



HAL
open science

FDR control for Online Anomaly Detection

Etienne Krönert, Alain Célisse, Dalila Hattab

► **To cite this version:**

Etienne Krönert, Alain Célisse, Dalila Hattab. FDR control for Online Anomaly Detection. 2024. hal-04321622v2

HAL Id: hal-04321622

<https://hal.science/hal-04321622v2>

Preprint submitted on 16 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

FDR control for Online Anomaly Detection

Etienne Krönert¹, Alain Céliste² and Dalila Hattab¹

¹*FS Lab, Worldline, France, e-mail:*

etienne.kronert@worldline.com; dalila.hattab@worldline.com

²*SAMM, Paris 1 Panthéon-Sorbonne University, France, e-mail:*

alain.celisse@univ-paris1.fr

Abstract: A new online multiple testing procedure is described in the context of anomaly detection, which controls the False Discovery Rate (FDR). An accurate anomaly detector must control the false positive rate at a prescribed level while keeping the false negative rate as low as possible. However in the online context, such a constraint remains highly challenging due to the usual lack of FDR control: the online framework makes it impossible to use classical multiple testing approaches such as the Benjamini-Hochberg (BH) procedure, which would require knowing the entire time series. The developed strategy relies on exploiting the local control of the “modified FDR” (mFDR) criterion. It turns out that the local control of mFDR enables global control of the FDR over the full series up to additional modifications of the multiple testing procedures. An important ingredient in this control is the cardinality of the calibration dataset used to compute the empirical p-values. A dedicated strategy for tuning this parameter is designed for achieving the prescribed FDR control over the entire time series. The good statistical performance of the full strategy is analyzed by theoretical guarantees. Its practical behavior is assessed by several simulation experiments which support our conclusions.

MSC2020 subject classifications: Anomaly detection, Time series, Multiple testing.

1. Introduction

1.1. Context

By monitoring indicators in real time to check the system health, online anomaly detection aims at raising an alarm when abnormal patterns are detected [1, 20]. This contrasts with offline anomaly detection, which can only analyze historical data retrospectively. One motivation for automatic anomaly detection is to reduce the workload of operations teams by allowing them to prioritize their

efforts where necessary. This is typically accomplished through the use of statistical and machine learning models [10, 9, 7]. However, a poorly calibrated anomaly detector leads to alarm fatigue. An overwhelming number of alarms desensitizes the people tasked with responding to them, leading to missed or ignored alarms or delayed responses [11, 8]. One of the reasons for alarm fatigue is the high number of false positives which take time to manage [26, 19]. The main goal of the present work is to design a new (theoretically grounded) strategy allowing to control the number of false positives when performing online anomaly detection.

1.2. Related works

Due to the lack and cost of labeled data, anomaly detection is most often formulated in the context of unsupervised learning [10, 7, 25]. The unlabeled training set is used to learn atypicality scores that are difficult to interpret, such as those based on likelihood or distance [10]. Furthermore, due to the lack of labels, test data is rarely available to evaluate an anomaly detector. This makes empirical evaluation of anomaly detectors difficult [29]. Conformal Anomaly Detection (CAD) [18] is a method based on Conformal Prediction [2]. The goal of CAD is to give a probabilistic interpretation of the atypicality score using a calibration set and a conformal p -value estimator. CAD allows control of the expected number of false positives within a time period [18]. However, it does not provide control over the false discovery rate (FDR), i.e., the proportion of false positives among all detections, which is the focus of the present work. The Benjamini-Hochberg (BH) procedure [5, 6] is a multiple testing procedure that controls the proportion of false positives among rejections, the False Discovery Rate (FDR). BH has been used to control FDR for anomaly detection in the offline context [4, 21]. In online anomaly detection, however, the decision to classify a new observation as an anomaly must be made instantaneously. Therefore, BH cannot be applied as it requires the entire time series. The [16, 24, 31] introduce, in a different context than anomaly detection, an online multiple testing procedure based on alpha investing. The p -value is compared to an adaptive threshold which depends on previous decisions. However, this method is difficult to apply to conformal p -values due to its low power. In [28] the author suggests using the principle of local FDR. At each observation, a decision is taken depending on the estimation of the local FDR. Unfortunately, this approach relies on a Gaussian assumption that limits its generality. Approaches based on online breakpoint detection are particularly well suited for collective anomaly detection, thanks to methods such as CUSUM [22, 3], which can accumulate the atypicality scores of several consecutive observations, thus increasing the power of the test. There are numerous theoretical results concerning this method about the average run length [30] or the probability of false negatives and false positives [17]. However, throughout this paper, anomalies are assumed to be punctual. In this case, online breakpoint detectors do not improve detection power. For this reason, these detectors are not studied in this paper.

Controlling false positives for online anomaly detection remains a difficult task. The core of the paper is a theoretical analysis which aims to control FDR for online anomaly detection in the case where the observations are iid. Although this is a strong assumption, this is the only situation for which there are theoretical results is the case of offline multiple testing [21]. In particular two challenges arise with online anomaly detection. First, the true p -values are unknown and need to be estimated. Second, the decisions are made in an online context, whereas most multiple testing methods are done in the offline context. The main contributions are to address these challenges. More precisely it is established that it is possible to design online anomaly detectors controlling FDR of the time series.

- This paper investigates the relationship between FDR and the cardinality of the calibration set used to estimate p -values. To ensure FDR control, a calibration set cardinality tuning method is proposed.
- This paper describes a data-driven threshold for anomaly detection based on multiple testing in order to control the False Discovery Rate (FDR) of the entire time series. It explains how global control of the time series FDR can be obtained from local control of a modified version of the FDR criterion. Then, a modified version of the Benjamini-Hochberg procedure is suggested to achieve local control of the modified FDR.

1.3. Description of the paper

First, an anomaly detector controlling the global FDR is introduced in Section 2. Second, Section 3 deals with conditions on empirical p -values to ensure local control of FDR at a desired level α . Third, algorithms that allow global control of the FDR time series from local control of the modified FDR are developed in Section 4. Then, the anomaly detector is empirically studied in Section 5. Finally, it is evaluated against a competitor from the literature in Section 6.

2. FDR control for Online Anomaly Detection

2.1. Statistical framework

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, with Ω the set of all possible outcomes, \mathcal{F} a σ algebra on Ω , and \mathbb{P} a probability measure on \mathcal{F} . Assume a sequential realization of the random variables $(X_t)_{1 \leq t \leq T}$, where X_t takes values in a set \mathcal{X} for all t . $T \in \mathbb{N} \cup \{\infty\}$ is the size of the time series. Let \mathcal{P}_0 be a probability distribution, called the reference distribution, on the space \mathcal{X} . For each instant t , the observation X_t is called “normal” if $X_t \sim \mathcal{P}_0$. Otherwise X_t is an “anomaly”. The aim of an anomaly detector is to find all anomalies among the new observations along the time series $(X_t)_{t \geq 1}$. In *online* anomaly detection, for each time $t \geq 1$, a decision about the status of X_t is made based on *past observations*:

$(X_s)_{1 \leq s \leq t}$. This contrasts with the offline context, where the complete series is known when a decision is made.

Since the present goal is to use FDR, a natural strategy is to rephrase the online anomaly detection problem as a multiple testing problem: at each step $1 \leq t \leq T$, a statistical test is performed on the hypotheses:

$$\mathcal{H}_{0t}, \text{“}X_t \text{ is not an anomaly”} \quad \text{against} \quad \mathcal{H}_{1t} \text{“}X_t \text{ is an anomaly”}.$$

Each statistical test is performed by comparing an estimation of the p -value under \mathcal{H}_{0t} , noted \hat{p}_t , with a threshold $\hat{\varepsilon}_t$. Details about how to calculate these quantities are given in the next section. If the p -value \hat{p}_t is below the threshold $\hat{\varepsilon}_t$, then \mathcal{H}_{0t} is rejected, X_t is detected as an anomaly. In the online context, \hat{p}_t and $\hat{\varepsilon}_t$ are computed sequentially when X_t is observed without knowledge of the next observations. A classical criterion controlling the proportion of type I errors (False Positives) in the whole time series is FDR [5]. For a given data-driven threshold $\hat{\varepsilon}$ and a set of *estimated* p -values $\hat{p} = (\hat{p}_t)_{t \geq 1}$, the FDR criterion of the sequence from 1 to T is given by

$$FDR_1^T(\hat{\varepsilon}, \hat{p}) = \mathbb{E}[FDP_1^T(\hat{\varepsilon}, \hat{p})],$$

$$\text{with } FDP_1^T(\hat{\varepsilon}, \hat{p}) = \frac{\sum_{t \in \mathcal{H}_0} \mathbb{1}[\hat{p}_t \leq \hat{\varepsilon}_t]}{\sum_{t=1}^T \mathbb{1}[\hat{p}_t \leq \hat{\varepsilon}_t]},$$

with the convention that $0/0 = 0$. In the above expression, FDP_1^T denotes the *False Discovery Proportion* (FDP) of the time series from 1 to T . Also $\mathcal{H}_0 = \{t \in \mathbb{N}^* | \mathcal{H}_{0t} \text{ is true}\}$ is called the set of null hypotheses. Let us emphasize that the anomalies satisfy $\mathbb{1}[\hat{p}_t \leq \hat{\varepsilon}_t] = 1$. The main objective of the present work is to define a data-driven sequence $\hat{\varepsilon} : t \mapsto \hat{\varepsilon}_t$ such that, for a given control level $\alpha \in [0, 1]$, under weak assumptions on the sequence $\hat{p} : t \mapsto \hat{p}_t$,

$$FDR_1^T(\hat{\varepsilon}, \hat{p}) \leq \alpha \tag{2.1}$$

Note that the FDR_1^T must be controlled on the whole time series, even if decisions are made locally. The control is said exact when “ \leq ” is replaced with “ $=$ ”. Such a control would imply that for a level $\alpha = 0.1$, at most 10% of the detected anomalies along the entire time series are false positives. The detection power of the anomaly detector is measured by means of the *False Negative Rate* defined, for the sequence from 1 to T , by

$$FNR_1^T(\hat{\varepsilon}, \hat{p}) = \mathbb{E}[FNP_1^T(\hat{\varepsilon}, \hat{p})], \tag{2.2}$$

$$\text{with } FNP_1^T(\hat{\varepsilon}, \hat{p}) = \frac{\sum_{t \in \mathcal{H}_1} \mathbb{1}[\hat{p}_t > \hat{\varepsilon}_t]}{|\mathcal{H}_1|}, \tag{2.3}$$

where FNP_1^T denotes the *False Negative Proportion* (FNP) of the sequence from 1 to T and $\mathcal{H}_1 = \{t \in \mathbb{N}^* | \mathcal{H}_{1t} \text{ is true}\}$ is the set of alternative hypotheses.

The Benjamini-Hochberg procedure, introduced in [5] and described in Definition 1, allows the control of the at a desired level α , for a family of m p -values.

Definition 1 (Benjamini-Hochberg [5]). Let m be an integer and $\alpha \in [0, 1]$. Let $(p_i)_{1 \leq i \leq m} \in [0, 1]^m$ be a family of p -values. The Benjamini-Hochberg (BH) procedure, denoted by BH_α , is given by

- a data-driven threshold:

$$\varepsilon_{BH_\alpha} = \max\left\{\frac{\alpha k}{m}; p_{(k)} \leq \frac{\alpha k}{m}, k \in \llbracket 1, m \rrbracket\right\},$$

- a set of rejected hypotheses:

$$BH_\alpha((p_i; i \in \llbracket 1, m \rrbracket)) = \{i; p_i \leq \varepsilon_{BH_\alpha}, i \in \llbracket 1, m \rrbracket\}.$$

The intuition behind this procedure consists in drawing the ordered statistics $i \mapsto p_{(i)}$ (Figure 1) with $p_{(1)} \leq \dots \leq p_{(n)}$ and the straight line $i \mapsto \frac{\alpha i}{m}$. Then the BH procedure will reject all hypotheses corresponding to p -values smaller than the last crossing point between the straight line and the ordered p -values curve.




Figure 1: Illustration of the Benjamini-Hochberg procedure. The p -values are sorted in ascending order. The threshold is the largest p -value that is lower than $\alpha k/m$, where k is the rank of the p -value.

However, a crucial remark at this stage is that controlling FDR on the full time series is a very challenging task in the current online context for at least two reasons:

- The main existing approaches to control FDR are described in an “offline” framework where the whole series is observed first, and decisions are taken afterwards [5, 21]. This makes these approaches useless in the present context.
- The existing approaches designed in the online context [16, 31] are difficult to parameterize and hard to apply with empirical p -values. Realistic scenarios usually exclude the knowledge of the true probability distribution of the test statistic, leading to the approximation or estimation of the corresponding p -values in practice.

The new anomaly detector introduced in the next section provides solutions to these problems and allows global control of FDR.

2.2. Description of the method

Like many other anomaly detectors, our detector is based on the notions of atypicality score, p -value, and threshold. The novelty of our approach lies in the use of a data-driven threshold, which allows to control FDR at a desired α level. This threshold is calculated by applying the Benjamini-Hochberg procedure to a subseries of length m with a carefully chosen α' level. At the same time, the size of the calibration set must be correctly specified. More precisely, the new anomaly detector described in Algorithm 1 relies on:

1. **Atypicality score:** A score $a : \mathcal{X} \rightarrow \mathbb{R}$ is a function reflecting the atypicality of an observation X_t . To be more specific, the further \mathcal{P}_{X_t} is from \mathcal{P}_0 , the larger $a(X_t)$ is expected to be. It is often implemented using a non-conformity measure (NCM) [18], which measures how different a point X_t is from a training set \mathcal{X}^{train} , $a(X_t) = \bar{a}(X_t, \mathcal{X}^{train})$.
2. **p -value:** It is the probability of observing $a(X)$ higher than $a(X_t)$ if $X \sim \mathcal{P}_0$. It is estimated using the empirical p -value, by the following equation:

$$\hat{p}_e(s_t, \mathcal{S}_t^{cal}) = \frac{1}{|\mathcal{S}_t^{cal}|} \sum_{s \in \mathcal{S}_t^{cal}} \mathbb{1}[s_t > s] \quad (2.4)$$

Where $\mathcal{S}_t^{cal} = \{a(X_{u_1}), \dots, a(X_{u_n})\}$ is the calibration set with n data points. The p -value enables an interpretable criterion measuring how unlikely $X_t \sim \mathcal{P}_0$ is. The empirical p -value is chosen because it is agnostic to the true and unknown distribution, which is not the case for the Gaussian p -value estimator. The performance of the BH procedure applied to empirical p -values depends strongly on the cardinality of the calibration set ($|\mathcal{S}_t^{cal}| = n$). Section 3.2 investigates how to optimally choose this number.

3. **Detection threshold:** $\varepsilon_t \in [0, 1]$, it discriminates observations considered as abnormal from others. The observations considered as anomalies are X_t whose (estimated) p -value is smaller than the threshold ε_t . To control FDR of the entire time series, a data-driven threshold is computed using the m most recent p -values. This detection threshold is computed using a multiple testing procedure inspired by Benjamini-Hochberg (BH) and described in Section 4.2. This procedure requires the calculation of an α' value estimated from a training set or using heuristics. See Appendix A.1.1 for more details.

In the next sections, the design choices of Algorithm 1 are specified and justified by a theoretical study of the detector. Section 3 studies the local control of FDR when BH is applied to subseries of m empirical p -values. This allows to specify the choice of the calibration set cardinality used in Algorithm 1. Then, in Section 4, the FDR control on the complete time series is obtained from the local control of some modified FDR. This allows to specify the procedure for selecting α' in Algorithm 1.

Algorithm 1 FDR Control Online Anomaly Detection

Require: T length of the time series, $(X_t)_{1 \leq t \leq T}$ time series, α desired FDR, m subseries length, ν integer to tune the calibration set cardinality

Require: Either (Z_t) an historical dataset or π the proportion of anomalies

- 1: **if** historical dataset **then**
- 2: $\alpha' \leftarrow \arg \max_{\tilde{\alpha}} \left(\frac{\mu_{R_{\tilde{\alpha}}}^{**}}{\mu_{R_{\tilde{\alpha}}}} \tilde{\alpha} \leq \alpha \right)$ ▷ Estimate α' with training set
- 3: **else**
- 4: $\alpha' \leftarrow \frac{\alpha}{1 + \frac{1-\alpha}{m\pi}}$ ▷ Estimate α' with heuristics
- 5: **end if**
- 6: $n \leftarrow \nu \cdot \frac{m}{\alpha'} - 1$ ▷ Get the calibration set cardinality
- 7: **for** t in $[1, T]$ **do**
- 8: $s_t \leftarrow a(X_t)$
- 9: $\hat{p}_t \leftarrow \hat{p}_e(s_t, S_t^{cal})$ ▷ Compute empirical p -value
- 10: $\hat{\varepsilon}_t \leftarrow \hat{\varepsilon}_{BH_{\alpha'}}(\hat{p}_{t-m}, \dots, \hat{p}_t)$ ▷ Get the threshold using Benjamini-Hochberg
- 11: **if** $\hat{p}_t < \hat{\varepsilon}_t$ **then** ▷ Retrieve anomalies
- 12: $d_t = 1$
- 13: **else**
- 14: $d_t = 0$
- 15: **end if**
- 16: **end for**
- 17: **Output:** $(d_t)_{t=1}^T$ boolean list that represent the detected anomalies.

3. Local control of FDR with Empirical p -values

The goal here is to describe a strategy achieving the desired FDR control for a time series of length m when using empirical p -values. A motivating example is first introduced for emphasizing the issue in Section 3.1. Then a theoretical understanding is provided along Section 3.2 which results in a new solution which applies to independent empirical p -values. An extension to some dependency setups is then discussed in Section 3.3.

3.1. Motivating example

The purpose here is to further explore the effect of the calibration set cardinality on the actual FDR control when using empirical p -values. This gives us more insight on how to find mathematical solutions. Let us start by generating observations using two distributions. The reference distribution is $\mathcal{P}_0 = \mathcal{N}(0, 1)$ and the alternative distribution is $\mathcal{P}_1 = \mathcal{N}(4, 10^{-4})$. The anomalies are located in the right tail of the reference distribution. The length m of the signal is $m = 100$. The number of observations under \mathcal{P}_0 is $m_0 = 99$. The experiments have been repeated $K = 10^4$ times. Figure 2 displays the actual value of FDR as a function of the cardinality n of the calibration set $\{X_1, \dots, X_n\}$ used to compute the empirical p -values (see Eq. 2.4). One clearly see that FDR is not uniformly controlled at level $m_0/m\alpha$. However there exist particular values of n for which this level of control is nevertheless achieved. As long as n has become large enough ($n \geq 500$), repeated picks can be observed with a decreasing height as n grows.




Figure 2

3.2. FDR control: main results for i.i.d. p -values

The present section aims first to explain the shape of the curve shown in Figure 2. This will help to specify conditions on the calibration set cardinality to get FDR control for subseries of length m .

