



HAL
open science

Post hoc false discovery proportion inference under a Hidden Markov Model

Marie Perrot-Dockès, Gilles Blanchard, Pierre Neuvial, Etienne Roquain

► **To cite this version:**

Marie Perrot-Dockès, Gilles Blanchard, Pierre Neuvial, Etienne Roquain. Post hoc false discovery proportion inference under a Hidden Markov Model. *Test*, 2023, 10.1007/s11749-023-00886-7. hal-03214472

HAL Id: hal-03214472

<https://hal.science/hal-03214472>

Submitted on 1 May 2021

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.



Distributed under a Creative Commons Attribution 4.0 International License

Post hoc false discovery proportion inference under a Hidden Markov Model

Marie Perrot-Dockès*

Université de Paris, CNRS

MAP5 UMR 8145, Paris, France.

and

Gilles Blanchard

Université Paris-Saclay, CNRS, Inria

Laboratoire de mathématiques d'Orsay

91405, Orsay, France.

and

Pierre Neuvial

Institut de Mathématiques de Toulouse

UMR 5219, Université de Toulouse, CNRS

UPS, F-31062 Toulouse Cedex 9, France.

and

Etienne Roquain

Laboratoire de Probabilités, Statistique et Modélisation,

Sorbonne Université, Université de Paris, CNRS, Paris, France.

May 1, 2021

Abstract

We address the multiple testing problem under the assumption that the true/false hypotheses are driven by a Hidden Markov Model (HMM), which is recognized as a fundamental setting to model multiple testing under dependence since the seminal work of Sun and Cai (2009). While previous work has concentrated on deriving specific procedures with a controlled False Discovery Rate (FDR) under this model, following a recent trend in selective inference, we consider the problem of establishing confidence

*marie.perrot-dockees@u-paris.fr

bounds on the false discovery proportion (FDP), for a user-selected set of hypotheses that can depend on the observed data in an arbitrary way. We develop a methodology to construct such confidence bounds first when the HMM model is known, then when its parameters are unknown and estimated, including the data distribution under the null and the alternative, using a nonparametric approach. In the latter case, we propose a bootstrap-based methodology to take into account the effect of parameter estimation error. We show that taking advantage of the assumed HMM structure allows for a substantial improvement of confidence bound sharpness over existing agnostic (structure-free) methods, as witnessed both via numerical experiments and real data examples.

Keywords: post hoc bounds; hidden Markov model; false discovery proportion; posterior distribution, bootstrap.

1 Introduction

1.1 Context and motivation

To analyze large, heterogeneous and complex data, the analyst often adopts a *post hoc* approach, by choosing methods, formulating questions, selecting models or variables, on the basis of the data set. In particular, performing statistical inference after selection is a flourishing research field, often named as *selective inference*. The main challenge is to avoid the selection bias, by properly calibrating the error probabilities or risks.

Observations arising from applications are generally not independent. Hidden Markov Models (HMM) are a tool of choice to model stochastic processes with temporal or spatial dependence, and they have been widely successfully used in various areas including signal processing (Gales and Young, 2008) economics (Kim et al., 1999), or computational biology (Koski, 2001). This paper is motivated by a specific use case in genomics: the differential analysis of DNA copy number alterations (CNA) in cancer cells. Cancer cells are characterized by structural changes in the number of gene copies along the genome, and modern biotechnologies such as microarrays and sequencing are commonly used to quantify such changes at high resolution (Albertson et al., 2003). DNA copy number (CN) profiles are generally modeled as piece-wise constant signals, and HMM have been extensively used for this purpose (see e.g. Fridlyand et al., 2004; Shah et al., 2009; Zhang, 2010). Given

the observation of CN profiles along the genome for individuals classified in two groups corresponding to different types of cancer, differential analysis aims at identifying regions of the genome for which the CN profiles differ “significantly” between the two cancer types. As a case in point, we consider a study of ovarian cancers where a differential analysis of 117 patients with or without endometriosis is performed (Okamoto et al., 2015), see Section 5.2. This study comprises CNA measurements for 236,385 genomic locations (loci) for 63 patients without endometriosis and 54 patients with endometriosis. Figure 1 displays two-sided Wilcoxon test statistics of the null hypothesis of no signal difference between the two cancer types, for 4,799 loci on chromosome 7. We model these test statistics

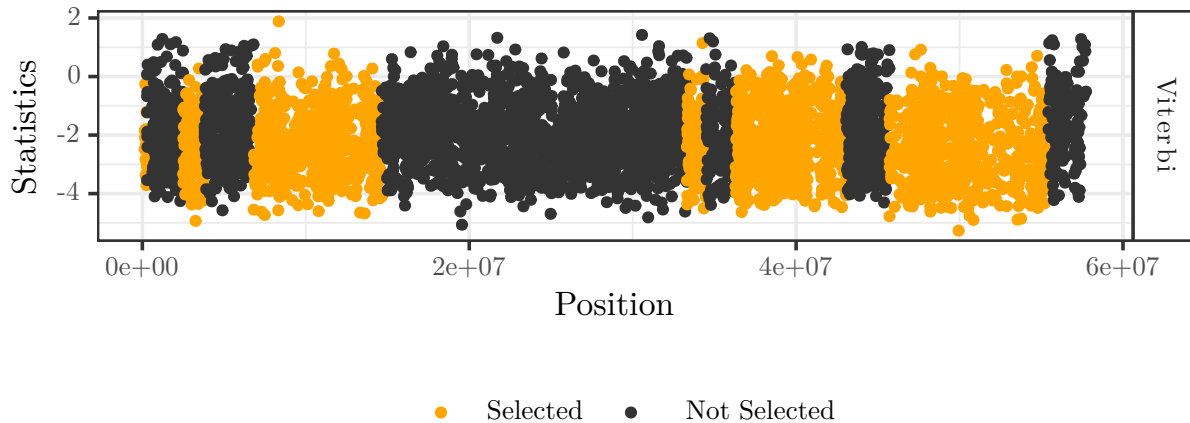


Figure 1: Wilcoxon test statistics for 13,239 loci on chromosome 7 in the Okamoto data set (Okamoto et al., 2015). Loci highlighted in orange correspond to a specific data-driven selection of regions.

as an HMM with two states, corresponding to the null and alternative hypotheses for the test. Loci highlighted in orange correspond to a specific data-driven selection of regions. The methods developed in this paper make it possible to construct confidence bounds on the false positives in such data-driven selections. They are both more informative than approaches based on the control of local False Discovery Rates (Sun and Cai, 2009), which only provide posterior point estimates, and more powerful than agnostic post hoc bounds (Blanchard et al., 2020) which do not take the dependency structure into account in the

inference. While the advantage of our approach will be put forward specifically with this data set (see Section 5.2), comprehensive numerical experiments will show that this superiority holds in a wide spectrum of cases, see Sections 4 and 5; typically anytime assuming an HMM structure is reasonable.

1.2 Post hoc bounds

The model considered here can be recast in the following more general setting: let us assume that the truth is carried by a *configuration vector* $\theta \in \{0, 1\}^m$ (which may be random or not), for which $\theta_i = 0$ if and only if the i -th variable is not active (in the above example $\theta_i = 1$ then means chromosomal aberration at position i). Then, for an observation X , and any set $R \subset \mathbb{N}_m := \{1, \dots, m\}$, we consider the following quantity of interest

$$\text{FDP}(\theta, R) = \frac{\sum_{i \in R} \mathbf{1}\{\theta_i = 0\}}{|R| \vee 1},$$

which is called the false discovery proportion. Since its introduction by Benjamini and Hochberg (1995), it has become a classical way of accounting for the errors made by the selection $R = S(X)$ in a multiple testing context.

Several ways of building confidence bounds for $\text{FDP}(\theta, S(X))$ have been considered in the literature, with varying interpretations and degrees of generality for the possible selection policies $S(\cdot) : x \mapsto S(x) \subset \mathbb{N}_m$ under scrutiny. Confidence bounds valid simultaneously over all subsets $R \subset \mathbb{N}_m$ have been considered by Genovese and Wasserman (2006); Goe-man and Solari (2011); Blanchard et al. (2020). Doing so, these bounds are also valid for any selection policy $S(\cdot)$. A counterpart of this uniform guarantee is conservativeness, that is, the obtained bound can be far from the true value $\text{FDP}(\theta, S(X))$ for a given particular selection policy $S(\cdot)$. Some improvements have been proposed in the literature by adding some local structure on the null hypotheses (Durand et al., 2020) or on considering subset restricted to a “path” of procedures of interest (Katsevich and Ramdas, 2018).

Other approaches have been developed specifically in the linear model with various error rates concerning inferences after a model selection policy $S(\cdot)$ (Scheffé, 1959; Berk et al.,

2013; Bachoc et al., 2018; Tibshirani et al., 2018; Bachoc et al., 2019). The LASSO policy is more specifically considered by Lee et al. (2016). Let us also mention that suitable confidence interval adjustments after selection have been proposed by Benjamini and Yekutieli (2005); Benjamini and Bogomolov (2014); Weinstein and Ramdas (2019).

1.3 Post hoc bounds in latent variables models

In this paper, we follow an approach very much in line with empirical Bayes methods for multiple testing, and in particular with the widely used two-group model and so-called “local fdr” method, see Efron et al. (2001); Efron (2008); Sun and Stephens (2018); Jin and Cai (2007); Sun and Cai (2009); Cai and Sun (2009); Cai and Jin (2010); Heller and Yekutieli (2014); Cai et al. (2019); Rebafka et al. (2019), among others (more details about these studies can be found in Sections 1.5 and 1.6). To this end, the configuration vector θ is assumed to be a vector of random latent variables. We show that this modeling can greatly simplify the construction of post hoc bounds, by using the distribution of the latent variable conditionally on X , that we call the posterior distribution in analogy with Bayesian terminology.

In the main part of the paper, following the seminal work of Sun and Cai (2009), we will assume that (θ, X) follows an HMM structure, with model parameter Γ . This model will be specified explicitly in Section 2. For the general considerations in the remainder of the present section though, it is sufficient to assume that we have a model parametrized by Γ for the joint distribution of (θ, X) , denoted by P_Γ , typically specified through the marginal distribution of the configuration vector θ , and the conditional distribution of the observation X conditionally on θ . The underlying distribution on the underlying probability space is denoted by \mathbb{P}_Γ as usual.

We consider the following aim: for any selection policy $S(\cdot) : x \in \mathcal{X} \mapsto S(x) \subset \mathbb{N}_m$, build a functional $I_\alpha(X, S(\cdot))$ valued in the intervals of $[0, 1]$, such that

$$\mathbb{P}_\Gamma (\text{FDP}(\theta, S(X)) \in I_\alpha(X, S(\cdot))) \geq 1 - \alpha. \quad (1)$$

We call $I_\alpha(X, S(\cdot))$ a post hoc interval family with selection policy $S(\cdot)$. Sometimes, the

post hoc interval $I_\alpha(X, S(\cdot))$ only requires knowledge of the selected region $S(X)$ on the observed data, in which case we will denote it by $I_\alpha(X, S(X))$ with an overload of notation.

As a particular case, intervals of the form $I_\alpha(X, S(\cdot)) = [0, \overline{\text{FDP}}_\alpha(X, S(\cdot))]$ can be considered, in which case (1) reduces to a post hoc upper-bound $\overline{\text{FDP}}_\alpha(X, S(\cdot))$ on $\text{FDP}(\theta, S(X))$. However, (1) also allows for considering post-hoc lower-bounds, which is also informative in practice, see Section 5.

1.4 Contribution: posterior post hoc bounds

A standard approach in multiple testing to tackle (1) is to derive a bound which holds conditionally on any value θ of the latent true null hypothesis configuration, i.e.

$$\mathbb{P}_\Gamma (\text{FDP}(\theta, S(X)) \in I_\alpha(X, S(\cdot)) \mid \theta) \geq 1 - \alpha, \quad \mathbb{P}_\Gamma - \text{ a.s.} \quad (2)$$

Note that this approach does not use at all the assumed model on the latent variables. To exploit this structural assumption, we consider instead solving (1) via conditioning with respect to X and solving

$$\mathbb{P}_\Gamma (\text{FDP}(\theta, S(X)) \in I_\alpha(X, S(\cdot)) \mid X) \geq 1 - \alpha, \quad \mathbb{P}_\Gamma - \text{ a.s.} , \quad (3)$$

which can be achieved by considering the *posterior distribution*, that is, the distribution of θ conditionally on X under \mathbb{P}_Γ .

If the parameter Γ governing the distribution of the latent variables is known, one can build a posterior interval only depending on $S(X)$ that fulfills (1): denoting $R = S(X)$, and provided $|R| > 0$,

$$I_\alpha(X, R) = [L_{\alpha\gamma}(X, R; \Gamma), U_{\alpha(1-\gamma)}(X, R; \Gamma)]; \quad (4)$$

$$U_\beta(X, R; \Gamma) = |R|^{-1} \min \left\{ n \in \{0, \dots, m\} : \mathbb{P}_\Gamma \left(\sum_{i \in R} (1 - \theta_i) \leq n \mid X \right) \geq 1 - \beta \right\}; \quad (5)$$

$$L_\beta(X, R; \Gamma) = |R|^{-1} \max \left\{ n \in \{0, \dots, m\} : \mathbb{P}_\Gamma \left(\sum_{i \in R} (1 - \theta_i) \geq n \mid X \right) \geq 1 - \beta \right\}, \quad (6)$$

for some $\gamma \in (0, 1)$ balancing the errors between the upper and lower bounds. We will sometimes drop the X in the notation for brevity.

This interval accounts for the particular HMM modeling via the posterior distribution, and thus is expected to be considerably sharper than the other intervals described in Section 1.2 and/or based on (2), that ignore this structure. Unfortunately, the functionals $L_\beta(X, S(X); \Gamma)$ and $U_\beta(X, S(X); \Gamma)$ are not directly accessible because the model parameter Γ is typically unknown.

We propose the following approximations of $U_\beta(S(X); \Gamma)$ (similar for $L_\beta(S(X); \Gamma)$):

- Plug-in: $U_\beta^{\text{PI}}(X, S(X)) = U_\beta(X, S(X); \hat{\Gamma})$, where $\hat{\Gamma}$ is an estimator of Γ . We will consider an estimator $\hat{\Gamma}$ based upon an iterative EM-type algorithm, see Section 2.2;
- Bootstrap 1: $U_\beta^{\text{boot1}}(X, S(\cdot))$, correcting the fluctuations of the above plug-in bounds by using a bootstrap approach generating resampled data X^* under $P_{\hat{\Gamma}}$. This bootstrap process requires the knowledge of the whole selection policy $S(\cdot)$, see Section 3.3.
- Bootstrap 2: $U_\beta^{\text{boot2}}(X, S(X))$, a heuristic approximation of $U_\beta^{\text{boot1}}(X, S(\cdot))$, which follows the same scheme, except that only $\hat{\Gamma}$ is resampled; in particular the selection set $S(X)$ is kept fixed during the bootstrap process. As a result, it does not require knowledge of the full selection policy $S(\cdot)$ but only of $R = S(X)$, the value of the selection policy for the observation X , see Section 3.4.
- Bootstrap 3: $U_\beta^{\text{boot3}}(X, S(\cdot))$, a bootstrap bound based on generating resampled data (θ^*, X^*) under $P_{\hat{\Gamma}}$ to approximate the distribution of $\text{FDP}(\theta, S(X))$, recentered with the plug-in bound $U_\beta(X, S(X); \hat{\Gamma})$, see Section 3.5. This process also requires the knowledge of the full selection policy.

These bounds have been implemented in the R package `SansSouci`, which is available from <https://pneuvial.github.io/sanssouci/>.

The coverage of these bounds is evaluated via extensive numerical experiments in Section 4, which typically reflect the impact of the estimation error of Γ in the different bounds. Our conclusion is that while the existing approaches, ignoring the latent HMM structure, are over-pessimistic, the plug-in approach can be slightly over-optimistic. A good trade-off

is provided by bootstrap-based strategies, which ensure a correct coverage while taking full advantage of the HMM structure.

1.5 Relation to posterior point estimates of the FDR

In previous literature, under a joint latent configuration/observation variable model P_Γ such as the one considered here, a commonly considered goal is to focus on the the FDR of a given selection policy $S(\cdot) : x \in \mathcal{X} \mapsto S(x) \subset \mathbb{N}_m$, that is,

$$\text{FDR}(S(\cdot), \Gamma) = \mathbb{E}_{(\theta, X) \sim P_\Gamma}(\text{FDP}(\theta, S(X))).$$

Observe that, since $\text{FDR}(S(\cdot), \Gamma) = \mathbb{E}_\Gamma(\mathbb{E}_\Gamma(\text{FDP}(\theta, S(X)) \mid X))$, the quantity

$$\widehat{\text{FDR}}(S(X), \Gamma) = \mathbb{E}_\Gamma(\text{FDP}(\theta, S(X)) \mid X) = |S(X)|^{-1} \sum_{i \in S(X)} \mathbb{P}_\Gamma(\theta_i = 0 \mid X) \quad (7)$$

is an unbiased estimator for $\text{FDR}(S(\cdot), \Gamma)$, for any selection policy $S(\cdot)$.

