only depends on the p-values or e-values that are known up to step t .

In statistical hypothesis testing, the null hypothesis H_t usually states the status quo which is challenged by testing it. If H_t is rejected, this can be seen as evidence that H_t is not true. This evidence is used to draw conclusions and derive recommendations for the future. However, if H_t is accepted, this is typically not seen as evidence that H_t is true, but just as insufficient evidence to reject H_t . Therefore, actions are typically not changed based on accepted hypotheses. For these reasons, we believe that reversing rejections should be forbidden, while accept-to-reject changes are useful in many applications.

For example, an important online multiple testing application is A/B testing in tech companies [18, 19]. Here, a rejection might lead to an adjustment of a website or an app. However, after an acceptance there won't be any changes. Consequently, reversing rejections can lead to undoing unnecessary adjustments made to a website, which can cost money and time, whereas turning accepted hypotheses into rejected ones may lead to additional improvements of the website at a later point. The same logic applies to many other online multiple testing tasks such as platform trials [23, 35] and genomic studies [2, 1].

Online multiple testing procedures are a special case of online ARC procedures, where we disallow all changes to decisions. Sometimes, online multiple testing procedures are suitable for the applications, leading to a rich literature on the topic, but ARC procedures also have their place, and this particular paper focuses on designing online ARC procedures.

1.2 Problem Setup

Let $I_0 \subseteq \mathbb{N}$ be the index set of true hypotheses. The false discovery proportion (FDP) of a rejection set R is defined as

$$\text{FDP}(R) := \frac{|R \cap I_0|}{|R| \vee 1}.$$

The false discovery rate (FDR) is the expected FDP: $\text{FDR}(R) = \mathbb{E}[\text{FDP}(R)]$. We let FDP_t and FDR_t be shorthand for $\text{FDP}(R_t)$ and $\text{FDR}(R_t)$, respectively, for each $t \in \mathbb{N}$. The type FDR control usually considered in online multiple testing is defined as:

$$\text{OnlineFDR} := \sup_{t \in \mathbb{N}} \text{FDR}_t.$$

Javanmard and Montanari [17] provided an algorithm that controlled OnlineFDR below a fixed level $\alpha \in [0, 1]$ for independent p-values. However, this does not guarantee FDR control at any data-adaptive stopping time (like the time of the hundredth rejection); this is severe restriction in the online setting. For this reason, Xu and Ramdas [32] introduced the SupFDR, defined as

$$\text{SupFDR} := \mathbb{E} \left[\sup_{t \in \mathbb{N}} \text{FDP}_t \right].$$

We can also define a notion of FDR for arbitrary (data-dependent) stopping times

$$\text{StopFDR} := \sup_{\tau \in \mathcal{T}} \text{FDR}_\tau,$$

where \mathcal{T} denotes the set of all stopping times. In this paper, we will generally use $\alpha \in (0, 1]$ to denote a constant used to induce a fixed level of error control unless otherwise specified. We note that enforcing control of SupFDR and StopFDR have the following relationship,

$$\begin{aligned} \text{SupFDR} \leq \alpha &\implies \text{StopFDR} \leq \alpha \\ \text{StopFDR} \leq \alpha &\implies \text{SupFDR} \leq \alpha(1 + \log(\alpha^{-1})). \end{aligned}$$

Controlling the SupFDR ensures valid FDR control at arbitrary (data-dependent) stopping times since the supremum of FDP_t will exceed FDP_τ at any single time τ . Interestingly, a reverse relationship also holds, meaning control of the FDR at stopping times implies control of the SupFDR, albeit at an inflated level (we prove this result in Appendix G.4).

In addition to the aforementioned false discovery metrics that involve all time steps, we will also consider the SupFDR until some fixed time K , defined as

$$\text{SupFDR}^K := \mathbb{E} \left[\sup_{1 \leq t \leq K} \text{FDP}_t \right].$$

The SupFDR^K can be useful if there is a fixed maximum number of hypotheses K , but we may want to stop early.

1.3 Our contributions

In this paper we introduce the online e-BH and online BH procedure, online ARC versions of the e-BH and BH procedures [31, 3].

The e-BH procedure was recently proposed by Wang and Ramdas [31] and provides an e-value analog of the seminal p-value based Benjamini-Hochberg (BH) procedure [3]. We show that online e-BH controls the SupFDR under arbitrary dependence between the e-values (Section 2.1). To the best of our knowledge, this is the first nontrivial online procedure with SupFDR control under arbitrary dependence. In addition, it uniformly improves e-LOND, the current state-of-the-art online procedure

for FDR control with arbitrarily dependent e-values. Next, we demonstrate how the “boosting” approach for the *offline* e-BH procedure [31] can be extended to the *online* e-BH procedure (Section 2.3). We also propose a new boosting technique for locally dependent e-values which improves on boosting under arbitrary dependence. A comparison of the guarantees provided by the e-BH and online e-BH procedure is given in Table 1.

Assumption	e-BH	Online e-BH
Independence or PRDS (boosted)	$\text{FDR} \leq \alpha$ [31]	$\text{OnlineFDR} \leq \alpha$ (Sec. 2.3)
Arbitrary dependence	$\text{FDR} \leq \alpha$ [31]	$\text{SupFDR} \leq \alpha$ (Sec. 2.1)
Local dependence (boosted)	Not considered before	$\text{SupFDR} \leq \alpha$ (Sec. 2.3)

Table 1: Comparison of e-BH and online e-BH guarantees.

In contrast, we show that the online BH procedure controls the OnlineFDR under positive regression dependence on a subset (PRDS) (Section 3.1). Furthermore, it controls the SupFDR at slightly inflated levels under weak positive (PRDN) and negative (WNDN) dependence assumptions (Sections 3.2 and 3.3). Under arbitrary dependence, online BH provides SupFDR^K control at an inflated level if a maximum number of hypotheses K is specified (Section 3.4). The guarantees of the BH and online BH procedure are summarized in Table 2.

Assumption	BH	Online BH
Independence	$\text{FDR} \leq \alpha$ [3]	$\text{OnlineFDR} \leq \alpha$ (Sec. 3.1)
PRDS	$\text{FDR} \leq \alpha$ [4]	$\text{OnlineFDR} \leq \alpha$ (Sec. 3.1)
PRDN	$\text{FDR} \leq \alpha(1 + \log(\alpha^{-1}))$ [27]	$\text{SupFDR} \leq \alpha(1 + \log(\alpha^{-1}))$ (Sec. 3.2)
WNDN	$\text{FDR} \leq \alpha(3.18 + \log(\alpha^{-1}))$ [6]	$\text{SupFDR} \leq \alpha(3.18 + \log(\alpha^{-1}))$ (Sec. 3.3)
Arbitrary dependence	$\text{FDR} \leq \alpha\ell_K$ [4]	$\text{SupFDR}^K \leq \alpha\ell_K$ (Sec. 3.4)

Table 2: Comparison of BH and online BH guarantees.

As shown in the Tables 1 and 2, both online procedures, online e-BH and online BH, provide OnlineFDR control whenever their offline counterparts provide FDR control and in some cases even control the SupFDR where the offline methods were only known to control the FDR. In addition, we will show that the *online* versions are natural *generalizations* of the *offline* procedures, meaning e-BH and BH can be obtained by online e-BH and BH for a specific choice of weighting parameters (see Sections 2 and 3). Therefore, we believe the names *online e-BH* and *online BH* are justified for our procedures.

Along the way, we also prove SupFDR control for all other existing online procedures with FDR control under arbitrary dependence. This includes the reshaped LOND procedure [37] and e-LOND [34]. Previously, these methods were only known to control the FDR at fixed times. In addition, we prove SupFDR control of LOND [13] under PRDN and WNDN at slightly inflated levels. We also provide new guarantees for the TOAD algorithm by Fisher [13].

1.4 The TOAD algorithm

While ARC procedures have not been explicitly named and conceptualized before, the only online ARC procedure that we are aware of (which is not already a fully online procedure) is the TOAD algorithm by Fisher [13]. The TOAD algorithm was introduced in the context of p-values and “decision deadlines” where one must make a decision about H_t by a deadline d_t based on a p-value P_t . One can view online procedures as having a deadline for making a decision as setting $d_t = t$, and online ARC procedures as setting $d_t = \infty$. In Fisher’s framework, rejection-to-acceptance changes before the deadline were not explicitly disallowed, although not used by the procedures. In the Appendix B we present the e-TOAD algorithm, an e-value analog of the TOAD procedure with decision deadlines. Then, our online e-BH

procedure can be viewed as an instance of our e-TOAD framework with unbounded deadlines, where e-TOAD extends TOAD and its guarantees from p-values to e-values. Further, while Fisher [13] showed that TOAD for arbitrarily dependent p-values controls the OnlineFDR, we extend his result to show that it actually controls the SupFDR. We also improve the power of TOAD under arbitrary dependence slightly by allowing individual shape functions for the hypotheses. Furthermore, we prove that TOAD for positively dependent p-values controls the SupFDR under PRDN and WNDN at a slightly inflated level.

2 The online e-BH procedure

Let us first review the e-BH procedure for offline multiple testing. Let K e-values E_1, \dots, E_K for the hypotheses H_1, \dots, H_K be given. The (base) e-BH procedure [31] rejects hypotheses with the k^* largest e-values, where

$$k^* = \max \left\{ k \in \{1, \dots, K\} : \sum_{j=1}^K \mathbb{1}\{E_j \geq K/(k\alpha)\} \geq k \right\},$$

with the convention $\max(\emptyset) = 0$. For the online e-BH procedure we weight each hypothesis H_i with a nonnegative constant γ_i such that $\sum_{i \in \mathbb{N}} \gamma_i \leq 1$. Let $\pi_0 := \sum_{i \in I_0} \gamma_i$ be the sum of the null weights. The online e-BH procedure is defined by the rejection sets

$$R_t^{\text{oe-BH}} := \left\{ i \leq t : E_i \geq \frac{1}{k_t^* \alpha \gamma_i} \right\}, t \in \mathbb{N}, \text{ where}$$

$$k_t^* = \max \left\{ k \in \{1, \dots, t\} : \sum_{j=1}^t \mathbb{1}\{E_j \geq 1/(k\alpha\gamma_j)\} \geq k \right\},$$

with the conventions $\max(\emptyset) = 0$ and $1/0 = \infty$. Note that k_t^* is nondecreasing in t , which implies that $(R_t^{\text{oe-BH}})_{t \in \mathbb{N}}$ is a nested sequence of rejection sets such that the online e-BH procedure is indeed an online ARC procedure. Also note that the naive procedure that sets R_t to be the rejection set of the offline e-BH procedure applied to E_1, \dots, E_t is not an online ARC procedure: it may change its mind in both directions (possibly more than once).

If we choose $\gamma_i = 1/K$ for $i = 1, \dots, K$ and $\gamma_i = 0$ otherwise, the online e-BH becomes the e-BH procedure, and in that sense it is an *online generalization of the e-BH procedure*. In particular, the online e-BH allows to choose different weights for the hypotheses and also to test an infinite number of hypotheses. Wang and Ramdas [31] have already formulated a weighted e-BH procedure which is similar to this method; however, their approach does not immediately apply to an infinite number of hypotheses while we claim that our approach defines an online ARC procedure with SupFDR control.

2.1 SupFDR control under arbitrary dependence

To show SupFDR control, we first define the following class of procedures.

Definition 2. A discovery set of finite cardinality, $R \subseteq \mathbb{N}$, is considered *self-consistent* if it satisfies $E_t \geq (\alpha \gamma_t |R|)^{-1}$ for e-values (or $P_t \leq \alpha \gamma_t |R|$ for p-values) for each $t \in R$. We denote by $\mathcal{R}(\alpha)$ the set of discovery sets that are self-consistent at level α .

The above definition is consistent with the definition of self-consistency as defined by Blanchard and Roquain [5] and generalizes the definition (referred to as “compliance”) introduced in Su [27]. Wang and Ramdas [31] showed that self-consistent e-value procedures control FDR under arbitrary dependence in the offline multiple testing setting with uniform weights. We now show an extension of this result for general weights $(\gamma_t)_{t \in \mathbb{N}}$ and any self-consistent rejection set in an infinite stream of hypotheses.

Proposition 2.1. *For a stream of infinite hypotheses and its corresponding arbitrarily dependent e-values $(E_t)_{t \in \mathbb{N}}$, $\mathbb{E} \left[\sup_{R \in \mathcal{R}(\alpha)} \text{FDP}(R) \right] \leq \alpha$.*

Proof. We first prove it for the case where we have K hypotheses. Since $\mathcal{R}(\alpha)$ is finite, there exists a R^* that maximizes the FDP. Thus we have that

$$\begin{aligned} \mathbb{E} \left[\sup_{R \in \mathcal{R}(\alpha)} \text{FDP}(R) \right] &= \text{FDR}(R^*) \\ &= \sum_{i \in I_0} \mathbb{E} \left[\frac{\mathbb{1}\{E_i \geq 1/(\alpha\gamma_i|R^*|)\}}{|R^*| \vee 1} \right] \\ &\leq \alpha \cdot \sum_{i \in I_0} \gamma_i \mathbb{E}[E_i] \leq \pi_0 \alpha \leq \alpha. \end{aligned} \quad (1)$$

Where the first inequality follows because $\mathbb{1}\{x \geq 1\} \leq x$ for all $x \geq 0$, the second follows from E_i being an e-value for each $i \in I_0$, and the last is due to $\pi_0 \leq 1$.

When we have an infinite stream of hypotheses, we know the following is true

$$\sup_{R \in \mathcal{R}(\alpha)} \text{FDP}(R) = \lim_{t \rightarrow \infty} \sup_{R \in \mathcal{R}(\alpha), R \subseteq [t]} \text{FDP}(R).$$

Thus, by taking an expectation over all the terms in the above equation, we get

$$\mathbb{E} \left[\sup_{R \in \mathcal{R}(\alpha)} \text{FDP}(R) \right] = \lim_{t \rightarrow \infty} \mathbb{E} \left[\sup_{R \in \mathcal{R}(\alpha), R \subseteq [t]} \text{FDP}(R) \right] \leq \pi_0 \alpha \leq \alpha,$$

where the equality is by the Fatou-Lebesgue theorem and the last two inequalities are by (1). \square

This immediately implies SupFDR control for online e-BH, since the rejection sets $R_t^{\text{oe-BH}}$ are self-consistent for all t .

