

Admissible online closed testing must employ e-values

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

July 7, 2025

Abstract

In contemporary research, data scientists often test an infinite sequence of hypotheses H_1, H_2, \dots one by one, and are required to make real-time decisions without knowing the future hypotheses or data. In this paper, we consider such an online multiple testing problem with the goal of providing simultaneous lower bounds for the number of true discoveries in data-adaptively chosen rejection sets. Employing the recent online closure principle, we show that for this task it is necessary to use an anytime-valid test for each intersection hypothesis. This connects two distinct branches of the literature: online testing of multiple hypotheses (where the hypotheses appear online), and sequential anytime-valid testing of a single hypothesis (where the data for a fixed hypothesis appears online). Motivated by this result, we construct a new online closed testing procedure and a corresponding short-cut with a true discovery guarantee based on multiplying sequential e-values. This general but simple procedure gives uniform improvements over the state-of-the-art methods but also allows to construct entirely new and powerful procedures. Finally, we also propose the first online true discovery procedures for exchangeable and arbitrarily dependent e-values.

1 Introduction

Online multiple testing is a framework in which a potentially infinite stream of hypotheses H_1, H_2, \dots is tested one by one over time [12, 24]. At each step $t \in \mathbb{N}$ it needs to be decided on the current hypothesis H_t without knowing how many hypotheses are to be tested in the future, what those hypotheses are, and without having access to all the data relevant to testing them (indeed, that data may only be collected in the future). This setting occurs in the tech industry [28, 36], machine learning [9, 64], open data repositories as used in genomics [1, 33, 6] and other data science tasks where flexible and real-time decision making is required.

A common error metric for a chosen rejection set R_t at time t is the *false discovery proportion*

$$\text{FDP}(R_t) = \frac{\text{Number of true hypotheses in } R_t}{\text{Size of } R_t}. \quad (1)$$

The usual approach, both in online and classical offline testing, is to control the *expected value* of $\text{FDP}(R_t)$, also known as the false discovery rate (FDR) [3], below some level $\alpha \in (0, 1)$. However, since the FDP may have high variance, controlling its expectation may not be enough.

In a seminal work, Goeman and Solari [18] proposed to control the tail probabilities of $\text{FDP}(S)$ *simultaneously over all possible sets* S instead. That means they suggest to provide an upper bound $q(S)$ for $\text{FDP}(S)$ such that the probability that there is any set S with $\text{FDP}(S) > q(S)$, is less or equal than α .

A major advantage of such simultaneous bounds on the FDP, compared to FDR control, is that the final rejection set(s) can be chosen post-hoc, meaning after looking at the data and calculating




Figure 1: Illustration of using an online procedure with simultaneous true discovery guarantee. At each time t , (1) the scientist proposes a hypothesis H_t for testing, possibly based on the data used for testing the previous hypotheses; (2) the scientist collects the data X_t required for testing H_t which may consist of new data and/or the reuse of old data; (3) the scientist chooses based on the data observed so far (possibly multiple) query sets $S \subseteq \{1, \dots, t\}$ of interest — if the scientist is interested in rejections, these query sets could also be interpreted as candidates for rejection; (4) the statistician employs the online procedure \mathbf{d} to provide lower bounds $\mathbf{d}(S)$ for the number of true discoveries in the query sets S requested by the scientist that hold simultaneously with probability of at least $1 - \alpha$.

the bounds $\mathbf{q}(S)$, without violating this error control. In other words, a scientist is allowed to query many (or all) such sets S , examine the reported bounds, and later choose one or a few final sets (and bounds) to report or follow up on. Since many more sets S will be queried than will be rejected, we call these sets S as *query* sets (rather than, say, *rejection* sets).

A bound that holds with high probability can be advantageous in applications where an inflated FDP has severe consequences. However, there are also connections between FDR and simultaneous FDP controlling procedures as illustrated by Goeman et al. [19], Katsevich and Ramdas [26], Meah et al. [32]. As done by Goeman et al. [20], we will consider simultaneous *lower bounds on the number of true discoveries* $\mathbf{d}(S) := (1 - \mathbf{q}(S))|S|$ in this paper, which is mathematically equivalent to *upper bounding the FDP*, but easier to handle. Their work focused on offline multiple testing, while we study this error metric in the online setting.

At each time t , an *online true discovery procedure* \mathbf{d} allows the scientist to pick any (and possibly multiple) query sets $S \subseteq \{1, \dots, t\}$ of interest. Then $\mathbf{d}(S)$ immediately provides a lower bound for the number of false hypotheses in S , which holds true with probability at least $1 - \alpha$ simultaneously over all sets that have been queried or that might be queried in the future. Note that this simultaneity permits to stop (or continue) the testing process data-adaptively at any time, e.g. after the 50th discovery. This is an additional benefit compared to online FDR control, which is usually only provided at fixed times [65, 62]. In Figure 1, we illustrate the scientific discovery process when using online true discovery procedures.

A recently popularized approach for online testing of a single hypothesis, where not the hypotheses but the data itself comes in sequentially, is the e-value. An e-value E_t for a hypothesis H_t is a nonnegative random variable which has expected value less or equal than one if H_t is true. The e-value (or its sequential extension, the e-process) is an alternative to the well-known p-value but more suitable for settings where early stopping of the sampling process or optional continuation is desired [42, 21, 38]. In this paper, we exploit this sequential suitability of the e-value to design online multiple testing procedures with simultaneous true discovery guarantees. This connects two mostly separate areas of the literature: online testing of multiple hypotheses and sequential testing of a single hypothesis; we elaborate on this below.

1.1 Related literature

Our work mixes ingredients from different subfields of sequential and multiple hypothesis testing.

The e-value has recently emerged as a fundamental concept in composite hypothesis testing and

underlying a universal approach to anytime-valid inference [60, 41, 50, 21, 37], but the roots can be traced back to the works of Ville [48], Wald [57] and Robbins [7]. A recent overview of the e-value literature is given by Ramdas et al. [38] and Ramdas and Wang [35, Chapter 1].

Interest in e-values has grown rapidly in recent years, including in particular multiple testing with e-values. Wang and Ramdas [59] introduced and analyzed the e-BH procedure, an e-value variant of the popular Benjamini-Hochberg (BH) procedure [3] for FDR control. Vovk and Wang [50] explored the possibility of combining several e-values by averaging and multiplication. They also used this to derive multiple testing procedures with familywise error rate (FWER) control by applying the closure principle [31] with these combination rules. The FWER is a strict error criterion defined as the probability of rejecting any true null hypothesis. Vovk and Wang [51, 53] extended these ideas to obtain procedures with a true discovery guarantee. All the aforementioned approaches consider classical *offline* multiple testing.

Online multiple testing initially focused on procedures for p-values [12, 1, 24, 36]. An overview of this literature was recently provided by Robertson et al. [39]. Xu and Ramdas [63] is the sole paper to consider online multiple testing with e-values, focusing on FDR control for dependent e-values.

A related line of work investigates simultaneous true discovery guarantees by closed testing, mostly focusing on offline settings with p-values. The closure principle was initially proposed and analyzed for FWER control [31, 43, 40]. However, Goeman and Solari [18] noted that the same principle can be applied to obtain simultaneous true discovery bounds, a more general and less conservative task than FWER control, although a similar approach was proposed earlier by Genovese and Wasserman [16, 17]. Many works have since built on these [19, 47, 22, 30]. Importantly, Goeman et al. [20] proved that all admissible procedures for bounding the true discovery proportion must employ closed testing.

The recent work of Fischer et al. [11] showed how the closure principle can also be used for *online* multiple testing. However, their investigation of admissibility and construction of concrete procedures is restricted to FWER control. The current work extends their ideas to lower bounds on the true discovery proportion.

Another related work to ours is by Katsevich and Ramdas [26]. They proposed various p-value based true discovery procedures for structured, knockoff and also online settings exploiting martingale techniques. Meah et al. [32] modified and improved some of their methods with a focus on m-consistency, a property that relates true discovery procedures to FDR. Our work will uniformly improve the methods by Katsevich and Ramdas [26] and Meah et al. [32] for the online setting.

1.2 Our contribution

We provide new insights about the central role of e-values in online multiple testing. In particular, we show that in order to derive admissible online procedures with simultaneous true discovery guarantee, we must necessarily employ *anytime-valid tests* for each intersection hypothesis, which in turn *must* employ test martingales [37]. By transitivity, the construction of these online true discovery procedures must rely on test martingales, which are sequential generalizations of e-values [38]. Thus, e-values enter naturally into the construction of admissible online procedures. Remarkably, even if an online procedure is constructed solely based on p-values, it must implicitly employ anytime-valid tests and can be reconstructed or improved using e-values. Guided by these theoretical results, we construct **SeqE-Guard**, a powerful and computationally efficient online true discovery procedure based on multiplying sequential e-values.

In Table 1, we give a brief comparison of our theoretical results to the related works. Our paper is the first connecting online multiple testing to anytime-valid testing of a single hypothesis. Furthermore, our setting generalizes the ones considered by Goeman et al. [20] (since the online setting generalizes the offline setting) and Fischer et al. [11] (since true discovery guarantee generalizes FWER control). Consequently, our results immediately apply in their settings, but their results do not imply ours.

Paper	Key theoretical messages
Goeman et al. [20]	<ul style="list-style-type: none"> • Admissible closed testing is sufficient for admissibility of monotone stacks of true discovery procedures • All admissible true discovery procedures are closed procedures
Fischer et al. [11]	<ul style="list-style-type: none"> • Increasing families of online intersection tests are sufficient to obtain online closed procedures • All admissible FWER controlling online procedures are online closed procedures
Ramdas et al. [37]	<ul style="list-style-type: none"> • Sufficient conditions for admissibility of anytime-valid tests based on test martingales • All admissible anytime-valid tests must employ test martingales
This paper	<ul style="list-style-type: none"> • Online closed procedures must employ anytime-valid tests for the intersection hypotheses • All admissible coherent online true discovery procedures are online closed procedures.

Table 1: Comparison of the key messages of closely related theoretical works. Our main claim, the first shown above, connects the work by Ramdas et al. [37] with the online multiple testing literature [11, 26, 32] and thus allows to apply the powerful anytime-valid testing concept of test martingales to online multiple testing. Furthermore, our second claim implies and unifies the second claims of Goeman et al. [20] and Fischer et al. [11], since the online setting generalizes the offline setting and true discovery guarantee generalizes FWER control (and all admissible procedures are coherent in these settings).

In Table 2 we list all the existing online procedures with true discovery guarantee and show the sequential e-values that yield uniform improvements when plugged into our **SeqE-Guard** algorithm. In addition to larger bounds, our approach reduces these complex methods with multiple parameters to simple e-values that can only take two different values (or three in case of the **online-adaptive** procedure). This facilitates the interpretation and comparison of the different methods.

1.3 Paper outline

In Section 2, we define the online setting formally and recap concepts like coherence (Section 2.1) and (online) closed testing (Section 2.2). Afterwards, we introduce a general approach to online true discovery guarantee based on test martingales and prove that every procedure must be constructed in that way (Section 2.3).

In Section 3, we consider online true discovery guarantee with sequential e-values and propose our **SeqE-Guard** algorithm for this task. By plugging specific sequential e-values into **SeqE-Guard** we immediately obtain uniform improvements of the state-of-the-art methods by Katsevich and Ramdas [26], Meah et al. [32] (Section 3.2). Separately, we investigate the use of **SeqE-Guard** with growth rate optimal (GRO) e-values [41, 21] and propose a hedging approach to increase the power to detect false hypotheses (Sections 3.3). In Section 4, we perform simulations to compare the proposed methods and to quantify the gain in power obtained by our improvements¹. Proofs for all theoretical results are provided in the Supplementary Material S.5.

Finally, we provide new true discovery procedures for online settings where the e-values are not sequential but exchangeable (Section 5) or arbitrarily dependent (Section 6).

¹The code for the simulations is available at github.com/fischer23/online_true_discovery.

Improved method	Sequential e-values	Algorithm
online-simple [26]	$E_i^{\text{os}} = \exp[\theta_c(\mathbb{1}\{P_i \leq \alpha_i\} - c\alpha_i)]$	SeqE-Guard
	$a > 0$	
	$c = \frac{\log(1/\alpha)}{a \log(1 + \log(1/\alpha)/a)}$	
	$\theta_c = \log(1/\alpha)/(ca)$	
online-adaptive [26]	$E_i^{\text{ad}} = \exp[\theta_c(\mathbb{1}\{P_i \leq \alpha_i\} - c_{\frac{\alpha_i}{1-\lambda_i}} \mathbb{1}\{P_i > \lambda_i\})]$	SeqE-Guard
	$\lambda_i \in [\alpha_i, 1)$	
	$a > 0$	
	$c = \frac{\log(1/\alpha)}{a \log(1 + (1 - \alpha^B/a)/B)}$	
	$B = \sup_{i \in \mathbb{N}} \frac{\alpha_i}{1 - \lambda_i}$	
	$\theta_c = \log(1/\alpha)/(ca)$	
u-online-simple [32]	$E_i^{\text{os},a} = \exp[\theta_{c_a}(\mathbb{1}\{P_i \leq \alpha_i\} - c_a \alpha_i)]$	SeqE-Guard + union over $a \in \mathbb{N}$
	$c_a = \frac{\log(1/\alpha(a))}{a \log(1 + \log(1/\alpha(a))/a)}$	
	$\alpha(a) = \frac{6\alpha}{a^2 \pi^2}$	
	$\theta_{c_a} = \log(1/\alpha(a))/(c_a a)$	
u-online-Freedman [32]	$E_i^{\text{Freed},a} = \exp[\lambda_a(\mathbb{1}\{P_i \leq \alpha_i\} - \alpha_i) - \psi(\lambda_a)\alpha_i(1 - \alpha_i)]$	SeqE-Guard + union over $a = 2^{j/2}$, $j \in \mathbb{N} \cup \{0\}$
	$\lambda_a = \log(1 + \frac{\kappa_a}{a})$	
	$\psi(\lambda_a) = \exp(\lambda_a) - \lambda_a - 1$	
	$\kappa_a = \sqrt{2a \log(1/\alpha(a))} + \frac{\log(1/\alpha(a))}{2}$	
	$\alpha(a) = \alpha \left(\frac{6}{\max(2 \log_2(a), 1)^2 (\pi^2 + 6)} \right)$	

Table 2: List of methods that can be uniformly improved by our SeqE-Guard algorithm. The left column shows the name of the existing method proposed in past work. The middle column gives the sequential e-value, which if plugged into the algorithm in the right column, yields a uniform improvement of the existing method in the left column. To the best of our knowledge this list includes all existing nontrivial online true discovery procedures. Thus, we claim that our approach uniformly improves all existing procedures. It should be noted that the shown sequential e-values may be further improved by ensuring that their expected value equals exactly 1 in case of uniformly distributed p-values (this can be accomplished by dividing the shown e-values by their expected value but we omit it here to keep the expressions short).

2 Online true discovery guarantee

In this section, we introduce general notation and recall concepts like simultaneous (online) true discovery guarantee and (online) closed testing. Then, we prove that any admissible procedure for delivering an online true discovery guarantee must rely on test martingales (sequential generalizations of e-values).

We consider the general online multiple testing setting described in [11]. Let (Ω, \mathbb{F}) , where $\mathbb{F} = (\mathcal{F}_i)_{i \in \mathbb{N}_0}$ be a filtered measurable space and \mathcal{P} some set of probability distributions on (Ω, \mathbb{F}) . The σ -field \mathcal{F}_i defines the information that can be used for testing hypothesis H_i ($\mathcal{F}_0 = \{\emptyset, \Omega\}$). Hence, in online multiple testing every hypothesis test is only allowed use some partial information which is increasing over time. One can think of \mathcal{F}_i as the data that is available at time i . However, we might want to add external randomization or coarsen the filtration. For example, many existing works on online multiple testing consider $\mathcal{F}_i = \sigma(P_1, \dots, P_i)$ [12, 24, 11], where each P_j is a p-value calculated for hypothesis H_j .

We assume that the data follows some unknown distribution $\mathbb{P} \in \mathcal{P}$. A null hypothesis H is simply a collection of probability distributions (a subset of \mathcal{P}); we are effectively testing whether $\mathbb{P} \in H$ or not. When $\mathbb{P} \in H$, we say that H is true null, and otherwise we call it a false null. We

define $I_0^{\mathbb{P}} := \{i \in \mathbb{N} : \mathbb{P} \in H_i\}$ and $I_1^{\mathbb{P}} := \mathbb{N} \setminus I_0^{\mathbb{P}}$ as the index sets of true and false null hypotheses, respectively.

As defined in [20, 11], a procedure with *simultaneous true discovery guarantee* is a random function $\mathbf{d} : 2^{\mathbb{N}_f} \rightarrow \mathbb{N} \cup \{0\}$, where $2^{\mathbb{N}_f}$ is the set of all finite subsets of \mathbb{N} (analogously we use $2^{\mathbb{N}-f}$ for the set of all infinite subsets of \mathbb{N}), such that for all $\mathbb{P} \in \mathcal{P}$:

$$\mathbb{P}(\mathbf{d}(S) \leq |S \cap I_1^{\mathbb{P}}| \text{ for all } S \in 2^{\mathbb{N}_f}) \geq 1 - \alpha.$$

Clearly, if \mathbf{d} always outputs 0, it is a valid procedure. Thus, implicitly, the larger \mathbf{d} is, the better (here, larger is meant componentwise; $\mathbf{d} \geq \mathbf{d}'$ if $\mathbf{d}(S) \geq \mathbf{d}'(S)$ for all $S \in 2^{\mathbb{N}_f}$).

\mathbf{d} is called an *online* true discovery procedure if $\mathbf{d}(S)$ is measurable with respect to $\mathcal{F}_{\max(S)}$ for all $S \in 2^{\mathbb{N}_f}$ [11]. This ensures that at any time $t \in \mathbb{N}$ the procedure \mathbf{d} provides a lower bound for the number of false hypotheses in every set $S \subseteq \{1, \dots, t\}$ with $\max(S) = t$ that remains valid no matter how many hypotheses will be tested in the future.

Note that one does not have to consider an infinite number of hypotheses but could just stop at some finite time $i \in \mathbb{N}$ by setting $H_j = \mathcal{P}$ for all $j > i$ and $\mathbf{d}(S) = \mathbf{d}(S \cap \{1, \dots, i\})$ for all $S \in 2^{\mathbb{N}_f}$ with $\max(S) > i$. With this, the online setting becomes classical offline testing in the case of $\mathcal{F}_1 = \mathcal{F}_2 = \dots$ and thus online multiple testing can be seen as a true generalization of classical multiple testing [11]. Although we are mainly interested in the strict online case $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \dots$, this particularly implies that all online procedures constructed in this paper also apply in the offline setting.

We would also like to point out that the hypotheses do not have to be prespecified. Hence, in practice one is allowed to data-adaptively construct each H_i based on \mathcal{F}_{i-1} .

Note that simultaneous lower bounds for the number of true discoveries instantly provide bounds for many other error rates such as the k-FWER or false discovery exceedance (FDX) [20]. Furthermore, true discovery guarantee is equivalent to controlling the false discovery proportion defined in (1) [20]. For $\alpha = 0.5$, this yields the noteworthy special case of median FDP control (as compared to FDR control, which bounds the mean of the FDP).

In Figure 2, we connect the results that will be given in the rest of the section and the related works [37, 11]. The left-hand side provides a general approach to construct coherent online true discovery procedures based on test martingales and the right-hand side proves that every admissible online procedure must be constructed in that way. While most of the upper and lower connections are known, the main insights provided by this paper are the middle connections between increasing families of online intersection tests and anytime-valid tests, connecting two distinct branches of the literature. All terms and results will be clarified in the following subsections from top to bottom, starting with *coherent online procedures*.