The marginal conditional probabilities $\mathbb{P}_\Gamma(\theta_i = 0 \mid X)$, $1 \leq i \leq m$, are generally referred to in the literature as the local FDR (Efron, 2008) or the ℓ -values (Castillo and Roquain, 2020), but without the explicit purpose of being applied to any selection policy to our knowledge. In fact, this quantity can be used for several purposes:

- for *any given* selection policy $S(\cdot)$, to compute an unbiased estimate of the FDR. Observe that this estimate only depends on the selected region $S(X)$, and not on the full selection policy $S(\cdot)$.
- to *design* a selection policy $S(\cdot)$ such that $\text{FDR}(S(\cdot), \Gamma)$ is equal to a specified level α . The simplest way to achieve this is to construct $S(X)$ so that $\widehat{\text{FDR}}(S(X), \Gamma)$ is constant equal to α . This is, for instance, the method proposed by Sun and Cai (2007, 2009); Cai and Sun (2009); Cai et al. (2019), where rejection regions taking the form of level sets of a relevant statistic $S_t(X) = \{i \in \mathbb{N}_m : T_i(X) \geq t\}$ are considered, and the threshold $t = t(X)$ is chosen also depending on X as the smallest value so that $\widehat{\text{FDR}}(S_{t(X)}(X), \Gamma) = \alpha$.

In practice, the parameter Γ is not known, so after introducing some suitable estimator $\widehat{\Gamma}$ of Γ (which can be built via an empirical Bayes approach), the plug-in quantity $\widehat{\text{FDR}}(S(X), \widehat{\Gamma})$ is used. In general, theoretical studies then ensure that the above properties hold in a suitable asymptotical sense, see Sun and Cai (2007, 2009); Cai and Sun (2009); Cai et al. (2019).

However, this approach only provides an (asymptotically) unbiased point estimate, and not a confidence interval for the FDP. The starting point of the present work is that we can use the same principle to derive confidence bounds for the FDP as in (4)-(5)-(6), which are more informative for practice, see Figures 4 and 5 in Section 5. Observe that, analogously to the two points above, the derived bounds can be used both to evaluate and design a policy.

1.6 Relation to other works

Multiple testing via a latent structure. Latent variables are widely used in statistics to design models with specific structures. In multiple testing, they have been used to model external “confounding factors” or “systemic effects” that induce dependencies between the tests, see Leek and Storey (2008); Friguet et al. (2009); Fan and Han (2017); Fan et al. (2019) and references of Sun et al. (2012) for genetic or genomic applications. In the present paper, the philosophy is rather different: latent variables are used to model dependencies between the null hypothesis configurations (here, between the coordinates of the configuration vector θ). While the unstructured case where the θ_i are i.i.d. can be traced back to the two-group model (Efron et al., 2001; Sun and Cai, 2007), the structured HMM case has been shown to improve over the independent case by Sun and Cai (2009). More recent examples of structures include two-sample sparsity (Cai et al., 2019) or stochastic block-models for graph-structured nulls (Rebafka et al., 2019), for which substantial improvements are also shown with respect to the unstructured case.

Empirical Bayes FDR methods. Our method, as all methods cited above, has a Bayesian flavor, which is the case in general for all latent-based multiple testing based on the use of “local fdr”, that is, the posterior probability that an item was generated under the null. While oracle versions (with a known model parameter) of such methods are known to be optimal in some way (see Sun and Cai, 2007 or more recently Heller and Rosset, 2021), a challenging part is to evaluate how the methods can handle parameter estimation, which is generally shown in an asymptotic manner. As already underlined in Section 1.5, all these methods have been developed for the FDR metric, not for FDP confidence intervals. Since the FDP metric considered in our post hoc bounds is considerably more difficult to analyze than the FDR one, obtaining a consistency result for the coverage of our bounds is out of scope here. It is left as a challenge for future investigations, see Section 6.2 for a discussion on this issue.

Non parametric inference in HMM. As said above, the feasibility of parameter estimation plays a crucial role in empirical Bayes multiple testing. The quality of estimation heavily relies on the structural assumptions made for the latent variable distribution and for the considered model for the observations conditionally on the latent variables. Adding structure allows to increase performance. For an HMM with non parametric emission densities, the situation is more favorable than in the unstructured two-group model, see Gassiat et al. (2016); Alexandrovich et al. (2016), with provable consistency guarantees for the “local fdr” values (De Castro et al., 2017; Abraham et al., 2021a).

Null estimation. Estimating the null density is both crucial and difficult in the two group model, as shown empirically by Efron (2004a, 2007, 2008, 2009) and further studied by Schwartzman (2010); Azriel and Schwartzman (2015); Stephens (2017); Sun and Stephens (2018). In particular, the recent work of Roquain and Verzelen (2020) theoretically shows that some sparsity is needed to obtain valid FDR inference when the null is estimated. By contrast, and as said above, the HMM assumption made in our setting makes this estimation much easier. In order to stick with the usual multiple testing line

of research, we will make a distinction between the situation where the null distribution is assumed to be known (which is studied in the core of this paper as it is the most standard case), and the case where it is not (which is studied more specifically in Appendix D). The numerical experiments of Section 4 confirm the validity of our bounds even in the latter case.

2 Hidden Markov modeling

In this section, we define the HMM modeling of our problem and add some more material that will be useful to deal with the post hoc bounds.

2.1 Model, notation and posterior distribution

We use here a setting close to the model of Sun and Cai (2009). Let us consider a hidden Markov process $(\theta, X) = ((\theta_i, X_i)_{1 \leq i \leq m})$, where

- $\theta = (\theta_1, \dots, \theta_m) \in \{0, 1\}^m$ is a unobserved latent variable sequence following a stationary Markov chain with matrix transition

$$A = \begin{pmatrix} a_{0,0} & a_{0,1} \\ a_{1,0} & a_{1,1} \end{pmatrix},$$

where $a_{q,\ell} \in (0, 1)$ for $q, \ell \in \{0, 1\}$, $a_{q,0} + a_{q,1} = 1$ for $q \in \{0, 1\}$ with $a_{0,0} \neq a_{1,0}$ to ensure that A is full rank.

- conditionally on θ , the observation $X = (X_i)_{1 \leq i \leq m} \in \mathbb{R}^m$ has independent coordinates and for all $i \in \{1, \dots, m\}$,

$$X_i | \theta_i \sim \begin{cases} f_0 & \text{if } \theta_i = 0 \\ f_1 & \text{if } \theta_i = 1 \end{cases},$$

where f_0 and f_1 are two distinct densities on \mathbb{R} (with respect to the Lebesgue measure).

Following Gassiat et al. (2016), this model is identifiable (which is a remarkable result since f_0, f_1 can be arbitrary densities; it is a nonparametric model). Moreover, since the sequence θ is assumed to be stationary, this implicitly means that θ_1 is generated according to the marginal stationary distribution π on $\{0, 1\}$ associated to A . Also, since we consider a testing paradigm where f_0 is the null distribution whereas f_1 is the alternative distribution, the roles of f_0 and f_1 are not exchangeable. We will consider both of the following cases:

- f_0 is known, in which case the parameter is $\Gamma = (A, f_1)$. In that case, the multiple testing task aims at finding observations X_i that depart from the distribution f_0 ;
- f_0 is unknown, in which case the parameter is $\Gamma = (A, f_0, f_1)$. In this case, since we do not know the distribution of the nulls, the task is detect ‘outliers’, that are X_i ’s with abnormal behavior. To distinguish f_0 from f_1 , a classical assumption is $\mathbb{P}(\theta_1 = 0) > \mathbb{P}(\theta_1 = 1)$, that is, f_0 is the predominant class in the sample X .

In our model, the distribution of θ conditionally on X (the posterior distribution), is well known and is based on a quantity $\ell_{1,1}(\Gamma)$ and transition matrices $\Pi_i(\Gamma)$, $2 \leq i \leq m$, whose explicit values are deferred to Appendix A.

Proposition 2.1 (Proposition 3.3.2 of Cappé et al., 2006). *In the model described in Section 2.1, the distribution of $\theta | X$ is a heterogeneous Markov Chain with transition matrices $\Pi_i(\Gamma)$, $2 \leq i \leq m$ given by (32) and with initial distribution a Bernoulli distribution with parameter $\ell_{1,1}(\Gamma)$ given by (27).*

2.2 Parameter estimation

Since our posterior bounds depend on the unknown parameter Γ , building an estimator $\widehat{\Gamma} = (\widehat{A}, \widehat{f}_1)$ of $\Gamma = (A, f_1)$ is crucial to obtain a computable bound. In this section, we only deal with the case where f_0 is known, the other case being similar and postponed to Appendix D.

For this, we make use of a pseudo Expectation-Maximization (EM) algorithm, combined with a weighted kernel estimator of f_1 , as proposed by Robin et al. (2007) and studied by

Nguyen and Matias (2014). In our framework, this estimator can be written as follows: starting from an initial guess $\widehat{\Gamma}^{(t-1)}$ of Γ , let

$$\widehat{f}_1^{(t)}(x) = \sum_{i=1}^m \ell_{i,1}(\widehat{\Gamma}^{(t-1)}) \frac{K((x - X_i)/h)}{h} \Big/ \sum_{i=1}^m \ell_{i,1}(\widehat{\Gamma}^{(t-1)}), \quad (8)$$

where K is a kernel function. We will typically use the Gaussian kernel $K(u) = (2\pi)^{-1/2}e^{-u^2/2}$, $u \in \mathbb{R}$. The rationale behind this is that weighting by $\ell_{i,1}(\widehat{\Gamma}^{(t-1)})$ will put more weights to the observations X_i for which θ_i is likely to be equal to 1, that is, the observation under the alternative. Doing so, $\widehat{f}_1^{(t)}(x)$ is close to an ideal standard kernel estimator of the density of an i.i.d. sample that would be based on the observations X_i under the alternative only. This gives rise to Algorithm 1.

Algorithm 1: EM-type algorithm to derive $\widehat{\Gamma} = (\widehat{A}, \widehat{f}_1)$

Input: $X = (X_1, \dots, X_m)$

Output: $\widehat{\Gamma} = (\widehat{A}, \widehat{f}_1)$ estimator of $\Gamma = (A, f_1)$

- Initialization: first guess $\widehat{\Gamma}^{(0)} = (\widehat{A}^{(0)}, \widehat{f}_1^{(0)})$ of Γ
 - Loop: at step $t \geq 1$, given $\widehat{\Gamma}^{(t-1)}$, do
 - Compute
 1. $\alpha_i(q) = \alpha_i(q; \widehat{\Gamma}^{(t-1)})$ and $\beta_i(q) = \beta_i(q; \widehat{\Gamma}^{(t-1)})$, $1 \leq i \leq m$, $q \in \{0, 1\}$, using the forward-backward algorithm (28), (29);
 2. $\ell_{i,q}(\widehat{\Gamma}^{(t-1)})$, $1 \leq i \leq m$, $q \in \{0, 1\}$, by (27);
 3. $\ell_{i,q,q'}(\widehat{\Gamma}^{(t-1)})$, $1 \leq i \leq m$, $q, q' \in \{0, 1\}$, by (30).
 - Update
 1. $\widehat{a}_{q,q'}^{(t)} = \sum_{i=1}^{m-1} \ell_{i,q,q'}(\widehat{\Gamma}^{(t-1)}) / \sum_{i=1}^{m-1} \ell_{i,q}(\widehat{\Gamma}^{(t-1)})$, for $0 \leq q, q' \leq 1$;
 2. $\widehat{A}^{(t)} = \begin{pmatrix} \widehat{a}_{0,0}^{(t)} & \widehat{a}_{0,1}^{(t)} \\ \widehat{a}_{1,0}^{(t)} & \widehat{a}_{1,1}^{(t)} \end{pmatrix}$;
 3. $\widehat{f}_1^{(t)}$ by (8);
 4. $\widehat{\Gamma}^{(t)} = (\widehat{A}^{(t)}, \widehat{f}_1^{(t)})$.
 - Stop the loop if convergence, that is, $\widehat{\Gamma}^{(t)}$ close enough to $\widehat{\Gamma}^{(t-1)}$.
 - Return $\widehat{\Gamma} = \widehat{\Gamma}^{(t)}$
-

3 Posterior post hoc confidence bounds

We first present the oracle bound, that is, the one using the true model parameter Γ . Then, we introduce several options for estimating this oracle, firstly based on a rough “plug-in” via the estimator $\widehat{\Gamma}$ and then refinements based on bootstrap approaches. Also, we only present formulae for the upper bounds U_β for brevity. The corresponding lower bounds L_β

are similar and quickly described in Appendix C.

3.1 Oracle bound

In this section, we fix any non-empty selected set $R = S(X) = \{j_1, \dots, j_s\} \subset \mathbb{N}_m$, for some indices $1 \leq j_1 < \dots < j_s \leq m$ and $s = |R|$. Here, while the case where R consists of contiguous indices is typical, we consider the more general situations where R is arbitrary. For notation brevity, we let $\theta_t^R = \theta_{j_t}$ for $t \in \{1, \dots, s\}$.

According to (5), we let

$$U_\beta(X, R; \Gamma) = s^{-1} \min \left\{ n \in \{0, \dots, m\} : \mathbb{P}_\Gamma \left(\sum_{t=1}^s (1 - \theta_t^R) \leq n \mid X \right) \geq 1 - \beta \right\}. \quad (9)$$

Proposition 2.1 ensures that, conditionally on X , $(\theta_t^R)_{1 \leq t \leq s}$ follows an heterogeneous Markov chain with initial probability $\ell_{j_1, 1}(\Gamma)$ and transition matrices

$$\Pi_t^R(\Gamma) = \prod_{i=j_{t-1}+1}^{j_t} \Pi_i(\Gamma), \quad \Pi_t^R(\Gamma) = \begin{pmatrix} \Pi_{t,0,0}^R(\Gamma) & \Pi_{t,0,1}^R(\Gamma) \\ \Pi_{t,1,0}^R(\Gamma) & \Pi_{t,1,1}^R(\Gamma) \end{pmatrix}, \quad t \in \{2, \dots, s\}. \quad (10)$$

As such, $U_\beta(X, R; \Gamma)$ is not explicit. In the sequel, we provide an algorithm to compute $U_\beta(X, R; \Gamma)$. For this, let for $1 \leq k \leq s$, $0 \leq \ell \leq s$,

$$B_{k,\ell,0} = \mathbb{P}_\Gamma \left(\sum_{t=1}^k (1 - \theta_t^R) \leq \ell, \theta_k^R = 0 \mid X \right); \quad (11)$$

$$B_{k,\ell,1} = \mathbb{P}_\Gamma \left(\sum_{t=1}^k (1 - \theta_t^R) \leq \ell, \theta_k^R = 1 \mid X \right). \quad (12)$$

In words, $B_{k,\ell,0}$ is the posterior probability that there are at most ℓ zero-occurrences in $\theta_{1:k}^R$, with a zero in the last position. Similarly, $B_{k,\ell,1}$ is the posterior probability that there are at most ℓ zero-occurrences in $\theta_{1:k}^R$, with a one in the last position. Since $B_{s,\ell,0} + B_{s,\ell,1}$ is the posterior probability that at most ℓ zero-occurrences occurs in the whole sequence $\theta^R = \theta_{1:s}^R$, the following holds.

Proposition 3.1. *The quantity $U_\beta(X, R; \Gamma)$ defined by (9) can be computed as*

$$U_\beta(X, R; \Gamma) = s^{-1} \min \{ n \in \{0, \dots, m\} : B_{s,n,0} + B_{s,n,1} \geq 1 - \beta \}, \quad (13)$$

where $B_{k,\ell,0}$ and $B_{k,\ell,1}$ are defined by (11) and (12), respectively.

We present below an explicit algorithm to compute the quantities $B_{k,\ell,0}$ and $B_{k,\ell,1}$, see Algorithm 2. In words, (14) comes from the fact that having at most ℓ zero-occurrences in $\theta_{1:k}^R$ with a zero in the last position means that we have at most $\ell - 1$ zero-occurrences in $\theta_{1:(k-1)}^R$ with either a zero or a one in position $k - 1$. As for (15), the fact that having at most ℓ zero-occurrences in $\theta_{1:k}^R$ with a one in the last position means that we have at most ℓ zero-occurrences in $\theta_{1:(k-1)}^R$ with either a zero or a one in position $k - 1$.

Algorithm 2: Computation of $B_{k,\ell,0}$ and $B_{k,\ell,1}$.

Input: $\Pi_k^R(\Gamma)$, $1 \leq k \leq s$; $\ell_{j_k,q}(\Gamma)$, $1 \leq k \leq s$, $q \in \{0, 1\}$

Output: $B_{k,\ell,0}$ and $B_{k,\ell,1}$, $1 \leq k \leq s$, $0 \leq \ell \leq s$

Initialization $B_{k,0,0} = 0$, $1 \leq k \leq s$, $B_{1,\ell,0} = \ell_{j_1,0}(\Gamma)$, $1 \leq \ell \leq s$, $B_{1,\ell,1} = \ell_{j_1,1}(\Gamma)$,
 $0 \leq \ell \leq s$.

for $2 \leq k \leq s$ **do**

$B_{k,0,1} = B_{k-1,0,1} \Pi_{k,1,1}^R(\Gamma)$

for $1 \leq \ell \leq s$ **do**

$$B_{k,\ell,0} = B_{k-1,\ell-1,0} \Pi_{k,0,0}^R(\Gamma) + B_{k-1,\ell-1,1} \Pi_{k,1,0}^R(\Gamma); \quad (14)$$

$$B_{k,\ell,1} = B_{k-1,\ell,0} \Pi_{k,0,1}^R(\Gamma) + B_{k-1,\ell,1} \Pi_{k,1,1}^R(\Gamma). \quad (15)$$

end

end

3.2 Plug-in bound

The first non-oracle bound that is proposed is simply obtained by plugging the estimator derived in the Section 2.2 into the oracle bound, that is,

$$U_\beta^{\text{PI}}(X, S(X)) = U_\beta(X, S(X); \widehat{\Gamma}), \quad (16)$$

where $\widehat{\Gamma}$ comes from Algorithm 1 and $U_\beta(X, R; \Gamma)$ is given by (9) and (13). Since the oracle bound is based on the conditional distribution of the latent variable given the observation, the above bound can be interpreted as an ‘‘empirical Bayes credible set’’ for the FDP.