Theorem 2.2. *The online e-BH procedure controls the SupFDR at level α under arbitrary dependence between the e-values.*

The existing state-of-the-art method for OnlineFDR control under arbitrary dependence of the e-values is e-LOND [34]. In the following subsection, we show that e-LOND even provides SupFDR control and that online e-BH uniformly improves e-LOND.

2.2 Online e-BH uniformly improves e-LOND

The e-LOND algorithm [34] rejects hypothesis H_t , if $E_t \geq (\alpha_t^{\text{e-LOND}})^{-1}$, where

$$\alpha_t^{\text{e-LOND}} := \alpha \gamma_t (|R_{t-1}^{\text{e-LOND}}| + 1) \quad \text{and} \quad R_t^{\text{e-LOND}} := \{i \leq t : E_i \geq (\alpha_i^{\text{e-LOND}})^{-1}\}.$$

It follows immediately that $R_t^{\text{e-LOND}} \subseteq R_t^{\text{e-BH}}$ for all $t \in \mathbb{N}$, where $R_t^{\text{e-BH}}$ is the rejection set of the online e-BH procedure at step t . Hence, online e-BH uniformly improves e-LOND, meaning that the former makes at least the same discoveries as the latter on every run. Xu and Ramdas [34] only proved FDR control for e-LOND at fixed times. We note that e-LOND is also a self-consistent procedure so its SupFDR control follows immediately from Proposition 2.1.

Proposition 2.3. *The e-LOND procedure [34] controls SupFDR at level α under arbitrary dependence between the e-values.*

The power gain of the online e-BH method compared to e-LOND is illustrated in Figure 1. The left plot shows that the online e-BH method substantially improves e-LOND when the signal of the alternative hypotheses is weak. This improvement is slightly smaller when the signal strong (right plot). The simulation setup is in Appendix F.

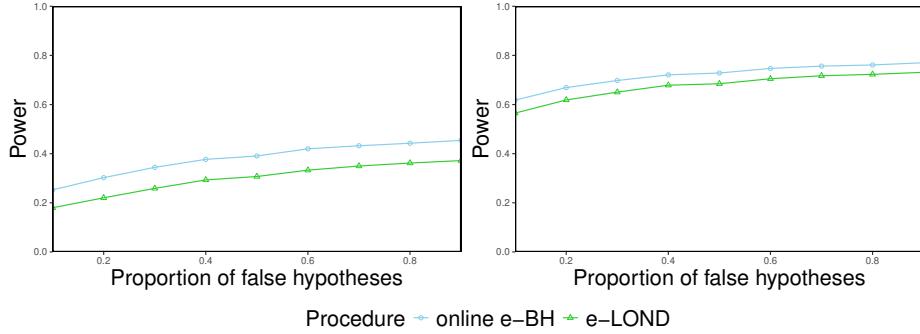


Figure 1: Power comparison of online e-BH and e-LOND for different proportions of false hypotheses. In the left (right) plot, the signal of the alternative is weak (strong).

Remark 1. Xu and Ramdas [33, 34] improved the e-BH and e-LOND procedure, respectively, using a randomization approach. The idea is to define a randomized e-value as

$$S_{\hat{\alpha}_t}(E_t) := \max(\mathbb{1}\{E_t \geq \hat{\alpha}_t\}E_t, \mathbb{1}\{U \leq E_t \hat{\alpha}_t\}\hat{\alpha}_t^{-1}),$$

where $\hat{\alpha}_t$ is a potentially data-dependent significance level obtained by e-BH or e-LOND and U a uniform random variable on $[0, 1]$ that is independent of E_t and $\hat{\alpha}_t$. Xu and Ramdas [33] showed that if E_t is an e-value, then $S_{\hat{\alpha}_t}(E_t)$ is a valid e-value as well. Furthermore, $S_{\hat{\alpha}_t}(E_t) \geq \hat{\alpha}_t^{-1}$ iff $E_t \geq \hat{\alpha}_t^{-1}U$, implying that e-BH and e-LOND are uniformly more powerful when applied to $S_{\hat{\alpha}_t}(E_t)$ instead of E_t . One could consider applying the same idea to online e-BH. However, note that we do not have one data dependent level with the e-BH but a sequence of levels that is increasing overtime due to future rejections. Hence, one would need to fix a time at which the stochastic rounding approach is applied. For example, one could set

$$\hat{\alpha}_t := \alpha\gamma_t(k_{t-1}^* + 1).$$

In this case stochastic rounding would increase the probability of rejecting H_t at time t . However, if we do not reject, then $S_{\hat{\alpha}_t}(E_t) = 0$ and there would not be any chance to obtain a rejection at a later stage. In this case, online e-BH becomes equivalent to the randomized e-LOND procedure. One could also consider applying stochastic rounding at the (potentially data-adaptive) time τ at which entire testing process was stopped. In this case, stochastic rounding could only lead to additional rejections at the very end and online e-BH becomes a weighted generalization of the randomized e-BH procedure by Xu and Ramdas [33]. In the next section, we consider a deterministic rounding approach which also leads to improved e-values. However, in contrast to stochastic rounding, it does not allow to round to data-adaptive levels.

2.3 Boosting the online e-BH procedure

If additional information about the marginal or joint distribution of the e-values is available, then the online e-BH procedure can be improved. We call this process “boosting” in line with the same term being used for e-BH [31]. We demonstrate the following three forms of boosting in this paper.

1. Boosting when information about the marginal distribution is available.
2. Boosting when information about the marginal distribution is available and the e-values are locally dependent.
3. Boosting when information about the marginal distribution is available and the e-values are positive regression dependent on a subset (PRDS).

While the first and third form are online extensions of the boosting techniques by Wang and Ramdas [31], the second form of boosting is new and particularly useful in the online setting. Online e-BH controls the SupFDR under the first two boosting types and the OnlineFDR under the third boosting type. In the remaining part of the section, we describe the first form of boosting, which is the cornerstone for all three methods. The details for the second and third form are provided in Appendix A.

We first reformulate the boosting method by Wang and Ramdas [31] which makes it easier to use it for our purposes; however, our approach is equivalent to theirs. Let K be the number of hypotheses. Define the truncation function $T : [0, \infty] \rightarrow [0, K/\alpha]$ as

$$T(x) := \sum_{k=1}^K \mathbb{1} \left\{ \frac{K}{k\alpha} \leq x < \frac{K}{(k-1)\alpha} \right\} \frac{K}{k\alpha} \quad \text{with } T(\infty) := \frac{K}{\alpha},$$

where we again use the convention $1/0 = \infty$. Note that $T(x)$ only takes values in the set $\{K/(\alpha k) : k \in \{1, \dots, K\}\} \cup \{0\}$. Let X_1, \dots, X_K be any nonnegative random variables (not necessarily e-values). The idea of boosting is that e-BH applied on X_1, \dots, X_K rejects the same hypotheses as if applied on $T(X_1), \dots, T(X_K)$. Hence, if $T(X_1), \dots, T(X_K)$ are valid e-values for H_1, \dots, H_K , e-BH applied on X_1, \dots, X_K controls the SupFDR (even though the latter may not be e-values). We call such random variables X_t *boosted e-values*. A simple approach is to set $X_t = b_t E_t$, where E_t is a valid e-value, for some boosting factor $b_t \geq 1$. Note that such a boosting factor always exists, since $T(E_t) \leq E_t$.

For a given $(\gamma_t)_{t \in \mathbb{N}}$ we can simply extend this approach to the online e-BH procedure by defining the truncation function for the t -th e-value as

$$T_t(x) := \sum_{k=1}^{\infty} \mathbb{1} \left\{ \frac{1}{k\alpha\gamma_t} \leq x < \frac{1}{(k-1)\alpha\gamma_t} \right\} \frac{1}{k\alpha\gamma_t}, \quad \text{with } T_t(\infty) := \frac{1}{\alpha\gamma_t}, \quad (2)$$

and the convention $0 \cdot \infty = 0$. In case of $\gamma_t = 1/K$, $t \leq K$, we have $T_t = T$ for all $t \leq K$. We capture the result about boosted e-values in the following proposition.

Proposition 2.4. *Let X_1, X_2, \dots be a sequence of nonnegative random variables such that $\mathbb{E}_{H_t}[T_t(X_t)] \leq 1$ for all $t \in \mathbb{N}$. Then the online e-BH procedure applied to X_1, X_2, \dots controls the SupFDR at level α .*

Due to the infinite sum in (2) it might be difficult to determine boosted e-values exactly. In Appendix A.1, we propose two different approaches to handle this and demonstrate the application in a Gaussian testing setup.

A power comparison of online e-BH with boosted and non-boosted e-values is shown in Figure 2. The left plot shows the results for weak signals and the right plot for strong signals. Boosting improves the power in both cases.

2.4 Online e-BH with p-values

In the following, we consider applying the online e-BH procedure to $\psi_1(P_1), \psi_2(P_2), \dots$, where each P_t is a p-value for H_t and $\psi_t : [0, 1] \rightarrow [0, \infty]$, $t \in \mathbb{N}$ is a function with $\psi_t(0) = \infty$. Wang and Ramdas [31] called ψ_t a *decreasing transform* and (for simplicity) assumed that $\psi_1 = \psi_2 = \dots = \psi$ for a strictly decreasing and continuous function ψ . Here, we allow a different function ψ_t for each hypothesis H_t and only assume that ψ_t is nonincreasing and left-continuous. To clarify this change of definition, we use the term *nonincreasing transform* instead. Define the generalized inverse of a nonincreasing transform ψ_t as

$$\psi_t^{-1}(x) := \max\{u \in [0, 1] : \psi_t(u) \geq x\}.$$

With this, boosting (see Section 2.3) can be applied to derive the following result, which is formally proven in the Appendix G.1.

Proposition 2.5. *Let ψ_1, ψ_2, \dots be nonincreasing transforms.*

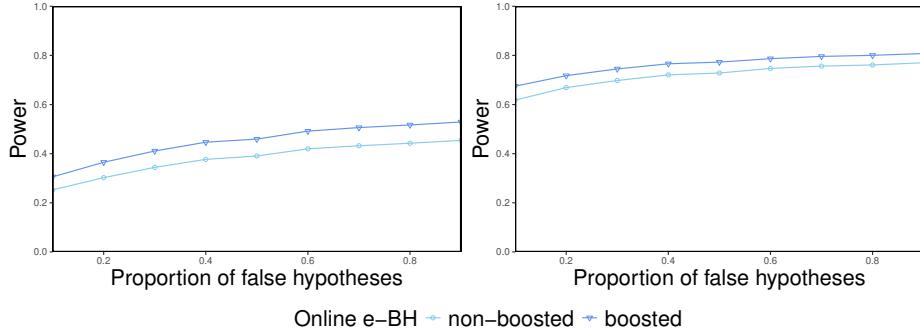


Figure 2: Power comparison of online e-BH with non-boosted and boosted e-values for different proportions of false hypotheses. In the left (right) plot, the signal of the alternative is weak (strong). The simulation setup is described in Appendix F.

(i) If for every t it holds that

$$\sum_{k=1}^{\infty} \frac{1}{k\alpha\gamma_t} \left(\psi_t^{-1} \left(\frac{1}{k\alpha\gamma_t} \right) - \psi_t^{-1} \left(\frac{1}{(k-1)\alpha\gamma_t} \right) \right) \leq 1,$$

then the online e-BH procedure applied to the random variables $\psi_1(P_1), \psi_2(P_2), \dots$ controls the SupFDR at level α .

(ii) If the p-values P_1, P_2, \dots are PRDS (see Appendix E for a formal definition of PRDS) and for all t it holds that

$$\sup_{k \in \mathbb{N}} \frac{1}{k\alpha\gamma_t} \psi_t^{-1} \left(1/(k\alpha\gamma_t) \right) \leq 1,$$

then the online e-BH procedure applied to the random variables $\psi_1(P_1), \psi_2(P_2), \dots$ controls the OnlineFDR at level α .

In the offline case, the Benjamini-Hochberg (BH) procedure [3] is equivalent to the e-BH procedure applied to the random variables $1/P_1, \dots, 1/P_k$ [31]. One could derive an online BH procedure in the same way by setting the decreasing transforms to $\psi_t(x) = 1/x = \psi_t^{-1}(x)$. Proposition 2.5 immediately implies that this online BH procedure controls the OnlineFDR, if the p-values are PRDS. In the next section, we study the behavior of the online BH procedure more deeply.

Remark 2. Benjamini and Yekutieli [4] introduced an offline p-value based multiple testing procedure with FDR control (BY procedure) and Blanchard and Roquain [5] generalized this using a shape function β (BR procedure). Let $\beta(k) = \int_0^k x d\nu(x)$, $k \in \{1, \dots, K\}$, for some probability measure ν on $(0, \infty)$. Blanchard and Roquain [5] proved that FDR control is guaranteed under arbitrary dependence of the p-values, if the hypotheses with the k^* smallest p-values are rejected, where

$$k^* = \max \left\{ k \in \{1, \dots, K\} : \sum_{j=1}^K \mathbb{1}\{P_j \leq \beta(k)\alpha/K\} \geq k \right\}.$$

The BY procedure is obtained by setting $\beta(k) = k/\ell_K$, where $\ell_K := \sum_{i=1}^K \frac{1}{i} \approx \log(K)$. Fisher [13] generalized the method by Blanchard and Roquain [5] to the online ARC case in the same manner as we did for the e-BH procedure by using general weights $(\gamma_t)_{t \in \mathbb{N}}$ instead of $1/K$ for each hypothesis (see (6) in Appendix for $d_t = \infty$). We call this the online BR procedure in the following. However,

Fisher [13] only showed FDR control for the online BR procedure at fixed times. In Appendix G.2, we show how the online BR procedure can be written as a special case of the online e-BH procedure using Proposition 2.5 (i) and therefore provides SupFDR control.

Remark 3. Javanmard and Montanari [16] introduced a version of their LOND algorithm with OnlineFDR control under arbitrary dependence and later Zrnic et al. [37] generalized this method to the reshaped LOND (r-LOND) procedure. The r-LOND procedure rejects hypothesis H_t , if $P_t \leq \alpha_t^{\text{r-LOND}}$, where

$$\alpha_t^{\text{r-LOND}} := \alpha \gamma_t \beta_t (|R_{t-1}^{\text{r-LOND}}| + 1) \quad \text{and} \quad R_{t-1}^{\text{r-LOND}} := \{i \leq t-1 : P_i \leq \alpha_i^{\text{r-LOND}}\}$$

and β_t is some shape function. In Appendix G.2 we show that r-LOND could be derived as a special case of a boosted e-LOND procedure and therefore also provides SupFDR control under arbitrary dependence of the p-values.

