

A UNIFIED TREATMENT OF MULTIPLE TESTING WITH PRIOR KNOWLEDGE USING THE P-FILTER¹

BY AADITYA K. RAMDAS^{*}, RINA F. BARBER[†],
 MARTIN J. WAINWRIGHT[‡] AND MICHAEL I. JORDAN[‡]

Carnegie Mellon University^{}, University of Chicago[†] and
 University of California, Berkeley[‡]*

There is a significant literature on methods for incorporating knowledge into multiple testing procedures so as to improve their power and precision. Some common forms of prior knowledge include (a) beliefs about which hypotheses are null, modeled by nonuniform prior weights; (b) differing importances of hypotheses, modeled by differing penalties for false discoveries; (c) multiple arbitrary partitions of the hypotheses into (possibly overlapping) groups and (d) knowledge of independence, positive or arbitrary dependence between hypotheses or groups, suggesting the use of more aggressive or conservative procedures. We present a unified algorithmic framework called *p-filter* for global null testing and false discovery rate (FDR) control that allows the scientist to incorporate all four types of prior knowledge (a)–(d) simultaneously, recovering a variety of known algorithms as special cases.

1. Introduction. Multiple hypothesis testing is both a classical and highly active research area, dating back (at least) to an initially unpublished 1953 manuscript by Tukey entitled “The Problem of Multiple Comparisons” [45, 46]. Given a large set of null hypotheses, the goal of multiple testing is to decide which subset to reject, while guaranteeing some notion of control on the number of false rejections. It is of practical importance to incorporate different forms of prior knowledge into existing multiple testing procedures; such prior knowledge can yield improvements in power and precision, and can also provide more interpretable results. Accordingly, we study methods that control the false discovery rate (FDR) or test the global null (GN) hypothesis while incorporating any number of the following strategies for incorporating prior knowledge: (a) the use of prior weights, (b) the use of penalty weights, (c) the partitioning of the hypotheses into groups, (d) the incorporation of knowledge of the dependence structure within the data, including options such as estimating and adapting to

Received April 2017; revised September 2018.

¹Supported in part by the Office of Naval Research GRant W911NF-16-1-0368, Air Force Office of Scientific Research Grant AFOSR-FA9550-14-1-0016, NSF Grant DMS-1654076 and by an Alfred P. Sloan fellowship.

MSC2010 subject classifications. Primary 62J15, 60G10; secondary 62F03.

Key words and phrases. Multiple testing, false discovery rate, prior knowledge, Simes, Benjamini–Hochberg–Yekutieli, adaptivity, group FDR.

the unknown number of nulls under independence, or reshaping rejection thresholds to preserve error-control guarantees in the presence of arbitrary dependence. It is a challenge to incorporate all of these forms of structure while maintaining internal consistency (coherence and consonance) among the pattern of rejections and acceptances, and most existing work has managed to allow only one or two of the four strategies (a), (b), (c), (d) to be employed simultaneously. We present a general unified framework, called *p-filter*, for integrating these four strategies while performing a GN test or controlling the FDR. The framework is accompanied by an efficient algorithm, with code publicly available at <https://www.stat.uchicago.edu/~rina/pfilter.html>. This framework allows scientists to mix and match techniques, and use multiple different forms of prior knowledge simultaneously. As a by-product, our proofs often simplify and unify the analysis of existing procedures, and generalize the conditions under which they are known to work.

Organization. The rest of this paper is organized as follows. In Section 2, we begin with an example to provide intuition, and we discuss the contributions of this paper. In Section 3, we describe the general *p-filter* framework, along with its associated theoretical guarantees; this section lays out the central contribution of our work. In Section 4, we present three lemmas that provide valuable intuition and are central to the proof of Theorem 1; see Section 5 for the proof itself. We prove the three aforementioned lemmas in Section 6, and prove some related propositions in Section 7. While directly related work is discussed immediately when referenced, we end by overviewing other related work in Section 8.

2. An example and our contributions. The various kinds of prior information considered in this paper have been studied in earlier works and repeatedly motivated in applied settings, and our focus is accordingly on the conceptual and mathematical aspects of multiple decision-making with prior knowledge. Before beginning our formal presentation, we consider a simple example, illustrated in Figure 1, in order to provide intuition.

Consider a set of sixteen hypotheses arranged in a 4×4 grid, as displayed in the first panel of Figure 1. One may imagine that one coordinate refers to spatial locations, and the other to temporal locations, so that each square represents an elementary null hypothesis $H_{s,t}$, stating that there is nothing of interest occurring at spatial location s at time t . As displayed in the leftmost panel of the first row, the nonnulls may be expected to have some spatiotemporal structure. In order to exploit this structure, the scientist may choose to group the hypotheses a priori in three ways: by spatial location, by temporal location and by spatiotemporal blocks, as displayed by the other three panels in the first row. Each such group can be associated with a group null hypothesis, which states that there is nothing of interest occurring within that group.




FIG. 1. Consider a set of $n = 16$ hypotheses arrayed in a 4×4 grid, with four different partitions into groups: elementary, rows, columns and blocks. On the top row is the underlying truth, with the leftmost panel showing the hypothesis-level nonnulls, and the other three panels showing which groups in each partition are hence identified as nonnull. On the bottom row is an example of a set of discoveries, with the leftmost panel showing the hypothesis-level rejections, and the other three panels showing which groups are correspondingly rejected (light-gray for correct rejections, black for false rejections). The false discovery proportions (FDP) in each partition are 0.2, 0, 0.33, 0.5, respectively. The true discovery proportions (empirical power) in each partition are 0.8, 0.66, 1, 0.5, respectively.

As displayed in the first row, the group-level nonnulls are simply implied by the elementary nonnulls. The second row of the figure displays the results of a hypothetical procedure that makes some elementary rejections (first panel), and hence some corresponding group level rejections (other three panels). The scientist may wish to not report too many false elementary rejections, but also not report discoveries at spurious locations, times or space-time blocks. One way of enabling this wish is to enforce group FDR constraints, in addition to overall FDR control. This would correspond to controlling the “spatial” FDR, “temporal” FDR, and “spatiotemporal” FDR, in addition to the overall FDR. Requiring the rejected hypotheses to satisfy additional constraints may reduce power, but it may also increase interpretability, and result in higher precision (achieved FDP).

For example, Barber and Ramdas [1] consider an example from neuroscience where each null hypothesis $H_{v,t}$ states that at t seconds after the presentation of the stimulus, a chosen feature of the stimulus is unrelated to (a specific measure of) the brain activity at voxel v in the brain. One may spatially group these hypotheses according to predefined regions of interest (ROIs) such as the visual cortices V1 to V5 and the left/right temporal cortices. This spatial partition allows us to capture the idea that either a ROI is unrelated to the stimulus, or many of its voxels v will be related. Similarly, for a fixed voxel, one may temporally group the hypotheses, to capture the idea that either a voxel v will remain unrelated to the stimulus at various delays, or it will be related at several consecutive delays (usually $t = 4, 6, 8, 10$ seconds after stimulus onset).

The earlier *p-filter* algorithm [1] simultaneously guarantees FDR control for multiple arbitrary and possibly nonhierarchical partitions, while ensuring that the elementary and group rejections are “internally consistent.” It is a multidimensional step-up procedure, one which reduces to the BH step-up procedure or Simes’ GN test in special cases. The reader may refer to the original paper for numerical simulations and more details on the neuroscience example.