Unfortunately, this plug-in bound can be anti-conservative, meaning that it can violate (1), as we will see in the simulations of Section 4. An intuitive explanation is that $U_\beta(X, S(X); \widehat{\Gamma})$ is expected to fluctuate on both sides of $U_\beta(X, S(X); \Gamma)$. and thus the event $\text{FDP}(\theta, S(X)) \leq U_\beta(X, S(X); \widehat{\Gamma})$ may not be true when $U_\beta(X, S(X); \widehat{\Gamma})$ is smaller than $U_\beta(X, S(X); \Gamma)$.

In other words, the variability w.r.t. the posterior distribution of θ conditional on X is taken into account by U_β , but it is still needed to evaluate the variability in X of the bounds, which also relies both on the variability of the estimate $\widehat{\Gamma}$ of Γ and the policy $S(\cdot)$. Our first bootstrap bound takes both variations into account, while the second one focuses solely on the variation of $\widehat{\Gamma}$.

3.3 First bootstrap bound

Let us consider the deterministic quantity

$$q_{1,\gamma}(\beta, S(\cdot); \Gamma) = \min \left\{ x \in \mathbb{R} : \mathbb{P}_\Gamma \left(U_\beta(X, S(X); \Gamma) - U_\beta(X, S(X); \widehat{\Gamma}) \leq x \right) \geq 1 - \gamma \right\} \quad (17)$$

which corresponds to the $(1-\gamma)$ -quantile of the distribution of $U_\beta(X, S(X); \Gamma) - U_\beta(X, S(X); \widehat{\Gamma})$ when X is generated according to the true model parameter Γ . Then, we have by definition, for all $\delta \in (0, 1)$,

$$\begin{aligned} \mathbb{P}_\Gamma(\text{FDP}(\theta, S(X)) \leq U_{\beta(1-\delta)}(X, S(X); \Gamma) \mid X) &\geq 1 - \beta(1 - \delta) \\ \mathbb{P}_\Gamma(U_{\beta(1-\delta)}(X, S(X); \Gamma) - U_{\beta(1-\delta)}(X, S(X); \widehat{\Gamma}) \leq q_{1,\beta\delta}(\beta(1 - \delta), S(\cdot); \Gamma)) &\geq 1 - \beta\delta. \end{aligned}$$

Note that the first bound concerns the distribution of θ conditionally on X while the second one concerns only the marginal distribution of X . Therefore, this immediately implies a bound with respect to the joint distribution of (θ, X) .

Proposition 3.2. *For any selection policy $S(\cdot) : x \in \mathcal{X} \mapsto S(x) \subset \mathbb{N}_m$ and $\delta \in (0, 1)$, for any model parameter Γ , we have*

$$\mathbb{P}_\Gamma(\text{FDP}(\theta, S(X)) \leq U_{\beta(1-\delta)}(X, S(X); \widehat{\Gamma}) + q_{1,\beta\delta}(\beta(1 - \delta), S(\cdot); \Gamma)) \geq 1 - \beta. \quad (18)$$

Hence, the RHS of the event in (18) is a post hoc bound in the sense of (1). However, again, it depends on the unknown parameter Γ , although it is only via the quantity $q_{1,\gamma}(\beta, S(\cdot); \Gamma)$, that we would like to think of as a “second order” term.

Now, plugging the estimate $\widehat{\Gamma}$ of Γ in the latter leads to the bootstrap-type approximation $q_\gamma(\beta, S(\cdot); \Gamma) \approx q_\gamma(\beta, S(\cdot); \widehat{\Gamma})$ which corresponds to consider the $(1 - \gamma)$ -quantile of the distribution of $U_\beta(X^*, S(X^*); \widehat{\Gamma}) - U_\beta(X^*, S(X^*); \widehat{\Gamma}^*)$ under the model parameter $\widehat{\Gamma}$, where X^* is an independent sample generated under $P_{\widehat{\Gamma}}$ and $\widehat{\Gamma}^* = \widehat{\Gamma}(X^*)$, see (17). As usual, the bootstrap quantity $q_\gamma(\beta, S(\cdot); \widehat{\Gamma})$ is in turn approximated by a quantity $\tilde{q}_{1,\gamma}^{(B)}(\beta, S(\cdot); \widehat{\Gamma})$ obtained via a Monte-Carlo scheme that generates B times the sequence $(X_i^*)_{1 \leq i \leq m}$ according to the model with parameter $\widehat{\Gamma}$.

We are now in position to introduce our first bootstrap bound:

$$U_{\beta,\delta}^{\text{boot1}}(X, S(\cdot)) = U_{\beta(1-\delta)}(X, S(X); \widehat{\Gamma}) + \tilde{q}_{1,\beta\delta}^{(B)}(\beta(1-\delta), S(\cdot); \widehat{\Gamma}), \quad (19)$$

Algorithm 3 describes how this bound may be computed.

Let us make some comments on this bound:

- This bootstrap bound has a semi-parametric flavor: while the θ sequence is generated via a parametric Markov chain using the estimated parameters, the resampled data X_i^* are obtained by weighted smooth bootstrap (Efron, 1979), since they are drawn from a weighted kernel density estimator (see Appendix B.3 for more details).
- The term δ balances the way the errors are distributed and should be chosen to derive an appropriate trade-off: a small δ will sharpen the bound $U_{\beta(1-\delta)}(X, S(X); \widehat{\Gamma})$ but makes $\tilde{q}_{1,\beta\delta}^{(B)}(\beta(1-\delta), S(\cdot); \widehat{\Gamma})$ larger. In the numerical experiments, we have observed that the impact of δ is moderate.
- In this bound, we should compute $S(X^*)$, see Algorithm 4. This means that the user should provide their whole selection policy $S : x \in \mathcal{X} \mapsto S(x) \subset \mathbb{N}_m$, and not only the selection $S(X)$ on the observed data set X . This can be seen as a constraint in some cases, and the next paragraph provides a solution to circumvent this limitation.

- Whenever $\tilde{q}_{1,\beta\delta}^{(B)}(\beta(1-\delta), S(\cdot); \hat{\Gamma})$ is nonnegative (which happens with high probability for suitable values of β and δ), we have $U_{\beta,\delta}^{\text{boot1}}(X, S(\cdot)) \geq U_{\beta(1-\delta)}(S(X); \hat{\Gamma})$ so the obtained bound is in general more conservative than the plug-in bound.

Algorithm 3: Core algorithm for computing the bootstrap bounds

Input: Common input for all the bounds : Data X ; Estimator $\widehat{\Gamma}$ via Algorithm 1
if f_0 is known (or by Algorithm 4 if f_0 is unknown); $\beta \in (0, 1)$

- Input U^{boot1} : Selection policy $S(\cdot) : x \in \mathcal{X} \mapsto S(x) \subset \mathbb{N}_m$; , $\delta \in (0, 1)$
- Input U^{boot2} : $\delta \in (0, 1)$, $S(X)$
- Input U^{boot3} : Selection policy $S(\cdot) : x \in \mathcal{X} \mapsto S(x) \subset \mathbb{N}_m$

Output: $U_{\beta, \delta}^{\text{boot1}}(S(\cdot), \widehat{\Gamma})$ or $U_{\beta, \delta}^{\text{boot2}}(S(X), \widehat{\Gamma})$ or $U_{\beta}^{\text{boot3}}(S(\cdot), \widehat{\Gamma})$

1. Generate independently B bootstrap samples as follows: for $1 \leq b \leq B$
 - (a) Draw $\theta^{*(b)}$ and $X^{*(b)}$ from $\widehat{\Gamma}$
 - (b) Compute $\widehat{\Gamma}^{*(b)}$ using Algorithm 1 with the sequence $X^{*(b)}$ (or using Algorithm 4 if f_0 is unknown);
 - (c) Compute $D_i^{(b)}$, for the appropriate $i \in (1, 2, 3)$ and $\lambda \in (\beta(1 - \delta), \beta(1 - \delta), \beta)$;
 - U^{boot1} : $D_1^{(b)} = U_{\lambda}(X^{*(b)}, S(X^{*(b)}); \widehat{\Gamma}) - U_{\lambda}(X^{*(b)}, S(X^{*(b)}); \widehat{\Gamma}^{*(b)})$;
 - U^{boot2} : $D_2^{(b)} = U_{\lambda}(X, S(X); \widehat{\Gamma}) - U_{\lambda}(X, S(X); \widehat{\Gamma}^{*(b)})$;
 - U^{boot3} : $D_3^{(b)} = \text{FDP}(\theta^{*(b)}, S(X^{*(b)})) - U_{\lambda}(X^{*(b)}, S(X^{*(b)}); \widehat{\Gamma}^{*(b)})$,

where $U_{\lambda}(X, S(X); \Gamma)$ is given by (9) and (13)

2. Compute $\tilde{q}_{i, \gamma}^{(B)}(\lambda)$ as the empirical γ -quantile of $D_i = (D_i^{(1)}, \dots, D_i^{(B)})$ for the appropriate $\gamma \in (\beta\delta, \beta\delta, \beta)$;
3. Return the corresponding bound

- $U_{\beta, \delta}^{\text{boot1}}(X, S(\cdot), \widehat{\Gamma}) = U_{\beta(1-\delta)}(X, S(X), \widehat{\Gamma}) + \tilde{q}_{1, \beta\delta}^{(B)}(\beta(1 - \delta))$
 - $U_{\beta, \delta}^{\text{boot2}}(X, S(X), \widehat{\Gamma}) = U_{\beta(1-\delta)}(X, S(X), \widehat{\Gamma}) + \tilde{q}_{2, \beta\delta}^{(B)}(\beta(1 - \delta))$
 - $U_{\beta}^{\text{boot3}}(X, S(X), \widehat{\Gamma}) = U_{\beta}(X, S(X), \widehat{\Gamma}) + \tilde{q}_{3, \beta}^{(B)}(\beta)$
-

3.4 Second bootstrap bound

As mentioned above, computing $U_{\beta}^{\text{boot1}}(\cdot)$ requires that the user provides $S(X^{*(b)})$ for every single bootstrap sample $X^{*(b)}$, $b = 1, \dots, B$. This could be inconvenient if the user does not want to provide the whole selection policy, but only $S(X)$, the desired selection on the observed data set and not on other virtual data sets.

Here, we circumvent this limitation with a twist: imagine first that we have at hand a sample Y that is an independent copy of X . Then one could estimate Γ by an estimator $\widehat{\Gamma}(Y)$ computed from the sample Y . Therefore,

$$\begin{aligned} U_{\beta(1-\delta)}(X, S(X); \Gamma) &= U_{\beta(1-\delta)}(X, S(X); \widehat{\Gamma}(Y)) + \left(U_{\beta(1-\delta)}(X, S(X); \Gamma) - U_{\beta(1-\delta)}(X, S(X); \widehat{\Gamma}(Y)) \right) \\ &\leq U_{\beta(1-\delta)}(X, S(X); \widehat{\Gamma}(Y)) + q_{2, \beta\delta}(\beta(1-\delta), S(X); \Gamma), \end{aligned}$$

with probability larger than $1 - \beta\delta$, where we denote

$$\begin{aligned} q_{2, \gamma}(\beta, S(X); \Gamma) \\ = \min \left\{ x \in \mathbb{R} : \mathbb{P}_{Y \sim P_{\Gamma}} \left(U_{\beta(1-\delta)}(X, S(X); \Gamma) - U_{\beta(1-\delta)}(X, S(X); \widehat{\Gamma}(Y)) \leq x \mid X \right) \geq 1 - \gamma \right\}. \end{aligned}$$

the $(1 - \gamma)$ -quantile of the distribution of $U_{\beta(1-\delta)}(X, S(X); \Gamma) - U_{\beta(1-\delta)}(X, S(X); \widehat{\Gamma}(Y))$ conditionally on X , when $Y \sim P_{\Gamma}$. Similarly to above, the latter can be approximated via a bootstrapped quantity $q_{2, \gamma}(\beta, S(X); \Gamma) \approx q_{2, \gamma}(\beta, S(X); \widehat{\Gamma}(Y))$ which is the γ -quantile of the distribution of $U_{\beta(1-\delta)}(X, S(X); \widehat{\Gamma}(Y)) - U_{\beta(1-\delta)}(X, S(X); \widehat{\Gamma}(Y^*))$ conditionally on X, Y , when $Y^* \sim P_{\widehat{\Gamma}}$, which itself is approximated by applying a Monte-Carlo approximation scheme, which gives rise to $\tilde{q}_{2, \gamma}^{(B)}(\beta, S(X); \widehat{\Gamma}(Y))$. Now, since we do not have access to a different sample Y , we propose to use X in place of Y , which leads to the following new bootstrap bound (denoting, again, $\widehat{\Gamma} = \widehat{\Gamma}(X)$):

$$U_{\beta, \delta}^{\text{boot2}}(X, S(X)) = U_{\beta(1-\delta)}(X, S(X); \widehat{\Gamma}) + \tilde{q}_{2, \beta\delta}^{(B)}(\beta(1-\delta), S(X); \widehat{\Gamma}), \quad (20)$$

Algorithm 3 gives full details about the computation of this bound.

Let us make the following comments on the second bootstrap bound:

- $\tilde{q}_{2, \beta\delta}^{(B)}(\beta(1-\delta), S(X); \widehat{\Gamma})$ does not depend on the full selection policy $S(\cdot) : x \in \mathcal{X} \mapsto S(x) \subset \mathbb{N}_m$, but only on the set $S(X)$;

- Whenever $\hat{q}_{2,\beta\delta}^{(B)}(\beta(1-\delta), S(X); \hat{\Gamma})$ is nonnegative the obtained bound is more conservative than the plug-in bound. However, while $U_{\beta,\delta}^{\text{boot}2}$ includes the variability in $\hat{\Gamma}$, this bound ignores the variations in $S(X)$, so is in general less conservative than $U_{\beta,\delta}^{\text{boot}1}$.

3.5 Third bootstrap bound

An elementary point is that the interval $I_\alpha(S(\cdot)) = [0, V_\alpha(S(\cdot); \Gamma)]$ satisfies the unconditional coverage (1) when choosing

$$q_\beta^{\text{naive}}(S(\cdot); \Gamma) = \min \{x \in \mathbb{R} : \mathbb{P}_\Gamma (\text{FDP}(\theta, S(X)) \leq x) \geq 1 - \beta\}.$$

Note that this bound is “unconditional” and based on the distribution of $\text{FDP}(\theta, S(X))$ when drawing $(\theta, X) \sim P_\Gamma$. Hence, it uses the selection policy $S(\cdot)$. This leads to choose

$$U_\beta^{\text{naive}}(S(\cdot)) = q_\beta^{\text{naive}}(S(\cdot); \hat{\Gamma}).$$

This bound thus relies on drawing independent couples (θ^*, X^*) from the distribution $P_{\hat{\Gamma}}$. This is different from the above-mentioned bounds (bootstrap 1, bootstrap 2), which are based only on the marginal X^* .

As we will see in Section 4, using U_β^{naive} is generally too conservative, in particular if the true $\text{FDP}(\theta, S(X))$ has large variability from a realization of (θ, X) to another, since it is based on an unconditional quantile of its distribution. To alleviate this effect, an improvement can be derived by using a proper re-centering by the plug-in bound $U_\beta(X, S(X); \hat{\Gamma})$ given in (16), that acts like a stabilization:

$$V_\beta(X, S(\cdot); \Gamma) = U_\beta(X, S(X); \hat{\Gamma}(X)) + q_{3,\beta}(\beta, S(\cdot); \Gamma) \tag{21}$$

$$q_{3,\beta}(\beta, S(\cdot); \Gamma) = \min \left\{ x \in \mathbb{R} : \mathbb{P}_{(\theta, X) \sim P_\Gamma} \left(\text{FDP}(\theta, S(X)) - U_\beta(X, S(X); \hat{\Gamma}(X)) \leq x \right) \geq 1 - \beta \right\}, \tag{22}$$

where $U_\beta(X, S(X); \Gamma)$ is given by (9) and (13). By definition, $V_\beta(X, S(\cdot); \Gamma) - U_\beta(X, S(X); \hat{\Gamma}(X))$

is deterministic and we have

$$\begin{aligned}
& 1 - \beta \\
& \leq \mathbb{P}_{(\theta, X) \sim P_T} \left(\text{FDP}(\theta, S(X)) - U_\beta(X, S(X); \hat{\Gamma}(X)) \leq V_\beta(X, S(\cdot); \Gamma) - U_\beta(X, S(X); \hat{\Gamma}(X)) \right) \\
& = \mathbb{P}_{(\theta, X) \sim P_T} (\text{FDP}(\theta, S(X)) \leq V_\beta(X, S(\cdot); \Gamma)).
\end{aligned}$$

This gives that $V_\beta(X, S(\cdot); \Gamma)$ is a valid post hoc bound, which can be approximated by the bootstrap bound

$$V_\beta(X, S(\cdot); \hat{\Gamma}) = U_\beta(X, S(X); \hat{\Gamma}(X)) + q_{3,\beta}(\beta, S(\cdot); \hat{\Gamma}).$$

Given (22), this bound also relies on drawing the couple (θ^*, X^*) from the distribution $P_{\hat{\Gamma}}$. Now, $q_{3,\beta}(\beta, S(\cdot); \hat{\Gamma})$ is approximated by $\tilde{q}_{3,\beta}^{(B)}(\beta, S(\cdot); \hat{\Gamma})$ via a Monte Carlo scheme that generates B times the couple (θ^*, X^*) according to the model with parameter $\hat{\Gamma}$. This leads to the final bound:

$$U_\beta^{\text{boot3}}(X, S(\cdot)) = U_\beta(X, S(X); \hat{\Gamma}) + \tilde{q}_{3,\beta}^{(B)}(\beta, S(\cdot); \hat{\Gamma}) \quad (23)$$

Algorithm 3 gives full details about the computation of this bound.