2.1 Coherent online true discovery procedures

An important property of multiple testing procedures is *coherence* [14, 43]. A true discovery procedure \mathbf{d} is called *coherent* [20], if for all disjoint $S, U \in 2^{\mathbb{N}_f}$ it holds that

$$\mathbf{d}(S) + \mathbf{d}(U) \leq \mathbf{d}(S \cup U) \leq \mathbf{d}(S) + |U|. \quad (2)$$

Coherence ensures consistent decisions or bounds of the multiple testing procedure and is therefore a desirable property. A procedure \mathbf{d} is admissible if there is no other procedure $\tilde{\mathbf{d}}$ that uniformly improves \mathbf{d} , where $\tilde{\mathbf{d}}$ is said to uniformly improve \mathbf{d} , if $\tilde{\mathbf{d}} \geq \mathbf{d}$ and $\mathbb{P}(\tilde{\mathbf{d}}(S) > \mathbf{d}(S)) > 0$ for at least one $\mathbb{P} \in \mathcal{P}$ and $S \in 2^{\mathbb{N}_f}$. Equivalently, \mathbf{d} is admissible if $\tilde{\mathbf{d}} \geq \mathbf{d}$ implies $\tilde{\mathbf{d}} = \mathbf{d}$.

Goeman et al. [20] showed that in the offline setting, all admissible true discovery procedures must be coherent. However, it turns out that this result is not true in the online case.

Example 1. Consider a setting with only two hypotheses H_1 and H_2 with independent p -values P_1 and P_2 that are uniformly distributed under the null hypothesis. Let $\mathbf{d}(\{1\}) = 1$, if $P_1 \leq \alpha/2$ (and 0

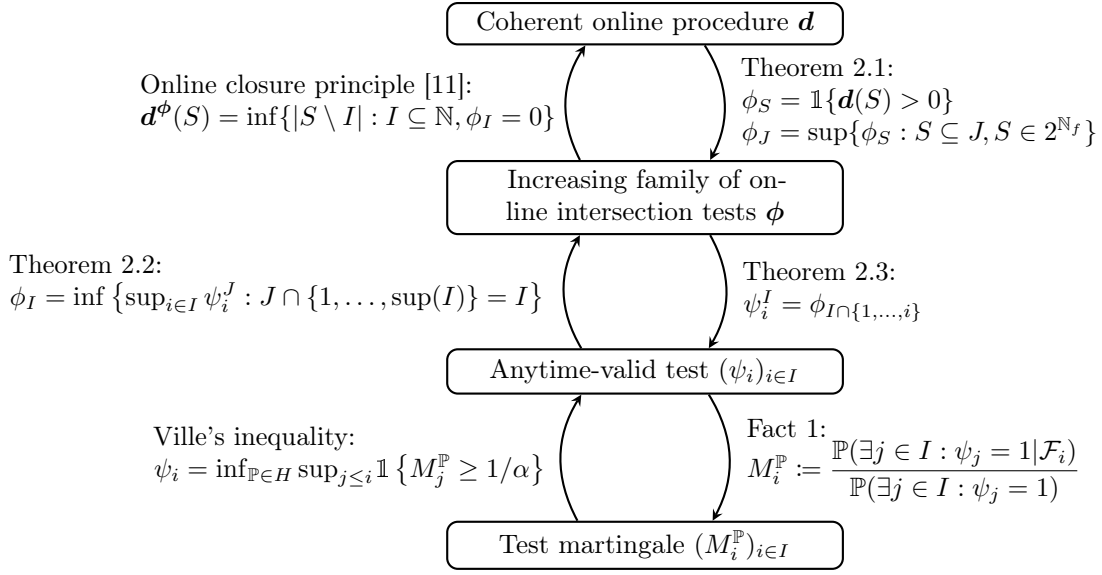


Figure 2: Illustration of the relation between online true discovery procedures, increasing families of online intersection tests, anytime-valid tests and test martingales. The left path from bottom to top provides a general approach for the construction of coherent online procedures with true discovery guarantee based on test martingales. The right path from top to bottom shows that coherent online procedure implicitly define test martingales. Taken together, the loop shows, in principle, how to improve any given coherent online procedure: we take the right path down, and then take the left path up. In particular, if the original procedure is admissible, following the loop must leave the procedure unchanged.

otherwise), $\mathbf{d}(\{2\}) = 1$, if $P_2 \leq \alpha/2$ (and 0 otherwise) and

$$\mathbf{d}(\{1, 2\}) = \begin{cases} 2, & (P_1 \leq \alpha/2 \wedge P_2 \leq \alpha) \vee (P_2 \leq \alpha/2 \wedge P_1 \leq \alpha) \\ 1, & (P_1 \leq \alpha/2 \wedge P_2 > \alpha) \vee (P_2 \leq \alpha/2 \wedge P_1 > \alpha) \\ 0, & \text{otherwise.} \end{cases}$$

Then \mathbf{d} is an incoherent online procedure, since $\mathbf{d}(\{1, 2\}) = 2$ and $\mathbf{d}(\{1\}) = 0$ if $P_2 \leq \alpha/2$ and $\alpha/2 < P_1 \leq \alpha$. However, it is not possible to improve \mathbf{d} by a coherent online procedure $\tilde{\mathbf{d}} \geq \mathbf{d}$. To see this, note that in order to be coherent, $\tilde{\mathbf{d}}(\{1\})$ needs to equal 1, if $P_2 \leq \alpha/2 \wedge P_1 \leq \alpha$. Since $\tilde{\mathbf{d}}$ is supposed to be an online procedure and thus $\tilde{\mathbf{d}}(\{1\})$ must not use information about P_2 , we must have $\mathbf{d}(\{1\}) = 1$, if $P_1 \leq \alpha$. However, this implies

$$\mathbb{P}(\{\tilde{\mathbf{d}}(\{1\}) = 0\} \cap \{\tilde{\mathbf{d}}(\{2\}) = 0\}) \leq \mathbb{P}(\{P_1 > \alpha\} \cap \{P_2 > \alpha/2\}) = (1 - \alpha)(1 - \alpha/2) < 1 - \alpha$$

for all $\mathbb{P} \in H_1 \cap H_2$ such that true discovery guarantee is not provided.

Although not all admissible online true discovery procedures must be coherent, we think it is sensible to focus on coherent online procedures as incoherent results are difficult to interpret and communicate.

2.2 Online closed testing

The closure principle was originally proposed for FWER control [31]. However, Goeman and Solari [18] noted that it can also be used for the more general task of providing a simultaneous (offline) true

discovery guarantee. For each intersection hypothesis $H_I = \bigcap_{i \in I} H_i$, let ϕ_I be an intersection test and $\phi = (\phi_I)_{I \subseteq \mathbb{N}}$ denote the family of intersection tests. Throughout the paper, it is understood that all tests are α -level tests, meaning $\mathbb{P}(\phi_I = 1) \leq \alpha$ for all $\mathbb{P} \in H_I$. Goeman and Solari [18] showed that the closed procedure defined by

$$\mathbf{d}^\phi(S) := \inf\{|S \setminus I| : I \subseteq \mathbb{N}, \phi_I = 0\} \quad (S \in 2^{\mathbb{N}_f}) \quad (3)$$

provides simultaneous guarantee of the number of true discoveries over all $S \in 2^{\mathbb{N}_f}$. Technically, Goeman and Solari [18] only considered a finite number of hypotheses, but the method and its guarantees extend to a countable number of hypotheses [11], and so we present that version for easier connection to the online setting. It should also be noted that Genovese and Wasserman [16, 17] introduced an equivalent procedure that was not derived by the closure principle [31].

Goeman et al. [20] proved an important result: every (offline) coherent true discovery procedure is equivalent to or uniformly improved by a closed procedure of the form (3). Therefore, the closure principle allows to construct and analyze all admissible true discovery procedures based on single tests for the intersection hypotheses which are usually much easier to handle.

Fischer et al. [11] showed that the closed procedure \mathbf{d}^ϕ is an online procedure, if the following two assumptions are fulfilled.

1. Every intersection test ϕ_I , $I \subseteq \mathbb{N}$, is an *online intersection test*, meaning ϕ_I is measurable with respect to $\mathcal{F}_{\sup(I)}$.
2. The family of intersection tests $\phi = (\phi_I)_{I \subseteq \mathbb{N}}$ is *increasing*², which means that for all $i \in \mathbb{N}$ and $I \subseteq \{1, \dots, i\}$ it holds

$$\phi_I \leq \phi_{I \cup K} \text{ for all } K \subseteq \{k \in \mathbb{N} : k > i\}. \quad (4)$$

A closed procedure \mathbf{d}^ϕ , where ϕ satisfies these two conditions, is called an *online closed procedure*. Every online closed procedure is a coherent online procedure, which follows immediately from the same result in the offline case [20]. Fischer et al. [11] proved that all online procedures with FWER control can be written as a closed procedure where the intersection tests satisfy the conditions 1 and 2. Hence we know that the closure principle is admissible for offline true discovery control [20] and online FWER control [11], but not yet for coherent online true discovery control. We now fill this gap with the following result which shows that any coherent online procedure can be recovered or improved by an online closed procedure.

Theorem 2.1. *Let \mathbf{d} be a coherent online procedure. Define*

$$\phi_S = \mathbb{1}\{\mathbf{d}(S) > 0\} \quad \forall S \in 2^{\mathbb{N}_f} \quad \text{and} \quad \phi_J = \sup\{\phi_S : S \subseteq J, S \in 2^{\mathbb{N}_f}\} \quad \forall J \in 2^{\mathbb{N}_f}. \quad (5)$$

Then $\phi = (\phi_I)_{I \subseteq \mathbb{N}}$ is an increasing family of online intersection tests and $\mathbf{d}^\phi \geq \mathbf{d}$.

With this result, we can focus on online closed procedures when considering coherent online true discovery guarantees. In the following subsection, we introduce our main result and show that to construct online closed procedures, it is necessary to define anytime-valid tests and therefore one must rely on test martingales. This yields the connected graph in Figure 2 and motivates the construction of e-value based online closed procedures, which we will focus on afterwards.

Note that Theorem 2.1 particularly holds in the offline case $\mathcal{F}_1 = \mathcal{F}_2 = \dots$ and therefore immediately yields the aforementioned result by Goeman et al. [20] as a corollary, implying that the requirement of 1 and 2 is not a restriction.

²Fischer et al. [11] used the term *predictable* for (4), but we use *increasing* in order to avoid confusion with the measure-theoretic definition of predictability.

Remark 1. *There are indeed cases where the closed procedure based on the intersection tests defined in (5) dominates the original procedure. For example, consider three hypotheses and the coherent true discovery procedure \mathbf{d} with $\mathbf{d}(\{1, 2, 3\}) = 1$, $\mathbf{d}(\{1, 2\}) = 1$, $\mathbf{d}(\{1, 3\}) = 1$, $\mathbf{d}(\{2, 3\}) = 1$ and $\mathbf{d}(\{i\}) = 0$ for $i \in \{1, 2, 3\}$. The corresponding closed procedure would give the same bounds except for further concluding that $\mathbf{d}(\{1, 2, 3\}) = 2$. This makes sense, since if there is at least one true discovery in $\{1, 2\}$, one in $\{1, 3\}$ and one in $\{2, 3\}$, there should be at least two true discoveries in $\{1, 2, 3\}$.*

2.3 Admissible coherent online procedures for true discovery guarantees must rely on test martingales

In this section, we make connections between increasing families of online intersection tests and anytime-valid tests, which then reveal close relations of online procedures with true discovery guarantee and test martingales.

Suppose we have an anytime-valid test $(\psi_i^I)_{i \in I}$ for each intersection hypothesis H_I . Following [38], $(\psi_i^I)_{i \in I}$ is an anytime-valid test for H_I , if ψ_i^I is measurable with respect to \mathcal{F}_i and we have $\mathbb{P}(\exists i \in I : \psi_i^I = 1) \leq \alpha$ for all $\mathbb{P} \in H_I$. The following theorem shows how we can use such anytime-valid tests to construct an increasing family of online intersection tests.

Theorem 2.2. *For all $I \subseteq \mathbb{N}$ let $(\psi_i^I)_{i \in I}$ be an anytime-valid test for H_I . Then $\phi = (\phi_I)_{I \subseteq \mathbb{N}}$, where*

$$\phi_I = \inf \left\{ \sup_{i \in I} \psi_i^J : J \cap \{1, \dots, \sup(I)\} = I \right\}, \quad (6)$$

is an increasing family of online intersection tests. In particular, for infinite $I \subseteq \mathbb{N}$, we just have $\phi_I = \sup_{i \in I} \psi_i^I$.

In the following theorem, we show a converse relationship, meaning that increasing families of online intersection tests implicitly define anytime-valid tests for the intersection hypotheses. In addition, we prove that every increasing family of online intersection tests can be constructed by anytime-valid tests using (6).

Theorem 2.3. *Let $\phi = (\phi_I)_{I \subseteq \mathbb{N}}$ be an increasing family of online intersection tests. Then $(\psi_i^I)_{i \in I}$, where*

$$\psi_i^I = \phi_{I \cap \{1, \dots, i\}}, \quad (7)$$

is an anytime-valid test for H_I for all $I \subseteq \mathbb{N}$. Furthermore, let $\tilde{\phi} = (\tilde{\phi}_I)_{I \subseteq \mathbb{N}}$ be defined by $(\psi_i^I)_{i \in I}$, $I \subseteq \mathbb{N}$, using (6). Then $\tilde{\phi}_S = \phi_S$ for all $S \in 2^{\mathbb{N}}$ and $\mathbf{d}^{\tilde{\phi}} = \mathbf{d}^{\phi}$.

Together, Theorems 2.1 and 2.3 imply that in order to define admissible coherent online true discovery procedures, we need to construct anytime-valid tests for the intersection hypotheses. In addition, they imply that every coherent online true discovery procedure implicitly defines nontrivial anytime-valid tests by (5) and (7) for each intersection hypothesis.

These results are particularly interesting, since Ramdas et al. [37] gave a precise characterization of anytime-valid tests, proving that every anytime-valid test can be reconstructed or uniformly improved using test martingales. To state their result more precisely, we need to introduce some terminology.

An anytime-valid test $(\tilde{\psi}_i)_{i \in I}$ uniformly improves $(\psi_i)_{i \in I}$, if $\tilde{\psi}_i \geq \psi_i$ for all $i \in I$ and $\mathbb{P}(\tilde{\psi}_i > \psi_i) > 0$ for some $i \in I$ and $\mathbb{P} \in \mathcal{P}$. Furthermore, a nonnegative process $(M_i)_{i \in I \cup \{0\}}$ adapted to the filtration $(\mathcal{F}_i)_{i \in I \cup \{0\}}$, where $\mathcal{F}_0 = \emptyset$, is a test (super)martingale for \mathbb{P} , if $M_0 \stackrel{(\leq)}{\leq} 1$ and $\mathbb{E}_{\mathbb{P}}[M_i | \mathcal{F}_{i-}] \stackrel{(\leq)}{\leq} M_{i-}$ for all $i \in I$, where $i- = \max\{j \in I \cup \{0\} : j < i\}$. We call $(M_i)_{i \in I \cup \{0\}}$ a test (super)martingale for a null hypothesis H_I , if the above holds for all $\mathbb{P} \in H_I$. In the following we formally state the result by Ramdas et al. [37] (adapted to our setup) and provide a self-contained proof in the Supplementary Material S.5.

Fact 1. Let $(\psi_i)_{i \in I}$ be an anytime-valid test for H_I , $I \subseteq \mathbb{N}$, and define for all $\mathbb{P} \in H_I$:

$$M_i^{\mathbb{P}} := \frac{\mathbb{P}(\exists j \in I : \psi_j = 1 | \mathcal{F}_i)}{\mathbb{P}(\exists j \in I : \psi_j = 1)} \quad (i \in I \cup \{0\})$$

with $M_i^{\mathbb{P}} = 1$ if the denominator equals zero. Then $(M_i^{\mathbb{P}})_{i \in I \cup \{0\}}$ is a test martingale for \mathbb{P} . Furthermore, let

$$\tilde{\psi}_i = \inf_{\mathbb{P} \in H_I} \sup_{j \leq i} \mathbb{1}\{M_j^{\mathbb{P}} \geq 1/\alpha\} \quad (i \in I).$$

Then $(\tilde{\psi}_i)_{i \in I}$ is an anytime-valid test for H_I and either equals or uniformly improves $(\psi_i)_{i \in I}$.

Fact 1 closes the loop illustrated in Figure 2. This particularly shows that every coherent online true discovery procedure \mathbf{d} implicitly constructs test martingales for the intersection hypotheses. Furthermore, taking the left path in Figure 2 with these test martingales always yields a procedure which either equals or uniformly improves the original procedure \mathbf{d} .

3 Online true discovery guarantee with sequential e-values

In the previous section we showed that we need to construct anytime-valid tests, and thus test martingales, for the intersection hypotheses when constructing (admissible) coherent online procedures with a true discovery guarantee. This general martingale-based approach is illustrated in the left path from bottom to top in Figure 2. Since every step involves taking the infimum over a large set, it seems computationally inefficient. However, in practice one can avoid this by using the same test martingale for all $\mathbb{P} \in H$, the same anytime-valid test $(\psi_i^J)_{i \in I}$ for all J with $J \cap \{1, \dots, \sup(I)\} = I$ and constructing ϕ in a way that permits a short-cut for \mathbf{d}^ϕ . In this section, we apply all this to derive a computationally efficient and powerful online true discovery procedure based on sequential e-values. We first introduce our general **SeqE-Guard** algorithm (Section 3.1), then show that it allows to uniformly improve all existing online true discovery procedures by using specific sequential e-values (Section 3.2) and finally analyze its behavior when using GRO e-values (Section 3.3).

3.1 The SeqE-Guard algorithm

Let $(M_t)_{t \in \mathbb{N}_0}$ be a test supermartingale with respect to $(\mathcal{F}_t)_{t \in \mathbb{N}_0}$ for some hypothesis H . We can break down $(M_t)_{t \in \mathbb{N}_0}$ into its individual factors $E_t = \frac{M_t}{M_{t-1}}$, $t > 0$, with the convention $0/0 = 0$. Due to the supermartingale property, E_t is nonnegative, measurable with respect to \mathcal{F}_t and

$$\mathbb{E}_{\mathbb{P}}[E_t | \mathcal{F}_{t-1}] = \mathbb{E}_{\mathbb{P}}[M_t / M_{t-1} | \mathcal{F}_{t-1}] = \mathbb{E}_{\mathbb{P}}[M_t | \mathcal{F}_{t-1}] / M_{t-1} \leq 1 \quad (\mathbb{P} \in H). \quad (8)$$

Hence, each of these random variables E_t is an e-value for H conditional on the past, sometimes called a sequential e-value in the literature [50]. Thus, every test supermartingale can be written as the product of sequential e-values. Conversely, every product of sequential e-values defines a test supermartingale (just multiply M_{t-1} on both sides of (8)). For these reasons, sequential e-values are potentially the perfect tool to define online closed procedures, which is why we will focus on them in the following.