Our contributions. Consider a collection $\{H_1, \dots, H_n\}$ of n unordered hypotheses, along with associated p -values $\{P_1, \dots, P_n\}$. It is convenient to introduce the shorthand notation $[n] := \{1, \dots, n\}$. Further, consider M arbitrary unordered partitions of these hypotheses into groups, where the m th partition (“layer”) contains an unordered set of $G^{(m)}$ groups:

$$A_1^{(m)}, \dots, A_{G^{(m)}}^{(m)} \subseteq [n] \quad \text{for } m = 1, \dots, M.$$

It may help the reader to imagine the first partition as being the elementary or finest partition, meaning that it contains n groups of one hypothesis each, but it is important to note that this partition is entirely optional, and can be dropped if there is no desire of controlling the overall FDR. If $\mathcal{H}_0 \subseteq [n]$ are the true nulls, then we call a group g null if $A_g^{(m)} \subseteq \mathcal{H}_0$. Let the set of null groups in partition m be denoted by $\mathcal{H}_0^{(m)}$.

Although we continue to use the name *p-filter* for the algorithm that we discuss in this paper, the algorithm goes significantly beyond the original algorithm; in particular, in our general setting of M arbitrary partitions there are seven ways in which the new procedure to be developed here goes beyond the original framework.

1. *Overlapping groups.* We allow the groups in any partition to overlap. An elementary hypothesis need not be part of just a single group $g \in [G^{(m)}]$, and we let $g^{(m)}(i)$ denote the set of groups in the m th partition to which P_i belongs namely to

$$g^{(m)}(i) = \{g \in [G^{(m)}] : P_i \in A_g^{(m)}\}.$$

For example, in the neuroscience example introduced earlier, if the scientist is unsure about the accuracy of the ROI borders, they may place boundary hypotheses into two or more ROIs to reflect this uncertainty.

2. *Incomplete partitions.* We allow partitions to be incomplete. Let the m th partition’s leftover set $L^{(m)} \subset [n]$ represent all elementary hypotheses that do not belong to any group in the m th partition:

$$L^{(m)} = [n] \setminus \bigcup_g A_g^{(m)}.$$

This gives additional flexibility to the user who may not want to assign some hypotheses to any groups. Note that $L^{(m)}$ is not just another group; this set is not

counted when calculating the group-level FDR in layer m , meaning that elementary discoveries within $L^{(m)}$ do not alter the group FDR at layer m . Hence, hypotheses in this leftover set have no internal consistency constraints imposed by layer m . For instance, in the neuroscience example, if it is determined (e.g., due to brain damage or surgery) that some voxels may not naturally fit into any ROI, then they can be left out of that partition.

3. *Internal consistency (IC)*. In order to maintain interpretability when dealing with overlapping groups, it is convenient to introduce two natural notions of internal consistency of the group rejections and elementary rejections:

- *Weak IC*. We reject H_i if and only if in every partition, either there is at least one rejected group containing i , or $i \in L^{(m)}$.
- *Strong IC*. We reject H_i if and only if in every partition, either every group that contains i is rejected, or $i \in L^{(m)}$.

These definitions² of IC are extensions to the multilayer setting of the notions of *coherence* and *consonance* as defined by Gabriel [14], and explored in the FWER literature by Sonnemann and Finner [39–41] and Romano et al. [32]. In the aforementioned neuroscience example, weak internal consistency may be more appropriate.

4. *Weights*. The m th partition can be associated with two sets of positive weights, one pair for each group g in that partition:

$$\text{Penalties } \{u_g^{(m)}\} \text{ and priors } \{w_g^{(m)}\}, \text{ such that } \sum_{g=1}^{G^{(m)}} u_g^{(m)} w_g^{(m)} = G^{(m)}.$$

This generalizes work on doubly weighted procedures by Blanchard and Roquain [10], who considered a single partition. Their work in turn generalizes earlier work using prior weights [16] and penalty weights Benjamini and Hochberg [4] separately. Large prior weights indicate beliefs that the hypotheses are more likely to be nonnull, and large penalties reflect which hypotheses are more scientifically important. For example, in the neuroscience example, weights can also be used to take differing ROI sizes into account, or prior knowledge of when and where effects are expected to be found.

5. *Reshaping*. Reshaping functions β are used to guard against possible dependence among the p -values by *undercounting* the size (or weight) of rejected sets. Reshaping makes it possible to handle arbitrary dependence; on the other hand, this favorable robustness property is accompanied by a loss of power. Following Blanchard and Roquain [10], for any probability measure ν on $[0, \infty)$, we define

²We remark that these are not the only two notions of internal consistency that can fit into our framework: any *monotone* notion of IC can be handled, where *monotone* means that decreasing the p -values can only possibly increase the number of rejections at all layers.

the reshaping function

$$(2.1) \quad \beta(k) := \int_0^k x \, dv(x) \leq k.$$

If the p -values within or across layers are arbitrarily dependent, we may use reshaping functions $\beta^{(m)}$ to reshape thresholds in layer m . In the special case that Simes' p -values are used to form the group-level p -values $P_g^{(m)}$, we may use reshaping functions $\tilde{\beta}_g^{(m)}$ to protect the Simes' procedure from arbitrary dependence within the group. The original procedure of Benjamini and Yekutieli [9] corresponds to choosing the reshaping function $\beta_{BY}(k) = \frac{k}{\sum_{i=1}^n \frac{1}{i}}$. Many other examples and their connections to other formulations of multiple testing methods can be found in the literature [10, 34, 35]. In contrast to the discrete distributions which have been the focus of past work, in the current paper it is necessary to consider continuous measures since the penalty weight of rejected hypotheses, unlike their count, can be fractional.

6. *Adaptivity*. For any partition whose group p -values are known to be independent (i.e., independence *between* groups, but not necessarily *within* each group), we can incorporate “null-proportion adaptivity” for that partition [5, 18]. For partition m , we fix a user-defined constant $\lambda^{(m)} \in (0, 1)$, and define a weighted null-proportion estimator:

$$(2.2) \quad \hat{\pi}^{(m)} := \frac{|u^{(m)} \cdot w^{(m)}|_\infty + \sum_g u_g^{(m)} w_g^{(m)} \mathbf{1}\{P_g^{(m)} > \lambda^{(m)}\}}{G^{(m)}(1 - \lambda^{(m)})}.$$

The use of null-proportion adaptivity in any *one layer* may improve the power in *all layers*, since more groups being discovered in one layer leads to more elementary discoveries, and hence more discovered groups in other layers. For a single group with no weights, our approach reduces to the original suggestion of Storey et al. [42, 43].

7. *Arbitrary group p -values*. Our new p -filter algorithm is no longer tied to the use of Simes' p -values at the group layers, unlike the original algorithm. In other words, each group-level p -value at each layer can be formed by combining the elementary p -values within that group [17, 47]. When the p -values are independent, some options include Fisher's $-2 \sum_i \ln P_i$ and Rosenthal's $\sum_i \Phi^{-1}(P_i)$, where Φ is the Gaussian CDF (originally proposed by Stouffer et al. [44]). When there are very few nonnulls, the Bonferroni correction is known to be more powerful, and it also works under arbitrary dependence, as does Rüschendorf's proposal of $2 \sum_i P_i / n$, and Rüger's proposal of $P_{(k)} \cdot n/k$ for a fixed k . Alternately, the group p -values can be constructed directly from raw data. Accordingly, we can appropriately use adaptivity or reshaping, as needed, depending on the induced dependence.