The difference between $U_\beta^{\text{boot3}}(X, S(\cdot))$ and $U_\beta^{\text{naive}}(S(\cdot))$ is that our plug-in bound $U_\beta(X, S(X); \hat{\Gamma})$ is used to recenter the FDP. As a result, it “stabilizes” the bound $U_\beta^{\text{naive}}(S(\cdot))$: while $U_\beta^{\text{naive}}(S(\cdot))$ depends on X only via $\hat{\Gamma}$, the bound $U_\beta^{\text{boot3}}(X, S(\cdot))$ also depends on the plug-in bound $U_\beta(X, S(X); \hat{\Gamma}(X))$. Note that knowledge of the full selection policy $S(\cdot)$ is required to compute both bounds $U_\beta^{\text{naive}}(S(\cdot))$ and $U_\beta^{\text{boot3}}(X, S(\cdot))$.

4 Numerical experiments

This section summarizes numerical experiments performed in order to assess the quality of the bounds introduced in Section 3.

Post hoc bounds. These bounds will be compared to the Simes post hoc bound, a baseline competitor which does not take the latent structure of the model into account.

The Simes post hoc bound has been introduced by Goeman and Solari (2011); it provides an upper bound on the false positives, but no corresponding lower bound. It is valid as soon as the Simes inequality holds, which is the case here because the p -values are independent conditionally on θ . Following Blanchard et al. (2020), the Simes post hoc bound can be formulated as:

$$U^{\text{Simes}}(S(X)) = |S(X)|^{-1} \min_{1 \leq k \leq m} \left\{ \sum_{i \in S(X)} \mathbf{1}\{p_i(X) > \beta k/m\} + k - 1 \right\}. \quad (24)$$

Here, since the null distribution P_0 is assumed to be known, we consider the p -values $p_i(X) = \overline{F}_0(X_i)$, $1 \leq i \leq m$, where $\overline{F}_0(x) = \mathbb{P}(Z \geq x)$ for any $Z \sim P_0$. The bounds considered are summarized in Table 1. Note that for all but the Simes bound, we also have a corresponding lower bound.

Name	Definition	Reference
Oracle	$U_\beta(S(X), \Gamma)$	(13)
Plug-in	$U_\beta(S(X), \widehat{\Gamma})$	(16)
Boot1	$U_{\beta, \delta}^{\text{boot1}}(S(\cdot), \widehat{\Gamma}), \delta \in \{0.1; 0.5; 0.9\}$	Section 3.3
Boot2	$U_{\beta, \delta}^{\text{boot2}}(S(X), \widehat{\Gamma}), \delta = 0.5$	Section 3.4
Naive	$U_\beta^{\text{naive}}(S(\cdot))$	Section 3.5
Boot3	$U_\beta^{\text{boot3}}(S(X), \widehat{\Gamma})$	Section 3.5
Simes	$U^{\text{Simes}}(S(X))$	(24)

Table 1: (Upper) bounds considered in the numerical experiments

To ensure that the bootstrap bounds are never less conservative than the Plug-in, these bounds have been slightly modified and \tilde{q}_β has been replaced by its positive part \tilde{q}_β^+ . This modification generally has no effect on the upper bounds in practice, but does affect the corresponding lower bounds, for which we have observed a tendency of the bootstrap towards overcompensation.

Throughout this section, the target risk level is set to $\beta = 0.1$. Therefore, our proposed upper bounds are supposed to satisfy $\mathbb{P}(U_\beta(X, S(X); \Gamma) < \text{FDP}(\theta, S(X))) \leq 0.1$ and the

corresponding lower bounds are supposed to satisfy $\mathbb{P}(L_\beta(X, S(X); \Gamma) > \text{FDP}(\theta, S(X))) \leq 0.1$.

Selection policies. We consider three different selection policies $S(\cdot)$, labeled as follows in the figures:

- “ $S(X) = \{p_i < 0.05\}$ ”: the p -value level set associated to the threshold 0.05.
- “ $S(X) = SC(0.05)$ ”: items selected by the procedure introduced by Sun and Cai (Sun and Cai, 2009) in the HMM model (see Section 1.5) to control FDR at level 0.05 using the parameter estimators $\hat{\Gamma}$ of Section 2.2.
- “ $S(X) = \text{Viterbi}$ ”: items i selected by the Viterbi algorithm, that is, such that $\hat{\theta}_i = 1$, where $\hat{\theta}$ is the estimation of θ using the Viterbi algorithm with the parameter estimators $\hat{\Gamma}$ of Section 2.2.

Section 4.1 illustrates the behavior of the considered bounds in a setting where our assumptions are met. The robustness of the method is then studied in Section 4.2, where we report the results of numerical experiments in the case where the HMM is not or poorly identifiable, and where the selected set $S(X)$ depends on prior knowledge not included in X . Finally, in Section 4.3 we report the results of further numerical experiments on DNA copy number data from a cancer study with known truth.

4.1 Results in a typical within-model scenario

The Markov chain $(\theta_i)_{1 \leq i \leq m}$ is generated from transition matrix :

$$A = \begin{pmatrix} 0.95 & 0.05 \\ 0.2 & 0.8 \end{pmatrix}.$$

with the stationary distribution $(0.8, 0.2)$ as initial state.

The X variables are generated such that $X_i | \theta_i = 0 \sim \mathcal{N}(0, 1)$ and $X_i | \theta_i = 1 \sim P_1$. Here, $P_1 = \mathcal{N}(3, 1)$.

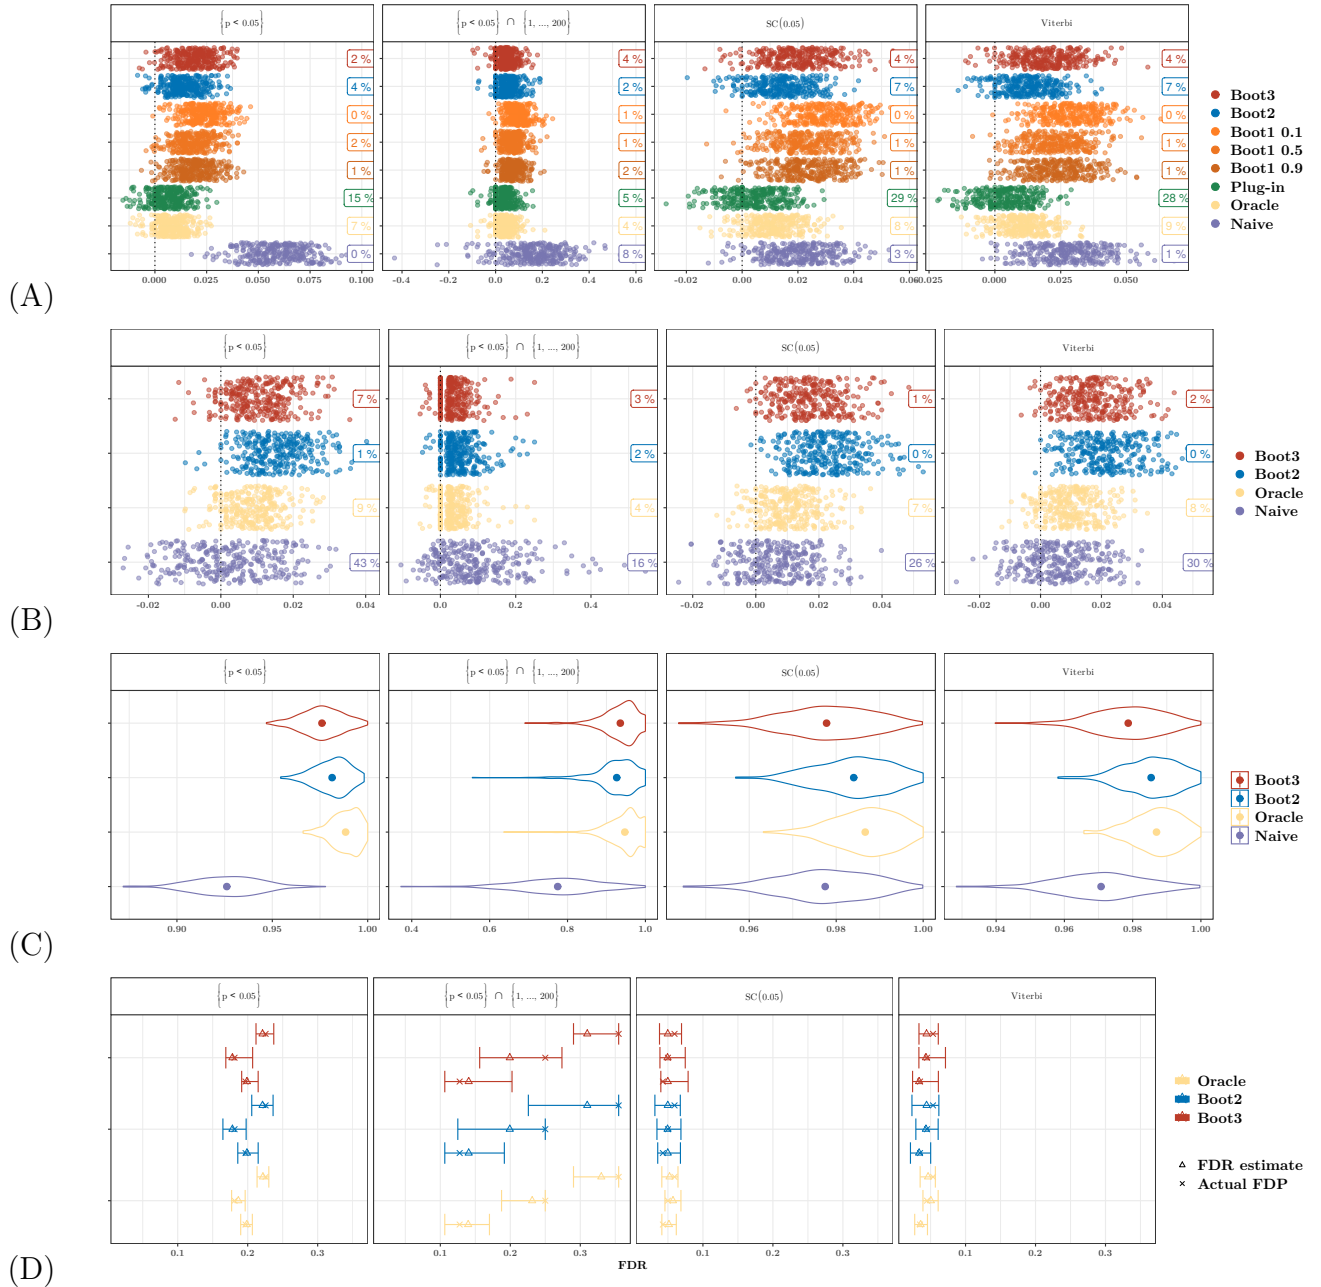


Figure 2: (A-B-C): Results for 300 simulation runs with the model parameters described in Section 4.1 (A): Difference $\Delta = U(\cdot) - \text{FDP}(\theta, S(X))$ for each upper bound (rows) and each selection policy $S(X)$ (columns), with the empirical violation probability of the bound displayed within rectangles. (B): Difference $\Delta' = \text{FDP}(\theta, S(X)) - L(\cdot)$ for the corresponding lower bounds. (C): Power of the different bounds. (D): Realizations of 80% FDP post hoc intervals for three illustrative simulation runs.

Figure 2 (A) displays the value of the difference Δ between the upper bound $U(\cdot)$ and the true FDP value $\text{FDP}(\theta, S(X))$ for each upper bound (in rows) and selection policy (in columns), for 300 simulation runs. The empirical violation probability of the bound, that is, the proportion of simulation runs for which a given bound is lower than the true FDP (i.e. $\Delta < 0$), is displayed within rectangles. This proportion is expected to be lower than $\beta = 10\%$. We chose not to display the classical Simes bound because, as expected, it is far more conservative than the proposed bounds that take into account the latent structure of the model.

As stated earlier, the Plug-in bound is an estimator of the oracle bound and not an upper bound of it, hence it is not surprising that its empirical violation probability exceeds the target. All of the other bounds have empirical violation probability below the target in this setting. This is expected for the bootstrap-corrected bounds Boot1, Boot2, and Boot3. The fact that the Naive bound is also below the target risk was not guaranteed, because (similarly to $U(\beta, S(X); \hat{\Gamma})$) this bound is only an estimator of $U(\beta, S(X); \Gamma)$. We observed that, the impact of the choice of δ in Boot1 seems moderate. Nevertheless Boot3 seems always less conservative than Boot1, which could be explained by the fact that Boot3 does not use such split of the confidence budget. Overall, Boot2 and Boot3 do not uniformly dominate each other, which is well expected because they come from two different bootstrap strategies. For the sake of readability of the illustrations, Boot1 and the Plug-in bound will not be displayed in the remainder of this section.

Figure 2 (B) displays, for the corresponding lower bound the value of the difference Δ' between the FDP and the lower bound $L(\cdot)$. Similarly to (A), valid lower bounds are expected to be above the true FDP (i.e. $\Delta' < 0$) in less than 10% of the simulations runs. This is the case for all the proposed bounds, except for the naive one. Figure 2 (C) displays the power of the different bounds, which is defined consistently with Blanchard et al. (2020):

$$\text{Power} = \mathbb{E} \left(\frac{|S(X)| - U(X)}{|S(X) \cap \mathcal{H}_1|} \mid |S(X) \cap \mathcal{H}_1| \neq 0, |S(X) \cap \mathcal{H}_0| \leq U(X) \right) \quad (25)$$

In this setting, the bounds are almost as powerful as the oracle.

Finally, in order to emphasize the interest of FDP post hoc interval compared to point-wise FDR estimate $\widehat{\text{FDR}}(S(X), \widehat{\Gamma})$, see (7), we have displayed both in Figure 2 (D) for 3 arbitrarily chosen simulation runs. The FDP post hoc intervals are clearly more informative than the corresponding point estimate since they quantify its uncertainty, which can be widely different according to the scenario: this is reflected in the interval lengths. In two of the three displayed simulation the estimated FDR is quite far from the true value, whereas the true value still lies in the post hoc interval.

4.2 Challenging our assumptions

This section briefly summarizes further numerical experiments carried out in order to test the robustness of the proposed bounds either to violations of the model assumptions, or to departures from a mild scenario towards more challenging settings. The corresponding illustrations are postponed to Appendix E.

Invalid selection policies. One of the assumptions of the method is that the selected set $S(X)$ cannot depend on an additional prior knowledge not included in the observation X . In terms of the modeling via the HMM, violating this assumption would correspond to a selection policy which would have access to “insider information” about the latent configuration vector θ under one form or another (e.g. an ancillary statistic providing some additional information). To assess how important this constraint is for the validity of the method, we have considered three selection policies that include full knowledge of \mathcal{H}_0 (this of course a very unrealistic situation, only considered here to illustrate the point). The results are displayed in Figure 6 (Appendix E) and show that the oracle bound (as well as the second bootstrap bound which tries to mimic the oracle) are too liberal in this case. The third bootstrap bound, which corrects the plug-in bound to try to match the FDP, respects the risk. However, it requires knowledge of the full selection policy $S(\cdot)$. Depending on the situation, this might be realistic (e.g. if the additional information stems from an ancillary statistic that can also be simulated) or not (e.g. if the additional information comes from

vaguely defined “insider expert knowledge”).

Identifiability issues. When the latent variables are independent, the model is not identifiable (Alexandrovich et al., 2016; Gassiat et al., 2016). When the latent variables are close to independence (that is, when $\det(A)$ gets close to 0), the model is close to singular. The bounds obtained for transition matrices A with various values of $\det(A)$ are displayed in Figure 7 (Appendix E). As expected, the bounds are too liberal in the independent case. However, they appear to be valid even for small determinants.

Unknown f_0 . Finally, we have also computed the bounds in a case where f_0 is unknown. We have considered two options for initializing the estimation of f_0 in the corresponding EM-type algorithm, see Algorithm 4: (i) using the true f_0 , and (ii) estimating f_0 using local FDR algorithm Efron (2004b). In both cases, the proposed bounds remain below the target risk, as illustrated in Figure 8 (Appendix E).

4.3 Semi-simulated data based on DNA copy numbers

In order to test the robustness of our methodology we have considered a more realistic scenario, which does not rely on a probabilistic simulation model. Using the R package `jointseg` (Pierre-Jean et al., 2019), we have generated synthetic copy number profiles as proposed (and further described) in Pierre-Jean et al. (2015): given a number m of loci and a number K of regions, we draw uniformly $K - 1$ breakpoint positions, thus defining K regions. Then, we draw K region labels, corresponding to the number of DNA copies for each parent (a.k.a. parental CN). Finally, for a region of size m_k , we draw m_k samples by resampling from a real DNA copy number data corresponding to this type of region. These data are available in the R package `acnr` (Pierre-Jean and Neuvial, 2017), which contain annotated CN profiles from several cancer data sets. Importantly, these data sets correspond to dilution series where tumor and normal cells are mixed in known proportion. Therefore, signal to noise ratio of the corresponding CN profile is implicitly controlled by the fraction of tumor cells.