3.2.1. Compute FDR when BH is applied to empirical p -values

The main focus is now given to independent p -values. In what follows, the classical proof (see Section 1.1 in the Supplementary Materials) of the FDR control by the BH-procedure is revisited, leading to the next result. Its main merit is to provide the mathematical expression of the plain blue curve observed in Figure 2.

Theorem 1. *Let n be the cardinality of the calibration set and m be that of the set of tested hypotheses where $\{X_1, \dots, X_m\}$ denotes a set of random variables. Let $m_0 \leq m$ be the cardinality of the random variables from the reference distribution \mathcal{P}_0 . Let the empirical p -value be denoted, for any $i \in \llbracket 1, m \rrbracket$, by $\hat{p}_i = \hat{p}_e(X_i, \{Z_{i,1}, \dots, Z_{i,n}\})$, where the calibration set is $\{Z_{i,1}, \dots, Z_{i,n}\}$ and each $Z_{i,j} \sim \mathcal{P}_0$. Let the random variable $R(i)$ be the number of detections, when*

BH_α is applied on the family p_1^m , defined by $p_i' = 0$ and $p_j' = p_j$ for $j \neq i$, for some i in \mathcal{H}_0 . Then, for every $\alpha \in]0, 1]$, the FDR value over the sequence from 1 to m is given by

$$FDR_1^m(\hat{\varepsilon}_{BH_\alpha}, \hat{p}) = m_0 \sum_{k=1}^m \frac{\lfloor \frac{\alpha k n}{m} \rfloor + 1}{n+1} \frac{1}{k} \mathbb{P}(R(i) = k),$$

where $\hat{\varepsilon}_{BH_\alpha}$ denotes the BH_α threshold from Definition 1 when the BH-procedure is applied to the empirical p -values $\hat{p} = (\hat{p}_i; 1 \leq i \leq m)$.

The proof is delayed to Appendix B.1. In general it is not possible to compute the exact value of the FDR without knowing the distribution of the random variables $R(i)$. This is in contrast with the case of true p -values where $\mathbb{P}(p_i \leq \frac{\alpha k}{m}) = \frac{\alpha k}{m}$, whereas with empirical p -values $\mathbb{P}(\hat{p}_i \leq \frac{\alpha k}{m}) = \frac{\lfloor \frac{\alpha k n}{m} \rfloor + 1}{n+1}$, which prevents any simplification of the final bound. Nevertheless, this value still suggests a solution to circumvent this difficulty: requiring conditions on α , m , and n such that $\frac{\lfloor \frac{\alpha k n}{m} \rfloor + 1}{n+1} = \frac{\alpha k}{m}$, for all k . This is precisely the purpose of the next Corollary 1.

3.2.2. Tuning of the calibration set cardinality

Corollary 1. *Under the same notations and assumptions as Theorem 1, the next two results hold true.*

1. Assume that there exists an integer $1 \leq \nu$ such that $\frac{\nu m}{\alpha}$ is an integer. If the cardinality n of the calibration set satisfies $n = n_\nu - 1 = \nu m / \alpha - 1$, then

$$FDR_1^m(\hat{\varepsilon}_{BH_\alpha}, \hat{p}) = \frac{m_0 \alpha}{m}.$$

2. For every $\alpha \in]0, 1]$, assume that the cardinality of the calibration set satisfies $n = n_\nu - 1 = \lceil \frac{\nu m}{\alpha} \rceil - 1$, for any integer $\nu \geq 1$. Then,

$$\frac{n}{(n+1)} \frac{m_0 \alpha}{m} \leq FDR_1^m(\hat{\varepsilon}_{BH_\alpha}, \hat{p}) \leq \frac{m_0 \alpha}{m}.$$

The proof is moved to Appendix B.3. The first statement in Corollary 1 establishes that it is possible to recover the desired control of FDR at the exact prescribed level α , provided that the cardinality of the calibration set is equal to $n = \nu m / \alpha - 1$. The (mild) restriction on the values of α reflects that the empirical p values do not satisfy the super uniformity property. By contrast, the second statement gives the desired control at the level $\alpha m_0 / m$ by means of lower and upper bounds. In particular, the lower bound tells us that the FDR value cannot be smaller than the desired level $\alpha m_0 / m$ up to a multiplicative factor equal to $1 - 1/n$, which goes to 1 as n grows. For instance, with $\alpha = 0.1$

and $m = 100$, $n_\nu = 1000$ would result in $FDR_1^m(\hat{\varepsilon}_{BH_\alpha}, \hat{p}) \cdot m / (m_0 \alpha) \in [0.999, 1]$. This small lack of control is the price of allowing any value of $\alpha \in]0, 1]$. It is also important to remember that in anomaly detection, abnormal events are expected to be rare. As a consequence, $\frac{m_0}{m}$ is close to 1 and the actual FDR is close to the desired α . However, in situations where m_0/m could deviate too far from 1, it would be helpful to include an estimator of m_0/m .

3.3. Extension to dependent p -values

In Section 3.2, Corollary 1 states that FDR is controlled at a prescribed level α with empirical p -values, as long as the calibration set cardinality is correctly tuned, even if the super uniformity property is not satisfied. This result was established in the case where the calibration sets of the different p -values are independent. The purpose of the present section is to extend these results to dependent p -values.

A classical result for BH, established in [6], proves that FDR is upper bounded by $\alpha m_0/m$ provided the p -value family satisfies the PRDS and super uniformity properties. The PRDS property is a form of positive dependence between p -values where all pairwise p -value correlations are positive. It results that a small p -value for a given observation makes other p -values for all considered observations simultaneously small as well, and vice-versa [4]. One important example of p -values family that are PRDS are conformal p -values sharing a unique calibration set. Furthermore, another important result from the literature gives that FDR can also be lower bounded in the case where the calibration set is the same for all (conformal) p -values [21]. It turns out that all these results can be extended to empirical p -value estimators (see Eq 2.4) when the calibration set cardinality is the one discussed in Corollary 1. Furthermore, in the context of online anomaly detection, moving windows are classically used to capture and process the incoming data. In results, the calibration sets of the p -value family will overlap. Thus, FDR control is also studied in the case of overlapping calibration sets. In the rest of this section, the results obtained for these different cases are presented.

PRDS p -values For this extension, the concept of positive regression dependency on each one from a subset called PRDS [6] is recalled.

Definition 2 (PRDS property). *A family of p -values \hat{p}_1^m is PRDS on a set $I_0 \subset \{1, \dots, m\}$, if for any $i \in I_0$ and any increasing set A , the probability $\mathbb{P}[\hat{p}_1^m \in A | \hat{p}_i = u]$ is increasing in u .*

The classical result in [6] can be extended to empirical p -values with the same choice of calibration set cardinality as in Corollary 1. It ensures that FDR is upper bounded by $\frac{m_0}{m} \alpha$ in the case where the p -value family is PRDS.

Corollary 2 (Corollary of Theorem 1.2 in [6]). *Suppose the family of p -values \hat{p}_1^m is PRDS on the set \mathcal{H}_0 of true null hypotheses, and suppose that \hat{p}_1^m respects*

super uniformity on all thresholds that can result from BH

$$\forall i \in \mathcal{H}_0, \forall k \in \llbracket 1, m \rrbracket, \quad \mathbb{P}(\hat{p}_i < \frac{\alpha k}{m}) \leq \frac{\alpha k}{m},$$

Then, the FDR is upper-bounded by α

$$FDR(\hat{\varepsilon}_{BH}, \hat{p}) \leq \frac{m_0 \alpha}{m}.$$

The proof is not reproduced here since it is almost the same as the one of Theorem 1.2 in [6].

Unique calibration set Moreover, in the case where the calibration set is the same for all p -values, FDR can also be lower bounded. This result extends the one from in [21], to empirical p -values using a calibration set with cardinality tuned as proposed in Corollary 1.

Corollary 3 (Corollary of Theorem 3.4 in [21]). *Assuming the following conditions: Let n be the cardinality of the calibration set, m the cardinality of the active set and m_0 the number of normal observations. Let \mathcal{P}_0 be the reference distribution. Let Z_i for i in $\llbracket 1, m \rrbracket$ be independent random variables, following \mathcal{P}_0 . Let X_i for i in $\llbracket 1, m \rrbracket$ be independent random variables and independents of (Z_j) . There are exactly m_0 random variables following the \mathcal{P}_0 distribution. Let a be a scoring function. For all i in $\llbracket 1, m \rrbracket$, let \hat{p}_i be the empirical p -values associated with the random variables X_i and computed as follows, $\hat{p}_i = p_e(a(X_i), \{a(Z_1), \dots, a(Z_m)\}, a)$.*

If the cardinality of the calibration set is a multiple of $n = n_\nu = \nu m / \alpha - 1$, then FDR using BH_α on $(\hat{p}_i)_{1 \leq i \leq m}$ is equal to $\frac{m_0 \alpha}{m}$:

$$FDR_1^m(\hat{\varepsilon}_{BH_\alpha}, \hat{p}) = \frac{m_0 \alpha}{m}$$

The proof is not provided in this paper since it is too similar to the one of Theorem 3.4 in [21].

Overlapping calibration set In the context of online anomaly detection, the calibration sets of the p -value family will often overlap, since moving windows are classically used to capture and process the incoming data. To have a perfect control of FDR, an upper and lower bound is needed. According to Proposition 1, p -values with overlapping calibration sets are PRDS.

Proposition 1 (PRDS property for overlapping calibration sets). *Let X_i for i in $\llbracket 1, m \rrbracket$ be independent random variables. There are exactly m_0 random variables that follow the \mathcal{P}_0 distribution and thus belong to \mathcal{H}_0 . Let \mathbf{Z} be the random vector that combines all calibration sets, all elements of \mathbf{Z} are generated from \mathcal{P}_0 . The set of n indices defining the elements of the calibration set related to \hat{p}_i in*

\mathbf{Z} is noted as \mathcal{D}_i . The calibration set related to X_1 is noted $\mathbf{Z}_{\mathcal{D}_1} = (\mathbf{Z}_{i_1}, \dots, \mathbf{Z}_{i_n})$. For all i in $\llbracket 1, m \rrbracket$: $\hat{p}_i = p_e(X_i, \mathbf{Z}_{\mathcal{D}_i})$.

Under these conditions, the set of p -values is PRDS on \mathcal{H}_0 .

The proof of the proposition is in delayed to Appendix B.2. Since such p -values are PRDS, it gives an upper bound control of FDR using Corollary 4.

Corollary 4 (PRDS property for overlapping calibration sets). *Under the same conditions as Proposition 1 and the following condition on calibration set cardinality: $\exists \nu \geq 1, n = \nu \frac{m}{\alpha} - 1$:*

$$FDR(\hat{\varepsilon}_{BH}, \hat{p}) \leq \frac{m_0 \alpha}{m}.$$

For perfect FDR control, a lower bound is required in addition to the upper bound proven above. However, no theoretical results exist to compute the lower bound. Indeed, the existing proof in [21] did not extend to overlapping calibration sets. Therefore, the experiment in Appendix C.1 suggests to establish this lower bound empirically. Table 1 shows the p -value associated with hypothesis \mathcal{H}_0 “FDR is identical whatever the overlap proportion between calibration sets”. Please note that the p -values displayed in this table should not be confused with \hat{p}_t which is associated with the H_{0t} hypothesis “ X_t is generated by \mathcal{P}_0 ”. All tested hypotheses have a p -values greater than the (Bonferroni) threshold equal to 0.00625. According to this experiment, there are no significant differences in the resulting FDR between the different proportions of overlap in the calibration sets. This would suggest that considering overlapping calibration sets should not worsen the control of false positives and negatives too much.

TABLE 1
p-values resulting from permutation tests

n	249	250	499	500	749	750	999	1000
p -value of the test	0.300	0.0326	0.572	0.313	0.588	0.435	0.735	0.690

4. Global control of FDR over the full time series

In the previous Section 3, it was seen that the BH procedure provides control on FDR of subseries of length m if the calibration set cardinality is carefully tuned. While working with streaming time series data, the anomaly detection problem requires to control the FDR value of the full time series to ensure that the global false detection rate (FDR) remains under control at the end of the iterative process. The final criterion to control is then the global FDR criterion given by

$$FDR_1^\infty(\hat{\varepsilon}, \hat{p}),$$

where $\hat{\varepsilon} = (\hat{\varepsilon}_t)_{t \geq 1}$ denotes a sequence of data-driven thresholds, and \hat{p} stands for a sequence of empirical p -values (see Section 3 for further details). However,

the control of FDR of subseries does not imply control of the global FDR of the entire time series. This fact is illustrated in Figure 3. This figure is obtained by splitting a series of length 1000 into 10 subseries of length 100. The BH procedure is then applied to each subseries. FDP is calculated for each subseries and for each complete time series. The results are displayed in two boxplots. The left boxplot shows that BH_α provides the desired control at the level $\alpha = 10\%$ for FDR computed over the subseries of length m . However, the right boxplot clearly deviates from α , meaning that the actual FDR value for the full time series of length 1000 is much larger than α (more than 20% on average), leading to more false positives at the level of the full time series.




Figure 3: Comparison of the calculation of FDR computed locally on a subseries and FDR computed globally on the entire time series using the Benjamini-Hochberg procedure applied to a subseries.

This implies that controlling FDR on subseries is not a good criterion. However, in contrast to the global objective, anomaly detection still requires a decision at each time step, i.e., for each new observation, without knowing what the next ones will be. This justifies the need for another local criterion that is used to make a decision at each iteration, leading to the sequence of data-driven thresholds $\hat{\varepsilon} = (\hat{\varepsilon}_t)_{t \geq 1}$ computed over subseries. An additional difficulty arises from the connection that must be made between this local criterion and FDR of the full time series.

To this end, Section 4.1 explores the connection between FDR of the full time series and the so-called modified-FDR (mFDR) of the subseries. In particular, it turns out that controlling the mFDR value for all subseries of a given length m yields the desired FDR control for the full time series. Section 4.2 then explains how the classical BH-procedure can be modified to get the mFDR control for subseries of length m .

4.1. How mFDR can help in controlling FDR of the full time series?

Mixture model and time series In this section one assumes that the time series is generated from a mixture process between a reference distribution \mathcal{P}_0 and an alternative distribution \mathcal{P}_1 . The anomaly locations are supposed to be independent and generated by a Bernoulli distribution. This is a common assumption usually made in the literature [15, 27] for simplicity.

Definition 3 (Time series process with anomalies). *Let $\pi \in [0, 1]$ be the anomaly proportion and \mathcal{P}_0 and \mathcal{P}_1 be two probability distributions on the observation domain \mathcal{X} . \mathcal{P}_0 is the reference distribution and \mathcal{P}_1 denotes the alternative distribution. The generation process of a time series containing anomalies $(A_t)_{t \geq 0}$ is given, for each $t \geq 0$, by*

- $A_t \sim B(\pi)$ (Bernoulli distribution)
- if $A_t = 0$, then $X_t \sim \mathcal{P}_0$.
- if $A_t = 1$, then $X_t \sim \mathcal{P}_1$.

Moreover, given the above scheme, $(X_t)_{t \geq 0}$ is a random process with independent and identically distributed random variables $X_t \sim (1 - \pi)\mathcal{P}_0 + \pi\mathcal{P}_1$.

This definition details how anomalies are generated. In particular, it assumes that anomalies are independent from each other. It should be noted that this does not prevent us from observing a sequence of successive anomalies along the time series. However, this scheme differs substantially from the case analyzed by [14], where specific patterns of successive anomalies are searched for.

Threshold on subseries Let us start with a subseries of length m in which each observation is summarized by its corresponding empirical p -value, and let us assume that there exists a function $f_m : [0, 1]^m \rightarrow [0, 1]$ that maps a set of m empirical p -values onto a real-valued variable. This random variable corresponds to the data-driven threshold. This function f_m is called a *local threshold* function because it returns a threshold when applied to a subset of length m . Given the notation above, the threshold sequence $\hat{\varepsilon}_o = (\hat{\varepsilon}_{o,t})_t$ can be defined as follows.

Threshold on (overlapping) subseries: $\hat{\varepsilon}_o : t \mapsto \hat{\varepsilon}_{o,t}$ is given by

$$\forall t \geq m, \quad \hat{\varepsilon}_{o,t} = f_m(\hat{p}_{t-m+1}, \dots, \hat{p}_t). \quad (4.1)$$

Figure 4 illustrates this data-driven threshold. Because two successive subseries differ from each other by two observations, the thresholds are different at each time step. It shows that the variables $\hat{\varepsilon}_{o,t}$ and $\hat{\varepsilon}_{o,t-1}$ are dependent because they share $m - 1$ common observations. But for any t , $\hat{\varepsilon}_{o,t}$ is independent from $\hat{\varepsilon}_{o,t-m-1}$. This can be reformulated as $\hat{\varepsilon}_{o,t_1}, \hat{\varepsilon}_{o,t_2}$ are independent if and only if $|t_1 - t_2| > m$. The local threshold function is said to be permutation invariant




Figure 4: Illustration of a threshold on (overlapping) subseries.

if its values does not depends on the p -value order. For any ρ permutation of $\llbracket 1, m \rrbracket$,

$$f_m(p_{\rho(1)}, \dots, p_{\rho(m)}) = f_m(p_1, \dots, p_m). \quad (4.2)$$

In Section 1.2 of the Supplementary Materials, a different scheme for defining a data-driven threshold from subseries is presented. The time series is divided into disjoint subseries and the same threshold is applied to all points belonging to that subseries. The term “overlapping subseries” is used to distinguish it from the disjoint case. In the present online anomaly detection context, considering the overlapping case sounds more convenient since the detection threshold can be updated at each time step (as soon as a new observation is given), which makes the anomaly detector more versatile. However, disjoint are easier to study.

FDR control with threshold on (overlapping) subseries As illustrated in Figure 3, controlling FDR on each subseries of length m (locally) is not equivalent to controlling FDR (globally) on the full time series. However in online anomaly detection, a decision has to be made at each time step regarding the potential anomalous status of each new observation. This requires a criterion to be controlled locally (on subseries) in such a way that the resulting global FDR value (the one of the full time series) can be proven to be controlled at the desired level α . This requirement for a local criterion justifies the introduction of the modified FDR criterion, denoted by mFDR [31, 13], which is defined as follows.

Definition 4 (mFDR). *With the previous notations, the mFDR expression of the subseries from $t - m + 1$ to t is given by*

$$mFDR_{t-m+1}^t(\hat{\epsilon}, \hat{p}) = \frac{\mathbb{E} \left[\sum_{u \in \mathcal{H}_0, t-m+1 \leq u \leq t} \mathbb{1}[\hat{p}_u \leq \hat{\epsilon}_u] \right]}{\mathbb{E} \left[\sum_{u=t-m+1}^t \mathbb{1}[\hat{p}_u \leq \hat{\epsilon}_u] \right]},$$

where $\hat{\epsilon} = (\hat{\epsilon}_u)_{t-m+1 \leq u \leq t}$ denotes a sequence of thresholds, \hat{p} is a sequence of empirical p -values evaluated at each observation of the subseries from $t - m + 1$ to t .