Suppose a procedure rejects a subset $\widehat{\mathcal{S}} \subseteq [n]$ of hypotheses and a subset $\widehat{\mathcal{S}}^{(m)} \subseteq [G^{(m)}]$ of groups in partition m . Then we may define the penalty-weighted FDR as

$$\text{FDR}_u^{(m)} = \mathbb{E} \left[\frac{\sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} \mathbf{1}\{g \in \widehat{\mathcal{S}}^{(m)}\}}{\sum_{g \in [G^{(m)}]} u_g^{(m)} \mathbf{1}\{g \in \widehat{\mathcal{S}}^{(m)}\}} \right].$$

The rejections made by `p-filter` will be internally consistent, and satisfy

$$\text{FDR}_u^{(m)} \leq \alpha^{(m)} \quad \text{simultaneously for all } m = 1, \dots, M.$$

In order to handle the possibility of a ratio “ $\frac{0}{0}$ ” in this definition, and in our later work, we adopt the “dotfraction” notation

$$(2.3) \quad \frac{a}{\dot{b}} := \begin{cases} 0 & \text{if } a = 0, \\ \frac{a}{b} & \text{if } a \neq 0, b \neq 0, \\ \text{undefined} & \text{if } a \neq 0, b = 0. \end{cases}$$

Dotfractions behave like fractions whenever the denominator is nonzero. The use of dotfractions simplifies the presentation: note that $\frac{a}{\dot{b}} \neq \frac{a}{\max(b, 1)}$ since b may be fractional due to the use of penalty weights. We formally derive properties of dotfractions in the Supplementary Material (Appendix C) [31].

When there is only one partition, and the weights equal one, the quantity $\text{FDR}_u^{(m)}$ reduces to the usual FDR defined by Benjamini and Hochberg [3] and `p-filter` reduces to the BH procedure.³ Many other procedures are recovered as special cases of `p-filter`, as detailed after Theorem 1.

3. FDR control and internal consistency for multiple partitions. Even in the case of just two partitions—one partition of groups and the elementary partition of individual hypotheses—it is nontrivial to provide a guarantee of internal consistency while controlling both group-level and individual-level FDR. For example, a sequential procedure, of first rejecting groups at a target FDR level α_1 and then rejecting individual hypotheses within rejected groups at level α_2 , may not control the elementary FDR (due to not accounting for selection bias), and may not be internally consistent (because there might be a group rejected in the first round, with none of the elementary hypotheses in this group rejected in the second round). Further, such a method is not easily generalized to nonhierarchical partitions. Similarly, a parallel procedure that independently runs FDR procedures on the groups and on the individuals, may also fail to be internally consistent. Naively intersecting these rejections—that is, rejecting those hypotheses whose groups are rejected at every layer—may also fail to control the FDR (see Barber and Ramdas [1] for explicit examples).

³For a review of its history that involves Eklund and Seeger in the 1960s, and Simes, Hommel, Soric, Benjamini and Hochberg in the 1980s and 1990s, see [5, 37].

The p -filter algorithm is a multivariate extension of classical step-up procedures that is roughly based on the following sequence of steps:

- Select all hypotheses in each layer whose p -values are smaller than some initial layer-specific threshold.
- Reject an elementary hypothesis if it is contained in a selected group in every layer.
- In each layer, reject a group hypothesis if it contains a rejected elementary hypothesis. Then estimate the group-FDP in each layer.
- Lower the initial thresholds at each layer, and repeat the steps above, until the group-FDP is below the desired level for all partitions.

Next, we discuss the necessary dependence assumptions and then derive the p -filter algorithm that implements the above scheme.

3.1. Marginal and joint distributional assumptions on p -values. We assume that the marginal distribution of each null p -value is stochastically larger than the uniform distribution, referred to as *superuniform* for brevity. Formally, for any index $i \in \mathcal{H}_0$, we assume that

$$(3.1) \quad \Pr\{P_i \leq t\} \leq t \quad \text{for all } t \in [0, 1].$$

Of course, uniformly distributed p -values trivially satisfy this condition. We use the phrase *under uniformity* to describe the situation in which the null p -values are marginally exactly uniform. If this phrase is not employed, it is understood that the null p -values are marginally superuniform.

Regarding assumptions on the joint distribution of p -values, three possible kinds of dependence will be considered in this paper: independence, positive dependence or arbitrary dependence. In the independent setting, null p -values are assumed to be mutually independent and independent of nonnulls. In the arbitrary dependence setting, no joint dependence assumptions are made on the p -values. The last case is that of positive dependence, as formalized by the *Positive Regression Dependence on a Subset* (PRDS) condition [9, 25, 36]. In order to understand its definition, it is helpful to introduce some basic notation. For a pair of vectors $x, y \in [0, 1]^n$, we use the notation $x \preceq y$ to mean that $x \leq y$ in the orthant ordering, that is, $x_i \leq y_i$ for all $i \in \{1, \dots, n\}$. A set $\mathcal{D} \subseteq [0, 1]^n$ is said to be *nondecreasing* if $x \in \mathcal{D}$ implies $y \in \mathcal{D}$ for all $y \succeq x$. We say that a function $f : [0, 1]^n \mapsto [0, \infty)$ is *nonincreasing*, if $x \preceq y$ implies $f(x) \geq f(y)$.

DEFINITION 1 (Positive dependence, PRDS). We say that the vector P satisfies PRDS if for any null index $i \in \mathcal{H}_0$ and nondecreasing set $\mathcal{D} \subseteq [0, 1]^n$, the function $t \mapsto \Pr\{P \in \mathcal{D} \mid P_i \leq t\}$ is nondecreasing over $t \in (0, 1]$.

The original positive regression dependence assumption was introduced by Lehmann [25] in the bivariate setting and by Sarkar [36] in the multivariate setting,

and extended to the PRDS assumption first made by Benjamini and Yekutieli [9]. These previous papers used the equality $P_i = t$ instead of the inequality $P_i \leq t$ in the definitions, but one can prove that both conditions are essentially equivalent.

The PRDS condition holds trivially if the p -values are independent, but also allows for some amount of positive dependence. For intuition, suppose that $Z = (Z_1, \dots, Z_n)$ is a multivariate Gaussian vector with covariance matrix Σ ; the null components correspond to Gaussian variables with zero mean. Letting Φ be the CDF of a standard Gaussian, the vector of p -values $P = (\Phi(Z_1), \dots, \Phi(Z_n))$ is PRDS on P_i for every index i if and only if all entries of the covariance matrix Σ are nonnegative; see Benjamini and Yekutieli [9] for additional examples of this type.

3.2. Specifying the `p-filter` algorithm. In order to run the `p-filter` algorithm, we need to search for rejection thresholds for each layer. These thresholds will be parametrized by *weighted* discovery counts $k^{(m)} \in [0, G^{(m)}]$ for each layer $m = 1, \dots, M$. The reader is cautioned that each $k^{(m)}$ need not be an integer but instead should be viewed as a real number corresponding to the total rejected penalty weight. If the weights $u_g^{(m)}$ are all set equal to one, then $k^{(m)}$ corresponds to the number of groups in layer m that are rejected. Given some prototypical vector $\vec{k} := (k^{(1)}, \dots, k^{(M)})$, we first perform an initial screening on each layer separately:

$$(3.2) \quad \hat{\mathcal{S}}_{\text{init}}^{(m)}(\vec{k}) = \left\{ g \in [G^{(m)}] : P_g^{(m)} \leq \min \left\{ \frac{w_g^{(m)} \alpha^{(m)} \beta^{(m)}(k^{(m)})}{\hat{\pi}^{(m)} G^{(m)}}, \lambda^{(m)} \right\} \right\}.$$

If the groups in partition m are independent, we replace $\beta^{(m)}(k^{(m)})$ by just $k^{(m)}$, and set $\hat{\pi}^{(m)}$ using (2.2); on the other hand, if they are arbitrarily dependent, we set $\hat{\pi}^{(m)} = 1$ and $\lambda^{(m)} = 1$. This convention allows the same expressions to be used in all settings.

