



Universiteit
Leiden
The Netherlands

Only closed testing procedures are admissible for controlling false discovery proportions

Goeman, J.J.; Hemerik, J.; Solari, A.

Citation

Goeman, J. J., Hemerik, J., & Solari, A. (2021). Only closed testing procedures are admissible for controlling false discovery proportions. *Annals Of Statistics*, 49(2), 1218-1238.
doi:10.1214/20-AOS1999

Version: Submitted Manuscript (under Review)
License: [Leiden University Non-exclusive license](#)
Downloaded from: <https://hdl.handle.net/1887/3273770>

Note: To cite this publication please use the final published version (if applicable).

Only Closed Testing Procedures are Admissible for Controlling False Discovery Proportions

Jelle J. Goeman*

Jesse Hemerik[†]Aldo Solari[‡]

February 6, 2019

Abstract

We consider the class of all multiple testing methods controlling tail probabilities of the false discovery proportion, either for one random set or simultaneously for many such sets. This class encompasses methods controlling familywise error rate, generalized familywise error rate, false discovery exceedance, joint error rate, simultaneous control of all false discovery proportions, and others, as well as seemingly unrelated methods such as gene set testing in genomics and cluster inference methods in neuroimaging. We show that all such methods are either equivalent to a closed testing method, or are uniformly improved by one. Moreover, we show that a closed testing method is admissible as a method controlling tail probabilities of false discovery proportions if and only if all its local tests are admissible. This implies that, when designing such methods, it is sufficient to restrict attention to closed testing methods only. We demonstrate the practical usefulness of this design principle by constructing a uniform improvement of a recently proposed method.

1 Introduction

Closed testing (Marcus et al., 1976) is known to be a fundamental principle of familywise error rate (FWER) control in multiple hypothesis testing. Indeed, almost every well-known procedure controlling FWER has been shown to be a special case of closed testing, and many procedures have even been explicitly constructed as such. This is natural from a theoretical perspective, as it has been shown by Sonnemann (1982, 2008), and Sonnemann and Finner (1988) that closed testing is necessary for FWER control: every admissible procedure that controls FWER is a special case of closed testing. Romano et al. (2011) extended the results of Sonnemann and Finner. They proved that from a FWER perspective not every closed testing procedure is admissible, but only consonant procedures. These results are valuable for designers of FWER controlling methods, who can rely exclusively on closed testing as a general design principle. Alternative design principles exist, such as the partitioning principle (Finner and Strassburger, 2002) and sequential rejection (Goeman and Solari, 2010), but these are equivalent to closed testing.

Rather than only for FWER control, Goeman and Solari (2011) showed that closed testing may also be used to obtain simultaneous confidence bounds for the false discovery proportion (FDP) of all subsets within a family of hypotheses. Used in this way, closed testing allows a form of post-selection inference. It allows users to look at the data prior to choosing thresholds and criteria for significance, while still keeping control of tail probabilities of the FDP. The approach of Goeman and Solari (2011) is equivalent to an earlier approach by Genovese and Wasserman (2004, 2006) that did not explicitly use closed testing. A natural question that arises is whether similar results to those of Sonnemann (1982), Sonnemann and Finner (1988) and Romano et al.

*Dept. of Biomedical Data Sciences, Leiden University Medical Center, Leiden, The Netherlands

[†]Oslo Centre for Biostatistics and Epidemiology, University of Oslo, Oslo, Norway

[‡]Dept. of Economics, Management and Statistics, University of Milano-Bicocca, Milan, Italy

(2011) also hold for this novel use of closed testing. When controlling FDP, is it sufficient to look only at closed testing-based methods? Which closed testing-based methods are admissible? These are the questions we will address in this paper.

We will look at these questions within the class of procedures that control tail probabilities of FDP, both for a single (random) set and simultaneously over many such sets. Many procedures, some of which seem to target control of other quantities than FDP at first sight, will turn out to be within this general class. These include all methods with regular FWER control; FWER control of intersection hypotheses; k -FWER control; simultaneous k -FWER control; False Discovery Exceedance control; control of recently proposed error rates such as the Joint Error Rate; Simultaneous Selective Inference; and even methods constructing confidence intervals for the overall proportion of true (or false) hypotheses. A detailed overview with references is given in the next section. We show how all of these methods can be written as procedures that simultaneously control FDP over all subsets within a family of hypotheses, although they typically focus the power of the procedure on a limited number of such subsets.

After formally defining the class of procedures we will study, we follow the lead of Sonnemann (1982) and Sonnemann and Finner (1988), generalizing their concept of *coherence* to FDP control. We show that every incoherent procedure can be uniformly improved by a coherent one, and that every coherent procedure is either a closed testing procedure or can be uniformly improved by one. This shows that only closed testing procedures are admissible as FDP controlling procedures. Next, we show that every closed testing procedure with admissible local tests is admissible as an FDP controlling procedure. Together, these results lay out design principles for construction of FDP controlling procedures. To design admissible procedures it is sufficient to create a closed testing procedure with admissible local tests. To show admissibility for a procedure designed in a different way, it is sufficient to show that the procedure is equivalent to such a procedure. We will discuss practical implications of our results for researchers seeking to develop new methods, and illustrate them by constructing a uniform improvement of a method recently proposed by Katsevich and Ramdas (2018).

2 Inference on false discovery proportions

Assume that we have data \mathbf{X} distributed according to P_θ with $\theta \in \Theta$ unknown. About θ we may formulate hypotheses of the form $H \subseteq \Theta$. Let the family of hypotheses of interest be $(H_i)_{i \in I}$, where $I \subseteq C \subseteq \mathbb{N}$ is finite. The maximal set C , possibly infinite, is arbitrary here, but will become important in Section 5. For any finite $S \subseteq C$, let $S_0 = \{i \in S : \theta \in H_i\}$ be the index set of the true hypotheses and $S_1 = S \setminus S_0$ the index set of the false hypotheses. We will make no further model assumptions in this paper: any models, any test statistics, and any dependence structures will be allowed. Throughout the paper we will denote all random quantities in boldface. Equalities and inequalities between random variables should be read as holding almost surely for all $\theta \in \Theta$ unless otherwise stated. Proofs of all theorems, lemmas and propositions are in Appendix C.

We will be studying procedures with FDP control. The FDP of a finite set S is given by

$$\pi_0(S) = \frac{|S_0|}{|S| \vee 1}.$$

We define a procedure with FDP control on I (i.e. on $(H_i)_{i \in I}$) as a random function $\mathbf{q}^I : 2^I \rightarrow [0, 1]$, where 2^I is the power set of I , such that for all $\theta \in \Theta$ it satisfies

$$P_\theta(\pi_0(S) \leq \mathbf{q}^I(S) \text{ for all } S \subseteq I) \geq 1 - \alpha. \quad (1)$$

It will be more convenient to use an equivalent representation that gives a simultaneous lower $(1 - \alpha)$ -confidence bound for $|S_1|$, the number of true discoveries. We say that a random function $\mathbf{d}^I : 2^I \rightarrow \mathbb{R}$ has *true discovery guarantee* on I if, for all $\theta \in \Theta$,

$$P_\theta(\mathbf{d}^I(S) \leq |S_1| \text{ for all } S \subseteq I) \geq 1 - \alpha. \quad (2)$$

To see that the class of methods of FDP control and the class of methods with true discovery guarantee are equivalent, note that if \mathbf{q}^I fulfils (1), then

$$\mathbf{d}^I(S) = (1 - \mathbf{q}^I(S))|S|,$$

fulfils (2) and, if \mathbf{d}^I fulfils (2), then

$$\mathbf{q}^I(S) = \frac{|S| - \mathbf{d}^I(S)}{|S| \vee 1}$$

fulfils (1). In the rest of the paper we will focus on methods with true discovery guarantee, which are mathematically easier to work with than methods with FDP control, e.g. because they automatically avoid issues with empty sets S . Without loss of generality we may assume that $\mathbf{d}^I(S)$ takes integer values, and $0 \leq \mathbf{d}^I(S) \leq |S|$. If $\mathbf{d}^I(S)$ is not integer, we may freely replace $\mathbf{d}^I(S)$ by $\lceil \mathbf{d}^I(S) \rceil$.

The class of FDP control (cf. true discovery guarantee) procedures encompasses seemingly diverse methods. Only few authors (Blanchard et al., 2017; Genovese and Wasserman, 2006; Goeman et al., 2019; Goeman and Solari, 2011) have explicitly proposed procedures that target control of FDP for all sets S simultaneously as implied by (1). However, many other well-known types of multiple testing procedures turn out to be special cases of FDP control procedures, even if they were not directly formulated to control (1) or its equivalent. We will review these procedures briefly in the rest of this section in order to emphasize the wide range of applications of the results of this paper. We will reformulate such procedures in terms of \mathbf{d}^I .

Procedures that control FWER (e.g. Berk et al., 2013; Bretz et al., 2009; Janson et al., 2016; Westfall and Young, 1993) are usually defined as producing a random set \mathbf{K} (possibly empty) for which it is guaranteed that, for all $\theta \in \Theta$,

$$P_\theta(|\mathbf{K}_0| = 0) \geq 1 - \alpha.$$

A generalization, k -FWER (Finos and Farcomeni, 2011; Guo et al., 2010; Hommel and Hoffmann, 1988; Lehmann and Romano, 2005; Romano and Shaikh, 2006; Sarkar, 2007), makes sure that, for all $\theta \in \Theta$,

$$P_\theta(|\mathbf{K}_0| < k) \geq 1 - \alpha,$$

a statement that reduces to regular FWER if $k = 1$ is chosen. It is easily seen that this is equivalent to requiring (2) if we take

$$\mathbf{d}^I(S) = \begin{cases} |S| - k + 1 & \text{if } S = \mathbf{K}, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

Free additional statements may be obtained from (3) by direct logical implication. For example, if $\mathbf{d}^I(S) = |S| - k + 1$ then we may immediately set $\mathbf{d}^I(U) = |U| - k + 1$, if positive, for all $U \subseteq S$ without compromising (2). We will come back to such implications in Section 4.

Related to k -FWER are methods controlling False Discovery Exceedance (FDX), also known as γ -FDP, at level γ (Delattre et al., 2015; Dudoit et al., 2004; Farcomeni, 2009; Korn et al., 2004; Romano and Shaikh, 2006; Sun et al., 2015). Such methods find a random set \mathbf{K} (possibly empty) such that, for all $\theta \in \Theta$,

$$P_\theta(|\mathbf{K}_0| \leq \gamma|\mathbf{K}|) \geq 1 - \alpha,$$

which is equivalent to (2) with

$$\mathbf{d}^I(S) = \begin{cases} \lceil (1 - \gamma)|S| \rceil & \text{if } S = \mathbf{K}, \\ 0 & \text{otherwise.} \end{cases}$$

In most methods controlling FDX the control level γ is fixed, but it may also be random, as e.g. in the permutation-based method of Hemerik and Goeman (2018). Variants, such as kFDP (Guo et al., 2014), which allow a minimum number of false discoveries regardless of the size of \mathbf{K} , also fit (2).

Other methods allow γ to be chosen post-hoc by controlling FDX simultaneously over several values of γ . One way to achieve this is by control of the Joint Error Rate (JER). The JER (Blanchard et al., 2017) constructs a sequence of $\mathbf{m} \geq 0$ distinct random sets $\mathbf{K}_1, \dots, \mathbf{K}_m \subseteq I$ and corresponding random bounds $\mathbf{k}_1, \dots, \mathbf{k}_m$, such that, for all $\theta \in \Theta$,

$$P_\theta(|\{\mathbf{K}_i\}_0| < \mathbf{k}_i \text{ for all } 1 \leq i \leq \mathbf{m}) \geq 1 - \alpha. \quad (4)$$

This is a special case of (2) if we set

$$\mathbf{d}^I(S) = \begin{cases} |\mathbf{K}_i| - \mathbf{k}_i + 1 & \text{if } S = \mathbf{K}_i \text{ for some } 1 \leq i \leq \mathbf{m}, \\ 0 & \text{otherwise.} \end{cases}$$

Joint error rate control may be used with nested sets (Blanchard et al., 2017) or tree-structured sets (Durand et al., 2018), and is meant to be combined with interpolation (see Section 4). Similar approaches were used by e.g. the permutation-based methods of Meinshausen (2006) and Hemerik et al. (2019). Also the Simultaneous Selective Inference approach of Katsevich and Ramdas (2018), discussed in detail in Section 9, can be seen as controlling JER with nested sets.

A different category of methods involves FWER control of many intersection hypotheses, as e.g. used in gene set testing in genomics and in cluster inference in neuroimaging. In genomics, a collection of distinct sets $K_1, \dots, K_m \subseteq I$ is given a priori, and the procedure generates corresponding random indicators $\mathbf{k}_1, \dots, \mathbf{k}_m \in \{0, 1\}$ for detection of signal in the corresponding set. FWER is controlled over all statements made, i.e., for all $\theta \in \Theta$,

$$P_\theta(|\{K_i\}_1| \geq \mathbf{k}_i \text{ for all } i = 1, \dots, m) \geq 1 - \alpha. \quad (5)$$

This corresponds to (2) with

$$\mathbf{d}^I(S) = \begin{cases} \mathbf{k}_i & \text{if } S = K_i \text{ for some } 1 \leq i \leq m, \\ 0 & \text{otherwise.} \end{cases}$$

Examples of such methods include Meinshausen (2008), Goeman and Mansmann (2008), Goeman and Finos (2012), Meijer and Goeman (2015b), Meijer et al. (2015), and Meijer and Goeman (2015a). In the latter two papers a connection with FDP control was already noted. In neuroimaging, cluster inference methods are similar except that in this case the sets $\mathbf{K}_1, \dots, \mathbf{K}_m$ and their number $\mathbf{m} \geq 0$ are random, and $k_i = 1$ for $i = 1, \dots, \mathbf{m}$ is fixed (Forman et al., 1995; Friston et al., 1996; Poline and Mazoyer, 1993). FWER control (5) is guaranteed by Gaussian random field theory. Such control translates to a true discovery guarantee requirement (2) in the same way as in the genomics case.