Assume that for each hypothesis H_t an e-value E_t is available and the e-values $(E_t)_{t \in \mathbb{N}}$ are sequentially valid with respect to $(\mathcal{F}_t)_{t \in \mathbb{N}_0}$, meaning $E_t \in \mathcal{F}_t$ and $\mathbb{E}_{\mathbb{P}}[E_t | \mathcal{F}_{t-1}] \leq 1$ for all $\mathbb{P} \in H_t$. For each intersection hypothesis H_I , we construct a process $(W_I^t)_{t \in I \cup \{0\}}$ with $W_I^0 = 1$ and

$$W_I^t = \prod_{i \in I \cap \{1, \dots, t\}} E_i \quad (t \in I). \quad (9)$$

Following the above argumentation, $(W_I^t)_{t \in I \cup \{0\}}$ is a test supermartingale for H_I with respect to $(\mathcal{F}_t)_{t \in I \cup \{0\}}$. By Ville's inequality [23], it follows that

$$\phi_I = \mathbb{1}\{\exists t \in I : W_I^t \geq 1/\alpha\} = \mathbb{1}\left\{\sup_{t \in I} W_I^t \geq 1/\alpha\right\} \quad (10)$$

is an intersection test. Furthermore, ϕ_I is an online intersection test and $\phi = (\phi_I)_{I \subseteq \mathbb{N}}$ is increasing such that the closed procedure \mathbf{d}^ϕ (3) is indeed an online procedure (see Section 2.2). Note that due to the supremum involved in Ville's inequality, the intersection tests ϕ_I , $I \subseteq \mathbb{N}$, are not symmetric, and thus very different from the ones considered by Tian et al. [45].

Sequential e-values arise naturally in a variety of settings [41, 21, 38, 52]. For example, note that many online multiple testing tasks are performed in an adaptive manner. That means, the hypotheses to test, the design of a study or the strategy to calculate an e-value, depend on the data observed so far. This is very natural, since the hypotheses are tested over time and therefore the statistician performing these tests automatically learns about the context of the study and the true distribution of the data during the testing process. Indeed, to avoid such data-adaptive designs one would need to ignore all the previous data when making design decisions for the future testing process and therefore could be required to prespecify all hypotheses that are going to be tested, the sample size for these tests, decide which tests to apply and fix many further design parameters at the very beginning of the testing process. This would take away much of the flexibility of online multiple testing procedures. However, if one uses past information for the design of an e-value and calculates the e-value on the same data that design information is based on, one would need to have knowledge about the conditional distribution of the e-value given that design information. This is often not available. Therefore, one usually uses independent and fresh data for each of the e-values. In this way, no matter in which way an e-value depends on the past data, it remains valid conditional on it. Consequently, we obtain sequential e-values well suited to our online true discovery procedure.

Computational shortcuts. One problem with procedures based on the (online) closure principle is that at each step $t \in \mathbb{N}$ up to 2^{t-1} additional intersection tests must be considered. However, for specific intersection tests this number can be reduced drastically. In Algorithm 1 we introduce a shortcut for the intersection tests in (10) that only requires one calculation per individual hypothesis but provides the same bounds as the entire closed procedure would. We call this Algorithm **SeqE-Guard** (guarantee for true discoveries with sequential e-values).

The **SeqE-Guard** algorithm will provide a sequence of lower bounds d_t on the number of true discoveries in an adaptively chosen sequence of query sets $S_t \subseteq \{1, \dots, t\}$, simultaneously over all t . The S_t and d_t sequences are (nonstrictly) increasing with t since by default the statistician only decides at step t whether to include the index t within S_t or not (based on E_1, \dots, E_t), but typically does not omit hypotheses that were deemed interesting at an earlier point (see Remark below). So it is understood below that $S_t \supseteq S_{t-1}$ for all t . As the multiple testing process continues, the statistician can report/announce one or more (S_t, d_t) pairs that they deem interesting thus far, and our theorem below guarantees that all such announcements will be accurate with high probability.

A short description of SeqE-Guard. At each step $t \in \mathbb{N}$, the statistician decides based on the e-values E_1, \dots, E_t whether the index t should be included in the query set S_t . If it is to be included, **SeqE-Guard** calculates the product of all e-values in S_t (that were not already excluded) and all e-values smaller than 1 that are not in S_t . If that product is larger than or equal to $1/\alpha$, we increase the true discovery bound by 1 and *exclude* the current largest e-value from the future analysis. In Supplementary Material S.1 we provide a simple calculation example for the **SeqE-Guard** algorithm.

Theorem 3.1. *Let E_1, E_2, \dots be sequential e-values for H_1, H_2, \dots . The true discovery bounds d_t for S_t defined by **SeqE-Guard** (Algorithm 1) are the same as the ones obtained by the online*

closure principle with the intersection tests in (10). In particular, S_t and d_t satisfy for all $\mathbb{P} \in \mathcal{P}$: $\mathbb{P}(d_t \leq |S_t \cap I_1^{\mathbb{P}}| \text{ for all } t \in \mathbb{N}) \geq 1 - \alpha$.

Algorithm 1 SeqE-Guard: Online true discovery guarantee with sequential e-values

Input: Sequence of sequential e-values E_1, E_2, \dots

Output: Query sets $S_1 \subseteq S_2 \subseteq \dots$ and true discovery bounds $d_1 \leq d_2 \leq \dots$

```

1:  $d_0 = 0$ 
2:  $S_0 = \emptyset$ 
3:  $U = \emptyset$ 
4:  $A = \emptyset$ 
5: for  $t = 1, 2, \dots$  do
6:    $S_t = S_{t-1}$ 
7:    $d_t = d_{t-1}$ 
8:   Statistician observes  $E_t$  and chooses whether index  $t$  should be included in  $S_t$ .
9:   if  $t \in S_t$  then
10:     $A = A \cup \{t\}$ 
11:    if  $\prod_{i \in A \cup U} E_i \geq 1/\alpha$  then
12:       $d_t = d_{t-1} + 1$ 
13:       $A = A \setminus \{\text{index of largest e-value in } A\}$ 
14:    end if
15:  else if  $E_t < 1$  then
16:     $U = U \cup \{t\}$ 
17:  end if
18:  return  $S_t, d_t$ 
19: end for

```

Remark 2. Note that closed procedures provide simultaneous true discovery guarantee simultaneously over all $S \in 2^{\mathbb{N}^f}$, while *SeqE-Guard* only gives a simultaneous lower bound on the number of true discoveries for a path of query sets $(S_t)_{t \in \mathbb{N}}$ with $S_1 \subseteq S_2 \subseteq \dots$. However, at some step $t \in \mathbb{N}$ one could also obtain a lower bound for the number of true discoveries in any other set $S \subseteq \{1, \dots, t\}$ that is not on the path. We formulated *SeqE-Guard* for single query paths due to computational convenience and as we think this reflects the proceeding in many applications. This was also done in previous works on online true discovery guarantee [26, 32]. It should be noted that if we apply *SeqE-Guard* to multiple query paths, we must use the same e-values for every query path we consider.

Example 2 (Nontrivial true discovery bounds with weak signals). Let $\alpha = 0.05$ and suppose we test 1000 hypotheses using sequential e-values. Further suppose that 300 of those e-values equal 0.9 and 700 of those e-values equal 1.1. Due to the small e-values, it seems like there is no or at most very small evidence against the null hypotheses. Indeed, standard multiple testing procedures like the e-BH procedure for FDR control [59] would not make any rejections (even if applied at larger levels like $\alpha = 0.5$). However, plugging these e-values into the *SeqE-Guard* algorithm yields a bound of $d_{1000} \geq 306$ for the query set $S_{1000} = \{i \leq 1000 : E_i > 1\}$, since $0.9^{300} 1.1^{395} \geq 20$. Consequently, the *SeqE-Guard* algorithm can claim with a probability of 0.95 that there are at least 306 false hypotheses among the 700 hypotheses with e-value greater than 1. This demonstrates the capability of the *SeqE-Guard* algorithm to detect small but frequent signals in the data.

3.2 Improving existing methods

Existing online procedures with simultaneous true discovery guarantees were proposed by Katsevich and Ramdas [26] and Meah et al. [32]. Their setting involved observing one p-value for each hypothesis

(as is standard in multiple testing) but in this section, we will show that these methods can be uniformly improved by **SeqE-Guard** by employing specific choices of the sequential e-values.

We start with the **online-simple** procedure by Katsevich and Ramdas [26]. Suppose that p-values P_1, P_2, \dots for the hypotheses H_1, H_2, \dots are available such that P_i is measurable with respect to \mathcal{F}_i and let $\alpha_1, \alpha_2, \dots$ be nonnegative thresholds such that α_i is measurable with respect to \mathcal{F}_{i-1} . It is assumed that the null p-values are valid conditional on the past, meaning $\mathbb{P}(P_i \leq x | \mathcal{F}_{i-1}) \leq x$ for all $i \in I_0^{\mathbb{P}}, x \in [0, 1]$. Katsevich and Ramdas [26] showed that

$$d^{\text{os}}(S_t) = \left\lceil -ca + \sum_{i=1}^t \mathbb{1}\{P_i \leq \alpha_i\} - c\alpha_i \right\rceil \quad (t \in \mathbb{N}) \quad (11)$$

provides simultaneous true discovery guarantee over all sets $S_t = \{i \leq t : P_i \leq \alpha_i\}$, $t \in \mathbb{N}$, where $a > 0$ is some parameter and $c = \frac{\log(1/\alpha)}{a \log(1 + \log(1/\alpha)/a)}$. On closer examination of their proof, one can observe that they proved their guarantee by implicitly showing that

$$E_i^{\text{os}} = \exp[\theta_c(\mathbb{1}\{P_i \leq \alpha_i\} - c\alpha_i)] \quad (i \in \mathbb{N}) \quad (12)$$

define sequential e-values, where $\theta_c = \log(1/\alpha)/(ca)$. We now propose to simply plug these sequential e-values into **SeqE-Guard**; this leads to Algorithm 2, which we will refer to as **closed online-simple** procedure in the following. To see this, note that

$$\prod_{i \in I} E_i^{\text{os}} \geq 1/\alpha \Leftrightarrow \sum_{i \in I} \mathbb{1}\{P_i \leq \alpha_i\} - c\alpha_i \geq \log(1/\alpha)/\theta_c = ca \quad (I \subseteq \mathbb{N}).$$

The superscript of the set A_t^c in Algorithm 2 is used to reflect the fact that A_t^c contains all indices of queried hypotheses that were excluded from the analysis until step t , and therefore can be seen as complement of the set A in Algorithm 1 with respect to S_t .

Algorithm 2 Closed online-simple

Input: Sequence of p-values P_1, P_2, \dots and sequence of (potentially data-dependent) individual significance levels $\alpha_1, \alpha_2, \dots$.

Output: Query sets $S_1 \subseteq S_2 \subseteq \dots$ and true discovery bounds $d_1 \leq d_2 \leq \dots$.

```

1:  $d_0 = 0$ 
2:  $S_0 = \emptyset$ 
3:  $A_0^c = \emptyset$ 
4: for  $t = 1, 2, \dots$  do
5:   if  $P_t \leq \alpha_t$  then
6:      $S_t = S_{t-1} \cup \{t\}$ 
7:   else
8:      $S_t = S_{t-1}$ 
9:   end if
10:  if  $\sum_{i \in \{1, \dots, t\} \setminus A_{t-1}^c} \mathbb{1}\{P_i \leq \alpha_i\} - c\alpha_i \geq ca$  then
11:     $d_t = d_{t-1} + 1$ 
12:     $A_t^c = A_{t-1}^c \cup \{\text{index of smallest individual significance level in } S_t \setminus A_{t-1}^c\}$ 
13:  else
14:     $d_t = d_{t-1}$ 
15:     $A_t^c = A_{t-1}^c$ 
16:  end if
17:  return  $S_t, d_t$ 
18: end for

```

Let A_t^c and d_t be defined as in **closed online-simple**, then

$$\begin{aligned} d_t &\geq |A_{t-1}^c| + \left[-ca + \sum_{i \in \{1, \dots, t\} \setminus A_{t-1}^c} \mathbb{1}\{P_i \leq \alpha_i\} - c\alpha_i \right] \\ &= \left[-ca + \sum_{i=1}^t \mathbb{1}\{P_i \leq \alpha_i\} - \sum_{i \in \{1, \dots, t\} \setminus A_{t-1}^c} c\alpha_i \right] \geq \mathbf{d}^{\text{os}}(S_t), \end{aligned} \quad (13)$$

which shows that our **closed online-simple** method uniformly improves the **online-simple** procedure by Katsevich and Ramdas [26]. The improvement can be divided into two parts. First, the **closed online-simple** procedure is coherent, providing that $d_{t-1} \leq d_t$ for all $t \in \mathbb{N}$. Second, every time the bound d_t is increased by one, the summand $-c\alpha_i$ is excluded from the bound, where α_i is the smallest significance level with index in $S_t \setminus A_{t-1}^c$. This shows that the (online) closure principle and **SeqE-Guard** automatically adapt to the number of discoveries and thus the proportion of false hypotheses.

Furthermore, note that E_i^{os} only takes two values. It takes $\exp[\theta_c(1 - c\alpha_i)]$ if $P_i \leq \alpha_i$ and $\exp[-\theta_c c\alpha_i]$ if $P_i > \alpha_i$, where $\mathbb{P}(P_i \leq \alpha_i | \mathcal{F}_{i-1}) \leq \alpha_i$ for all $\mathbb{P} \in H_i$. Hence, we have for all $\mathbb{P} \in H_i$,

$$\mathbb{E}_{\mathbb{P}}[E_i^{\text{os}} | \mathcal{F}_{i-1}] \leq \alpha_i \exp[\theta_c(1 - c\alpha_i)] + (1 - \alpha_i) \exp[-\theta_c c\alpha_i] =: u_i. \quad (14)$$

For example, $\alpha_i = \alpha = 0.1$ and $a = 1$ yield $u_i = 0.977$ (u_i is increasing in a ; for $a = 3$ we obtain $u_i = 0.997$), which shows that E_i^{os} is not admissible and thus can be improved. A simple improvement can be obtained by plugging the e-value $\tilde{E}_i^{\text{os}} = E_i^{\text{os}}/u_i$ instead of E_i^{os} into **SeqE-Guard**. Doing this at every step, we obtain a further improvement; we call this the **admissible online-simple** method in the following.

Proposition 3.2. *By plugging in the sequential e-values $(E_i^{\text{os}})_{i \in \mathbb{N}}$ (12) into the **SeqE-Guard** algorithm, we obtain multiple uniform improvements (that can be applied together) over the **online-simple** method by Katsevich and Ramdas [26]:*

1. *The lower bound d_t is nondecreasing in t (**coherent online-simple**).*
2. *Every time the bound d_t is increased by one, the summand $-c\alpha_i$ is excluded from the bound, where α_i is the smallest threshold with index in $S_t \setminus A_{t-1}^c$ (**closed online-simple**).*

Furthermore, since the expected value of E_i^{os} is strictly smaller than 1 under H_0 , an additional improvement can be obtained:

3. *The sequential e-values E_i^{os} , $i \in \mathbb{N}$, can be replaced by $\tilde{E}_i^{\text{os}} = E_i^{\text{os}}/u_i$, where $u_i < 1$ is given by (14) (**admissible online-simple**).*

Katsevich and Ramdas [26] introduced one further online procedure with simultaneous true discovery guarantee, the **online-adaptive** method. A uniform improvement can be obtained in the exact same manner as for the **online-simple** method above; we present this in Supplementary Material S.2.1. Inspired by the methods of Katsevich and Ramdas [26], two further online procedures with simultaneous true discovery guarantee were proposed by Meah et al. [32]. The first procedure is obtained by taking a union of **online-simple** bounds for different parameters a . The second procedure exploits Freedman's inequality [13] and a union bound. In Supplementary Material S.2.2, we show how both these recent methods can be uniformly improved by our e-value based approach. The proposed improvements are technically not instances of the **SeqE-Guard** algorithm, but can be obtained by the union of **SeqE-Guard** bounds (or using the average of multiple test martingales). See Table 2 for the concrete sequential e-values yielding the mentioned improvements.

In the following section we consider the **SeqE-Guard** algorithm with GRO e-values.

3.3 Adaptively hedged GRO e-values

The most common strategy to calculate e-values in practice is based on variants of the GRO-criterion [41, 21, 61], which dates back to the Kelly criterion [27, 5]. Here, we assume that each null hypothesis H_i comes with an alternative H_i^A . Suppose H_i^A contains a single distribution \mathbb{Q}_i . Then the growth rate optimal (GRO) e-value E_i^{GRO} is defined as the e-value E_i that maximizes the growth rate under \mathbb{Q}_i , given by $\mathbb{E}_{\mathbb{Q}_i}[\log(E_i)]$, over all e-values for H_i . If H_i^A is composite, one can define a prior over the distributions contained in H_i^A and then calculate the GRO e-value according to the mixture distribution based on that prior. If the null hypothesis is simple, the GRO e-value is given by the likelihood ratio (LR) of the alternative over the null distribution [41]. If H_i is composite, the GRO e-value takes the form of a LR of the alternative against a specific (sub-) distribution [21, 29]. The GRO e-value is particularly powerful when many e-values are combined by multiplication [41] and therefore seems to be a reasonable choice for our **SeqE-Guard** algorithm. Furthermore, since the growth rate is the standard measure of performance for e-values [38], there might be applications where do not have access to the data to calculate our own e-value, but just get a GRO e-value for each individual hypothesis. For example, this could be the case in meta-analyses, where each study just reported the GRO e-value. In the following, we will discuss how the GRO concept transfers to online true discovery guarantee based on sequential e-values.

A naive approach would be to just plugin the GRO e-values into **SeqE-Guard**. However, the problem with this is that the GRO e-values only maximize the growth rate under their alternatives. This makes sense when testing a single hypothesis. In our setting, the product of GRO e-values would only maximize the growth rate of $W_t = \prod_{i=1}^t E_i$ if all hypotheses are false — a very unlikely scenario. Indeed, GRO e-values can be small if the null hypothesis is true, so directly using GRO e-values in **SeqE-Guard** can lead to low power. Hence, when considering products of multiple e-values for different hypotheses one needs to incorporate the possibility that a hypothesis is true. One approach is to hedge the GRO e-values by defining

$$\tilde{E}_i^{\text{GRO}} = 1 - \lambda_i + \lambda_i E_i^{\text{GRO}},$$

where $\lambda_i \in [0, 1]$ is measurable with respect to \mathcal{F}_{i-1} . To see that $\tilde{E}_1^{\text{GRO}}, \tilde{E}_2^{\text{GRO}}, \dots$ indeed define sequential e-values if the GRO e-values are sequential, just note that $\mathbb{E}_P[\tilde{E}_i^{\text{GRO}} | \mathcal{F}_{i-1}] = 1 - \lambda_i + \lambda_i \mathbb{E}_P[E_i^{\text{GRO}} | \mathcal{F}_{i-1}] \leq 1$ for all $\mathbb{P} \in H_i$.

In order to derive a reasonable choice for λ_i , we consider a specific Bayesian two-groups model. Suppose $H_i = \{\mathbb{P}_i\}$ and $H_i^A = \{\mathbb{Q}_i\}$ are both simple hypotheses and which of these hypotheses is true is random, where τ_i gives the probability that the alternative hypothesis H_i^A is true. A reasonable approach would be to choose the e-value E_i for H_i that maximizes the growth rate under the true distribution $\mathbb{E}_{(1-\tau_i)\mathbb{P}_i + \tau_i\mathbb{Q}_i}[\log(E_i)]$.