3 The online BH procedure

In the same manner as for e-values, one can define an online ARC version of the p-value based Benjamini-Hochberg (BH) procedure [3], or the *online BH procedure* for short. Recall, as before, that $(\gamma_t)_{t \in \mathbb{N}}$ is a nonnegative sequence that sums to one. Then, define:

$$R_t^{\text{OBH}} := \{i \leq t : P_i \leq k_t^* \alpha \gamma_i\}, t \in \mathbb{N}, \text{ where}$$

$$k_t^* = \max \left\{ k \in \{1, \dots, t\} : \sum_{j=1}^t \mathbb{1}\{P_j \leq k \alpha \gamma_j\} \geq k \right\}$$

with the convention $\max(\emptyset) = 0$. Again, note that applying the offline BH procedure to the p-values P_1, \dots, P_t at each time t would not define an online ARC procedure.

The online BH procedure is a special instance of the TOAD framework under positive dependence by Fisher [13], which is obtained if no decision deadlines are specified for the TOAD algorithm, meaning that $d_t = \infty$ for all $t \in \mathbb{N}$. However, the great generality of the TOAD framework leads to a certain notational overhead and complexity that is unnecessary to define and analyze the online BH procedure, leading to this special case being lost. The online BH procedure, we argue, is a simple and important algorithm which deserves particular emphasis.

As with the online e-BH procedure, if we set $\gamma_1, \dots, \gamma_K$ to equal $1/K$ and the rest of the γ sequence to equal zero, one simply obtains the BH procedure at time K , so online BH is a *generalization* of the BH procedure. Although OnlineFDR control of the online BH procedure under PRDS is implicitly proven by Fisher [13], we also provide a short, direct proof in the next subsection. Furthermore, we prove that the online BH procedure additionally controls the SupFDR in case of PRDN (Section 3.2) and WNDN (Section 3.3) p-values, and SupFDR K for arbitrary p-values (Section 3.4).

3.1 OnlineFDR control under positive dependence

The most general condition under which FDR control of the BH procedure is usually proven is *positive regression dependence on a subset* (PRDS) [4] (see Definition 3 in Appendix E). The online BH procedure also controls the OnlineFDR under PRDS.

Proposition 3.1. *If the null p-values are PRDS, the online BH procedure ensures that OnlineFDR $\leq \alpha$.*

Proof. For any $t \in \mathbb{N}$, we have

$$\text{FDR}_t = \sum_{i \leq t, i \in I_0} \mathbb{E} \left[\frac{\mathbb{1}\{P_i \leq k_t^* \alpha \gamma_i\}}{k_t^*} \right] \leq \sum_{i \leq t, i \in I_0} \alpha \gamma_i \leq \alpha,$$

where the first inequality follows from Lemma 1 (b) of Ramdas et al. [21]. \square

In the next subsection, we prove that online BH controls the SupFDR at a slightly inflated level under a weaker assumption than PRDS.

Remark 4. In the same manner as online e-BH uniformly improves e-LOND (Section 2.2), online BH uniformly improves the LOND procedure introduced by Javanmard and Montanari [16] which controls the OnlineFDR under PRDS [37]. The LOND procedure is exactly defined as e-LOND but with p-values instead of e-values, meaning H_t is rejected if $P_t \leq \alpha \gamma_t (|R_{t-1}^{\text{LOND}}| + 1)$ where $|R_{t-1}^{\text{LOND}}|$ is the number of previous rejections. Another uniformly more powerful online procedure than LOND is provided by the LORD algorithm [17, 19], which is only known to control the OnlineFDR under independence of the p-values. In Appendix C we compare LORD to our online BH procedure.

Remark 5. It is easy to see from the proof above that online BH even controls the OnlineFDR at level $\alpha \pi_0$. In order to avoid this conservatism in the offline case, Storey and colleagues [25, 26] introduced an adaptive BH procedure by estimating π_0 . In Appendix D, we introduce the natural online generalization of this Storey-BH procedure and prove its OnlineFDR control under independence of the p-values.

3.2 SupFDR control under positive dependence

We will now show that the online BH procedure can also control the SupFDR, albeit at a slightly inflated level. Note that R_t^{oBH} is clearly a self-consistent set (just like online e-BH) — it is the largest such set when R is restricted to be a subset of $\{1, \dots, t\}$. Recall that $\mathcal{R}(\alpha)$ is the set of self-consistent rejection sets. The following theorem provides a weighted version of the **FDR-linking** theorem (i.e., weighted analog of Theorems 1 and 2 in Su [27]).

Theorem 3.2 (Weighted FDR-linking theorem). *Under any dependence structure among p-values $(P_t)_{t \in [K]}$ for fixed $K \in \mathbb{N}$, we can ensure the following is true:*

$$\mathbb{E} \left[\sup_{R \in \mathcal{R}(\alpha)} \text{FDP}(R) \right] \leq \pi_0 \alpha + \pi_0 \alpha \int_{\pi_0 \alpha}^1 \frac{\text{FDR}_0(x)}{x^2} dx,$$

where $\text{FDR}_0(x)$ is the FDR of the weighted BH procedure applied only to the p-values corresponding to the null hypotheses, i.e., I_0 , with weights $(\gamma_t / \pi_0)_{t \in I_0}$.

The proof of Theorem 3.2 is provided in Appendix G.3. We can use the weighted **FDR-linking** theorem to prove an upper bound of the expected maximum FDP over all self-consistent discovery sets under *positive regression dependence on the nulls* (PRDN). PRDN is a positive dependence assumption that is restricted to the dependency between null p-values and is therefore weaker than PRDS (see Appendix E for a definition of PRDN).

Proposition 3.3. *For PRDN p-values $(P_t)_{t \in \mathbb{N}}$, we have that*

$$\mathbb{E} \left[\sup_{R \in \mathcal{R}(\alpha)} \text{FDP}(R) \right] \leq \pi_0 \alpha (1 + \log((\pi_0 \alpha)^{-1})) \leq \alpha (1 + \log(\alpha^{-1})).$$

Proof. Similar to the proof of Proposition 2.1, we first prove the above result restricted to a finite set of K hypotheses. Since under positive regression dependence and the global null, the weighted Simes p-value is a valid p-value, we have that $\text{FDR}_0(x) = x$ under PRDN. Then, we get the upper bounds in the proposition statement for the finite K hypotheses by carrying out the integration in the upper bound of the weighted **FDR-linking** theorem and the fact that the function $f(x) = x(1 + \log(1/x))$ is increasing in $(0, 1]$. Now, we apply the same limit argument as seen in Proposition 2.1 to show this upper bound extends to the infinite hypotheses case as well. Thus, we have shown our desired result. \square

The SupFDR control of the online BH procedure is now a direct result of Proposition 3.3.

Theorem 3.4. *The online BH procedure ensures that $\text{SupFDR} \leq \alpha(1 + \log(\alpha^{-1}))$ when the p-values $(P_t)_{t \in \mathbb{N}}$ are PRDN.*

Note that the bound in Theorem 3.4 is sharp in the following sense.

Theorem 3.5. *For every $\varepsilon > 0$, there exists a $K_0 \in \mathbb{N}$, and $\alpha' \in (0, 1]$ such that, for online BH,*

$$\text{SupFDR} \geq \text{StopFDR} > (1 - \varepsilon)\alpha(1 + \log(\alpha^{-1}))$$

for all $\alpha \leq \alpha'$ where the first K_0 hypotheses are null and have p-values that are i.i.d. uniform random variables, and the remaining hypotheses are non-null with p-values set to 0. Note that this p-value distribution satisfies the PRDN condition, meaning that the SupFDR upper bound in Theorem 3.4 is unimprovable.

We defer the proof of this result to Appendix G.5.

Notably, the p-value configuration constructed in Theorem 3.5 ensures that all p-values are independent. Hence, it is not even possible to improve the SupFDR bound in Theorem 3.4 under stronger assumptions like PRDS or independence.

3.3 SupFDR control under negative dependence

We can also show that online BH controls the SupFDR under some kind of negative dependence, following a similar approach to Chi et al. [6]. Like positive dependence, there are many notions of negative dependence. A comprehensive overview of the different notions and a comparison between these is given by Chi et al. [6]. In addition, they showed that the BH procedure controls the FDR at a slightly inflated level, if the p-values are *weakly negatively dependent on the nulls* (WNDN) (see Appendix E for a formal definition of WNDN). The following result can be shown about weighted self-consistent procedures under negative dependence.

Proposition 3.6. *For WNDN p-values $(P_t)_{t \in \mathbb{N}}$, we have that*

$$\mathbb{E} \left[\sup_{R \in \mathcal{R}(\alpha)} \text{FDP}(R) \right] \leq \alpha(3.18 + \log(\alpha^{-1})).$$

The above result, when restricted to a finite set of hypotheses, follows from the proof of Theorem 19 in Chi et al. [6], as the proof utilizes an application of the **FDR-linking** theorem with a bound on the type I error of the Simes p-value under WNDN. Prop. 12 in Chi et al. [6] shows that the weighted Simes p-value has identical type I error under WNDN as unweighted Simes, so we can apply the weighted **FDR-linking** theorem to obtain the upper bounds in the proposition for finite hypotheses using an identical argument to the aforementioned proof. Combining this result with the limit argument seen in Proposition 3.3, we can obtain the full result for an infinite stream of hypotheses.

The following theorem shows that online BH even controls SupFDR under WNDN at an inflated level and is a direct result of Proposition 3.6.

Theorem 3.7. *The online BH procedure ensures that $\text{SupFDR} \leq \alpha(3.18 + \log(\alpha^{-1}))$ when the p-values $(P_t)_{t \in \mathbb{N}}$ are WNDN.*

The sharpness result in Theorem 3.5 holds on an instance of independent p-values (which are WNDN) that achieve SupFDR and StopFDR which approaches $\alpha(1 + \log(\alpha^{-1}))$. Thus, the dependence on $\log(\alpha^{-1})$ cannot be improved, but it may still be possible to show a sharper constant than 3.18.

Propositions 3.3 and 3.6 directly imply that the LOND algorithm [16, 37] (see Remark 4) and the TOAD algorithm for positively dependent p-values [13] (see (6) in Appendix B with $\beta_t(k) = k$) also have the prescribed SupFDR bounds, since both also satisfy the weighted self-consistency condition.

Proposition 3.8. *LOND and TOAD ensure that $\text{SupFDR} \leq \alpha(1 + \log(\alpha^{-1}))$ when $(P_t)_{t \in \mathbb{N}}$ are PRDN, and $\text{SupFDR} \leq \alpha(3.18 + \log(\alpha^{-1}))$ when $(P_t)_{t \in \mathbb{N}}$ are WNDN.*

3.4 SupFDR^K control under arbitrary dependence

In Remark 2 we mentioned that the online BR procedure controls the SupFDR, which is formally proven in Appendix G.2. Suppose we have some fixed number $K \in \mathbb{N}$ of hypotheses and let $\beta(k) = k/\ell_K$ for $k \in \{1, \dots, K\}$, where $\ell_K = \sum_{i=1}^K \frac{1}{i} \approx \log(K)$. It follows immediately, that online BR applied at level $\alpha\ell_K$ rejects exactly the same hypotheses as online BH applied at level α . Hence, online BH controls the SupFDR^K at the inflated level $\alpha\ell_K$.

Theorem 3.9. *The online BH procedure ensures that $\text{SupFDR}^K \leq \alpha\ell_K$ for all $K \in \mathbb{N}$ under arbitrarily dependent p-values.*

Online BH does not control the SupFDR under arbitrary dependence if the number of hypotheses is infinite. To see this, note that Guo and Rao [15, Theorem 5.1 (iii)] showed that there exist p-value configurations such that the FDR of the BH procedure equals exactly $\min(\alpha\ell_K, 1)$, where K is the number of total hypotheses. Hence, for every α there is a K such that the FDR of the BH procedure equals 1. Since the BH procedure is a special case of online BH for every K , the online BH procedure can neither control the OnlineFDR, the StopFDR nor the SupFDR at a level below 1 under arbitrary dependence.

4 Related literature

In the following, we give a brief overview of the existing online multiple testing literature divided according to the different settings and assumptions considered in this paper.

Independence. The online multiple testing framework was introduced by Foster and Stine [14] and initially focused on online FDR control at fixed times and with independent p-values [17, 20, 19]. The most popular procedures in this setting are the LORD algorithm, which was introduced by Javanmard and Montanari [17] and later generalized and improved by Ramdas et al. [19], and the SAFFRON algorithm, which was introduced by Ramdas et al. [20]. We compare the (improved) LORD algorithm to online BH in Appendix C and SAFFRON to online Storey-BH in Appendix D. We conclude that the existing methods and our online ARC procedures lead to similar power, however, our online ARC procedures provide a more balanced and better interpretable weighting. In addition, online BH even controls the OnlineFDR under PRDS and allows to control the SupFDR at a slightly inflated level. Such results are not known for LORD nor SAFFRON.

PRDS. The LOND algorithm was introduced by Javanmard and Montanari [16] and proven to control the OnlineFDR under PRDS by Zrnic et al. [37]. The TOAD algorithm (for PRDS p-values) by Fisher [13] generalizes LOND by allowing delayed decisions. We proved that LOND and TOAD also control the SupFDR at a slightly inflated level (under PRDN and WNDN), allowing them to stop data-adaptively at any time (Section 3.3).

Arbitrary dependence. Javanmard and Montanari [16] introduced a modified version of the LOND algorithm that provides OnlineFDR control under arbitrary dependence, which was later generalized by Zrnic et al. [37] to the r-LOND algorithm. With TOAD (for arbitrarily dependent p-values), Fisher [13] introduced a modification of r-LOND that allows to incorporate decision deadlines. Recently, Xu and Ramdas [34] introduced e-LOND, providing OnlineFDR control with arbitrarily dependent e-values. In this paper we showed that r-LOND (Appendix G.2), TOAD (Appendix B) and e-LOND (Section 2.2) even provide SupFDR control without any additional assumptions and therefore can be stopped at arbitrary stopping times.