Mathematically the difference between mFDR and FDR is that the expectation is no longer on the ratio but independently on the numerator and the

denominator. The main interest for mFDR is clarified by Theorem 2, which establishes that the control of the latter at the α level provides a global control of FDR at the same level under simple conditions.

Theorem 2 (Global FDP control using local threshold). *Let $\hat{p} = (\hat{p}_t)_{t \geq 1}$ be the p -value random process, for a time series that follows the scheme detailed in Definition 3. Let $\hat{\mathbf{P}} = (\hat{\mathbf{P}}_{t,k})_{t \geq 1, 1 \leq k \leq m}$ a process of p -values vectors, such that $\hat{\mathbf{P}}_{t,m} = \hat{p}_t$. Assume that $\hat{\varepsilon}_o : t \mapsto \varepsilon_{o,t}$ is given by $\hat{\varepsilon}_{o,t} = f_m(\hat{\mathbf{P}}_t)$, for any $t \geq 1$, with $f_m : [0, 1]^m \rightarrow [0, 1]$ permutation invariant (see Eq. (4.2)). Let us also assume that there exists n such that $|t_1 - t_2| > n$ implies that \hat{p}_{t_1} and \hat{p}_{t_2} are independent, and $|t_1 - t_2| > n + m$ implies that $\hat{\mathbf{P}}_{t_1}$ and $\hat{\mathbf{P}}_{t_2}$ are independent. For all t , $\hat{\mathbf{P}}_t = (\hat{\mathbf{P}}_{t,1}, \dots, \hat{\mathbf{P}}_{t,m})$ is exchangeable. Then, the global FDR (and FDP) value of the full (infinite) time series is equal to the local mFDR value of the any subseries of length m computed at time $t \in \mathbb{N}^*$. More specifically*

$$FDP_1^\infty(\hat{\varepsilon}_o, \hat{p}) = FDR_1^\infty(\hat{\varepsilon}_o, \hat{p}) = mFDR_{t-m+1}^t(\hat{\varepsilon}_o, \hat{\mathbf{P}}_t) = mFDR_1^m(\hat{\varepsilon}_o, \hat{\mathbf{P}}_m)$$

The proof is postponed to Appendix B.4. An important consequence is that the desired control for FDR of the full (infinite) time series at level α can be achieved if one can control the successive mFDR of all (shifted) subseries of length m at level α . This point is not obvious at all and constitutes the main concern of Section 4.2 where a new multiple testing procedure is designed to yield the desired control of the mFDR criterion. The main limitation of Theorem 2 is the requirement that f_m has to be permutation invariant. Note that this property holds for the BH procedure, for example. Let us also mention that the empirical p -values computed as $\hat{p}_t = \hat{p}\text{-value}(a(X_t), \{a(X_{t-n}), \dots, a(X_{t-1})\})$ and vector p -values defined as $\hat{\mathbf{P}}_t = (\hat{p}_{t-m+1}, \dots, \hat{p}_t)$ actually satisfy the requirements of Theorem 2 regarding the independence and the stationarity.

4.2. Modified BH-procedure and mFDR control

As shown in Section 4.1, controlling the FDR value of the full time series is equivalent to controlling mFDR of subseries of length m . The strategy then consists in controlling the mFDR criterion of all successive subseries of length m along the full time series at level α . The main challenge addressed in the present section is to design a new multiple testing procedure that controls the local mFDR criterion at a prescribed level α .

In Section 4.2.1, it is shown that applying BH_α to a subseries of size m does not result in mFDR being controlled to the desired level α . However, mFDR is upper bounded by a quantity that depends on the level of control α used in the BH procedure. This leads to a strategy of using the BH procedure, but with a different, well-chosen α' , to allow control of mFDR at the desired level α . This so-called modified BH-procedure is introduced in Section 4.2.2.

4.2.1. mFDR control with the BH-procedure

The next Proposition 2 establishes the actual mFDR level achieved by the BH-procedure.

Proposition 2. *Let $(X_i)_1^m$ satisfy the requirements detailed by Definition 3, with an abnormality proportion π . m_0 is the random variable representing the number of data points generated by \mathcal{P}_0 . Let $p = (p_1, \dots, p_m)$ be the associated p -values. Let α belong to $[0, 1]$. Suppose there are integers ν and n such that $n = \nu \frac{m}{\alpha} - 1$. Let $R_{\alpha,1}^m$ be the random variable representing the number of rejections when BH_α is applied to p . Suppose one of the following three statements is true:*

1. (p_1, \dots, p_m) are true p -values that is, for any $1 \leq i \leq m$, $p_i = \mathbb{P}_{X \sim \mathcal{P}_0}(a(X) \geq a(X_i))$.
2. (p_1, \dots, p_m) are empirical p -values with independent calibration sets of cardinality n , $p_i = \hat{p}_e(a(X_i), \{a(Z_{i,1}), \dots, a(Z_{i,n})\})$. With $Z_{i,j}$ are iid random variables generated by \mathcal{P}_0 .

In these two cases, let choose i in \mathcal{H}_0 , the sequence $(p'_j)_{1 \leq j \leq m}$ is defined by $p'_i = 0$ and $p'_j = p_j$. Let $R_{\alpha,1}^{*,m}$ define the number of rejections when applying BH_α on p'

3. (p_1, \dots, p_m) are empirical p -values with a unique calibration set of cardinality n , $p_i = \hat{p}_e(a(X_i), \{a(Z_1), \dots, a(Z_n)\})$. With Z_j are iid random variables generated by \mathcal{P}_0 .

Let choose i in \mathcal{H}_0 , the sequence $(p'_j)_{1 \leq j \leq m}$ is defined by $p'_i = 0$ and $p'_j = p_j - \frac{1}{n} \mathbb{1}[p_j < p_i]$. Let $R_{\alpha,1}^{*,m}$ define the number of rejections when applying BH_α on p' .

Then, applying BH_α to the p -values $(p_i)_{1 \leq i \leq m}$ leads to

$$mFDR_1^m(p) = \alpha \frac{\mathbb{E} \left[\frac{|\mathcal{H}_0|}{m} R_{1,\alpha}^{*,m} \right]}{\mathbb{E} R_{1,\alpha}^m} \leq \alpha \frac{\mathbb{E} R_{1,\alpha}^{*,m}}{\mathbb{E} R_{1,\alpha}^m},$$

Furthermore, if $\mathbb{E}[R_{\alpha,1}^{*,m} | m_0]$ is decreasing:

$$mFDR_1^m(p) \leq \alpha(1 - \pi) \frac{\mathbb{E} R_{1,\alpha}^{*,m}}{\mathbb{E} R_{1,\alpha}^m},$$

The proof is moved to Appendix B.5. The upper bound is generally not too large since, the proportion of normal data $\frac{|\mathcal{H}_0|}{m}$ is close to 1. The decreasing property of $\mathbb{E}[R_{\alpha,1}^{*,m} | m_0]$ is verified as soon as the data points generated by \mathcal{P}_1 have a higher probability of being detected as anomalies than those generated by \mathcal{P}_0 . This is verified whenever the atypicality score is chosen correctly. Then, to get control of mFDR, it is sufficient to find α' such that $\frac{\mathbb{E} R_{\alpha'}^{*,m}}{\mathbb{E} R_{\alpha'}^m} = \alpha$, and then

apply BH with level α' to get the data driven threshold $\hat{\varepsilon}_{BH_{\alpha'}}$. This method is called “modified Benjamini-Hochberg” and is described in Definition 5. There is no theoretical result in the case where the p -values p_1, \dots, p_m are PRDS. In particular, there are no theoretical results if the calibration sets overlap as described in Section 3.3.

4.2.2. Modified BH

From the previous Section 4.2.1, it is now possible to suggest and analyze the new modified BH-procedure (mBH in the sequel).

Definition 5 (Modified BH-procedure (mBH)). *Let m be an integer and $\alpha \in [0, 1]$. Let us introduce the level $\alpha' = \alpha'(\alpha)$ an estimate of $\arg \max_{\tilde{\alpha}} \left\{ \frac{\mathbb{E}[R_{1,\tilde{\alpha}}^{*,m}]}{\mathbb{E}[R_{1,\tilde{\alpha}}^m]} \tilde{\alpha} \leq \alpha \right\}$, which can be given by some procedure. Then the modified BH-procedure, denoted by mBH_{α} , is given for all true p -values $(p_1, \dots, p_m) \in [0, 1]^m$ by,*

$$mBH_{\alpha}(p_1, \dots, p_m) = BH_{\alpha'}(p_1, \dots, p_m).$$

The related mBH_{α} threshold at level α is defined as

$$\varepsilon_{mBH_{\alpha}} = \varepsilon_{BH_{\alpha'}},$$

when computed with true p -values, and $\hat{\varepsilon}_{mBH_{\alpha}} = \hat{\varepsilon}_{BH_{\alpha'}}$ when used with empirical p -values.

The above definition defines the mBH_{α} in terms of the BH-procedure by simply changing the level of control α' . There are several ways to estimate α' . It can be done using heuristics, which give a formula to directly calculate α' as a function of α , m and π . However, the heuristics rely on assumptions that are difficult to verify and on the proportion of anomalies, which is often unknown. Another solution is to estimate $\mathbb{E}[R_{1,\tilde{\alpha}}^{*,m}]$ and $\mathbb{E}[R_{1,\tilde{\alpha}}^m]$ on a grid of $\tilde{\alpha}$ using a training set. These two methods are described in more detail in Appendix A.1.

Combining the results of Theorem 2 and Proposition 2, the procedure mBH allows to get the control of FDR of the complete series at the desired level α . This property is described in Corollary 5.

Corollary 5 (Control of FDR using mBH). *Under the same notations and assumptions as Theorem 2. Let m and n be two integers. Suppose one of the following statements if true:*

1. For all t p -values of $\hat{\mathbf{P}}_t$ are true p -values.
2. For all t p -values of $\hat{\mathbf{P}}_t$ are empirical p -values with independent calibration sets of cardinality n .

Let i be a true null hypothesis in $\hat{\mathbf{P}}_t$, the sequence $\hat{\mathbf{P}}'_t$ is defined by $\hat{P}'_{t,i} = 0$ and $\hat{P}'_{t,j} = \hat{P}_{t,j}$ for $j \neq i$. Let $R_{\alpha,1}^{*,m}$ be the number of rejections when applying BH_{α} to $\hat{\mathbf{P}}'_t$.

3. For all t , p -values of $\hat{\mathbf{P}}_t$ are empirical p -values with a unique calibration set of cardinality n .

Let i be a true null hypothesis in $\hat{\mathbf{P}}_t$, the sequence $\hat{\mathbf{P}}'_t$ is defined by $\hat{P}'_{t,i} = 0$ and $\hat{P}'_{t,j} = \hat{P}_{t,j} - \frac{1}{n}\mathbb{1}[\hat{P}_{t,j} < \hat{P}_{t,i}]$ for $j \neq i$. Let $R_{\alpha,1}^{*,m}$ be the number of rejections when applying BH_α to $\hat{\mathbf{P}}'_t$.

Suppose there are ν and α' such that:

$$\frac{\mathbb{E}R_{1,\alpha'}^{*,m}}{\mathbb{E}R_{1,\alpha'}^m} \alpha' = \alpha$$

Then, FDR of the entire time series can be controlled at level α by using $\hat{\varepsilon}_{BH_{\alpha'},t} = f_m(\hat{\mathbf{P}}_t)$

$$FDR_1^\infty(\hat{\varepsilon}_{BH_{\alpha'},\hat{p}}) \leq (1 - \pi)\alpha \quad (4.3)$$

The proof is moved to Appendix B.7. Corollary 5 describes the properties that the sequence of p -values \hat{p} and the sequence of vector p -values $\hat{\mathbf{P}}$ should satisfy in order to get the control on FDR. Theorem 3 gives practical ways to compute \hat{p} and $\hat{\mathbf{P}}$ that satisfy these requirements. It gives the ingredient to build an anomaly detector controlling FDR at a desired level α . Our Algorithm 1 implements this result.

Theorem 3 (Global FDR control using mBH_α). *Let (X_t) be a mixture process introduced in Definition 3. Let $\alpha \in [0, 1]$ be the desired FDR level for the full time series. Let m and n be integers.*

If one of the three statements is true:

1. $\forall t, \hat{p}_t = 1 - \mathbb{P}_{X \sim \mathcal{P}_0}(a(X) > a(X_t))$
2. $\forall t, \hat{p}_t = \hat{p}_e(a(X_t), \{a(Z_{t,1}), \dots, a(Z_{t,n})\})$ with $Z_{t,i} \sim \mathcal{P}_0$

In this two cases, let $\hat{\mathbf{P}}_t = (\hat{p}_{t-m+1}, \dots, \hat{p}_t)$. Let choose i in $\mathcal{H}_0 \cap \llbracket t-m+1, t \rrbracket$, the sequence $\hat{\mathbf{P}}'_t$ is defined by $\hat{P}'_{t,i} = 0$ and $\hat{P}'_{t,j} = \hat{P}_{t,j}$ for $j \neq i$. Let $R_{\alpha,1}^{*,m}$ be the number of rejections when applying BH_α to $\hat{\mathbf{P}}'_t$.

3. $\forall k \in \llbracket 1, m \rrbracket, (\hat{\mathbf{P}}_t)_{t,k} = \hat{p}_e(a(X_{t-m+k}), \mathcal{S}^{cal})$, with the calibration set

$$\mathcal{S}^{cal} = \{(1 - A_{t-n+1-m})a(X_{t-n+1-m}) + A_{t-n+1-m}a(Z_{t,1}), \dots, (1 - A_{t-m})a(X_{t-m}) + A_{t-m}a(Z_{t,n})\}$$

with $Z_{t,i} \sim \mathcal{P}_0$ and $\hat{p}_t = \hat{\mathbf{P}}_{t,m}$.

Let choose i in $\mathcal{H}_0 \cap \llbracket t-m+1, t \rrbracket$, the sequence $\hat{\mathbf{P}}'_t$ is defined by $\hat{P}'_{t,i} = 0$ and $\hat{P}'_{t,j} = \hat{P}_{t,j} - \frac{1}{n}\mathbb{1}[\hat{P}_{t,j} < \hat{P}_{t,i}]$ for $j \neq i$. Let $R_{\alpha,1}^{*,m}$ define the number of rejections when applying BH_α to $\hat{\mathbf{P}}'_t$.

Suppose there are α' and ν such that:

$$\frac{\mathbb{E}R_{1,\alpha'}^{*,m}}{\mathbb{E}R_{1,\alpha'}^m}\alpha' \leq \alpha \quad \text{and} \quad n = \nu m/\alpha' - 1.$$

Then, with $\hat{\varepsilon}_{BH_{\alpha'},t} = \hat{\varepsilon}_{BH_{\alpha'}}(\hat{\mathbf{P}})_t$:

$$FDR_1^\infty(\hat{\varepsilon}_{BH_{\alpha'}}, \hat{p}) \leq (1 - \pi)\alpha.$$

The proof is deferred to Appendix B.8. The main merit of Theorem 3 is to establish the actual level of control for the global FDR of the full time series depending on the type of empirical p -value used in the anomaly detection process. The last type of empirical p -values is (almost) the one used in practice in the present work. More specifically, Section 6.2 describes empirical p -values based on a “Sliding Calibration Set”.

5. Empirical results

In this section, properties related to the anomaly detector are empirically evaluated to illustrate and complete the various theoretical results set out in the previous sections. First, from Corollary 1, it is known that when BH is applied to empirical p -value subseries, there are specific values for the calibration set cardinality at which FDR is controlled to the desired value α . In Section 5.1, to illustrate this property and to further study the impact of cardinality, FDR and FNR are calculated as a function of calibration set cardinality in several scenarios. Then, according to Proposition 2, the mBH procedure (unlike the BH procedure) allows to control mFDR of subseries. However, this hides the fact that the mBH procedure requires the estimation of a parameter α' . In case this parameter is estimated by the heuristics described in Appendix A.1.1, it is necessary to check the assumptions **Heuristic** and **Power**. These assumptions are difficult to ensure in practice. In Section 5.2, the evaluation is done on simulated data where the level of atypicality of the anomalies (affecting the power of the detector) varies from one sample to another. Different scenarios are tested to check if the mFDR control holds. Finally, Theorem 2 gives asymptotic FDR control over the full time series. But there is no result about the speed of convergence, which is necessary when applied to finite time series. In Section 5.3, FDR for the full time series is calculated in different situations, as a function of the time series size. It is possible to determine when the entire series reaches control of FDR.

5.1. Assessing the local control of FDR with empirical p -values

The purpose of the present section is to compute the actual FDR value when empirical p -values are used instead of true ones. The question raised here is

whether FDR of the subseries is really controlled at a given level α . The empirical results must be compared with the theoretical FDR expression established in Theorem 1.

Experiment description Two scenarios have been considered to explore how much the thickness of the distribution tails can affect the results.

1. **Thin tails:**

The reference probability distribution is $\mathcal{P}_0 = \mathcal{N}(0, 1)$ for normal observations and $\mathcal{P}_1 = \delta_{\Delta_{\mathcal{N}}}$ for anomalies, where $\Delta_{\mathcal{N}} \in \mathbb{R}$ is a parameter encoding the strength of the shift. Here $\delta_{\Delta_{\mathcal{N}}}$ denotes the Dirac measure such that $\delta_{\Delta_{\mathcal{N}}}(z) = 1$ if $z = \Delta_{\mathcal{N}}$ and 0 otherwise.

2. **Thick tails:**

$\mathcal{P}_0 = \mathcal{T}(5)$ is a Student probability distribution with 5 degrees of freedom and $\mathcal{P}_1 = \delta_{\Delta_t}$ denotes the alternative distribution of anomalies, where $\Delta_{\mathcal{N}} \in \mathbb{R}$ is a parameter encoding the strength of the shift.

Regarding the value of the shift strength in Scenarios 1 and 2, two values of $\Delta_{\mathcal{N}}$ have been considered 3.5 and 4. The values of $\Delta_{\mathcal{T}}$ have been chosen such that

$$\mathbb{P}_{X \sim \mathcal{N}(0,1)}(X > \Delta_{\mathcal{N}}) = \mathbb{P}_{X \sim \mathcal{T}(5)}(X > \Delta_t)$$

for each choice of $\Delta_{\mathcal{N}}$. This avoids any bias in the comparison of the detection power of the considered strategy depending on the ongoing scenario.