Proposition 3.3. *Suppose \mathbb{Q}_i is absolutely continuous with respect to \mathbb{P}_i . Then the e-value E_i that maximizes the growth rate under the true distribution $(1 - \tau_i)\mathbb{P}_i + \tau_i\mathbb{Q}_i$ is given by $1 - \tau_i + \tau_i E_i^{\text{GRO}}$, where E_i^{GRO} maximizes the growth rate under the alternative \mathbb{Q}_i .*

Proposition 3.3 shows that in the case of simple null and alternative hypotheses the optimal e-value under the true distribution is the same as if we maximize the growth rate under the alternative and then hedge the resulting e-value according to the probability that the alternative is true. We think it is possible that something analogous holds for composite null hypotheses as well. However, even if this is not the case, it seems to be a reasonable strategy in general.

Therefore, our approach is to specify an estimate $\hat{\tau}_i$ for the probability τ_i that H_i is false, where $\hat{\tau}_i$ can either depend on prior information or the past data, and then set $\lambda_i = \hat{\tau}_i$. We propose to set

$$\hat{\tau}_i = \frac{1/2 + \sum_{j=1}^{i-1} \mathbb{1}\{E_j^{\text{GRO}} > 1\}}{i} \quad (i \in \mathbb{N}), \quad (15)$$




Figure 3: Comparing $\hat{\tau}_i$ (solid line) defined in (15) to the true τ_i (dashed line) for different scenarios. The simulation design is described in Section 4. The estimate $\hat{\tau}_i$ has a tendency to bias towards 50% but it estimates the proportion of false hypotheses very well when the signal is strong.

if no prior information is available. The reasoning for this choice is that the GRO e-value can be interpreted as a (generalized) LR and if $E_j^{\text{GRO}} > 1$ this means that the data prefers the alternative over the null distribution.

We illustrate the behavior of $\hat{\tau}_i$ defined in (15) for different scenarios in Figure 3. We consider the simulation setup described in Section 4, but with $n = 100$ and proportion of false hypotheses $\pi_A \in \{0.2, 0.5, 0.8\}$. If the signal is weak, $\hat{\tau}_i$ lies between the true proportion and 0.5 and if the signal is strong, the estimate $\hat{\tau}_i$ is close to the true proportion. We think that this behavior is desirable, since a tendency towards 0.5 is not that hurtful if the alternative and null distribution are close as the GRO e-values have a small variance in this case.

Hedging the e-values before multiplying them does not only apply to GRO e-values. Such strategies for merging sequential e-values have been used by many preceding authors like Waudby-Smith and Ramdas [61], Vovk and Wang [52]. However, the above argumentation provides a reasonable choice for the parameter λ_i in our setting.

We analyze this approach experimentally in Section 4.2.

4 Simulations

In this section we numerically calculate the true discovery proportion (TDP) bound, which is defined as the true discovery bound for S_t divided by the size of S_t . We compare TDP bounds obtained by applying SeqE-Guard to the different sequential e-values proposed in the previous sections. In Subsection 4.1, we compare the **online-simple** method by Katsevich and Ramdas [26] with its uniform improvement. In Subsection 4.2 we demonstrate how hedging GRO e-values improve the true discovery bound. Finally, in Subsection 4.3, we compare the proposed e-values to decide which is best suited for practice.

We consider the same simulation setup in all subsections. We sequentially test $n = 1000$ null hypotheses H_i , $i \in \{1, \dots, n\}$, of the form $H_i : X_i \sim \mathcal{N}(0, 1)$ against the alternative $H_i^A : X_i \sim$

$\mathcal{N}(\mu_A, 1)$ for some $\mu_A > 0$, where X_1, \dots, X_n are independent data points or test statistics. The probability that the alternative hypothesis is true is set by a parameter $\pi_A \in (0, 1)$ and the desired guarantee is set to $\alpha = 0.1$. For all comparisons we consider a grid of simulation parameters $\mu_A \in \{2, 3, 4\}$ and $\pi_A \in \{0.1, 0.3, 0.5\}$, where we refer to $\mu_A = 2$ as weak signal, $\mu_A = 3$ as medium signal and $\mu_A = 4$ as strong signal. The p-values are calculated by $\Phi(-X_i)$, where Φ is the CDF of a normal distribution. The raw GRO e-values are given by the likelihood ratio $E_i^{\text{GRO}} = p_{\mu_A}(X_i)/p_0(X_i)$, where p_{μ_A} and p_0 are the densities of a normal distribution with variance 1 and mean μ_A and 0, respectively. The query sets S_t , $t \in \{1, \dots, t\}$, are defined as $S_t = \{i \in \{1, \dots, t\} : P_i \leq \alpha\}$. All of the results in the following are obtained by averaging over 1000 independent trials.

4.1 Comparing the online-simple method [26] with its improvements

In Section 3.2 we showed that the **online-simple** method by Katsevich and Ramdas [26] can be uniformly improved by the **closed online-simple** procedure (Algorithm 2). We also showed that this closed procedure can be further uniformly improved by the **admissible online-simple** procedure, which ensures that the expected value of each sequential e-value is exactly one. In this section, we aim to quantify the gain in power for making true discoveries by using these improvements instead of the **online-simple** method. Although Katsevich and Ramdas [26] proposed $a = 1$ as default parameter, we found $a = 3$ to perform better which is why we use it here.

The results are illustrated in Figure 4. It can be seen that the **closed online-simple** procedure leads to a substantial improvement of the **online-simple** procedure in all cases. Of the queried hypotheses, the former approximately identifies 10%–20% more as false. The additional improvement of the admissible online-simple method is quite small in this case, however, it is potentially larger for smaller parameters a (see Section 3.2).

In Section S.2.2 we also provide simulations for our improvements over the methods proposed by Meah et al. [32]. The behavior of both, the initially proposed methods and our improvements, is quite similar as in the **online-simple** case.

4.2 GRO e-values

In Section 3.3 we argued that the raw GRO e-values should be hedged to account for the probability that a null hypothesis is true. Now we compare the true discovery bounds obtained by applying **SeqE-Guard** to raw and hedged GRO e-values. For the hedged GRO e-values we chose the predictable parameter proposed in (15). The results are illustrated in Figure 5, showing that GRO e-values lead to very low bounds which can be increased substantially by hedging.

4.3 Which sequential e-values should we choose?

In this section, we compare the **SeqE-Guard** procedure when applied with the best versions of the proposed sequential e-values to derive recommendations for practice. More precisely, we compare the **admissible online-simple** method with the hedged GRO e-value.

The results are depicted in Figure 6. The procedures perform quite different in the various settings. When the proportion of false hypotheses is small, the **SeqE-Guard** algorithm performs best with hedged GRO e-values, particularly, if the signal is strong. However, if the proportion of false hypotheses is large, the **admissible online-simple** method clearly outperforms the **SeqE-Guard** with hedged GRO e-values. Hence, if we expect sparse but strong signal, applying **SeqE-Guard** with hedged GRO e-values is the best choice. In contrast, for dense but weak signals, the **admissible online-simple** method should be preferred.

Furthermore, it should be noted that the only reasonable query path for the **online-simple** method and its improvements is given by $S_t = \{i \leq t : P_i \leq \alpha_i\}$ (which we use in the simulations for $\alpha_i = \alpha$). To see this, note that all e-values E_i^{os} , $i \leq t$, with $i \notin S_t$ are smaller than 1 and




Figure 4: True discovery proportion bounds obtained by the `online-simple` method [26], the `closed online-simple` method and the `admissible online-simple` method. The `closed online-simple` method leads to substantially larger bounds than the original `online-simple` method. An additional improvement can be obtained by the `admissible online-simple` method, which is closest to the true proportion (top line) in all figures.

therefore their inclusion in S_t would not increase the lower bound for the number of true discoveries. Furthermore, excluding e-values E_i^{os} , $i \leq t$, with $P_i \leq \alpha_i$ would be difficult to communicate, as those e-values reached their maximum possible value. Hence, hedged GRO e-values are better suited for an exploratory analysis where the scientist might be interested in several different query paths (see Remark 2). For example, `SeqE-Guard` with hedged GRO e-values can provide a (nontrivial) query path with online FWER control (by including $t \in S_t$ iff this implies $d_t = d_{t-1} + 1$), while simultaneously providing (nontrivial) real-time lower bounds for the number of false hypotheses among all hypotheses with e-values greater than, say, 2. For such an exploratory proceeding, we would recommend the hedged GRO e-values, since they showed good performance (Figure 5) without adapting to the query path at all. In Supplementary Material S.4, we show how GRO e-values could be adapted to the chosen query path using a boosting approach [59] that uniformly improves the bounds for that specific query path (but may reduce the bounds for other query paths). In Supplementary Material S.3 we show how one could investigate multiple query paths when only p-values are available, using other p-to-e calibrators than the simple binary calibrator applied for the `online-simple` method.

5 Online true discovery guarantee with exchangeable e-values

In the previous sections we considered sequential e-values which naturally arise if the data used for the different e-values is independent, but the hypotheses to test, the study design or the strategy to calculate the e-value depend on the previous (independent) data. However, there are situations where the data, or at least some part of the data, must be reused for the different hypotheses. One such example is online outlier detection with conformal e-values.

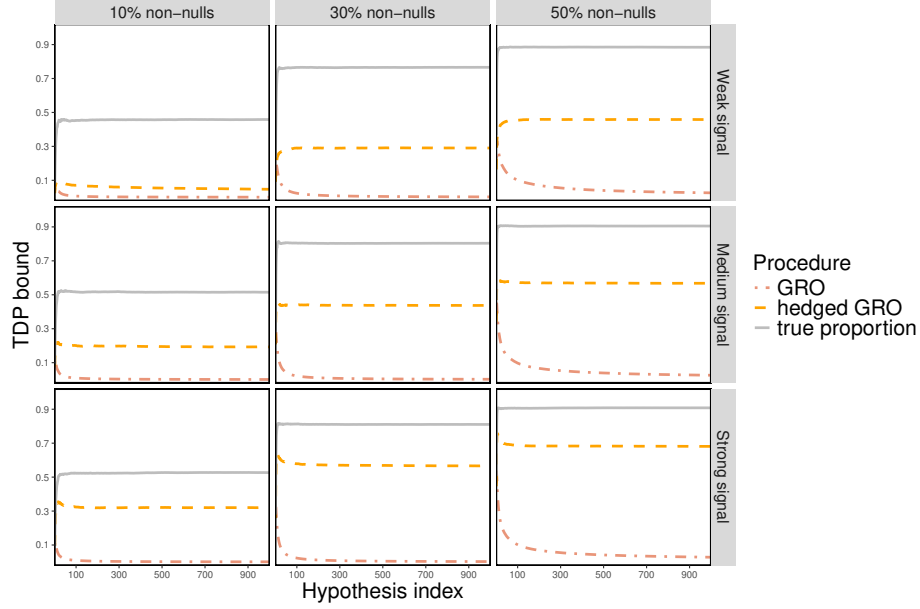


Figure 5: True discovery proportion bounds obtained by applying **SeqE-Guard** to GRO e-values and hedged GRO e-values. The hedged GRO e-values improve the GRO e-values substantially.

Example 3 (Online outlier detection with soft-rank e-values). Consider an outlier detection problem, where a calibration data set $\mathcal{D} = \{X'_t\}_{t=1}^n$ is given that contains n i.i.d. data points $X'_t \in \mathbb{R}^d$ drawn from an unknown distribution \mathbb{P}_X [2]. Furthermore, suppose we have a possibly infinite sequence X_1, X_2, \dots of independent test data points coming in over time and for which we want to test whether they are drawn from \mathbb{P}_X as well, yielding the null hypothesis $H_t : X_t \sim \mathbb{P}_X$. Points that are drawn from \mathbb{P}_X are called inliers and points that are not drawn from \mathbb{P}_X are called outliers. A modern approach for this problem is based on conformal prediction [2, 56, 54]. For this, suppose a score function $\hat{s} : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ is given, which maps a data point X_t to a nonnegative scalar $\hat{s}(X_t)$ such that a larger score $\hat{s}(X_t)$ indicates that it is more likely that X_t is an outlier. The score function is usually determined based on some training data that is independent of the calibration and test data. The existing approach calculates a rank-based p-value P_t^{perm} equaling the rank of the score of X_t amongst the scores of calibration data set (and X_t). However, we consider the “soft-rank permutation e-value” introduced in an early preprint of [59], which is defined as

$$E_t^{\text{soft}} = (n+1) \frac{\hat{s}(X_t)}{\hat{s}(X_t) + \sum_{i=1}^n \hat{s}(X'_i)}.$$

To see that this yields a valid e-value just note that X'_1, \dots, X'_n, X_t are exchangeable under H_t , which means that the joint distribution of these random variables remains the same under every permutation, and therefore $\mathbb{E}_{\mathbb{P}} [\hat{s}(X_t) | \hat{s}(X_t) + \sum_{i=1}^n \hat{s}(X'_i)] = (\hat{s}(X_t) + \sum_{i=1}^n \hat{s}(X'_i)) / (n+1)$ for all $\mathbb{P} \in H_t$. Another way to obtain a valid e-value would be to calibrate P_t^{perm} into an e-value (see [10] for admissible calibrators of the permutation p-value). Note that no matter which of these two paths is chosen, the e-values $E_1^{\text{soft}}, E_2^{\text{soft}}, \dots$ are dependent, since they are all based on the same calibration data set. However, also note that the sequence of e-values corresponding to true hypotheses $(E_t^{\text{soft}})_{t \in I_0^{\mathbb{P}}}$ is exchangeable, since $X_t \sim \mathbb{P}_X$ for all $t \in I_0^{\mathbb{P}}$.

Motivated by the online outlier detection problem, we construct an online procedure with true discovery guarantee for the setting where the e-values corresponding to the true hypotheses $(E_t)_{t \in I_0^{\mathbb{P}}}$

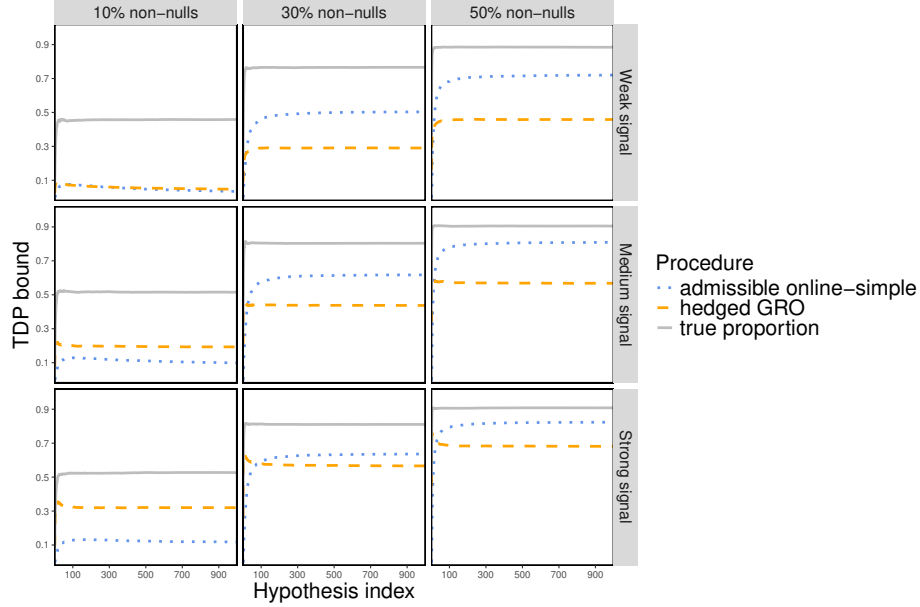


Figure 6: True discovery proportion bounds obtained by the **admissible online-simple** method and applying **SeqE-Guard** to hedged GRO e-values. If the proportion of false hypotheses is small, the hedged GRO e-values perform well, particularly if the signal is strong. If the proportion of false hypotheses is large, the **admissible online-simple** method should be preferred. However, the **admissible online-simple** method adapts to the specific query path, while the hedged GRO e-values are better suited for exploring multiple query paths.

are exchangeable for all $\mathbb{P} \in \mathcal{P}$. Note that this condition neither implies the sequential e-value condition of Section 3 nor is it implied by it. A useful property of exchangeable random variables is that their average forms a reverse martingale [25]. Ramdas and Manole [34] used this to prove the following Ville’s inequality for exchangeable random variables Y_1, Y_2, \dots :

$$\mathbb{P} \left(\exists t \geq 1 : \frac{1}{t} \sum_{i=1}^t |Y_i| \geq 1/\alpha \right) \leq \alpha \mathbb{E}[|Y_1|].$$

Hence, given a sequence of e-values E_1, E_2, \dots for the hypotheses H_1, H_2, \dots where the e-values corresponding to true hypotheses are exchangeable, we can define an intersection test for H_I as

$$\phi_I = \mathbb{1} \left\{ \exists t \in I : \frac{1}{|\{i \in I : i \leq t\}|} \sum_{i \in I, i \leq t} E_i \geq 1/\alpha \right\}. \quad (16)$$

It immediately follows that $\phi = (\phi_I)_{I \subseteq \mathbb{N}}$ defines an increasing family of online intersection tests. Hence, the resulting closed procedure (3) defines an online procedure with true discovery guarantee. We provide an exact short-cut for this closed procedure in Algorithm 3 (we call this Algorithm **ExE-Guard**). The short-cut is very similar to the one for sequential e-values (**SeqE-Guard**). The only difference is that we consider the average instead of the product and therefore have to take into account all e-values smaller than $1/\alpha$ instead of 1 from the e-values that are not contained in the query set. Since the product is usually larger than the average, we expect **SeqE-Guard** to be more powerful than **ExE-Guard**. Hence, if both conditions are met, we should use the former. For example, this can be the case if the e-values are independent and identically distributed under the null hypothesis.

Theorem 5.1. *Let E_1, E_2, \dots be e-values for H_1, H_2, \dots such that the e-values corresponding to true hypotheses are exchangeable. The true discovery bounds d_t for S_t defined by **ExE-Guard** (Algorithm 3) are the same as the ones obtained by the online closure principle with the intersection tests defined in (16). In particular, S_t and d_t satisfy for all $\mathbb{P} \in \mathcal{P}$: $\mathbb{P}(d_t \leq |S_t \cap I_1^{\mathbb{P}}| \text{ for all } t \in \mathbb{N}) \geq 1 - \alpha$.*

Proof. The proof is exactly the same as for Theorem 3.1. We just replaced the product with the average and for this reason defined U as the set of all e-values not in the query set that are smaller than $1/\alpha$. \square

Algorithm 3 ExE-Guard: Online true discovery guarantee with exchangeable null e-values

Input: Sequence of e-values E_1, E_2, \dots where the null e-values are exchangeable.

Output: Query sets $S_1 \subseteq S_2 \subseteq \dots$ and true discovery bounds $d_1 \leq d_2 \leq \dots$.

```

1:  $d_0 = 0$ 
2:  $S_0 = \emptyset$ 
3:  $U = \emptyset$ 
4:  $A = \emptyset$ 
5: for  $t = 1, 2, \dots$  do
6:    $S_t = S_{t-1}$ 
7:    $d_t = d_{t-1}$ 
8:   Observe  $E_t$  and decide whether  $t$  should be included in  $S_t$ .
9:   if  $t \in S_t$  then
10:     $A = A \cup \{t\}$ 
11:    Define  $\bar{E}_t$  as the average of e-values with indices in  $A \cup U$ .
12:    if  $\bar{E}_t \geq 1/\alpha$  then
13:       $d_t = d_{t-1} + 1$ 
14:       $A = A \setminus \{\text{index of largest e-value in } A\}$ 
15:    end if
16:  else if  $E_t < 1/\alpha$  then
17:     $U = U \cup \{t\}$ 
18:  end if
19:  return  $S_t, d_t$ 
20: end for

```

Remark 3. Note that if $E_t \geq t/\alpha$, the bound d_t will be increased by one and the e-value E_t excluded from the future analysis. Hence, we could truncate E_t at the value $\frac{1}{\alpha t}$ and exploit this for boosting E_t similarly as described for sequential e-values in Section S.4. However, note that we are not allowed to use any information about the previous e-values to increase the boosting factor due to the dependency between the e-values. Also note that t/α is usually much larger than the cutoff m_t (S.10) defined for sequential e-values, which is why the boosting factors will be smaller.