In partial conjunction testing (Benjamini and Heller, 2008; Wang and Owen, 2018), the hypothesis $H_0^{k/n} : |I_1| < k$ is tested for some $1 \leq k \leq n$. The requirement that $\delta \in \{0, 1\}$ is a valid test of $H_0^{k/n}$ is equivalent to (2) with

$$\mathbf{d}^I(S) = \begin{cases} \delta k & \text{if } S = I, \\ 0 & \text{otherwise.} \end{cases}$$

Finally, related to partial conjunction methods are methods for aim to make one-sided confidence intervals for $|I_0|/|I|$, the proportion of true null hypotheses in the testing problem as a whole (Ge and Li, 2012; Meinshausen and Rice, 2006). Here, the requirement that $[0, \mathbf{u}]$ is a valid confidence interval for $|I_0|/|I|$ is equivalent to demanding (2) with

$$\mathbf{d}^I(S) = \begin{cases} (1 - \mathbf{u})|I| & \text{if } S = I, \\ 0 & \text{otherwise.} \end{cases}$$

This listing of the different types of methods that may be written as true discovery guarantee methods is certainly not exhaustive, but a general pattern emerges. Any method controlling a $(1 - \alpha)$ -tail probability of the

number or proportion of true discoveries (from below) or false discoveries (from above) either in one subset of I , or in several subsets simultaneously, are special cases of general discovery guarantee procedures. The sets and bounds are all allowed to be random; only α must be fixed.

Writing procedures as true discovery control procedures, even when the rewriting is trivial, may bring a new perspective to the use of the procedure. As proposed by Goeman and Solari (2011), procedures that fulfil (1) or (2) allow a different, flexible way of using multiple testing methods. In flexible multiple testing the user may look at the data before choosing post hoc one or several sets $S \subseteq I$ of interest, based on any desired criteria, and find their $\mathbf{d}^I(S)$. Regardless of this data peeking the bounds on the selected sets are simultaneously valid due to the simultaneity in (2). Writing procedures in this form, therefore, in principle opens the way to their use as post-selection inference methods (see Rosenblatt et al., 2018, for an application). Of course, this is only useful if the user has some real choice, i.e. if $\mathbf{d}^I(S) \neq 0$ for a decent number of sets S . We will see in Section 4 how to get rid of some of the zeros in the definitions above.

3 True discovery guarantee using closed testing

A general way to construct procedures with true discovery guarantee is using closed testing, a method introduced by Marcus et al. (1976) as a means of constructing methods for FWER control. Genovese and Wasserman (2006) and Goeman and Solari (2011) adapted closed testing to make it useable for true discovery guarantee and FDP control. We will briefly review these methods here.

For every finite set $S \subseteq C$ we define a corresponding intersection hypothesis as $H_S = \bigcap_{i \in S} H_i$. This hypothesis is true if and only if all $H_i, i \in S$ are true. We have $H_\emptyset = \Theta$, which is always true. For every such intersection hypothesis H_S we may choose a *local test* ϕ_S taking values in $\{0, 1\}$, with 1 indicating rejection of H_S . This is a valid statistical test for H_S if it has the property that

$$\sup_{\theta \in H_S} \mathbb{P}_\theta(\phi_S = 1) \leq \alpha.$$

We always choose $\phi_\emptyset = 0$ surely. Choosing a local test for every finite $S \subseteq C$ will yield a *suite of local tests* $\phi = (\phi_S)_{S \subseteq C, |S| < \infty}$. To deal with restricted combinations (Shaffer, 1986) efficiently, if present, we demand that identical hypotheses have identical tests, i.e. if for $U, V \subseteq C$ we have $H_U = H_V$, then we also have $\phi_U = \phi_V$. If $H_U = \emptyset$ for some $U \subseteq C$, we may take $\phi_U = 1$ surely.

From such a suite of local tests we may obtain a true discovery guarantee procedure in two simple steps. First, we need to correct the tests for multiple testing. We define the *effective local test* within the family I by

$$\phi_S^I = \min\{\phi_U : S \subseteq U \subseteq I\}.$$

As shown by Marcus et al. (1976), the effective local tests have FWER control over all intersection hypotheses $H_S, S \subseteq I$, i.e., for all $\theta \in \Theta$,

$$\mathbb{P}_\theta(\phi_S^I \leq |S| \text{ for all } S \subseteq I) \geq 1 - \alpha.$$

Next, we calculate $\mathbf{d}^I(S)$. We see that the procedure defined by $\mathbf{d}^I(S) = \phi_S^I$ already fulfils (2). More recently, however, Goeman and Solari (2011) showed that closed testing may also be used for more powerful FDP control. For any suite of local tests ϕ , these authors defined the associated procedure

$$\mathbf{d}_\phi^I(S) = \min_{U \in 2^S} \{|S \setminus U| : \phi_U^I = 0\}, \quad (6)$$

and proved that it has true discovery guarantee. Note that the minimum is always defined since $\phi_\emptyset^I = \phi_\emptyset = 0$ surely.

An earlier general approach to developing procedures with true discovery guarantee was developed without reference to closed testing by Genovese and Wasserman (2006). Starting from a suite of local tests, they proved true discovery guarantee for the procedure

$$\mathbf{g}_\phi^I(S) = \min_{V \in 2^I} \{|S \setminus V| : \phi_V = 0\}. \quad (7)$$

This procedure is equivalent to the procedure (6), as was shown by Hemerik et al. (2019). This result was hidden in the supplemental information to that paper, and it deserves some more prominence. The proof is short. We repeat it here.

Lemma 1. $\mathbf{g}_\phi^I = \mathbf{d}_\phi^I$.

The expressions (6) and (7) are very useful for constructing methods with true discovery guarantee. Local tests tend to be easy to specify in most models, as each local test is a test of a single hypothesis, so that standard statistical test theory may be used. Given a suite of local tests, (6) or (7) takes care of the multiplicity. A computational problem remains, of course, if $|I|$ is large, since direct application of (6) and (7) takes exponential time. Often, however, shortcuts are available that allow faster computation (Goeman et al., 2019; Goeman and Solari, 2011). We'll see an example in Section 9.

Comparing (6) and (7), the single step expression of Genovese and Wasserman (2006) is clearly more elegant. However, the link of (6) to closed testing is valuable because it connects true discovery guarantee procedures to the enormous literature on closed testing (see Henning and Westfall, 2015, for an overview). We found that the detour via effective local tests is often profitable in practice because expressions for $\mathbf{d}_\phi^I(S)$ can be easier to derive through expressions for ϕ_S^I , as we shall see in Section 9 (also Goeman et al., 2019; Hemerik and Goeman, 2018).

4 Coherence and interpolation

By viewing methods in terms of true discovery guarantee, as we have done in Section 2, they are upgraded from making a confidence statement about discoveries in a limited number of sets $S \subseteq I$ to doing the same for all subsets of I . However, in the definitions above, most of these statements are the trivial $\mathbf{d}^I(S) = 0$. Often, however, some of the statements can be uniformly improved by a process called *interpolation*. In this section we discuss interpolation and how it can improve true discovery guarantee procedures. We will define coherent procedures as procedures that cannot be improved by interpolation.

Let \mathbf{d}^I be some procedure with true discovery guarantee. We define the *interpolation* $\bar{\mathbf{d}}^I$ of \mathbf{d}^I as

$$\bar{\mathbf{d}}^I(S) = \max_{U \in 2^I} \left\{ \mathbf{d}^I(U) - |U \setminus S| + \mathbf{d}^I(S \setminus U) \right\}. \quad (8)$$

Interpolation was used in weaker versions or in specific cases by several authors (Blanchard et al., 2017; Durand et al., 2018; Genovese and Wasserman, 2006; Meinshausen, 2006). Taking $U = S$, we see that $\bar{\mathbf{d}}^I(S) \geq \mathbf{d}^I(S)$. Moreover, the improvement from \mathbf{d}^I to $\bar{\mathbf{d}}^I$ is for free, as noted in the following lemma.

Lemma 2. *If \mathbf{d}^I has true discovery guarantee then also $\bar{\mathbf{d}}^I$.*

Intuitively, the rationale for interpolation is as follows. If $\mathbf{d}^I(U)$ is large, and S has so much overlap with U that the signal $\mathbf{d}^I(U)$ in U does not fit in $U \setminus S$, then the remaining signal must be in S . Since this reasoning follows by direct logical implication, it will not increase the occurrence of type I error: we can only make an erroneous statement about S if we had already made one about U . As an example, consider interpolation for k -FWER controlling procedures. The interpolated version of (3) is simply

$$\bar{\mathbf{d}}^I(S) = 0 \vee (|S \cap \mathbf{K}| - k + 1), \quad (9)$$

an expression that simplifies even further to $\bar{\mathbf{d}}^I(S) = |S \cap \mathbf{K}|$ with regular FWER when $k = 1$.

Interpolation is not necessarily a one-off process, and interpolated procedures may sometimes be further improved by another round of interpolation. We call a procedure *coherent* if it cannot be improved by interpolation, i.e. if

$$\bar{\mathbf{d}}^I(S) = \mathbf{d}^I(S) \quad \text{for all } S \subseteq I. \quad (10)$$

We can characterize coherent procedures further with the following lemma.

Lemma 3. \mathbf{d}^I is coherent if and only if for every disjoint $V, W \subseteq I$ we have

$$\mathbf{d}^I(V) + \mathbf{d}^I(W) \leq \mathbf{d}^I(V \cup W) \leq \mathbf{d}^I(V) + |W|. \quad (11)$$

We intentionally use the same term *coherent* that was used by Sonnemann (1982) in the context of FWER control of intersection hypotheses. Looking only at FWER control of intersection hypotheses is equivalent to looking only at $\mathbb{1}\{\mathbf{d}^I(S) > 0\}$ for every S , where $\mathbb{1}\{\cdot\}$ denotes an indicator function. In that case (10) reduces to simply requiring that $U \subseteq V$ and $\mathbf{d}^I(U) > 0$ implies that $\mathbf{d}^I(V) > 0$, which is exactly Sonnemann’s definition of coherence.

Methods that are created through closed testing are automatically coherent, as the following lemma claims.

Lemma 4. The procedure \mathbf{d}_ϕ^I is coherent.

Since an incoherent procedure can always be replaced by a coherent procedure that is at least as good, we will restrict attention to coherent procedures for the rest of this paper.

5 Monotone procedures

The methods from the literature discussed in Sections 2 and 3 are typically not defined for a specific family I of hypotheses, but are usually formulated as generic procedures that can be used for any family, large or small. Researchers developing methods are usually not looking for good properties for a specific family at a specific scale $|I|$, but for methods that are generally applicable and have good properties whatever I .

We can embed the procedure \mathbf{d}^I into a stack of procedures $\mathbf{d} = (\mathbf{d}^I)_{I \subseteq C, |I| < \infty}$, where we may have some maximal family $C \subseteq \mathbb{N}$. We will briefly call \mathbf{d} a *monotone procedure* if it fulfils the three criteria below. In contrast, we call \mathbf{d}^I for a specific I a *local procedure*, or a local member of \mathbf{d} .

1. *true discovery guarantee*: \mathbf{d}^I has true discovery guarantee for every finite $I \subseteq C$;
2. *coherence*: \mathbf{d}^I is coherent for every finite $I \subseteq C$;
3. *monotonicity*: $\mathbf{d}^I(S) \geq \mathbf{d}^J(S)$ for every finite $S \subseteq I \subseteq J \subseteq C$.

The first two criteria are no more than natural. We demand true discovery guarantee for every member of the monotone procedure, and we demand coherence for every local member since otherwise we may always improve it by a coherent procedure. The monotonicity requirement relates local procedures at different scales to each other. It says that inference on the number of discoveries in a set S should never get better if we embed S in a larger family J rather than in a smaller family I . As the multiple testing problem gets larger, inference should get more difficult. This requirement relates closely to the “subsetting property” of Goeman and Solari (2014) and the monotonicity property of various FWER control procedures (e.g. Bretz et al., 2009; Goeman and Solari, 2010). It is also a natural requirement also for FDP controlling procedures, and the procedures cited in Section 2 generally adhere to it by construction.

There are a few notable exceptions to the rule that method designers tend to design monotone rather than local procedures. The examples we are aware of are all FWER-controlling procedures. Rosenblum et al. (2014)

proposed a local procedure for $|I| = 2$ hypotheses that optimizes the power for rejecting at least one of these. Their method is specific for the scale $|I|$ it was defined for; extensions to $|I| > 2$ do not exist (Rosset et al., 2018), and natural extensions to $|I| < 2$ do not fulfill the monotonicity requirement. Rosset et al. (2018) developed methods that optimize any-power for specific scales $|I|$ under an exchangeability assumption, and that also have non-monotone behavior. We remark, however, that every coherent local procedure \mathbf{d}^I with true discovery guarantee may be trivially embedded in a monotone procedure with $C = I$ (or even $C = \mathbb{N}$) by setting

$$\mathbf{d}^J(S) = \begin{cases} \mathbf{d}^I(S) & \text{if } S \subseteq I, \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

This embedding is in itself not useful, but it allows translation of some properties we will derive of monotone procedures to properties of their local members.

