Tuning the false discovery rate in the closed eBH procedure

Vladimir Vovk

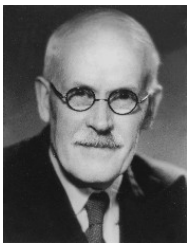

Users of these tests speak of the
5 per cent. point [p-value of 5%]
in much the same way as I should
speak of the $K = 10^{-1/2}$ point
[e-value of $10^{1/2}$], and of the 1
per cent. point [p-value of 1%]
as I should speak of the
$K = 10^{-1}$ point [e-value of 10].

Project "Hypothesis testing with e-values"

Working Paper #11

First posted November 21, 2026. Last revised November 24, 2026.

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

Abstract

The closed eBH procedure controls the false discovery rate under arbitrary dependence between the input e-values. It takes the target false discovery rate α as its input and outputs a rejection set R, the usual goal being to maximize its size $|R|$. This note discusses a modification of the closed eBH procedure that allows choosing the target false discovery rate post hoc: the rejection set and its level α can be chosen after seeing the data, and the expected ratio of the false discovery proportion to α is at most 1. We adapt Vovk and Wang's (2023) algorithm for computing discovery e-matrices to computing, in quadratic time, the frontier of achievable combinations $(|R|, \alpha)$ of the size of a rejection set R and its level α. The frontier allows us to compute an optimal rejection set in view of a given utility function F, $F(|R|, \alpha) \to \max$, striking a balance between the goals of reducing α and increasing $|R|$.

Contents

1 Introduction **1**

2 The closed eBH procedure **1**

3 The relative eBH procedure **4**

4 Computing the frontier **7**

5 Conclusion **11**

References **12**

1 Introduction

The closed eBH procedure in [7, (7)] (see also [3, Definition 9.31]) is an elegant improvement of the original eBH procedure [6], but its disadvantage is that we need to decide in advance on the target false discovery rate α. For example, what if we can slightly exceed α but then reject many more null hypotheses? This note discusses a computationally efficient way of choosing both the rejection set R and its *level* α that plays the role of the target false discovery rate.

We start in Sect. 2 with fixing our terminology and notation and a short description of the closed eBH procedure and its standard property of validity. The next section, Sect. 3, gives a straightforward post hoc modification of the closed eBH procedure, which we call the "relative eBH procedure"; it is used, more or less explicitly, in several publications (see, e.g., [7, Sect. 8.3] and [3, Sect. 9.6]). The levels α output by the relative eBH procedure are valid in relative terms: the expected ratio of the false discovery proportion to the reported level is at most 1. Section 4 gives an efficient algorithm for the relative eBH procedure by adapting an algorithm in [5, Algorithm 4] for computing discovery e-matrices. For both closed and relative eBH procedures we give time-uniform bounds illustrating their validity for sequences of studies, and Sect. 5 summarizes some advantages of the relative eBH procedure.

We will introduce most of our terminology and notation at the point where we need them. The notation for the natural numbers is $\mathbb{N} := \{1, 2, \dots\}$ and $\mathbb{N}_0 := \{0, 1, \dots\}$. The terms "increasing" and "decreasing" will be used in the wide sense allowing intervals of constancy.

2 The closed eBH procedure

Our sample space $(\Omega, \mathcal{A})$ is arbitrary, and our statistical model is the set $\mathfrak{P}(\Omega)$ of all probability measures on $(\Omega, \mathcal{A})$; since this model contains all probability measures on Ω, there is no real loss of generality. For $\theta \in \mathfrak{P}(\Omega)$ and a nonnegative random variable E, we set $\mathbb{E}^\theta(E) := \int E \, \mathrm{d}\theta$ (this is always defined, perhaps with $\mathbb{E}^\theta(E) = \infty$).

Suppose that we are given K *e-variables* $E^{(1)}, \dots, E^{(K)}$ for testing hypotheses $H_1, \dots, H_K \subseteq \mathfrak{P}(\Omega)$, which are our *base hypotheses*: each $E^{(k)}$ is a nonnegative random variable satisfying $\mathbb{E}^\theta(E^{(k)}) \le 1$ for all $\theta \in H_k$. We would like to reject some of the base hypotheses (in fact, as many of them as possible under a validity constraint). A *rejection set* is a set $R \subseteq [K]$, where $[K] := \{1, \dots, K\}$; this is the set of base hypotheses, represented by their indices, that the researcher chooses to reject. For each $\theta \in \mathfrak{P}(\Omega)$, let

$$I_\theta := \{k \in [K] \mid \theta \in H_k\}$$

be the indices of the hypotheses containing θ. If the researcher rejects H_k, we refer to this decision as a *discovery*. We say that the discovery is *true* if $\theta \notin H_k$, and it is *false* if $\theta \in H_k$, where θ is the true (unknown) probability measure governing the data generation.