For weak internal consistency, we define the elementary rejections as

$$(3.3a) \quad \begin{aligned} \hat{\mathcal{S}}(\vec{k}) &= \hat{\mathcal{S}}_{\text{weak}}(\vec{k}) = \bigcap_{m=1}^M \left(\left[\bigcup_{g \in \hat{\mathcal{S}}_{\text{init}}^{(m)}(\vec{k})} A_g^{(m)} \right] \cup L^{(m)} \right) \\ &= \{ P_i : \forall m, \text{ either } P_i \in L^{(m)}, \text{ or } \exists g \in g^{(m)}(i), A_g^{(m)} \in \hat{\mathcal{S}}_{\text{init}}^{(m)}(\vec{k}) \}. \end{aligned}$$

Alternately, for strong internal consistency, we may instead define

$$(3.3b) \quad \begin{aligned} \hat{\mathcal{S}}(\vec{k}) &= \hat{\mathcal{S}}_{\text{strong}}(\vec{k}) = \bigcap_{m=1}^M \left([n] \setminus \bigcup_{g \in [G^{(m)}] \setminus \hat{\mathcal{S}}_{\text{init}}^{(m)}(\vec{k})} A_g^{(m)} \right) \\ &= \{ P_i : \forall m, \text{ either } P_i \in L^{(m)}, \text{ or } \forall g \in g^{(m)}(i), A_g^{(m)} \in \hat{\mathcal{S}}_{\text{init}}^{(m)}(\vec{k}) \}. \end{aligned}$$

Finally, using either $\widehat{\mathcal{S}}(\vec{k}) = \widehat{\mathcal{S}}_{\text{weak}}(\vec{k})$ or $\widehat{\mathcal{S}}(\vec{k}) = \widehat{\mathcal{S}}_{\text{strong}}(\vec{k})$, we redefine the set of groups in layer m which are rejected as

$$(3.4) \quad \widehat{\mathcal{S}}^{(m)}(\vec{k}) = \{g \in [G^{(m)}] : A_g^{(m)} \cap \widehat{\mathcal{S}}(\vec{k}) \neq \emptyset \text{ and } g \in \widehat{\mathcal{S}}_{\text{init}}^{(m)}(\vec{k})\}.$$

Examining these definitions, it may be verified that (weak or strong) internal consistency is satisfied by the rejections $\widehat{\mathcal{S}}(\vec{k})$, $\widehat{\mathcal{S}}_m(\vec{k})$.

Of course, these definitions depend on the initial choice of the vector \vec{k} . Since we would like to make a large number of discoveries, we would like to use a \vec{k} that is as large as possible (coordinatewise), while at the same time controlling the layer-specific FDRs, which are the expectations of

$$\text{FDP}_u^{(m)}(\vec{k}) := \frac{\sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} \mathbf{1}\{g \in \widehat{\mathcal{S}}^{(m)}(\vec{k})\}}{\sum_{g \in [G^{(m)}]} u_g^{(m)} \mathbf{1}\{g \in \widehat{\mathcal{S}}^{(m)}(\vec{k})\}}.$$

Now, define the data-dependent set of feasible vectors \vec{k} as

$$(3.5) \quad \widehat{\mathcal{K}} = \left\{ \vec{k} \in [0, G_1] \times \cdots \times [0, G_M] : \sum_{g \in \widehat{\mathcal{S}}^{(m)}(\vec{k})} u_g^{(m)} \geq k^{(m)} \text{ for all } m \right\},$$

where we suppress the implicit dependence of $\widehat{\mathcal{K}}$ on input parameters such as $\alpha^{(m)}$, $\lambda^{(m)}$, $\{w_g^{(m)}\}$, $\{u_g^{(m)}\}$. In particular, if the penalty weights are all equal to one, then the consistency condition defining the “feasible” \vec{k} ’s is equivalent to requiring that $|\widehat{\mathcal{S}}^{(m)}(\vec{k})| \geq k^{(m)}$ for all $m = 1, \dots, M$; that is, the numbers of rejections in each layer at the vector \vec{k} are elementwise $\geq \vec{k}$. This condition can be viewed as a generalization, to the multipartition setting, of the “self-consistency” condition described by Blanchard and Roquain [10].

It is also worth noting that the p-filter algorithm in Barber and Ramdas [1] was derived in terms of thresholds \vec{t} instead of number of rejections \vec{k} , and there the corresponding feasibility condition was that $\text{FDP}_m(\vec{t}) \leq \alpha^{(m)}$, where $\text{FDP}_m(\vec{t})$ is an empirical-Bayes type estimate of the FDP. Indeed, if we avoid $\hat{\pi}^{(m)}$, $\beta^{(m)}$, $w^{(m)}$, $u^{(m)}$ for simplicity, then associating $\hat{t}^{(m)}$ to $\alpha^{(m)} \hat{k}^{(m)} / G^{(m)}$ and comparing our derivation to that of Barber and Ramdas [1], we can see that the “self-consistency” viewpoint and the “empirical-Bayes” viewpoint are equivalent and lead to the same algorithm. However, when dealing with reshaping under arbitrary dependence, the proofs are simpler in terms of \vec{k} than in terms of \vec{t} , explaining our switch in notation.

As with the BH and BY procedures, we then choose the largest feasible thresholds $k^{(m)}$, given by

$$(3.6) \quad \hat{k}^{(m)} = \max\{k^{(m)} : \exists k^{(1)}, \dots, k^{(m-1)}, k^{(m+1)}, \dots, k^{(M)} \text{ s.t. } \vec{k} \in \widehat{\mathcal{K}}\}.$$

This choice defines our algorithm: the p-filter algorithm rejects the hypotheses $\widehat{\mathcal{S}}(\hat{k}^{(1)}, \dots, \hat{k}^{(M)})$, as defined in (3.3a) or (3.3b), with rejections at layer m given by $\widehat{\mathcal{S}}^{(m)}(\hat{k}^{(1)}, \dots, \hat{k}^{(M)})$ as defined in (3.4). Next, we present the theoretical guarantees associated with p-filter.

3.3. *Theoretical guarantees.* The following proposition states that the set of feasible vectors $\widehat{\mathcal{K}}$ actually has a well-defined “maximum” corner.

PROPOSITION 1. *Let the set of feasible vectors $\widehat{\mathcal{K}}$ be defined as in equation (3.5), and let the partition-specific maximum feasible vector $\widehat{k}^{(m)}$ be defined as in equation (3.6). Then we have*

$$(3.7) \quad (\widehat{k}^{(1)}, \dots, \widehat{k}^{(M)}) \in \widehat{\mathcal{K}}.$$

The proof is provided in Section 7.1; it is a generalization of the corresponding result for the original p-filter algorithm [1].