We proceed as follows to obtain two groups of samples with known differential regions. First, we generate $n_1 = 50$ CN profiles as described in the preceding paragraph, with the same regions (here, 10 regions). Then we generate $n_2 = 50$ samples from the same regions, but modify the label of two regions in such a way that this group of samples only has 8 different regions instead of 10. In this setting, $\theta_i = 1$ if the position i belongs to one of the two modified regions and $\theta_i = 0$ otherwise. We then compute Wilcoxon statistics and scale them using their limit law. More precisely, to compare a group 1 of size n_1 to a group 2 of size n_2 at position i we use $T_i^{(s)}$

$$T_i^{(s)} = \left(T_i - \frac{n_1 n_2}{2} \right) / \sqrt{\frac{n_1 n_2 (n_1 + n_2 + 1)}{12}}, \quad (26)$$

where T_i is the classical Wilcoxon statistic, the sum of the rank of group 1. This process is illustrated by Figure 9 in Appendix E.4. In this particular example, the proportion of tumor cells has been set to 70%, corresponding to a moderate to high signal to noise ratio.

The results of these numerical experiments are summarized in Figure 3, where the difference Δ between the FDP upper bound and the true FDP is displayed for each of 100 simulation runs. The lines in this figure correspond to increasing values for the SNR (as governed by the fraction of tumor cells), while the columns correspond to the selection policies introduced at the beginning of Section 4.

All the bounds appear to be valid in this settings, since their empirical violation probability is less than the target level $\beta = 10\%$ (i.e. $\Delta < 0$ in less than 10% of the simulation runs). Overall, the bounds are quite conservative in all settings, with empirical violation probabilities often closer to 0 than 10%. This conservativeness is partly explained by the fact that for high SNR values, the problem is so easy that the upper bounds often match the true FDP exactly, as illustrated by the presence of a mode at 0. The Simes bound is the most conservative, with null empirical violation probability in all but one scenario. As SNR decreases, we observe for all bounds a shift of the empirical distribution of Δ toward positive values. The Simes bounds is highly conservative already at tumor fraction 0.7, and Δ even concentrates at 1 for three of the four selection policies at tumor fractions 0.3 and 0.5: this corresponds to a true FDP close to 0 and an upper bound close to 1, meaning that

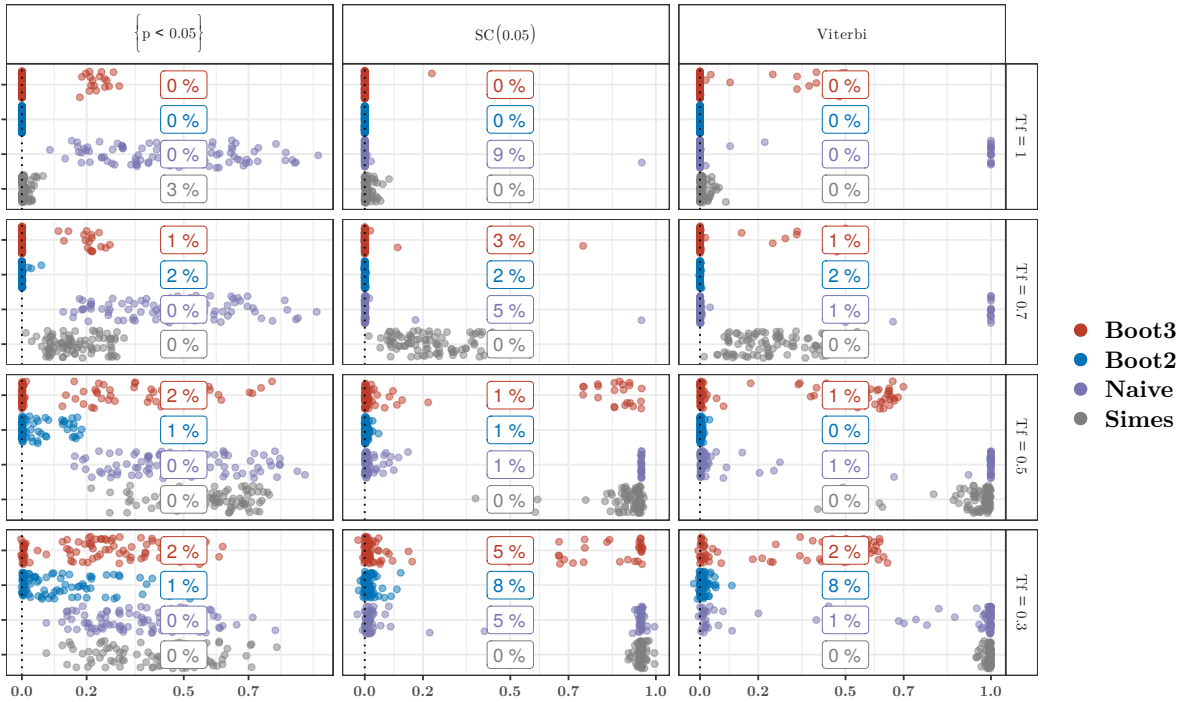


Figure 3: Summary of 100 numerical experiments with semi-simulated CN data. Each dot corresponds to a realization of the difference $\Delta = U(\cdot) - \text{FDP}(\theta, S(X))$ for each upper bound (colors) and each selection policy $S(X)$ (columns) for different values of the SNR (rows). The empirical violation probability of the bound is displayed within a rectangle.

the Simes bound is not informative at all in these settings. The first bootstrap bound has the same tendency, but to a lesser extent. The second and third bootstrap bounds show a remarkable behavior, with Δ remaining very close to 0 for the vast majority of simulation runs for all selection policies, even for small SNR values. In addition, Boot2 and Boot3 have also a high power (25), as shown in Figure 10 (Appendix E.5). These results strengthen our confidence in the applicability of our methods and that even in practical applications, our proposed bounds (and especially the second and third bootstrap bounds) will be able to evaluate the amount of true signal after selection with accuracy and correct coverage.

5 Applications

5.1 Influenza-like illness (ILI)

We apply the proposed method to the weekly incidence rates of influenza-like illness (ILI). This data were collected from the Sentinelles Network, a national surveillance system in France. Sun and Cai (2009) studied this data set between January 1985 and February 2008, to be comparable we will restrict ourselves to this period. They stated that the incidence rates can be classified into one of the two categories: aberration or usual.

The usual is the null states and the aberration is the alternatives one. The weekly ILI incidence rates are standardized according to the sizes of the underlying population and the representativeness of the participating physicians. Sun and Cai (2009) also applied a log transformation to reduce the skewness of the original data, hence so will we. In this particular example, the law under the null hypothesis (P_0) is unknown. Therefore, we estimate \hat{f}_0 using Algorithm 4, with $\mathcal{N}(2.37, 0.76^2)$, the estimation of Sun and Cai (2009), as initial value. Accordingly, the p -values are replaced by the empirical p -values \hat{p} defined in Appendix E.3.

To emphasize the advantages of bounds on the FDP compared to the FDR estimate, we displayed the 80% FDP post hoc intervals on the different sets $S(X)$ described in the numerical experiments. For instance, for the selection $SC(0.05)$, it is interesting to know

not only that the estimated FDR is 5% but also that the 90% upper bound is smaller than 0.08. This application also underlines the interest of developing lower bounds, for instance in the first set ($\{\hat{p} < 0.05\}$) we not only know, with high probability, that the FDP is smaller than 0.054 but also that it is higher than 0.028, which narrows down the probable true value of the FDP.

For completeness, we also added two selections policies :

- $SC(FDR_p)$ which uses Sun and Cai (2009) algorithm (with our parameter estimator) at a level corresponding to the estimated FDR of the set $\{\hat{p} < 0.05\}$. Doing so, the selections $SC(FDR_p)$ and $\{\hat{p} < 0.05\}$ have the same estimated FDR.
- $\{\hat{p} < th\}$ which selects the $|SC(0.05)|$ smallest \hat{p} -values. Doing so, $\{\hat{p} < th\}$ and $SC(0.05)$ select the same number of null hypotheses.

We added the selection policies $SC(FDR_p)$ to compare two selection sets that have the same FDR estimate but not the same size. The selection $\{\hat{p} < th\}$ has been added to compare two sets that have the same size but not the same estimated FDR. The result are displayed Figure 4. All the intervals are sharp around the FDP, which means that the FDP variance is estimated to be low. However, the post hoc intervals of $SC(FDR_p)$ is wider than the one of $\{\hat{p} < 0.05\}$. This information is not provided by FDR estimates. The set $\{\hat{p} < th\}$ has the same size that $|SC(0.05)|$ but a larger FDR. It emphasizes that Sun and Cai (2009) is a better selection policy. Overall, this reinforces the interest in using a post hoc interval rather than a pointwise FDR estimate.

5.2 Copy Number Alterations Associated with Endometriosis in Ovarian Clear Cell Adenocarcinoma

In this section our proposed bounds are applied to a study of 117 ovarian cancer patients briefly mentioned in the introduction (Okamoto et al., 2015). We focus on the part of the study that aims at comparing, for 13,239 loci located on chromosome 7, the DNA copy numbers of 54 patients with endometriosis (a common gynecological disorder characterized

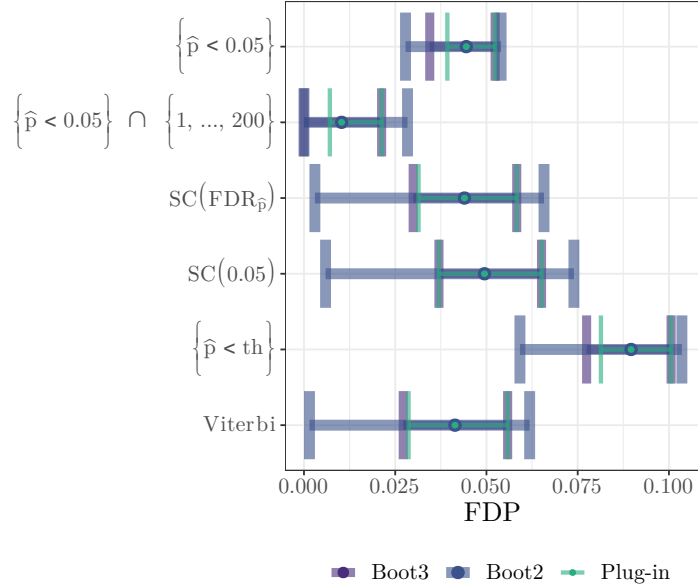


Figure 4: 80 % FDP post hoc interval intervals for different selection policies on the ILI incidence rates

by ectopic growth of endometrial glands and stroma) to that of 63 patients without endometriosis. As in the semi-simulated setting described in Section 4.3 we compare the two groups using Wilcoxon tests and scale the statistics using Equation (26). The obtained statistics for the different positions of the genome are displayed in the left part of Figure 5 (A) for the full chromosome 7 and (B) for the “short arm” of chromosome 7 (corresponding to the first 4,799 loci).

A first important observation is that the test statistics are not centered at 0 but globally shifted toward negative values across the entire chromosome 7, and more prominently so in chromosome 7p. This indicates that patients with endometriosis have on average a larger CN than the others, which may be due to an increase prevalence of trisomy 7 in these patients, or to a larger proportion of tumor cells in the biological samples corresponding to these patients. Here, our goal is not to detect such macroscopic changes but to pinpoint chromosomal regions that deviate from the rest of the chromosome. Indeed, such regions could indicate the presence of “driver” genes more actively or more early implicated in the

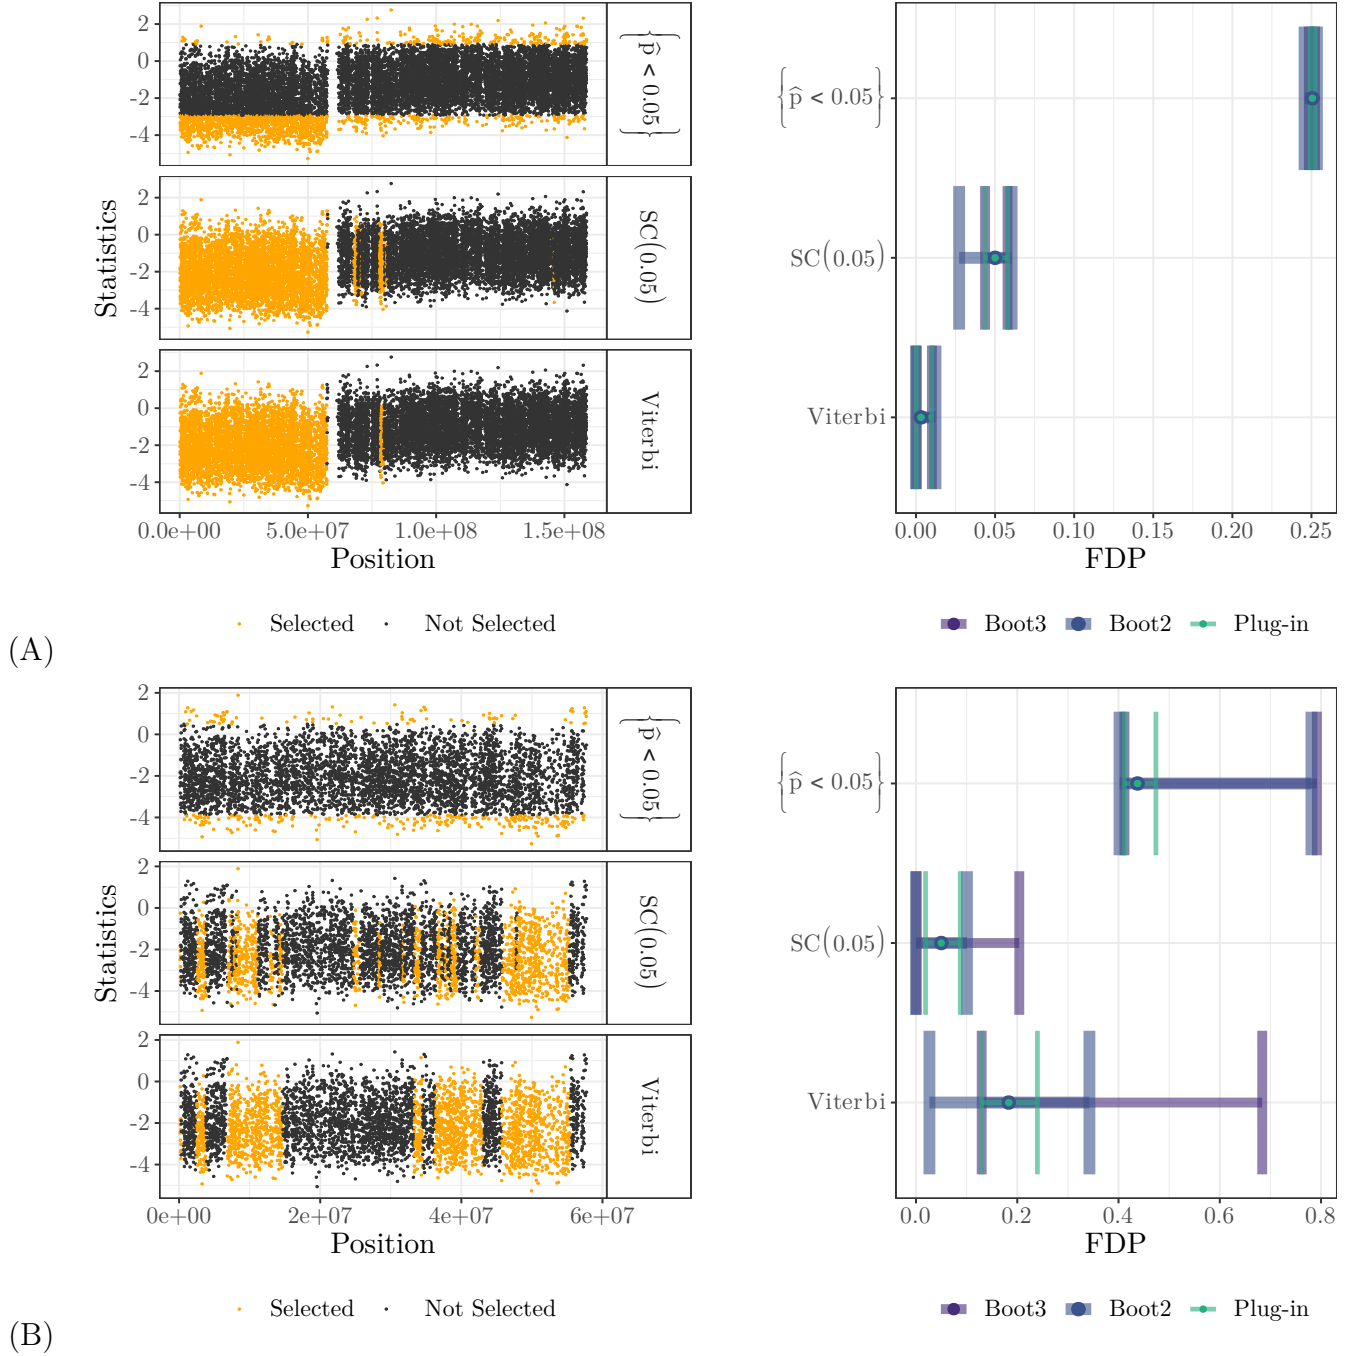


Figure 5: Left : Wilcoxon test statistic $T_i^{(s)}$ (see Equation 26) comparing the copy numbers between patients with and without endometriosis at different loci on chromosome 7 for (A) and 7p for (B). Right : 80% FDP post hoc interval different selection policies. Since f_0 is unknown, empirical p -values defined in Appendix E.3 are used.

tumor process compared to other “passenger” regions.