FDR control at stopping times. Xu and Ramdas [32] were the first considering FDR control at arbitrary stopping times. They introduced SupLORD and showed that it provides valid SupFDR control. However, they rely on the assumption that the null p-values are valid conditional on all past p-values, which is a strong assumption close to independence [12]. Other works also considered FDR control at particular data-adaptive stopping times [37, 13], all relying on some assumption of

independence or conditional validity. Our paper is the first bringing together arbitrary/positive/negative dependence and stopping times.

Batching and delayed decisions. Zrnic et al. [36] and Fisher [13] explored the possibility of increasing power of online procedures by batching multiple hypotheses and allowing delayed decisions, respectively. They introduced online frameworks where decisions, both acceptances and rejections, can be reversed until some prespecified decision deadlines. Although Fisher [13] noted that his TOAD algorithm will never turn a rejected hypothesis into an accepted one, the online ARC framework was not explicitly mentioned before.

E-values. Multiple testing with e-values was initially considered in the offline setting. Wang and Ramdas [31] provided with e-BH an analog of the popular Benjamini-Hochberg (BH) procedure [3]. Vovk and Wang [30, 29] considered merging functions for e-values and applying the closure principle with these. Recently, e-values were also brought up in online multiple testing. Xu and Ramdas [34] introduced the aforementioned e-LOND and Fischer and Ramdas [10] considered online closed testing with e-values, proving that e-values are essential for an online true discovery guarantee.

5 Conclusion

Online ARC multiple testing is a flexible and powerful concept. It allows to test hypotheses one at a time, while ensuring that a rejection remains, regardless of the future. However, by changing the decision for accepted hypotheses based on future information online ARC procedures are more powerful than classical online procedures. We showed that the e-BH and BH procedure possess natural online ARC generalizations.

The online e-BH method controls the SupFDR under arbitrary dependence between the e-values. This allows to apply the online e-BH procedure at any e-values and to stop data-adaptively at any time while ensuring valid FDR control. This offers a much higher flexibility than classical online FDR control at fixed times. We proved that the same flexibility is provided by the existing online procedures r-LOND [37], e-LOND [34], online BR [5, 13] and TOAD for arbitrarily dependent p-values [13]. Previously, data-adaptive stopping was only proven under some independence assumption about the joint distribution of the test statistics.

The p-value based online BH procedure controls the SupFDR at a slightly inflated level when the p-values are PRDN or WNDN. The same holds for the LOND [17] and the TOAD (for positively dependent p-values) [13] algorithm.

In Table 3 we list all online procedures with proven SupFDR control under different dependence assumptions between the p-values or e-values.

		Arb. dep.	Pos. dep.	Neg. dep.	Cond. val.
e-value	Online e-BH	Section 2.1	\implies	\implies	\implies
	e-LOND [34]	Section 2.2	\implies	\implies	\implies
	e-TOAD	Section B	\implies	\implies	\implies
p-value	Online BH	-	Section 3.2*	Section 3.3*	-
	Online BR	Section G.2	\implies	\implies	\implies
	r-LOND [37]	Section G.2	\implies	\implies	\implies
	TOAD (arb. dep.) [13]	Section B	\implies	\implies	\implies
	TOAD (pos. dep.) [13]	-	Section 3.3*	Section 3.3*	-
	LOND [17]	-	Section 3.3*	Section 3.3*	-
	SupLORD [32]	-	-	-	Xu and Ramdas [32]

Table 3: Procedures with proven SupFDR control under different dependence assumptions. The “implies” symbol indicates that the guarantee is implied by control under arbitrary dependence and the asterisk that the SupFDR guarantee is provided at a slightly inflated level.

Acknowledgments

LF acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project number 281474342/GRK2224/2. AR was funded by NSF grant DMS-2310718.

Supplementary Material

The R code to reproduce all simulations is available at https://github.com/fischer23/online_e-BH.

References

- [1] 1000 Genomes Project Consortium. A global reference for human genetic variation. *Nature*, 526(7571):68, 2015.
- [2] Ehud Aharoni and Saharon Rosset. Generalized α -investing: definitions, optimality results and application to public databases. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 76(4):771–794, 2014.
- [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 Series B: Statistical Methodology*, 57(1):289–300, 1995.
- [4] Yoav Benjamini and Daniel Yekutieli. The control of the false discovery rate in multiple testing under dependency. *The Annals of Statistics*, 29(4):1165–1188, 2001.
- [5] Gilles Blanchard and Etienne Roquain. Two simple sufficient conditions for FDR control. *Electronic Journal of Statistics*, 2:963–992, 2008.
- [6] Ziyu Chi, Aaditya Ramdas, and Ruodu Wang. Multiple testing under negative dependence. *Bernoulli*, 31(2):1230–1255, 2025.
- [7] Yo Joong Choe and Aaditya Ramdas. Combining Evidence Across Filtrations. arXiv:2402.09698, 2024.
- [8] A. Philip Dawid, Steven de Rooij, Peter Grunwald, Wouter M. Koolen, Glenn Shafer, Alexander Shen, Nikolai Vereshchagin, and Vladimir Vovk. Probability-free pricing of adjusted American lookbacks. arXiv:1108.4113, 2011.
- [9] A. Philip Dawid, Steven de Rooij, Glenn Shafer, Alexander Shen, Nikolai Vereshchagin, and Vladimir Vovk. Insuring against loss of evidence in game-theoretic probability. *Statistics & Probability Letters*, 81(1):157–162, 2011.
- [10] Lasse Fischer and Aaditya Ramdas. Online closed testing with e-values. arXiv:2407.15733, 2024.
- [11] Lasse Fischer, Marta Bofill Roig, and Werner Brannath. The online closure principle. *The Annals of Statistics*, 52(2):817–841, 2024.
- [12] Aaron Fisher. Online false discovery rate control for LORD++ and SAFFRON under positive, local dependence. *Biometrical Journal*, 66(1):2300177, 2024.
- [13] Aaron J Fisher. Online control of the false discovery rate under “Decision Deadlines”. In *International Conference on Artificial Intelligence and Statistics*, 2022.
- [14] Dean P Foster and Robert A Stine. α -investing: a procedure for sequential control of expected false discoveries. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 70(2):429–444, 2008.

- [15] Wenge Guo and M Bhaskara Rao. On control of the false discovery rate under no assumption of dependency. *Journal of Statistical Planning and Inference*, 138(10):3176–3188, 2008.
- [16] Adel Javanmard and Andrea Montanari. On online control of false discovery rate. arXiv:1502.06197, 2015.
- [17] Adel Javanmard and Andrea Montanari. Online rules for control of false discovery rate and false discovery exceedance. *The Annals of Statistics*, 46(2):526–554, 2018.
- [18] Ron Kohavi, Alex Deng, Brian Frasca, Toby Walker, Ya Xu, and Nils Pohlmann. Online controlled experiments at large scale. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2013.
- [19] Aaditya Ramdas, Fanny Yang, Martin J Wainwright, and Michael I Jordan. Online control of the false discovery rate with decaying memory. *Neural Information Processing Systems*, 30, 2017.
- [20] Aaditya Ramdas, Tijana Zrnic, Martin Wainwright, and Michael Jordan. SAFFRON: an adaptive algorithm for online control of the false discovery rate. In *International Conference on Machine Learning*, 2018.
- [21] 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. *The Annals of Statistics*, 47 (5):2790–2821, 2019.
- [22] David S Robertson, Jan Wildenhain, Adel Javanmard, and Natasha A Karp. onlineFDR: an R package to control the false discovery rate for growing data repositories. *Bioinformatics*, 35(20): 4196–4199, 2019.
- [23] David S Robertson, James MS Wason, Franz König, Martin Posch, and Thomas Jaki. Online error rate control for platform trials. *Statistics in Medicine*, 42(14):2475–2495, 2023.
- [24] Glenn Shafer, Alexander Shen, Nikolai Vereshchagin, and Vladimir Vovk. Test Martingales, Bayes Factors and p -Values. *Statistical Science*, 26(1), 2011.
- [25] John D Storey. A direct approach to false discovery rates. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 64(3):479–498, 2002.
- [26] John D Storey, Jonathan E Taylor, and David Siegmund. Strong control, conservative point estimation and simultaneous conservative consistency of false discovery rates: a unified approach. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 66(1):187–205, 2004.
- [27] Weijie J. Su. The FDR-Linking Theorem. arXiv:1812.08965, 2018.
- [28] V. G. Vovk. A Logic of Probability, with Application to the Foundations of Statistics. *Journal of the Royal Statistical Society. Series B (Methodological)*, 55(2):317–351, 1993.
- [29] Vladimir Vovk and Ruodu Wang. E-values: Calibration, combination and applications. *The Annals of Statistics*, 49(3):1736–1754, 2021.
- [30] Vladimir Vovk and Ruodu Wang. Confidence and discoveries with e-values. *Statistical Science*, 38 (2):329–354, 2023.
- [31] Ruodu Wang and Aaditya Ramdas. False discovery rate control with e-values. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(3):822–852, 2022.
- [32] Ziyu Xu and Aaditya Ramdas. Dynamic algorithms for online multiple testing. In *Mathematical and Scientific Machine Learning*, 2021.

- [33] Ziyu Xu and Aaditya Ramdas. More powerful multiple testing under dependence via randomization. arXiv:2305.11126, 2023.
- [34] Ziyu Xu and Aaditya Ramdas. Online multiple testing with e-values. In *International Conference on Artificial Intelligence and Statistics*, 2024.
- [35] Sonja Zehetmayer, Martin Posch, and Franz Koenig. Online control of the false discovery rate in group-sequential platform trials. *Statistical Methods in Medical Research*, 31(12):2470–2485, 2022.
- [36] Tijana Zrnic, Daniel Jiang, Aaditya Ramdas, and Michael Jordan. The power of batching in multiple hypothesis testing. In *International Conference on Artificial Intelligence and Statistics*, 2020.
- [37] Tijana Zrnic, Aaditya Ramdas, and Michael I Jordan. Asynchronous online testing of multiple hypotheses. *Journal of Machine Learning Research*, 22(33):1–39, 2021.

Here we give a brief overview of the appendix. We begin with a detailed explanation of boosting under local dependence and PRDS e-values in Section A. In Section B, we introduce e-TOAD, an e-value version of the TOAD procedure [13], which can be applied in online settings with decision deadlines. We also show that e-TOAD and TOAD (for arbitrarily dependent p-values) provide SupFDR control under arbitrary dependence. In Section C, we compare the LORD algorithm [17, 19] with the online BH procedure. Afterwards, we introduce an online version of the Storey-BH procedure [25, 26] and compare it to the SAFFRON algorithm [20] (Section D). In Section E, we discuss the considered notions of positive and negative dependence. In Section F, we describe the simulation setup used in the paper. In Section G, we provide omitted proofs and derivations for results stated in the main paper.

A Details about boosting with online e-BH

In this section, we provide several details regarding the boosting techniques briefly described in Section 2.3. In Section A.1, we demonstrate how the infinite sum in (2) can be handled. In Section A.2, we show how the boosting approach can be further improved when information about a local dependence structure is available while maintaining SupFDR control. We show in Section A.3 how OnlineFDR control can be guaranteed while employing an improved boosting technique for PRDS e-values.

A.1 Handling the infinite sum in the truncation function used for boosting

In this section we propose two approaches to handle the infinite sum in (2) and demonstrate their application in a Gaussian testing problem. The first approach is a conservative one. Here, for each $t \in \mathbb{N}$ we define some natural number $s \in \mathbb{N}$ and set

$${}^+T_t^s(x) := x \mathbb{1} \left\{ x < \frac{1}{s\alpha\gamma_t} \right\} + \sum_{k=1}^s \mathbb{1} \left\{ \frac{1}{k\alpha\gamma_t} \leq x < \frac{1}{(k-1)\alpha\gamma_t} \right\} \frac{1}{k\alpha\gamma_t}. \quad (3)$$

Obviously, we have $T_t(x) \leq {}^+T_t^s(x)$ for all $s \in \mathbb{N}$ and therefore $\mathbb{E}_{H_t}[{}^+T_t^s(X_t)] \leq 1$ implies that X_t is a boosted e-value. We still have that ${}^+T_t^s(x) \leq x$ such that boosting with ${}^+T_t^s$ can only increase power compared to no boosting.

Another approach is to set

$${}^-T_t^s(x) := \sum_{k=1}^s \mathbb{1} \left\{ \frac{1}{k\alpha\gamma_t} \leq x < \frac{1}{(k-1)\alpha\gamma_t} \right\} \frac{1}{k\alpha\gamma_t} = {}^+T_t^s(x) - x \mathbb{1} \left\{ x < \frac{1}{s\alpha\gamma_t} \right\} \quad (4)$$

for some natural number $s \in \mathbb{N}$. Note that we have ${}^-T_t^s(x) = T_t(x)$ for all $x \geq 1/(s\alpha\gamma_t)$ and ${}^-T_t^s(x) < T_t(x)$ for all $x < 1/(s\alpha\gamma_t)$. Therefore, applying online e-BH to X_t with $\mathbb{E}_{H_t}[{}^-T_t^s(X_t)] \leq 1$ does not necessarily control SupFDR. However, we can apply online e-BH to ${}^-T_t^s(X_t)$ instead, which provides SupFDR guarantee since ${}^-T_t^s(X_t)$ is a valid e-value. Furthermore, note that $T_t[{}^-T_t^s(x)] = {}^-T_t^s(x)$ for all x , implying that the e-value ${}^-T_t^s(X_t)$ cannot be further improved by boosting if $\mathbb{E}_{H_t}[{}^-T_t^s(X_t)] = 1$. With this approach we do not necessarily improve non-boosted e-values, since $E_t \leq X_t$ does not imply that $E_t \leq {}^-T_t^s(X_t)$, where \leq is to be read in an almost sure sense. However, in some situations this is even more powerful than boosting with (2). For example, suppose we know in advance that the total number of rejections won't be larger than r . Then boosting with ${}^-T_t^r$ is uniformly more powerful than boosting with T_t , since e-values smaller than $1/(r\alpha\gamma_t)$ cannot lead to a rejection anyway. If such prior information is not available one could also adapt s to t , for example, by setting $s = \lceil 1/\gamma_t \rceil$. In this case we could not reject hypotheses with $X_t < \alpha^{-1}$, however, all the mass of ${}^-T_t^s(X_t)$ is put to values greater than α^{-1} , increasing the probability for a rejection at larger levels.