Remark 4. One can convert any finite sequence of e-values into a sequence of exchangeable e-values by randomly permuting them [34]. This was used by Gasparin et al. [15] to improve combination rules for p-values. One could exploit this in offline multiple testing by randomly permuting (possibly arbitrarily dependent) e-values before plugging them into **ExE-Guard**. Note that this is not possible in the online setting as the order of the hypotheses (and thus the order of the e-values) is fixed.

6 Online true discovery guarantee under arbitrary dependence

In the Sections 3 and 5 we introduced two martingale based true discovery procedures that work under different assumptions about the joint distribution of the e-values. While we believe that

these are very common situations for many applications, there are also cases where no information about the dependence structure is available. This typically occurs when data is reused for multiple hypotheses. For example, this is the case when open data repositories are evaluated in different studies [33, 6] or in Kaggle competitions [4]. In addition, arbitrarily dependent test statistics occur when the data for multiple hypotheses is collected from the same subpopulation, as it can be the case in online A/B testing [28]. For this reason, in this section we propose a method for online true discovery guarantee that works under arbitrary dependence of the e-values. We just assume that each e-value E_i for H_i is measurable with respect to \mathcal{F}_i , which is needed to define an online procedure based on these e-values.

While (hedged) multiplication is the admissible way of combining sequential e-values [52], the average is the only admissible merging function for arbitrarily dependent e-values [58, 49]. Therefore, it seems evident to look at averages when constructing intersection tests for arbitrarily dependent e-values. However, since we consider online true discovery guarantee, it is not possible to use the symmetric average.

Example 4. Consider the intersection hypotheses $H_{\{1\}}$ and $H_{\{1,2\}}$ and suppose we test each of these using the (unweighted) average of two e-values E_1 and E_2 . That means, $\phi_{\{1\}} = 1$ if $E_1 \geq 1/\alpha$ and $\phi_{\{1,2\}} = 1$ if $(E_1 + E_2)/2 \geq 1/\alpha$. Suppose $1/\alpha \leq E_1 < 2/\alpha - E_2$. Then $\phi_{\{1\}} = 1$ but $\phi_{\{1,2\}} = 0$, implying that $\phi = (\phi_I)_{I \subseteq \mathbb{N}}$ is not increasing and the resulting closed procedure is not an online procedure. Also note that it is not possible to apply the unweighted average for an infinite number of hypotheses.

For this reason, we will focus on weighted averages. To this regard, prespecify a nonnegative sequence $(\gamma_i)_{i \in \mathbb{N}}$ with $\sum_{i \in \mathbb{N}} \gamma_i = 1$. Then for each $I \subseteq \mathbb{N}$, define the intersection test

$$\phi_I = \mathbb{1} \left\{ \sum_{i \in I} E_i \gamma_{t(i;I)} \geq 1/\alpha \right\}, \quad (17)$$

where $t(i; I) = |\{j \in I : j \leq i\}|$. For every $I \subseteq \mathbb{N}$, we have

$$\mathbb{E}_{\mathbb{P}} \left[\sum_{i \in I} E_i \gamma_{t(i;I)} \right] = \sum_{i \in I} \gamma_{t(i;I)} \mathbb{E}_{\mathbb{P}} [E_i] \leq \sum_{i \in I} \gamma_{t(i;I)} \leq 1 \quad (\mathbb{P} \in H_I),$$

implying that $\sum_{i \in I} E_i \gamma_{t(i;I)}$ is a valid e-value for H_I . By Markov's inequality, ϕ_I in (17) defines an intersection test. Furthermore, if $I \subseteq \{1, \dots, i\}$ and $K = I \cup J$ for some $J \subseteq \{j \in \mathbb{N} : j > i\}$, then $t(l; I) = t(l; K)$ for all $l \in I$, which implies that $\phi = (\phi_I)_{I \subseteq \mathbb{N}}$ is increasing. Since each ϕ_I is additionally an online intersection test, the closed procedure based on the intersection tests in (17) is an online true discovery procedure under arbitrary dependence of the e-values.

One problem with such weighted methods is that it is difficult to find a short-cut for the corresponding closed procedure in general. In Algorithm 4 we provide a conservative short-cut, if $(\gamma_i)_{i \in \mathbb{N}}$ is nonincreasing (we call the Algorithm **ArbE-Guard**). That means, the true discovery bounds provided by the short-cut are always smaller than or equal to the bounds of the exact closed procedure and therefore it also provides simultaneous true discovery guarantee. The idea of the short-cut is to make the conservative assumption that all e-values that are not contained in the query set equal 0. Hence, if we choose loose criteria for including an index in the query set, the short-cut will be close to the full closed procedure based on (17). For example, if we just want a lower bound for the number of false hypotheses among all hypotheses under consideration, the short-cut is exact. However, if we choose strict criteria for including indices in the query set, e.g., if we desire FWER control, then the short-cut is conservative. The proof of the following theorem can be found in Supplementary Material S.5.

Theorem 6.1. Let E_1, E_2, \dots be arbitrarily dependent e-values for H_1, H_2, \dots . The true discovery bounds d_t for S_t defined by **ArbE-Guard** (Algorithm 4) satisfy $d_t \leq \mathbf{d}^\phi(S_t)$ for all $t \in \mathbb{N}$, where \mathbf{d}^ϕ is

the online closed procedure based on the intersection tests defined in (17). In particular, S_t and d_t satisfy for all $\mathbb{P} \in \mathcal{P}$: $\mathbb{P}(d_t \leq |S_t \cap I_1^{\mathbb{P}}| \text{ for all } t \in \mathbb{N}) \geq 1 - \alpha$.

Algorithm 4 ArbE-Guard: Online true discovery guarantee with arbitrarily dependent e-values

Input: Nonnegative and nonincreasing sequence $(\gamma_i)_{i \in \mathbb{N}}$ with $\sum_{i \in \mathbb{N}} \gamma_i \leq 1$ and sequence of (possibly) arbitrarily dependent e-values E_1, E_2, \dots .

Output: Query sets $S_1 \subseteq S_2 \subseteq \dots$ and true discovery bounds $d_1 \leq d_2 \leq \dots$.

```

1:  $d_0 = 0$ 
2:  $S_0 = \emptyset$ 
3:  $A^c = \emptyset$ 
4: for  $t = 1, 2, \dots$  do
5:    $S_t = S_{t-1}$ 
6:    $d_t = d_{t-1}$ 
7:   Observe  $E_t$  and decide whether  $t$  should be included in  $S_t$ .
8:   if  $t \in S_t$  then
9:     if  $\sum_{i \in S_t \setminus A^c} E_i \gamma_{t(i; \{1, \dots, t\} \setminus A^c)} \geq 1/\alpha$  then
10:       $d_t = d_{t-1} + 1$ 
11:       $A^c = A^c \cup \left\{ \operatorname{argmin}_{i \in S_t \setminus A^c} \sum_{j \in S_t \setminus [A^c \cup \{i\}]} E_j \gamma_{t(j; \{1, \dots, t\} \setminus [A^c \cup \{i\}])} \right\}$ 
12:    end if
13:  end if
14:  return  $S_t, d_t$ 
15: end for

```

In the following proposition we show that for nonincreasing $(\gamma_i)_{i \in \mathbb{N}}$ the intersection tests defined in (17) are uniformly improved by the ones in (16) (see Supplementary Material S.5 for the proof). Hence, we should always prefer the **ExE-Guard** over the **ArbE-Guard** if the null e-values are exchangeable.

Proposition 6.2. *Let $(\gamma_i)_{i \in \mathbb{N}}$ be a nonincreasing sequence with $\sum_{i \in \mathbb{N}} \gamma_i \leq 1$. Furthermore, let $\phi_I^{\text{AE}}, I \subseteq \mathbb{N}$, be the intersection test defined in (17) and ϕ_I^{ExE} be the intersection test in (16). Then $\phi_I^{\text{ExE}} \geq \phi_I^{\text{AE}}$ for all $I \in 2^{\mathbb{N}_f}$. Hence, **ExE-Guard** uniformly improves **ArbE-Guard** if the e-values corresponding to true hypotheses are exchangeable.*

Remark 5. *In the same manner as described in Remark 3 for exchangeable e-values, one could boost each e-value $E_t, t \in \mathbb{N}$, before plugging it into **ArbE-Guard** by truncating it at $\frac{1}{\alpha \gamma_t}$.*

7 Discussion

In this paper, we proposed a new closed testing based online true discovery procedure for sequential e-values and derived a general short-cut that only requires one calculation per hypothesis. Although the **SeqE-Guard** algorithm is restricted to sequential e-values, it is a general procedure for the task of online true discovery guarantee, since there are many different ways to construct sequential e-values. In particular, it yields uniform improvements of the state-of-the-art methods by Katsevich and Ramdas [26] and Meah et al. [32], although they were not explicitly constructed using e-values.

From a theoretical point of view this paper gives new insights about the role of e-values in multiple testing by showing that every admissible coherent online true discovery procedure must be based on sequential e-values. From a practical point of view, we constructed a powerful and flexible multiple testing procedure, which allows to observe hypotheses one-by-one over time and make fully data-adaptive decisions about the hypotheses and stopping time while bounding the number of true discoveries or equivalently, the false discovery proportion. On the way, we introduced a new idea for hedging of GRO e-values in multiple testing.

Although not being the main focus of this paper, we also introduced new methods with online true discovery guarantee for the setting of exchangeable and arbitrarily dependent e-values. To the best of our knowledge, these are the first approaches in these settings and therefore provide the new benchmark for future work.

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. The authors thank Etienne Roquain for a helpful comment.

References

- [1] 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.
- [2] Stephen Bates, Emmanuel Candès, Lihua Lei, Yaniv Romano, and Matteo Sesia. Testing for outliers with conformal p-values. *The Annals of Statistics*, 51(1):149–178, 2023.
- [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] Casper Solheim Bojer and Jens Peder Meldgaard. Kaggle forecasting competitions: An overlooked learning opportunity. *International Journal of Forecasting*, 37(2):587–603, 2021.
- [5] Leo Breiman. Optimal gambling systems for favourable games. In *Fourth Berkeley Symposium on Mathematical Statistics and Probability*, pages 65–78, 1961.
- [6] 1000 Genomes Project Consortium. A global reference for human genetic variation. *Nature*, 526(7571):68, 2015.
- [7] Donald A Darling and Herbert Robbins. Confidence sequences for mean, variance, and median. *Proceedings of the National Academy of Sciences*, 58(1):66–68, 1967.
- [8] Boyan Duan, Aaditya Ramdas, Sivaraman Balakrishnan, and Larry Wasserman. Interactive martingale tests for the global null. *Electronic Journal of Statistics*, 14(2):4489–4551, 2020.
- [9] Jean Feng, Scott Emerson, and Noah Simon. Approval policies for modifications to machine learning-based software as a medical device: A study of bio-creep. *Biometrics*, 77(1):31–44, 2021.
- [10] Lasse Fischer and Aaditya Ramdas. Sequential Monte-Carlo testing by betting. *arXiv preprint arXiv:2401.07365*, 2024.
- [11] Lasse Fischer, Marta Bofill Roig, and Werner Brannath. The online closure principle. *The Annals of Statistics*, 52(2):817–841, 2024.
- [12] 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.
- [13] David A Freedman. On tail probabilities for martingales. *The Annals of Probability*, pages 100–118, 1975.

- [14] K Ruben Gabriel. Simultaneous test procedures—some theory of multiple comparisons. *The Annals of Mathematical Statistics*, 40(1):224–250, 1969.
- [15] Matteo Gasparin, Ruodu Wang, and Aaditya Ramdas. Combining exchangeable p-values. *arXiv preprint arXiv:2404.03484*, 2024.
- [16] Christopher Genovese and Larry Wasserman. A stochastic process approach to false discovery control. *The Annals of Statistics*, 32(3):1035–1061, 2004.
- [17] Christopher R Genovese and Larry Wasserman. Exceedance control of the false discovery proportion. *Journal of the American Statistical Association*, 101(476):1408–1417, 2006.
- [18] Jelle J Goeman and Aldo Solari. Multiple testing for exploratory research. *Statistical Science*, 26(4):584–597, 2011.
- [19] Jelle J Goeman, Rosa J Meijer, Thijmen JP Krebs, and Aldo Solari. Simultaneous control of all false discovery proportions in large-scale multiple hypothesis testing. *Biometrika*, 106(4):841–856, 2019.
- [20] Jelle J Goeman, Jesse Hemerik, and Aldo Solari. Only closed testing procedures are admissible for controlling false discovery proportions. *The Annals of Statistics*, 49(2):1218–1238, 2021.
- [21] Peter Grünwald, Rianne de Heide, and Wouter M Koolen. Safe testing. *Journal of the Royal Statistical Society Series B: Statistical Methodology (with discussion)*, 2024.
- [22] Jesse Hemerik and Jelle J Goeman. False discovery proportion estimation by permutations: confidence for significance analysis of microarrays. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 80(1):137–155, 2018.
- [23] Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform chernoff bounds via nonnegative supermartingales. *Probability Surveys*, 17:257–317, 2020.
- [24] 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.
- [25] Olav Kallenberg. *Probabilistic symmetries and invariance principles*, volume 9. Springer, 2005.
- [26] Eugene Katsevich and Aaditya Ramdas. Simultaneous high-probability bounds on the false discovery proportion in structured, regression and online settings. *The Annals of Statistics*, 48(6):3465–3487, 2020.
- [27] John L Kelly. A new interpretation of information rate. *The Bell System Technical Journal*, 35(4):917–926, 1956.
- [28] Ron Kohavi, Alex Deng, Brian Frasca, Toby Walker, Ya Xu, and Nils Pohlmann. Online controlled experiments at large scale. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1168–1176, 2013.
- [29] Martin Larsson, Aaditya Ramdas, and Johannes Ruf. The numeraire e-variable and reverse information projection. *The Annals of Statistics*, 2025.
- [30] Jinzhou Li, Marloes H Maathuis, and Jelle J Goeman. Simultaneous false discovery proportion bounds via knockoffs and closed testing. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, page qkae012, 2024.
- [31] Ruth Marcus, Peritz Eric, and K Ruben Gabriel. On closed testing procedures with special reference to ordered analysis of variance. *Biometrika*, 63(3):655–660, 1976.

- [32] Iqraa Meah, Gilles Blanchard, and Etienne Roquain. False discovery proportion envelopes with m-consistency. *Journal of Machine Learning Research*, 25(270):1–52, 2024.
- [33] Violeta Muñoz-Fuentes, Pilar Cacheiro, Terrence F Meehan, Juan Antonio Aguilar-Pimentel, Steve DM Brown, Ann M Flenniken, Paul Flicek, Antonella Galli, Hamed Haseli Mashhadi, Martin Hrabě de Angelis, et al. The international mouse phenotyping consortium (IMPC): a functional catalogue of the mammalian genome that informs conservation. *Conservation genetics*, 19(4):995–1005, 2018.
- [34] Aaditya Ramdas and Tudor Manole. Randomized and exchangeable improvements of Markov’s, Chebyshev’s and Chernoff’s inequalities. *Statistical Science*, 2025.
- [35] Aaditya Ramdas and Ruodu Wang. Hypothesis testing with e-values. *arXiv preprint arXiv:2410.23614*, 2024.
- [36] Aaditya Ramdas, Fanny Yang, Martin J Wainwright, and Michael I Jordan. Online control of the false discovery rate with decaying memory. *Advances in Neural Information Processing Systems*, 30, 2017.
- [37] Aaditya Ramdas, Johannes Ruf, Martin Larsson, and Wouter Koolen. Admissible anytime-valid sequential inference must rely on nonnegative martingales. *arXiv preprint arXiv:2009.03167*, 2020.
- [38] Aaditya Ramdas, Peter Grünwald, Vladimir Vovk, and Glenn Shafer. Game-theoretic statistics and safe anytime-valid inference. *Statistical Science*, 38(4):576–601, 2023.
- [39] David S Robertson, James MS Wason, and Aaditya Ramdas. Online multiple hypothesis testing. *Statistical Science*, 38(4):557, 2023.
- [40] Joseph P Romano, Azeem Shaikh, and Michael Wolf. Consonance and the closure method in multiple testing. *The International Journal of Biostatistics*, 7(1):Art. 12, 27, 2011.
- [41] Glenn Shafer. Testing by betting: A strategy for statistical and scientific communication. *Journal of the Royal Statistical Society Series A: Statistics in Society (with discussion)*, 184(2):407–431, 2021.
- [42] Glenn Shafer, Alexander Shen, Nikolai Vereshchagin, and Vladimir Vovk. Test martingales, Bayes factors and p-values. *Statistical Science*, 2011.
- [43] Eckart Sonnemann. *Allgemeine Lösungen multipler Testprobleme*. Universität Bern. Institut für Mathematische Statistik und Versicherungslehre, 1982.
- [44] Samuel A Stouffer, Edward A Suchman, Leland C DeVinney, Shirley A Star, and Robin M Williams Jr. *The American soldier: Adjustment during army life (studies in social psychology in World War II)*, volume 1. Princeton Univ. Press, 1949.
- [45] Jinjin Tian, Xu Chen, Eugene Katsevich, Jelle Goeman, and Aaditya Ramdas. Large-scale simultaneous inference under dependence. *Scandinavian Journal of Statistics*, 50(2):750–796, 2023.
- [46] Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12:389–434, 2012.
- [47] Anna Vesely, Livio Finos, and Jelle J Goeman. Permutation-based true discovery guarantee by sum tests. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 85(3): 664–683, 2023.

- [48] Jean Ville. *Etude critique de la notion de collectif*. Gauthier-Villars Paris, 1939.
- [49] Vladimir Vovk. Testing randomness online. *Statistical Science*, 36(4):595–611, 2021.
- [50] Vladimir Vovk and Ruodu Wang. E-values: Calibration, combination and applications. *The Annals of Statistics*, 49(3):1736–1754, 2021.
- [51] Vladimir Vovk and Ruodu Wang. Confidence and discoveries with e-values. *Statistical Science*, 38(2):329–354, 2023.
- [52] Vladimir Vovk and Ruodu Wang. Merging sequential e-values via martingales. *Electronic Journal of Statistics*, 18(1):1185–1205, 2024.
- [53] Vladimir Vovk and Ruodu Wang. True and false discoveries with independent and sequential e-values. *Canadian Journal of Statistics*, 52(4):e11833, 2024.
- [54] Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. *Algorithmic learning in a random world*, volume 29. Springer, 2005.
- [55] Vladimir G Vovk. A logic of probability, with application to the foundations of statistics. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 55(2):317–341, 1993.
- [56] Volodya Vovk, Alexander Gammerman, and Craig Saunders. Machine-learning applications of algorithmic randomness. In *International Conference on Machine Learning*, pages 444–453, 1999.
- [57] A Wald. Sequential tests of statistical hypotheses. *The Annals of Mathematical Statistics*, 16(2): 117–186, 1945.
- [58] Ruodu Wang. The only admissible way of merging e-values. *arXiv preprint arXiv:2409.19888*, 2024.
- [59] 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.
- [60] Larry Wasserman, Aaditya Ramdas, and Sivaraman Balakrishnan. Universal inference. *Proceedings of the National Academy of Sciences*, 117(29):16880–16890, 2020.
- [61] Ian Waudby-Smith and Aaditya Ramdas. Estimating means of bounded random variables by betting. *Journal of the Royal Statistical Society Series B: Statistical Methodology (with discussion)*, 2023.
- [62] Ziyu Xu and Aaditya Ramdas. Dynamic algorithms for online multiple testing. In *Mathematical and Scientific Machine Learning*, pages 955–986. PMLR, 2022.
- [63] Ziyu Xu and Aaditya Ramdas. Online multiple testing with e-values. In *International Conference on Artificial Intelligence and Statistics*, pages 3997–4005. PMLR, 2024.
- [64] 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*, pages 3806–3815. PMLR, 2020.
- [65] Tijana Zrnic, Aaditya Ramdas, and Michael I Jordan. Asynchronous online testing of multiple hypotheses. *Journal of Machine Learning Research*, 22(33):1–39, 2021.