Procedures created from a suite of local tests are automatically monotone, as formalized in the following lemma.

Lemma 5. *The procedure $\mathbf{d}_\phi = (\mathbf{d}_\phi^I)_{I \subseteq C, |I| < \infty}$ is a monotone procedure.*

We will mostly be studying monotone procedures in this paper, but investigate implications for local procedures where appropriate.

The property of primary interest is admissibility. Let us formally define admissibility for true discovery guarantee procedures. Recall that a statistical test δ of a hypothesis H is uniformly improved by a statistical test $\tilde{\delta}$ of the same hypothesis if (1.) $\tilde{\delta} \geq \delta$; and (2.) $P_\theta(\tilde{\delta} > \delta) > 0$ for some $\theta \in \Theta$. A statistical test is *admissible* if no test exists that uniformly improves it (Lehmann and Romano, 2006, Section 6.7). We call a suite of local tests ϕ admissible if ϕ_S is admissible for all finite $\emptyset \subset S \subseteq C$. We note that existence of admissible tests is not assured in all models, but that under a weak condition all tests that exhaust the α -level are admissible. We discuss these technical issues in Appendix A, where we also motivate our definition of admissibility compared to alternatives in the literature.

Analogously to admissibility of single tests we define admissibility for procedures with true discovery guarantee. A uniform improvement of a monotone procedure \mathbf{d} is a monotone procedure $\tilde{\mathbf{d}}$ such that (1.) $\tilde{\mathbf{d}}^I(S) \geq \mathbf{d}^I(S)$ for all finite $S \subseteq I \subseteq C$; and (2.) $P_\theta(\tilde{\mathbf{d}}^I(S) > \mathbf{d}^I(S)) > 0$ for some $\theta \in \Theta$ and some finite $S \subseteq I \subseteq C$. A uniform improvement of a local procedure \mathbf{d}^I is a local procedure $\tilde{\mathbf{d}}^I$ such that (1.) $\tilde{\mathbf{d}}^I(S) \geq \mathbf{d}^I(S)$ for all $S \subseteq I$; and (2.) $P_\theta(\tilde{\mathbf{d}}^I(S) > \mathbf{d}^I(S)) > 0$ for some $\theta \in \Theta$ and some $S \subseteq I$. We call a local or monotone procedure that cannot be uniformly improved admissible. If all local members of a monotone procedure are admissible, then the monotone procedure is admissible, but the converse is not necessarily true, as illustrated in Appendix B.

6 All admissible procedures are closed testing procedures

Theorem 1, below, claims that every monotone procedure that has discovery control is either equivalent to a closed testing procedure or can be uniformly improved by one. We already know from Lemma 3 that every incoherent procedure can be uniformly improved by a coherent procedure. It follows that every procedure that is not equivalent to a closed testing procedure is inadmissible: the class of all closed testing procedures is essentially complete (Lehmann and Romano, 2006, Section 1.8) for true discovery guarantee. This is the first main result of this paper.

Theorem 1. *Let \mathbf{d} be a monotone procedure. Then, for every finite $S \subseteq C$,*

$$\phi_S = \mathbf{1}\{\mathbf{d}^S(S) > 0\}$$

is a valid local test of H_S . For the suite $\phi = (\phi_S)_{S \subseteq C, |S| < \infty}$ we have, for all $S \subseteq I \subseteq C$ with $|I| < \infty$,

$$\mathbf{d}_\phi^I(S) \geq \mathbf{d}^I(S).$$

Coherence is necessary but not sufficient to guarantee admissibility. The procedure $\mathbf{d}_\phi^I(S)$ implied by Theorem 1 may in some cases be truly a uniform improvement over the original, coherent $\mathbf{d}^I(S)$. To see a classical example in which a coherent procedure can uniformly improved by a closed testing argument, think of Bonferroni. Combined with (9), Bonferroni is coherent. However, it is uniformly improved by Holm's procedure that follows from a well-known step-down argument. This stepping-down can be seen as a direct application of closed testing with the local test defined in Theorem 1. Step-down arguments are standard for FWER control and have been applied to several FDP controlling methods in the past (Blanchard et al., 2017; Goeman et al., 2019; Hemerik et al., 2019).

It should be noted that in case of a monotone procedure, the local test ϕ_S defined in Theorem 1 is truly local, in the sense that it uses only the information used by the restricted testing problem \mathbf{d}^S about the hypotheses H_i , $i \in S$. For example, in a testing problem based on p -values, the local test would use only the p -values p_i , $i \in S$. In other testing problems, some global information may be used, e.g. the overall estimate of σ^2 in a large one-way ANOVA, but still in such situations the local test is very natural: as a local test for H_S we use the test for discovery of signal in hypotheses H_i , $i \in S$, that we would use in the situation where the hypotheses H_i , $i \notin S$ are not of interest to us. Such a local test is implicitly defined by the local procedure \mathbf{d}^S .

The result of the theorem is formulated in terms of monotone procedures. It applies immediately to local procedures as well if we use the trivial embedding (12) of a local procedure into a monotone one. With this embedding we even have $\mathbf{d}_\phi^I(S) = \mathbf{d}^I(S)$. This leads to the following corollary.

Corollary 1. *Let \mathbf{d}^I be a coherent procedure. Then, for every $S \subseteq I$,*

$$\phi_S = \mathbb{1}\{\mathbf{d}^I(S) > 0\}$$

is a valid local test of H_S . For the suite $\phi = (\phi_S)_{S \subseteq I}$ we have, for all $S \subseteq I$,

$$\mathbf{d}_\phi^I(S) = \mathbf{d}^I(S).$$

Corollary 1 shows that every coherent procedure with true discovery guarantee is equivalent to a closed testing procedure. It may possibly be uniformly improved by another closed testing procedure if the suite of local tests ϕ is not admissible, as we shall see in the next section. The embedding into a monotone procedure used in Theorem 1 helps to define the local tests in terms of the local procedure \mathbf{d}^S , rather than in terms of the original \mathbf{d}^I , allowing for more intuitive, and possibly more powerful local tests.

7 All closed testing procedures are admissible

So far we have seen that a procedure with true discovery guarantee may be uniformly improved by interpolation to coherent procedures. This procedure in turn may be uniformly improved by closed testing procedure. Clearly, equivalence to a closed testing procedure is necessary for admissibility. Are all closed testing procedures admissible? In this section we derive a simple condition for admissibility of monotone procedures that is both necessary and sufficient. We show that admissibility of the monotone procedure \mathbf{d}_ϕ follows directly from admissibility of its local tests. This is the second main result of this paper.

Theorem 2. *\mathbf{d}_ϕ is admissible if and only if the suite ϕ is admissible.*

We have already seen from Theorem 1 that only closed testing procedures are admissible. Theorem 2 says that all closed testing procedures are admissible, provided they fulfil the reasonable demand that they are built

from admissible local tests. To check admissibility of the local tests, Appendix A shows that under a weak assumption it is sufficient to check that the local tests exhaust the α -level. Theorem 2 thus makes it easy to guarantee admissibility of monotone procedures.

Unlike Theorem 1, the result of Theorem 2 does not immediately translate to local procedures: even if ϕ is admissible, it may happen for some finite $I \subseteq C$ that \mathbf{d}_ϕ^I can be uniformly improved by some other procedure \mathbf{d}^I . About such local improvements we have the following proposition.

Proposition 1. *If $\mathbf{d}^I \geq \mathbf{d}_\phi^I$ is admissible, then there is an admissible ψ such that $\mathbf{d}^I = \mathbf{d}_\psi^I$ and, for all $S \subseteq I$, $\psi_S \geq \phi_S^I$.*

Proposition 1 limits the available room for local improvements of admissible monotone procedures. Combining Proposition 1 and Theorem 2 we see that such improvements have to be admissible monotone procedures, and therefore closed testing procedures, themselves. The difference between ϕ and ψ , if both are admissible, is that for every $S \subseteq I$, ϕ_S uses only the local information in $\mathbf{d}_\phi^S(S)$, but the same does not necessarily hold for ψ_S .

In Appendix B we give an example of a local improvement of a monotone procedure. In this example Θ is restricted in such a way that the local tests are known to have limited power even under the alternative. We show that this knowledge can be exploited to obtain local improvements. In the example it is crucial that, for some $S \subset I$,

$$\sup_{\theta \in H_S} P_\theta(\phi_S^I = 1) = \tilde{\alpha} < \alpha,$$

Since the effective local test ϕ_S^I never exhausts the α -level there is room $\alpha - \tilde{\alpha}$ to be exploited by ψ_S , which can be constructed to create a local improvement. Local improvements are also possible in case null hypotheses are composite, using the Partitioning Principle, as shown in Finner and Strassburger (2002), examples 4.1–4.3, and Goeman and Solari (2010), section 4. For many well-known procedures, e.g. Holm’s procedure under arbitrary dependence, local improvements do not exist. However, we have no general theory on the relationship between admissibility of a monotone procedure and admissibility of its local members. We leave this as an open problem.

8 Consonance and familywise error

Theorem 2 establishes a necessary and sufficient condition for admissibility of monotone true discovery guarantee procedures, and therefore of FDP-controlling procedures. At first sight, our results may seem at odds with those of Romano et al. (2011), who proved that for FWER control, which is a special case of true discovery guarantee, only consonant procedures are admissible. However, this seeming contradiction disappears when we realize that admissibility of a procedure as a true discovery control procedure does not automatically imply admissibility as a FWER controlling procedure and vice versa. In this section we take a sidestep to FWER control, investigating the concept of consonance, and extending some of the results of Romano et al. (2011) on admissibility of FWER controlling procedures.

We call a procedure \mathbf{d}^I *consonant* if it has the property that for every $S \subseteq I$, $\mathbf{d}^I(S) > 0$ implies that for at least one $i \in S$ we have $\mathbf{d}^I(\{i\}) = 1$, almost surely for all $\theta \in \Theta$. If $\mathbf{d}^I = \mathbf{d}_\phi^I$, this is equivalent to the more usual formulation in terms of the suite ϕ , that $\phi_S^I = 1$ implies that for at least one $i \in S$ we have $\phi_{\{i\}}^I = 1$, almost surely for all $\theta \in \Theta$. We call a monotone procedure \mathbf{d} consonant if all local members \mathbf{d}^I , $I \subseteq C$ finite, are consonant. We call a suite ϕ consonant if all ϕ_S , $S \subseteq C$ finite, are consonant.

Conceptually, consonant procedures allow pinpointing of effects. If $\mathbf{d}^I(S) > 0$, signal has been detected somewhere in S . A consonant procedure in this case can always find at least one elementary hypothesis to pin the effect down on. This is a desirable property, as it can be unsatisfactory for a researcher to know that an effect exist but not where it can be found. On the other hand, Goeman et al. (2019) argued that non-consonant procedures can be far more powerful in large-scale multiple testing procedures than consonant ones.

For consonant procedures a stronger version of Lemma 3 holds.

Lemma 6. \mathbf{d}^I is consonant and coherent if and only if, for every disjoint $V, W \subseteq I$,

$$\mathbf{d}^I(V \cup W) = \mathbf{d}^I(V) + \mathbf{d}^I(W). \quad (13)$$

Classically, focus in the literature on closed testing has been on FWER controlling procedures (Henning and Westfall, 2015). An FWER-controlling procedure on a finite $I \subseteq C$ returns a set $\mathbf{R}^I \subseteq I$ such that, for all $\theta \in \Theta$,

$$P_\theta(|\mathbf{R}_0^I| = 0) \geq 1 - \alpha.$$

As argued in Section 2, we can relate FWER controlling procedures to true discovery guarantee procedures and vice versa. If \mathbf{R}^I is a FWER controlling procedure, then \mathbf{r}^I with

$$\mathbf{r}^I(S) = |S \cap \mathbf{R}^I|,$$

for all $S \subseteq I$, is a coherent procedure with true discovery guarantee, as we know from (9). Conversely, if \mathbf{d}^I is a coherent procedure with true discovery guarantee, then

$$\mathbf{R}^I = \{i \in I : \mathbf{d}^I(\{i\}) = 1\}$$

is a FWER controlling procedure. Both types of procedures may be created from local tests. The FWER controlling procedure from the suite ϕ is given by

$$\mathbf{R}_\phi^I = \{i \in I : \phi_{\{i\}}^I = 1\} = \{i \in I : \mathbf{d}_\phi^I(\{i\}) = 1\}. \quad (14)$$

We can compare the procedure \mathbf{d}_ϕ defined from ϕ through (6) with the procedure

$$\mathbf{r}_\phi^I(S) = |S \cap \mathbf{R}_\phi^I|,$$

indirectly defined through \mathbf{R}_ϕ^I . This is the procedure that discards all information in \mathbf{d}_ϕ^I that is not contained in \mathbf{R}_ϕ^I . Lemma 7 describes consonance as the property that no information is lost in the process.

Lemma 7. If \mathbf{d}_ϕ^I is consonant, $\mathbf{d}_\phi^I = \mathbf{r}_\phi^I$; otherwise \mathbf{d}_ϕ^I uniformly improves \mathbf{r}_ϕ^I .