Different cardinalities have been considered for the calibration set following the mathematical expression


$$n \in \{k \cdot 10, k \in \llbracket 1, 200 \rrbracket\} \cup \{\ell \cdot 10 - 1, \ell \in \llbracket 1, 200 \rrbracket\}.$$

In particular, all integers between 10 and 2 000 are explored with a step size equal to 10, as well as all integers between 9 and 1 999 with a step size of 10. This choice is justified by the particular expression of FDR value provided by Corollary 1. All the n elements of the calibration set are generated from the reference distribution that is, $\{Z_1, \dots, Z_n\} \sim \mathcal{P}_0$. Concerning the m observations corresponding to the tested hypotheses $\{X_1, \dots, X_m\}$, the m_0 normal observations are generated according to \mathcal{P}_0 and the m_1 anomalies from \mathcal{P}_1 . Here $m = 100$ and $m_0 = 99$. Each simulation condition has been repeated $B = 10^4$ times. For each repetition $1 \leq b \leq B$, the observations are indexed by b such that $X_{b,j} \sim \mathcal{P}_1$ for each $j \in \llbracket 1, m_1 \rrbracket$, and $X_{b,j} \sim \mathcal{P}_0$ for $j \in \llbracket m_1 + 1, m \rrbracket$. In the present scenarios, the anomalies are all in the right tail of the reference probability distribution. Therefore, the empirical p -values are computed according to Eq. 2.4 with the scoring function $a(x) = x$. For each repetition $1 \leq b \leq B$,


$$\forall 1 \leq j \leq m, \quad \hat{p}_{b,j} = \hat{p}_e(X_{b,j}, \{Z_{b,1}, \dots, Z_{b,n}\}).$$

After computing the empirical p -values, the BH_α procedure (see Definition 1) is applied to this sequence. Then FDP and FNP of the sequence are calculated.


The results are averaged over all sequences to get an estimation of FDR and FNR.



(a) FDR depending on n with 4 sigmas anomalies (b) FNR depending on n with 4 sigmas anomalies



(c) FDR depending on n with 3.5 sigmas anomalies (d) FNR depending on n with 3.5 sigmas anomalies



(e) FDR depending on n with 4 sigmas anomalies and $m = 150$ (f) FNR depending on n with 4 sigmas anomalies and $m = 150$

Figure 5: Effect of calibration set cardinality and abnormality score on the FDR control and FNR

Results and analysis Figure 5 displays FDR (left panel) and FNR (right panel) as a function of the calibration set cardinality for the two scenarios (Gaussian and Student) described in the previous paragraph. The blue (respectively orange) curve corresponds to the Gaussian (resp. Student) reference distribution. The horizontal line is the desired FDR control level $\alpha = 0.1$. Figures 5a and 5b are obtained with $\Delta_{\mathcal{N}} = 4$, while Figures 5c and 5d are obtained with $\Delta_{\mathcal{N}} = 3.5$.

According to these plots, the behavior of both FDR and FNR does not exhibit a strong dependence with respect to the reference probability distribution. The results are very close for both the Gaussian and Student distributions. As shown in Figures 5a and 5c, FDR control is achieved at the prescribed level for particular values of the calibration set cardinality. These values coincide with the ones given by Theorem 1, which are of the form $\nu\alpha/m - 1 = \nu 10^3 - 1$ for an integer ν . Notice that the FNR curve rises sharply from 1 to $n = 999$. This reflects that although the FDR value becomes (close to) optimal as n increases from 1 to $n = 999$, the proportion of false negatives increases at the same time, leading to suboptimal statistical performance (because of too many false negatives). Fortunately, a larger cardinality n of the calibration set, e.g. $n = 1999$, would greatly improve the results at the cost of a larger calibration set, which also increases the computational cost.

The FDR value seems not to depend on the strength of the distribution shift Δ , as illustrated by Figures 5a and 5c. Let us mention that the focus of the expectation does not say anything about the probability distribution of FDP, which can be influenced by the strength of the shift. In contrast, the comparison of Figures 5b and 5d clearly shows the impact of the shift strength on the FNR value. As the shift strength becomes lower, anomalies become more difficult to detect, inflating FNR.

According to Figure 5a, FDR can be controlled for particular values of the calibration set cardinalities that are not predicted by Corollary 1. For example, $n = 1499$ controls FDR at 0.1, while FDR is controlled for $n = 999$ and $n = 1999$ according to Corollary 1. Furthermore, complementary experiments (summarized by Figure 5 in Section 3.1 of the Supplementary Materials) illustrate that these particular values of n depend on the number of anomalies m_1 , in contrast to the values derived from Corollary 1. Using Theorem 1, their existence can be explained by the specific distribution of the number of detections. For example, Figure 5d shows a high probability of detecting 3 anomalies. More generally, assuming that there exists $k^* \in \llbracket 1, m \rrbracket$ such that $\mathbb{P}(R(i) = k^*) \approx 1$, Theorem 1 justifies that

$$\begin{aligned} FDR_1^m(\hat{\varepsilon}_{BH_\alpha}, \hat{p}) &= \frac{n}{n+1} \cdot \alpha \frac{m_0}{m} + \frac{m_0}{n+1} \sum_{k=1}^m \frac{(1 - q_{n,k})}{k} \mathbb{P}(R(i) = k) \\ &\approx \frac{n}{n+1} \cdot \alpha \frac{m_0}{m} + \frac{m_0}{n+1} \frac{(1 - q_{n,k^*})}{k^*}. \end{aligned}$$

Then the proof detailed in Appendix B.3 yields that $1 - q_{n,k^*} = \frac{\alpha}{m(n+1)}$ can be reached for all $\nu \geq 1$, such that $n = \nu \frac{m}{\alpha k^*} - 1$. This allows to conclude that $FDR_1^m(\hat{\varepsilon}_{BH_\alpha}, \hat{p}) \approx \frac{m_0 \alpha}{m}$.

How to choose the right cardinality of the calibration set? Intuitively, an optimal choice of the cardinality n of the calibration set should enable the FDR control while minimizing the number of false negatives and avoiding excessive computation time. To achieve this objective, the first part of Corollary 1

explains that n must be chosen from the set $\mathcal{N} = \{\nu m/\alpha - 1, \nu \geq 1\}$. Using the simulation scenarios described in the paragraph “Experiment description”, FNR is estimated for different values n of the calibration set cardinality when $n \in \mathcal{N}$. The results are visualized in Figure 6 where the FNR value is represented as a function of n . For all the considered scenarios (Fig. 6a, 6b, 6c, 6d), the FNR




Figure 6: FNR as a function of the calibration set cardinality constrained to belong to \mathcal{N} .

value converges to the value reached with true p -values (horizontal dashed line) as n grows. From Figures 6a and 6b, the convergence speed depends on the “difficulty” of the problem.

In practice, however, the lack of labeled data prevents the computation of the actual FNR value, making it difficult to choose the optimal value of n . To address this challenge, our suggestion is to choose the largest possible value of n that does not exceed the computation time limit. This would result in a value of n that minimizes the FNR criterion while satisfying the computational constraints. However, following this suggestion does not prevent us from computational drawbacks, as shown in Figure 6a, where the optimal FNR value is reached for $n = 3999$, while choosing a larger n does not bring any gain (but still increases the computational cost).

5.2. Control of $mFDR$ on subseries

Experiment description Let say the parameter α' is chosen using the heuristic strategy provided in Appendix A.1.1, it comes from Corollary 5, the mBH_α

procedure controls the $mFDR_1^m$ only if the **Power** assumption is satisfied. Since the power of the anomaly detector depends on how easy it is to detect anomalies, the level of atypicality δ is introduced. To quantify the atypicality of a data point X_t , the true p -value is computed as $p_t = \mathbb{P}_{X \sim \mathcal{P}_0}(X > X_t)$, and the atypicality level is defined as the inverse of the p -value: $\delta_t = 1/p_t$. The atypicality level is preferred to the p -values because it is easier to plot on the x-axis of the graph when the p -value is small. To evaluate the effect of power, for each sample, all anomalies have their atypicality level lower bounded by a given parameter δ . Therefore, it is possible to observe the effect of a variation in the level of atypicality on $mFDR_1^m$, FDR_1^m , and FNR_1^m . For a given scenario - i.e. a proportion of anomalies π , a level of atypicality δ , and a desired level of mFDR noted α - the actual mFDR, FDR, and FNR are estimated. These quantities are estimated using $K = 50$ samples of m data points. To control for the estimation error introduced by estimating with a finite number of samples, each estimation is repeated $B = 100$ times. The estimation proceeds as follows: For each repetition b , for each sample k , m data points are generated. m_0 normal data points are generated according to the reference distribution $U([0, 1])$. m_1 anomalies are generated according to the alternative distribution $U([0, 1/\delta])$, where δ is the degree of atypicality of the anomalies. Two procedures are applied to each of this sequence of m points: BH_α and mBH_α . For each procedure, FDP, FNP as well as the number of detections and false positives are calculated. By averaging FDP and FNP over the K samples, FDR and FNR are estimated. By averaging the number of detections and false positives, mFDR can be estimated. These steps are then repeated over the different scenarios.

Results and Analysis The results are shown in Figure 7 by varying δ , α , and m_1 . In Figure 7, the level of atypicality δ is shown on the x-axis. The y-axis represents the estimated mFDR (in Figures 7a or 7c) or FNR (in Figures 7b or 7d). Different colors are used to distinguish BH and mBH procedures. For a low level of atypicality δ , FNR and mFDR are high because the anomalies are difficult to detect. By increasing δ , FNR and mFDR decrease. As shown in Figure 7b, for values of δ around 100, FNR is equal to 0, which can also generate a constant mFDR as shown in Figure 7a. For the mBH-procedure, mFDR is constant and equal to α . This is consistent with Theorem 3, which guarantees control at level α if all anomalies are detected. Figure 7d shows that all anomalies are detected for $\delta = 2000$. The same result in Figure 7b with $\delta = 100$. This is explained by the different parameters of the experiment. The easier the anomalies are detected, the faster $FNR = 0$ is reached for a small δ and therefore the easier it is to guarantee $mFDR = \alpha$.

Conclusion In order to control mFDR at the desired level α using mBH, FNR has to be equal to 0. The capacity of mBH to control mFDR depends on the difficulty of the problem. When the proportion of anomalies and the level of atypicality are lower, the power of mBH decreases and mFDR is harder to control. The results of this experiment give an idea of the atypicality level that the detector can find.

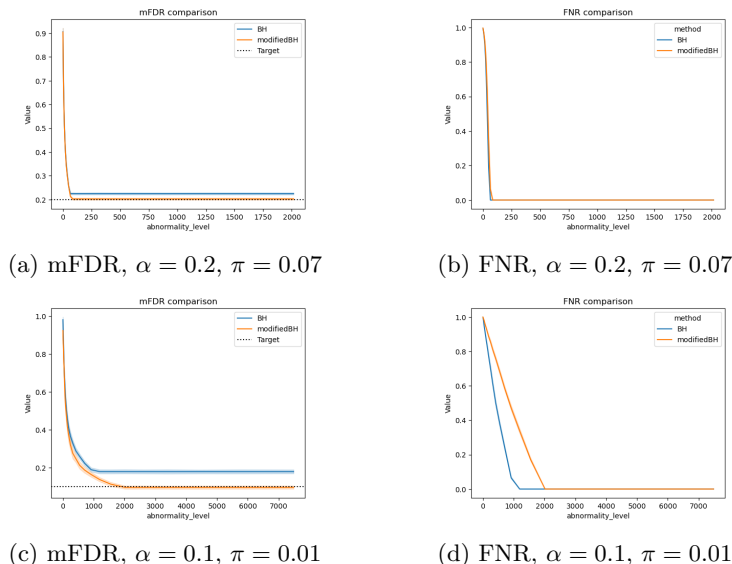


Figure 7: mFDR and FNR as a function of level of atypicality across different scenarios

5.3. Convergence of false discovery rate control

This section studies the convergence rate of FDR over the full time series using mBH_α .

Experiment Description The theoretical results obtained in Theorem 2 only guarantee an asymptotic control of FDR on the whole time series. In practice, it is more useful to have a control of FDR at any time, i.e., on subseries of finite size. The question is studied empirically by observing the speed of convergence of FDR to the level α . FDR of the full time series is calculated across different scenarios as a function of time series size. In order to get the distribution of FDR, the experiment is repeated on $B = 100$ time series. The maximal time series size explored is $T = 10^4$. For each time series b , for each time step t , the data point is generated according to the following mixture:

$$A_{b,t} \sim B(\pi)$$

$$p_{b,t} \sim \begin{cases} U([0, 1]), & \text{if } A_{b,t} = 1 \\ U([0, 1/\delta]), & \text{else} \end{cases}$$

Then the data driven threshold is computed by applying the $\hat{\varepsilon}_{mBH_\alpha}$ to the subseries $\{p_{t-m}, \dots, p_t\}$. The points where p_t is lower than the threshold are detected as anomalies. FDP is measured for each sub-series $\llbracket 0, t \rrbracket$ by comparing the detected anomalies with the true anomalies ($A_t = 1$). This allows us to track

the evolution of FDP while observing new data points. Different scenarios are generated by varying the proportion of anomalies π and the atypicality level δ .

Results and analysis In Figure 8, the false discovery proportion is represented on the y-axis according to the size of the time series given on the x-axis. Different levels of α , used to compute the mBH threshold, are tested, with the results of the median FDP and its 95% band shown in different colors. Different scenarios are represented by varying the proportion of anomalies between the subfigures. It can be observed that the convergence is quite fast from a size of 2000 data points, since for a α of 0.05, there is a 95% chance to have a false positive rate between 0.04 and 0.06, on Figure 8. Thus, the control of the false positive rate can be ensured with a high probability for a series of one data point per minute recorded over a few days.




Figure 8: FDR over the full time series as a function of the time series size.

Conclusion This ensures that control at the desired level α is achieved not only for infinite time series, but also for finite time series, which allows our model to be used in practice.

5.4. Anomaly detection in presence of time series dependency

The theoretical results that are presented in this paper require the time series to be iid. In practice, time series contain dependency. To apply our anomaly detector preprocessing needs to be applied in order to decorrelate the data points. In this section three examples are presented: iid gaussian and non-iid AR(3) and ARCH(3).

- iid: $\forall t \in \llbracket 1, T \rrbracket, A_t \sim B(\pi)$ and $X_t = \begin{cases} \mathcal{N}(0, 1) & \text{if } A_t = 0 \\ \Delta & \text{else} \end{cases}$
- AR(3): $\forall t \in \llbracket 1, T \rrbracket, A_t \sim B(\pi), z_t = \begin{cases} \mathcal{N}(0, 1) & \text{if } A_t = 0 \\ \Delta & \text{else} \end{cases}$
and $X_t = \phi_1 X_{t-1} + \phi_2 X_{t-2} + \phi_3 X_{t-3} + z_t$

model	p -value	FDR	FNR	model	p -value	FDR	FNR
iid	Sliding Cal.-*	0.19	0.01	iid	Sliding Cal.-*	0.12	0.01
iid	Sliding Cal.	0.41	0.00	iid	Sliding Cal.	0.32	0.00
AR(3)	Sliding Cal.-*	0.19	0.01	AR(3)	Sliding Cal.-*	0.11	0.01
AR(3)	Sliding Cal.	0.41	0.00	AR(3)	Sliding Cal.	0.33	0.00
ARCH(2)	Sliding Cal.-*	0.20	0.01	ARCH(2)	Sliding Cal.-*	0.13	0.04
ARCH(2)	Sliding Cal.	0.44	0.00	ARCH(2)	Sliding Cal.	0.31	0.00

(a) $\alpha = 0.2$ (b) $\alpha = 0.1$

TABLE 2
FDR and FNR results depending on different dependency models.

- ARCH(2): $\forall t \in \llbracket 1, T \rrbracket, A_t \sim B(\pi), z_t = \begin{cases} \mathcal{N}(0, 1) & \text{if } A_t = 0 \\ \Delta & \text{else} \end{cases},$
 $\sigma_t^2 = \theta_0 + \theta_1 X_{t-1}^2 + \theta_2 X_{t-2}^2$ and $X_t = \sigma_t z_t$

For each of these three examples, 100 time series of 3000 data points are generated. The first half of the time series is used to learn the parameters of the statistical model. For example, in the case of AR(3), the AR(3) model is fitted and $\hat{\phi}_1, \hat{\phi}_2$ and $\hat{\phi}_3$ are estimated. This model is then used to obtain the uncorrelated time series: $\hat{z}_t = X_t - \hat{\phi}_1 X_{t-1} + \hat{\phi}_2 X_{t-2} + \hat{\phi}_3 X_{t-3}$. The anomaly detector is applied to this extracted time series. The aim of the experiment is to assess whether the estimation of the model is accurate enough to find the anomalies in the series of \hat{z}_t . When the anomaly detector is applied, the p -value of \hat{z}_t is estimated by comparing its value with that of the n points preceding t with an estimated normal status. This estimation method is called the ‘‘Sliding Calibration set’’. In doing so, the estimation of the p -value may be biased by previous detection errors. An oracle version of our detector is also applied, where the p -value \hat{p}_t is estimated by comparing the value \hat{z}_t with that of the n points preceding t with a *true* normal status. This estimation method is called ‘‘Sliding Calibration set- \star ’’.

The results are summarized in Table 2, each line represent a different dependency model, with the two p -value estimators. FDR and FNR are presented in columns. When the ‘‘Sliding Cal.’’ estimator is used, the proportion of anomalies is not controlled to the desired level. For example in the AR(3) case FDR is measured at 0.4 instead of 0.2. However, when the ‘‘Sliding Cal.- \star ’’ estimator is used, FDR is controlled to the desired level in all the cases tested. This shows that the lack of control is due to the contamination of the calibration set and not to the dependence between observations This issue is discussed more deeply in Section 6.5.

This section shows that it is possible to apply our detector to non-independent time series. The difficulty lies in the possibility of whitening the data.

6. Comparison against competitor

Corollary 5 enables us to conclude that the mBH_α procedure can be used to control FDR of the full time series. If the parameter α' of mBH is computed using the heuristics provided in Appendix A.1.1, the assumptions **Heuristic** and **Power** must be verified. Even though these assumptions are hard to verify theoretically, the experiment in Section 5.2 shows that $mFDR_1^m$ is controlled for the tested scenario, provided that the anomalies are sufficiently atypical. The experiment in Section 5.3 shows that control of FDR is possible even when the time series is not infinite, as required by Theorem 2. These different results provide the conditions for building an anomaly detector that controls FDR of the time series by controlling mFDR on the subseries and the p -empirical value. Our anomaly detector is compared against Levels based On Recent Discovery (LORD), an online multiple testing procedure introduced in [16] to control FDR. The anomaly detectors are evaluated under different scenarios by varying the generated anomalies and the targeted FDR. To understand the source of the difficulties that the anomaly detectors may encounter, different sequences of p -values with oracle information are introduced.

6.1. Data

The synthetic data are generated from a Gaussian distribution. With the use of the empirical p -value estimator, there is no need to evaluate on other data distributions. Only the anomaly proportion and the distribution shift associated with anomalies affect the performance of the anomaly detector. Data are generated according to Definition 3 with Gaussian reference distribution and anomaly spike as in Section 5.1. The strength of the distribution shift noted by Δ takes value in $\{3\sigma, 3.5\sigma, 4\sigma\}$ and the anomaly proportion noted by π is equal to 0.01. Each generated time series contains $T = 10^4$ data points. Each experiment is repeated over 100 time series. For t in $\llbracket 1, T \rrbracket$:

$$A_t \sim B(\pi)$$

$$X_t = \begin{cases} \mathcal{N}(0, \sigma^2) & \text{if } A_t = 0 \\ \Delta\sigma & \text{else} \end{cases}$$

The value of Δ represents the atypicality score of the anomalies. Anomalies with higher Δ are easier to detect. In this experiment, the standard deviation σ is set to 1.