Therefore, we chose to estimate f_0 from the data, as described in Appendix D. Note that \hat{f}_0 will not estimate the distribution of the null hypothesis “the two groups have the same number of DNA copy at this position” but the distribution of the most frequent type of difference between the two groups. For instance if all patients without endometriosis have two copies of chromosome 7 (not affected by their ovarian cancer) and patients with endometriosis have a third copy of chromosome 7 and in some rare position they have a fourth copy, the estimated f_0 will be the law of the statistics comparing copy numbers two and three. Because of the marked shift between the short (left) and long (right) arms of chromosome 7, the estimation of f_0 is quite different depending on whether the entire chromosome or only one arm is analyzed. Therefore, we have carried out an analysis of the entire chromosome (Figure 5 (A)) and an analysis of its short arm, chromosome 7p (Figure 5 (B)):

Full chromosome 7: By construction, the selection policies that take into account the position mostly select regions on chromosome 7p. The post hoc intervals corresponding to each selection policy are represented in the right panel of Figure 5 (A). In this scenario the post hoc intervals are very tight around the pointwise FDR estimate, which reflects the fact that \hat{f}_0 and \hat{f}_1 are quite distinct here.

Chromosome 7p: When focusing on chromosome 7p, \hat{f}_0 and \hat{f}_1 are much closer to each other, resulting in wider FDP post hoc intervals than for the full chromosome. For instance, in the sets SC(0.05) the FDR is controlled at 5% but the upper bounds goes up to 20%. Even the plug-in bound, which has been shown in the simulation not to be conservative enough, goes up to 10%.

These examples emphasize the added value of post hoc intervals on the FDP over FDR point estimates, since a given value of the point estimate will be interpreted very differently depending on the width of the interval.

6 Discussion

6.1 Detecting close to singular scenarios

The HMM model becomes singular in the situation where $f_0 = f_1$, or when the transition matrix A is of rank 1 (i.e. the coordinates of the configuration vector θ are drawn i.i.d. from a Bernoulli distribution), in which case the model is not identifiable: the same data distribution for the observable X can be obtained for several sets of parameters, in particular different joint distributions of (θ, X) so that the notion of ground truth for the FDP is questionable.

The behavior of our approach in a situation close to singular was discussed in Section 4.2; in a truly singular situation, almost surely the estimator \hat{A} is not of rank 1, so that the estimation error (estimated at $\hat{\Gamma}$) will surely be underestimated by the bootstrap procedure.

It would be therefore in principle necessary to have a test for the singular case and stop the procedure if that test is not rejected (using an agnostic multiple testing procedure would then be more appropriate; see also the discussion on null estimation in the independent case, Section 1.6). While we do not cover the precise design of such a test in the present work, a suitable heuristic procedure in practice can be to monitor if $|\det \hat{A}|$ is too close to 0, or \hat{f}_1 is too close to f_0 .

6.2 Asymptotic consistency of plug-in

An important theoretical insight for FDR control methods using empirical Bayes-type approaches is that asymptotic consistency is usually granted for plug-in methods (Sun and Cai, 2009). In the case of nonparametric HMMs, consistent estimation with quantitative guarantees for the HMM parameters and the local FDR have been obtained recently, see discussion in Section 1.6. In the setting considered here, since we are considering in principle arbitrary policies and therefore selection sets $S(X)$, it seems we would need a stronger convergence property for parameter estimation, e.g. in the sense of total variation distance for the full joint distribution of (θ, X) . If such a convergence holds, we can guarantee the

asymptotic of the plug-in approach for FDP bounds (see Appendix F.)

However, obtaining such a consistency result for convergence of the full estimated joint distribution seems a tall order, since the asymptotics would be in the increasing size of the observation space, which will also make TV distance larger. (It might be more plausible for theory to assume that the parameter is estimated using an observation of larger size than the one it is used to perform multiple testing on.) Furthermore, our experiments show that the simple plug-in approach is generally unsatisfactory in practice and that the proposed bootstrap-based bounds are more appropriate.

6.3 Bayesian post hoc versus frequentist post hoc

The posterior (oracle) bound (3) can accommodate absolutely any selection policy $S(\cdot)$. Thus it seems that the bound holds in fact for any selected set $R = S(X)$. This can appear surprising at first, since in a frequentist setting, uniform agnostic post-hoc bounds (holding for all subsets R and over all latent parameter configurations θ) generally have a price for complexity (e.g. a union bound over a reference family of candidate rejection sets), while the Bayesian approach seems to offer a free lunch in that regard.

A key to understand this apparent conundrum is to insist that

1. strictly speaking, the guarantees only hold if the assumed structural model (prior) on the latent variable θ is correct. In Bayesian terminology, if \mathbb{P}_T is considered a prior, then the proposed (oracle) bounds are credible sets on the posterior FDP.
2. the posterior bound (3) is *not* a uniform statement over all possible rejection sets, but should be interpreted as conditioning with respect to X of an arbitrary but given selection policy $S(\cdot)$ that must *only depend on the observation X* .

This excludes to use any form of “insider” or “leaked” information on the latent variable configuration to be used for to determine the selected set, such that:

- Ancillary statistics conveying additional information;

- “Expert knowledge” that would not be incorporated in the prior, e.g. an expert communicating that a certain region has a higher chance of containing false hypotheses, or that false hypothesis configurations should have a certain shape, while the prior does not distinguish a priori such configurations. This would be akin to using a prior that we know from the expert information is wrong (we can also interpret the expert knowledge as “insider information” on the configuration under the original prior).

As discussed in Section 4.2, violation of these model assumptions can result in the oracle bound to be incorrect. From a Bayesian point of view, the prior structural distribution assumption on θ should reflect the entirety of the available prior information, and not be used merely as a convenient default.

Connected to this, it would be interesting to study whether “frequentist Bayes” approaches based on posterior concentration to the true latent parameter (see, e.g., Castillo and Roquain, 2020, for a recent review) would be applicable here. While recent progress has been achieved to analyze the FDR from this perspective (Castillo and Roquain, 2020; Abraham et al., 2021a,b), the same question relative to post hoc bounds has not been explored yet up to our knowledge.

Acknowledgements

This work has been supported by ANR-16-CE40-0019 (SansSouci), ANR-17-CE40-0001 (BASICS), ANR-19-CHIA-0021-01 (BiSCottE), the UPSaclay Excellency Chair REC-2019-044, and by the GDR ISIS through the "projets exploratoires" program (project TASTY).

Appendix

A Posterior distribution

We provide here the details for the quantities involved in the posterior distribution, see Proposition 2.1. We recall that $\pi = (\pi_0, \pi_1)$ denotes the marginal stationary distribution

of the underlying Markov Chain. Let for $i \in \mathbb{N}_m$ and $q \in \{0, 1\}$,

$$\begin{aligned}\alpha_i(q) &= \mathbb{P}_\Gamma(X_{1:i}, \theta_i = q); \\ \beta_i(q) &= \mathbb{P}_\Gamma(X_{(i+1):m} \mid \theta_i = q); \\ \boldsymbol{\ell}_{i,q}(\Gamma) &= \mathbb{P}_\Gamma(\theta_i = q \mid X_{1:m}) = \frac{\alpha_i(q)\beta_i(q)}{\sum_{k \in \{0,1\}} \alpha_i(k)\beta_i(k)},\end{aligned}\tag{27}$$

with the common notation for which $\mathbb{P}_\Gamma(X_{1:i}, \theta_i = q)$ denotes the density of $(X_{1:i}, \theta_i)$ taken at point $(X_{1:i}, q)$ and $\mathbb{P}_\Gamma(X_{(i+1):m} \mid \theta_i = q)$ denotes the density of $X_{(i+1):m}$ conditionally on $\theta_i = q$ taken at point $X_{(i+1):m}$. These quantities can be obtained through the standard forward-backward algorithm, that is,

$$\alpha_1(q) = \pi_q f_q(X_1), \quad \alpha_{i+1}(q) = f_q(X_{i+1}) \sum_{\ell \in \{0,1\}} a_{\ell,q} \alpha_i(\ell), \quad 1 \leq i \leq m-1 \tag{28}$$

$$\beta_m(q) = 1, \quad \beta_{i-1}(q) = \sum_{\ell \in \{0,1\}} f_\ell(X_i) a_{q,\ell} \beta_i(\ell), \quad 2 \leq i \leq m. \tag{29}$$

Further, we let for $i \in \{2, \dots, m\}$ and $q, q' \in \{0, 1\}$,

$$\boldsymbol{\ell}_{i,q,q'}(\Gamma) = \mathbb{P}_\Gamma(\theta_i = q', \theta_{i-1} = q \mid X_{1:m}) = \frac{\beta_i(q') \alpha_{i-1}(q) f_{q'}(X_i) a_{q,q'}}{\sum_{k \in \{0,1\}} \beta_i(k) \alpha_i(k)}; \tag{30}$$

$$\Pi_{i,q,q'}(\Gamma) = \mathbb{P}_\Gamma(\theta_i = q' \mid \theta_{i-1} = q, X_{1:m}) = \frac{\boldsymbol{\ell}_{i,q,q'}(\Gamma)}{\boldsymbol{\ell}_{i-1,q}(\Gamma)} = \frac{\beta_i(q') f_{q'}(X_i) a_{q,q'}}{\beta_{i-1}(q)}; \tag{31}$$

$$\Pi_i(\Gamma) = \begin{pmatrix} \Pi_{i,0,0}(\Gamma) & \Pi_{i,0,1}(\Gamma) \\ \Pi_{i,1,0}(\Gamma) & \Pi_{i,1,1}(\Gamma) \end{pmatrix}. \tag{32}$$

B Details of the implementation

B.1 Details on Algorithm 1

The initialization $\widehat{\Gamma}^{(0)}$ of $\widehat{\Gamma}$ in Algorithm 1 is done as follows: $\widehat{A}^{(0)}$ is computed first by estimating the stationary distribution, that is, $(\pi_0, 1 - \pi_0)$, by using a standard Storey estimator $\widehat{\pi}_0^{(0)}$ ($\lambda = 0.8$) Storey (2002). Then, the parameter $\widehat{A}_{1,1}^{(0)}$ is taken uniformly over the possible range for this parameter, intersected with $[0.6, 1]$ to ensure that the null class

is predominant. On the other hand, $\hat{f}_1^{(0)} = (\hat{f} - \hat{\pi}_0^{(0)} f_0)/(1 - \hat{\pi}_0^{(0)})$, where \hat{f} is a Kernel-based estimation of the density of the measurements. The algorithm is stopped if the distance between to successive iteration of $\hat{\Gamma}$ (measured in infinite norm, restricted on the observations for the densities) is below 10^{-4} .

B.2 Details on Algorithm 2

In practice for a large s Algorithm 2 can be time consuming because it requires the estimation of $2s \times s$ matrices. To speed up the algorithm we notice that it is not necessary to compute the matrices $B_{k,\ell,0}$ and $B_{k,\ell,1}$ for all $\ell \in \{1, \dots, s\}$ to get $U_\beta(X, S(X), \Gamma)$. Indeed, we can stop for u such that $B_{k,u,1} + B_{k,u,0} \geq 1 - \beta$: as shown in Proposition 3.1 $U_\beta(X, S(X), \Gamma) = u$.

B.3 Details on Algorithm 3

Drawing $\theta^{*(b)}$ and $X^{*(b)}$ from $\hat{\Gamma}$

1. Draw a Markov chain $\theta^{*(b)}$ of size m with transition matrix \hat{A} , and initial distribution the stationary distribution, namely $(\frac{\hat{a}_{1,0}}{\hat{a}_{0,1} + \hat{a}_{1,1}}, 1 - \frac{\hat{a}_{1,0}}{\hat{a}_{0,1} + \hat{a}_{1,1}})$.
2. For $1 \leq i \leq m$ draw independently $X_i^{*(b)}$ as (case f_0 known):

$$X_i^{*(b)} \sim \begin{cases} f_0 & \text{if } \theta_i^{*(b)} = 0; \\ \hat{f}_1 & \text{if } \theta_i^{*(b)} = 1; \end{cases}$$

Drawing from \hat{f}_1 can be done easily by noting that \hat{f}_1 is the density of a mixture of Gaussian distributions:

$$\hat{f}_1 = \sum_{i=1}^m w_i \mathcal{N}(X_i, h^2), \quad w_i = \frac{\ell_{i,1}(X; \hat{\Gamma})}{\sum_{j=1}^m \ell_{j,1}(X; \hat{\Gamma})}, \quad 1 \leq i \leq m.$$

If f_0 is unknown, the only difference is that $X_i^{*(b)} \sim \hat{f}_0$ when $\theta_i^{*(b)} = 0$. Drawing from \hat{f}_0 is made similarly to the case of \hat{f}_1 .

Empirical quantiles To obtain the empirical quantiles $\tilde{q}_{i,\gamma}$ of $D_i = (D_i^{(1)}, \dots, D_i^{(B)})$ we start by ordering D_i . More precisely let b_1, \dots, b_B such that $D_i^{(b_1)} \leq D_i^{(b_2)} \leq \dots \leq D_i^{(b_B)}$ then we define $\tilde{q}_{i,\gamma} = D_i^{(b_j)}$, where j is the smallest integer larger or equal to γB .

C Computation of the lower bounds

Expressions for lower bounds can be obtained similarly to those of upper bounds. We provide them in this section for completeness.

Oracle and plug-in lower bounds. According to (6), we have (using the notation of Section 3.1, e.g., $R = S(X)$, $s = |S(X)|$):

$$L_\beta(X, R; \Gamma) = s^{-1} \max \left\{ n \in \{0, \dots, m\} : \mathbb{P}_\Gamma \left(\sum_{t=1}^s (1 - \theta_t^R) \geq n \mid X \right) \geq 1 - \beta \right\}.$$

Similarly to Proposition 3.1 the quantity $L_\beta(X, S(X); \Gamma)$ can be computed as

$$L_\beta(X, R; \Gamma) = s^{-1} \max \{ n \in \{0, \dots, m\} : B_{s,n-1,0} + B_{s,n-1,1} \leq \beta \}.$$

Adding the convention $B_{s,-1,0} = B_{s,-1,1} = 0$. The plug-in lower bound is given by $L_\beta(X, S(X); \hat{\Gamma})$.

First bootstrap lower bound. Proceeding as for the upper bound (see Section 3.3), we obtain

$$\begin{aligned} \mathbb{P}_\Gamma(\text{FDP}(\theta, S(X)) \geq L_{\beta(1-\delta)}(X, S(X); \Gamma) \mid X) &\geq 1 - \beta(1 - \delta) \\ \mathbb{P}_\Gamma(L_{\beta(1-\delta)}(X, S(X); \Gamma) - L_{\beta(1-\delta)}(X, S(X); \hat{\Gamma}) \geq q_{1,\beta\delta}^\ell(\beta(1 - \delta), S(\cdot); \Gamma)) &\geq 1 - \beta\delta. \end{aligned}$$

by letting

$$q_{1,\gamma}^\ell(\beta, S(\cdot); \Gamma) = \max \{ x \in \mathbb{R} : \mathbb{P}_\Gamma(L_{\beta(1-\delta)}(X, S(X); \Gamma) - L_{\beta(1-\delta)}(X, S(X); \hat{\Gamma}) \geq x) \geq 1 - \gamma \}.$$

This gives

$$\mathbb{P}_\Gamma(\text{FDP}(\theta, S(X)) \geq L_{\beta(1-\delta)}(X, S(X); \hat{\Gamma}) + q_{1,\beta\delta}^\ell(\beta(1 - \delta), S(\cdot); \Gamma)) \geq 1 - \beta. \quad (33)$$

and leads to the boot1 lower bound:

$$L_{\beta,\delta}^{\text{boot1}}(X, S(\cdot)) = L_{\beta(1-\delta)}(X, S(X); \widehat{\Gamma}) + \tilde{q}_{1,\beta\delta}^{\ell,(B)}(\beta(1-\delta), S(\cdot); \widehat{\Gamma}),$$

where $\tilde{q}_{1,\beta\delta}^{\ell,(B)}(\beta(1-\delta), S(\cdot); \widehat{\Gamma})$ is the Monte-Carlo approximation of the bootstrap quantile $q_{1,\beta\delta}^{\ell}(\beta(1-\delta), S(\cdot); \widehat{\Gamma})$, itself being an approximation of $q_{1,\beta\delta}^{\ell}(\beta(1-\delta), S(\cdot); \Gamma)$. In practice, $\tilde{q}_{1,\beta\delta}^{\ell,(B)}(\beta(1-\delta), S(\cdot); \widehat{\Gamma})$ can be derived as $E_1^{(b_i)}$ with $i = \lfloor B\beta\delta \rfloor + 1$ where $E_1^{(b_1)} \leq E_1^{(b_2)} \leq \dots \leq E_1^{(b_B)}$ with

$$E_1^{(b)} = L_{\beta(1-\delta)}(X^{*(b)}, S(X^{*(b)}), \widehat{\Gamma}) - L_{\beta(1-\delta)}(X^{*(b)}, S(X^{*(b)}), \widehat{\Gamma}^{*(b)}), 1 \leq b \leq B.$$