The vector $(\widehat{k}^{(1)}, \dots, \widehat{k}^{(M)})$ is not just feasible from the perspective of self-consistency as captured by $\widehat{\mathcal{K}}$, but it is also feasible from the perspective of FDR control. Specifically, the next theorem establishes that—assuming for now that we can find $(\widehat{k}^{(1)}, \dots, \widehat{k}^{(M)})$ —selecting the set $\widehat{\mathcal{S}}(\widehat{k}^{(1)}, \dots, \widehat{k}^{(M)})$ guarantees simultaneous control of $\text{FDR}_u^{(m)}$ for all M partitions. In this theorem, the notation $\text{Simes}_w(P_{A_g^{(m)}})$ refers to the weighted Simes’ p -value (see Appendix A in the Supplementary Material [31] for details).

THEOREM 1. *Any procedure that computes the vector $(\widehat{k}^{(1)}, \dots, \widehat{k}^{(M)})$ according to definition (3.6) satisfies the following properties, for all partitions $m = 1, \dots, M$ simultaneously:*

(a) *If the base p -values are independent, and all group p -values are given by $P_g^{(m)} = \text{Simes}_w(P_{A_g^{(m)}})$, then employing adaptivity by defining $\widehat{\pi}^{(m)}$ as in (2.2) guarantees that $\text{FDR}_u^{(m)} \leq \alpha^{(m)}$.*

(b) *If base p -values are positively dependent (PRDS) and group p -values are given by $P_g^{(m)} = \text{Simes}_w(P_{A_g^{(m)}})$, then without adaptivity or reshaping, we have that $\text{FDR}_u^{(m)} \leq \alpha^{(m)} \frac{\sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)}}{G^{(m)}} \leq \alpha^{(m)}$.*

(c) *When all p -values are arbitrarily dependent, and are constructed arbitrarily (under the assumption that $P_g^{(m)}$ is superuniform for any null group $g \in \mathcal{H}_0^{(m)}$, meaning it is a valid p -value), then using reshaping as in (2.1) guarantees that $\text{FDR}_u^{(m)} \leq \alpha^{(m)} \frac{\sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)}}{G^{(m)}} \leq \alpha^{(m)}$.*

(d) *In the setting of part (c), if additionally the groups at layer m are independent (i.e., $P_{A_g^{(m)}}$ is independent from $P_{-A_g^{(m)}}$, for each $g \in [G^{(m)}]$), then using reshaping as in (2.1) and adaptivity for layer m as in (2.2), guarantees that $\text{FDR}_u^{(m)} \leq \alpha^{(m)}$.*

The proof, given in Section 5, uses three interpretable lemmas that we first discuss in Section 4. It also introduces several new ideas to handle overlapping

groups with dependent p -values. To remark on the difference between parts (c) and (d), what these two results guarantee is that if we use adaptivity for some set $\mathcal{M}_{\text{adapt}} \subset [M]$ of layers, and do not use adaptivity (i.e., set $\hat{\pi}^{(m)} = 1$) for the remaining layers, then FDR control is maintained across *all* layers as long as, for each $m \in \mathcal{M}_{\text{adapt}}$, the layer-specific independence statement holds— $P_{A_g^{(m)}}$ is independent from $P_{-A_g^{(m)}}$, for each $g \in [G^{(m)}]$. If this condition fails for some $m \in \mathcal{M}_{\text{adapt}}$, the FDR control in other layers will in fact not be affected. One application of statement (d) is when the base p -values are independent, there are no overlapping groups, and group p -values are formed using a Fisher, Rosenthal or other combinations of the base p -values. Recently, Katsevich and Sabatti [23] proved that in case (d), the FDR is controlled even without using reshaping, albeit at a constant factor larger than the target level.

In practice, if we have accurate side information about group structures that the rejected hypotheses likely respect, then we may significantly improve our *precision*, achieving a lower FDR than the theoretical bound, without affecting our power much. However, inaccurate side information may significantly lower our power, since each p -value would have additional misguided constraints to meet. These issues were explored in simulations by Barber and Ramdas [1].

Special cases. The setting with a single partition ($M = 1$) recovers a wide variety of known algorithms. Considering only the finest partition with n groups containing one hypothesis each, the `p-filter` algorithm and associated Theorem 1 together recover known results about (a) the BH procedure of Benjamini and Hochberg [3] when weights, reshaping and adaptivity are not used, (b) the BY procedure of Benjamini and Yekutieli [9] when reshaping is used, (c) the prior-weighted BH procedure of Genovese et al. [16] when only prior weights are used, (d) the penalty-weighted BH procedure of Benjamini and Hochberg [4] when only penalty weights are used, (e) the doubly weighted BH procedure of Blanchard and Roquain [10] when both sets of weights and reshaping are used and (f) the Storey-BH procedure of Storey et al. [42, 43] when only adaptivity is used.

When we instantiate `p-filter` with the coarsest partitions with a single group containing all n hypotheses, we recover exactly (g) the Simes' test [38] without weights, and (h) a variant by Hochberg and Liberman [19] if prior weights are used. We recover the results of (i) the `p-filter` by Barber and Ramdas [1] under positive dependence, when we do not use weights, adaptivity, reshaping, overlapping groups, leftover sets and restrict ourselves to Simes' p -values. We also recover a host of new procedures: for example, while the past literature has not yet shown how to use either prior or penalty weights together with adaptivity, `p-filter` reduces to (j) a doubly weighted adaptive procedure for the finest partition under independence. Also, while the aforementioned procedures were each proved under one form of dependence or the other, we recover results for all three forms of dependence at one go, with a single unified proof technique.

Algorithm 1 The p-filter for multilayer FDR control

Input: M possibly incomplete partitions of possibly overlapping groups of indices $[n]$;

A vector of base p -values $P \in [0, 1]^n$;

Group p -values $P_g^{(m)}$ for each group $g = 1, \dots, G^{(m)}$ in layers $m = 1, \dots, M$;

M target FDR levels $\{\alpha^{(m)}\}$;

M sets of prior weights and/or penalty weights $\{w_g^{(m)}, u_g^{(m)}\}$;

M thresholds for adaptive null proportion estimation $\{\lambda^{(m)}\}$;

M reshaping functions $\{\beta^{(m)}\}$, if desired.

Initialize: Set $k^{(m)} = G^{(m)}$, and $\hat{\pi}^{(m)}$ as in Definition (2.2).

repeat

for $m = 1, \dots, M$ **do**

Update the m th vector: defining $\widehat{\mathcal{S}}^{(m)}(\vec{k})$ as in equation (3.4) (using weak or strong internal consistency, as desired), let

$$k^{(m)} \leftarrow \max \left\{ k'^{(m)} \in [0, G^{(m)}] : \sum_{g \in \widehat{\mathcal{S}}^{(m)}(k^{(1)}, \dots, k^{(m-1)}, k'^{(m)}, k^{(m+1)}, \dots, k^{(M)})} u_g^{(m)} \geq k'^{(m)} \right\}$$

end for

until the vectors $k^{(1)}, \dots, k^{(M)}$ are all unchanged for one full cycle.

Output: Vector $\hat{k} = (k^{(1)}, \dots, k^{(M)})$, rejected hypotheses $\widehat{\mathcal{S}}(\hat{k})$, and rejected groups $\widehat{\mathcal{S}}^{(m)}(\hat{k})$ in each partition.