6.2. Threshold and p -value estimators description

Our detector mBH on overlapping subseries Using the p -value with the empirical estimator, the anomalies are detected by using mBH as the threshold estimator on overlapping subseries in the Algorithm 1. For each time t , the

threshold is computed as: $\hat{\varepsilon}_{mBH\alpha,t} = f_m(\hat{p}_{t-m}, \dots, \hat{p}_t)$, where f_m is the mBH-procedure. To ensure FDR control according to Theorem 1, the cardinality of the calibration is set equal to $n = \frac{m}{\alpha} - 1$. In this experiment, m is equal to 100 and α takes values 0.1 and 0.2 depending on the tested scenario. So the calibration set takes values of 999 or 1999.

LORD LORD introduced in [16] is based on alpha-investing rules to define a threshold on p -values. For each time t , the threshold is computed according to the alpha-investing rules, depending on the previous decisions made by the algorithm. For more precision, see the original article [16]. The empirical p -value given in Eq. 2.4 does not respect the super uniformity property required by LORD to control FDR. The conformal p -value defined in Equation 6.1 does.

$$\hat{p}_c(s_t, \mathcal{S}_t^{cal}) = \frac{1}{|\mathcal{S}_t^{cal}| + 1} \left(1 + \sum_{s \in \mathcal{S}_t^{cal}} \mathbb{1}[s_t > s] \right) \quad (6.1)$$

Using conformal p -values to apply the LORD algorithm leads to a weak power to detect anomalies. The issue is that $\hat{p}_c \geq \frac{1}{n+1}$ is always checked, and the threshold sequence $\hat{\varepsilon}_t$ decreases rapidly when no rejections are made. No anomaly can be detected. For these reasons, the empirical p value is used when applying LORD and mBH. This experiment uses LORD3 from [16] with the same parameters as in the original paper.

p -value estimation Different sequences of p -values are used to understand the limitations of our anomaly detector. The true p -values are used to evaluate the case where the only limitation comes from the multiple testing procedure. One can thus understand how the estimation of p -values affects the detection of anomalies. One way to estimate p -values in practice is to use the same calibration set for all p -values. This is referred to as a fixed calibration set. However, the p -values may be biased in that particular calibration set. In practice, the usual way to estimate p -values is to use a sliding calibration set. To estimate the p -value of a data point X_t , n preceding data points are used as a calibration set. To avoid bias in the estimation, the points detected as abnormal cannot be part of the calibration set and are replaced by other data points. However, the calibration set may be biased by undetected anomalies. To evaluate this effect, the sliding calibration set \rightarrow is introduced, where the knowledge oracle of the labels is used to construct the calibration set from the previous data points.

The different p -value sequences are computed as follows:

- **Oracle:** The true p -value is used instead of the estimated one.

$$\forall t \in \llbracket 1, T \rrbracket, \hat{p}_t = \Phi(X_t)$$

- **Fixed calibration set (Fixed Cal.):** The p -value is estimated using the

same calibration set $\{Z_i, i \in [1, n]\}$ for all observations.

$$\forall t \in \llbracket 1, T \rrbracket, \hat{p}_t = \frac{1}{n} \sum_{i=1}^n \mathbb{1}[Z_i > X_t]$$

- **Sliding Calibration set- \star (Sliding Cal.- \star):** The p -value is estimated using a calibration that is a sliding windows containing the n previous true normal data.

$$\forall t \in \llbracket 1, T \rrbracket, \hat{p}_t = \frac{1}{n} \sum_{i=1}^n \mathbb{1}[X_{h(t,i)} > X_t]$$

With h the function that selects observations that respect \mathcal{H}_0 . For each t and i , $h(t, i)$ returns the i -th observation lower than t that satisfies the \mathcal{H}_0 hypothesis.

- **Sliding Calibration set (Sliding Cal.):** The calibration set is a sliding windows containing the n previous estimated normal data.

$$\forall t \in \llbracket 1, T \rrbracket, \hat{p}_t = \frac{1}{n} \sum_{i=1}^n \mathbb{1}[X_{\hat{h}(t,i)} > X_t]$$

With \hat{h} the function that estimates the function h . For each t and i , $\hat{h}(t, i)$ gives the i -th observation lower than t and that satisfies $d_{\hat{h}(t,i)} = 0$.

6.3. Performance metrics

The anomaly detectors are evaluated by their ability to control FDR and minimize FNR of the full time series. Therefore, the two applied metrics are FDP and FNP calculated as follows:

$$FDP = \frac{\sum_{t=1}^T \mathbb{1}[\hat{p}_t < \hat{\varepsilon}_t](1 - A_t)}{\sum_{t=1}^T \mathbb{1}[\hat{p}_t < \hat{\varepsilon}_t]}$$

and

$$FNP = \frac{\sum_{t=1}^T \mathbb{1}[\hat{p}_t > \hat{\varepsilon}_t]A_t}{\sum_{t=1}^T A_t}$$

where $\hat{\varepsilon}_t$ is estimated using mBH or LORD and \hat{p}_t is estimated using one of the estimators defined in Section 6.2.

6.4. Results

The box plots shown in Figures 9-12 represent FDP and FNP distributions for 1000 repetitions. Within each subfigure (a, b, c, d, e,...), the box plots are

shown according to the multiple testing procedure mBH or LORD, the p -value estimator and the distribution shift between normal data and anomalies.

Table 3 gives a summary using FDR and FNR estimates. It enables an easy comparison between the values obtained with the different strategies combining the multiple testing procedure mBH or LORD, the choice of the level α varying from 0.1 to 0.2, the p -value estimator and the distribution shift between the normal data and the anomalies.

FDR, $\alpha = 0.1$				FNR, $\alpha = 0.1$			
	$\Delta = 4\sigma$	$\Delta = 3.5\sigma$	$\Delta = 3\sigma$		$\Delta = 4\sigma$	$\Delta = 3.5\sigma$	$\Delta = 3\sigma$
mBH with Oracle PV	0.101	0.113	0.281	mBH with Oracle PV	0.020	0.151	0.793
mBH with Fixed Cal.	0.100	0.109	0.348	mBH with Fixed Cal.	0.026	0.135	0.669
mBH with Sliding Cal.-*	0.100	0.113	0.256	mBH with Sliding Cal.-*	0.019	0.140	0.669
mBH with Sliding Cal.	0.335	0.222	0.346	mBH with Sliding Cal.	0.040	0.217	0.694
LORD with Oracle PV	0.106	0.115	0.367	LORD with Oracle PV	0.033	0.260	0.905
LORD with Fixed Cal.	0.111	0.277	0.736	LORD with Fixed Cal.	0.070	0.340	0.896
LORD with Sliding Cal.-*	0.070	0.190	0.841	LORD with Sliding Cal.-*	0.781	0.845	0.978
LORD with Sliding Cal.	0.075	0.098	0.627	LORD with Sliding Cal.	0.052	0.327	0.907

(a) FDR, $\alpha = 0.1$				(b) FNR, $\alpha = 0.1$			
	$\Delta = 4\sigma$	$\Delta = 3.5\sigma$	$\Delta = 3\sigma$		$\Delta = 4\sigma$	$\Delta = 3.5\sigma$	$\Delta = 3\sigma$
mBH with Oracle PV	0.200	0.208	0.277	mBH with Oracle PV	0.009	0.062	0.395
mBH with Fixed Cal.	0.206	0.211	0.301	mBH with Fixed Cal.	0.014	0.045	0.355
mBH with Sliding Cal.-*	0.210	0.219	0.283	mBH with Sliding Cal.-*	0.008	0.059	0.339
mBH with Sliding Cal.	0.833	0.815	0.761	mBH with Sliding Cal.	0.003	0.018	0.101
LORD with Oracle PV	0.211	0.216	0.290	LORD with Oracle PV	0.016	0.117	0.610
LORD with Fixed Cal.	0.210	0.263	0.665	LORD with Fixed Cal.	0.04	0.144	0.689
LORD with Sliding Cal.-*	0.061	0.149	0.625	LORD with Sliding Cal.-*	0.805	0.835	0.941
LORD with Sliding Cal.	0.117	0.133	0.321	LORD with Sliding Cal.	0.026	0.168	0.692

(c) FDR, $\alpha = 0.2$				(d) FNR, $\alpha = 0.2$			
	$\Delta = 4\sigma$	$\Delta = 3.5\sigma$	$\Delta = 3\sigma$		$\Delta = 4\sigma$	$\Delta = 3.5\sigma$	$\Delta = 3\sigma$
mBH with Oracle PV	0.101	0.113	0.281	mBH with Oracle PV	0.020	0.151	0.793
mBH with Fixed Cal.	0.100	0.109	0.348	mBH with Fixed Cal.	0.026	0.135	0.669
mBH with Sliding Cal.-*	0.100	0.113	0.256	mBH with Sliding Cal.-*	0.019	0.140	0.669
mBH with Sliding Cal.	0.335	0.222	0.346	mBH with Sliding Cal.	0.040	0.217	0.694
LORD with Oracle PV	0.106	0.115	0.367	LORD with Oracle PV	0.033	0.260	0.905
LORD with Fixed Cal.	0.111	0.277	0.736	LORD with Fixed Cal.	0.070	0.340	0.896
LORD with Sliding Cal.-*	0.070	0.190	0.841	LORD with Sliding Cal.-*	0.781	0.845	0.978
LORD with Sliding Cal.	0.075	0.098	0.627	LORD with Sliding Cal.	0.052	0.327	0.907

TABLE 3
Comparison of mBH versus LORD for online anomaly detection in Gaussian white noise with different abnormality levels.

6.5. Analysis

Effect of the strength of the distribution shift Δ According to the assumption **Power** of the Theorem 3, mBH_α allows control of FDR at the level α if all anomalies are detected. To test this assertion, the different columns of the Table 3a, are compared. In the row “mBH with Oracle PV”, with $\Delta = 4\sigma$ FDR is estimated at 0.101, which is close to the desired level $\alpha = 0.1$. While with $\Delta = 3\sigma$ the FDR level is estimated at 0.281 which is almost three times the desired level α . The FNR results in Table 3b need to be taken into consideration. When $\Delta = 4\sigma$, FNR is close to 0, while when $\Delta = 3\sigma$, FNR is 0.793. Similar results are obtained with other test configurations in Table 3c and Table 3d. To control FDR at the desired level, FNR needs to be close to 0.

Effect of p -value estimation To understand how the p -value estimation can affect the control of FDR, the first four rows of Table 3a are compared. In the column “ 4σ ”, the FDR values for the configurations “Oracle PV”, “Fixed Cal.” and “Sliding Cal.-*” are very close to the desired level $\alpha = 0.1$. This control is enabled by Theorem 3, since the p -values verify all hypotheses, in particular all data in the calibration sets are generated according to the reference distribution. However, in the case of “Sliding Cal.”, FDR increases at a value of 0.335. For

the same configurations, FNR remains low, between 0.019 and 0.040, as shown in Table 3b. The increase of FDR when using “Sliding Cal.” instead of “Sliding Cal.- \star ” is a consequence of calibration set contamination. Indeed, according to the procedure used to build the calibration sets described in Section 6.2, any detected anomalies are removed from the calibration sets used to estimate the next p -values. If an observation is incorrectly detected as an anomaly, that data point cannot be part of the calibration set in future steps of online detection. Instead, it is replaced by another data point having a statistically lower atypicality score (false positives have high atypicality scores). As a result, the calibration set contains data points with lower scores than if it had been generated under \mathcal{P}_0 . This leads to an underestimation of the p -values and an increase in the number of false positives. This illustrates the major drawback of mBH: it is highly sensitive to the lack of robustness of the p -value estimator. Figure 9a shows that using a fixed calibration instead of a sliding calibration- \star gives a larger variance on FDP while FDR is the same. Using a single calibration set for the entire time series means that FDP is highly dependent on the start of the time series. By changing the calibration set at each time step, the statistical fluctuations in FDP are smoothed over the course of the time series analysis.

Comparison with LORD This section compares the results obtained with mBH and those obtained with LORD. As known from the literature, LORD controls FDR of super-uniform p -values. In this experiment, the question is the ability of the LORD method to control FDR of empirical p -values that have no theoretical guarantees. As shown in Figure 3a, LORD is able to ensure control of FDR for all calibration set definitions when anomalies are easier to detect, such as for $\Delta = 4\sigma$ or $\Delta = 3.5\sigma$. In particular, unlike mBH, LORD is able to control FDR in the case of the sliding calibration set. However, the mBH method has a lower FNR than the LORD method, as shown in 3a and 3b. For example, Table 3b shows that FNR is equal to 0.019 with mBH, while it is equal to 0.781 with LORD, in the case of using the Sliding Calibration set- \star on data with $\Delta = 3\sigma$. Nevertheless, in the Sliding Calibration Set case, the LORD method has approximately the same FNR, but with a lower FDR (0.335 versus 0.075). The contamination problem of mBH offsets the superior performance observed in the Sliding Calibration Set- \star .

7. Conclusion

A new online anomaly detector has been described that aims at controlling FDR of the full time series at a prescribed level α . The main challenges tackled by the new detector are:

- Empirical p -values: it clearly identifies and fulfills conditions on the calibration set cardinality (used for computing empirical p -values) ensuring the FDR control. By contrast in situations where empirical p -values are used with no care about these conditions, then FDR is not be controlled at a prescribed level.

- Online detection: it ensures the FDR control of the full time series (global control) by means of the local control of mFDR of sub-series, using a modified version of the BH-procedure.

The theoretical guarantees developed in the present work are mainly established under the independence assumption, where the tightest control is made possible. However some results on the FDR control are also derived under dependence assumptions such as the PRDS property. The reported empirical experiments highlight the ability of the new procedure to control FDR of the full time series, while it also clearly emphasizes the sensitivity of the procedure with respect to its hyper-parameters or the features of the underlying probability distribution of the data.

8. Acknowledgments

The authors would like to thank Cristian Preda, director of the MODAL team at Inria, for valuable discussions.

References

- [1] AHMED, M., MAHMOOD, A. N. and HU, J. (2016). A survey of network anomaly detection techniques. *Journal of Network and Computer Applications* **60** 19–31.
- [2] ANGELOPOULOS, A. N. and BATES, S. (2021). A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv preprint arXiv:2107.07511*.
- [3] AUE, A. and KIRCH, C. (2024). The state of cumulative sum sequential changepoint testing 70 years after Page. *Biometrika* **111** 367–391.
- [4] BATES, S., CANDÈS, E., LEI, L., ROMANO, Y. and SESIA, M. (2023). Testing for outliers with conformal p-values. *The Annals of Statistics* **51** 149–178.
- [5] BENJAMINI, Y. and HOCHBERG, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)* **57** 289–300.
- [6] BENJAMINI, Y. and YEKUTIELI, D. (2001). The control of the false discovery rate in multiple testing under dependency. *Annals of statistics* 1165–1188.
- [7] BLÁZQUEZ-GARCÍA, A., CONDE, A., MORI, U. and LOZANO, J. A. (2021). A review on outlier/anomaly detection in time series data. *ACM Computing Surveys (CSUR)* **54** 1–33.
- [8] BLUM, J. M. and TREMPER, K. K. (2010). Alarms in the intensive care unit: too much of a good thing is dangerous: is it time to add some intelligence to alarms? *Critical care medicine* **38** 702–703.
- [9] BRAEI, M. and WAGNER, S. (2020). Anomaly detection in univariate time-series: A survey on the state-of-the-art. *arXiv preprint arXiv:2004.00433*.
- [10] CHANDOLA, V., BANERJEE, A. and KUMAR, V. (2009). Anomaly detection: A survey. *ACM computing surveys (CSUR)* **41** 1–58.

- [11] CVACH, M. (2012). Monitor alarm fatigue: an integrative review. *Biomedical instrumentation & technology* **46** 268–277.
- [12] FISHER, R. (1951). The Design of Experiments, volume 6th Ed. *Hafner, New York, NY*.
- [13] FOSTER, D. P. and STINE, R. A. (2008). α -investing: a procedure for sequential control of expected false discoveries. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **70** 429–444.
- [14] GRETTON, A., BORWARDT, K. M., RASCH, M. J., SCHÖLKOPF, B. and SMOLA, A. (2012). A kernel two-sample test. *The Journal of Machine Learning Research* **13** 723–773.
- [15] HUBER, P. J. (1992). Robust estimation of a location parameter. In *Breakthroughs in statistics: Methodology and distribution* 492–518. Springer.
- [16] JAVANMARD, A. and MONTANARI, A. (2018). Online rules for control of false discovery rate and false discovery exceedance. *The Annals of statistics* **46** 526–554.
- [17] KIRCH, C. and STOEHR, C. (2022). Sequential change point tests based on U-statistics. *Scandinavian Journal of Statistics* **49** 1184–1214.
- [18] LAXHAMMAR, R. (2014). Conformal anomaly detection: Detecting abnormal trajectories in surveillance applications, PhD thesis, University of Skövde.
- [19] LEWANDOWSKA, K., MĘDRZYCKA-DĄBROWSKA, W., TOMASZEK, L. and WUJTEWICZ, M. (2023). Determining Factors of Alarm Fatigue among Nurses in Intensive Care Units—A Polish Pilot Study. *Journal of Clinical Medicine* **12** 3120.
- [20] LI, Y., ZHANG, X., HE, S., CHEN, Z., KANG, Y., LIU, J., LI, L., DANG, Y., GAO, F., XU, Z. et al. (2022). An Intelligent Framework for Timely, Accurate, and Comprehensive Cloud Incident Detection. *ACM SIGOPS Operating Systems Review* **56** 1–7.
- [21] MARANDON, A., LEI, L., MARY, D. and ROQUAIN, E. (2022). Machine learning meets false discovery rate. *arXiv preprint arXiv:2208.06685*.
- [22] PAGE, E. S. (1954). Continuous inspection schemes. *Biometrika* **41** 100–115.
- [23] PHIPSON, B. and SMYTH, G. K. (2010). Permutation P-values should never be zero: calculating exact P-values when permutations are randomly drawn. *Statistical applications in genetics and molecular biology* **9**.
- [24] RAMDAS, A., ZRNIC, T., WAINWRIGHT, M. and JORDAN, M. (2018). SAF-FRON: an adaptive algorithm for online control of the false discovery rate. In *International conference on machine learning* 4286–4294. PMLR.
- [25] RUFF, L., KAUFFMANN, J. R., VANDERMEULEN, R. A., MONTAVON, G., SAMEK, W., KLOFT, M., DIETTERICH, T. G. and MÜLLER, K.-R. (2021). A unifying review of deep and shallow anomaly detection. *Proceedings of the IEEE* **109** 756–795.
- [26] SOLET, J. M. and BARACH, P. R. (2012). Managing alarm fatigue in cardiac care. *Progress in Pediatric Cardiology* **33** 85–90.
- [27] STAERMAN, G. (2022). Functional anomaly detection and robust estimation, PhD thesis, Institut polytechnique de Paris.