If FWER control is what we are after, however, we must look at admissibility of \mathbf{R}^I directly. As with procedures with true discovery guarantee, we will focus on monotone (stacks of) procedures defined for all finite $I \subseteq C$. We call a procedure $\mathbf{R} = (\mathbf{R}^I)_{I \subseteq C, |I| < \infty}$ monotone if for all finite $J \subseteq I \subseteq C$ we have

$$\mathbf{R}^J \supseteq \mathbf{R}^I \cap J.$$

As above for true discovery guarantee procedures, it asserts that enlarging the multiple testing problem from J to I will never increase the number of rejections in J (Bretz et al., 2009; Goeman and Solari, 2010). Analogous to the definition in Section 5, we define a uniform improvement of a monotone FWER control procedure \mathbf{R} as a monotone FWER control procedure $\tilde{\mathbf{R}}$ such that (1.) $\tilde{\mathbf{R}}^J \supseteq \mathbf{R}^I$ for all finite $I \subseteq C$; and (2.) $P_\theta(\tilde{\mathbf{R}}^I \supset \mathbf{R}^I) > 0$ for some $\theta \in \Theta$ and some finite $I \subseteq C$. A procedure is admissible if no uniform improvement exists. What can we say about admissibility of FWER control procedures?

Romano et al. (2011) showed that consonance is necessary for admissibility of FWER controlling procedures. Proposition 2 is a variant of the result of Romano et al. (2011) for monotone procedures.

Proposition 2. If \mathbf{R} is admissible, then a consonant suite ϕ exists such that $\mathbf{R} = \mathbf{R}_\phi$.

We also have a second necessary condition for admissibility.

Proposition 3. *If \mathbf{R} is admissible, then an admissible suite ψ exists such that $\mathbf{R} = \mathbf{R}_\psi$.*

It would be tempting to conclude from Propositions 2 and 3 that if \mathbf{R} is admissible, then $\mathbf{R} = \mathbf{R}_\phi$ with ϕ consonant and admissible. Certainly under weak assumptions we may choose $\phi = \psi$. For example, if some ϕ_S is inadmissible because it fails to exhaust the α -level we may choose ψ_S as ϕ_S plus a randomized multiple of $\phi_{\{i\}}$ for some $i \in S$, if we allow randomized tests, and use Lemma 9. However, we were unable to prove in full generality that $\phi = \psi$ is always possible. Perhaps in some awkward models admissible local tests cannot be consonant, and consonant tests cannot be admissible. We leave the question open when $\phi = \psi$ is possible. In converse however, if we can find a ϕ that is both admissible and consonant, we have an admissible procedure:

Proposition 4. *If ϕ is consonant and admissible, then \mathbf{R}_ϕ is admissible.*

9 Application: improving an existing method

We will now illustrate how existing methods may be improved by embedding them in a closed testing procedure and using the results of this paper. We chose a method recently proposed by Katsevich and Ramdas (2018). This elegant method, which we abbreviate K&R, allows users to choose a p -value cutoff for significance post hoc, and uses stochastic process arguments to control both FDP and FDR. We focus on the FDP control property here.

Take $I = \{1, \dots, m\}$. Let hypotheses H_1, \dots, H_m have corresponding p -values $\mathbf{p}_1, \dots, \mathbf{p}_m$ independent, and with \mathbf{p}_i standard uniform, or possibly stochastically larger, if H_i is true. Katsevich and Ramdas proposed a method that controls (4) for sets $\mathbf{K}_1 \subset \dots \subset \mathbf{K}_m$, where \mathbf{K}_i consists of the indices of the i hypotheses with smallest p -values, with ties broken arbitrarily. As shown by Katsevich and Ramdas for these sets we have, if $\alpha \leq 0.31$,

$$P(|\{\mathbf{K}_i\}_0| \leq c(1 + m\mathbf{p}_{(i)}) \text{ for all } 1 \leq i \leq m) \geq 1 - \alpha,$$

where $c = -\log(\alpha)/\log(1 - \log(\alpha))$, and $\mathbf{p}_{(i)}$ is the i th smallest p -value. For $\alpha = 0.05$ we have $c \approx 2.163$. As in Section 2 we can write this as a procedure with true discovery guarantee on I by writing

$$\mathbf{d}^I(S) = \begin{cases} 0 \vee [i - c(1 + m\mathbf{p}_{(i)})] & \text{if } S = \mathbf{K}_i \text{ for some } 1 \leq i \leq m, \\ 0 & \text{otherwise,} \end{cases} \quad (15)$$

where we round up to ensure that $\mathbf{d}^I(S)$ is always an integer.

Is the procedure (15) admissible, and if not, how can we improve it? We apply the results of this paper. First, we remark that the method as defined is not coherent. To interpolate, we introduce the notation $\mathbf{p}_{(i:S)}$, with $1 \leq i \leq |S|$, for the i th smallest p -value among the multiset $\{p_i : i \in S\}$, so that in particular $\mathbf{p}_{(i)} = \mathbf{p}_{(i:I)} = \mathbf{p}_{(i:\mathbf{K}_i)}$. The interpolation of the procedure follows directly from (8), yielding

$$\mathbf{d}^I(S) = 0 \vee \max_{i=1, \dots, m} [|\mathbf{K}_i \cap S| - c(1 + m\mathbf{p}_{(i:I)})]. \quad (16)$$

To simplify this expression, call $g(i) = |\mathbf{K}_i \cap S| - c(1 + m\mathbf{p}_{(i:I)})$. Let π_1, \dots, π_m be the permutation such that $\mathbf{K}_i = \{\pi_1, \dots, \pi_i\}$ for all $1 \leq i \leq m$. If $\pi_i \notin S$, then either $i > 1$ and $g(i) \leq g(i-1)$ or $i = 1$ and $g(i) < 0$, so we may restrict the maximum in (16) to values of i with $\pi_i \in S$. We have $S = \{\pi_{j_1}, \dots, \pi_{j_{|S|}}\}$ with $j_1 < \dots < j_{|S|}$. If $i = \pi_{j_k}$, then $g(i) = k - c(1 + m\mathbf{p}_{(k:S)})$. Therefore, (16) reduces to

$$\mathbf{d}^I(S) = 0 \vee \max_{k=1, \dots, |S|} [k - c(1 + m\mathbf{p}_{(k:S)})], \quad (17)$$

taking $\mathbf{d}^I(\emptyset) = 0$ implicitly. We note that interpolated method makes non-trivial statements for sets S not of the form \mathbf{K}_i , and may even improve $\mathbf{d}^I(\mathbf{K}_i)$ for some i . It may be checked using Lemma 3 that the procedure

(17) is coherent, so no further rounds of interpolation are needed. The K&R procedure was not developed for a specific scale m . Writing $|I|$ for m in (17) we have a procedure that is defined for general I , and it is easy to check that $(\mathbf{d}^I)_{I \subseteq \mathbb{N}, |I| < \infty}$ is monotone.

Next, we use Theorem 1 to embed the method in a closed testing procedure, which hopefully results in further improvement of the procedure. By the theorem, the local test for finite $\emptyset \neq S \subseteq \mathbb{N}$ is given by

$$\phi_S = \mathbb{1}\{\mathbf{p}_{(i:S)} \leq (i - c)/c|S| \text{ for at least one } i = 1, \dots, |S|\}. \quad (18)$$

We will construct the closed testing procedure based on this local test. To do this, we note that the local test is very similar to the Simes (1986) local test studied in the context of FDP control by Goeman et al. (2019). We generalize a result from that paper to local tests of the form

$$\phi_S = \mathbb{1}\{\mathbf{p}_{(i:S)} \leq l_{i:|S|} \text{ for at least one } i = 1, \dots, |S|\}, \quad (19)$$

where we assume that $l_{i:m} \geq l_{i:n}$ for all $m \leq n$. For convenience, we let $l_{i:n}$ be defined for all $i = 1, 2, \dots$ and $n = 0, 1, \dots$, i.e. even if $i > n$. Without loss of generality we can take $l_{i:n} = l_{i:i}$ for $i > n \neq 0$, and $l_{i:0} = 1$ for all i .

The general form (19) encompasses the local test (18), taking

$$l_{i:m} = \frac{i - c}{cm} \quad (20)$$

if $m \neq 0$, and the Simes test, taking $l_{i:m} = i\alpha/m$ if $m \neq 0$. It can also be used for a much broader range of tests, e.g. to the local tests implied by the False Discovery Rate controlling procedures of Blanchard and Roquain (2009), to higher criticism (Donoho et al., 2004), to the local tests implied by the Dvoretzky-Kiefer-Wolfowitz inequality (Genovese and Wasserman, 2004; Meinshausen, 2006), to the local tests implied by second and higher order generalized Simes constants (Cai and Sarkar, 2008; Gou and Tamhane, 2014), and to the local tests implied by the FDR controlling procedures of Benjamini and Liu (1999), and Romano and Shaikh (2006, equation 4.1).

Generalizing a result of Goeman et al. (2019) for the Simes local test we prove that calculation of $\mathbf{d}_\phi^I(S)$ for local tests of type (19) can be done in quadratic time. We use (6), so we first characterize the effective local test ϕ_S^I before deriving an expression for $\mathbf{d}^I(S)$.

Lemma 8. *If $\phi_S, \emptyset \neq S \subseteq I$, is of the form (19), with $l_{i:m} \geq l_{i:n}$ for all $i \geq 1$ and $m, n \geq 0$, then*

$$\phi_S^I = \mathbb{1}\{\mathbf{p}_{(i:S)} \leq l_{i:\mathbf{h}_I} \text{ for at least one } i = 1, \dots, |S|\},$$

and

$$\mathbf{d}^I(S) = \max_{1 \leq u \leq |S|} 1 - u + |\{i \in S : \mathbf{p}_i \leq l_{u:\mathbf{h}_I}\}|, \quad (21)$$

where

$$\mathbf{h}_I = \max \{n \in \{0, \dots, |I|\} : \mathbf{p}_{(|I|-n+i:I)} > l_{i:n}, \text{ for } i = 1, \dots, n\}.$$

We note that \mathbf{h}_I can be calculated in $O(|I|^2)$ time, and, given \mathbf{h}_I , we can calculate $\mathbf{d}^I(S)$ in $O(|S|^2)$ time. In the case of the Simes test the calculation times of \mathbf{h}_I and $\mathbf{d}^I(S)$ are even linear after sorting the p -values, as shown by Meijer et al. (2019) and Goeman et al. (2019).

Our result is related to the results on FDP control based on bounding functions (Blanchard et al., 2017; Genovese and Wasserman, 2006). However, earlier authors have focused on proving type I error only, in which case it was sufficient to give a computable lower bound for $\mathbf{d}^I(S)$. From the perspective of admissibility, in view of Theorems 1 and 2, it is relevant that we compute $\mathbf{d}^I(S)$ exactly in Theorem 8.

By Theorem 1, the method (21) is everywhere at least as powerful as the interpolated method (17). In fact, it is a uniform improvement of that method as we shall see in the simulation experiment below. The next

question is whether the method defined by (21) with (20) is admissible, or whether it can be further improved. We can verify this using Theorem 2 by checking whether the local tests are admissible. It is immediately obvious that this is not the case. Taking e.g. $|S| = 1$, we see that at $\alpha = 0.05$ with $c \approx 2.163$ we have $\phi_S = \mathbb{1}\{\mathbf{p}_{(1:S)} \leq (1 - c)/c < 0\} = 0$, which is clearly not admissible. We may freely decrease c to $c_1 = 1/(1 + \alpha) \approx 0.952$ to obtain the uniformly more powerful local test $\phi_S = \mathbb{1}\{\mathbf{p}_{(1:S)} \leq \alpha\}$. We can use the same reasoning for $|S| = 2, 3, \dots$, decreasing the value of c to the minimal value that guarantees type I error control. This value may easily be calculated numerically since the worst case distribution of $(\mathbf{p}_i)_{i \in S}$ under H_S is the easy independent uniform case. We obtain a new local test of the form (19) with

$$l_{i:m} = \frac{i - c_m}{mc_m}. \quad (22)$$

We tabulated the minimal values of c_m (taking $\alpha = 0.05$) for some values of m in Table 1. Note that $c_m \leq c$ for all m , since $l_{i:m}$ is monotone in c_m and the local test with $c_m = c$ is valid by the results of Katsevich and Ramdas (2018), so the new local test uniformly improves the old one. We note that with these choices of c_m the critical values $l_{i:m}$ cannot be further increased without destroying type I error control of the local tests, so we conclude that the resulting local tests are admissible provided that the test $\mathbb{1}\{\mathbf{p}_i \leq \alpha\}$ is admissible as an α -level local test of H_i for all i and α . Assuming this, by Theorem 2 the resulting true discovery guarantee procedure is admissible. We note that, since c_m is increasing in m , (22) still fulfils the conditions of Lemma 8, so that the admissible method is still computable in quadratic time.

Table 1: Values of c_m calculated by Monte Carlo integration (10^6 samples)

m	1	2	3	4	5	7	10	15	20	50	100	500	1000
c_m	0.95	1.38	1.55	1.64	1.71	1.78	1.84	1.90	1.92	1.98	2.00	2.01	2.02

We have started with the procedure of Katsevich and Ramdas (2018) and improved it uniformly in three steps: the method was first improved by interpolation. The resulting coherent method was further improved by embedding it in a closed testing procedure, and finally that closed testing procedure was improved to an admissible method by improving its local tests. This way we obtained a sequence of four methods, each uniformly improving the previous one. We will call them the *original* (15), *coherent* (17), *closed*, defined by (21) with (20), and *admissible* method, defined by (21) with (22). We performed a small simulation experiment to assess the relative improvement made with each of the three steps. We used $m = 1000$ hypotheses, of which m_0 were true, and $m_1 = m - m_0$ false. We sampled p -values independently. For true null hypotheses, we used $\mathbf{p}_i \sim \mathcal{U}(0, 1)$. For false null hypotheses, we used $\mathbf{p}_i \sim \Phi^{-1}(-\gamma\mathbf{Z})$, where Φ is the standard normal distribution function, and $\mathbf{Z} \sim \mathcal{N}(0, 1)$. We took values $m_1 = 8, 40, 200$ and $\gamma = 2, 3, 4$ for $n_1 = 8, 40$ and $\gamma = 1, 2, 3$ for $m_1 = 200$. A true discovery guarantee procedure gives exponentially many output values. We report only results for sets \mathbf{K}_i of the i smallest p -values, as the original method did. We used $i = 5, 10, 20, 50, 200$.