3.4. An efficient implementation. Although one can employ a brute-force grid search to find $(\hat{k}^{(1)}, \dots, \hat{k}^{(M)})$, the p-filter algorithm presented in Algorithm 1 is able to find this vector efficiently using a coordinate-descent style procedure, and is a strict generalization of the algorithm by the same name in Barber and Ramdas [1].

The following proposition provides a correctness guarantee for Algorithm 1.

PROPOSITION 2. *The output of Algorithm 1 is the maximum feasible corner $(\hat{k}^{(1)}, \dots, \hat{k}^{(M)})$ defined in equations (3.6) and (3.7).*

This result was proved by Barber and Ramdas [1] in the setting of the original p-filter algorithm, where the $k^{(m)}$'s take only integer values; here, the algorithm is slightly more subtle, with real-valued $k^{(m)}$'s due to the presence of penalty weights $u_g^{(m)}$. The proof of the proposition for this more general setting is given in Section 7.2.

4. Three lemmas. In this section, we present three lemmas that lie at the heart of the succinct proofs of the theorems in this paper. Our motivation for presenting these lemmas here is that they are interpretable, and provide valuable intuition for the proofs that follow.

4.1. *A superuniformity lemma for FDR control.* In order to develop some intuition for the lemma that follows, we note that our superuniformity assumption (3.1) on the null p -values can be reformulated as

$$(4.1) \quad \text{For any } i \in \mathcal{H}_0, \quad \mathbb{E} \left[\frac{\mathbf{1}\{P_i \leq t\}}{t} \right] \leq 1 \quad \text{for any fixed } t \in [0, 1].$$

Of course, if P_i is uniform then the above inequality holds with equality.

The following lemma guarantees that property (4.1) continues to hold for certain random thresholds $f(P)$. Recall that the term “nonincreasing” is interpreted coordinatewise, with respect to the orthant ordering.

LEMMA 1 (Superuniformity lemma). *Let $i \in \mathcal{H}_0$ be a null hypothesis with p -value P_i , and let P^{-i} denote the other $n - 1$ p -values.*

(a) *For any nonincreasing function $f : [0, 1]^n \rightarrow [0, \infty)$, if P_i is independent of P^{-i} , then we have*

$$\mathbb{E} \left[\frac{\mathbf{1}\{P_i \leq f(P)\}}{f(P)} \mid P^{-i} \right] \leq 1.$$

Furthermore, if we additionally assume that f has range $[0, 1]$ and satisfies the LOOP condition (Supplementary Material, Appendix B [31]), and that P_i is uniformly distributed, then the inequality is replaced with equality

$$\mathbb{E} \left[\frac{\mathbf{1}\{P_i \leq f(P)\}}{f(P)} \mid P^{-i} \right] = 1.$$

(b) *For any nonincreasing function $f : [0, 1]^n \rightarrow [0, \infty)$, if P is PRDS with respect to P_i , then*

$$\mathbb{E} \left[\frac{\mathbf{1}\{P_i \leq f(P)\}}{f(P)} \right] \leq 1.$$

(c) *For any constant $c \geq 0$, any function $f : [0, 1]^n \rightarrow [0, \infty)$, and any reshaping function β , under arbitrary dependence of the p -values,*

$$\mathbb{E} \left[\frac{\mathbf{1}\{P_i \leq c \cdot \beta(f(P))\}}{c \cdot f(P)} \right] \leq 1.$$

(d) *For any constant $c \geq 0$, any functions $f_1, \dots, f_m : [0, 1]^n \rightarrow [0, \infty)$, and any reshaping functions $\beta_1, \dots, \beta^{(m)}$, under arbitrary dependence of the p -values,*

$$\mathbb{E} \left[\frac{\mathbf{1}\{P_i \leq c \cdot \prod_{\ell=1}^m \beta_\ell(f_\ell(P))\}}{c \cdot \prod_{\ell=1}^m f_\ell(P)} \right] \leq 1.$$

The proofs of statement (a) with equality, and of statement (d), are given in Section 6.1. Statement (a) with inequality is recovered as a special case of statement (b), which was proved by Blanchard and Roquain [10], who also proved (c). The more general statement (d), with more than one reshaping function present in the bound, will be required in the proof of the following novel group superuniformity Lemma 2.

4.2. A group-level superuniformity lemma. In analogy to the superuniformity Lemma 1, we present the following lemma, which contains analogous bounds under the settings of independent or positively dependent base p -values (in which case the group p -value is constructed with a Simes' p -value), and in the setting of arbitrarily dependent base p -values (in which case the group p -value can be constructed by any method—reshaped Simes, Fisher or others—as long as it is a valid p -value).

LEMMA 2 (Group-level superuniformity lemma). *Let $g \in \mathcal{H}_0^{\text{grp}}$ be a null group, that is, $A_g \subseteq \mathcal{H}_0$. Let P_{A_g} denote the p -values in this group, $P_{A_g} = (P_j)_{j \in A_g}$, and let P_{-A_g} denote the remaining p -values, $P_{-A_g} = (P_j)_{j \notin A_g}$.*

(a) *If $f : [0, 1]^n \rightarrow [0, \infty)$ is a nonincreasing function, and the base p -values P_1, \dots, P_n are independent, then*

$$\mathbb{E} \left[\frac{\mathbf{1}\{\text{Simes}_w(P_{A_g}) \leq f(P)\}}{f(P)} \mid P_{-A_g} \right] \leq 1.$$

(b) *If $f : [0, 1]^n \rightarrow [0, \infty)$ is a nonincreasing function, and the base p -values P_1, \dots, P_n are positively dependent (PRDS), then*

$$\mathbb{E} \left[\frac{\mathbf{1}\{\text{Simes}_w(P_{A_g}) \leq f(P)\}}{f(P)} \right] \leq 1.$$

(c) *If the base p -values P_1, \dots, P_n are arbitrarily dependent, then for any constant $c > 0$, any reshaping function β , and any function $f : [0, 1]^n \rightarrow [0, \infty)$, we have*

$$\mathbb{E} \left[\frac{\mathbf{1}\{T(P_{A_g}) \leq c\beta(f(P))\}}{cf(P)} \right] \leq 1,$$

where $T : [0, 1]^{|A_g|} \rightarrow [0, 1]$ is any valid group p -value; that is, any function with the property that $T(P_{A_g})$ is superuniform whenever g is null.

(d) *Let g_1, \dots, g_k be a set of k possibly overlapping null groups, meaning $A_{g_1}, \dots, A_{g_k} \subseteq \mathcal{H}_0$, and S_1, \dots, S_k represent the corresponding Simes' p -values. If $f : [0, 1]^n \rightarrow [0, \infty)$ is a nonincreasing function, and the base p -values P_1, \dots, P_n are positively dependent (PRDS), then*

$$\mathbb{E} \left[\frac{\mathbf{1}\{\text{Simes}(S_1, \dots, S_k) \leq f(P)\}}{f(P)} \right] \leq 1.$$

The proof of this lemma relies on Lemma 1 and can be found in Section 6.2. We remark that statement (d) is different from statement (b) applied to the null group $g = \bigcup_{i=1}^k g_i$; indeed, in statement (d), the arguments to the Simes' procedure are themselves Simes' p -values, and not the original base p -values. If desired, statement (d) can be further bootstrapped to apply to the root of an entire tree of null groups, where each internal node stores the Simes' p -value calculated on its children.

As an aside, one may wonder whether the Simes' p -values are themselves positively dependent (PRDS), given that they satisfy a superuniformity lemma much like the PRDS p -values. We have neither been able to prove nor disprove such a claim, and it may be of independent interest to do so.