- [28] WANG, W., LIU, Z., SHI, X. and PIERCE, L. (2019). Online fdr controlled anomaly detection for streaming time series. In *5th Workshop on Mining and Learning from Time Series (MiLeTS)*.
- [29] WU, R. and KEOGH, E. J. (2021). Current time series anomaly detection benchmarks are flawed and are creating the illusion of progress. *IEEE transactions on knowledge and data engineering* **35** 2421–2429.
- [30] WU, Y. (2007). *Inference for change point and post change means after a CUSUM test* **180**. Springer Science & Business Media.
- [31] XU, Z. and RAMDAS, A. (2022). Dynamic algorithms for online multiple testing. In *Mathematical and Scientific Machine Learning* 955–986. PMLR.

Appendix A: More precision on the method

A.1. Evaluating the ratio of rejection numbers

A.1.1. Using heuristic arguments

The Section 4.2.1 raises the importance of the ratio $\mathbb{E}R_{1,\alpha}^{*,m}/\mathbb{E}R_{1,\alpha}^m$ of rejection numbers. The present section aims at deriving a numeric approximation to this ratio. In a first step, a first result details the value of the denominator. In a second step, an approximation to the numerator is derived based on a heuristic argument and also empirically justified on simulation experiments.

When $mFDR$ is assumed to equal α , the expected number of rejections can be made explicit.

Proposition 3. *With the previous notation, let (X_1, \dots, X_m) be given by Definition 3, where π denotes the unknown proportion of anomalies, and assume that $mFDR_1^m = \alpha$ and $FNR_1^m = \beta \in [0, 1]$. Then*

$$\mathbb{E}[R_{1,\alpha}^m] = \frac{m\pi(1 - \beta)}{1 - \alpha}. \quad (\text{A.1})$$

The proof is postponed to Appendix B.6. For instance, Eq. (A.1) establishes that the expected number of rejection output by BH_α increases with π , the unknown proportion of anomalies along the signal. This makes sens since the more anomalies, the more expected rejections. The expected number of rejection is also increasing with α : the larger α , the less restrictive the threshold, and the more rejections should be made. However the number of rejection decreases with the FNR value β . As β increases, the proportion of false negatives grows meaning that fewer alarms are raised, which results in a smaller number of rejections.

In what follows, the assumption is made that anomalies are easy to detect, meaning that the FNR_1^m value β is negligible compared to 1. In this context, Proposition 3 would yield that

$$\mathbb{E}[R_{1,\alpha}^m] \approx \frac{m\pi}{1 - \alpha}. \quad (\text{Power})$$

Another assumption is also made about the relationship between $\mathbb{E}[R_{1,\alpha}^m]$ and $\mathbb{E}[R_{1,\alpha}^{*,m}]$. This assumption is based on a heuristic argument supported by the results of numerical experiments as reported in Table 4. In what follows, it is assumed that

$$\mathbb{E}[R_{1,\alpha}^{*,m}] = \mathbb{E}[R_{1,\alpha}^m] + 1. \quad (\text{Heuristic})$$

No mathematical proof of this statement is given in the present paper. However, Table 4 displays numerical values which empirically support this approximation, whereas further analyzing the connection between these quantities should be necessary.

BH_α	0.05	0.1	0.2
$\mathbb{E}[R_{1,\alpha}^m]$	2.14	2.32	2.78
$\mathbb{E}[R_{1,\alpha}^m(i)]$	3.18	3.44	3.99

TABLE 4

Numerical evaluations for different values of α (10^3 repetitions)

Let us emphasize that Table 4 has been obtained with Gaussian data (generated similarly to those detailed in Section 5.1). For all the three considered values of α , one observes that $\mathbb{E}[R^*]$ remains close to (but also slightly larger than) $\mathbb{E}[R] + 1$.

These two assumptions give rise to a strategy for computing the ratio $\frac{\mathbb{E}[R^*]}{\mathbb{E}[R]}$. So all ingredient to build a procedure that control mFDR are given.

$$\frac{\mathbb{E}[R^*]}{\mathbb{E}[R]} = 1 + \frac{1 - \alpha}{m\pi} \quad (\text{A.2})$$

$$\alpha' = \alpha \left(1 + \frac{1 - \alpha}{m\pi} \right)^{-1} \quad (\text{A.3})$$

$$(\text{A.4})$$

The value α' can be used with the mBH procedure. Under (**Heuristic**) and (**Power**), according to Corollary 5, mBH allows global control of FDR at the level α .

A.1.2. Estimation on a training set

The true proportion of anomalies is usually not known, and estimating the proportion of anomalies is error-prone. Assumptions are also difficult to ensure. The number of detections is known by the user. To estimate $\mathbb{E}[R_{\hat{\alpha}}]$, the procedure $BH_{\hat{\alpha}}$ is applied to each of the subseries of length m from the training set. The average number of detections is computed and noted as $\hat{\mu}_{R_{\hat{\alpha}}}$. However, it is not possible to compute $R_{\hat{\alpha}'}^*$, the number of detections after one p -value associated with normal data is set to 0 since the true labels are unknown. But since the

proportion of true anomalies π is close to 0, $\mathbb{E}[R_{\tilde{\alpha}'}^*]$ can be approximate by $\mathbb{E}[R_{\tilde{\alpha}'}^{**}]$, the number of detections after one p -values (possible abnormal) is set to 0. Then, to estimate $\mathbb{E}[R_{\tilde{\alpha}'}^{**}]$, the procedure $BH_{\tilde{\alpha}}$ is applied again to the same subseries but after a randomly chosen p -value is replaced by 0. Here again, the average number of detections is computed and noted $\hat{\mu}_{R_{\tilde{\alpha}}}^{**}$. By varying $\tilde{\alpha}$ it is possible to estimate:

$$\alpha' = \arg \max_{\tilde{\alpha}} \left\{ \frac{\hat{\mu}_{R_{\tilde{\alpha}}}^{**}}{\hat{\mu}_{R_{\tilde{\alpha}}}^{**}} \tilde{\alpha} \leq \alpha \right\} \quad (\text{A.5})$$

This estimator of α' can be used for the modified BH procedure. According to Corollary 5, the global control of FDR is ensured by $\hat{\varepsilon}_{BH_{\alpha'}}$ when the size of the training set goes to infinity.

Appendix B: Proofs

B.1. Proof of Theorem 1

Proof of Theorem 1. The proof of BH applied to true p -values is reproduced (see Section 1.1 of the Supplementary Materials), but in the case of empirical p -values. The only modification is that $\mathbb{P}(\hat{p}_i \leq \frac{\alpha k}{m})$ is not equal to $\frac{\alpha k}{m}$ since \hat{p}_i now follows the discrete uniform distribution

$$\mathbb{P}(\hat{p}_i \leq \frac{\alpha k}{m}) = \sum_{\ell=0}^{\lfloor n\alpha k/m \rfloor} \mathbb{P}(n\hat{p}_i = \ell) = \frac{\lfloor \frac{\alpha k n}{m} \rfloor + 1}{n + 1}.$$

Plugging this in the FDR expression, it gives

$$FDR_1^m(\hat{\varepsilon}_{BH_{\alpha}}, \hat{p}) = m_0 \sum_{k=1}^m \frac{\lfloor \frac{\alpha k n}{m} \rfloor + 1}{k} \mathbb{P}(R(i) = k).$$

Recall that \hat{p}_i follows $U(0, 1/n, 2/n, \dots, 1)$ entails that $n\hat{p}_i$ follows $U(0, 1, 2, \dots, n)$. \square

B.2. PRDS property for p -values having overlapping calibration set

The following construction is used to describe a family of p -values with overlapping calibration set. Let Z the vector that combine all calibration set, the Z_i are i.i.d. with marginal probability \mathcal{P}_0 . The set of the n indices defining the elements of the calibration set related to \hat{p}_i in Z is noted \mathcal{D}_i . The calibration related to X_1 is noted $Z_{\mathcal{D}_1} = (Z_{i_1}, \dots, Z_{i_n})$. For all i in $\llbracket 1, m \rrbracket$: $\hat{p}_i = p\text{-value}(X_i, Z_{\mathcal{D}_i})$.

To proof that p -values with overlapping calibration set are PRDS as described in Definition 2, the methodology used in [4] to be extended in the case of overlapping calibration set. For i in $\llbracket 1, m \rrbracket$ the calibration set associated to X_i is noted $Z_{\mathcal{D}_i}$. The law of total probabilities gives:

$$\begin{aligned}\mathbb{P}[\hat{p}_1^m \in A | \hat{p}_i = u] &= \int \mathbb{P}[\hat{p}_1^m \in A | \hat{p}_i = u | Z_{\mathcal{D}_i} = z] \mathbb{P}[Z_{\mathcal{D}_i} = z] dz \\ &= \mathbb{E}_{Z_{\mathcal{D}_i} | \hat{p}_i = u} \mathbb{P}[\hat{p}_1^m \in A | \hat{p}_i = u | Z_{\mathcal{D}_i} = z]\end{aligned}$$

If these two lemma are suppose to be true, the PRDS property is verified.

Lemma 3.1. *For non-decreasing set A and vectors z, z' such that $z \succeq z'$, then*

$$\mathbb{P}[\hat{p}_1^m \in A | Z_{\mathcal{D}_i} = z] \geq \mathbb{P}[\hat{p}_1^m \in A | Z_{\mathcal{D}_i} = z'] \quad (\text{B.1})$$

Lemma 3.2. *For $u \geq u'$, if i belongs to the set of inliers, there exists $Z_{\mathcal{D}_i,1} \sim Z_{\mathcal{D}_i} | \hat{p}_i = u$ and $Z_{\mathcal{D}_i,2} \sim Z_{\mathcal{D}_i} | \hat{p}_i = u'$ such that $\mathbb{P}[Z_{\mathcal{D}_i,1}] \succeq \mathbb{P}[Z_{\mathcal{D}_i,2}]$*

Indeed, take $i \in \llbracket 1, m \rrbracket$ and $u \geq u'$ and define $Z_{\mathcal{D}_i,1}$ and $Z_{\mathcal{D}_i,2}$ as in the statement of Lemma 3.2.

$$\begin{aligned}\mathbb{P}[\hat{p}_1^m \in A | p_i = u] &= \mathbb{E}_{Z_{\mathcal{D}_i,1}}[\mathbb{P}[\hat{p}_1^m \in A | Z_{\mathcal{D}_i} = Z_{\mathcal{D}_i,1}]] \quad (\text{Lemma 3.2}) \\ &\geq \mathbb{E}_{Z_{\mathcal{D}_i,2}}[\mathbb{P}[\hat{p}_1^m \in A | Z_{\mathcal{D}_i} = Z_{\mathcal{D}_i,2}]] \quad (\text{Lemma 3.1}) \\ &\geq \mathbb{P}[\hat{p}_1^m \in A | p_i = u'] \quad (\text{Lemma 3.2})\end{aligned}$$

It shows that, when $u \geq u'$ then $\mathbb{P}[\hat{p}_1^m \in A | p_i = u] \geq \mathbb{P}[\hat{p}_1^m \in A | p_i = u']$, which means $\mathbb{P}[\hat{p}_1^m \in A | p_i = u]$ is increasing in u . The PRDS property is satisfied. To complete the proof, the introduced lemmas are proven.

Proof of Lemma 3.1. Let be i in $\llbracket 1, m \rrbracket$ and vectors z, z' and \bar{z} vectors such that $z \succeq z'$. The vectors z, z' are used to define the calibration set related to the p -values \hat{p}_i and \bar{z} is used to define elements of calibrations sets that are not in the calibration set of \hat{p}_i . By conditioning on the calibration sets defined by (z, \bar{z}) and (z', \bar{z}) it gives:

$$\mathbb{P}[\hat{p}_1^m \in A | Z_{\mathcal{D}_i} = z, Z_{\overline{\mathcal{D}_i}} = \bar{z}] \geq \mathbb{P}[\hat{p}_1^m \in A | Z_{\mathcal{D}_i} = z', Z_{\overline{\mathcal{D}_i}} = \bar{z}] \quad (\text{B.2})$$

This result comes from the decomposition the following decomposition, for all j in $\llbracket 1, m \rrbracket$

$$\begin{aligned}\hat{p}_j &= \frac{1}{n} \sum_{k \in \mathcal{D}_j} \mathbb{1}[a(Z_k) \geq a(X_j)] \\ &= \frac{1}{n} \left(\sum_{k \in \mathcal{D}_j \cap \mathcal{D}_i} \mathbb{1}[a(Z_k) \geq a(X_j)] + \sum_{k \in \mathcal{D}_j \setminus \mathcal{D}_i} \mathbb{1}[a(Z_k) \geq a(X_j)] \right)\end{aligned}$$

The conclusion comes from $Z_{\mathcal{D}_i} \succeq Z'_{\mathcal{D}_i}$ which implies $Z_{\mathcal{D}_i \cap \mathcal{D}_j} \succeq Z'_{\mathcal{D}_i \cap \mathcal{D}_j}$.

Since $Z_{\mathcal{D}_j \setminus \mathcal{D}_i} \perp Z_{\mathcal{D}_i}$, Eq. B.2 can be integrated over $Z_{\overline{\mathcal{D}_i}}$ to give:

$$\mathbb{P}[\hat{\rho}_1^m \in A | Z_{\mathcal{D}_i} = z] \geq \mathbb{P}[\hat{\rho}_1^m \in A | Z_{\mathcal{D}_i} = z'] \quad (\text{B.3})$$

□

Proof of Lemma 3.2. Let $S'_{i,(1)} \leq S_{i,(2)} \leq \dots \leq S_{i,(n)}$ the order statistics of $(a(Z_{\mathcal{D}_i,1}), \dots, a(Z_{\mathcal{D}_i,n}))$. Let $S'_{i,(1)} \leq S'_{i,(2)} \leq \dots \leq S'_{i,(n+1)}$ the order statistics of $(a(Z_{\mathcal{D}_i,1}), \dots, a(Z_{\mathcal{D}_i,n}), a(X_i))$. And R_i the rank of $a(X_i)$ among these.

$$\{(S_{(1)}, \dots, S_{(n)}) | R_i = k, S'_{i,(1)}, \dots, S'_{i,(n+1)}\} = (S'_{(1)}, \dots, S'_{(k-1)}, S'_{(k+1)}, \dots, S'_{(n+1)}) \quad (\text{B.4})$$

Using that R_i is independent of $S'_{i,(1)}, \dots, S'_{i,(n+1)}$:

$$\{(S_{(1)}, \dots, S_{(n)}) | R_i = k\} = (S'_{(1)}, \dots, S'_{(k-1)}, S'_{(k+1)}, \dots, S'_{(n+1)}) \quad (\text{B.5})$$

The right-hand side is not increasing with k and $\hat{\rho}_i = \frac{R_i - 1}{n}$ □

B.3. Proof of Corollary 1

Proof of Corollary 1. To get a deeper understanding of the FDR expression obtained in Theorem 1, $q_{n,k}$ the fractional part of $\frac{\alpha kn}{m}$ is introduced:

$$q_{n,k} = \frac{\alpha kn}{m} - \left\lfloor \frac{\alpha kn}{m} \right\rfloor$$

When plugged into the FDR expression, it gives:

$$\begin{aligned} FDR &= m_0 \sum_{k=1}^m \frac{\frac{\alpha kn}{m} + 1 - q_{n,k}}{n+1} \frac{1}{k} \mathbb{P}(R(1) = k) \\ FDR &= \frac{m_0 \alpha}{m} \frac{n}{n+1} + \frac{m_0}{n+1} \sum_{k=1}^m \frac{1 - q_{n,k}}{k} \mathbb{P}(R(1) = k) \end{aligned} \quad (\text{B.6})$$

In order to get lower and upper bounds of FDR, the value of $q_{n,k}$ should be expressed as a function of α , k , n and m .

For the next part of the proof, it is useful to express the relation between $q_{n,k}$ and $q_{n+1,k}$. It gives the effect of increasing the cardinality of the calibration by one. Using the definition of the fractional part:

$$\begin{aligned} q_{n+1,k} - q_{n,k} &= \frac{\alpha k(n+1)}{m} - \left\lfloor \frac{\alpha k(n+1)}{m} \right\rfloor - \frac{\alpha kn}{m} + \left\lfloor \frac{\alpha kn}{m} \right\rfloor \\ q_{n+1,k} - q_{n,k} &= \frac{\alpha k}{m} - \left\lfloor \frac{\alpha k(n+1)}{m} \right\rfloor + \left\lfloor \frac{\alpha kn}{m} \right\rfloor \end{aligned}$$

Which can be expressed as a congruence relation:

$$q_{n+1,k} - q_{n,k} \equiv \frac{\alpha k}{m} \pmod{1} \quad (\text{B.7})$$

Two cases are studied:

1. Particular case: there exists an integer $1 \leq \nu$ such that $\frac{\nu m}{\alpha}$ is an integer. the notation $n_\nu = \frac{\nu m}{\alpha}$ is introduced. Since: $\frac{\alpha k n_\nu}{m} = \frac{\alpha k \nu m / \alpha}{m} = k \nu$ is an integer, then the fractional part is null:

$$q_{n_\nu, k} = 0$$

If the calibration set cardinality n is equal to $n = n_\nu - 1 = \frac{\nu m}{\alpha} - 1$. Then, the congruence relation in Eq. B.7 gives:

$$\begin{aligned} q_{n_\nu-1, k} &\equiv q_{n, k} - \alpha k / m \pmod{1} \\ q_{n_\nu-1, k} &\equiv 0 - \alpha k / m \pmod{1} \end{aligned}$$

Using the fact that fractional part of a number belongs to $[0, 1[$, the only possible value to $q_{n_\nu-1, k}$ is:

$$q_{n_\nu-1, k} = 1 - \alpha k / m$$

Plugging the value of $q_{n_\nu-1, k}$ into Eq. B.6, it gives:

$$FDR = \frac{m_0 \alpha n}{m(n+1)} + \frac{m_0}{n+1} \sum_{k=1}^m \frac{\alpha k}{km} \mathbb{P}(R(1) = k)$$

Simplifying by k and using that $\sum_{k=1}^m \mathbb{P}(R(i) = k) = 1$, the result is obtained:

$$\begin{aligned} FDR &= \frac{m_0 \alpha n}{m(n+1)} + \frac{m_0 \alpha}{(n+1)m} \\ FDR &= \frac{m_0 \alpha}{m} \end{aligned}$$

2. General case: With $\alpha \in]0, 1]$, for each ν the notation $n_\nu = \lceil \frac{\nu m}{\alpha} \rceil$ is introduced. Notice that this definition is consistent with the particular case. The ceiling function definition gives:

$$\left\lceil \frac{\nu m}{\alpha} \right\rceil - 1 < \frac{\nu m}{\alpha} \leq \left\lceil \frac{\nu m}{\alpha} \right\rceil$$

Multiplying by αk on each side and the n_ν notation:

$$\frac{\alpha k (n_\nu - 1)}{m} < k \nu \leq \frac{\alpha k (n_\nu)}{m}$$

It implies that $\lfloor \frac{\alpha k(n_\nu - 1)}{m} \rfloor < \lfloor \frac{\alpha k(n_\nu)}{m} \rfloor$. Also, Eq. B.7 is expressed as $q_{n_\nu, k} - q_{n_\nu - 1, k} \equiv \frac{\alpha k}{m} \pmod{1}$:

$$1 - \frac{\alpha k}{m} \leq q_{n_\nu - 1, k} < 1 \quad (\text{B.8})$$

Indeed, the fractional part of a number has to be larger than $1 - \alpha k/m$ so that adding $\alpha k/m$ increases the integer part.