Second bootstrap lower bound. The same heuristic as for U^{boot2} (see Section 3.4) gives:

$$L_{\beta,\delta}^{\text{boot2}}(X, S(X)) = L_{\beta(1-\delta)}(X, S(X); \widehat{\Gamma}) + \tilde{q}_{2,\beta\delta}^{\ell,(B)}(\beta(1-\delta), S(X); \widehat{\Gamma}),$$

where $\tilde{q}_{2,\beta\delta}^{\ell,(B)}(\beta(1-\delta), S(X); \widehat{\Gamma}) = E_2^{(b_i)}$ with $i = \lfloor B\beta\delta \rfloor + 1$ where $E_2^{(b_1)} \leq E_2^{(b_2)} \leq \dots \leq E_2^{(b_B)}$ with

$$E_2^{(b)} = L_{\beta(1-\delta)}(X, S(X), \widehat{\Gamma}) - L_{\beta(1-\delta)}(X, S(X), \widehat{\Gamma}^{*(b)}), 1 \leq b \leq B.$$

Third lower bootstrap bound. Using similar arguments as for U^{boot3} (see Section 3.5), we obtain the bound

$$L_{\beta}^{\text{boot3}}(X, S(\cdot)) = L_{\beta}(X, S(X), \widehat{\Gamma}) + \tilde{q}_{3,\beta}^{\ell,(B)}(\beta, S(\cdot); \widehat{\Gamma}),$$

where $\tilde{q}_{3,\beta}^{\ell,(B)}(\beta, S(X); \widehat{\Gamma}) = E_3^{(b_i)}$ with $i = \lfloor B\beta \rfloor + 1$ where $E_3^{(b_1)} \leq E_3^{(b_2)} \leq \dots \leq E_3^{(b_B)}$ with

$$E_3^{(b)} = \text{FDP}(X^{*(b)}, S(X^{*(b)})) - L_{\beta}(X^{*(b)}, S(X^{*(b)}), \widehat{\Gamma}^{*(b)}), 1 \leq b \leq B.$$

D Estimation of f_0

When f_0 is unknown, we should estimate it along with A, f_1 in Algorithm 1. The following algorithm builds an estimator \widehat{f}_0 in the same spirit as \widehat{f}_1 (8), by replacing the $\ell_{i,1}$ by $\ell_{i,0}$, that is,

$$\widehat{f}_0^{(t)}(x) = \sum_{i=1}^m \ell_{i,0}(\widehat{\Gamma}^{(t-1)}) \frac{K((x - X_i)/h)}{h} / \sum_{i=1}^m \ell_{i,0}(\widehat{\Gamma}^{(t-1)}). \quad (34)$$

Algorithm 4: EM-type algorithm to derive $\widehat{\Gamma} = (\widehat{A}, \widehat{f}_0, \widehat{f}_1)$ (with f_0 unknown).

Input: $X = (X_1, \dots, X_m)$

Output: $\widehat{\Gamma} = (\widehat{A}, \widehat{f}_0, \widehat{f}_1)$ estimator of $\Gamma = (A, f_0, f_1)$

- Initialization: first guess $\widehat{\Gamma}^{(0)} = (\widehat{A}^{(0)}, \widehat{f}_0^{(0)}, \widehat{f}_1^{(0)})$ of Γ
- Loop: at step $t \geq 1$, given $\widehat{\Gamma}^{(t-1)}$, do
 - Compute
 1. $\alpha_i(q) = \alpha_i(q; \widehat{\Gamma}^{(t-1)})$ and $\beta_i(q) = \beta_i(q; \widehat{\Gamma}^{(t-1)})$, $1 \leq i \leq m$, $q \in \{0, 1\}$, using the forward-backward algorithm (28), (29);
 2. $\ell_{i,q}(X; \widehat{\Gamma}^{(t-1)})$, $1 \leq i \leq m$, $q \in \{0, 1\}$, by using (27);
 3. $\ell_{i,q,q'}(X; \widehat{\Gamma}^{(t-1)})$, $1 \leq i \leq m$, $q, q' \in \{0, 1\}$, by using (30).
 - Update
 1. $\widehat{a}_{q,q'}^{(t)} = \sum_{i=1}^{m-1} \ell_{i,q,q'}(\widehat{\Gamma}^{(t-1)}) / \sum_{i=1}^{m-1} \ell_{i,q}(\widehat{\Gamma}^{(t-1)})$, for $0 \leq q, q' \leq 1$;
 2. $\widehat{A}^{(t)} = \begin{pmatrix} \widehat{a}_{0,0}^{(t)} & \widehat{a}_{0,1}^{(t)} \\ \widehat{a}_{1,0}^{(t)} & \widehat{a}_{1,1}^{(t)} \end{pmatrix}$;
 3. $\widehat{f}_0^{(t)}$ using (34);
 4. $\widehat{f}_1^{(t)}$ using (8);
 5. $\widehat{\Gamma}^{(t)} = (\widehat{A}^{(t)}, \widehat{f}_0^{(t)}, \widehat{f}_1^{(t)})$.
 - Stop the loop if convergence, that is, $\widehat{\Gamma}^{(t)}$ close enough to $\widehat{\Gamma}^{(t-1)}$.
- Return $\widehat{\Gamma} = \widehat{\Gamma}^{(t)}$

Note that, at the end of Algorithm 4, the user can define the “null” state according to their preference. For instance:

- define the "null" state as the predominant one, that is, the most probable one according to the stationary distribution of \widehat{A} ;
- define the "null" state according to the density among $\{\widehat{f}_0, \widehat{f}_1\}$ whose mean is closer to 0.

When f_0 is estimated, we obtain as a by-product empirical p -values, denoted by \hat{p} : these are based on the estimated empirical cumulative function under \mathcal{H}_0 , namely $\hat{p}_i = 2(\min(1 - \hat{F}_0(t_i), \hat{F}_0(t_i)))$, where

$$\hat{F}_0(x) = \frac{\sum_{x_i < x} \ell_{i,0}(\hat{\Gamma})}{\sum_{i=1}^m \ell_{i,0}(\hat{\Gamma})}.$$

E Additional numerical experiments

In this section we present additional numerical experiments. We have modified the simulation described Section 4 by changing either the model parameters $\hat{\Gamma}$ or the selection policies $S(\cdot)$. In all these additional experiments, the number of hypotheses is $m = 3200$, the number of runs is 300 and the risk is $\beta = 10\%$.

E.1 Invalid selection policies

The selection policies used in Figure (6) produce invalid bounds (even the oracle bound) as they depend of some knowledge of \mathcal{H}_0 , the set of true null hypotheses.

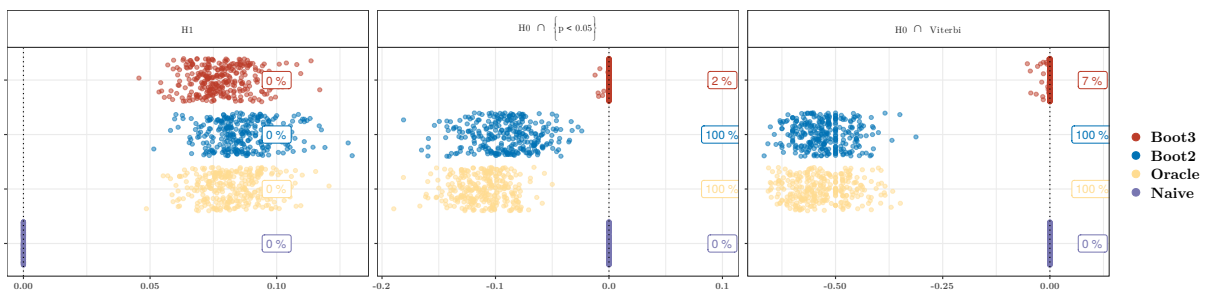


Figure 6: Plot similar to Figure 2 (A) (same model and simulation parameters) but when the selection policy is of the form $S(X, \mathcal{H}_0)$, that is, depends on some knowledge of \mathcal{H}_0 . Recall that the targeted level β is 10%.

E.2 Independent states or small determinant

In Figure 7, we set $a_{0,0} = 0.95$, and the value of $a_{1,1}$ is modified in order to achieve the desired value of $\det(A)$.

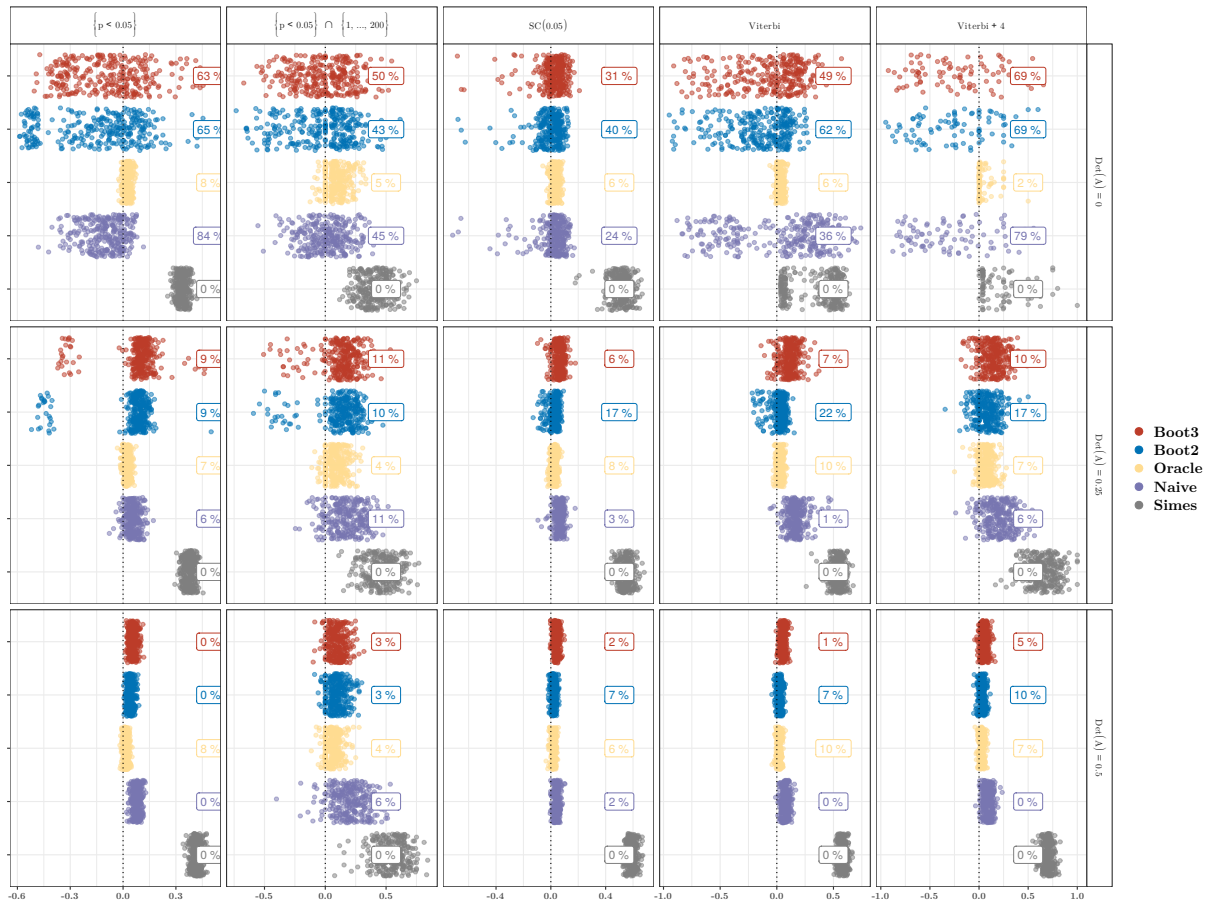


Figure 7: Plot similar to Figure 2 (A) (same densities f_0 and f_1 , and same simulation parameters), for different model parameters making $\det(A)$ small (rows). Recall that the targeted level β is 10%.

E.3 Unknown f_0

Figure 8 presents the results of numerical experiments in which f_0 is unknown and has to be estimated from the data as well.

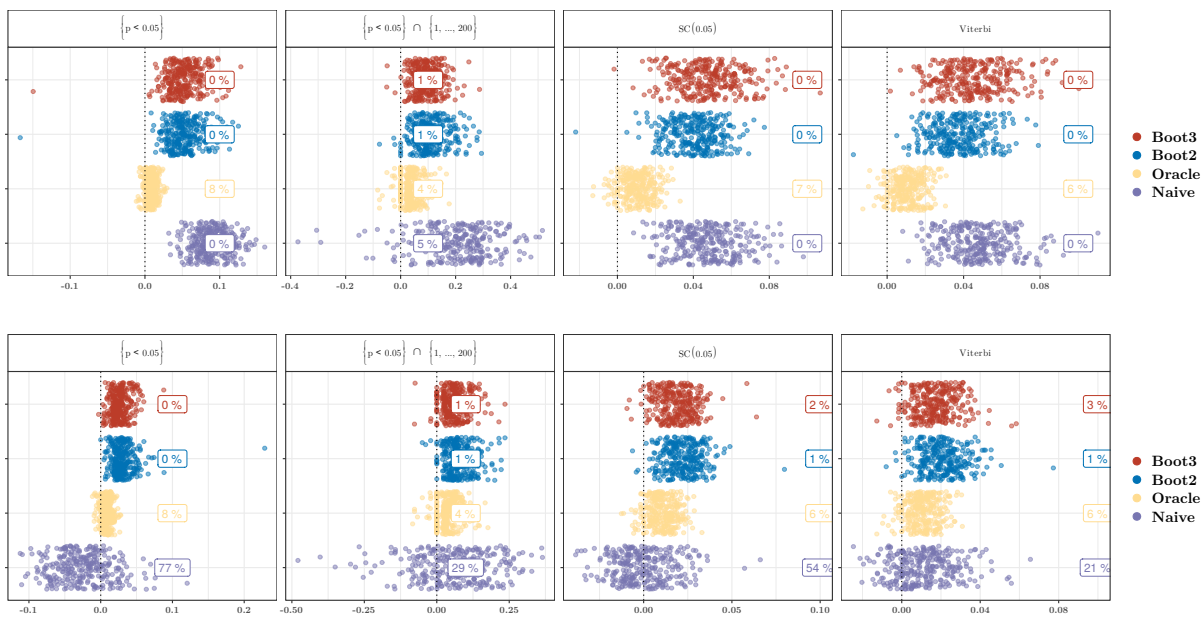


Figure 8: Plot similar to Figure 2 (A) (same model and simulation parameters) when f_0 is unknown, that is, when the bounds also use a f_0 estimator, see Algorithm 4. Top panel: initialization with the true f_0 ; Bottom panel: initialization using local FDR algorithm (Efron, 2004b)

E.4 Semi-simulated copy-number data

Figure 9 illustrates the semi-simulated data set analyzed in Section 4.3.

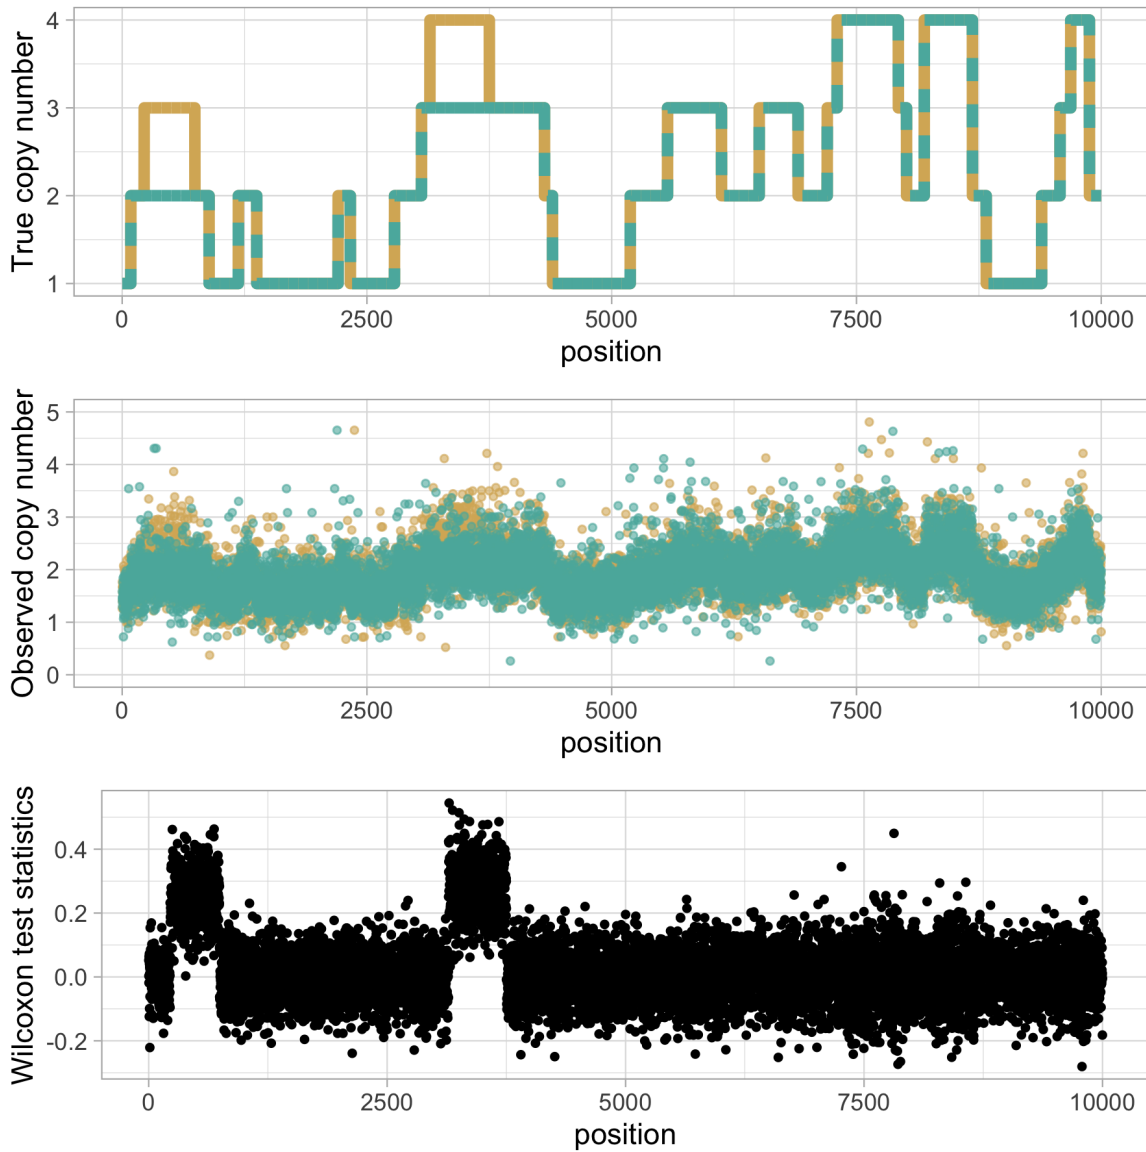


Figure 9: Illustration of the semi-simulated data set analyzed in Section 4.3. Top: true CN regions; middle: CN signal for one sample for group 1 and one sample for group 2; bottom: Wilcoxon test statistics for the comparison of $n_1 = 50$ samples from group 1 and $n_2 = 50$ samples from group 2. Here the proportion of tumor cells is set to 70%.