Let

$$\mathrm{FDP}_I(R) := \frac{|R \cap I|}{|R| \vee 1},$$

where $I, R \subseteq [K]$. In particular, $\mathrm{FDP}_{I_\theta}(R)$ is the *false discovery proportion* of a rejection set R when the data are generated from θ. Our informal goal is to make $\mathrm{FDP}_{I_\theta}(R)$ small for a large R; the obvious obstacle is, of course, that we do not know θ.

For a nonempty $I \subseteq [K]$, the arithmetic mean

$$E_I := \frac{1}{|I|} \sum_{k \in I} E^{(k)} \tag{1}$$

is an e-variable for testing $\cap_{k \in I} H_k$. Following [3, Sect. 9.6], we also set $E_\emptyset := 1$. Therefore, $\mathbb{E}^\theta(E_{I_\theta}) \le 1$ for all $\theta \in \mathfrak{P}(\Omega)$. The family $(E_I)_{I \subseteq [K]}$ is our *e-collection*.

The *closed eBH procedure* of [7, (7)] (see also [3, Definition 9.31]) takes $\alpha \in (0, 1]$ (the target false discovery rate) as an input and outputs a (largest) set R such that

$$\alpha E_I \ge \mathrm{FDP}_I(R) \quad \text{for all } I \subseteq [K]. \tag{2}$$

The validity of the closed eBH procedure follows immediately from $\mathbb{E}^\theta(E_{I_\theta}) \le 1$. The rejection set in the following proposition does not have to be a largest one: the researcher may choose any rejection set satisfying (2), possibly after looking at the data.

Proposition 1. *Let $\alpha \in (0, 1]$. For every $\theta \in \mathfrak{P}(\Omega)$ and every random set $\hat{R} \subseteq [K]$ satisfying (2),*

$$\mathbb{E}^\theta \left(\mathrm{FDP}_{I_\theta}(\hat{R}) \right) \le \alpha.$$

Proof. Taking $I := I_\theta$ in (2) gives $\mathrm{FDP}_{I_\theta}(\hat{R}) \le \alpha E_{I_\theta}$, and we know that $\mathbb{E}^\theta(E_{I_\theta}) \le 1$. □

The left-hand side of the inequality in Proposition 1 is the false discovery rate of $\hat{R}$ under θ, and so Proposition 1 says that the closed eBH procedure controls the false discovery rate at level α [7].

Proposition 1 is about a single study. Let us now consider a sequence of studies indexed by $n = 1, 2, \dots$, each of the kind described above (with the number of base hypotheses that may vary from study to study). In the nth study, the researcher chooses a rejection set R_n, possibly after looking at the data, and reports a level α_n for it. The studies are *sequential*: the base hypotheses and the e-variables of the nth study, as well as the probability measure θ_n generating its data, may depend on the outcomes of the earlier studies. Let $\mathcal{F}_n$ be the σ-algebra generated by the outcomes of the first n studies, including the researcher's choices ($\mathcal{F}_0$ is the trivial σ-algebra), and let $\mathbb{P}$ and $\mathbb{E}$ refer to the joint distribution of all studies. We assume that θ_n is the conditional distribution of the data of the nth study given $\mathcal{F}_{n-1}$. Let $I_n := I_{\theta_n}$ be the set of indices of the true base hypotheses in the nth study, E_n be the e-variable E_{I_n}

of the nth study (i.e., (1) for $I := I_n$), and $\mathrm{FDP}_n := \mathrm{FDP}_{I_n}(R_n)$ be its false discovery proportion. Our assumptions imply

$$\mathbb{E}(E_n \mid \mathcal{F}_{n-1}) \le 1, \quad n \in \mathbb{N}, \tag{3}$$

and this is all we will need; (3) holds, in particular, when the studies are independent.

Proposition 2. *Suppose that the level $\alpha_n \in (0,1]$ for the nth study is chosen before the nth study is carried out (formally, it is $\mathcal{F}_{n-1}$-measurable) and that R_n satisfies (2) with α_n in place of α (and with the e-variables of the nth study). For each $\delta \in (0,1]$, with probability at least $1-\delta$,*

$$\forall N \in \mathbb{N}: \quad \sum_{n=1}^{N} (\mathrm{FDP}_n - \alpha_n) \le 3\sqrt{N+1}\, \ln \frac{2(N+1)^2}{\delta}. \tag{4}$$

In particular,

$$\limsup_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} (\mathrm{FDP}_n - \alpha_n) \le 0 \quad \text{a.s.}$$

According to (4), if the closed eBH procedure is used at the same level α in all studies, with probability at least $1-\delta$, the average false discovery proportion of the first N studies is at most $\alpha + O(N^{-1/2} \ln \frac{N}{\delta})$ as $N \to \infty$, and its lim sup is at most α a.s. These are frequentist interpretations of the closed eBH procedure's validity. The rejection sets R_n may be chosen after looking at the data, but only among the sets satisfying (2) at the level α_n chosen before the study; the next section relaxes this requirement.