Wang and Ramdas [31] have introduced several concrete examples how the boosted e-values X_1, X_2, \dots could be chosen in specific testing problems. In the following we extend one of them to the online case.

Example 1. Consider an e-value E_t that is obtained from a likelihood ratio between two normal distributions with variance 1 but a difference $\delta > 0$ in means

$$E_t = \exp(\delta Z_t - \delta^2/2),$$

where Z_t follows a standard normal distribution under the null hypothesis H_t . Hence, E_t is a log-normally distributed random variable with parameters $(-\delta^2/2, \delta)$ if H_t is true. Suppose we are looking for a boosted e-value of the type $X_t = b_t E_t$ for some boosting factor $b_t \geq 1$. Using (3), we can calculate

$$\begin{aligned} \mathbb{E}_{H_t} [{}^+T_t^s(b_t E_t)] &= b_t \left[1 - \Phi \left(\frac{\delta}{2} + \log(s\alpha\gamma_t b_t)/\delta \right) \right] \\ &+ \sum_{k=1}^s [\Phi(\delta/2 - \log([k-1]\alpha\gamma_t b_t)/\delta) - \Phi(\delta/2 - \log(k\alpha\gamma_t b_t)/\delta)] / (k\alpha\gamma_t), \end{aligned}$$

where Φ is the CDF of a standard normal distribution. Thus, we can set this equation to 1 and solve for b_t numerically. For $\alpha = 0.05$, $\gamma_t = 0.01$, $\delta = 3$ and $s = 10$, we obtain a boosting factor of $b_t = 1.165$. For increasing s , the method becomes less conservative and the boosting factor increases. For $s = 100$ we obtain $b_t = 1.174$, but increasing s further only leads to marginal improvements. Using (4), we obtain

$$\mathbb{E}_{H_t} [{}^-T_t^s(b_t E_t)] = \sum_{k=1}^s [\Phi(\delta/2 - \log([k-1]\alpha\gamma_t b_t)/\delta) - \Phi(\delta/2 - \log(k\alpha\gamma_t b_t)/\delta)] / (k\alpha\gamma_t).$$

In this case, the same parameters as before yield a boosting factor of $b_t = 3.071$. This boosting factor decreases if s increases. For $s = 100$, we obtain $b_t = 1.73$. The boosting factors obtained by these two approaches converge to the same limit for $s \rightarrow \infty$. Note that with the latter approach we need to plugin the e-value ${}^-T_t^s(b_t E_t)$ instead of $b_t E_t$ into the online e-BH algorithm. However, if we expect the number of rejections to be small, we can increase power substantially using ${}^-T_t^s$.

In the following section, we show how the boosting approach can be further improved if additional information about the joint distribution of the e-values is available.

A.2 Boosting under local dependence

In many online applications not all e-values are arbitrarily dependent. Indeed, e-values that lie far in the past often have no influence on the current testing process. For this reason, Zrnic et al. [37] introduced a *local dependence* structure that allows arbitrary dependence for e-values close together in time while e-values in the distant past should be independent. Precisely, $(E_t)_{t \in \mathbb{N}}$ are called locally dependent with lags $(L_t)_{t \in \mathbb{N}}$, if for all $t \in I_0$ it holds that

$$E_t \perp\!\!\!\perp E_{t-L_t-1}, \dots, E_1.$$

If $L_t = t - 1$, the e-values are arbitrarily dependent. If $L_t < t - 1$, the e-value E_t is independent of some of the past e-values.

In the following, we will show that such local dependence information can be exploited to further boost the e-values and increase power. For this, we define a lag-dependent truncation function.

$$\begin{aligned} T_t^{L_t}(x) &= \frac{1}{(k_{t-L_t-1}^* + 1)\alpha\gamma_t} \mathbb{1} \left\{ x \geq \frac{1}{(k_{t-L_t-1}^* + 1)\alpha\gamma_t} \right\} \\ &+ \sum_{k=k_{t-L_t-1}^* + 2}^{\infty} \mathbb{1} \left\{ \frac{1}{k\alpha\gamma_t} \leq x < \frac{1}{(k-1)\alpha\gamma_t} \right\} \frac{1}{k\alpha\gamma_t} \\ &= \min \left(T_t(x), \frac{1}{(k_{t-L_t-1}^* + 1)\alpha\gamma_t} \right), \end{aligned}$$

where T_t is the truncation function in (2). Note that $T_t^{L_t}(x) \leq T_t(x)$ for all x and L_t . Hence, the boosted e-values using $T_t^{L_t}$ can potentially be larger than the ones obtained by T_t . The idea of $T_t^{L_t}$ is that if we know that $k_{t-L_t-1}^*$ of the independent past hypotheses were rejected, H_t can be rejected if $E_t \geq [(k_{t-L_t-1}^* + 1)\alpha\gamma_t]^{-1}$. Hence, there is no additional gain if E_t is strictly larger than $[(k_{t-L_t-1}^* + 1)\alpha\gamma_t]^{-1}$. We can then use this information for improved boosting.

Proposition A.1. *Let X_1, X_2, \dots be a sequence of nonnegative random variables such that $\mathbb{E}_{H_t}[T_t^{L_t}(X_t)] \leq 1$ for all $t \in \mathbb{N}$. Then the online e-BH procedure applied with X_1, X_2, \dots controls the SupFDR at level α .*

Finding such X_1, X_2, \dots is particularly simple if a sequence of locally dependent e-values $(E_t)_{t \in \mathbb{N}}$ with lags $(L_t)_{t \in \mathbb{N}}$ and known marginal distributions is available. In this case one can simply calculate $\mathbb{E}_{H_t}[T_t^{L_t}(b_t E_t) | k_{t-L_t-1}^*]$, since E_t is independent of $k_{t-L_t-1}^*$.

Analogously to (3) and (4), we can also define ${}^+T_t^{L_t,s}$ and ${}^-T_t^{L_t,s}$, respectively. In the following we continue Example 1 to quantify the improvement gain due to local dependence.

Example 2. Consider Example 1 but suppose the e-value E_t is independent of the e-values E_1, \dots, E_{t-L_t-1} for some $L_t \in \{1, \dots, t-1\}$. We obtain

$$\begin{aligned} & \mathbb{E}_{H_t}[{}^+T_t^{L_t,s}(b_t E_t) | k_{t-L_t-1}^*] \\ &= b_t [1 - \Phi(\delta/2 + \log(s\alpha\gamma_t b_t)/\delta)] \\ &+ [1 - \Phi(\delta/2 - \log((k_{t-L_t-1}^* + 1)\alpha\gamma_t b_t)/\delta)] / [(k_{t-L_t-1}^* + 1)\alpha\gamma_t] \\ &+ \sum_{k=k_{t-L_t-1}^*+2}^s [\Phi(\delta/2 - \log([k-1]\alpha\gamma_t b_t)/\delta) - \Phi(\delta/2 - \log(k\alpha\gamma_t b_t)/\delta)] / (k\alpha\gamma_t). \end{aligned}$$

For the parameters $\alpha = 0.05$, $\gamma_t = 0.01$, $\delta = 3$, $s = 100$ and $k_{t-L_t-1}^* = 2$, this leads to the boosting factor $b_t = 1.265$. For a larger number of previous (independent) rejections $k_{t-L_t-1}^* = 10$, the boosting factor also increases to $b_t = 1.541$. Furthermore, we have

$$\mathbb{E}_{H_t}[{}^-T_t^{L_t,s}(b_t E_t) | k_{t-L_t-1}^*] = \mathbb{E}_{H_t}[{}^+T_t^{L_t,s}(b_t E_t) | k_{t-L_t-1}^*] - b_t [1 - \Phi(\delta/2 + \log(s\alpha\gamma_t b_t)/\delta)].$$

In this case we obtain a boosting factor of $b_t = 1.940$ for $k_{t-L_t-1}^* = 2$ and $b_t = 2.639$ for $k_{t-L_t-1}^* = 10$.

The gain in power obtained by exploiting information about the local dependence structure to improve boosting is illustrated in Figure 2. It is seen that using information about locally dependent e-values improves power significantly.

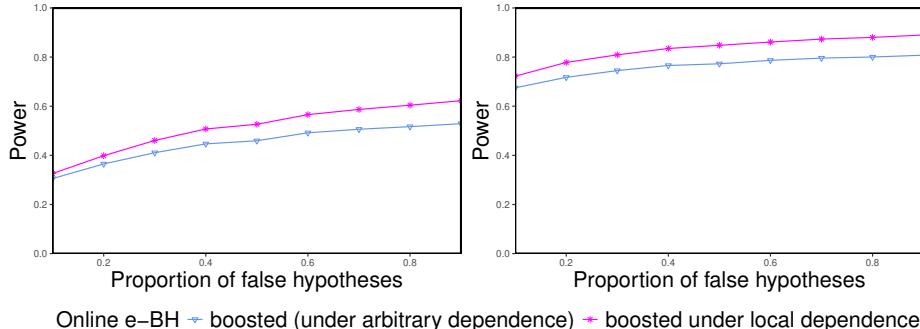


Figure 3: Power comparison of online e-BH with boosted and boosted e-values under local dependence for different proportions of false hypotheses. In the left (right) plot the signal of the alternative is weak (strong). The simulation setup is described in Section F.

In the p-value literature, it is usually started with a procedure under independence which is made conservative when local dependence is present. Interestingly, we take the other direction and start with a procedure that works under arbitrary dependence and improve it when information about a local dependence structure is available. A similar behavior for the BH and e-BH procedure was noted by Wang and Ramdas [31].

Remark 6. Another advantage of local dependence is that γ_t can be measurable with respect to $\mathcal{F}_{t-L_t-1} = \sigma(E_1, \dots, E_{t-L_t-1})$ and thus depend on the previous data. In fact, to choose γ_t predictable it is even sufficient that E_t is valid conditional on \mathcal{F}_{t-L_t-1} , which is slightly weaker than independence. To see this, just note that for all $t \in I_0$ it holds that

$$\mathbb{E}[\gamma_t E_t] = \mathbb{E}[\mathbb{E}[\gamma_t E_t | \mathcal{F}_{t-L_t-1}]] = \mathbb{E}[\gamma_t \mathbb{E}[E_t | \mathcal{F}_{t-L_t-1}]] \leq \mathbb{E}[\gamma_t].$$

A.3 Boosting with PRDS e-values

Wang and Ramdas [31] also introduced a more powerful boosting technique than the one we described in Section A.1 which works under positive regression dependence on a subset (PRDS) [4] (see Section E for a formal definition of PRDS for e-values). In the following we extend this method to the online setting. However, the proof technique of FDR control is different in this case such that SupFDR control can no longer be guaranteed and we only show OnlineFDR control.

Let the truncation function T_t be as in (2). Instead of bounding the expected value of the truncated e-value $T_t(X_t)$, Wang and Ramdas [31] proposed to bound the probability

$$\sup_{k \in \mathbb{N}} \frac{1}{k\alpha\gamma_t} \mathbb{P}_{H_t}(X_t \geq 1/(k\alpha\gamma_t)) = \sup_{x \geq 0} x \mathbb{P}_{H_t}(T_t(X_t) \geq x) \leq 1, \quad (5)$$

if the X_t , $t \in \mathbb{N}$, are PRDS. In the following, we show that this condition can also be used for boosting in the online case.

Proposition A.2. *Let X_1, X_2, \dots be a sequence of nonnegative random variables such that (5) holds for all $t \in \mathbb{N}$. Then the online e-BH procedure applied with X_1, X_2, \dots controls the OnlineFDR at level α , if the random variables X_1, X_2, \dots are PRDS.*

Proof. Let $t \in \mathbb{N}$ be fixed and $\mathbf{X} = (X_1\alpha\gamma_1, \dots, X_t\alpha\gamma_t)$. Then $f(\mathbf{X}) = 1/(k_t^* \vee 1)$, where

$$k_t^* = \max \left\{ k \in \{1, \dots, t\} : \sum_{j=1}^t \mathbb{1}\{X_j \geq 1/(k\alpha\gamma_j)\} \geq k \right\},$$

is a nonincreasing function. Furthermore, the range of f is given by $I_f = \{1, 1/2, 1/3, \dots\}$ and

$$\sup\{x \in I_f : x \leq X_i\alpha\gamma_i\} = T_i(X_i)\alpha\gamma_i.$$

With this, Lemma 1 (ii) in [31] implies that

$$\mathbb{E}[f(\mathbf{X}) \mathbb{1}\{X_i\alpha\gamma_i \geq f(\mathbf{X})\}] \leq \sup_{x \geq 0} x \mathbb{P}(T_i(X_i)\alpha\gamma_i \geq x)$$

for all $i \in I_0$, $i \leq t$. Since

$$\sup_{x \geq 0} x \mathbb{P}(T_i(X_i)\alpha\gamma_i \geq x) \leq \alpha\gamma_i \Leftrightarrow \sup_{x \geq 0} x \mathbb{P}(T_i(X_i) \geq x) \leq 1,$$

we obtain

$$\begin{aligned} \text{FDR}_t &= \sum_{i \in I_0, i \leq t} \mathbb{E} \left[\frac{\mathbb{1}\{X_i \geq 1/(\alpha\gamma_i(k_t^* \vee 1))\}}{(k_t^* \vee 1)} \right] \\ &\leq \sum_{i \in I_0, i \leq t} \sup_{x \geq 0} x \mathbb{P}(T_i(X_i)\alpha\gamma_i \geq x) \leq \alpha \sum_{i \in I_0, i \leq t} \gamma_i \leq \alpha. \end{aligned}$$

□

By Markov's inequality, $\mathbb{E}_{H_t}[T_t(X_t)] \leq 1$ implies that (5) is satisfied. Hence, boosting with (5) is more powerful than boosting with $\mathbb{E}_{H_t}[T_t(X_t)] \leq 1$.

B The e-TOAD procedure

First, we review the TOAD procedure by Fisher [13] under arbitrary dependence. Afterwards, we introduce e-TOAD, an e-value analog of the TOAD procedure and show that e-TOAD provides SupFDR control. By showing that TOAD is a special case of e-TOAD we also conclude SupFDR control for the TOAD procedure.