By plugging the bounds of $q_{n_\nu - 1, k}$ into Eq. B.6, it can give the bounds of FDR. At first, to compute the upper bound of FDR the lower bound of $q_{n_\nu - 1, k}$ is used:

$$FDR \leq \frac{m_0(n_\nu - 1)\alpha}{mn_\nu} + \frac{m_0}{n+1} \sum_{k=1}^m \frac{\alpha k}{km} \mathbb{P}(R(1) = k)$$

With the same calculations as for the ‘‘Particular case’’, it gives:

$$FDR \leq \frac{m_0\alpha}{m}$$

Similarly, the lower bound of FDR can be obtained using the $q_{n, k}$ upper bound from Eq. B.8 plugged into Eq. B.6:

$$\begin{aligned} \frac{m_0(n_\nu - 1)\alpha}{mn_\nu} + \frac{m_0}{n+1} \sum_{k=1}^m \frac{(1-k)}{k} \mathbb{P}(R_1 = k) &< FDR \\ \frac{m_0(n_\nu - 1)\alpha}{mn_\nu} &< FDR \end{aligned}$$

□

B.4. Proof of Theorem 2

Proof of Theorem 2. Let us start with the FDP expression for a time series of length T .

$$\begin{aligned} FDP_{t=1}^T &= \frac{\sum_{t=1}^T \mathbb{1}[\hat{p}_t < \hat{\varepsilon}_{o,t}](1 - A_t)}{\sum_{t=1}^T \mathbb{1}[\hat{p}_t < \hat{\varepsilon}_{o,t}]} \\ &= \frac{\frac{1}{T} \sum_{t=1}^T \mathbb{1}[\hat{p}_t < f_m(\hat{\mathbf{P}}_t)](1 - A_t)}{\frac{1}{T} \sum_{t=1}^T \mathbb{1}[\hat{p}_t < f_m(\hat{\mathbf{P}}_t)]}, \end{aligned}$$

The decision process $(\mathbb{1}[\hat{p}_t < \hat{\varepsilon}_{o,t}])_t$ and the false positives process $(\mathbb{1}[\hat{p}_t < \hat{\varepsilon}_{o,t}]A_t)_t$ are not independent, therefore it is not possible to use the Law of

Large Numbers directly. The alternative strategy consists first in splitting the numerator and denominator into several disjoint subseries corresponding to independent and identically distributed processes. Then partitioning the times series of length $T = T'(n + m)$ into T' subseries, each of length $n + m$, it results that

$$\begin{aligned} & \frac{1}{T} \sum_{t=1}^T \mathbb{1}[\hat{\rho}_t < f_m(\hat{\mathbf{P}}_t)](1 - A_t) \\ &= \frac{1}{n + m} \sum_{k=1}^{n+m} \left(\frac{1}{T'} \sum_{t=0}^{T'-1} \mathbb{1}[\hat{\rho}_{t(n+m)+k} < f_m(\hat{\mathbf{P}}_{t(n+m)+k})](1 - A_{t(n+m)+k}) \right). \end{aligned} \quad (\text{B.9})$$

Interestingly for each k from 1 to $m+n$, the summands within the brackets do all belong to different subseries, which makes the sum over t a sum of independent and identically distributed random variables. It results that, for each $1 \leq k \leq n + m$, the average within the brackets is converging to its expectation by the LLN theorem.

Since the limit of a (finite) sum is equal to the sum of the limits, the average in Eq. (B.9) is converging and

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{1}[\hat{\rho}_t < f_m(\hat{\mathbf{P}}_t)](1 - A_t) = \sum_{k=1}^m \mathbb{E} \left[\mathbb{1}[\hat{\rho}_k < f_m(\hat{\mathbf{P}}_k)](1 - A_k) \right] \text{ a.s.} \quad (\text{B.10})$$

$$= m \mathbb{E} \left[\mathbb{1}[\hat{\rho}_m < f_m(\hat{\mathbf{P}}_m)](1 - A_m) \right] \text{ a.s.} \quad (\text{B.11})$$

Furthermore,

$$FP_1^m(\hat{\varepsilon}_o, \hat{\mathbf{P}}_m) = \sum_{k=1}^m \mathbb{1}[\hat{\mathbf{P}}_{m,k} < f_m(\hat{\mathbf{P}}_m)](1 - A_k)$$

Using the exchangeability of $\hat{\mathbf{P}}_m$ and the permutation invariance of f_m is comes:

$$\mathbb{E}FP_1^m(\hat{\varepsilon}_o, \hat{\mathbf{P}}_1) = m \times \mathbb{E}[\mathbb{1}[\hat{\mathbf{P}}_{m,m} < f_m(\hat{\mathbf{P}}_m)](1 - A_m)]$$

Finally with $\hat{\mathbf{P}}_{m,m} = \hat{\rho}_m$ and Eq. B.11 it gives:

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{1}[\hat{\rho}_t < f_m(\hat{\mathbf{P}}_m)](1 - A_t) = \mathbb{E}FP_1^m(\hat{\varepsilon}_o, \hat{\mathbf{P}}_m)$$

Then after applying the same reasoning on the denominator, it gives:

$$FDP_1^\infty(\hat{\varepsilon}_o, \hat{\rho}) = \frac{\mathbb{E}FP_1^m(\hat{\varepsilon}_o, \hat{\mathbf{P}}_m)}{\mathbb{E}R_1^m(\hat{\varepsilon}_o, \hat{\mathbf{P}}_m)}$$

□

B.5. Proof of Proposition 2

Lemma 3.3. Let $(p_i)_{1 \leq i \leq m}$ be a sequence of m p -values with m_0 true null hypothesis. Suppose i belong to the set of true negative \mathcal{H}_0 , and D_i the random variable equal to 1 if i is detected by the BH_α procedure.

- If the p -values are independent and verify that: $\forall k \in \llbracket 1, m \rrbracket, \mathbb{P}(p_i \leq \frac{\alpha k}{m})$, then:

$$\mathbb{P}(D_i) = \frac{\alpha \mathbb{E}[R^*]}{m} \quad (\text{B.12})$$

Where R is the number of hypotheses rejected by BH_α and R^* is the number of hypotheses rejected after p_i is set to 0.

- If the p -values are empirical p -values using a unique calibration set with cardinality $\nu_\alpha^m - 1$, then:

$$\mathbb{P}(D_i) = \frac{\alpha \mathbb{E}[\tilde{R}^*]}{m} \quad (\text{B.13})$$

Where \tilde{R}^* is the number of rejected hypothesis by applying BH_α after on $(p'_j)_{1 \leq j \leq m}$, where $p'_i = 0$, and for $j \neq i, p'_j = p_j - \frac{1}{n} \mathbb{1}[p_j < p_i]$.

Proof of Lemma 3.3. Proof of the first statement:

Using, the random variable R representing the number of rejection of BH_α , i is rejected if p_i is below the threshold $\frac{\alpha R}{m}$.

$$\mathbb{P}[D_i] = \mathbb{E}[\mathbb{1}[p_i \leq \frac{\alpha R}{m}]] \quad (\text{B.14})$$

Let $(p'_j)_{1 \leq j \leq m}$ be defined by $p'_i = 0$ and $p'_j = p_j$. The conditions of Lemma D6 from [21] are satisfied, it follows:

$$\mathbb{P}[D_i] = \mathbb{E}[\mathbb{1}[p_i \leq \frac{\alpha R(i)}{m}]] \quad (\text{B.15})$$

Where $R(i)$ is the number detection when applying BH_α on $(p'_j)_{1 \leq j \leq m}$.

According to the law of total expectation:

$$\mathbb{P}[D_i] = \mathbb{E}[\mathbb{E}[\mathbb{1}[p_i \leq \frac{\alpha R(i)}{m}] | p \setminus p_i]]. \quad (\text{B.16})$$

Since $R(i)$ is measurable is $p \setminus p_i$ and p_i is independent from $p \setminus p_i$, it gives:

$$\mathbb{P}[D_i] = \mathbb{E}[\mathbb{P}_{p_i}(p_i \leq \frac{\alpha R(i)}{m})] \quad (\text{B.17})$$

By hypothesis $\mathbb{P}_{p_i}(p_i \leq \frac{\alpha R(i)}{m}) = \frac{\alpha R(i)}{m}$, then:

$$\mathbb{P}[D_i] = \frac{\alpha}{m} \mathbb{E}[R(i)] \quad (\text{B.18})$$

which conclude the proof of the first statement.

Proof of the second statement:

Let $W_i, C_{i,j}$ defined as:

$$W_i = (\{s_1, \dots, s_n, s_{n+i}\}, (s_i, i \in \mathcal{H}_0, i \neq j), (s_i, i \in \mathcal{H}_1)) \quad (\text{B.19})$$

$$C_{i,j} = \frac{1}{n} \left(\sum_{s \in \{s_1, \dots, s_n, s_{n+i}\}} \mathbb{1}[s > s_{n+j}] - 1 \right) \quad (\text{B.20})$$

1. $p_j = C_{i,j} + \frac{1}{n} \mathbb{1}[s_{n+j} > s_{n+i}]$
2. p_i independent of W_i
3. p_i follow uniform distribution in $\{0, \frac{1}{n}, \dots, 1\}$

Lemma D.6 from [21] is applied. Let $(p'_j)_{1 \leq j \leq m}$ be defined by $p'_i = 0$ and for $j \neq i$, $p'_j = C_{i,j}$. For all j , $p'_j \leq p_j$ and if $p_j > p_i$ then $p'_j = p_j$. Thus, the conditions of the Lemma D.6 from [21] are verified, which gives:

$$\mathbb{1}[p_i \leq \frac{\alpha R}{m}] = \mathbb{1}[p_i \leq \frac{\alpha \tilde{R}(i)}{m}] \quad (\text{B.21})$$

$$\mathbb{P}[D_i] = \mathbb{E}[\mathbb{E}[p_i \leq \frac{\alpha \tilde{R}(i)}{m} | W_i]] \quad (\text{B.22})$$

Since, $\tilde{R}(i)$ in measurable in W_i and p_i is independent of W_i .

$$\mathbb{P}[D_i] = \mathbb{E}[\mathbb{P}_{p_i}(p_i \leq \frac{\alpha \tilde{R}(i)}{m})] \quad (\text{B.23})$$

By hypothesis, p_i is a empirical p -value with calibration set verified that there exist an integer ν such that $n = \nu \frac{m}{\alpha} - 1$. So according to Corollary 1, $\mathbb{P}_{p_i}(p_i \leq \frac{\alpha \tilde{R}(i)}{m}) = \frac{\alpha \tilde{R}(i)}{m}$.

Finally, the second statement is verified with:

$$\mathbb{P}[D_i] = \frac{\alpha}{m} \mathbb{E}[\tilde{R}(i)] \quad (\text{B.24})$$

□

Lemma 3.4. *Let X and Y be two random variables. Suppose that $k \mapsto \mathbb{E}[X|Y = k]$ is decreasing, then:*

$$\mathbb{E}[XY] \leq \mathbb{E}[X]\mathbb{E}[Y]. \quad (\text{B.25})$$

Proof of Lemma 3.4. Let Z be a random variable that follows the same law than Y but is independent. Since $k \mapsto \mathbb{E}[X|Y = k]$ is decreasing:

$$(Y - Z)(\mathbb{E}[X|Y] - \mathbb{E}[X|Z]) \leq 0 \quad (\text{B.26})$$

$$\mathbb{E}[(Y - Z)(\mathbb{E}[X|Y] - \mathbb{E}[X|Z])] \leq 0 \quad (\text{B.27})$$

$$(\text{B.28})$$

By distributing the product and using that Y and Z follow the same law, this gives:

$$2\mathbb{E}[Y\mathbb{E}[X|Y]] - 2\mathbb{E}[Y\mathbb{E}[X|Z]] \leq 0 \quad (\text{B.29})$$

$$(\text{B.30})$$

Finally using $\mathbb{E}[Y\mathbb{E}[X|Y]] = \mathbb{E}[XY]$ and independence of Y and Z :

$$\mathbb{E}[XY] \leq \mathbb{E}[X]\mathbb{E}[Y] \quad (\text{B.31})$$

$$(\text{B.32})$$

□

Proof of Proposition 2. The mFDR formula is given by

$$mFDR_1^m(p) = \frac{\mathbb{E}[FP_{1,\alpha}^m(p)]}{\mathbb{E}[R_{1,\alpha}^m(p)]}.$$

Let us compute the numerator $\mathbb{E}[FP_1^m(p)]$ value after applying the BH_α . Keep in mind that here the family of true null hypotheses is random, generated by $(A_i)_{1 \leq i \leq m}$. In order to meet the conditions of Lemma 3.3, it is possible to condition with respect to \mathcal{H}_0 . Using $D_i = \mathbb{1}[p_i \leq \frac{R}{m}\alpha]$, it appears that

$$FP_1^m(p) = \sum_{i \in \mathcal{H}_0} D_i.$$

Lemma 3.3 allows to calculate its conditional expectation:

$$\begin{aligned} \mathbb{E}[FP_{1,\alpha}^m | \mathcal{H}_0] &= \sum_{i \in \mathcal{H}_0} \frac{\alpha}{m} \mathbb{E}[R(i) | \mathcal{H}_0] \\ &= \frac{\alpha m_0}{m} \mathbb{E}[R^* | \mathcal{H}_0]. \end{aligned}$$

With R^* the number of rejection where one p -values is set to 0. Integrating with respect to \mathcal{H}_0 . Then using the fact that $m_0 = |\mathcal{H}_0|$ is measurable with respect to \mathcal{H}_0 .

$$\begin{aligned} \mathbb{E}[FP_{1,\alpha}^m] &= \mathbb{E} \left[\frac{\alpha m_0}{m} \mathbb{E}[R^* | \mathcal{H}_0] \right] \\ &= \alpha \mathbb{E} \left[\frac{m_0}{m} R^* \right] \end{aligned}$$

Finally, if $\mathbb{E}[R^*|m_0]$ is decreasing, Lemma 3.4 gives,

$$\mathbb{E}[FP_{1,\alpha}^m] \leq \alpha(1 - \pi)\mathbb{E}[R^*]$$

with $1 - \pi = \mathbb{E}\frac{m_0}{m}$, the proportion of data generated by the reference distribution. \square

B.6. Proof of Proposition 3

Proof of Proposition 3. By definition $mFDR_1^m = \frac{\mathbb{E}[FP_1^m]}{\mathbb{E}R_1^m}$, and $R_1^m = FP_1^m + TP_1^m$. With hypothesis the $mFDR$ is equal to α , this gives:

$$\begin{aligned} \alpha &= \frac{\mathbb{E}[FP_1^m]}{\mathbb{E}R_1^m} \\ \alpha &= \frac{\mathbb{E}[FP_1^m]}{\mathbb{E}[FP_1^m + TP_1^m]} \\ \alpha(\mathbb{E}[FP_1^m] + \mathbb{E}[TP_1^m]) &= \mathbb{E}[FP_1^m] \\ (\alpha - 1)\mathbb{E}[FP_1^m] &= -\alpha\mathbb{E}[TP_1^m] \\ \mathbb{E}[FP_1^m] &= \frac{\alpha}{1 - \alpha}\mathbb{E}[TP_1^m] \end{aligned}$$

Then, the expectation of true positives is expressed using the proportion of false negatives β , the proportion of anomaly π in the m observations, A_i the random variable equal to 1 if the observation X_i is an anomaly and D_i the random variable equal to 1 if the observation X_i is detected as anomaly :

$$\begin{aligned} \mathbb{E}[TP_1^m] &= \sum_{i=1}^m \mathbb{P}[A_i = 1 \text{ and } D_i = 1] \\ &= \sum_{i=1}^m \mathbb{P}[A_i = 1] \mathbb{P}[D_i = 1 | A_i = 1] \\ &= m\pi(1 - \beta) \end{aligned}$$

Therefore, the $\mathbb{E}[FP_1^m]$ can be expressed as:

$$\mathbb{E}[FP_1^m] = \frac{\alpha m \pi (1 - \beta)}{1 - \alpha}$$

So the $\mathbb{E}[R_1^m]$ is expressed as follows:

$$\begin{aligned} \mathbb{E}[R_1^m] &= \frac{\alpha m \pi (1 - \beta)}{1 - \alpha} + m\pi(1 - \beta) \\ &= \frac{m\pi(1 - \beta)}{1 - \alpha} \end{aligned}$$

\square

B.7. Proof of Corollary 5

Proof of Corollary 5. All conditions being satisfied Theorem 2 gives that:

$$FDR_1^\infty(\hat{\varepsilon}_{BH_{\alpha'}}, \hat{p}) = mFDR_1^m(\hat{\varepsilon}_{BH_{\alpha'}}, \hat{\mathbf{P}}_m) \quad (\text{B.33})$$

According to Proposition 2, if one of the 3 statements is true:

$$mFDR_1^m(\hat{\varepsilon}_{BH_{\alpha'}}, \hat{p}) \leq (1 - \pi)\alpha' \frac{\mathbb{E}[R_{\alpha'}^*]}{\mathbb{E}[R_{\alpha'}]}$$

By hypothesis $\alpha' \frac{\mathbb{E}[R_{\alpha'}^*]}{\mathbb{E}[R_{\alpha'}]} = \alpha$ which allows to conclude.

$$mFDR_1^m(\hat{\varepsilon}_{BH_{\alpha'}}, \hat{p}) \leq (1 - \pi)\alpha$$

□

B.8. Proof of Theorem 3

Proof of Theorem 3. The Corollary 5 gives the two properties that the p -values families has to verify to control FDR of the time series:

- The $\hat{\mathbf{P}}_t$ are identically distributed and independent when time distance is larger than $n + m$.
- For each t , $\hat{\mathbf{P}}_t$ are either true p -values, or empirical p -values with independent or unique calibration set.

In the following, these properties are verified for the different p -values.