Proof of Proposition 2. Set

$$Y_n := \mathrm{FDP}_n - \alpha_n \in [-1,1], \quad S_N := \sum_{n=1}^{N} Y_n, \quad V_N := \sum_{n=1}^{N} Y_n^2.$$

Let us prove that, with probability at least $1-\delta$,

$$\forall N \in \mathbb{N}: \quad S_N \le 3\sqrt{V_N+1}\, \ln \frac{2(V_N+1)^2}{\delta},$$

which is stronger than (4). Taking $I := I_n$ in (2) gives $\mathrm{FDP}_n \le \alpha_n E_n$, i.e., $Y_n \le \alpha_n(E_n - 1)$.

For $\lambda \in [0,1]$, let

$$T_N^\lambda := \prod_{n=1}^{N} \big(1 + \lambda \alpha_n (E_n - 1)\big)$$

be the capital at time N of a gambler who starts from 1 and bets the fraction $\lambda\alpha_n$ of his current capital on each E_n. Each factor is nonnegative, since $\lambda\alpha_n \le 1$ and $E_n \ge 0$, and so, by (3) and the $\mathcal{F}_{n-1}$-measurability of α_n, T^λ is a test

supermartingale, i.e., a nonnegative supermartingale with respect to $(\mathcal{F}_N)$ with $T_0^\lambda = 1$. For $\lambda \in [0, 1/2]$, we have $\lambda Y_n \geq -1/2$, and so

$$T_N^\lambda \geq \prod_{n=1}^{N} (1 + \lambda Y_n) \geq \exp\left(\lambda S_N - \lambda^2 V_N\right), \tag{5}$$

where the first inequality follows from $Y_n \leq \alpha_n(E_n - 1)$, and the second from $1 + t \geq \exp(t - t^2)$ for $t \geq -1/2$. See, e.g., [4, (3.5)] for the second inequality; the expression $\exp(\dots)$ in (5) is sometimes called *Doléans's supermartingale* [4, Sect. 3.2].

Now let

$$\lambda_k := \frac{1}{2\sqrt{k}} \quad \text{and} \quad w_k := \frac{1}{k(k+1)}, \quad k \in \mathbb{N}; \tag{6}$$

notice that $\sum_k w_k = 1$. The idea is to use λ_k at the times N at which $V_N \approx k$; the two addends on the right-hand side of (7) below are then of the same order of magnitude (apart from the logarithm in the second addend), and the factor $1/2$ ensures $\lambda_k \leq 1/2$. The mixture $M_N := \sum_k w_k T_N^{\lambda_k}$ (the capital of a gambler who splits his initial capital among the strategies above in the proportions w_k) is also a test supermartingale, and so, by Ville's inequality, $M_N < 1/\delta$ for all N with probability at least $1 - \delta$. On this event, (5) (applicable as $\lambda_k \leq 1/2$) gives $w_k \exp(\lambda_k S_N - \lambda_k^2 V_N) \leq M_N < 1/\delta$, i.e.,

$$S_N < \lambda_k V_N + \frac{1}{\lambda_k} \ln \frac{k(k+1)}{\delta}, \tag{7}$$

for all N and k. Given N, take $k := \max(1, \lceil V_N \rceil)$, so that $V_N \leq k \leq V_N + 1$. Then $\lambda_k V_N \leq \sqrt{V_N}/2$ and $k(k+1) \leq (V_N + 1)(V_N + 2) \leq 2(V_N + 1)^2$, and so

$$S_N < \frac{1}{2}\sqrt{V_N} + 2\sqrt{V_N + 1} \ln \frac{2(V_N + 1)^2}{\delta} \leq 3\sqrt{V_N + 1} \ln \frac{2(V_N + 1)^2}{\delta},$$

where the last inequality uses $\ln(2(V_N + 1)^2/\delta) \geq \ln 2 > 1/2$.

For the last statement of the proposition, let $\delta \to 0$. □

See Algorithm 1 for an efficient implementation of the closed eBH procedure.

3 The relative eBH procedure

In this section we allow the level α to be chosen post hoc. For each rejection set $R \subseteq [K]$ define the random variable

$$\alpha_R := \min\{\alpha \in [0, 1] \mid \forall I \subseteq [K] : \alpha E_I \geq \mathrm{FDP}_I(R)\}, \tag{8}$$

where $\min \emptyset := \infty$. We will refer to the version of the closed eBH procedure that assigns the level α_R to each rejection set R as the *relative eBH procedure*.

Suppose we have decided on a utility function $F : \mathbb{N}_0 \times [0, 1] \to \mathbb{R}$ that is increasing in its first argument and decreasing in the second, and extend it

by $F(r,\infty) := -\infty$, $r \in \mathbb{N}_0$. We will assume that it is computable in time $O(1)$. The main topic of this note is to design an efficient algorithm for the optimization problem

$$F(|R|, \alpha_R) \to \max \tag{9}$$

over $R \subseteq [K]$. The algorithm is given as its input the e-values $E^{(k)}(\omega)$, $k \in [K]$, for the realized outcome $\omega \in \Omega$.