Supplementary material for “Admissible online closed testing must employ e-values”

S.1 Simple example: Lower bounds for the number of false hypotheses

A simple but interesting special case of the general online discovery process (see Figure 1) occurs if we choose $S = \{1, \dots, t\}$ at each time $t \in \mathbb{N}$. That means we observe a stream of hypotheses H_1, H_2, \dots with sequential e-values E_1, E_2, \dots and want to provide a real-time lower bound $\mathbf{d}(\{1, \dots, t\}) = d_t$ for the number of false hypotheses among H_1, \dots, H_t which holds true with high probability simultaneously over all times t :

$$\mathbb{P}(d_t \leq |\{i \leq t : H_i \text{ false}\}| \text{ for all } t \in \mathbb{N}) \geq 1 - \alpha \quad \text{for some } \alpha \in [0, 1]. \quad (\text{S.1})$$

In this case, the **SeqE-Guard** algorithm consists at each time t of two simple steps: (1) multiply the e-values up to step t ; (2) if the product is greater or equal than $1/\alpha$, increase the lower bound by 1 and exclude the largest e-value from the future analysis. More precisely, set $d_0 = 0$ and $A = \emptyset$, and then do for $t = 1, 2, \dots$:

1. Set $A = A \cup \{t\}$ and calculate $\Pi = \prod_{i \in A} E_i$.
2. If $\Pi \geq 1/\alpha$, then update $d_t = d_{t-1} + 1$ and $A = A \setminus \{\text{index of largest e-value in } A\}$; otherwise, set $d_t = d_{t-1}$.

For example, suppose $\alpha = 0.05$ and the first five e-values are

$$E_1 = 5, \quad E_2 = 4, \quad E_3 = 0.8, \quad E_4 = 0.5, \quad E_5 = 14.$$

At time $t = 2$ the product (of E_1 and E_2) equals 20 and therefore we can set $d_2 = 1$ and then exclude E_1 from the future analysis. Then, at time $t = 5$, the product (of E_2, E_3, E_4 and E_5) is again greater than 20 and therefore we can increase the lower bound and set $d_5 = 2$. Hence, in this case we can confidently claim (with probability 0.95) that there is at least one false hypothesis among H_1 and H_2 and at least two false hypotheses among H_1, \dots, H_5 . This claim remains valid regardless of how many hypotheses are tested in the future and what the e-values for these hypotheses look like.

Of course the claims above are too imprecise for many applications, as they only state that two of the five hypotheses are false, but not which ones. However, this is only a simple example. In general, the users of our **SeqE-Guard** algorithm can specify (based on the data) any subset of hypotheses in which they are interested and the **SeqE-Guard** algorithm will provide a lower bound for the number of false hypotheses in the subset that is valid simultaneously over all times and possible subsets. For example, a user might only be interested in the number of false hypotheses among H_1, H_2 and H_5 , since due to their small e-values H_3 and H_4 are unlikely to be false anyway. The **SeqE-Guard** algorithm would then still provide a lower bound of 2 for this subset, since excluding e-values smaller than 1 from the query set won't decrease the bound (see Algorithm 1). This bound would be much more informative than the same lower bound for all five hypotheses, although we cannot make claims about the truth of individual hypotheses either. However, note that this is the same for FDR procedures, as they just state that the expected proportion of true hypotheses among the rejected ones is below α .

S.2 Uniform improvements of existing methods

S.2.1 Uniform improvement of the online-adaptive method by Katsevich and Ramdas [26]

Let p-values P_1, P_2, \dots and significance levels $\alpha_1, \alpha_2, \dots$ be defined as for the **online-simple** algorithm (see Section 3.2), and the null p-values be valid conditional on the past. Furthermore, let $(\lambda_i)_{i \in \mathbb{N}}$ be additional parameters such that $\lambda_i \in [\alpha_i, 1]$ is measurable with respect to \mathcal{F}_{i-1} and $B := \sup_{i \in \mathbb{N}} \frac{\alpha_i}{1-\lambda_i} < \infty$. The **online-adaptive** bound by Katsevich and Ramdas [26]

$$d^{\text{ad}}(S_t) = \left\lceil -ca + \sum_{i=1}^t \mathbb{1}\{P_i \leq \alpha_i\} - c \frac{\alpha_i}{1-\lambda_i} \mathbb{1}\{P_i > \lambda_i\} \right\rceil$$

provides simultaneous true discovery guarantee over all sets $S_t = \{i \leq t : P_i \leq \alpha_i\}$, $t \in \mathbb{N}$, where $a > 0$ is some regularization parameter and $c = \frac{\log(1/\alpha)}{a \log(1+(1-\alpha^{B/a})/B)}$. Note that c has a different value than for the **online-simple** algorithm. Similar as demonstrated in Section 3.2, Katsevich and Ramdas [26] proved the error guarantee by showing that $E_i^{\text{ad}} = \exp[\theta_c(\mathbb{1}\{P_i \leq \alpha_i\} - c \frac{\alpha_i}{1-\lambda_i} \mathbb{1}\{P_i > \lambda_i\})]$, $i \in \mathbb{N}$, are sequential e-values, where $\theta_c = \log(1/\alpha)/(ca)$. Note that

$$\prod_{i \in I} E_i^{\text{ad}} \geq 1/\alpha \Leftrightarrow \sum_{i \in I} \mathbb{1}\{P_i \leq \alpha_i\} - c \frac{\alpha_i}{1-\lambda_i} \mathbb{1}\{P_i > \lambda_i\} \geq \log(1/\alpha)/\theta_c = ca \quad (I \in 2^{\mathbb{N}_f}).$$

With this, one can define a uniform improvement of the **online-adaptive** algorithm in the exact same manner as for the **online-simple** algorithm. Note that the **online-adaptive** method already adapts to the proportion of null hypotheses using the parameter λ_i and therefore cannot be further improved by the (online) closure principle in that direction. However, it still leads to a real uniform improvement by transforming it into a coherent procedure.

Proposition S.1. *The SeqE-Guard algorithm applied with the sequential e-values $(E_i^{\text{ad}})_{i \in \mathbb{N}}$ uniformly improves the **online-adaptive** method by Katsevich and Ramdas [26].*

Furthermore, the e-values E_i^{ad} are inadmissible if $\alpha_i/(1-\lambda_i)$ is not constant for all $i \in \mathbb{N}$ and thus can be improved. For this, note that $E_i^{\text{ad}} = \exp(\theta_c)$, if $P_i \leq \alpha_i$, $E_i^{\text{ad}} = 1$, if $\alpha_i < P_i \leq \lambda_i$ and $E_i^{\text{ad}} = \exp(-\theta_c c \alpha_i/(1-\lambda_i))$, if $P_i > \lambda_i$. Hence, for all $\mathbb{P} \in H_i$, we can provide a tight upper bound for the expectation of E_i^{ad} by

$$\begin{aligned} \mathbb{E}_{\mathbb{P}}[E_i^{\text{ad}} | \mathcal{F}_{i-1}] &= \exp(\theta_c) \mathbb{P}(P_i \leq \alpha_i | \mathcal{F}_{i-1}) \\ &\quad + \mathbb{P}(\alpha_i < P_i \leq \lambda_i | \mathcal{F}_{i-1}) + \exp\left(-\theta_c c \frac{\alpha_i}{1-\lambda_i}\right) \mathbb{P}(P_i > \lambda_i | \mathcal{F}_{i-1}) \\ &= (\exp(\theta_c) - 1) \mathbb{P}(P_i \leq \alpha_i | \mathcal{F}_{i-1}) + \mathbb{P}(P_i \leq \lambda_i | \mathcal{F}_{i-1}) \left[1 - \exp\left(-\theta_c c \frac{\alpha_i}{1-\lambda_i}\right)\right] \\ &\quad + \exp\left(-\theta_c c \frac{\alpha_i}{1-\lambda_i}\right) \\ &\leq (\exp(\theta_c) - 1) \alpha_i + \lambda_i \left[1 - \exp\left(-\theta_c c \frac{\alpha_i}{1-\lambda_i}\right)\right] + \exp\left(-\theta_c c \frac{\alpha_i}{1-\lambda_i}\right) \\ &=: u(\alpha, \alpha_i, \lambda_i, B, a), \end{aligned}$$

which can easily be calculated for given α , α_i , λ_i , B and a . If $B = \alpha_i/(1-\lambda_i)$, then $u(\alpha, \alpha_i, \lambda_i, B, a) = 1$. However, in practice $\alpha_i/(1-\lambda_i)$ may vary over time such that there are indices $j \in \mathbb{N}$ with $\alpha_j/(1-\lambda_j) < B$. In this case, the e-value E_j^{ad} becomes conservative. For example, if $\alpha = 0.1$, $\alpha_i = 0.1$, $\lambda_i = 0.5$, $B = 0.4$ and $a = 1$, then $u(\alpha, \alpha_i, \lambda_i, B, a) = 0.966$.

S.2.2 Uniform improvements of the methods by Meah et al. [32]

In this Subsection, we derive uniform improvements of the online procedures introduced by Meah et al. [32], the **u-online-simple** and **u-online-Freedman** method. It should be noted that the uniform improvements presented here are not instances of our **SeqE-Guard** algorithm, however, they are obtained by taking the union of many **SeqE-Guard** bounds or by taking the mean of many test martingales (instead of a single one) for each intersection test, and hence are closely related to **SeqE-Guard**.

S.2.2.1 Uniform improvement of the u-online-simple method

Meah et al. [32] have introduced a modified version of the **online-simple** method by Katsevich and Ramdas [26]. In the following, we show how this method can be uniformly improved by e-value based online closed testing.

Let p-values P_1, P_2, \dots and significance levels $\alpha_1, \alpha_2, \dots$ be defined as for the **online-simple** algorithm (see Section 3.2), and the null p-values be valid conditional on the past. Instead of fixing the parameter $a > 0$ in advance, Meah et al. [32] combined **online-simple** bounds for different parameters using a union bound. More precisely, for each $a \in \mathbb{N}$ let $\alpha(a) = \frac{6\alpha}{a^2\pi^2}$ and $c_a = \frac{\log(1/\alpha(a))}{a \log(1+\log(1/\alpha(a))/a)}$. By (11), we have that for all $\mathbb{P} \in \mathcal{P}$:

$$\mathbb{P}(\mathbf{d}_a^{\text{os}}(S_t) \leq |S_t \cap I_1^{\mathbb{P}}| \text{ for all } t \in \mathbb{N}) \geq 1 - \alpha(a), \quad (\text{S.2})$$

$$\text{where } \mathbf{d}_a^{\text{os}}(S_t) = \left\lceil -c_a a + \sum_{i=1}^t \mathbb{1}\{P_i \leq \alpha_i\} - c_a \alpha_i \right\rceil \quad (t \in \mathbb{N}), \quad (\text{S.3})$$

with $\alpha_1, \alpha_2, \dots$ being nonnegative thresholds and $S_t = \{i \leq t : P_i \leq \alpha_i\}$. Meah et al. [32] proposed the bound

$$\mathbf{d}^{\text{u-os}}(S_t) = \max_{a \in \mathbb{N}} \mathbf{d}_a^{\text{os}}(S_t) \quad t \in \mathbb{N}, \quad (\text{S.4})$$

whose true discovery guarantee follows by applying a union bound to (S.2). We refer to the procedure $\mathbf{d}^{\text{u-os}}$ as **u-online-simple** method in this paper.

In Section 3.2 we have shown that the bound in (S.3) can be uniformly improved by Algorithm 2 for all $a > 0$ and $\alpha(a) \in (0, 1)$. Hence, a simple uniform improvement of the **u-online-simple** method can be obtained by taking the maximum of these improved bounds. We call this improvement **closed u-online-simple** procedure.

More precisely, let $E_i^{\text{os},a} = \exp[\theta_{c_a}(\mathbb{1}\{P_i \leq \alpha_i\} - c_a \alpha_i)]$ and define

$$W_I^{t,a} := \prod_{i \in I \cap \{1, \dots, t\}} E_i^{\text{os},a}.$$

Then the **closed u-online-simple** method is given by $\mathbf{d}^{\text{cu-os}}(S_t) = \max_{a \in \mathbb{N}} \mathbf{d}^{\phi^a}(S_t)$, where \mathbf{d}^{ϕ^a} is the online closed procedure obtained by the intersection tests

$$\phi_I^a := \mathbb{1} \left\{ \exists t \in I : W_I^{t,a} \geq 1/\alpha(a) \right\} = \mathbb{1} \left\{ \exists t \in I : \frac{6}{a^2\pi^2} W_I^{t,a} \geq 1/\alpha \right\}, \quad a \in \mathbb{N}, I \subseteq \mathbb{N}.$$

Since $\mathbf{d}^{\phi^a}(S_t) \geq \mathbf{d}_a^{\text{os}}(S_t)$ for all $a \in \mathbb{N}$ (see Section 3.2), we also have $\mathbf{d}^{\text{cu-os}}(S_t) \geq \mathbf{d}^{\text{u-os}}(S_t)$. However, we can improve $\mathbf{d}^{\text{cu-os}}$ even further.

For this, note that the intersection test ϕ_I^a can be uniformly improved simultaneously for all a by

$$\phi_I^{\text{m-os}} := \mathbb{1} \left\{ \exists t \in I : \sum_{a \in \mathbb{N}} \frac{6}{a^2\pi^2} W_I^{t,a} \geq 1/\alpha \right\}.$$

Since (weighted) means of test supermartingales are test supermartingales again, Ville's inequality implies that $\phi_I^{\text{m-os}}$ is an intersection test and the online closed procedure $\mathbf{d}^{\text{m-os}} := \mathbf{d}^{\phi^{\text{m-os}}}$ provides simultaneous true discovery guarantee. We refer to $\mathbf{d}^{\text{m-os}}$ as **m-online-simple** procedure in the following. Since $\phi_I^{\text{m-os}} \geq \phi_I^a$ for all $a \in \mathbb{N}$, we have

$$\mathbf{d}^{\text{m-os}}(S_t) \geq \mathbf{d}^{\text{cu-os}}(S_t) \geq \mathbf{d}^{\text{u-os}}(S_t) \text{ for all } t \in \mathbb{N}.$$

Note that $E_i^{\text{os},a'} \geq E_j^{\text{os},a'}$ for some $a' \in \mathbb{N}$ and $i \neq j$ implies that $E_i^{\text{os},a} \geq E_j^{\text{os},a}$ for all $a \in \mathbb{N}$ and $E_i^{\text{os},a'} \geq 1$ for some $a' \in \mathbb{N}$ and $i \in \mathbb{N}$ implies that $E_i^{\text{os},a} \geq 1$ for all $a \in \mathbb{N}$. With this, we can derive a short-cut for $\mathbf{d}^{\text{m-os}}(S_t)$, $t \in \mathbb{N}$, in the same manner as we did in Algorithm 1 and Algorithm 2 (see Algorithm 5). We capture these results in the following proposition.

Proposition S.2. *The m-online-simple procedure (Algorithm 5) and the cu-online-simple procedure uniformly improve the u-online-simple method by Meah et al. [32].*

Algorithm 5 m-online-simple

Input: Sequence of p-values P_1, P_2, \dots and sequence of (potentially data-dependent) individual significance levels $\alpha_1, \alpha_2, \dots$.

Output: Query sets $S_1 \subseteq S_2 \subseteq \dots$ and true discovery bounds $d_1 \leq d_2 \leq \dots$.

```

1:  $d_0 = 0$ 
2:  $S_0 = \emptyset$ 
3:  $A_0^c = \emptyset$ 
4: for  $t = 1, 2, \dots$  do
5:   if  $P_t \leq \alpha_t$  then
6:      $S_t = S_{t-1} \cup \{t\}$ 
7:   else
8:      $S_t = S_{t-1}$ 
9:   end if
10:  if  $\sum_{a \in \mathbb{N}} \frac{6}{a^2 \pi^2} \prod_{i \in \{1, \dots, t\} \setminus A_{t-1}^c} E_i^{\text{os},a} \geq 1/\alpha$  then
11:     $d_t = d_{t-1} + 1$ 
12:     $A_t^c = A_{t-1}^c \cup \{\text{index of smallest individual significance level in } S_t \setminus A_{t-1}^c\}$ 
13:  else
14:     $d_t = d_{t-1}$ 
15:     $A_t^c = A_{t-1}^c$ 
16:  end if
17:  return  $S_t, d_t$ 
18: end for

```

In Figure S.1, we compare the u-online-simple method [32] with our m-online-simple procedure (Algorithm 5). The simulation setup is the same as described in Section 4. The plots show that the improvement obtained by the m-online-simple procedure is substantial.

S.2.2.2 Uniform improvement of the u-online-Freedman method

Similarly as the u-online-simple method (Section S.2.2.1), the u-online-Freedman procedure is obtained by taking the (weighted) union of many individual bounds. However, instead of using the online-simple bounds by Katsevich and Ramdas [26], the u-online-Freedman procedure is based on Freedman's inequality [13].

Let p-values P_1, P_2, \dots and significance levels $\alpha_1, \alpha_2, \dots$ be defined as for the online-simple algorithm (see Section 3.2), and the null p-values be valid conditional on the past. Meah et al. [32]




Figure S.1: True discovery proportion bounds obtained by the `u-online-simple` method [32] and our `m-online-simple` method. The `m-online-simple` procedure improves the `u-online-simple` method substantially.

showed (see Corollary 41 in [32]) that for all $\mathbb{P} \in \mathcal{P}$:

$$\mathbb{P}(\mathbf{d}_a^{\text{Freed}}(S_t) > |S_t \cap I_1^{\mathbb{P}}| \text{ and } A_t \leq a \text{ for some } t \in \mathbb{N}) \leq \alpha(a), \quad (\text{S.5})$$

$$\text{where } \mathbf{d}_a^{\text{Freed}}(S_t) = 1 + \left[-\kappa_a + \sum_{i=1}^t \mathbb{1}\{P_i \leq \alpha_i\} - \alpha_i \right], \quad (\text{S.6})$$

$$\text{and } B_t = \sum_{i=1}^t \alpha_i(1 - \alpha_i), \quad (\text{S.7})$$

with $S_t = \{i \leq t : P_i \leq \alpha_i\}$, $\kappa_a = \sqrt{2a \log(1/\alpha(a))} + \frac{\log(1/\alpha(a))}{2}$ and $\alpha(a) = \alpha\left(\frac{6}{\max(2 \log_2(a), 1)^2(\pi^2 + 6)}\right)$ for some parameter $a = 2^{j/2}$, $j \in \mathbb{N} \cup \{0\}$. The `u-online-Freedman` procedure is then obtained by applying a union to (S.5) over all $j \in \mathbb{N} \cup \{0\}$ (see Corollary 42 in [32]). In the following, we show that the `SeqE-Guard` algorithm allows to uniformly improve the bound $\mathbf{d}_a^{\text{Freed}}$ for each a .