E.5 Power in the semi-simulated case

Figure 10 displays the power of the different bounds in the semi-simulated data set analyzed in Section 4.3.

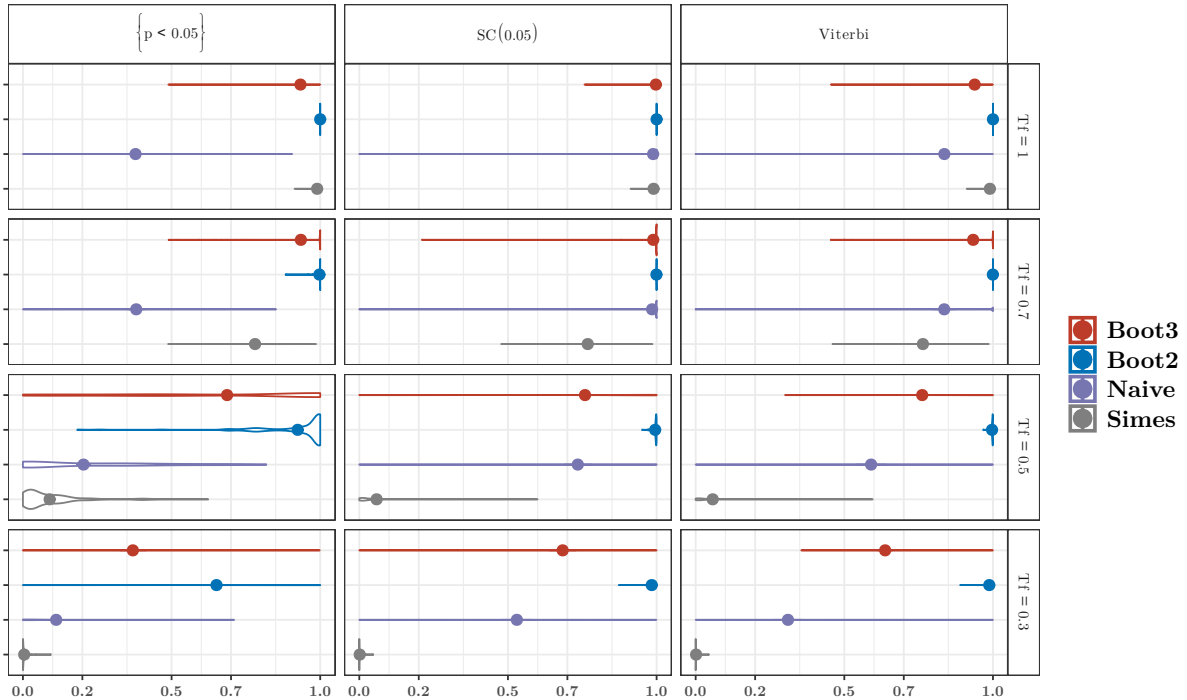


Figure 10: Summary of the power for the semi-simulated data set analyzed in Section 4.3.

F Towards plug-in consistency

The aim of this section is to provide sufficient conditions in order to ensure that the plug-in bound is asymptotically valid, as m tends to infinity. This supports the discussion made in Section 6.2.

Lemma F.1. *Assume that $\hat{\Gamma}$ is an estimator of Γ such that*

$$\forall \Gamma, d_{tv}(\mathcal{D}_{\hat{\Gamma}}(\theta | X), \mathcal{D}_{\Gamma}(\theta | X)) \text{ converges in probability to } 0 \text{ under } \Gamma, \quad (35)$$

where $\mathcal{D}_\Gamma(\theta \mid X)$ denotes the conditional distribution of θ conditionally on X under P_Γ and d_{tv} denotes the total variation distance. Then the plug-in bound $U_\beta^{PI}(X, S(X))$ (16) satisfies that

$$\forall \Gamma, \liminf_m \left\{ \mathbb{P}_\Gamma \left(\text{FDP}(\theta, S(X)) \leq U_\beta^{PI}(X, S(X)) \right) \right\} \geq 1 - \beta.$$

Proof. Since $U_\beta^{PI}(X, S(X)) = U_\beta(X, S(X); \hat{\Gamma})$, we have point-wise in X , for all Γ ,

$$\begin{aligned} & \mathbb{P}_\Gamma \left(\text{FDP}(\theta, S(X)) \leq U_\beta^{PI}(X, S(X)) \mid X \right) \\ &= \mathbb{P}_\Gamma \left(\text{FDP}(\theta, S(X)) \leq U_\beta(X, S(X); \hat{\Gamma}) \mid X \right) \\ &\geq \mathbb{P}_{\hat{\Gamma}} \left(\text{FDP}(\theta, S(X)) \leq U_\beta(X, S(X); \hat{\Gamma}) \mid X \right) \\ &\quad - \sup_{n \in [0, m]} \left\{ \mathbb{P}_\Gamma \left(\sum_{i \in S(X)} (1 - \theta_i) \leq n \mid X \right) - \mathbb{P}_{\hat{\Gamma}} \left(\sum_{i \in S(X)} (1 - \theta_i) \leq n \mid X \right) \right\} \\ &\geq 1 - \beta - d_{tv}(\mathcal{D}_{\hat{\Gamma}}(\theta \mid X), \mathcal{D}_\Gamma(\theta \mid X)), \end{aligned}$$

by definition (9) of the functional U_β . This entails the result by integrating over X . \square

References

- K. Abraham, I. Castillo, and E. Gassiat. Multiple testing in nonparametric hidden Markov models: An empirical Bayes approach. *arXiv:2101.03838*, 2021a.
- K. Abraham, I. Castillo, and E. Roquain. Empirical bayes cumulative ℓ -value multiple testing procedure for sparse sequences, *arXiv:2102.00929*, 2021b.
- D. G. Albertson, C. Collins, F. McCormick, and J. W. Gray. Chromosome aberrations in solid tumors. *Nature genetics*, 34(4):369–376, 2003.
- G. Alexandrovich, H. Holzmann, and A. Leister. Nonparametric identification and maximum likelihood estimation for hidden Markov models. *Biometrika*, 103(2):423–434, 2016.

- D. Azriel and A. Schwartzman. The empirical distribution of a large number of correlated normal variables. *Journal of the American Statistical Association*, 110(511):1217–1228, 2015.
- F. Bachoc, G. Blanchard, and P. Neuvial. On the post selection inference constant under restricted isometry properties. *Electron. J. Statist.*, 12(2):3736–3757, 2018.
- F. Bachoc, H. Leeb, and B. M. Pötscher. Valid confidence intervals for post-model-selection predictors. *The Annals of Statistics*, 47(3):1475 – 1504, 2019.
- Y. Benjamini and M. Bogomolov. Selective inference on multiple families of hypotheses. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(1):297–318, 2014.
- Y. Benjamini and Y. Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. Ser. B*, 57(1):289–300, 1995.
- Y. Benjamini and D. Yekutieli. False discovery rate-adjusted multiple confidence intervals for selected parameters. *J. Am. Stat. Assoc.*, 100(469):71–81, 2005.
- R. Berk, L. Brown, A. Buja, K. Zhang, and L. Zhao. Valid post-selection inference. *The Annals of Statistics*, 41(2):802 – 837, 2013.
- G. Blanchard, P. Neuvial, and E. Roquain. Post hoc confidence bounds on false positives using reference families. *Ann. Statist.*, 48(3):1281–1303, 06 2020.
- T. T. Cai and J. Jin. Optimal rates of convergence for estimating the null density and proportion of nonnull effects in large-scale multiple testing. *Ann. Statist.*, 38(1):100–145, 2010.
- T. T. Cai and W. Sun. Simultaneous testing of grouped hypotheses: finding needles in multiple haystacks. *J. Amer. Statist. Assoc.*, 104(488):1467–1481, 2009.

- T. T. Cai, W. Sun, and W. Wang. Covariate-assisted ranking and screening for large-scale two-sample inference. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 81(2):187–234, 2019.
- O. Cappé, E. Moulines, and T. Rydén. *Inference in hidden Markov models*. Springer Science & Business Media, 2006.
- I. Castillo and E. Roquain. On spike and slab empirical Bayes multiple testing. *Annals of Statistics*, 48(5):2548–2574, 2020.
- Y. De Castro, E. Gassiat, and S. Le Corff. Consistent estimation of the filtering and marginal smoothing distributions in nonparametric hidden Markov models. *IEEE Transactions on Information Theory*, 63(8):4758–4777, 2017.
- G. Durand, G. Blanchard, P. Neuvial, and E. Roquain. Post hoc false positive control for structured hypotheses. *Scandinavian Journal of Statistics*, 47(4):1114–1148, 2020.
- B. Efron. Bootstrap methods: Another look at the jackknife. *The Annals of Statistics*, 7(1):1 – 26, 1979.
- B. Efron. Large-scale simultaneous hypothesis testing: the choice of a null hypothesis. *J. Am. Stat. Assoc.*, 99(465):96–104, 2004a.
- B. Efron. Large-scale simultaneous hypothesis testing: the choice of a null hypothesis. *Journal of the American Statistical Association*, 99(465):96–104, 2004b.
- B. Efron. Doing thousands of hypothesis tests at the same time. *Metron - International Journal of Statistics*, LXV(1):3–21, 2007.
- B. Efron. Microarrays, empirical Bayes and the two-groups model. *Statist. Sci.*, 23(1):1–22, 2008.
- B. Efron. Empirical Bayes estimates for large-scale prediction problems. *J. Am. Stat. Assoc.*, 104(487):1015–1028, 2009.

- B. Efron, R. Tibshirani, J. D. Storey, and V. Tusher. Empirical Bayes analysis of a microarray experiment. *J. Amer. Statist. Assoc.*, 96(456):1151–1160, 2001.
- J. Fan and X. Han. Estimation of the false discovery proportion with unknown dependence. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(4):1143–1164, 2017.
- J. Fan, Y. Ke, Q. Sun, and W.-X. Zhou. Farmtest: Factor-adjusted robust multiple testing with approximate false discovery control. *Journal of the American Statistical Association*, pages 1–29, 2019.
- J. Fridlyand, A. M. Snijders, D. Pinkel, D. G. Albertson, and A. N. Jain. Hidden Markov models approach to the analysis of array CGH data. *Journal of multivariate analysis*, 90(1):132–153, 2004.
- C. Friguet, M. Kloareg, and D. Causeur. A factor model approach to multiple testing under dependence. *Journal of the American Statistical Association*, 104(488):1406–1415, 2009.
- M. Gales and S. Young. *The application of hidden Markov models in speech recognition*. Now Publishers Inc, 2008.
- É. Gassiat, A. Cleyne, and S. Robin. Inference in finite state space non parametric hidden Markov models and applications. *Statistics and Computing*, 26(1-2):61–71, 2016.
- C. R. Genovese and L. Wasserman. Exceedance control of the false discovery proportion. *Journal of the American Statistical Association*, 101(476):1408–1417, 2006.
- J. J. Goeman and A. Solari. Multiple testing for exploratory research. *Statistical Science*, 26(4):584 – 597, 2011.
- R. Heller and S. Rosset. Optimal control of false discovery criteria in the two-group model. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 83(1):133–155, 2021.

- R. Heller and D. Yekutieli. Replicability analysis for genome-wide association studies. *Ann. Appl. Stat.*, 8(1):481–498, 03 2014.
- J. Jin and T. T. Cai. Estimating the null and the proportional of nonnull effects in large-scale multiple comparisons. *J. Amer. Statist. Assoc.*, 102(478):495–506, 2007.
- E. Katsevich and A. Ramdas. Simultaneous high-probability bounds on the false discovery proportion in structured, regression, and online settings, arXiv:1803.06790, 2018.
- C.-J. Kim, C. R. Nelson, et al. *State-space models with regime switching: classical and Gibbs-sampling approaches with applications*, volume 1. The MIT press, 1999.
- T. Koski. *Hidden Markov models for bioinformatics*, volume 2. Springer Science & Business Media, 2001.
- J. D. Lee, D. L. Sun, Y. Sun, J. E. Taylor, et al. Exact post-selection inference, with application to the lasso. *The Annals of Statistics*, 44(3):907–927, 2016.
- J. T. Leek and J. D. Storey. A general framework for multiple testing dependence. *Proceedings of the National Academy of Sciences*, 105(48):18718–18723, 2008.
- V. H. Nguyen and C. Matias. Nonparametric estimation of the density of the alternative hypothesis in a multiple testing setup. application to local false discovery rate estimation. *ESAIM: PS*, 18:584–612, 2014.
- A. Okamoto, J. Sehouli, N. Yanaihara, Y. Hirata, I. Braicu, B.-G. Kim, S. Takakura, M. Saito, S. Yanagida, M. Takenaka, et al. Somatic copy number alterations associated with japanese or endometriosis in ovarian clear cell adenocarcinoma. *PLoS One*, 10(2): e0116977, 2015.
- M. Pierre-Jean and P. Neuvial. *acnr: Annotated Copy-Number Regions*, 2017. R package version 1.0.0.
- M. Pierre-Jean, G. Rigai, and P. Neuvial. Performance evaluation of DNA copy number segmentation methods. *Briefings in bioinformatics*, 16(4):600–615, 2015.

- M. Pierre-Jean, G. Rigaiil, and P. Neuvial. *jointseg: Joint segmentation of multivariate (copy number) signals*, 2019. R package version 1.0.2.
- T. Rebafka, E. Roquain, and F. Villers. Graph inference with clustering and false discovery rate control, arXiv:1907.10176, 2019.
- S. Robin, A. Bar-Hen, J.-J. Daudin, and L. Pierre. A semi-parametric approach for mixture models: Application to local false discovery rate estimation. *Computational statistics & data analysis*, 51(12):5483–5493, 2007.
- E. Roquain and N. Verzelen. False discovery rate control with unknown null distribution: is it possible to mimic the oracle?, arXiv:1912.03109, 2020.
- H. Scheffé. *The analysis of variance*. John Wiley & Sons, Inc., New York; Chapman & Hall, Ltd., London, 1959.
- A. Schwartzman. Comment: “Correlated z -values and the accuracy of large-scale statistical estimates” [mr2752597]. *J. Amer. Statist. Assoc.*, 105(491):1059–1063, 2010.
- S. P. Shah, K.-J. Cheung Jr, N. A. Johnson, G. Alain, R. D. Gascoyne, D. E. Horsman, R. T. Ng, and K. P. Murphy. Model-based clustering of array cgh data. *Bioinformatics*, 25(12):i30–i38, 2009.
- M. Stephens. False discovery rates: a new deal. *Biostatistics*, 18(2):275–294, 2017.
- J. D. Storey. A direct approach to false discovery rates. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(3):479–498, 2002.
- L. Sun and M. Stephens. Solving the empirical Bayes normal means problem with correlated noise, arXiv:1812.07488, 2018.
- W. Sun and T. T. Cai. Oracle and adaptive compound decision rules for false discovery rate control. *J. Amer. Statist. Assoc.*, 102(479):901–912, 2007.

- W. Sun and T. T. Cai. Large-scale multiple testing under dependence. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(2):393–424, 2009.
- Y. Sun, N. R. Zhang, and A. B. Owen. Multiple hypothesis testing adjusted for latent variables, with an application to the agemap gene expression data. *The Annals of Applied Statistics*, 6(4):1664–1688, 2012.
- R. J. Tibshirani, A. Rinaldo, R. Tibshirani, and L. Wasserman. Uniform asymptotic inference and the bootstrap after model selection. *Ann. Stat.*, 46(3):1255–1287, 2018.
- A. Weinstein and A. Ramdas. Online control of the false coverage rate and false sign rate, arXiv:1905.01059, 2019.
- N. R. Zhang. DNA copy number profiling in normal and tumor genomes. In J. Feng, W. Fu, and F. Sun, editors, *Frontiers in Computational and Systems Biology*, pages 259–281. Springer, 2010.