We expect a linear time algorithm on the sorted p -values, similar to that of Meijer et al. (2019) and Goeman et al. (2019) for Simes local tests, to be possible for the closed and admissible procedures. However, for this example we were satisfied with the naive quadratic time algorithm, which for $m = 1000$ already calculated \mathbf{h}_I and $\mathbf{d}^I(S)$ for the closed and admissible methods in less than 0.1 seconds on a standard PC.

The results are given in Table 2. For each setting and each method we report the average value of $\mathbf{d}^I(\mathbf{K}_i)$ over 10^6 simulations. Several things can be noticed about these simulation results. The most important finding for the message of this section is that all three improvement steps are substantial. The improvement from the original to the coherent procedure is perhaps largest. It is especially noticeable for large rejected sets, where the original method may all too often give $\mathbf{d}^I(S) = 0$, especially if $|S| \gg m_1$. The improvement from the coherent procedure to closed testing is most apparent if m_1 is large. This is natural because the improvement can be seen as a “step-down” argument, implicitly incorporating an estimate of m_1 in the procedure. The final

improvement from the initial closed testing to the admissible procedure is clear throughout the table. Although the improvement from the coherent to the closed procedure seems the smallest one, we emphasize that closed testing is also crucial for the construction of the admissible procedure.

We also see some of the properties of the K&R method. We have that $\mathbf{d}^I(S) \leq |S| - 2$ with probability 1 for the original method (15), since $c > 2$ if $\alpha = 0.05$. This also holds for the coherent method (17). For the closed method the same is not true, but we have that if $\max_{1 \leq i \leq m} \mathbf{p}_i > l_{3:3} \approx 0.004$ we have $\phi_I(S) = 0$ for every S with $|S| = 2$, since $l_{1:3} < l_{2:3} < 0$. Therefore

$$\mathbb{P}(\mathbf{d}^I(S) \leq |S| - 2 \text{ for all } S \subseteq I) \geq \mathbb{P}(\max_{1 \leq i \leq m} \mathbf{p}_i > l_{3:3}) \approx 1,$$

unless all hypotheses are false. For the admissible method, by an analogous reasoning using Lemma 8, the same holds if $c_{|I_0|} > 2$, since then $c_{\mathbf{h}_I} > 2$ with large probability. This happens from $|I_0| \approx 100$. It follows that none of the methods in this section, not even the admissible method, should be expected make any FWER-rejections in practical applications. The admissible method is (almost) fully non-consonant in the sense that for all $i \in I$, $\phi_{\{i\}}^I \approx 0$, and we have $\mathbf{R}_\phi^I = \emptyset$ with probability almost 1 unless $|I_1| \approx |I|$. By Romano et al. (2011) the method is clearly inadmissible as a FWER-controlling method. By Theorem 2 it is admissible, however, as a method with true discovery guarantee: its lack of power for FWER-type statements is compensated by larger power for non-FWER-type statements. Indeed, Katsevich and Ramdas have shown that their method may significantly outperform Simes-based closed testing (Goeman et al., 2019) in some scenarios, which in turn outperforms consonant FWER-based testing in terms of FDP in large-scale testing problems.

10 Discussion

We have studied the class of all methods controlling tail probabilities of false discovery proportions. This class encompasses very diverse methods, e.g. familywise error control procedures, false discovery exceedance procedures, simultaneous selective inference, and cluster inference. We have shown that all such procedures can be written as methods simultaneously controlling false discovery proportions over all subsets of the family of hypotheses. This rewrite, trivial as it may be in some cases, is valuable in its own right, because makes it possible to study methods jointly that seemed incomparable before, and takes a step in reducing the ‘phletora of error rates’ lamented by Benjamini (2010). Moreover, methods that were constructed to give non-trivial error bounds for only a single random hypothesis set of interest, now give simultaneous error bounds for all such sets, allowing their use in flexible selective inference in the sense advocated by Goeman and Solarì (2011).

We formulated all such procedures in terms of true discovery guarantee, i.e. giving a lower bound to the number of true discoveries in each set, because this representation is mathematically easier to work with. Also, by emphasizing true rather than false discoveries, it gives a valuable positive frame to the multiple testing problem. Otherwise, this change in representation is purely cosmetic; we may continue to speak of FDP control procedures.

Admissibility is a very weak requirement for statistical tests, as under a weak assumption all tests that exhaust their α -level are admissible. It is not so easy to achieve for FDP control procedures. We have formulated a condition for admissibility of FDP control procedures that is both necessary and sufficient. All admissible FDP control procedures are closed testing procedures, and all closed testing procedures are admissible as FDP control procedures, provided they are well-designed in the sense that all their local tests are admissible. Apparently, control of false discovery proportions and closed testing procedures are so closely tied together that the relationship seems almost tautological. Admissibility is closely tied to optimality. Since optimal methods must be admissible, and admissible methods must be closed testing procedures, we have shown that only closed testing procedures can be optimal.

This theoretical insight has great practical value for methods designers. It can be used to uniformly improve existing methods, as we have demonstrated on the method of Katsevich and Ramdas (2018). Given a procedure

		$m_1 = 8$			$m_1 = 40$			$m_1 = 200$		
$\gamma =$		2	3	4	2	3	4	1	2	3
$\mathbf{d}^I(\mathbf{K}_5)$	<i>original</i>	0.1	1.3	2.9	1.9	3.0	3.0	1.1	3.0	3.0
	<i>coherent</i>	0.2	1.5	2.9	2.0	3.0	3.0	1.1	3.0	3.0
	<i>closed</i>	0.2	1.5	2.9	2.0	3.0	3.0	1.2	3.0	3.0
	<i>admissible</i>	0.3	1.7	2.9	2.2	3.0	3.0	1.4	3.0	3.0
$\mathbf{d}^I(\mathbf{K}_{10})$	<i>original</i>	0.1	0.9	2.8	1.9	8.0	8.0	2.5	8.0	8.0
	<i>coherent</i>	0.3	2.2	4.8	2.0	8.0	8.0	2.8	8.0	8.0
	<i>closed</i>	0.4	2.2	4.8	2.0	8.0	8.0	2.8	8.0	8.0
	<i>admissible</i>	0.5	2.5	5.0	2.2	8.0	8.0	3.3	8.0	8.0
$\mathbf{d}^I(\mathbf{K}_{20})$	<i>original</i>	0.0	0.1	0.3	3.7	15.9	18.0	3.8	17.2	18.0
	<i>coherent</i>	0.4	2.3	5.0	5.8	15.9	18.0	4.9	17.2	18.0
	<i>closed</i>	0.4	2.3	5.0	5.9	16.0	18.0	5.0	17.3	18.0
	<i>admissible</i>	0.5	2.6	5.3	6.5	16.3	18.0	5.8	17.6	18.0
$\mathbf{d}^I(\mathbf{K}_{50})$	<i>original</i>	0.0	0.0	0.0	0.3	10.2	22.3	2.9	41.1	48.0
	<i>coherent</i>	0.4	2.3	5.0	6.4	21.3	32.4	7.0	41.1	48.0
	<i>closed</i>	0.4	2.3	5.0	6.4	21.5	32.6	7.2	41.7	48.0
	<i>admissible</i>	0.5	2.7	5.3	7.3	22.4	33.1	8.7	42.3	48.0
$\mathbf{d}^I(\mathbf{K}_{200})$	<i>original</i>	0.0	0.0	0.0	0.0	0.0	0.0	0.0	40.6	130.1
	<i>coherent</i>	0.4	2.3	5.0	6.4	21.3	32.4	7.3	71.8	140.6
	<i>closed</i>	0.4	2.3	5.0	6.4	21.5	32.6	7.4	76.4	146.3
	<i>admissible</i>	0.5	2.7	5.3	7.3	22.5	33.2	9.3	81.2	149.1

Table 2: Averages of $\mathbf{d}^I(\mathbf{K}_i)$ for several values of i over 10^6 simulations, relating to the method of Katsevich and Ramdas (2018) and its successive uniform improvements.

that controls FDP, we first make sure it is coherent. Next, we can explicitly construct the local tests implied by the procedure, and turn it closed testing procedure. To check admissibility, we now only need to check admissibility of the local tests. Each step may result in substantive improvement, as we have shown in simulations. Alternatively, when designing a method we may start for a suite of local tests that has good power properties. The options are virtually unlimited here. The validity of the local test as an α -level test guarantees control of FDP. Correlations between test statistics, that often complicate multiple testing procedures, are taken into account by the local test. Admissibility of the local tests guarantees admissibility of the resulting procedure. In both cases the computational problem remains that closed testing may require exponentially many tests, but this is the only remaining problem. Polynomial time shortcuts are possible. Ideally these are exact, as for K&R above, and admissibility is retained. If the full closed testing procedure is not computable for large testing problems, we may settle for an inadmissible but computable method, based on a conservative shortcut (e.g. Hemerik and Goeman, 2018; Hemerik et al., 2019). It may still be worthwhile to compare such a method to full closed testing in small-scale problems to see how much power is lost.

We defined admissibility in terms of simultaneous FDP control for all possible subsets of the family of hypotheses. In some cases we may not be interested in all of these sets, as e.g. when targeting FWER control exclusively. Even with FWER, we retain the result that admissible procedures must be closed testing procedures with admissible local tests. We lose, however, the property that all such procedures are automatically admissible.

With other loss functions, additional criteria might come in, as consonance in the case of familywise error control. Variants of consonance may be useful as well (Brannath and Bretz, 2010).

Our focus was mostly on monotone procedures. Such procedures are defined for multiple testing problems on different scales simultaneously. Connecting between different scales, they have the property that adding more hypotheses to the multiple testing problem will never result in stronger conclusions for the hypotheses that were already there. This is an intuitively desirable property by itself, which prevents some paradoxes (Goeman and Solari, 2014). Monotone procedures have additional valuable properties: viewed as closed testing procedures, they have local tests that are truly local: the local test on S uses only the information that the corresponding local procedure \mathbf{d}^S uses. Admissible monotone procedures, however, may sometimes be locally improved, and we have given an example of this. Such improvements, if admissible, must still be closed testing procedures with admissible local tests themselves.

We have restricted to finite testing problems. Extensions to countably infinite problems are of interest e.g. when considering online control (Javanmard et al., 2018). The results of this paper may trivially be extended to allow infinite $|I|$ if we are willing to assume that $|C_1| < \infty$, so that $\mathbf{d}^I < \infty$. If $|C_1|$ is unbounded, care must be taken to scale \mathbf{d} properly to keep it in the non-trivial range. This scaling adds some technical complexity, and is not assumption-free because $\mathbf{d}^I(S)$ scales with the unknown $|S_1|$. However, since most of the results of this paper compare competing methods, which obviously require the same scaling, we conjecture that the optimality of closed testing will translate to FDP control in countable and even uncountable multiple testing problems. We leave this to future research.

Finally, we remark that we have only considered procedures that control tail probabilities of the false discovery proportion. These methods can also be used for bounding the median FDP (Goeman and Solari, 2011). However, if there is interest in the central tendency of FDP it is more common to bound the mean FDP, better known as False Discovery Rate (FDR). Given the close connection we have established between closed testing and FDP tail probabilities, it is likely that there is also a connection between closed testing and FDR control. Some connections have already been found between Simes-based closed testing and the procedure of Benjamini and Hochberg (1995) by Goeman et al. (2019). It is likely that there are more such connections. Any procedure that controls FDR, since FDR control implies weak FWER control, implies a local test and can therefore be used to construct a closed testing procedure. Conversely, if FDP is controlled with $1 - \alpha$ -confidence at level γ , then FDR is controlled at $\alpha(1 - \gamma) + \gamma$, as Lehmann and Romano (2005) have shown. More profound relationships may be found in the future.

A Existence of admissible procedures

Admissibility as defined in Section 5 is known in the literature as α -admissibility on Θ . Alternative definitions of admissibility exist (Lehmann and Romano, 2006, Section 6.7). With α -admissibility on $\Theta_1 = \Theta \setminus H$, the value of θ under which $\tilde{\delta}$ improves δ with positive probability must be in Θ_1 . With d -admissibility, there is an additional requirement that $\tilde{\delta}$ may not improve δ with positive probability for $\theta \in H$. However, α -admissibility on Θ is most commonly considered in the multiple testing context (e.g. Lehmann and Romano, 2006, Section 9.3), because with multiple hypotheses there is no unique Θ_1 .

Admissible tests may not always exist. Consider for example the model where $\mathbf{X} \sim \mathcal{N}(\mu, \sigma^2)$, with parameter space $\Theta = (-\infty, \infty) \times [0, 1]$ for $\theta = (\mu, \sigma^2)$. Let δ be any valid test for $H : \mu = 0$, for example $\delta = \mathbb{1}\{\mathbf{X} > z_{1-\alpha}\}$, where $z_{1-\alpha}$ is the $1 - \alpha$ -quantile of the standard normal distribution. Then for any c outside the rejection region of δ ,

$$\tilde{\delta} = \delta + \mathbb{1}\{\mathbf{X} = c\}$$

improves δ with positive probability for $\mu = c$ and $\sigma^2 = 0$. Since any valid test may be improved in a similar way unless $\alpha = 1$, an admissible test does not exist.