Howard et al. [23] improved Freedman's inequality by a test martingale approach. Our improvement of the bound $\mathbf{d}_a^{\text{Freed}}$ is based on the same technique, which additionally uses the `SeqE-Guard` algorithm. For this, we define the sequential e-values

$$E_i^{\text{Freed}, a} := \exp(\lambda_a(\mathbb{1}\{P_i \leq \alpha_i\} - \alpha_i) - \psi(\lambda_a)\alpha_i(1 - \alpha_i)) \quad (i \in \mathbb{N}), \quad (\text{S.8})$$

where $\lambda_a = \log\left(1 + \frac{\kappa_a}{a}\right)$ and $\psi(\lambda_a) = \exp(\lambda_a) - \lambda_a - 1$. To see that $E_i^{\text{Freed}, a}$, $i \in \mathbb{N}$, are sequential e-values, define $X_i := \mathbb{1}\{P_i \leq \alpha_i\} - \alpha_i$ and note that for all $i \in I_0^{\mathbb{P}}$ it holds that $X_i \leq 1$, $\mathbb{E}_{\mathbb{P}}[X_i | \mathcal{F}_{i-1}] \leq 0$ and $\mathbb{E}_{\mathbb{P}}[X_i^2 | \mathcal{F}_{i-1}] \leq \alpha_i(1 - \alpha_i)$. Therefore, $\mathbb{E}_{\mathbb{P}}[\exp(\lambda_a X_i) | \mathcal{F}_{i-1}] \leq \exp(\psi(\lambda_a)\alpha_i(1 - \alpha_i))$ (e.g., Lemma 6.7 in [46]). Now let $d_t^{\text{Freed}, a}$ be the bound obtained by applying `SeqE-Guard` with

$E_1^{\text{Freed},a}, E_2^{\text{Freed},a}, \dots$. In the same manner as in (13), we get that

$$\begin{aligned} d_t^{\text{Freed},a} &\geq 1 + \left[-\frac{\log(1/\alpha(a))}{\lambda_a} + \sum_{i=1}^t \mathbb{1}\{P_i \leq \alpha_i\} - \alpha_i - \frac{\psi(\lambda_a)\alpha_i(1-\alpha_i)}{\lambda_a} \right] \\ &= 1 + \left[-\frac{\log(1/\alpha(a)) + \psi(\lambda_a)B_t}{\lambda_a} + \sum_{i=1}^t \mathbb{1}\{P_i \leq \alpha_i\} - \alpha_i \right]. \end{aligned}$$

Further, if $B_t \leq a$, we obtain

$$\begin{aligned} & -\frac{\log(1/\alpha(a)) + \psi(\lambda_a)B_t}{\lambda_a} \\ & \geq -\frac{\log(1/\alpha(a)) + \psi(\lambda_a)a}{\lambda_a} \\ & = -\frac{\log(1/\alpha(a)) + \kappa_a - \log\left(1 + \frac{\kappa_a}{a}\right)(a + \kappa_a) + \log\left(1 + \frac{\kappa_a}{a}\right)\kappa_a}{\log\left(1 + \frac{\kappa_a}{a}\right)} \\ & \geq -\frac{\log(1/\alpha(a)) - 2(\sqrt{a + \kappa_a} - \sqrt{a})^2 + \log\left(1 + \frac{\kappa_a}{a}\right)\kappa_a}{\log\left(1 + \frac{\kappa_a}{a}\right)} \\ & = -\frac{\log(1/\alpha(a)) - 2(\sqrt{a + \kappa_a} - \sqrt{a})^2}{\log\left(1 + \frac{\kappa_a}{a}\right)} - \kappa_a \\ & = -\frac{\log(1/\alpha(a)) - 2\left(\sqrt{a + \sqrt{2a\log(1/\alpha(a))} + \frac{\log(1/\alpha(a))}{2}} - \sqrt{a}\right)^2}{\log\left(1 + \frac{\kappa_a}{a}\right)} - \kappa_a \\ & = -\frac{\log(1/\alpha(a)) - 2\left(\sqrt{\left(\sqrt{a} + \sqrt{\frac{\log(1/\alpha(a))}{2}}\right)^2} - \sqrt{a}\right)^2}{\log\left(1 + \frac{\kappa_a}{a}\right)} - \kappa_a \\ & = -\frac{\log(1/\alpha(a)) - 2\left(\sqrt{\frac{\log(1/\alpha(a))}{2}}\right)^2}{\log\left(1 + \frac{\kappa_a}{a}\right)} - \kappa_a \\ & = -\kappa_a, \end{aligned}$$

where the second inequality follows by Lemma 43 of Meah et al. [32]. This shows that $d_t^{\text{Freed},a} \geq \mathbf{d}_a^{\text{Freed}}(S_t)$ for all $a \geq B_t$ and we additionally improve (S.5) by providing a nontrivial bound for $a < B_t$. Hence, a uniform improvement of the **u-online-Freedman** procedure can either be obtained by taking the maximum of the improved bounds $d_t^{\text{Freed},a}$ over all $a = 2^j$, $j \in \mathbb{N} \cup \{0\}$, or by the more powerful mean strategy described in Appendix S.2.2.1. In line with Appendix S.2.2.1, we call these two procedures **cu-online-Freedman** and **m-online-Freedman**, respectively.

Proposition S.3. *The m-online-Freedman procedure and the cu-online-Freedman procedure uniformly improve the u-online-Freedman method by Meah et al. [32].*

Note that the sequential e-values $E_i^{\text{Freed},a}$, $i \in \mathbb{N}$, are just binary e-values depending on $\mathbb{1}\{P_i \leq \alpha_i\}$. Hence, they are very similar to the ones defined for the **online-simple** method (12), the difference lies just in the weighting of the two cases $P_i \leq \alpha_i$ and $P_i > \alpha_i$. Therefore, the **SeqE-Guard** algorithm does not only improve these methods, it also facilitates their interpretation. We only need to ensure that the expected value of each sequential e-value is bounded by one under the null hypothesis, which is very easy to check. Indeed, one can check that $E_i^{\text{Freed},a}$ is slightly conservative as well and could be




Figure S.2: True discovery proportion bounds obtained by the **u-online-Freed** method [32] and our **m-online-Freed** method. The **m-online-Freed** procedure improves the **u-online-Freed** method substantially.

improved, we just chose the representation (S.8) as it simplifies the proof of the uniform improvement. Such looseness would be difficult to detect with the procedures by Katsevich and Ramdas [26] and Meah et al. [32], as their proofs of validity are based on far more complicated arguments.

In Figure S.2, we compare the **u-online-Freed** method [32] with our **m-online-Freed** procedure. The simulation setup is the same as described in Section 4. The behavior of the **u-online-Freed** and **m-online-Freed** method is similar as for the **u-online-simple** and **m-online-simple** method (see Figure S.1). The **m-online-Freed** method leads to a significant uniform improvement over the **u-online-Freed** method.

S.3 Calibrating sequential p-values into e-values

While e-values have been of recent interest, most studies still use p-values as measure of evidence against the null hypothesis. In this case, we can *calibrate* the p-values into e-values and then apply **SeqE-Guard**. In this context, a calibrator is a function from p-values to e-values, meaning it takes as an input a p-value and yields an e-value as an output.

For example, the **online-simple** method of Katsevich and Ramdas [26] introduced in Section 3.2 is implicitly based on calibrating each p-value into a simple binary e-value. However, there are infinitely many other calibrators that could be used. Assume that p-values $(P_t)_{t \in \mathbb{N}}$ for the hypotheses $(H_t)_{t \in \mathbb{N}}$ are given. A decreasing function $f : [0, 1] \rightarrow [0, \infty]$ is a calibrator if $\int_0^1 f(x) dx \leq 1$, and it is admissible if equality holds [55, 50]. Note that if the p-values are sequential p-values, meaning P_t is measurable with respect to \mathcal{F}_t and $\mathbb{P}(P_t \leq x | \mathcal{F}_{t-1}) \leq x$ for all $x \in [0, 1]$, then the calibrated e-values are sequential e-values.

An example calibrator is

$$h_x(p) = \exp(x\Phi^{-1}(1-p) - x^2/2), \quad (\text{S.9})$$

where Φ denotes the CDF of a standard normal distribution. To see that this is a valid calibrator, note that $\Phi^{-1}(U)$, where $U \sim [0, 1]$, follows a standard normal distribution and the moment generating function of a standard normal distribution for the real parameter $x > 0$ is given by $\exp(x^2/2)$. Duan et al. [8] used this calibrator to derive a martingale Stouffer [44] global test. This global test is based on a confidence sequence of Howard et al. [23].

We compare the application of **SeqE-Guard** with calibrated e-values and hedged GRO e-values experimentally in the next section.

S.4 Boosting of sequential e-values

Wang and Ramdas [59] introduced a way to *boost* e-values before plugging them into their e-BH procedure without violating the desired FDR control. In this section, we propose a similar (but simpler) approach for **SeqE-Guard** that will improve its power for a specific query path.

We begin by noting that whenever the bound d_t is increased by one, the largest e-value with index in A will not be considered in the following analysis. Hence, extremely large e-values will be excluded from the analysis anyway, which makes it possible to truncate the e-values at a specific threshold without changing its outcome. This makes the resulting truncated e-values conservative under the null, and one can improve the procedure by multiplying the e-value by a suitable constant larger than 1 to remove its conservativeness. This truncation+multiplication operation is what is referred to as *boosting* the e-value.

To this end, recall that $A_{t-1} \subseteq S_{t-1}$, $t \in \mathbb{N}$, is the index set of all previous e-values in the query set that were not already excluded by **SeqE-Guard** and $U_{t-1} \subseteq \{1, \dots, t-1\} \setminus S_{t-1}$ the index set of all previous e-values that are not contained in the query set and that are smaller than 1. Now, define

$$m_t := \max \left\{ \max_{i \in A_{t-1}} E_i, \frac{1}{\alpha \prod_{i \in A_{t-1} \cup U_{t-1}} E_i} \right\}, \quad (\text{S.10})$$

and note that m_t is predictable (measurable with respect to \mathcal{F}_{t-1}). Furthermore, if $E_t \geq m_t$ and $t \in S_t$, then $d_t = d_{t-1} + 1$ and E_t will be excluded in the further analysis. If $E_t \geq m_t$ and $t \notin S_t$, then $E_t \geq 1$, since $m_t \geq 1$ by definition, and E_t won't be considered in the analysis anyway. Hence, we define the truncation function $T_t : [0, \infty] \rightarrow [0, m_t]$ as

$$T_t(x) := x \mathbb{1}\{x \leq m_t\} + m_t \mathbb{1}\{x > m_t\} \quad (\text{S.11})$$

and then choose a boosting factor $b_t \geq 1$ as large as possible such that

$$\mathbb{E}_{\mathbb{P}}[T_t(b_t E_t) | \mathcal{F}_{t-1}] \leq 1 \text{ for all } \mathbb{P} \in H_t. \quad (\text{S.12})$$

Note that $b_t = 1$ always satisfies (S.12); so a boosting factor always exists and is always at least one. Condition (S.12) immediately implies that $T_t(b_t E_t)$ is a sequential e-value. Furthermore, using $b_t E_t$ in **SeqE-Guard** yields exactly the same results as using $T_t(b_t E_t)$. Therefore, applying **SeqE-Guard** to the boosted e-values $(b_t E_t)_{t \in \mathbb{N}}$ provides simultaneous true discovery guarantee and is uniformly more powerful than with non-boosted e-values, since $b_t \geq 1$. Note that in this case m_t should also be calculated based on the boosted e-values $b_1 E_1, \dots, b_{t-1} E_{t-1}$. As also mentioned by Wang and Ramdas [59], one could use different functions than $x \mapsto bx$ for some $b \geq 1$ to boost the e-values. In general, it is only required that each boosted e-value E_t^{boost} , $t \in \mathbb{N}$, satisfies $\mathbb{E}_{\mathbb{P}}[T_t(E_t^{\text{boost}}) | \mathcal{F}_{t-1}] \leq 1$ for all $\mathbb{P} \in H_t$. We summarize this result in the following theorem.

Proposition S.1. *Let $E_1^{\text{boost}}, E_2^{\text{boost}}, \dots$ be a sequence of nonnegative random variables such that $\mathbb{E}_{\mathbb{P}}[T_t(E_t^{\text{boost}}) | \mathcal{F}_{t-1}] \leq 1$ for all $\mathbb{P} \in H_t$, where T_t is given by (S.11) and m_t is calculated as in (S.10) based on $E_1^{\text{boost}}, \dots, E_{t-1}^{\text{boost}}$. Then, applying **SeqE-Guard** to the boosted sequential e-values $E_1^{\text{boost}}, E_2^{\text{boost}}, \dots$ provides simultaneous true discovery guarantee.*

It should be noted that boosting adapts the sequential e-values to the query path $(S_t)_{t \geq 1}$ and Proposition S.1 only holds for the query path for which the boosted e-values $E_1^{\text{boost}}, E_2^{\text{boost}}, \dots$ were constructed. If multiple query paths are investigated, one must use the sequential e-values $T_1(E_1^{\text{boost}}), T_2(E_2^{\text{boost}}), \dots$ for all other query paths. Therefore, it may happen that boosting weakens the performance of **SeqE-Guard** for the other query paths.

In the following we provide several examples that illustrate how the boosting factors can be determined in specific cases and demonstrate the possible gain in efficiency.

Example S.1. We consider Example 3 from Wang and Ramdas [59] adapted to our setting. For each $t \in \mathbb{N}$, we test the simple null hypothesis $H_t : X_t | \mathcal{F}_{t-1} \sim \mathcal{N}(\mu_0, 1)$ against the simple alternative $H_t^A : X_t | \mathcal{F}_{t-1} \sim \mathcal{N}(\mu_1, 1)$, where X_t denotes the data for H_t . In this case, the GRO e-value is given by the likelihood ratio between two normal distributions with variance 1 and means μ_1 and μ_0

$$E_t = \exp(\delta Z_t - \delta^2/2), \quad (\text{S.13})$$

where $\delta = \mu_1 - \mu_0 > 0$ and $Z_t = X_t - \mu_0$ follows a standard normal distribution conditional on \mathcal{F}_{t-1} under H_t . Hence, conditional on the past, each null e-value follows a log-normal distribution with parameter $(-\delta^2/2, \delta)$. With this, we obtain for all $t \in \mathbb{I}_0^{\mathbb{P}}$:

$$\begin{aligned} & \mathbb{E}_{\mathbb{P}}[b_t E_t \mathbb{1}\{b_t E_t \leq m_t\} + m_t \mathbb{1}\{b_t E_t > m_t\} | \mathcal{F}_{t-1}] \\ &= b_t \mathbb{E}_{\mathbb{P}}[E_t \mathbb{1}\{b_t E_t \leq m_t\} | \mathcal{F}_{t-1}] + m_t \mathbb{P}(b_t E_t > m_t | \mathcal{F}_{t-1}) \\ &= b_t \left[1 - \Phi \left(\frac{\delta}{2} - \frac{\log(m_t/b_t)}{\delta} \right) \right] + m_t \left[1 - \Phi \left(\frac{\log(m_t/b_t) + \delta^2/2}{\delta} \right) \right], \end{aligned}$$

where Φ is the CDF of a standard normal distribution. The last expression can be set equal to 1 and then be solved for b_t numerically. For example, for $\delta = 3$ and $m_t = 20$, we obtain $b_t = 3.494$. Hence, the e-value E_t could be multiplied by 3.494 without violating the true discovery guarantee, a substantial gain. In general, the larger m_t , the smaller is the boosting factor. For example, if $m_t = 5$, then $b_t = 11.826$ and if $m_t = 100$, then $b_t = 1.774$. Nevertheless, even the latter boosting factor would increase the power of the true discovery procedure significantly and we would usually expect m_t to be smaller than 100 in most settings. If we use, as described in Section 3.3, the e-value $1 - \lambda_t + \lambda_t E_t^{\text{GRO}}$, $\lambda_t \in (0, 1)$ instead, we need to solve

$$m_t + \Phi \left(\frac{\log(s_t) + \delta^2/2}{\delta} \right) [b_t(1 - \lambda_t) - m_t] + b_t \lambda_t \left[1 - \Phi \left(\frac{\delta}{2} - \frac{\log(s_t)}{\delta} \right) \right] = 1,$$

for $b_t \in [1, 1/(1 - \lambda_t))$, where $s_t = (\lambda_t - 1 + m_t/b_t)/\lambda_t$. In this case, $\delta = 3$, $\lambda_t = 0.5$ and $m_t = 20$ yield a boosting factor of $b_t = 1.354$.

Example S.2. Suppose we observe sequential p-values P_1, P_2, \dots and want to apply the calibrator (S.9). If the p-values are uniformly distributed conditional on the past, the resulting e-value has the exact same distribution as the e-value in (S.13) under the null hypothesis for $\delta = x$, where x is the freely chosen parameter for the calibrator. Hence, we can do the exact same calculations to obtain an appropriate boosting factor. If the p-values are stochastically larger than uniform, we could still use that same boosting factor, as the resulting e-values provide true discovery guarantee but might be conservative.

Example S.3. In case of the **closed online-simple** method (Algorithm 2) it is particular simple to “boost” the e-values. Since E_t^{os} only takes two different values, we can simply ensure $E_t^{\text{os}} \leq m_t$ by choosing α_t such that $E_t^{\text{os}} \leq m_t$ if $\mathbb{1}\{P_t \leq \alpha_t\} = 1$. Note that in case of $\alpha_t = \nu$ for all $t \in \mathbb{N}$ and some $\nu > 0$ such that $\exp[\theta_c(1 - c\nu)] \leq 1/\alpha$ it is not possible to improve the bounds of the **closed online-simple** method further by boosting, since we already have $E_t^{\text{os}} \leq m_t$ almost surely.

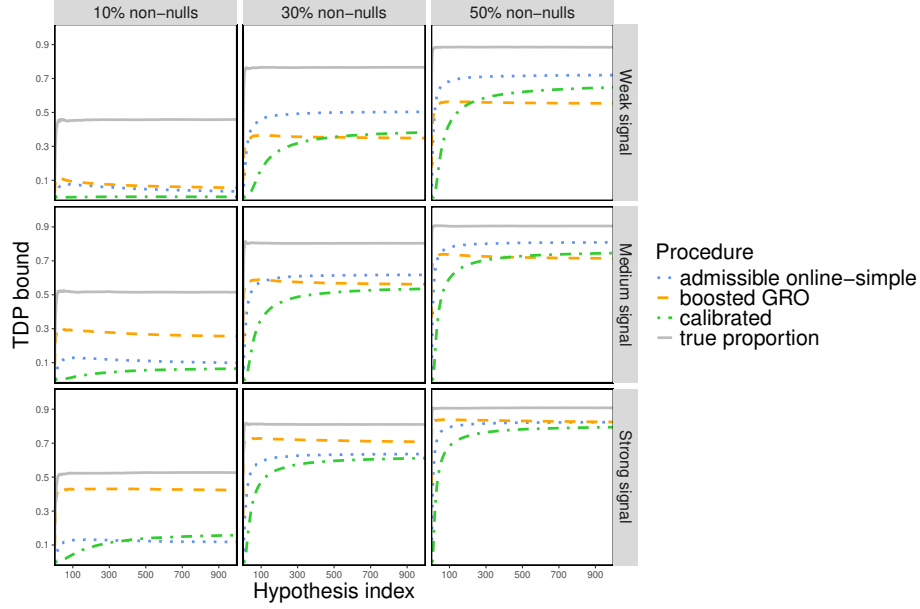


Figure S.3: True discovery proportion bounds obtained by the `admissible online-simple` as well as the `SeqE-Guard` algorithm with boosted GRO and (boosted) calibrated e-values. The boosted GRO e-values perform best if the proportion of false hypotheses is small and the `admissible online-simple` method performs best if the proportion is large.