By the definition (8) of α_R,

$$\mathrm{FDP}_I(R)/\alpha_R \le E_I \tag{10}$$

for all $I, R \subseteq [K]$, with the conventions $0/0 := 0$ and $c/\infty := 0$ (if $\alpha_R = 0$, then $\mathrm{FDP}_I(R) = 0$ for all I). The following property of validity is stated in [3, Theorem 9.32] and [7, Sect. 8.3].

Proposition 3. *For every $\theta \in \mathfrak{P}(\Omega)$ and every random set $\hat{R} \subseteq [K]$, in particular for a solution $\hat{R}$ of the optimization problem (9),*

$$\mathbb{E}^\theta\left(\frac{\mathrm{FDP}_{I_\theta}(\hat{R})}{\alpha_{\hat{R}}}\right) \le \mathbb{E}^\theta\left(\max_{R\subseteq[K]} \frac{\mathrm{FDP}_{I_\theta}(R)}{\alpha_R}\right) \le 1.$$

Proof. The inequality (10) with $I := I_\theta$ gives $\max_R \mathrm{FDP}_{I_\theta}(R)/\alpha_R \le E_{I_\theta}$, and we know that $\mathbb{E}^\theta(E_{I_\theta}) \le 1$. □

Proposition 3 says that $\alpha_{\hat{R}}$ is a valid post hoc level for $\hat{R}$ in the sense of [1] and [2]: the expected ratio of the false discovery proportion to the reported level is at most 1, even though $\hat{R}$, and therefore its level, is chosen after looking at the data. Since $\hat{R}$ can be any random set, the utility function F in (9) does not even have to be fixed in advance: one can compute the whole frontier (defined formally in the next section) and choose a point on it after inspecting it.

Remark 4. Proposition 3 says that the expected ratio of the false discovery proportion to the reported level is at most 1 even if the level is chosen post hoc. If the level α is fixed in advance, we also have the more familiar property that the expected difference between the false discovery proportion and α is at most 0, which is the property of FDR control stated in Proposition 1. This property, $\mathbb{E}^\theta(\mathrm{FDP}_{I_\theta}(\hat{R}) - \alpha_{\hat{R}}) \le 0$, may be violated when the level is chosen post hoc. For example, suppose that only the first of $K := 2$ base hypotheses is true and that the e-variables $(E^{(1)}, E^{(2)})$ take the values $(2, 100)$ and $(0, 100)$ with probability $1/2$ each. If the researcher chooses the largest rejection set whose level is at most $1/4$ and reports its level, she reports $\{1, 2\}$ with level $1/4$ in the first case and $\{2\}$ with level $1/50$ in the second case. The expected difference between the false discovery proportion and the reported level is then

$$\frac{1/2 - 1/4}{2} + \frac{0 - 1/50}{2} = \frac{23}{200} > 0,$$

whereas the expected ratio is $(2 + 0)/2 = 1$. (For checking calculations such as the preceding one and the one in Remark 8 below, it is useful to consult

Lemma 9.) The most obvious way to ensure that the expected false discovery proportion is at most the expected reported level is to fix, before seeing the data, a reference level $c \in (0,1]$ and to set

$$\alpha_R^c := \max_{I \subseteq [K]} \big(\mathrm{FDP}_I(R) - c(E_I - 1) \big); \tag{11}$$

this *absolute eBH procedure* will then satisfy $\mathbb{E}^\theta(\mathrm{FDP}_{I_\theta}(\hat{R}) - \alpha_{\hat{R}}^c) \le 0$. While the absolute eBH procedure contains the closed eBH procedure with the target false discovery rate of c as an admissible option (similarly to the relative eBH procedure), it might not contain much more in practically interesting cases as $\alpha_R^c \ge (1 - E^{\min})c$ for any rejection set R, where $E^{\min}$ is the smallest e-value among $E^{(1)}, \dots, E^{(K)}$. (The reference level c may depend on I, in particular on $|I|$, in (11), but it does not seem to help much.)

Proposition 3 is about a single study. Let us now consider a sequence of studies, as in Sect. 2, where in each study the researcher chooses both the rejection set R_n and a level α_n for it, allowing $\alpha_n > \alpha_{R_n}$ (to be discussed after the proof of Corollary 6). We have the following analogue of Proposition 2 for levels chosen post hoc.