In the online setup with “decision deadlines” by Fisher [13] each H_t comes with a decision deadline $d_t \geq t$, which means that the decision for hypothesis H_t can be changed until time d_t . For example, if $d_t = t$ for all $t \in \mathbb{N}$, the classical online setup is obtained. Furthermore, let $C_t = \{i \leq t : d_i \geq t\}$ be the set of currently “active” hypotheses. We will see that TOAD and e-TOAD never turn a rejected hypothesis into an accepted hypothesis and therefore are online ARC procedures.

Initialize $R_0 = \emptyset$. The TOAD algorithm under arbitrary dependence [13] is recursively defined as follows.

$$R_t^{\text{TOAD}} := \left\{ i \in \{1, \dots, t\} : P_i \leq \alpha \gamma_i \beta_i(k_{\min(d_i, t)}^*) \right\}, t \in \mathbb{N}, \text{ where} \quad (6)$$

$$k_t^* = |R_{t-1}^{\text{TOAD}} \setminus C_t| + \max \left\{ k \leq |C_t| : \sum_{j \leq t, j \in C_t} \mathbb{1}\{P_j \leq \alpha \gamma_j \beta_j(|R_{t-1}^{\text{TOAD}} \setminus C_t| + k)\} \geq k \right\}.$$

The shape functions β_i are defined the same as in Section 2.4. The TOAD procedure under arbitrary dependence is an online generalization of the BR procedure [5]. Note that Fisher [13] only allowed to use the same shape function $\beta = \beta_i$ for all $i \in \mathbb{N}$. This can cost a lot of power, particularly if $d_t < \infty$. For example, suppose $d_1 = 1$, which means that the decision for the first hypothesis must be made immediately. Then we could choose $\beta_1(1) = 1$, meaning to put all mass on $k_1^* = 1$ since this is the only relevant value. This would not be possible if the same shape function must be used for all hypotheses, which usually implies $\beta_1(1) < 1$. We will show that TOAD even controls the SupFDR for individual shape functions by proving that it is a special case of e-TOAD.

The e-TOAD algorithm is defined by

$$R_t^{\text{e-TOAD}} := \left\{ i \in \{1, \dots, t\} : E_i \geq (\alpha \gamma_i k_{\min(d_i, t)}^*)^{-1} \right\}, t \in \mathbb{N}, \text{ where}$$

$$k_t^* = |R_{t-1}^{\text{e-TOAD}} \setminus C_t| + \max \left\{ k \leq |C_t| : \sum_{j \leq t, j \in C_t} \mathbb{1}\{E_j \geq [\alpha \gamma_j (|R_{t-1}^{\text{e-TOAD}} \setminus C_t| + k)]^{-1}\} \geq k \right\}.$$

Note that e-TOAD becomes online e-BH in case of $d_t = \infty$ for all t and e-LOND if $d_t = t$ for all t .

Analogously to Section 2, the SupFDR control of the e-TOAD algorithm follows because the e-value E_i , $i \in \{1, \dots, t\}$, of every rejected hypothesis at time t satisfies

$$E_i \geq 1/(\alpha \gamma_i |R_t^{\text{e-TOAD}}|).$$

Proposition B.1. *The e-TOAD procedure provides SupFDR control at level α under arbitrary dependence between the e-values.*

Define the truncation function as

$$T_t^{\text{TOAD}}(x) = \sum_{k=1}^{d_t} \mathbb{1} \left\{ \frac{1}{k \alpha \gamma_t} \leq x < \frac{1}{(k-1) \alpha \gamma_t} \right\} \frac{1}{k \alpha \gamma_t} \quad \text{with } T_t(\infty) = \frac{1}{\alpha \gamma_t}.$$

Then e-TOAD applied to X_1, X_2, \dots rejects the same hypotheses as if applied to $T_1^{\text{TOAD}}(X_1), T_2^{\text{TOAD}}(X_2), \dots$. Therefore, applying e-TOAD to X_1, X_2, \dots ensures SupFDR control, if $\mathbb{E}[T_t^{\text{TOAD}}(X_t)] \leq 1$ for all $t \in I_0$. With this, we can prove the following proposition in the exact same manner as we proved Proposition 2.5.

Proposition B.2. *Let ψ_1, ψ_2, \dots be nonincreasing transforms such that for every t it holds that*

$$\sum_{k=1}^{d_t} \frac{1}{k\alpha\gamma_t} \left(\psi_t^{-1} \left(\frac{1}{k\alpha\gamma_t} \right) - \psi_t^{-1} \left(\frac{1}{(k-1)\alpha\gamma_t} \right) \right) \leq 1.$$

Then the e-TOAD procedure applied to the random variables $\psi_1(P_1), \psi_2(P_2), \dots$ controls the SupFDR at level α .

Define ψ_t such that $\psi_t^{-1}(1/(\alpha\gamma_t k)) = \alpha\beta_t(k)\gamma_t$. Then $\psi_t(P_t) \geq 1/(\alpha\gamma_t k)$ iff $P_t \leq \psi_t^{-1}(1/(\alpha\gamma_t k)) = \alpha\beta_t(k)\gamma_t$. Hence, e-TOAD applied to $\psi_1(P_1), \psi_2(P_2), \dots$ rejects the same hypotheses as TOAD applied to P_1, P_2, \dots . In the exact same manner as in Proposition G.1, we can show that ψ_1, ψ_2, \dots satisfy Proposition B.2. Hence, TOAD controls the SupFDR.

Proposition B.3. *The TOAD procedure (6) provides SupFDR control at level α under arbitrary dependence between the p-values.*

C Online BH versus LORD

In this section, we compare the LORD procedure [17, 19] with the online BH procedure. Javanmard and Montanari [17] introduced the LORD algorithm with OnlineFDR control which was later generalized and improved to the LORD++ algorithm by Ramdas et al. [19]. Here we follow the latter presentation, but call it LORD instead of LORD++ for brevity.

Let $\alpha_1^{\text{LORD}}, \alpha_2^{\text{LORD}}, \dots$ be individual levels such that

$$\frac{\sum_{i=1}^t \alpha_i^{\text{LORD}}}{|R_t^{\text{LORD}}| \vee 1} \leq \alpha \text{ for all } t \in \mathbb{N}, \text{ where } R_t^{\text{LORD}} := \{i \leq t : P_i \leq \alpha_i^{\text{LORD}}\}. \quad (7)$$

Ramdas et al. [19] showed that any online procedure defined by the rejection sets R_t^{LORD} , $t \in \mathbb{N}$, controls the OnlineFDR if the null p-values are independent from each other and the non-nulls; and the thresholds α_i^{LORD} , $i \in \mathbb{N}$, are nonrandom and nondecreasing functions of the past rejection indicators $\mathbb{1}\{P_1 \leq \alpha_1^{\text{LORD}}\}, \dots, \mathbb{1}\{P_{i-1} \leq \alpha_{i-1}^{\text{LORD}}\}$.

Let $\alpha_{i,t}^{\text{BH}} := k_t^* \alpha \gamma_i$, $t \in \mathbb{N}$ and $i \leq t$, be the individual significance level for hypothesis H_i at time t obtained by the online BH procedure (Section 3). Since $k_t^* = |R_t^{\text{BH}}|$, it is easy to see that the individual levels $\alpha_{i,t}^{\text{BH}}$, $i \leq t$, also satisfy the LORD condition (7). However, $\alpha_{i,t}^{\text{BH}}$ is not a function of only the first $i-1$ rejections, but depends on all rejections up to step t . Therefore, online BH is not a special instance of the LORD algorithm, but online BH and LORD are both allowed to spend the same amount of significance level up to each time t . For this reason, online BH and LORD should have similar power, and the difference between these procedures lies in a different weighting of the hypotheses (with online BH all individual levels gain from rejections, while with LORD only future levels gain from rejections).

This is verified via simulations in Figure 4, where we compare online BH with a special instance of the LORD algorithm implemented in the R package `onlineFDR` [22]. The simulation setup is described in Section F. Recall that $(\gamma_t)_{t \in \mathbb{N}}$ is a sequence that sums to one and which defines the weighting of the different hypotheses. In Figure 4, we set $\gamma_t = q^{t-1}(1-q)$ with $q = 0.99$ in the left plot and with $q = 0.999$ in the right plot for both procedures. That means, in the right plot the γ sequence is decreasing slower. It is easy to see that LORD performs better than online BH in the left plot and worse in the right plot. In addition, online BH with $q = 0.999$ performs nearly the exact same as

LORD with $q = 0.99$. This supports our theoretical reasoning that the difference between these two procedures lies in a different weighting of the hypotheses, but both procedures spend the same amount of significance level.

We think one advantage of online BH compared to LORD is that its weighting is easier to interpret. If $\gamma_i = 2\gamma_j$ for $i \neq j$, then H_i will be tested at twice the level compared to H_j with online BH, regardless of the number of rejections. In contrast, the weighting of LORD strongly depends on the number of rejections in a complex nonlinear manner. Another important advantage of online BH compared to LORD is that online BH controls the OnlineFDR under PRDS, while LORD is only known to control it under independence. In addition, we proved that online BH even controls the SupFDR at slightly inflated levels under PRDN and WNDN. Such results are not known for the LORD procedure. However, of course LORD is an online procedure in the strict sense, which can be desirable in certain applications, while online BH is an online ARC procedure. We conclude that there is much to be gained by relaxing the online requirement to an online ARC requirement: we get a procedure with seemingly comparable power, but with a much stronger suite of type-I error control guarantees.

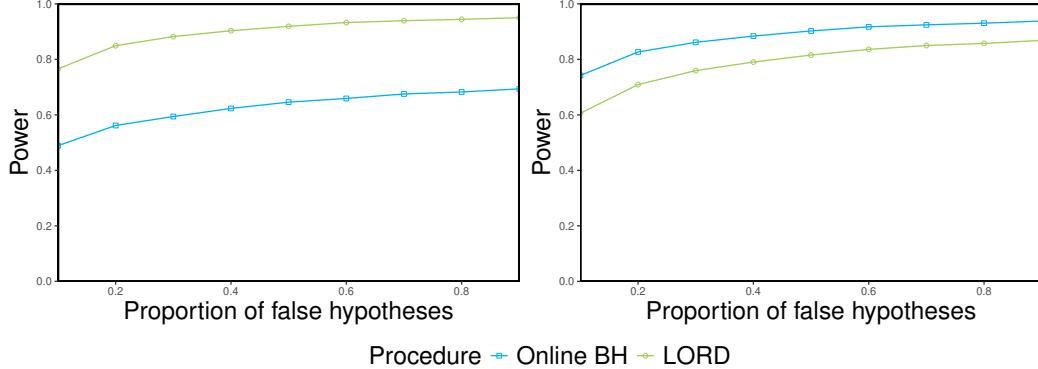


Figure 4: Power comparison of online BH and LORD for different proportions of false hypotheses. In the left plot the sequence $(\gamma_t)_{t \in \mathbb{N}}$ decreases fast ($q = 0.99$) and in the right plot it decreases slow ($q = 0.999$). The simulation setup is described in Section C.

D The online Storey-BH procedure

Similarly as for e-BH and BH, one can define an online version of the Storey-BH procedure [25, 26], which we call just SBH in the following. In this section, we introduce the online SBH procedure, prove its OnlineFDR guarantee and compare it with the existing SAFFRON [20] procedure.

Recall that $(\gamma_t)_{t \in \mathbb{N}}$ is a sequence that sums to one. In addition, let $\lambda \in [\alpha, 1)$ be a user-defined constant and

$$\hat{\pi}_0^t = \frac{\max_{i \in \mathbb{N}} \gamma_i + \sum_{i=1}^{\infty} \gamma_i \mathbb{1}\{P_i > \lambda \vee i > t\}}{(1 - \lambda)}. \quad (8)$$

Then the online SBH method is defined by the rejection sets

$$R_t^{\text{oSBH}} := \{i \leq t : P_i \leq \min(k_t^* \alpha \gamma_i / \hat{\pi}_0^t, \lambda)\}, t \in \mathbb{N},$$

where $k_t^* = \max \left\{ k \in \{1, \dots, t\} : \sum_{j=1}^t \mathbb{1}\{P_j \leq \min(k \alpha \gamma_j / \hat{\pi}_0^t, \lambda)\} \geq k \right\},$

with the convention $\max(\emptyset) = 0$. Note that for $\gamma_1 = \dots = \gamma_K = 1/K$, we have $\hat{\pi}_0^K = \frac{1 + \sum_{i=1}^K \mathbb{1}\{P_i > \lambda\}}{(1-\lambda)K}$. In this case the online SBH procedure recovers the SBH procedure [25, 26] and therefore is a true generalization of it. In addition, since $\hat{\pi}_0^t$ is nonincreasing in t , the online SBH procedure is indeed an online ARC procedure.