1. The sequence of true p -value $\mathbb{P}_{X \sim \mathcal{P}_0}(a(X) > a(X_t))$ is i.i.d., because the time series mixture is i.i.d. Then $\hat{\mathbf{P}}_t = (\hat{p}_{t-m+1}, \dots, p_t)$ are independent for a time distance larger than m . Using the first statement of Corollary 5 FDR of the whole time series is controlled at level $(1 - \pi)\alpha$.
2. The sequence of empirical p -value is i.i.d., because the time series mixture i and the calibration sets are i.i.d. Then $\hat{\mathbf{P}}_t = (\hat{p}_{t-m+1}, \dots, p_t)$ are independent for a time distance larger than m . Using the second statement of Corollary 5 FDR of the whole time series is controlled at level $(1 - \pi)\alpha$.
3. This p -value family is not i.i.d. However, because the calibration are build using a sliding window of size n , two p -values subseries of length m , $\hat{\mathbf{P}}_{t_1}$ and $\hat{\mathbf{P}}_{t_2}$, are independent when $|t_1 - t_2| > m + n$. Then the third statement of Corollary 5 ensure that FDR of the whole time series is controlled at level $(1 - \pi)\alpha$.

□

Appendix C: Supplementary experiments

C.1. Overlapping Calibration set and FDR control

The following experiments aims at drawing a comparison between the FDR values in three scenarios: independent calibration sets, partially overlapping calibration sets with an overlap size driven by the value of sn (size of the shift), and the same calibration set for all empirical p -values. To be more specific, the calibration sets (and corresponding empirical p -values) were generated according to the following scheme. Each calibration set is of cardinality n . When moving from one calibration set to the next one, the shift size is equal to sn , where s in $[0, 1]$ is the proportion of independent data between calibration sets, resulting in an overlap of cardinality $(1 - s)n$. Therefore an overlap occurs as long as $s < 1$. All these ways to build the calibration sets are called “calibration sets strategies”.

1. The independent p -values (iid Cal.) are generated according to

$$\forall i \in \llbracket 1, m \rrbracket, \quad \mathbf{Z}_i \sim \mathcal{P}_0^n, \quad \hat{p}_{1,i} = \hat{p}\text{-value}(X_i, \mathbf{Z}_i). \quad (\text{C.1})$$

2. The p -values with the same calibration set (Same Cal.) are generated by

$$\forall i \in \llbracket 1, m \rrbracket, \quad \mathbf{Z} \sim \mathcal{P}_0^n, \quad \hat{p}_{2,i} = \hat{p}\text{-value}(X_i, \mathbf{Z}). \quad (\text{C.2})$$

3. The p -values with overlapping calibration sets (Over. Cal.) are generated given, for $0 < s < n$, by

$$\begin{aligned} \forall i \in \llbracket 1, m \rrbracket, \quad \mathbf{Z}_i &= \{Z_{\lfloor isn \rfloor + 1}, \dots, Z_{\lfloor isn \rfloor + n}\}, \quad \hat{p}_{3,i} = \hat{p}\text{-value}(X_i, \mathbf{Z}_i), \\ \text{and} \quad \{Z_{s+1}, \dots, Z_{\lfloor 2sn \rfloor + 1}, \dots, Z_{\lfloor msn \rfloor + 1}, \dots, Z_{\lfloor msn \rfloor + n}\} &\sim \mathcal{P}_0^{ms+n}. \end{aligned} \quad (\text{C.3})$$

According to these three scenarios, as s increases, the overlap cardinality decreases, which results in more and more (almost) independent calibration sets. This is illustrated by the empirical results collected in Table 5. For each calibration set strategy, presented in row, and for each calibration set cardinality in column, the estimated FDR is shown. In this experiment, the reference distribution \mathcal{P}_0 is the Gaussian $\mathcal{N}(0, 1)$ and the anomalies are equal to $\Delta = 4$. The number of tested p -values, noted m , is equal to 100 and m_1 , the number of anomalies, is equal to 1. On each sample, BH-procedure is applied with $\alpha = 0.1$ and the FDP is computed. Each FDR is estimated over 10^3 repetitions.

TABLE 5
FDR results with overlapping calibration sets

n	249	250	499	500	749	750	999	1000
Same Cal.	0.164	0.175	0.112	0.183	0.137	0.131	0.097	0.154
Over Cal. (s=0.1%)	0.167	0.174	0.100	0.156	0.138	0.125	0.093	0.140
Over Cal. (s=0.2%)	0.162	0.176	0.095	0.170	0.124	0.127	0.109	0.143
Over Cal. (s=0.5%)	0.163	0.166	0.110	0.170	0.116	0.132	0.111	0.149
Over Cal. (s=1%)	0.151	0.180	0.094	0.177	0.127	0.128	0.099	0.143
Over Cal. (s=2%)	0.164	0.180	0.108	0.179	0.133	0.140	0.097	0.143
Over Cal. (s=5%)	0.168	0.172	0.108	0.169	0.125	0.130	0.096	0.144
Over Cal. (s=10%)	0.165	0.181	0.104	0.185	0.122	0.140	0.105	0.146
Over Cal. (s=20%)	0.173	0.207	0.109	0.171	0.136	0.149	0.101	0.140
Over Cal. (s=50%)	0.180	0.187	0.103	0.183	0.121	0.128	0.094	0.143
iid Cal.	0.171	0.188	0.115	0.174	0.138	0.143	0.104	0.132

The values of n in the columns of Table 5 are chosen such that, for each pair of columns, the FDR value is smaller for the left column and larger for the right column (see Figure 2 for a visual illustration of this phenomenon). Table 5 illustrates that, in the context of the present numerical experiments, the FDR estimation is not too strongly impacted by the value of s (proportion of the overlap). To assert that the observed differences between FDR estimations in each column are not significant, permutation tests [12, 23] are performed. Under \mathcal{H}_{0n} hypothesis, the FDR are the same across all calibration set strategies for the calibration set cardinality n . Under \mathcal{H}_{1n} there are at least two calibration strategies leading to different FDR. The FDP samples that have been used to estimate the FDR are reused. The maximal gap between sample means is used as statistic. The test is performed using the function “permutation_test” from the Python library called Scipy. The significance level is fixed at 0.05. Since multiple tests are performed over the different cardinalities, the threshold for rejecting a hypothesis is 0.00625, according Bonferroni correction. The results are display in Table 6. All tested hypotheses have a p -values greater than the threshold 0.00625. There are no significant difference in the resulting FDR between the different proportions of overlapping in calibration sets. This would suggest that considering overlapping calibration sets should not worsen too much the control of false positives and negatives.

TABLE 6
 p -values resulting from permutations test

n	249	250	499	500	749	750	999	1000
p -value of the test	0.300	0.0326	0.572	0.313	0.588	0.435	0.735	0.690

Appendix D: Figures

D.1. Figures related to the experiment of Section 6.4

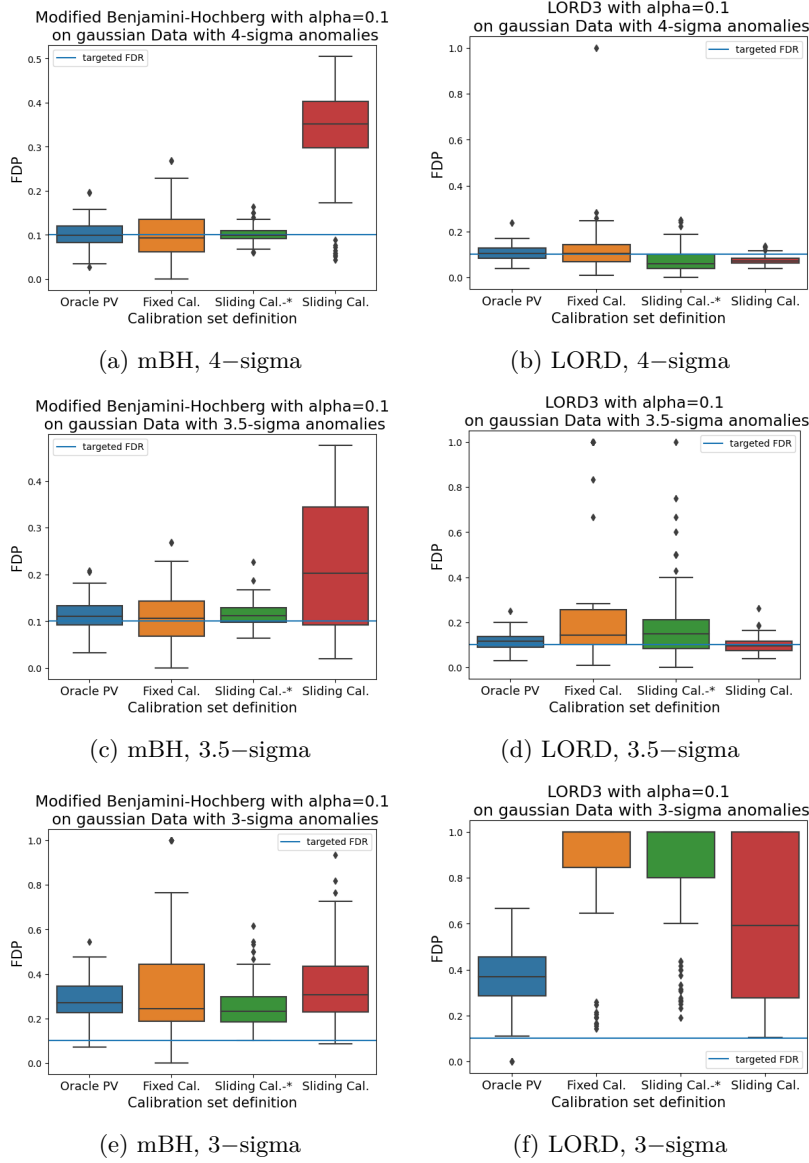


Figure 9: Comparison of FDPs acquired from using different multiple testing procedures, mBH or LORD, and from the way the p -values are calculated, in the case $\alpha = 0.1$.

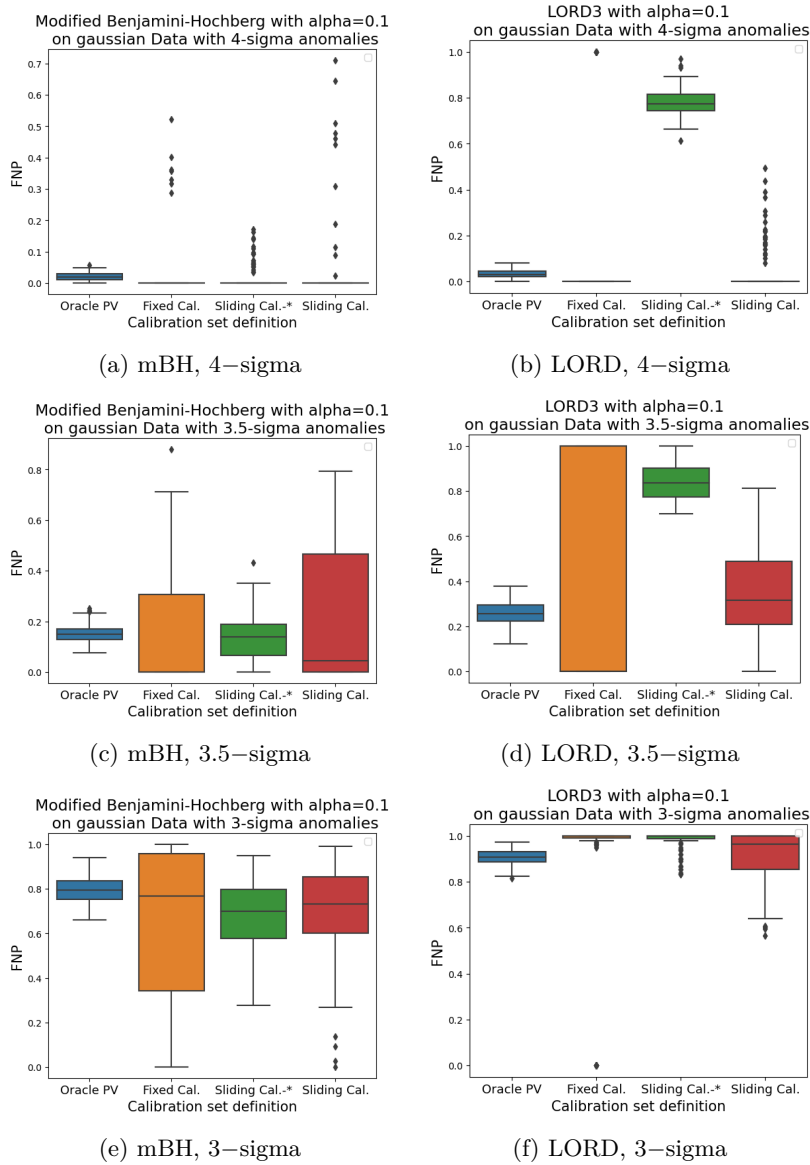


Figure 10: Comparison of FNPs acquired from using different multiple testing procedures, mBH or LORD, and from the way the p -values are calculated, in the case $\alpha = 0.1$.

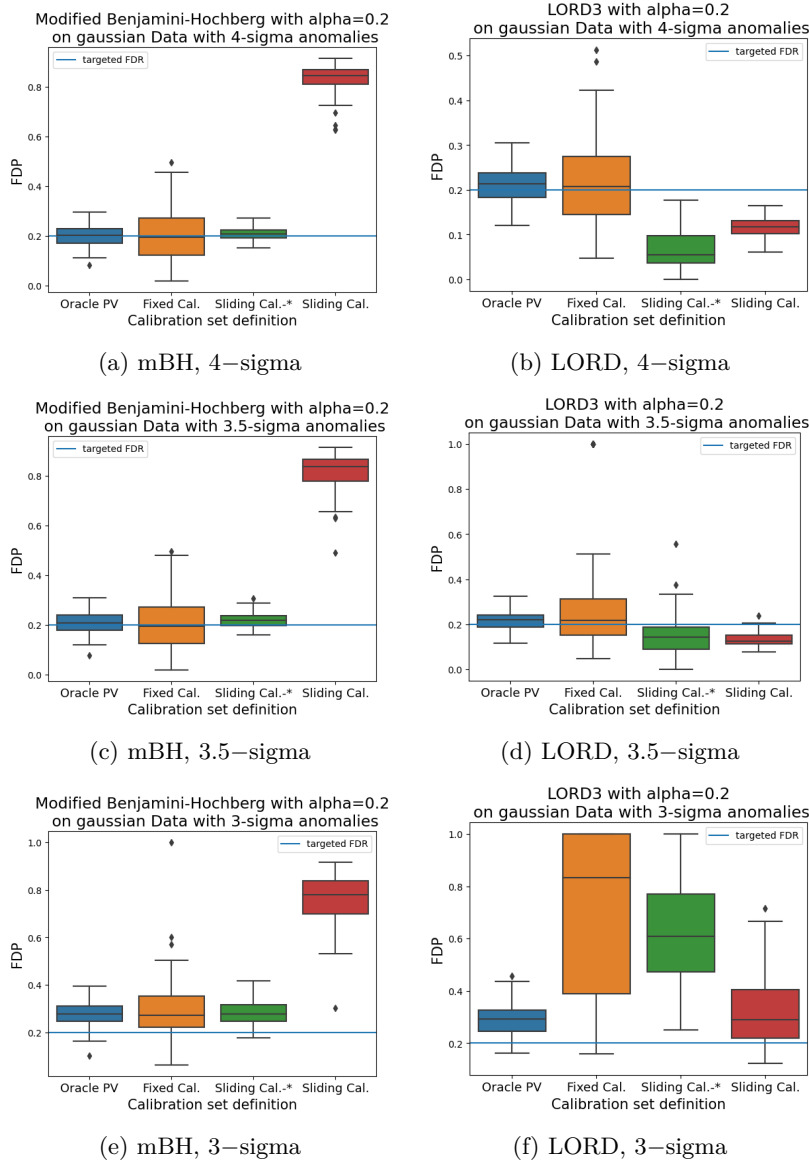


Figure 11: Comparison of FDPs acquired from using different multiple testing procedures, mBH or LORD, and from the way the p -values are calculated, in the case $\alpha = 0.2$

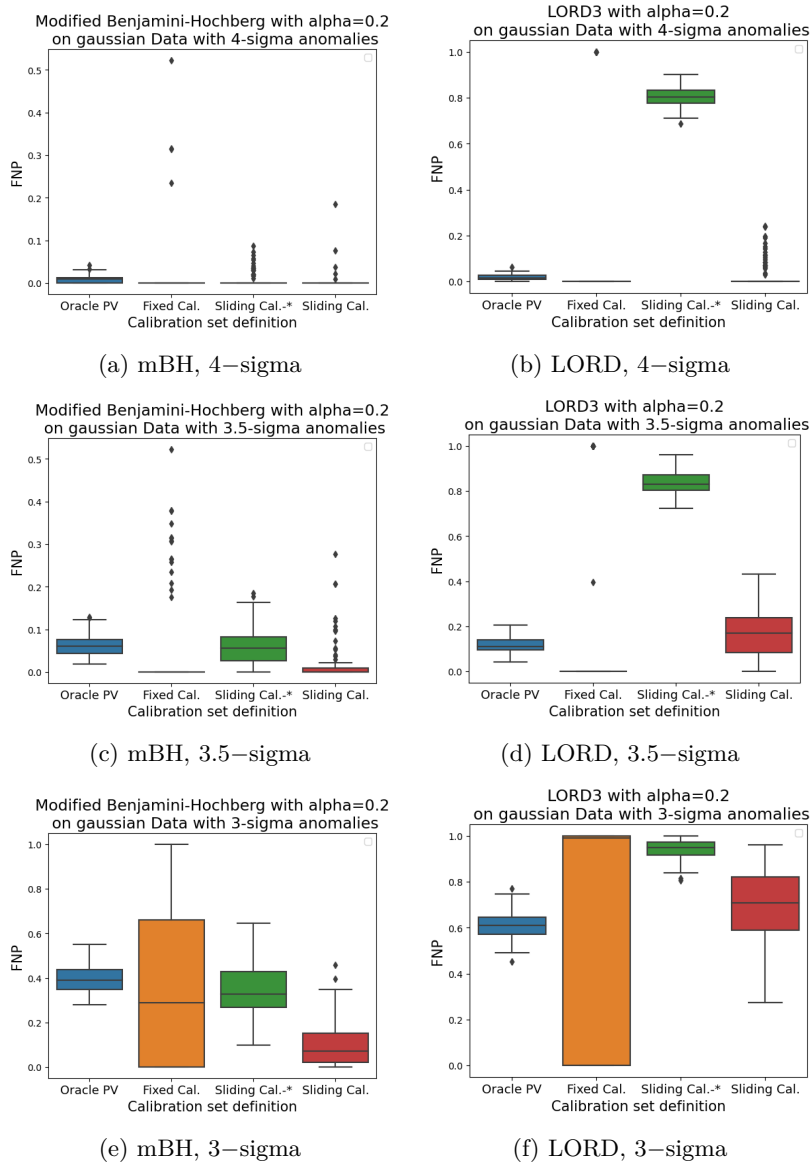


Figure 12: Comparison of FNPs acquired from using different multiple testing procedures, mBH or LORD, and from the way the p -values are calculated, in the case $\alpha = 0.2$.