Proposition 5. *Suppose that $\alpha_{R_n} \le \alpha_n \le 1$ for all $n \in \mathbb{N}$, where α_{R_n} is the level (8) of R_n in the nth study, and set*

$$S_N := \sum_{n=1}^N \left(\frac{\mathrm{FDP}_n}{\alpha_n} - 1 \right), \quad V_N := \sum_{n=1}^N \left(\frac{\mathrm{FDP}_n}{\alpha_n} - 1 \right)^2$$

(with $0/0 := 0$). For each $\delta \in (0,1]$, with probability at least $1 - \delta$,

$$\forall N \in \mathbb{N}: \quad S_N \le 3\sqrt{V_N + 1} \ln \frac{2(V_N + 1)^2}{\delta}. \tag{12}$$

Proof. Set $Y_n := \mathrm{FDP}_n/\alpha_n - 1$, so that $S_N = \sum_{n=1}^N Y_n$ and $V_N = \sum_{n=1}^N Y_n^2$. If $R_n = \emptyset$, then $Y_n = -1$; otherwise, $\mathrm{FDP}_n/\alpha_n \le \mathrm{FDP}_n/\alpha_{R_n} \le E_n$ by (10). Therefore, $-1 \le Y_n \le E_n - 1$ for all n. It remains to repeat the argument of the proof of Proposition 2, replacing α_n by 1 in the upper bound $\alpha_n(E_n - 1)$ for Y_n and in the definition of T_N^λ, which becomes $\prod_{n=1}^N (1 - \lambda + \lambda E_n)$; since α_n no longer enters T_N^λ, it does not need to be $\mathcal{F}_{n-1}$-measurable. □

Corollary 6. *Suppose, in addition, that $\alpha_n \ge \alpha_{\min}$ whenever $R_n \ne \emptyset$, for some $\alpha_{\min} \in (0,1]$. Then, for each $\delta \in (0,1]$, with probability at least $1 - \delta$,*

$$\forall N \in \mathbb{N}: \quad \sum_{n=1}^N \frac{\mathrm{FDP}_n}{\alpha_n} \le N + \frac{3}{\alpha_{\min}} \sqrt{N + 1} \ln \frac{2(N+1)^2}{\alpha_{\min}^4 \delta}. \tag{13}$$

In particular,

$$\limsup_{N \to \infty} \frac{1}{N} \sum_{n=1}^N \frac{\mathrm{FDP}_n}{\alpha_n} \le 1 \quad a.s.$$

Proof. Since $0 \le \mathrm{FDP}_n/\alpha_n \le 1/\alpha_{\min}$, we have $(\mathrm{FDP}_n/\alpha_n - 1)^2 \le 1/\alpha_{\min}^2$, and so $V_N + 1 \le (N+1)/\alpha_{\min}^2$. It remains to plug this into (12), whose right-hand side is increasing in V_N, and, for the last statement, to let $\delta \to 0$. □

Allowing $\alpha_n > \alpha_{R_n}$ in Proposition 5 might appear counterintuitive, since using α_n weakens the researcher's statement about the rejection set R_n. However, small α_n make the bound in (12) looser; in particular, we needed a lower bound on α_n in Corollary 6 in order to have a function of N as an upper bound on the average deviation of FDP_n/α_n from 1. A natural choice for α_n in Corollary 6 is $\alpha_n := \max(\alpha_{R_n}, \alpha_{\min})$; this ensures the existence of a positive floor for α_n used for controlling the term V_N in (12) while using the optimal level α_{R_n} above the floor.

The rejection sets R_n and the levels α_n in Proposition 5 and Corollary 6 can be chosen in any way, in particular post hoc (as long as $\alpha_{R_n} \le \alpha_n$), and so Corollary 6 gives an interpretation of post hoc levels in terms of the ratios FDP_n/α_n, provided the reported levels are bounded away from 0. For a level α fixed in advance ($\alpha_n := \alpha$ for all n), we can take $\alpha_{\min} := \alpha$, and then (13) multiplied by α becomes almost identical to (4) with $\alpha_n = \alpha$.

Remark 7. For a fixed δ, the bound in Corollary 6 exceeds N by $O(\sqrt{N}\ln N)$, whereas the law of the iterated logarithm (especially in its game-theoretic form [4, Chap. 5]) suggests $O(\sqrt{N \ln\ln N})$. The extra factor comes from the grid (6) in the proof of Proposition 2: λ_k is used when $V_N \approx k$, and so the cost $\ln(1/w_k)$ of using it is about $2\ln V_N$. It is cheaper to use the kth grid point when $V_N \approx 2^k$, which reduces this cost to about $2\ln\log_2 V_N$; λ_k should then be chosen by minimizing the right-hand side of (7) for $V_N := 2^k$ (subject to $\lambda_k \le 1/2$), as in [4, Sect. 5.1]. (There is no need to change the weights w_k, since we just need a slowly converging series $\sum_k w_k$.) For fixed δ and $\alpha_{\min}$, this gives us (13) with $O(\sqrt{N \ln\ln N})$ in place of the second addend on the right-hand side (the precise expression is, however, more complicated than (13)). This order cannot be improved: if the study in the example of Remark 4 is repeated independently, the ratios FDP_n/α_n are independent and take the values 0 and 2 with probability 1/2 each, and so, by the law of the iterated logarithm, for any $\epsilon > 0$, $S_N > (1-\epsilon)\sqrt{2N \ln\ln N}$ for infinitely many N almost surely.