Proposition D.1. *If the null p -values are independent from each other and the non-nulls, the online SBH procedure ensures that $\text{OnlineFDR} \leq \alpha$.*

Proof. Let $\hat{\pi}_0^{t,-j} = \frac{\max_{i \in \mathbb{N}} \gamma_i + \sum_{i=1, i \neq j}^{\infty} \gamma_i \mathbb{1}\{P_i > \lambda \vee i > t\}}{(1-\lambda)}$. For any $t \in \mathbb{N}$, we have

$$\begin{aligned} \text{FDR}_t &= \sum_{i \leq t, i \in I_0} \mathbb{E} \left[\frac{\mathbb{1}\{P_i \leq \min(k_t^* \alpha \gamma_i / \hat{\pi}_0^t, \lambda)\}}{k_t^*} \right] \\ &= \sum_{i \leq t, i \in I_0} \mathbb{E} \left[\frac{\mathbb{1}\{P_i \leq \min(k_t^* \alpha \gamma_i / \hat{\pi}_0^{t,-i}, \lambda)\}}{k_t^*} \right] \\ &\leq \sum_{i \leq t, i \in I_0} \mathbb{E} \left[\frac{\mathbb{1}\{P_i \leq k_t^* \alpha \gamma_i / \hat{\pi}_0^{t,-i}\}}{k_t^*} \right] \\ &= \sum_{i \leq t, i \in I_0} \mathbb{E} \left[\mathbb{E} \left(\frac{\mathbb{1}\{P_i \leq k_t^* \alpha \gamma_i / \hat{\pi}_0^{t,-i}\}}{k_t^*} \middle| (P_j)_{j \in \{1, \dots, t\} \setminus i} \right) \right] \\ &\stackrel{(i)}{\leq} \sum_{i \leq t, i \in I_0} \alpha \gamma_i \mathbb{E} \left[\frac{1}{\hat{\pi}_0^{t,-i}} \right] \\ &\stackrel{(ii)}{\leq} \sum_{i \leq t, i \in I_0} \alpha \gamma_i \frac{1}{\max_{j \leq t} \gamma_j + \sum_{j \leq t, j \in I_0, j \neq i} \gamma_j} \\ &= \alpha \sum_{i \leq t, i \in I_0} \frac{\gamma_i}{\max_{j \leq t} \gamma_j + \sum_{j \leq t, j \in I_0, j \neq i} \gamma_j} \leq \alpha. \end{aligned}$$

Inequality (i) follows by Lemma 1 (a) and inequality (ii) by Lemma 3 of Ramdas et al. [21]. \square

An existing online procedure that adapts to the proportion of true hypotheses is SAFFRON [20], which is defined as LORD (see Section C) but with the condition

$$\frac{\sum_{i=1}^t \alpha_i^{\text{SAFFRON}} \frac{\mathbb{1}\{P_i > \lambda\}}{1-\lambda}}{|R_t^{\text{SAFFRON}}| \vee 1} \leq \alpha \text{ for all } t \in \mathbb{N}, \text{ where } R_t^{\text{SAFFRON}} := \{i \leq t : P_i \leq \alpha_i^{\text{SAFFRON}}\} \quad (9)$$

instead of (7). The relation between SAFFRON and online SBH is similar to the relation between LORD and online BH (see Section C). That means, the levels defined by online SBH satisfy (9). However, in contrast to the LORD vs. online BH case, online SBH satisfies the SAFFRON condition (9) only in a conservative manner due to the additional summand $\max_{i \in \mathbb{N}} \gamma_i$ in (8). We think this is required, since SAFFRON only allows to use the information $P_i > \lambda_i$ for all levels $\alpha_j^{\text{SAFFRON}}$, $j > i$, while online SBH uses it for all $j \in \mathbb{N}$. Consequently, online SBH can lead to a more balanced weighting. Overall, we think both procedures lead to a comparable power and the main difference lies in a different weighting of the hypotheses, just as for online BH and LORD.

This can also be seen in Figure 5, where we compare online SBH with SAFFRON as implemented in the R package `onlineFDR` [22], which should be compared to Figure 4. Both procedures were applied with $\lambda = 0.5$, the remaining simulation setup is the same as in Section C.

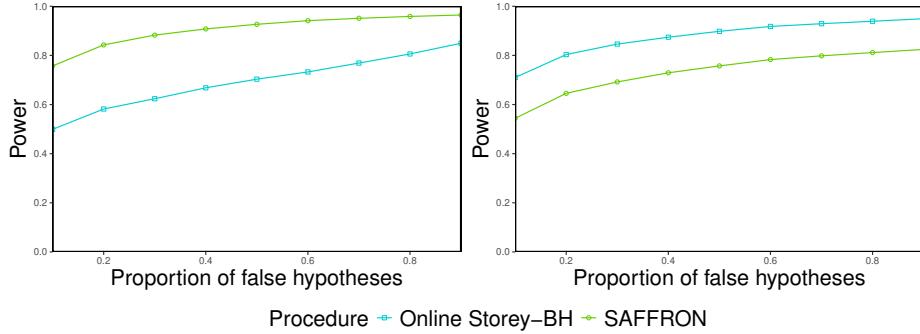


Figure 5: Power comparison of online SBH and SAFFRON for different proportions of false hypotheses. In the left plot the sequence $(\gamma_t)_{t \in \mathbb{N}}$ decreases fast ($q = 0.99$) and in the right plot it decreases slow ($q = 0.999$). The simulation setup is described in Section D.

E Positive and negative dependence

Here we give the formal definitions of the used positive and negative dependence assumptions made in the paper, which are online extensions of the classical definitions [4, 27, 6, 31].

Let us first define two different positive dependence assumptions that one can make on p-values. The first is the PRDS condition for FDR control of the BH procedure that was introduced by Benjamini and Yekutieli [4]. We generalize this typical PRDS definition to the online setting (i.e., applicable to countably infinite hypotheses) in the following way.

Definition 3 (PRDS). A set $D \subseteq \mathbb{R}^n$ is called *increasing* if $\mathbf{x} \in D$ implies that $\mathbf{y} \in D$ for all $\mathbf{y} \geq \mathbf{x}$. For a sequence of p-values $(P_t)_{t \in I}$ with a countable index set I , we say that they have *positive regression dependence on a subset (PRDS)* if, for any finite $I' \subseteq I$, increasing set $D \subseteq \mathbb{R}^{I'}$, and null index $i \in I_0$, we have that $\mathbb{P}((P_t)_{t \in I'} \in D \mid P_i \leq s)$ is increasing over $s \in (0, 1]$.

Note that this definition is equivalent to the standard PRDS definition when I is finite. Further, this is the PRDS definition implicitly used by Zrnic et al. [37]. For example, PRDS is satisfied if the p-values corresponding to true hypotheses are independent from all the other p-values. However, it also holds under some kind of positive dependence [4].

A slightly weaker notion of positive dependence is PRDN [27], which is restricted to the dependency between p-values that correspond to true hypotheses.

Definition 4 (PRDN). For a sequence of p-values $(P_t)_{t \in I}$ with a countable index set I , we say that they have *positive regression dependence within nulls (PRDN)* if, for any finite subset of the null hypotheses $I' \subseteq I_0$, increasing set $D \subseteq \mathbb{R}^{I'}$, and null index $i \in I_0$, we have that $\mathbb{P}((P_t)_{t \in I'} \in D \mid P_i \leq s)$ is a increasing over $s \in (0, 1]$.

Likewise, this is the generalization of PRDN from Su [27] to the online setting.

The only type of negative dependence we consider in this paper is WNDN.

Definition 5 (WNDN). We say that a sequence of p-values $(P_t)_{t \in I}$ are *weakly negatively dependent within nulls (WNDN)* if $\mathbb{P}(\bigcap_{i \in A} P_i \leq s) \leq \prod_{i \in A} \mathbb{P}(P_i \leq s)$ for any $A \subseteq I_0$ and $s \in [0, 1]$.

Similarly as for positive dependence, there are many types of negative dependence. A comparison of the different types is given by Chi et al. [6].

Since e-values and p-values are working on inverse scales, the PRDS condition is also flipped around when considering e-values [31].

Definition 6 (PRDS for e-values). A set $D \subseteq \mathbb{R}^n$ is called *decreasing* if $\mathbf{x} \in D$ implies that $\mathbf{y} \in D$ for all $\mathbf{y} \leq \mathbf{x}$. For a sequence of e-values $(E_t)_{t \in I}$ with a countable index set I , we say that they have *positive regression dependence on a subset (PRDS)* if, for any finite $I' \subseteq I$, decreasing set $D \subseteq \mathbb{R}^{I'}$, and null index $i \in I_0$, we have that $\mathbb{P}((E_t)_{t \in I'} \in D \mid E_i \geq s)$ is decreasing over $s \in (0, \infty)$.

This definition of PRDS for e-values ensures that if the p-values $(P_t)_{t \in \mathbb{N}}$ are PRDS and $(\psi_t)_{t \in \mathbb{N}}$ are nonincreasing transforms, then the random variables $(\psi_t(P_t))_{t \in \mathbb{N}}$ are PRDS in the e-value sense.

F Simulation setup

We simulated $m = 100$ independent trials, each consisting of $n = 1000$ hypothesis pairs of the form $H_t : X_t \sim N(0, 1)$ vs. $H_t^A : X_t \sim N(\mu_A, 1)$, $t \in \mathbb{N}$, where $\mu_A > 0$ is the signal strength. In case of $\mu_A = 3.5$ we speak of weak signal and in case of $\mu_A = 4.5$ of strong signal. To obtain locally dependent e-values, we first generated normally distributed random variables (Z_1, \dots, Z_n) with mean 0, variance 1 and a batch dependence structure. That means we have batches $B_i = \{Z_{(i-1)b+1}, \dots, Z_{ib}\}$, $i \in \{1, \dots, n/b\}$, such that random variables contained in the same batch have correlation 0.5 while random variables from different batches are independent. The final test statistics were obtained by $X_t = Z_t + \mu_A \Pi_A$, where $\Pi_A \sim \text{Ber}(\pi_A)$ and $\pi_A \in \{0.1, 0.2, \dots, 0.9\}$ gives the probability that H_t^A is true.

The e-values were calculated as likelihood ratio $\phi_{\mu_A}(X_t)/\phi_0(X_t)$, where ϕ_μ denotes the density of a normal distribution with variance 1 and mean μ . The boosted e-values were obtained as described in Examples 1 and 2 using the truncation functions \bar{T}_t^s and $\bar{T}_t^{L_t, s}$ with $s = n$, respectively. In all simulations concerning e-values we set the batch-size to $b = 20$ and applied all procedures with $\gamma_t = q^{t-1}(1 - q)$, $t \in \{1, \dots, n\}$, where $q = 0.99$.

The p-values were calculated by $P_t = \Phi(-Z_t)$. When considering p-values, we generated independent test statistics ($b = 1$) and only considered weak signals ($\mu_A = 3.5$).

The code for the simulations is available at https://github.com/fischer23/online_e-BH.

G Omitted proofs and derivations

In this section, we provide the omitted proofs or derivations for results that were mentioned in the main paper.

G.1 Online e-BH with p-values

In this subsection, we provide the formal proof of Proposition 2.5.

Proof. If we prove that $\mathbb{E}[T_t(\psi_t(U))] \leq 1$, where T_t is given by (2) and $U \sim U[0, 1]$, then the first claim follows by Proposition 2.4. For this, note that

$$\begin{aligned} \mathbb{E}[T_t(\psi_t(U))] &= \sum_{k=1}^{\infty} \mathbb{P} \left(\frac{1}{k\alpha\gamma_t} \leq \psi_t(U) < \frac{1}{(k-1)\alpha\gamma_t} \right) \frac{1}{k\alpha\gamma_t} \\ &= \sum_{k=1}^{\infty} \mathbb{P} \left(\psi^{-1} \left(\frac{1}{k\alpha\gamma_t} \right) \geq U > \psi^{-1} \left(\frac{1}{(k-1)\alpha\gamma_t} \right) \right) \frac{1}{k\alpha\gamma_t} \\ &= \sum_{k=1}^{\infty} \frac{1}{k\alpha\gamma_t} \left(\psi_t^{-1} \left(\frac{1}{k\alpha\gamma_t} \right) - \psi_t^{-1} \left(\frac{1}{(k-1)\alpha\gamma_t} \right) \right) \\ &\leq 1. \end{aligned}$$

The second claim follows immediately by Proposition A.2, since

$$\begin{aligned} \sup_{k \in \mathbb{N}} \frac{1}{k\alpha\gamma_t} \mathbb{P}(\psi_t(U) \geq 1/(k\alpha\gamma_t)) &= \sup_{k \in \mathbb{N}} \frac{1}{k\alpha\gamma_t} \mathbb{P}(U \leq \psi_t^{-1}(1/(k\alpha\gamma_t))) \\ &= \sup_{k \in \mathbb{N}} \frac{1}{k\alpha\gamma_t} \psi_t^{-1}(1/(k\alpha\gamma_t)). \end{aligned}$$

□

G.2 SupFDR control of online BR and r-LOND

In the following proposition, we show how the online BR procedure can be written as a special instance of the online e-BH method and therefore provides SupFDR control under arbitrary dependence between the e-values.

Proposition G.1. *The online BR procedure provides SupFDR control at level α under arbitrary dependence between the p-values.*

Proof. Let β_t be a shape function and ψ_t be chosen such that $\psi_t^{-1}(1/(\alpha\gamma_t k)) = \alpha\beta_t(k)\gamma_t$. Note that $\psi_t(P_t) \geq 1/(\alpha\gamma_t k)$ iff $P_t \leq \psi_t^{-1}(1/(\alpha\gamma_t k)) = \alpha\beta_t(k)\gamma_t$. Hence, online e-BH applied to $\psi_1(P_1), \psi_2(P_2), \dots$ rejects the same hypotheses as online BR applied to P_1, P_2, \dots . Furthermore, it holds that

$$\begin{aligned} \sum_{k=1}^{\infty} \frac{1}{k\alpha\gamma_t} \left(\psi_t^{-1}\left(\frac{1}{k\alpha\gamma_t}\right) - \psi_t^{-1}\left(\frac{1}{(k-1)\alpha\gamma_t}\right) \right) &= \sum_{k=1}^{\infty} \frac{1}{k} (\beta_t(k) - \beta_t(k-1)) \\ &= \sum_{k=1}^{\infty} \int_{k-1}^k x/k \, d\nu(x) \\ &\leq \int_0^{\infty} 1 \, d\nu(x) \\ &= 1. \end{aligned}$$

Consequently, the SupFDR control follows by Proposition 2.5. \square

In the same manner as we derived the online BR procedure as special case of the boosted online e-BH procedure, one could also derive r-LOND as a special case of a boosted e-LOND procedure. For this, one can just define the truncation function

$$T_t^{\text{LOND}} = \sum_{k=1}^t \mathbb{1} \left\{ \frac{1}{k\alpha\gamma_t} \leq x < \frac{1}{(k-1)\alpha\gamma_t} \right\} \frac{1}{k\alpha\gamma_t} \quad \text{with } T_t(\infty) = \frac{1}{\alpha\gamma_t}$$

and then do the same steps as above. Consequently, r-LOND also provides SupFDR control.

Proposition G.2. *The r-LOND procedure provides SupFDR control at level α under arbitrary dependence between the p-values.*

G.3 Weighted FDR-Linking theorem

Here we provide the proof of Theorem 3.2.