However, we can easily guarantee existence of admissible tests if we rule out degenerate models. A weak assumption for this is the following.

Assumption 1 (common null events). *For every measurable event E and for every $\theta_1, \theta_2 \in \Theta$ we have that $P_{\theta_1}(E) = 0$ implies that $P_{\theta_2}(E) = 0$.*

Under Assumption 1 the collection of null events is common to all parameter values; there are no events that happen with positive probability for some parameter values, but with probability zero for others. This is a weak assumption that holds for most models regularly used in applied statistics, both continuous and discrete, when we are willing to exclude deterministic corner cases from the parameter space. In the example above, Assumption 1 holds if we simply restrict the parameter space to $\Theta = (-\infty, \infty) \times (0, 1]$, excluding $\sigma^2 = 0$.

If we accept Assumption 1, Lemma 9 presents a very simple sufficient condition for admissibility: every statistical test of H that fully exhausts its α -level for some $\theta \in H$ is admissible.

Lemma 9. *If Assumption 1 holds, then a statistical test δ of hypothesis $H \subseteq \Theta$ is admissible if $\theta \in H$ exists such that $P_{\theta}(\delta = 1) = \alpha$.*

B A local improvement

In this section we construct a local improvement of an admissible monotone procedure to illustrate Proposition 1. Assume that for each $H_i, i \in C$ we have a p -value \mathbf{p}_i . Assume that each \mathbf{p}_i is standard uniform if H_i is true. Under these assumptions we can define the standard fixed sequence testing procedure, which starts testing H_1 using \mathbf{p}_1 at level α , continues one by one with H_2, H_3 in order, and stops when it fails to reject some hypothesis. It is well known that this procedure is a closed testing procedure. The local test is defined for all $S \subseteq C$ by

$$\phi_S = \mathbb{1}\{\mathbf{p}_{\min(S)} \leq \alpha\}.$$

If we assume that the test $\phi_{\{i\}} = \mathbb{1}\{\mathbf{p}_i \leq \alpha\}$ is admissible for all i , then, by Theorem 2, the fixed sequence procedure is admissible.

We will now make some additional assumptions that will allow a uniform improvement of the procedure at the fixed scale $I \subseteq C$. Assume for convenience that all p -values are independent. Next, assume that the distribution of every \mathbf{p}_i is constrained even under the alternative. Assume that some $P(x)$ exists such that, for every $\theta \in \Theta$,

$$P_{\theta}(\mathbf{p}_i \leq x) \leq P(x).$$

This means that the power for each test is inherently limited. Even under the alternative, we reject e.g. H_1 with probability at most $P(\alpha)$. Clearly $P(x) \geq x$ or we would not have uniformity under the null. We will now demonstrate that the fixed sequence procedure can be uniformly improved, locally at any I , if $P(\alpha) < 1$.

We start with the simple case $I = \{1, 2\}$. By Proposition 1, the improvement is a closed testing procedure that involves local tests $\psi_S \geq \phi_S^I$. Consider $S = \{2\}$. Then $\phi_S^I = \mathbb{1}\{\mathbf{p}_1 \leq \alpha, \mathbf{p}_2 \leq \alpha\}$. Under H_2 this has $P_{\theta}(\phi_S^I = 1) \leq \alpha P(\alpha) < \alpha$. Clearly, there is room for improvement. Let us consider the procedure with $\psi_S = \phi_S$, for all $S \subseteq C$ except $S = \{2\}$, when

$$\psi_S = \mathbb{1}\{\mathbf{p}_1 \leq \alpha, \mathbf{p}_2 \leq \alpha/P(\alpha)\}.$$

The latter is a valid local test of H_2 . The resulting procedure at $I = \{1, 2\}$ starts testing H_1 at level α , and continues, if H_1 is rejected, to test H_2 at level $\alpha/P(\alpha) > \alpha$. This is clearly a uniform improvement of the original procedure at $I = \{1, 2\}$. To see that this is not a counterexample to Theorem 2, consider $I = \{2\}$ instead. Clearly, we do not have

$$\mathbb{1}\{\mathbf{p}_1 \leq \alpha, \mathbf{p}_2 \leq \alpha/P(\alpha)\} = \mathbf{d}_{\psi}^{\{2\}}(\{2\}) \geq \mathbf{d}_{\phi}^{\{2\}}(\{2\}) = \mathbb{1}\{\mathbf{p}_2 \leq \alpha\}.$$

The local uniform improvement at $I = \{1, 2\}$ comes at the cost of a potential deterioration at $I = \{2\}$.

Similar local improvements actually exist for every finite $I \subseteq C$ except $I = \{1\}$. Define recursively

$$\alpha_1 = \alpha; \quad \alpha_{i+1} = \alpha_i / P(\alpha_i).$$

From this, fix some $I \subseteq C$, and define a local test as

$$\psi_S = \mathbb{1}\{\mathbf{p}_i \leq \alpha_i \text{ for all } i \in L_S\},$$

where $L_S = \{i \in I: i \leq \min(S)\}$. To check that this is a valid local test, we verify that for all $\theta \in H_S$

$$P_\theta(\psi_S = 1) \leq \alpha_l \prod_{i=1}^{l-1} P(\alpha_i) = \alpha,$$

where $l = |L_S|$. The resulting procedure is still a fixed sequence procedure that tests all $H_i, i \in I$ in order, stopping the first time it fails to reject. Only, rather than testing at level α every time, it tests at level α_i in step i . If $x < P(x) < 1$ for all $0 < x < 1$ the sequence $\alpha_1, \alpha_2, \dots$ is strictly increasing and approaches 1.

Crucial for this example is the assumption that we have limited power and, more importantly, that we know the limit to the power. If we are not willing to assume that $P(\alpha) < 1$, or if we do not know P , then the above local improvements are not possible. It is difficult to think of uniform local improvements in the case $P(x) = 1$, and we believe they do not exist. It may be worthwhile to think of adaptive procedures that learn P as the procedure moves along, but we will not pursue this direction here. In any case, due to the cost inherent to learning P , such a procedure would not uniformly improve \mathbf{d}_ϕ .

C Proofs

Proof of Lemma 1

Take any $S \in 2^I$. For any V with $\phi_V = 0$ there exists a $U = S \cap V \subseteq S$ which has $\phi_U^I \leq \phi_V = 0$ and $|S \setminus U| = |S \setminus V|$. Consequently, $\mathbf{d}_\phi^I(S) \leq \mathbf{g}_\phi^I(S)$.

For any $U \subseteq S$ with $\phi_U^I = 0$ there is a $V \supseteq U$ with $\phi_V = 0$ that has $|S \setminus V| \leq |S \setminus U|$. Consequently, $\mathbf{g}_\phi^I(S) \leq \mathbf{d}_\phi^I(S)$. \square

Proof of Lemma 2

Let E be the event that $\mathbf{d}^I(S) \leq |S_1|$ for all $S \subseteq I$. Suppose that E happened and choose any $S \subseteq I$ and $U \subseteq I$. Then

$$\mathbf{d}^I(U) - |U \setminus S| + \mathbf{d}^I(S \setminus U) \leq |U_1| - |U \setminus S| + |(S \setminus U)_1| \leq |S_1|.$$

Consequently, if E happened, $\bar{\mathbf{d}}^I(S) \leq |S_1|$ for all $S \subseteq I$. Since $P(E) \geq 1 - \alpha$, we have true discovery guarantee for $\bar{\mathbf{d}}^I$. \square

Proof of Lemma 3

Suppose \mathbf{d}^I is coherent and let V, W be disjoint. Then, taking $S = V \cup W$ and $U = V$ in (10), we have

$$\mathbf{d}^I(V \cup W) \geq \mathbf{d}^I(V) - |V \setminus S| + \mathbf{d}^I(S \setminus V) = \mathbf{d}^I(V) + \mathbf{d}^I(W).$$

Also, taking $S = V$ and $U = V \cup W$ in (10), we obtain $\mathbf{d}^I(V) \geq \mathbf{d}^I(V \cup W) - |W|$.

Next, suppose (11) holds for all disjoint V, W . For every $S \subseteq I$, we have, by (8) and the left-hand inequality of (11),

$$\bar{\mathbf{d}}^I(S) \leq \max_{U \subseteq 2^I} \left\{ \mathbf{d}^I(U \cup S) - |U \setminus S| \right\}.$$

Since $\mathbf{d}^I(U \cup S) \leq \mathbf{d}^I(S) + |U \setminus S|$ by the right-hand inequality of (11), we get $\bar{\mathbf{d}}^I(S) \leq \mathbf{d}^I(S)$. Since also $\bar{\mathbf{d}}^I(S) \geq \mathbf{d}^I(S)$, by taking $U = S$ in (8), we have $\bar{\mathbf{d}}^I(S) = \mathbf{d}^I(S)$ for all $S \subseteq I$, so \mathbf{d}^I is coherent. \square

Proof of Lemma 4

We use Lemma 3. Let $V, W \subseteq I$ be disjoint. Then some $U \subseteq V \cup W$ exists such that $\phi_U^I = 0$ and

$$|V \setminus U| + |W \setminus U| = |(V \cup W) \setminus U| = \mathbf{d}_\phi^I(V \cup W).$$

Since $\phi_{U \cap V}^I \leq \phi_U^I = 0$, we have $\mathbf{d}_\phi^I(V) \leq |V \setminus U|$. Similarly, $\mathbf{d}_\phi^I(W) \leq |W \setminus U|$, so we have $\mathbf{d}_\phi^I(V) + \mathbf{d}_\phi^I(W) \leq \mathbf{d}_\phi^I(V \cup W)$.

Also, there exists $T \subseteq V \subseteq (V \cup W)$ such that $\phi_T^I = 0$ and $|V \setminus T| = \mathbf{d}_\phi^I(V)$. Now

$$|(V \cup W) \setminus T| = |V \setminus T| + |W| = \mathbf{d}_\phi^I(V) + |W|,$$

so we have $\mathbf{d}_\phi^I(V \cup W) \leq \mathbf{d}_\phi^I(V) + |W|$. \square

Proof of Lemma 5

Genovese and Wasserman (2006) already proved that \mathbf{d}_ϕ^I has true discovery guarantee for all finite $I \subseteq C$, and we have coherence by Lemma 4, so we only need to prove monotonicity. This is trivial from (7) and Lemma 1. Take any finite $S \subseteq I \subseteq J \subseteq C$. Then

$$\mathbf{d}_\phi^I(S) = \min_{V \subseteq 2^I} \{ |S \setminus V| : \phi_V = 0 \} \geq \min_{V \subseteq 2^J} \{ |S \setminus V| : \phi_V = 0 \} = \mathbf{d}_\phi^J(S).$$

\square

Proof of Theorem 1

Take any finite $S \subseteq C$. Since \mathbf{d}^S has discovery control on S , we have

$$\max_{\theta \in H_S} \mathbb{P}_\theta(\phi_S = 1) = \max_{\theta \in H_S} \mathbb{P}_\theta(\mathbf{d}^S(S) > 0) = \max_{\theta \in H_S} \mathbb{P}_\theta(\mathbf{d}^S(S) > |S_1|) \leq \alpha,$$

so ϕ_S is a valid test of H_S . This proves the first statement.

Take any finite $S \subseteq C$ again. We have $\phi_S^I = 1$ if and only if $\mathbf{d}^W(W) > 0$ for all $S \subseteq W \subseteq I$. For all such W we have, by coherence of \mathbf{d}^W and monotonicity of \mathbf{d} ,

$$\mathbf{d}^W(W) \geq \mathbf{d}^W(S) \geq \mathbf{d}^I(S).$$

Consequently, $\phi_S^I \geq \mathbf{1}\{\mathbf{d}^I(S) > 0\}$. We obtain

$$\mathbf{d}_\phi^I(S) \geq \min_{U \subseteq 2^S} \{ |S \setminus U| : \mathbf{d}^I(U) = 0 \}. \quad (23)$$

By coherence of \mathbf{d}^I and Lemma 3, we have, for all $U \subseteq S \subseteq I$, that

$$|S \setminus U| \geq \mathbf{d}^I(S) - \mathbf{d}^I(U).$$

Combining this with (23) the second statement of the theorem follows. \square

Proof of Theorem 2

We prove the two counterpoints. Let \mathbf{d}_ϕ be inadmissible, and let $\tilde{\mathbf{d}}$ be a monotone procedure that uniformly improves it. By Theorem 1 there exists $\mathbf{d}_\psi \geq \tilde{\mathbf{d}}$ that also uniformly improves \mathbf{d}_ϕ . We have, for every finite $S \subseteq C$,

$$\psi_S = \mathbb{1}\{\mathbf{d}_\psi^S(S) > 0\} \geq \mathbb{1}\{\mathbf{d}_\phi^S(S) > 0\} = \phi_S,$$

Also, by Theorem 1 ψ_S is a valid local test for H_S .

Let $S \subseteq I \subseteq C$, $|I| < \infty$, and $\theta \in \Theta$ be such that $P_\theta(E) > 0$ for $E = \{\mathbf{d}_\psi^I(S) > \mathbf{d}_\phi^I(S)\}$. If E happened, by (7) there is a $U \subseteq I$ with $|S \setminus U| = \mathbf{d}_\phi^I(S)$, such that $\psi_U = 1$ and $\phi_U = 0$. Consequently,

$$P(\phi_U < \psi_U) \geq P(E) > 0,$$

so ϕ_U is inadmissible. Since $|S \setminus U| = \mathbf{d}_\phi^I(S) < \mathbf{d}_\psi^I(S) \leq |S|$, we have $U \neq \emptyset$.