4 Computing the frontier

In this section we fix the realized outcome $\omega \in \Omega$ and write $e_k := E^{(k)}(\omega)$, $e_I := E_I(\omega)$, and $\alpha_R := \alpha_R(\omega)$. Set

$$\beta_R := \max_{I \subseteq [K]} \frac{\mathrm{FDP}_I(R)}{e_I}, \quad R \subseteq [K],$$

with the conventions $0/0 := 0$ and $c/0 := \infty$ for $c > 0$. Then

$$\alpha_R = \begin{cases} \beta_R & \text{when } \beta_R \le 1 \\ \infty & \text{otherwise.} \end{cases}$$

The *frontier* is the function $r \mapsto \alpha^*_r$, where

$$\alpha^*_r := \min_{R \subseteq [K] : |R| = r} \alpha_R, \quad r \in \{0, \dots, K\}.$$

Since F is decreasing in its second argument,

$$F(|R|, \alpha_R) \le F(|R|, \alpha^*_{|R|}) \quad \text{for all } R \subseteq [K]. \tag{14}$$

Therefore, to solve the optimization problem (9) it suffices to know the frontier and, for each r, a set of size r whose level is α^*_r.

Remark 8. The sequence α^*_r is not necessarily monotone: e.g., for the e-values $(1, 10, 10)$ we have $(\alpha^*_0, \dots, \alpha^*_3) = (0, 2/11, 1/7, 1/3)$, and so rejecting two base hypotheses requires a smaller level than rejecting one. For solving optimization problems (9) we only need, of course, the Pareto frontier, i.e., the running minimum of our frontier from the right.

If we do not know anything about the nature of the base hypotheses, it makes sense to reject a number of them with the largest e-values (see, e.g., [5, Sect. 6] and [7, Lemma 52]), and Lemma 9 below shows that this is optimal in our current setting. From now on we assume, without loss of generality, that the e-values are sorted in ascending order,

$$e_1 \le \dots \le e_K, \tag{15}$$

and set $R_0 := \emptyset$ and

$$R_r := \{K - r + 1, \dots, K\}, \quad r \in [K];$$

R_r consists of the hypotheses with the r largest e-values. For $r \in [K]$ and $j \in \{0, \dots, r-1\}$, let

$$e_{r,j} := \min_{I \subseteq [K] : |R_r \setminus I| = j} e_I; \tag{16}$$

since $|R_r \setminus I_\theta| = j$ means that R_r contains exactly j true discoveries, $e_{r,j}$ reflects the implausibility of there being exactly j true discoveries among R_r.

Lemma 9. *For each $r \in [K]$, $\alpha_{R_r} = \alpha^*_r$ and $\beta_{R_r} = M_r/r$, where*

$$M_r := \max_{j \in \{0, \dots, r-1\}} \frac{r - j}{e_{r,j}} \tag{17}$$

(with $c/0 := \infty$ for $c > 0$). Therefore,

$$\alpha^*_r = \begin{cases} M_r/r & \text{if } M_r \le r \\ \infty & \text{otherwise.} \end{cases}$$

Proof. The definitions immediately imply that $\beta_{R_r} = M_r/r$. It remains to prove that $\beta_R \ge \beta_{R_r}$ for every $R \subseteq [K]$ of size r. If $R \ne R_r$, there are $i_* < i^*$ such that $i_* \in R$ and $i^* \notin R$; set $R' := (R \setminus \{i_*\}) \cup \{i^*\}$. It suffices to show that

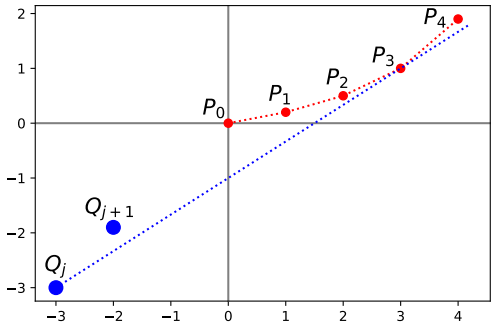

Figure 1: Geometry behind Algorithm 1 (reproduced from [5, Fig. 8])

$\beta_{R'} \le \beta_R$, since repeating such replacements transforms R into R_r. In other words, we need to check that for each $I \subseteq [K]$ there exists $I' \subseteq [K]$ such that $\mathrm{FDP}_I(R')/e_I \le \mathrm{FDP}_{I'}(R)/e_{I'}$. Let $I \subseteq [K]$. If $i^* \in I$ and $i_* \notin I$, we can set $I' := (I \setminus \{i^*\}) \cup \{i_*\}$, so that $|R \cap I'| = |R' \cap I|$ and $e_{I'} \le e_I$. Otherwise, $|R \cap I| \ge |R' \cap I|$, and we can set $I' := I$. □