Proof. We follow a similar proof structure as the proof of Theorem 1 in Su [27]. By the definition of self-consistency, we note that any $R \in \mathcal{R}(\alpha)$ and $t \in R$, the following inequality holds:

$$|R| \geq \lceil (\alpha\gamma_t)^{-1} P_t \rceil. \quad (10)$$

Thus, we get an upper bound on the supremum over FDP

$$\sup_{R \in \mathcal{R}(\alpha)} \text{FDP}(R) = \max_{R \in \mathcal{R}(\alpha)} \frac{|I_0 \cap R|}{|R| \vee 1} \leq \max_{R \in \mathcal{R}(\alpha)} \frac{|I_0 \cap R|}{\max\{\lceil (\alpha\gamma_t)^{-1} P_t \rceil : t \in R\} \cup \{1\}}.$$

The inequality is a result of applying the bound in (10) to every $t \in R$ and taking the best possible bound.

Let $K_0 := |I_0|$ and $P_1^{\text{null}}, \dots, P_{K_0}^{\text{null}}$ and $\gamma_1^{\text{null}}, \dots, \gamma_{K_0}^{\text{null}}$ denote the corresponding weights. Now, define $P_{(1)}^{\text{null}}, \dots, P_{(K_0)}^{\text{null}}$ and $\gamma_{(1)}^{\text{null}}, \dots, \gamma_{(K_0)}^{\text{null}}$ be the null p-values and weights sorted in ascending order based on $(\gamma_i^{\text{null}})^{-1} P_i^{\text{null}}$. We continue the derivation of the upper bound as follows.

$$\begin{aligned} \sup_{R \in \mathcal{R}(\alpha)} \text{FDP}(R) &\leq \max_{R \in \mathcal{R}(\alpha)} \frac{|I_0 \cap R|}{\lceil (\alpha \gamma_{(|I_0 \cap R|)}^{\text{null}})^{-1} P_{(|I_0 \cap R|)}^{\text{null}} \rceil} \leq \max_{j \in [K_0]} \frac{j}{\lceil (\alpha \gamma_{(j)}^{\text{null}})^{-1} P_{(j)}^{\text{null}} \rceil} \\ &\leq \frac{\alpha}{\min_{j \in [K_0]} (j \gamma_{(j)}^{\text{null}})^{-1} P_{(j)}^{\text{null}}} = \frac{\pi_0 \alpha}{\min_{j \in [K_0]} (j \gamma_{(j)}^{\text{null}} / \pi_0)^{-1} P_{(j)}^{\text{null}}}. \end{aligned} \quad (11)$$

The first inequality is noting that there exists a null p-value and corresponding weight of index $i \in R$ that satisfies $(\gamma_{(i)}^{\text{null}})^{-1} P_{(i)}^{\text{null}} \leq (\gamma_{(|I_0 \cap R|)}^{\text{null}})^{-1} P_{(|I_0 \cap R|)}^{\text{null}}$ simply by the cardinality of $I_0 \cap R$. The second inequality is by changing the indexing set of the maximum, and the third inequality is by dropping the ceiling function in the denominator.

Now we note that

$$P^{\text{null-Simes}} := \min_{j \in [K_0]} (j \gamma_{(j)}^{\text{null}} / \pi_0)^{-1} P_{(j)}^{\text{null}}$$

is precisely the weighted Simes p-value applied only to the p-values and weights of the null hypotheses, i.e., I_0 . One rejects the weighted Simes p-value if and only if the weighted BH procedure makes any discoveries at level $\alpha \in [0, 1]$.

We now observe that trivially, we can augment the bound in (11) to get that

$$\sup_{R \in \mathcal{R}(\alpha)} \text{FDP}(R) \leq \left(\frac{\pi_0 \alpha}{P^{\text{null-Simes}}} \right) \wedge 1. \quad (12)$$

Now, let F denote the c.d.f. of $P^{\text{null-Simes}}$. Thus, we get the following derivation:

$$\begin{aligned} \mathbb{E} \left[\sup_{R \in \mathcal{R}(\alpha)} \text{FDP}(R) \right] &\leq \mathbb{E} \left[\left(\frac{\pi_0 \alpha}{P^{\text{null-Simes}}} \right) \wedge 1 \right] \\ &= \mathbb{P} (P^{\text{null-Simes}} \leq \pi_0 \alpha) + \mathbb{E} \left[\frac{\pi_0 \alpha}{P^{\text{null-Simes}}} \cdot \mathbb{1} \{ P^{\text{null-Simes}} > \pi_0 \alpha \} \right] \\ &= F(\pi_0 \alpha) + \int_{\pi_0 \alpha}^1 \frac{\pi_0 \alpha}{x} dF(x). \end{aligned}$$

The first inequality is by plugging in (12). The second equality follows from casing on the value of $P^{\text{null-Simes}}$ into disjoint cases, and the last equality comes from the definition of F .

Using integration by parts, it follows that

$$\begin{aligned} \mathbb{E} \left[\sup_{R \in \mathcal{R}(\alpha)} \text{FDP}(R) \right] &\leq F(\pi_0 \alpha) + \pi_0 \alpha \left(F(1) - \frac{F(\pi_0 \alpha)}{\pi_0 \alpha} + \int_{\pi_0 \alpha}^1 \frac{F(x)}{x^2} dx \right) \\ &= \pi_0 \alpha + \pi_0 \alpha \int_{\pi_0 \alpha}^1 \frac{F(x)}{x^2} dx = \pi_0 \alpha + \pi_0 \alpha \int_{\pi_0 \alpha}^1 \frac{\text{FDR}_0(x)}{x^2} dx. \end{aligned}$$

The last equality follows from the equivalence between Type I error of weighted Simes p-value and FDR of weighted BH under the global null (i.e., when applied to only null p-values). Hence we have shown our desired result. \square

G.4 Relationship between SupFDR and StopFDR

We noted in the introduction that a bound on the FDR at all stopping times also implies a bound on the SupFDR. In this section we give a formal derivation of this result.

Definition 7. Recall that a filtration $(\mathcal{F}_t)_{t \in \mathbb{N}}$ is a sequence of nested sigma-algebras, and stopping times (w.r.t. (\mathcal{F}_t)) are random variables $\tau \in \mathbb{N}$ such that $\mathbb{1}\{\tau = t\}$ is measurable w.r.t. \mathcal{F}_t .

Theorem G.3. *Let $(W_t)_{t \in \mathbb{N}}$ be a nonnegative process that is adapted to some filtration (\mathcal{F}_t) , meaning W_t is measurable with respect to \mathcal{F}_t for all $t \in \mathbb{N}$. Furthermore, let $(W_t)_{t \in \mathbb{N}}$ be almost surely bounded (with respect to \mathbb{P}) by some constant $c \geq 1$ and $\mathbb{E}[W_\tau] \leq 1$ for all stopping times τ with respect to (\mathcal{F}_t) . Then, $\mathbb{E}[\sup_{t \in \mathbb{N}} W_t] \leq 1 + \log(c)$.*

Proof. By Ville's inequality, we have that

$$\mathbb{P}\left(\sup_{t \in \mathbb{N}} W_t > \alpha^{-1}\right) \leq \alpha$$

for all $\alpha \in [0, 1]$. Thus, we can use the tail bound integration formula for expectation of nonnegative random variables:

$$\begin{aligned} \mathbb{E}\left[\sup_{t \in \mathbb{N}} W_t\right] &\leq \int_0^\infty \mathbb{P}\left(\sup_{t \in \mathbb{N}} W_t \geq s\right) ds \\ &= \int_0^c \mathbb{P}\left(\sup_{t \in \mathbb{N}} W_t \geq s\right) ds \leq \int_0^c s^{-1} \wedge 1 ds \\ &= \int_0^1 ds + \int_1^c s^{-1} ds = 1 + \log(c) \end{aligned}$$

□

Corollary G.4. *If a sequence of discovery sets (R_t) ensures StopFDR control at level α , then SupFDR is controlled at level $\alpha(1 + \log(\alpha^{-1}))$.*

The above theorem can also be seen as similar to defining an adjuster [24, 9, 8, 7] for upper bounded e-processes or calibrator [28, 29] for lower bounded anytime-valid p-values.

G.5 Sharpness of SupFDR and StopFDR control of online BH

We will show that the SupFDR and StopFDR control of online BH under PRDN p-values in Theorem 3.4 is sharp, as claimed in Theorem 3.5. First, we cite the following construction from Su [27] for a worst-case construction of the FDP of picking non-nulls as a function of the realized null p-values. Let there be K fixed hypotheses, and without loss of generality let $I_0 = [|K_0|]$, i.e., the nulls come first. Now, define the following quantities:

$$\begin{aligned} j^* &:= \operatorname{argmax}_{j \in I_0} j / \lceil KP_{(j)}^{\text{null}} / \alpha \rceil, \\ K_1^* &:= (\lceil KP_{(j^*)}^{\text{null}} / \alpha \rceil - j^*) \vee 0, \end{aligned}$$

where $P_{(1)}^{\text{null}}, \dots, P_{(K_0)}^{\text{null}}$ are the p-values corresponding to null hypotheses ordered from smallest to largest. Consider the following adversarial construction of non-null p-values for $t \in \{K_0 + 1, \dots, K\}$.

$$P_t = 1 - \mathbb{1}\{t \leq K_0 + K_1^*\}. \quad (13)$$

Fact 1 (Lemma 2.4 of Su [27]). Let R^{BH} be the discovery set of the BH procedure at level α applied to any p-values where the non-null p-values are set in the manner of (13). When $K_1^* \leq K - K_0$, then $\text{FDP}(R^{\text{BH}}) = \max_{j \in I_0} (j / [K p_{(j^*)}^{\text{null}} / \alpha]) \wedge 1$.

Figure 6 shows how we derive our problem instance for the online ARC setting from the aforementioned offline construction. Let $\gamma_t = K^{-1}$ for $t \in [K]$ and $\gamma_t = 0$ for $t > K$. We let the first K_0 p-values be identical to the offline setting and the remaining hypotheses be non-null with p-values set to 0. Now we note the following.

Lemma G.5. *For the above choice of p-values and weights for the online setting, we have that $R_{K_0+K_1^*}^{\text{OBH}} = R^{\text{BH}}$ and as a result, we get the following identity for the FDP when $K_1^* \leq K - K_0$:*

$$\text{FDP}(R_{K_0+K_1^*}^{\text{OBH}}) = (\max_{j \in I_0} j / [K p_{(j^*)}^{\text{null}} / \alpha]) \wedge 1.$$

Proof. Note that we can map p-value realizations of our offline construction to our online construction simply by setting the null p-values to be equivalent but setting all non-null p-values to be 0 in the online construction. Applying weighted BH with weights of $\gamma_t = K^{-1}$ to the first $K_0 + K_1^*$ hypotheses is exactly equivalent to applying standard BH to K hypotheses where the first $K_0 + K_1^*$ hypotheses have identical p-values to the weighted BH p-values, and the remaining p-values are 1, as the hypotheses with p-values of 1 will never be rejected. Thus, when combined with Fact 1, we can justify the above lemma. \square

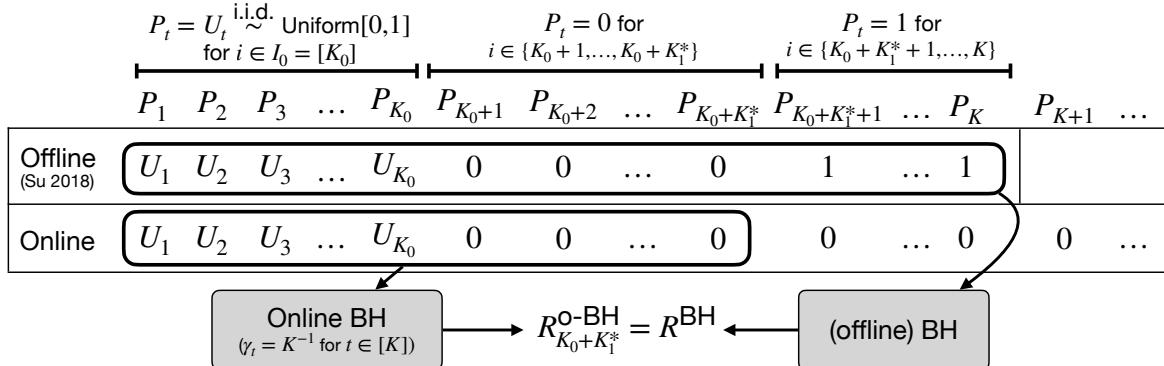


Figure 6: Visualization of our p-value construction for the online ARC setting that results in online BH producing discovery sets that are equal to the discovery sets in Su [27] produced by the (offline) BH procedure (which maximize the FDR). This occurs whenever $K_1^* \leq K - K_0$, and we formalize the discovery equivalence result in Lemma G.5.

Now, Su [27] showed that the FDP expression in the above lemma was sufficient to justify sharpness of the FDR bound when the null p-values are chosen to be i.i.d. uniform random variables.

Fact 2 (Theorem 6 of Su [27]). For every $\varepsilon > 0$, there exists a K , $K_0 < K$, and $\alpha' \in (0, 1)$ such that $\text{FDR}(R^{\text{BH}}) > (1 - \varepsilon)\alpha(1 + \log(\alpha^{-1}))$ for all $\alpha \leq \alpha'$ where the null p-values are i.i.d. uniform random variables, and the non-null p-values are set in accordance with (13).

Now, we note that $K_0 + K_1^*$ is a stopping time since it can be determined solely by the null p-values (which we will have observed by the K_0 th time step). Thus, if we apply our Lemma G.5 to the above fact, we have shown our desired result in Theorem 3.5.

As a result, SupFDR and StopFDR control of online BH is sharp in the PRDN case, and the $\log(\alpha^{-1})$ factor is unimprovable in the WNDN case.

Remark 7. We will also note what our example here means w.r.t. the sharpness of certain self-consistent procedures in the offline multiple testing setting.

- If we are thinking about the self-consistent set that has the maximum FDP over all self-consistent sets, then our construction does indicate that under i.i.d. null and non-null p-values, a FDR lower bound of $\alpha(1 + \log(\alpha^{-1}))$ is sharp.
- However, we can also restrict ourselves to discovery sets along the “BH path”, i.e., self-consistent sets that correspond to the smallest k p-values for some $k \in [K]$. The discovery sets used for our example would not be along the BH path, since the null p-values would not be 0 almost surely, and anything along the BH path would discover all the non-null p-values first (as they are all 0). In this case, we would still have i.i.d. null p-values, but some of the non-null p-values must be adversarially chosen in the sense of Su [27], i.e., setting the K_1^* non-null p-values to be 1, to restrict the “BH path” to reach the sharp lower bound of $\alpha(1 + \log(\alpha^{-1}))$.