4.3. An inverse binomial lemma for adaptivity with weights. The following lemma is required for the proof of adaptivity with weights; more specifically, we use it to bound the expected inverse of the doubly weighted null-proportion estimate.

LEMMA 3 (Inverse binomial lemma). *Given a vector $a \in [0, 1]^d$, constant $b \in [0, 1]$, and Bernoulli variables $Z_i \stackrel{\text{i.i.d.}}{\sim} \text{Bernoulli}(b)$, the weighted sum $Z := 1 + \sum_{i=1}^d a_i Z_i$ satisfies*

$$(4.2) \quad \frac{1}{1 + b \sum_{i=1}^d a_i} \leq \mathbb{E} \left[\frac{1}{Z} \right] \leq \frac{1}{b(1 + \sum_{i=1}^d a_i)}.$$

Since $\mathbb{E}[Z] = 1 + b \sum_{i=1}^d a_i$, the lower bound on $\mathbb{E}[1/Z]$ follows by Jensen's inequality. We include this bound to provide context for the upper bound on $\mathbb{E}[1/Z]$, whose proof can be found in Section 6.3. When $a_i = 1$ for all i and $b = 1$, the claim follows by a standard property of binomial distributions, as described in Benjamini et al. [6].

With these three lemmas in place, we now turn to the proof of the main theorem in the next section.

5. Proof of Theorem 1. In order to be able to handle all four cases of the theorem, we define a function $\gamma^{(m)}$ to be the identity if we are not using reshaping (theorem statements (a), (b)), or $\gamma^{(m)} = \beta^{(m)}$ if we are using reshaping (theorem statements (c), (d)). We also let $\hat{\pi}^{(m)} = 1$ and $\lambda^{(m)} = 1$ if we are not using adaptivity (theorem statements (b), (c)), or let $\hat{\pi}^{(m)}$ be defined as in equation (2.2) where adaptivity is used (theorem statements (a), (d)).

Fix any partition m . Since $\Pr\{P_i = 0\} = 0$ for any $i \in \mathcal{H}_0$ by assumption, we assume that $P_i \neq 0$ for any $i \in \mathcal{H}_0$ without further mention; this assumption then implies that if $g \in \hat{\mathcal{S}}^{(m)}(\hat{k}^{(1)}, \dots, \hat{k}^{(M)})$ for some null group $g \in \mathcal{H}_0^{(m)}$, we must

have $\widehat{k}^{(m)} > 0$. We can then calculate

$$\begin{aligned}
 \text{FDP}_u^{(m)}(\widehat{k}^{(1)}, \dots, \widehat{k}^{(M)}) &= \frac{\sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} \mathbf{1}\{g \in \widehat{\mathcal{S}}^{(m)}(\widehat{k}^{(1)}, \dots, \widehat{k}^{(M)})\}}{\sum_{g \in [G^{(m)}]} u_g^{(m)} \mathbf{1}\{g \in \widehat{\mathcal{S}}^{(m)}(\widehat{k}^{(1)}, \dots, \widehat{k}^{(M)})\}} \\
 &\leq \frac{\sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} \mathbf{1}\{g \in \widehat{\mathcal{S}}^{(m)}(\widehat{k}^{(1)}, \dots, \widehat{k}^{(M)})\}}{\widehat{k}^{(m)}} \\
 &\leq \frac{\sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} \mathbf{1}\{g \in \widehat{\mathcal{S}}_{\text{init}}^{(m)}(\widehat{k}^{(1)}, \dots, \widehat{k}^{(M)})\}}{\widehat{k}^{(m)}} \\
 &= \frac{\sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} \mathbf{1}\{P_g^{(m)} \leq \min\{w_g^{(m)} \frac{\alpha^{(m)} \gamma^{(m)}(\widehat{k}^{(m)})}{\widehat{\pi}^{(m)} G^{(m)}}, \lambda^{(m)}\}\}}{\widehat{k}^{(m)}},
 \end{aligned}$$

where the first inequality follows by definition (3.5) of the feasible set $\widehat{\mathcal{K}}$, the second follows since $\widehat{\mathcal{S}}^{(m)}(\vec{k}) \subseteq \widehat{\mathcal{S}}_{\text{init}}^{(m)}(\vec{k})$ for any \vec{k} by definition, and the last step uses the definition of $\widehat{\mathcal{S}}_{\text{init}}^{(m)}(\vec{k})$ in (3.2) (without reshaping, for theorem statements (a), (b), or with reshaping for theorem statements (c), (d)).

Multiplying the numerator and denominator of each term by $\frac{\alpha^{(m)} w_g^{(m)}}{G^{(m)}}$, and taking expectations on both sides, it follows that

$$\begin{aligned}
 \text{FDR}_u^{(m)} &\leq \frac{\alpha^{(m)}}{G^{(m)}} \sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)} \\
 (5.1) \quad &\times \mathbb{E} \left[\frac{\mathbf{1}\{P_g^{(m)} \leq \min\{\frac{w_g^{(m)} \alpha^{(m)} \gamma^{(m)}(\widehat{k}^{(m)})}{\widehat{\pi}^{(m)} G^{(m)}}, \lambda^{(m)}\}\}}{\frac{w_g^{(m)} \alpha^{(m)} \widehat{k}^{(m)}}{G^{(m)}}} \right].
 \end{aligned}$$

With these calculations in place, we now prove the four statements of the theorem. Given the suggestive form of the above expression, it is natural to anticipate the use of the two superuniformity lemmas.

Theorem statement (a). Define the function $f_g^{(m)}$ that maps the vector P to $\frac{w_g^{(m)} \alpha^{(m)} \widehat{k}^{(m)}}{\widehat{\pi}^{(m)} G^{(m)}}$. Note that $f_g^{(m)}$ is a nonincreasing function of P , since $\widehat{k}^{(m)}$ is a nonincreasing function of P by definition of our procedure, while $\widehat{\pi}^{(m)}$ is a nondecreasing function of P . We also define the quantity

$$(5.2) \quad \widehat{\pi}_{-g}^{(m)} := \frac{|u^{(m)} w^{(m)}|_\infty + \sum_{h \neq g} u_h^{(m)} w_h^{(m)} \mathbf{1}\{P_h^{(m)} > \lambda^{(m)}\}}{G^{(m)}(1 - \lambda^{(m)})}.$$