Let us fix $r \in [K]$. For given sizes of $I \cap R_r$ and $I \setminus R_r$, the arithmetic mean e_I is smallest when these two sets consist of the indices of the smallest e-values in R_r and outside R_r, respectively. Therefore, for each $j \in \{0, \dots, r-1\}$, the minimum in (16) is attained at a set of the form $I = S_{r,j} \cup \{1, \dots, i\}$, where $S_{r,j} := \{K - r + 1, \dots, K - j\}$ is R_r without its j largest elements and $i \in \{0, \dots, K - r\}$. Following [5, Sect. 9], we set

$$
\begin{aligned}
s_k &:= e_1 + \dots + e_k, && k = 0, \dots, K, \\
\sigma_j &:= e_{K-r+1} + \dots + e_{K-j} = s_{K-j} - s_{K-r}, && j = 0, \dots, r-1
\end{aligned}
$$

(the dependence of σ_j on r is not reflected in our notation), so that

$$
e_{r,j} = \min_{i \in \{0, \dots, K-r\}} \frac{s_i + \sigma_j}{i + r - j}. \tag{18}
$$

Computing M_r directly using (17) and (18) takes time $O(K^2)$, which gives the overall time $O(K^3)$ for computing the frontier. Algorithm 1, which is a minor modification of [5, Algorithm 4], computes M_r in time $O(K)$, which gives the overall time $O(K^2)$.

The geometry behind Algorithm 1 is shown in Fig. 1. The coordinates of each of the points P_k, $k = 0, \dots, K - r$, are (k, s_k), and the coordinates of the point Q_j, where $j \in \{0, \dots, r-1\}$, are $-(r - j, \sigma_j)$. The slope of the straight line going from Q_j to P_i is

$$
a_i := \frac{s_i + \sigma_j}{i + r - j},
$$

Algorithm 1 Computing the frontier

Input: An increasing sequence of e-values $e_1 \le \cdots \le e_K$.
Output: The frontier $\alpha_0^*, \dots, \alpha_K^*$.

1: $s_0 := 0$
2: **for** $k = 1, \dots, K$
3: $\quad s_k := s_{k-1} + e_k$
4: $\alpha_0^* := 0$
5: **for** $r = 1, \dots, K$
6: $\quad k := K - r$
7: $\quad M := 0$
8: $\quad$ **for** $j = 0, \dots, r-1$
9: $\quad\quad \sigma_j := s_{K-j} - s_{K-r}$
10: $\quad\quad$ `slope` $:= \frac{s_k + \sigma_j}{k+r-j}$
11: $\quad\quad$ **for** $i = k-1, \dots, 0$
12: $\quad\quad\quad$ `new_slope` $:= \frac{s_i + \sigma_j}{i+r-j}$
13: $\quad\quad\quad$ **if** `new_slope` $>$ `slope` **break**
14: $\quad\quad\quad k := i$
15: $\quad\quad\quad$ `slope` $:=$ `new_slope`
16: $\quad\quad M := \max(M, (r-j)/$`slope`$)$ $\quad\triangleright\ c/0 := \infty$
17: $\quad$ **if** $M \le r$
18: $\quad\quad \alpha_r^* := M/r$
19: $\quad$ **else**
20: $\quad\quad \alpha_r^* := \infty$

and so $e_{r,j}$ is the smallest of these slopes over $i \in \{0, \dots, K-r\}$. Since the sequence $e_1, \dots, e_{K-r}$ is increasing, connecting the points $P_0, P_1, \dots$ in this order (see the red line in Fig. 1) gives us the graph of a convex function. Therefore, $e_{r,j}$ is the slope of the straight line (shown in blue) passing through Q_j and touching the red line from below. Let P_{k_j} be the leftmost point at which the blue line touches the red line. As i increases from 0 to $K-r$, the slope a_i strictly decreases until $i = k_j$ and does not decrease after that.

The command **break** in line 13 of Algorithm 1 means leaving the innermost loop, as in Python or R; in this context, it is equivalent to "go to line 16". The loops over i are executed in the decreasing order of i (and are not executed at all when $k = 0$). The variable k in line 14 is the index of the "current vertex" P_k; we start from the rightmost P_k in line 6 and then keep moving left. Figure 1 illustrates the execution of Algorithm 1 when $k = 3$, so that the current vertex is P_3. For each j, the loop over i moves the current vertex to the left as long as this does not make the line going from Q_j to the current vertex steeper, and line 16 uses the slope of the resulting line, which will be shown to be $e_{r,j}$.

Theorem 10. *Given the e-values (15), Algorithm 1 computes the frontier in time $O(K^2)$ and memory $O(K)$. If $r^* \in \{0, \dots, K\}$ maximizes $F(r, \alpha_r^*)$, the set R_{r^*} is a solution of the optimization problem (9), and its level is $\alpha_{R_{r^*}} = \alpha_{r^*}^*$.*

Proof. Fix $r \in [K]$ and $j \in \{0, \dots, r-1\}$, and suppose that $k \geq k_j$ when the iteration of the loop over j for this j starts in line 9. As long as $i \geq k_j$, we have `new_slope` $\leq$ `slope` in line 13, and if $k_j > 0$, the iteration with $i = k_j - 1$ ends with **break**. Therefore, in line 16 we have $k = k_j$ and `slope` $= e_{r,j}$. The assumption $k \geq k_j$ holds for $j = 0$, since then $k = K - r$, and for $j > 0$ it follows from $k_j \leq k_{j-1}$, which we will now check.