Conversely, let ϕ_S be inadmissible for some finite $\emptyset \subset S \subseteq C$, and let $\phi'_S \geq \phi_S$ be a test that uniformly improves it, so that $E = \{\phi'_S > \phi_S\}$ has $P_\theta(E) > 0$ for some $\theta \in \Theta$. Define the suite ψ such that $\psi_S = \phi'_S$ and $\psi_I = \phi_I$ for $I \neq S$, and consider the monotone procedure \mathbf{d}_ψ . Since $\psi \geq \phi$ we also have $\mathbf{d}_\psi \geq \mathbf{d}_\phi$. Since E implies $\mathbf{d}_\psi^S(S) > \mathbf{d}_\phi^S(S)$, we see that \mathbf{d}_ψ uniformly improves \mathbf{d}_ϕ , so the latter is inadmissible. \square

Proof of Proposition 1

Using (12) we may embed \mathbf{d}^I in a monotone procedure. By Theorem 1 we have $\mathbf{d}^I \leq \mathbf{d}^I_{\chi}$, with, for all $S \subseteq I$, $\chi_S = \mathbb{1}\{\mathbf{d}^I(S) > 0\} \geq \mathbb{1}\{\mathbf{d}_\phi^I(S) > 0\} = \phi_S^I$. For every $S \subseteq I$, if χ_S is not admissible, let $\psi_S \geq \chi_S \geq \phi_S^I$ be a uniform improvement; otherwise, let $\psi_S = \chi_S$. Without loss of generality we may assume that ψ_S is admissible. Then $\mathbf{d}^I_\psi \geq \mathbf{d}^I$. Since \mathbf{d}^I is admissible, we have $\mathbf{d}^I_\psi = \mathbf{d}^I$. \square

Proof of Lemma 6

Choose any $V, W \subseteq I$ disjoint. Call $S = V \cup W$. We use complete induction on $|S|$. Suppose that (13) holds for all sets smaller than S . If $V = \emptyset$ or $W = \emptyset$ the result is trivial, so we assume and $V, W \neq \emptyset$. If $\mathbf{d}^I(S) = 0$ the result follows immediately from Lemma 3, so we may assume $\mathbf{d}^I(S) > 0$. By consonance there is an $i \in S$ such that $\mathbf{d}^I(\{i\}) = 1$. Without loss of generality, suppose that $i \in W$. By Lemma 3 and the induction hypothesis we have

$$\mathbf{d}^I(S) = \mathbf{d}^I(V \cup W \setminus \{i\}) + \mathbf{d}^I(\{i\}) = \mathbf{d}^I(V) + \mathbf{d}^I(W \setminus \{i\}) + \mathbf{d}^I(\{i\}).$$

Since $V \neq \emptyset$, $W \subset S$ and we may use the induction hypothesis once more, saying that $\mathbf{d}^I(W \setminus \{i\}) + \mathbf{d}^I(\{i\}) = \mathbf{d}^I(W)$ to obtain (13).

For the converse, suppose that (13) holds. By Lemma 3, \mathbf{d}^I is coherent. Choose some $\emptyset \neq S \subseteq I$ such that $\mathbf{d}^I(S) > 0$. We use complete induction on $|S|$ to show that $\mathbf{d}^I(\{i\}) = 1$ for some $i \in S$. Suppose the result holds for all sets smaller than S . Choose any $i \in S$ and let $V = \{i\}$ and $W = S \setminus \{i\}$. By (13), either $\mathbf{d}^I(V) > 0$ or $\mathbf{d}^I(W) > 0$. If the former, we have the result immediately. If the latter, we have use the induction hypothesis to conclude that $\mathbf{d}^I(\{i\}) = 1$ for some $j \in W \subset S$. \square

Proof of Lemma 7

Choose $S \subseteq I$ en let $\mathbf{V} = S \cap \mathbf{R}_\phi^I$ and $\mathbf{W} = S \setminus \mathbf{R}_\phi^I$. By definition of \mathbf{R}_ϕ^I , for every $i \in \mathbf{V}$ we have $\phi_{\{i\}}^I = 1$, so $\phi_J^I = 1$ for all $\emptyset \neq J \subseteq \mathbf{V}$, so $\mathbf{d}_\phi^I(\mathbf{V}) = |\mathbf{V}|$ by (6).

If \mathbf{d}_ϕ^I is consonant, by definition of \mathbf{R}_ϕ^I , for every $i \in \mathbf{W}$ we have $\phi_{\{i\}}^I = 0$. By consonance, we must have $\phi_{\mathbf{W}}^I = 0$, so $\mathbf{d}_\phi^I(\mathbf{W}) = 0$ by (6). By Lemma 6 we have $\mathbf{d}_\phi^I(S) = |\mathbf{V}|$, which proves the first statement.

If \mathbf{d}_ϕ^I is not consonant, by Lemma 3 we have $\mathbf{d}_\phi^I(S) \geq \mathbf{d}_\phi^I(\mathbf{V}) = \mathbf{r}_\phi^I(S)$. Moreover, $S \subseteq I$ exists such that for some $\theta \in \Theta$ we have with positive probability that $\mathbf{d}_\phi^I(S) > 0$ and $S \cap \mathbf{R}_\phi^I = \emptyset$, or \mathbf{d}_ϕ^I would be consonant. For this S with positive probability $\mathbf{d}_\phi^I(S) > 0 = \mathbf{r}_\phi^I(S)$. \square

Proof of Proposition 2

For all finite $S \subseteq C$ consider $\phi_S = \mathbb{1}\{\mathbf{R}^S \neq \emptyset\}$. Then for all finite $S \subseteq I \subseteq C$

$$\phi_S^I = \mathbb{1}\{\mathbf{R}^J \neq \emptyset \text{ for all } S \subseteq J \subseteq I\}. \quad (24)$$

We show that ϕ is consonant, i.e. if $\phi_S^I = 1$, there is an $i \in \mathbf{R}_S \neq \emptyset$ such that $\phi_{\{i\}}^I = 1$. We proceed by complete induction on $|S|$ downward from $|I|$. Assume that for all $V \subseteq I$ with $|V| > |S|$, it holds that $\phi_V^I = 1$ implies that an $i \in \mathbf{R}_V \neq \emptyset$ exists such that $\phi_{\{i\}}^I = 1$, but the same does not hold for $V = S$. We will derive a contradiction. Since $\phi_S^I = 1$, indeed, by (24), $\mathbf{R}^S \neq \emptyset$. Choose $i \in \mathbf{R}^S$ and $J \ni i$. If $J \subseteq S$, we have $i \in \mathbf{R}_S \cap J \subseteq \mathbf{R}^J$ by monotonicity, so $\mathbf{R}^J \neq \emptyset$. If not $J \subseteq S$, we have $\phi_{J \cup S}^I = 1$, since $\phi_S^I = 1$ and $S \subset J \cup S$. By the induction hypothesis there exists a $j \in \mathbf{R}^{J \cup S}$ such that $\phi_{\{j\}}^I = 1$. Now either $j \in S$ or $j \in J \setminus S$. In the former case $j \in S \cap \mathbf{R}^{J \cup S} \subseteq \mathbf{R}^S$ by monotonicity, and we have a contradiction. Therefore, we must have $j \in J \setminus S$, in which case $j \in \mathbf{R}^{J \cup S} \cap J \subseteq \mathbf{R}^J$, also by monotonicity, so $\mathbf{R}^J \neq \emptyset$. Since $\mathbf{R}^J \neq \emptyset$ for all $J \ni i$, we have $\phi_{\{i\}}^I = 1$ by (24). This proves consonance.

Since \mathbf{R} is monotone, we have

$$\phi_S^I \geq \mathbb{1}\{\mathbf{R}^I \cap J \neq \emptyset \text{ for all } S \subseteq J \subseteq I\} = \mathbb{1}\{\mathbf{R}^I \cap S \neq \emptyset\}.$$

Clearly, by (14), $\mathbf{R}_\phi^I \geq \{i \in I : \mathbf{R}^I \cap \{i\} \neq \emptyset\} = \mathbf{R}^I$ for all finite $I \subseteq C$. Since \mathbf{R} is admissible we must have $\mathbf{R}_\phi^I = \mathbf{R}^I$. \square

Proof of Proposition 3

Let \mathbf{R} be admissible. By Proposition 2, $\mathbf{R} = \mathbf{R}_\phi$, with ϕ consonant. If ϕ is admissible, we are done. Otherwise, let ψ uniformly improve ϕ . Without loss of generality, we may assume that ψ is admissible. Then $\mathbf{R}_\psi \geq \mathbf{R}_\phi$, and we must have equality since \mathbf{R} is admissible. \square

Proof of Proposition 4

Let ϕ be consonant. By Lemma 7, we have $\mathbf{d}_\phi = \mathbf{r}_\phi$. Suppose \mathbf{R} uniformly improves \mathbf{R}_ϕ . Then $\mathbf{r} = (\mathbf{r}^I)_{I \subseteq C, |I| < \infty}$ with $\mathbf{r}^I(S) = |S \cap \mathbf{R}^I|$ uniformly improves $\mathbf{r}_\phi = \mathbf{d}_\phi$, so ϕ is inadmissible by Theorem 2. \square

Proof of Lemma 8

Consider first the case $|S| > \mathbf{h}_I$. By definition of \mathbf{h}_I , there is an $1 \leq i \leq |S|$ such that

$$\mathbf{P}(i:S) \leq \mathbf{P}(|I| - |S| + i:I) \leq l_{i:|S|}.$$

Consequently, $\phi_S = \phi_S^I = 1$ for all S with $|S| > \mathbf{h}_I$. Since for such S also $l_{i:|S|} \leq l_{i:\mathbf{h}_I}$, the result of the lemma holds if $|S| > \mathbf{h}_I$.

Now consider the case $|S| \leq \mathbf{h}_I$. First, suppose that there exists some $1 \leq i \leq |S|$ with $\mathbf{P}(i:S) \leq l_{i:\mathbf{h}_I}$. Take any $V \supseteq S$. If $|V| > \mathbf{h}_I$, we have $\phi_V = 1$ as proved above. If $|V| \leq \mathbf{h}_I$, we have

$$\mathbf{P}(i:V) \leq \mathbf{P}(i:S) \leq l_{i:\mathbf{h}_I} \leq l_{i:|V|}$$

so that $\phi_V = 1$. Since this holds for all $V \supseteq S$, we have $\phi_S^I = 1$.

Next, suppose that there is no $1 \leq i \leq |S|$ with $p_{(i:S)} \leq l_{i:\mathbf{h}_I}$. For $j = 1, \dots, |I|$, define \mathbf{Z}_j as a set with $|\mathbf{Z}_j| = j$ such that for all $u \in \mathbf{Z}_j$ and $v \in I \setminus \mathbf{Z}_j$ we have $\mathbf{p}_u \geq \mathbf{p}_v$. Let $\mathbf{W} = S \cup \mathbf{Z}_j$ for some $0 \leq j \leq \mathbf{h}_I$ such that $|\mathbf{W}| = \mathbf{h}_I$. If $1 \leq i \leq \mathbf{h}_I - j$, by the assumption we have

$$\mathbf{P}_{(i:\mathbf{W})} = p_{(i:S)} > l_{i:\mathbf{h}_I} = l_{i:|\mathbf{W}|}.$$

If $\mathbf{h}_I - j \leq i \leq \mathbf{h}_I$, since $\mathbf{Z}_j \subseteq \mathbf{Z}_{\mathbf{h}_I}$ we have

$$p_{(i:\mathbf{W})} = p_{(i:\mathbf{Z}_{\mathbf{h}_I})} > l_{i:\mathbf{h}_I} = l_{i:|\mathbf{W}|}$$

because $\phi_{\mathbf{Z}_{\mathbf{h}_I}} = 0$ by definition of \mathbf{h}_I . Taken together, this implies that $\phi_{\mathbf{W}} = 0$, so $\phi_{\mathbf{W}}^I = 0$. This proves the statement about ϕ_S^I .