In Figure S.3 we compare the `admissible online-simple` method with the `SeqE-Guard` algorithm with boosted GRO (the GRO e-values were first hedged and then boosted as described in Example S.1) and calibrated e-values (using the calibrator in (S.9) with $x = 0.1$ and boosting as described in Example S.2). The calibrated e-values perform worst in most cases, boosted GRO e-values perform best if the proportion of false hypotheses is small and the `admissible online-simple` method performs best if the proportion of false hypotheses is large. Also note that compared to Figure 6, the bound of the hedged GRO e-values increased substantially using boosting. However, it adapts the e-values to the query paths and thus the bounds might decrease for other query paths.

S.5 Omitted proofs

Proof of Theorem 2.1. The simultaneous true discovery guarantee of \mathbf{d} implies that the ϕ_I define intersection tests. Furthermore, it follows that ϕ is increasing since \mathbf{d} is coherent and ϕ_S is an online intersection test since \mathbf{d} is an online procedure. Therefore, it only remains to show that $\mathbf{d}(S) \leq \mathbf{d}^\phi(S)$ for all $S \in 2^{\mathbb{N}_f}$. Suppose $\mathbf{d}(S) = s$. Due to the coherence of \mathbf{d} , we have $\phi_J = 1$ for all $J \subseteq S$ with $|J| > |S| - s$. The coherence further implies that $\phi_I = 1$ for all $I \subseteq \mathbb{N}$ with $|S \cap I| > |S| - s$ and hence $\mathbf{d}^\phi(S) \geq s$. \square

Proof of Theorem 2.2. Since $(\psi_i^I)_{i \in I}$, $I \subseteq \mathbb{N}$, is an anytime-valid test for H_I , it immediately follows that ϕ_I is an online intersection test. Furthermore, for $i \in \mathbb{N}$, $I \subseteq \{1, \dots, i\}$ and $K \subseteq \{k \in \mathbb{N} : k > i\}$

it holds that $J \cap \{1, \dots, \sup(I \cup K)\} = I \cup K$ implies that $J \cap \{1, \dots, \sup(I)\} = I$. Therefore,

$$\begin{aligned}\phi_I &= \inf \left\{ \sup_{i \in I} \psi_i^J : J \cap \{1, \dots, \sup(I)\} = I \right\} \\ &\leq \inf \left\{ \sup_{i \in I} \psi_i^J : J \cap \{1, \dots, \sup(I \cup K)\} = I \cup K \right\} \\ &\leq \inf \left\{ \sup_{i \in I \cup K} \psi_i^J : J \cap \{1, \dots, \sup(I \cup K)\} = I \cup K \right\} = \phi_{I \cup K},\end{aligned}$$

showing that ϕ is increasing. \square

Proof of Theorem 2.3. Since $\phi_{I \cap \{1, \dots, i\}}$ is an online intersection test, we have that ψ_i^I is measurable with respect to \mathcal{F}_i . Furthermore, since ϕ is increasing, it holds that $\mathbb{P}(\exists i \in I : \psi_i^I = 1) = \mathbb{P}(\phi_I = 1) \leq \alpha$ for all $\mathbb{P} \in H_I$. Hence, $(\psi_i^I)_{i \in I}$ is an anytime-valid test for H_I with respect to the filtration $(\mathcal{F}_i)_{i \in I}$. Now let $\tilde{\phi}$ be defined by (6). Then for all $S \in 2^{\mathbb{N}_f}$,

$$\begin{aligned}\tilde{\phi}_S &= \inf \left\{ \sup_{i \in S} \psi_i^J : J \cap \{1, \dots, \sup(S)\} = S \right\} \\ &= \inf \left\{ \sup_{i \in S} \phi_{J \cap \{1, \dots, i\}} : J \cap \{1, \dots, \sup(S)\} = S \right\} \\ &= \sup_{i \in S} \phi_{S \cap \{1, \dots, i\}} = \phi_S,\end{aligned}$$

where the last equality follows since ϕ is increasing. Note that this implies $\tilde{\mathbf{d}}^\phi = \mathbf{d}^\phi$, since $\tilde{\phi}$ and ϕ are increasing and therefore $\tilde{\mathbf{d}}^\phi(S)$ and $\mathbf{d}^\phi(S)$, $S \in 2^{\mathbb{N}_f}$, can be determined solely based on $\tilde{\phi}_I$ and ϕ_I , respectively, with $I \subseteq S$. \square

Proof of Fact 1. The tower property for conditional expectations immediately implies that $(M_i^\mathbb{P})_{i \in I \cup \{0\}}$ is a test martingale for \mathbb{P} and Ville's inequality shows that $(\tilde{\psi}_i)_{i \in I}$ is an anytime-valid test for H_I . Furthermore, $\psi_i = 1$ implies that $\mathbb{P}(\exists j \in I : \psi_j = 1 | \mathcal{F}_i) = 1$ and therefore $M_i^\mathbb{P} \geq 1/\alpha$ for all $\mathbb{P} \in H_I$. \square

Proof of Theorem 3.1. Let d_t , $t \in \mathbb{N}$, be the bounds of **SeqE-Guard**, $A_t \subseteq S_t$ be the set A at time t before checking whether $\Pi_{A \cup U} \geq 1/\alpha$, $U_t \subseteq \{1, \dots, t\} \setminus S_t$ the index set of e-values that are smaller than 1 and Π_Z , $Z \subseteq \mathbb{N}$, be the product of all e-values with index in Z . Since ϕ is increasing, we have $d_1 = \mathbf{d}^\phi(S_1)$. Now assume that $d_1 = \mathbf{d}^\phi(S_1), \dots, d_{t-1} = \mathbf{d}^\phi(S_{t-1})$ and $S_t = S_{t-1} \cup \{t\}$. Due to the coherence of \mathbf{d}^ϕ , it holds $d_{t-1} \leq \mathbf{d}^\phi(S_t) \leq d_{t-1} + 1$. In the following, we show that $\Pi_{A_t \cup U_t} \geq 1/\alpha$ implies that $\mathbf{d}^\phi(S_t) \geq d_{t-1} + 1$ and $\Pi_{A_t \cup U_t} < 1/\alpha$ implies that $\mathbf{d}^\phi(S_t) < d_{t-1} + 1$, which proves the assertion.

We first show that $\mathbf{d}^\phi(S_t) \geq d_{t-1} + 1$, if $\Pi_{A_t \cup U_t} \geq 1/\alpha$. For this, we prove that $\Pi_{A_t \cup U_t} \geq 1/\alpha$ implies $\phi_I = 1$ for all $I = V \cup W$, where $V \subseteq S_t$ and $W \subseteq \{1, \dots, t\} \setminus S_t$ with $|V| \geq |A_t|$. Since ϕ is increasing, we have $\mathbf{d}^\phi(S_t) = \min\{|S_t \setminus I| : I \subseteq \{1, \dots, t\}, \phi_I = 0\}$, which then implies that $\mathbf{d}^\phi(S_t) \geq |S_t| - |A_t| + 1 = d_{t-1} + 1$. First, note that it is sufficient to show the claim for all $I = V \cup U_t$ with $|V| \geq |A_t|$, since multiplication with e-values that are larger or equal than 1 cannot decrease the product. Now let such an I and V be fixed. Let t_1, \dots, t_m , where $t_m = t_{|S_t| - |A_t| + 1} = t$, be the times at which $\Pi_{A_{t_i} \cup U_{t_i}} \geq 1/\alpha$ and $\tilde{m} \in \{1, \dots, m\}$ be the smallest index such that $|V \cap \{1, \dots, t_{\tilde{m}}\}| > |S_{t_{\tilde{m}}}| - \tilde{m}$. Note that \tilde{m} always exists, because $|A_t| = |A_{t_m}| = |S_{t_m}| - m + 1$. With this, we have

$$\Pi_{(V \cap \{1, \dots, t_{\tilde{m}}\}) \cup U_{t_{\tilde{m}}}} \geq \Pi_{(A_t \cap \{1, \dots, t_{\tilde{m}-1}\}) \cup (\{t_{\tilde{m}-1} + 1, \dots, t_{\tilde{m}}\} \cap S_t) \cup U_{t_{\tilde{m}}}} = \Pi_{A_{t_{\tilde{m}}} \cup U_{t_{\tilde{m}}}} \geq 1/\alpha.$$

The first inequality follows since $A_t \cap \{1, \dots, t_{\tilde{m}-1}\}$ minimizes the product of the e-values over all subsets $J \subseteq S_{t_{\tilde{m}-1}}$ with $|J| = |S_{t_{\tilde{m}-1}}| - (\tilde{m} - 1)$ that satisfy $|J \cap \{1, \dots, t_i\}| \leq |S_{t_i}| - i$ for all $i \in \{1, \dots, \tilde{m} - 1\}$ and $(\{t_{\tilde{m}-1} + 1, \dots, t_{\tilde{m}}\} \cap S_t) \subseteq V$ (due to definition of \tilde{m}), where $t_0 = 0$.

Hence, it remains to show that $\mathbf{d}^\phi(S_t) < d_{t-1} + 1$, if $\Pi_{A_t \cup U_t} < 1/\alpha$. Since $\Pi_{(A_t \cap \{1, \dots, i\}) \cup U_t} < 1/\alpha$ for all $i \in \{1, \dots, t-1\}$, $\Pi_{A_t \cup U_t} < 1/\alpha$ implies that $\phi_{A_t \cup U_t} = 0$. Furthermore, since $|S_t \setminus A_t| = d_{t-1}$, the claim follows. \square

Proof of Proposition 3.3. Due to [41], the e-value maximizing the growth rate under $(1 - \tau_i)\mathbb{P}_i + \tau_i\mathbb{Q}_i$ is given by the likelihood ratio

$$\frac{d[(1 - \tau_i)\mathbb{P}_i + \tau_i\mathbb{Q}_i]}{d\mathbb{P}_i} = 1 - \tau_i + \tau_i \frac{d\mathbb{Q}_i}{d\mathbb{P}_i} = 1 - \tau_i + \tau_i E_i^{\text{GRO}},$$

where $\frac{d\mathbb{Q}_i}{d\mathbb{P}_i}$ denotes the Radon-Nikodym derivative. \square

Proof of Theorem 6.1. For each $I \subseteq \mathbb{N}$, we consider the intersection test

$$\phi_I^c = \mathbb{1} \left\{ \sum_{i \in I \cap S_{\sup(I)}} E_i \gamma_{t(i; I)} \geq 1/\alpha \right\},$$

where $t(i; I) = |\{j \in I : j \leq i\}|$ and $S_\infty = \bigcup_{t \in \mathbb{N}} S_t$. Obviously, this intersection test is more conservative than ϕ_I defined in (17) and therefore $\mathbf{d}^{\phi^c}(S_t) \leq \mathbf{d}^\phi(S_t)$ for all $t \in \mathbb{N}$. Let $d_t, t \in \mathbb{N}$, be the bounds of ArbE-Guard and $A_t^c \subseteq S_t$ be the set A^c at time t before checking whether $\sum_{i \in S_t \setminus A^c} E_i \gamma_{t(i; \{1, \dots, t\} \setminus A^c)} \geq 1/\alpha$. We will show that $\mathbf{d}^{\phi^c}(S_t) = d_t$ for all $t \in \mathbb{N}$. Let $t \in \mathbb{N}$ with $\phi_{\{1, \dots, t\} \setminus A_t^c}^c = 1$ be fixed. Similar as in the proof of Theorem 3.1, we need to show that this implies $\phi_I^c = 1$ for all $I = V \cup W$, where $V \subseteq S_t$ and $W \subseteq \{1, \dots, t\} \setminus S_t$ with $|V| \geq |S_t \setminus A_t^c|$. First, note that it is sufficient to show the claim for all $W = \{1, \dots, t\} \setminus S_t$, since $(\gamma_i)_{i \in \mathbb{N}}$ is nonincreasing. Now let a $V \subseteq S_t$ with $|V| \geq |S_t \setminus A_t^c|$ and $W = \{1, \dots, t\} \setminus S_t$ be fixed.

Let t_1, \dots, t_m , where $t_m = t_{|A_t^c|+1} = t$, be the times at which $\phi_{\{1, \dots, t_i\} \setminus A_{t_i}^c}^c = 1$ and $\tilde{m} \in \{1, \dots, m\}$ be the smallest index such that $|V \cap \{1, \dots, t_{\tilde{m}}\}| > |S_{t_{\tilde{m}}}| - \tilde{m}$. Note that \tilde{m} always exists, because $|S_t \setminus A_t^c| = |S_{t_m} \setminus A_{t_m}^c| = |S_{t_m}| - m + 1$. With this, we have

$$\begin{aligned} \sum_{i \in (V \cup W) \cap \{1, \dots, t_{\tilde{m}}\} \cap S_{t_{\tilde{m}}}} E_i \gamma_{t(i; (V \cup W) \cap \{1, \dots, t_{\tilde{m}}\})} &= \sum_{i \in V \cap \{1, \dots, t_{\tilde{m}}\}} E_i \gamma_{t(i; (V \cup W) \cap \{1, \dots, t_{\tilde{m}}\})} \\ &\geq \sum_{i \in S_{t_{\tilde{m}}} \setminus A_{t_{\tilde{m}}}^c} E_i \gamma_{t(i; [S_{t_{\tilde{m}}} \setminus A_{t_{\tilde{m}}}^c] \cup W_{t_{\tilde{m}}})} \\ &\geq 1/\alpha, \end{aligned}$$

where $W_{t_{\tilde{m}}} = W \cap \{1, \dots, t_{\tilde{m}}\}$. The first inequality follows since $(S_t \setminus A_t) \cap \{1, \dots, t_{\tilde{m}-1}\}$ minimizes $\sum_{i \in J} E_i \gamma_{t(i; J \cup W_{t_{\tilde{m}-1}})}$ among all $J \subseteq S_{t_{\tilde{m}-1}}$ with $|J| = |S_{t_{\tilde{m}-1}}| - (\tilde{m} - 1)$ that satisfy $|J| \leq |S_{t_i}| - i$ for all $i \in \{1, \dots, \tilde{m} - 1\}$ and $(\{t_{\tilde{m}-1}, \dots, t_{\tilde{m}}\} \cap S_t) \subseteq V$ (due to the definition of \tilde{m}), where $t_0 = 0$. \square

In order to prove Proposition 6.2, we first show the following three lemmas. The first two are used to prove the third one, which then implies the proposition.

Lemma S.1. Let $n \in \mathbb{N}$, $k < n$, and $a_1 \geq \dots \geq a_k \geq a_{k+1} = \dots = a_n = a \geq 0$ be some real numbers such that $\sum_{i=1}^n a_i = 1$. Furthermore, let $b_i, i \in \{1, \dots, n\}$, be non-negative real numbers such that

$$\frac{1}{k} \sum_{i=1}^k b_i \leq \frac{1}{n-k} \sum_{i=k+1}^n b_i.$$

Then there exist $a_1^* \geq \dots \geq a_{k-1}^* \geq a_k^* = \dots = a_n^* \geq a$ with $\sum_{i=1}^n a_i^* = 1$ and

$$\frac{1}{n} \sum_{i=1}^n a_i b_i \leq \frac{1}{n} \sum_{i=1}^n a_i^* b_i.$$

Proof. Define $a_i^* = a_i - \frac{n-k}{n}(a_k - a)$, $i \leq k$, and $a_i^* = a + (a_k - a)\frac{k}{n}$, $i > k$. Since $a_k \geq a$, we have

$$\sum_{i=1}^n a_i^* = \sum_{i=1}^n a_i - \frac{k(n-k)}{n}(a_k - a) + (a_k - a)\frac{k(n-k)}{n} = 1.$$

In addition, it holds for all $i > k$

$$a_k^* = a_k - \frac{n-k}{n}(a_k - a) = \frac{k}{n}a_k + a\frac{n-k}{n} = a_i^*.$$

Furthermore, we have

$$\begin{aligned} \sum_{i=1}^n a_i^* b_i &= \sum_{i=1}^k a_i b_i - \frac{k(n-k)}{n}(a_k - a)\frac{1}{k} \sum_{i=1}^k b_i + \sum_{i=k+1}^n a_i b_i + \frac{k(n-k)}{n}(a_k - a)\frac{1}{n-k} \sum_{i=k+1}^n b_i \\ &\geq \sum_{i=1}^n a_i b_i. \end{aligned}$$

□

Lemma S.2. Let $n \in \mathbb{N}$, $k < n$, and $a_1 \geq \dots \geq a_k \geq a_{k+1} = \dots = a_n = a \geq 0$ be some real numbers such that $\sum_{i=1}^n a_i = 1$. Furthermore, let b_i , $i \in \{1, \dots, n\}$, be non-negative real numbers such that

$$\frac{1}{k} \sum_{i=1}^k b_i \geq \frac{1}{n-k} \sum_{i=k+1}^n b_i.$$

Then there exist $a_1^* \geq \dots \geq a_k^* \geq 0$ with $\sum_{i=1}^k a_i^* = 1$ and

$$\frac{1}{n} \sum_{i=1}^n a_i b_i \leq \frac{1}{n} \sum_{i=1}^k a_i^* b_i.$$

Proof. Define $a_i^* = a_i + \frac{a(n-k)}{k}$ for all $i \in \{1, \dots, k\}$. Then,

$$\sum_{i=1}^k a_i^* = \sum_{i=1}^n a_i = 1.$$

Furthermore, we have

$$\begin{aligned} \sum_{i=1}^k a_i^* b_i &= \sum_{i=1}^k a_i b_i + \frac{a(n-k)}{k} \sum_{i=1}^k b_i \\ &\geq \sum_{i=1}^k a_i b_i + a \sum_{i=k+1}^{n-k} b_i \\ &= \sum_{i=1}^n a_i b_i. \end{aligned}$$

□

Lemma S.3. Let $n \in \mathbb{N}$ and $a_1 \geq \dots \geq a_n \geq 0$ be some real numbers such that $\sum_{i=1}^n a_i = 1$. Furthermore, let b_i , $i \in \{1, \dots, n\}$, be arbitrary non-negative real numbers. Then,

$$\sum_{i=1}^n a_i b_i \leq \max_{t \in \{1, \dots, n\}} \frac{1}{t} \sum_{i=1}^t b_i.$$

Proof. Start with $k = n - 1$ and apply Lemma S.1 and Lemma S.2 recursively for all $k = n - 1, \dots, 1$ (at each step k choose the lemma that applies for b_1, \dots, b_n). Then we end up with a $\tilde{k} \leq n$ and $a_1^* = \dots = a_{\tilde{k}}^* = \frac{1}{\tilde{k}}$ such that

$$\frac{1}{n} \sum_{i=1}^n a_i b_i \leq \sum_{i=1}^{\tilde{k}} a_i^* b_i = \frac{1}{\tilde{k}} \sum_{i=1}^{\tilde{k}} b_i.$$

□

Note that Lemma S.3 is a generalization of Chebyshev's sum inequality which applies in case of $b_1 \leq \dots \leq b_n$. In this case, $\max_{t \in \{1, \dots, n\}} \frac{1}{t} \sum_{i=1}^t b_i = \frac{1}{n} \sum_{i=1}^n b_i$.

Proof of Proposition 6.2. Let $I \in 2^{\mathbb{N}_f}$ be arbitrary. The proof follows immediately by Lemma S.3, if we replace a_i by $\gamma_{t(i;I)}$ and b_i by E_i . □