Let $j \in \{0, \dots, r-2\}$. First we check that the point Q_{j+1} lies on or above the blue line passing through Q_j, as shown in Fig. 1. The slope of the line going from Q_j to Q_{j+1} is $\sigma_j - \sigma_{j+1} = e_{K-j} \geq e_{K-r+1}$. If $k_j < K - r$, the slope of the blue line is at most $e_{k_j+1} \leq e_{K-r}$, since the blue line touches the red line from below at P_{k_j}. If $k_j = K - r$, the slope of the blue line is the average of $e_1, \dots, e_{K-j}$, which does not exceed e_{K-j}. In both cases, the line going from Q_j to Q_{j+1} is at least as steep as the blue line, and so Q_{j+1} lies on or above the blue line. Therefore, the line going from Q_{j+1} to P_{k_j} is not steeper than the blue line, and so $k_{j+1} \leq k_j$.

We can see that line 16 replaces M by $\max(M, (r-j)/e_{r,j})$, so that $M = M_r$ at the end of the loop over j, and α_r^* is computed correctly by Lemma 9. The variable k never increases inside the loop over j, and each iteration of the loop over i either decreases k or is the last one for the given j; therefore, the loop over j takes time $O(K)$. Computing $s_0, \dots, s_K$ takes time $O(K)$, and the memory used by the algorithm is $O(K)$. The last statement of the theorem follows from (14) and Lemma 9. □

If the e-values are not sorted, sorting them takes time $O(K \log K)$. For a fixed $\alpha \in (0, 1]$, the largest r such that $\alpha_r^* \leq \alpha$ is the largest number of rejections of the closed eBH procedure of [7, (7)] at level α. Algorithm 1 finds it for all α at once in time $O(K^2)$, whereas the algorithm in [7, Sect. 10] takes time $O(K^2 \log K)$ for one α in the worst case, although, as its authors note, its running time is often closer to $K \log K$ in practice; the algorithm mentioned in [3, Sect. 9.6.1] takes time $O(K^3)$.

5 Conclusion

This note discusses a procedure that assigns a level to every rejection set and so allows the researcher to choose the rejection set, together with its level, after seeing the data. The smallest level among the rejection sets of a given size is attained by rejecting the base hypotheses with the largest e-values, and the frontier can be computed in time $O(K^2)$.

A researcher who is interested in the false discovery rate and is happy to fix the level α in advance can simply use the closed eBH procedure at level α; see Sect. 2 (although she may miss much larger rejection sets available at a slightly larger level). The relative eBH procedure does not have any parameters, and its levels are valid in relative terms: the expected ratio of the false discovery proportion to the reported level is at most 1 (Proposition 3). It is an especially

natural choice when the levels that the researcher might report range widely (overwhelming evidence can be reported with a very small level).

Acknowledgments

I am grateful to Glenn Shafer for his encouragement and support over many years; his ideas have shaped my thinking about the foundations of probability, and about e-values in particular.

Claude Opus 5–5.5 and Fable 5.1–5.2 have been used in exploring proof ideas and developing proofs, which I reviewed carefully, and in checking the note. I take full responsibility for all claims and statements made here, including mathematical statements and their proofs.

References

[1] Peter D. Grünwald. Beyond Neyman–Pearson: E-values enable hypothesis testing with a data-driven alpha. *Proceedings of the National Academy of Sciences of the USA*, 121(39):e2302098121, 2024.

[2] Nick W. Koning. Post-hoc α hypothesis testing and the post-hoc p-value. Technical Report arXiv:2312.08040 [math.ST], arXiv.org e-Print archive, December 2025.

[3] Aaditya Ramdas and Ruodu Wang. Hypothesis testing with e-values. *Foundations and Trends in Statistics*, 1(1–2):1–390, 2025.

[4] Glenn Shafer and Vladimir Vovk. *Game-Theoretic Foundations for Probability and Finance*. Wiley, Hoboken, NJ, 2019.

[5] Vladimir Vovk and Ruodu Wang. Confidence and discoveries with e-values. *Statistical Science*, 38:329–354, 2023. Revised version: arXiv:1912.13292v5.

[6] Ruodu Wang and Aaditya Ramdas. False discovery rate control with e-values. *Journal of the Royal Statistical Society B*, 84:822–852, 2022.

[7] Ziyu Xu, Aldo Solari, Lasse Fischer, Rianne de Heide, Aaditya Ramdas, and Jelle Goeman. Bringing closure to false discovery rate control: A general principle for multiple testing. Technical Report arXiv:2509.02517 [stat.ME], arXiv.org e-Print archive, July 2026. This paper merges and subsumes arXiv:2504.11759 [stat.ME] and arXiv:2504.15946 [math.ST].