To prove the statement about $\mathbf{d}^I(S)$ we will use

$$\mathbf{d}^I(S) = \min_{U \in 2^S} \{|S \setminus U| : \phi_U^I = 0\} = \min_{U \in 2^S} \{|I| : \phi_{S \setminus U}^I = 0\}.$$

As shown above we have $\phi_{S \setminus U}^I = 1$ if and only if for some $1 \leq u \leq |S \setminus U|$ we have $|\{i \in S \setminus U : \mathbf{p}_i \leq l_{u:\mathbf{h}_I}\}| \geq u$, and we may trivially extend the range to $1 \leq u \leq |S|$. Thus, $\phi_{S \setminus U}^I = 0$ if and only if for all such u we have $|\{i \in S \setminus U : \mathbf{p}_i \leq l_{u:\mathbf{h}_I}\}| \leq u - 1$. That is, for all such u ,

$$|\{i \in S : \mathbf{p}_i \leq l_{u:\mathbf{h}_I}\}| - u + 1 \leq |\{i \in U : \mathbf{p}_i \leq l_{u:\mathbf{h}_I}\}|. \quad (25)$$

Denote the left-hand side of (25) by $\mathbf{g}(u)$ and the right-hand side by $\mathbf{f}(U, u)$. Let $\mathbf{d} = \max_{1 \leq u \leq |S|} \mathbf{g}(u)$. Since $0 \leq \mathbf{d} \leq |S|$ we can pick $\mathbf{U} \subset S$ with $|\mathbf{U}| = \mathbf{d}$ such that for all $i \in \mathbf{U}$ and $j \in S \setminus \mathbf{U}$, we have $\mathbf{p}_i \leq \mathbf{p}_j$. For this \mathbf{U} , for all $1 \leq u \leq |S|$, we have $\mathbf{f}(\mathbf{U}, u) \leq |\mathbf{U}| = \mathbf{d}$. If $\mathbf{f}(\mathbf{U}, u) < \mathbf{d} = |\mathbf{U}|$, then $\mathbf{g}(u) \leq \mathbf{f}(S, u) = \mathbf{f}(\mathbf{U}, u)$, where the latter step follows by construction of \mathbf{U} . If $\mathbf{f}(\mathbf{U}, u) = \mathbf{d}$, then $\mathbf{g}(u) \leq \mathbf{d} = \mathbf{f}(\mathbf{U}, u)$. We conclude that \mathbf{U} satisfies (25) for all $1 \leq u \leq |S|$. Obviously, (25) cannot hold for any U with $|U| < \mathbf{d}$. We conclude that $\mathbf{d}^I(S) = \mathbf{d}$. \square

Proof of Lemma 9

Suppose that that $\theta \in H$ exists such that $\mathbf{P}_\theta(\boldsymbol{\delta} = 1) = \alpha$, and that $\boldsymbol{\delta}'$ is a test of H that uniformly improves $\boldsymbol{\delta}$. We will derive a contradiction under Assumption 1. Because $\boldsymbol{\delta}'$ is a uniform improvement, some $\theta' \in \Theta$ exists such that

$$\mathbf{P}_{\theta'}(\boldsymbol{\delta}' > \boldsymbol{\delta}) > 0. \quad (26)$$

By Assumption 1, (26) remains valid if we replace θ' by θ . Consequently, since $\boldsymbol{\delta}' \geq \boldsymbol{\delta}$, and since $\{\boldsymbol{\delta}' > \boldsymbol{\delta}\}$ and $\{\boldsymbol{\delta} = 1\}$ are disjoint, we have

$$\mathbf{P}_\theta(\boldsymbol{\delta}' = 1) = \mathbf{P}_\theta(\boldsymbol{\delta}' > \boldsymbol{\delta}) + \mathbf{P}_\theta(\boldsymbol{\delta} = 1) > \alpha$$

which contradicts that $\boldsymbol{\delta}'$ is a valid test of H . \square

Acknowledgements

This paper was motivated by many inspiring discussions during the workshop ‘‘Post-selection Inference and Multiple Testing’’ in Toulouse, Februari 2018. We thank all participants, the organizers Pierre Neuvial, Etienne Roquain and Gilles Blanchard, and especially Ruth Heller for asking the question that triggered this research project. We thank Jonathan Rosenblatt and the Israel Science foundation for financing the computing equipment used for this research (grants 926/14 and 900/16). Jelle Goeman was supported by NWO VIDI grant 639.072.412.

References

- Benjamini, Y. (2010). Simultaneous and selective inference: Current successes and future challenges. *Biometrical Journal* 52(6), 708–721.
- Benjamini, Y. and R. Heller (2008). Screening for partial conjunction hypotheses. *Biometrics* 64(4), 1215–1222.
- Benjamini, Y. and Y. Hochberg (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)* 57(1), 289–300.
- Benjamini, Y. and W. Liu (1999). A step-down multiple hypotheses testing procedure that controls the false discovery rate under independence. *Journal of Statistical Planning and Inference* 82(1-2), 163–170.
- Berk, R., L. Brown, A. Buja, K. Zhang, L. Zhao, et al. (2013). Valid post-selection inference. *The Annals of Statistics* 41(2), 802–837.
- Blanchard, G., P. Neuvial, and E. Roquain (2017). Post hoc inference via joint family-wise error rate control. *arXiv preprint arXiv:1703.02307*.
- Blanchard, G. and É. Roquain (2009). Adaptive false discovery rate control under independence and dependence. *Journal of Machine Learning Research* 10, 2837–2871.
- Brannath, W. and F. Bretz (2010). Shortcuts for locally consonant closed test procedures. *Journal of the American Statistical Association* 105(490), 660–669.
- Bretz, F., W. Maurer, W. Brannath, and M. Posch (2009). A graphical approach to sequentially rejective multiple test procedures. *Statistics in medicine* 28(4), 586–604.
- Cai, G. and S. K. Sarkar (2008). Modified Simes’ critical values under independence. *Statistics & Probability Letters* 78(12), 1362–1368.
- Delattre, S., E. Roquain, et al. (2015). New procedures controlling the false discovery proportion via Romano-Wolf’s heuristic. *The Annals of Statistics* 43(3), 1141–1177.
- Donoho, D., J. Jin, et al. (2004). Higher criticism for detecting sparse heterogeneous mixtures. *The Annals of Statistics* 32(3), 962–994.
- Dudoit, S., M. J. van der Laan, and K. S. Pollard (2004). Multiple testing. Part I. Single-step procedures for control of general type I error rates. *Statistical Applications in Genetics and Molecular Biology* 3(1), 1–69.
- Durand, G., G. Blanchard, P. Neuvial, and E. Roquain (2018). Post hoc false positive control for spatially structured hypotheses. *arXiv preprint arXiv:1807.01470*.
- Farcomeni, A. (2009). Generalized augmentation to control the false discovery exceedance in multiple testing. *Scandinavian Journal of Statistics* 36(3), 501–517.
- Finner, H. and K. Strassburger (2002). The partitioning principle: a powerful tool in multiple decision theory. *Annals of Statistics* 30(4), 1194–1213.
- Finos, L. and A. Farcomeni (2011). k-FWER control without p-value adjustment, with application to detection of genetic determinants of multiple sclerosis in Italian twins. *Biometrics* 67(1), 174–181.
- Forman, S. D., J. D. Cohen, M. Fitzgerald, W. F. Eddy, M. A. Mintun, and D. C. Noll (1995). Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): use of a cluster-size threshold. *Magnetic Resonance in medicine* 33(5), 636–647.
- Friston, K. J., A. Holmes, J.-B. Poline, C. J. Price, and C. D. Frith (1996). Detecting activations in PET and fMRI: levels of inference and power. *Neuroimage* 4(3), 223–235.
- Ge, Y. and X. Li (2012). Control of the false discovery proportion for independently tested null hypotheses. *Journal of Probability and Statistics* 2012, article 320425.
- Genovese, C. and L. Wasserman (2004). A stochastic process approach to false discovery control. *Annals of Statistics* 32(3), 1035–1061.
- Genovese, C. R. and L. Wasserman (2006). Exceedance control of the false discovery proportion. *Journal of the American Statistical Association* 101(476), 1408–1417.
- Goeman, J., R. Meijer, T. Krebs, and A. Solari (2019). Simultaneous control of all false discovery proportions

- in large-scale multiple hypothesis testing. *Biometrika*, in press.
- Goeman, J. J. and L. Finos (2012). The inheritance procedure: multiple testing of tree-structured hypotheses. *Statistical Applications in Genetics and Molecular Biology* 11(1), 1–18.
- Goeman, J. J. and U. Mansmann (2008). Multiple testing on the directed acyclic graph of gene ontology. *Bioinformatics* 24(4), 537–544.
- Goeman, J. J. and A. Solari (2010). The sequential rejection principle of familywise error control. *The Annals of Statistics* 38(6), 3782–3810.
- Goeman, J. J. and A. Solari (2011). Multiple testing for exploratory research. *Statistical Science* 26(4), 584–597.
- Goeman, J. J. and A. Solari (2014). Multiple hypothesis testing in genomics. *Statistics in Medicine* 33(11), 1946–1978.
- Gou, J. and A. C. Tamhane (2014). On generalized Simes critical constants. *Biometrical Journal* 56(6), 1035–1054.
- Guo, W., L. He, S. K. Sarkar, et al. (2014). Further results on controlling the false discovery proportion. *The Annals of Statistics* 42(3), 1070–1101.
- Guo, W., M. B. Rao, et al. (2010). On stepwise control of the generalized familywise error rate. *Electronic Journal of Statistics* 4, 472–485.
- Hemerik, J. and J. J. Goeman (2018). False discovery proportion estimation by permutations: confidence for significance analysis of microarrays. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 80(1), 137–155.
- Hemerik, J., A. Solari, and J. J. Goeman (2019). Permutation-based simultaneous confidence bounds for the false discovery proportion. *Biometrika*, in press.
- Henning, K. S. and P. H. Westfall (2015). Closed testing in pharmaceutical research: Historical and recent developments. *Statistics in Biopharmaceutical Research* 7(2), 126–147.
- Hommel, G. and T. Hoffmann (1988). Controlled uncertainty. In *Multiple Hypothesenprüfung/Multiple Hypotheses Testing*, pp. 154–161. Springer.
- Janson, L., W. Su, et al. (2016). Familywise error rate control via knockoffs. *Electronic Journal of Statistics* 10(1), 960–975.
- Javanmard, A., A. Montanari, et al. (2018). Online rules for control of false discovery rate and false discovery exceedance. *The Annals of statistics* 46(2), 526–554.
- Katsevich, E. and A. Ramdas (2018). Towards “Simultaneous Selective Inference”: post-hoc bounds on the false discovery proportion. *arXiv preprint arXiv:1803.06790*.
- Korn, E. L., J. F. Troendle, L. M. McShane, and R. Simon (2004). Controlling the number of false discoveries: application to high-dimensional genomic data. *Journal of Statistical Planning and Inference* 124(2), 379–398.
- Lehmann, E. and J. Romano (2005). Generalizations of the familywise error rate. *Annals of Statistics* 33(3), 1138–1154.
- Lehmann, E. L. and J. P. Romano (2006). *Testing statistical hypotheses*. Springer.
- Marcus, R., E. Peritz, and K. Gabriel (1976). Closed testing procedures with special reference to ordered analysis of variance. *Biometrika* 63(3), 655–660.
- Meijer, R. J. and J. J. Goeman (2015a). A multiple testing method for hypotheses structured in a directed acyclic graph. *Biometrical Journal* 57(1), 123–143.
- Meijer, R. J. and J. J. Goeman (2015b). Multiple testing of gene sets from Gene Ontology: possibilities and pitfalls. *Briefings in Bioinformatics* 17(5), 808–818.
- Meijer, R. J., T. J. Krebs, and J. J. Goeman (2015). A region-based multiple testing method for hypotheses ordered in space or time. *Statistical Applications in Genetics and Molecular Biology* 14(1), 1–19.
- Meijer, R. J., T. J. Krebs, and J. J. Goeman (2019). Hommel’s procedure in linear time. *Biometrical Jour-*

- nal* 61(1), 73–82.
- Meinshausen, N. (2006). False discovery control for multiple tests of association under general dependence. *Scandinavian Journal of Statistics* 33(2), 227–237.
- Meinshausen, N. (2008). Hierarchical testing of variable importance. *Biometrika* 95(2), 265–278.
- Meinshausen, N. and J. Rice (2006). Estimating the proportion of false null hypotheses among a large number of independently tested hypotheses. *The Annals of Statistics* 34(1), 373–393.
- Poline, J.-B. and B. M. Mazoyer (1993). Analysis of individual positron emission tomography activation maps by detection of high signal-to-noise-ratio pixel clusters. *Journal of Cerebral Blood Flow & Metabolism* 13(3), 425–437.
- Romano, J. P., A. Shaikh, and M. Wolf (2011). Consonance and the closure method in multiple testing. *The International Journal of Biostatistics* 7(1), 1–25.
- Romano, J. P. and A. M. Shaikh (2006). Stepup procedures for control of generalizations of the familywise error rate. *The Annals of Statistics* 34(4), 1850–1873.
- Rosenblatt, J. D., L. Finos, W. D. Weeda, A. Solari, and J. J. Goeman (2018). All-resolutions inference for brain imaging. *NeuroImage* 181, 786–796.
- Rosenblum, M., H. Liu, and E.-H. Yen (2014). Optimal tests of treatment effects for the overall population and two subpopulations in randomized trials, using sparse linear programming. *Journal of the American Statistical Association* 109(507), 1216–1228.
- Rosset, S., R. Heller, A. Painsky, and E. Aharoni (2018). Optimal procedures for multiple testing problems. *arXiv preprint arXiv:1804.10256*.
- Sarkar, S. K. (2007). Stepup procedures controlling generalized FWER and generalized FDR. *The Annals of Statistics* 35(6), 2405–2420.
- Shaffer, J. P. (1986). Modified sequentially rejective multiple test procedures. *Journal of the American Statistical Association* 81(395), 826–831.
- Simes, R. (1986). An improved Bonferroni procedure for multiple tests of significance. *Biometrika* 73(3), 751–754.
- Sonnemann, E. (1982). *Allgemeine Lösungen multipler Testprobleme*. Universität Bern. Institut für Mathematische Statistik und Versicherungslehre.
- Sonnemann, E. (2008). General solutions to multiple testing problems. *Biometrical Journal* 50(5), 641–656.
- Sonnemann, E. and H. Finner (1988). Vollständigkeitssätze für multiple Testprobleme. In *Multiple Hypothesenprüfung/Multiple Hypotheses Testing*, pp. 121–135. Springer.
- Sun, W., B. J. Reich, T. Tony Cai, M. Guindani, and A. Schwartzman (2015). False discovery control in large-scale spatial multiple testing. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 77(1), 59–83.
- Wang, J. and A. B. Owen (2018). Admissibility in partial conjunction testing. *Journal of the American Statistical Association*, in press.
- Westfall, P. and S. Young (1993). *Resampling-Based Multiple Testing*. New York: Wiley.