Returning to expression (5.1), we may then deduce that

$$\begin{aligned}
 \text{FDR}_u^{(m)} &\leq \frac{\alpha^{(m)}}{G^{(m)}} \sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)} \mathbb{E} \left[\frac{\mathbf{1}\{P_g^{(m)} \leq \min\{\frac{w_g^{(m)} \alpha^{(m)} \hat{k}^{(m)}}{\hat{\pi}^{(m)} G^{(m)}}, \lambda^{(m)}\}\}}{\frac{\hat{\pi}^{(m)} w_g^{(m)} \alpha^{(m)} \hat{k}^{(m)}}{\hat{\pi}^{(m)} G^{(m)}}} \right] \\
 &= \frac{\alpha^{(m)}}{G^{(m)}} \sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)} \mathbb{E} \left[\mathbf{1}\{P_g^{(m)} \leq \lambda^{(m)}\} \cdot \frac{\mathbf{1}\{P_g^{(m)} \leq f_g^{(m)}(P)\}}{\hat{\pi}^{(m)} f_g^{(m)}(P)} \right] \\
 (5.3) \quad &\stackrel{(i)}{\leq} \frac{\alpha^{(m)}}{G^{(m)}} \sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)} \mathbb{E} \left[\frac{\mathbf{1}\{P_g^{(m)} \leq f_g^{(m)}(P)\}}{\hat{\pi}_{-g}^{(m)} f_g^{(m)}(P)} \right],
 \end{aligned}$$

where inequality (i) holds because the event $P_g^{(m)} \leq \lambda^{(m)}$ implies $\hat{\pi}^{(m)} = \hat{\pi}_{-g}^{(m)}$. Conditioning on $P_{-A_g^{(m)}}$ for each group g in expression (5.3), we get

$$\begin{aligned}
 \text{FDR}_u^{(m)} &\leq \frac{\alpha^{(m)}}{G^{(m)}} \sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)} \mathbb{E} \left[\frac{\mathbf{1}\{P_g^{(m)} \leq f_g^{(m)}(P)\}}{\hat{\pi}_{-g}^{(m)} f_g^{(m)}(P)} \right] \\
 &= \frac{\alpha^{(m)}}{G^{(m)}} \sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)} \mathbb{E} \left[\mathbb{E} \left[\frac{\mathbf{1}\{P_g^{(m)} \leq f_g^{(m)}(P)\}}{\hat{\pi}_{-g}^{(m)} f_g^{(m)}(P)} \mid P_{-A_g^{(m)}} \right] \right] \\
 &\stackrel{(ii)}{=} \frac{\alpha^{(m)}}{G^{(m)}} \sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)} \mathbb{E} \left[\frac{1}{\hat{\pi}_{-g}^{(m)}} \mathbb{E} \left[\frac{\mathbf{1}\{P_g^{(m)} \leq f_g^{(m)}(P)\}}{f_g^{(m)}(P)} \mid P_{-A_g^{(m)}} \right] \right] \\
 &\leq \frac{\alpha^{(m)}}{G^{(m)}} \sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)} \mathbb{E} \left[\frac{1}{\hat{\pi}_{-g}^{(m)}} \right],
 \end{aligned}$$

where equality (ii) holds because $\hat{\pi}_{-g}^{(m)}$ is a function of only the p -values outside of group g , that is, of $P_{-A_g^{(m)}}$, while the last inequality holds by Lemma 2(a).

Finally, observe that independence between the different groups of partition m implies that the indicator variables $\mathbf{1}\{P_h^{(m)} > \lambda^{(m)}\}$ are independent Bernoullis with probabilities $\leq 1 - \lambda^{(m)}$ of success. Thus, as a consequence of Lemma 3, we can prove that

$$(5.4) \quad \mathbb{E} \left[\frac{1}{\hat{\pi}_{-g}^{(m)}} \right] \leq \frac{G^{(m)}}{\sum_{h \in \mathcal{H}_0^{(m)}} u_h^{(m)} w_h^{(m)}}.$$

To establish property (5.4), let $b := (1 - \lambda^{(m)})$, $d := |\mathcal{H}_0^{(m)}| - 1$, and define

$$Z := 1 + \sum_{h \in \mathcal{H}_0^{(m)}, h \neq g} a_h \mathbf{1}\{P_h^{(m)} > \lambda^{(m)}\} \quad \text{where } a_h = \frac{u_h^{(m)} w_h^{(m)}}{|u^{(m)} \cdot w^{(m)}|_\infty}.$$

Since $Z \leq \frac{G^{(m)}(1 - \lambda^{(m)})}{|u^{(m)} \cdot w^{(m)}|_\infty} \hat{\pi}_{-g}^{(m)}$ as the right-hand side expression sums over more indices than the left, applying Lemma 3 guarantees that

$$\begin{aligned} & \mathbb{E} \left[\frac{|u^{(m)} \cdot w^{(m)}|_\infty}{G^{(m)}(1 - \lambda^{(m)}) \hat{\pi}_{-g}^{(m)}} \right] \\ & \leq \mathbb{E} \left[\frac{1}{Z} \right] \leq \frac{|u^{(m)} \cdot w^{(m)}|_\infty}{(1 - \lambda^{(m)}) (|u^{(m)} \cdot w^{(m)}|_\infty + \sum_{h \in \mathcal{H}_0^{(m)}, h \neq g} u_h^{(m)} w_h^{(m)})}. \end{aligned}$$

Some simple algebra then leads to property (5.4).

Plugging (5.4) back into our bounds on FDR, we finally obtain

$$\begin{aligned} \text{FDR}_u^{(m)} & \leq \frac{\alpha^{(m)}}{G^{(m)}} \sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)} \mathbb{E} \left[\frac{1}{\hat{\pi}_{-g}^{(m)}} \right] \\ & \leq \frac{\alpha^{(m)}}{G^{(m)}} \sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)} \frac{G^{(m)}}{\sum_{h \in \mathcal{H}_0^{(m)}} u_h^{(m)} w_h^{(m)}} \\ & \leq \alpha^{(m)}. \end{aligned}$$

Theorem statement (b). The proof of statement (b) follows the same steps as for (a), but without the need to condition on $P_{-A_g^{(m)}}$, since we do not use adaptivity.

Define the function $f_g^{(m)}(P) = \frac{w_g^{(m)} \alpha^{(m)} \hat{k}^{(m)}}{G^{(m)}}$. Then $f_g^{(m)}$ is a nonincreasing function of P , since $\hat{k}^{(m)}$ is a nonincreasing function of P .

Returning to (5.1), as in the proof of statement (a), we calculate

$$\text{FDR}_u^{(m)} \leq \frac{\alpha^{(m)}}{G^{(m)}} \sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)} \mathbb{E} \left[\frac{\mathbf{1}\{P_g^{(m)} \leq f_g^{(m)}(P)\}}{f_g^{(m)}(P)} \right].$$

By Lemma 2(b), we know that $\mathbb{E} \left[\frac{\mathbf{1}\{P_g^{(m)} \leq f_g^{(m)}(P)\}}{f_g^{(m)}(P)} \right] \leq 1$ and, therefore,

$$\text{FDR}_u^{(m)} \leq \frac{\alpha^{(m)}}{G^{(m)}} \sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)},$$

as claimed.

Theorem statement (c). We now turn to proving the method under reshaping. Define $f_g^{(m)}(P) = \hat{k}^{(m)}$, and define constant $c_g^{(m)} = \frac{w_g^{(m)} \alpha^{(m)}}{G^{(m)}}$. Returning to (5.1), as before, we calculate

$$\text{FDR}_u^{(m)} \leq \frac{\alpha^{(m)}}{G^{(m)}} \sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)} \mathbb{E} \left[\frac{\mathbf{1}\{P_g^{(m)} \leq c_g^{(m)} \cdot \beta^{(m)}(f_g^{(m)}(P))\}}{c_g^{(m)} \cdot f_g^{(m)}(P)} \right].$$

By Lemma 2(c), we know that $\mathbb{E} \left[\frac{\mathbf{1}\{P_g^{(m)} \leq c_g^{(m)} \cdot \beta^{(m)}(f_g^{(m)}(P))\}}{c_g^{(m)} \cdot f_g^{(m)}(P)} \right] \leq 1$ since $P_g^{(m)}$ is assumed to be superuniform for any null group $g \in \mathcal{H}_0^{(m)}$. Therefore,

$$\text{FDR}_u^{(m)} \leq \frac{\alpha^{(m)}}{G^{(m)}} \sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)}.$$

Theorem statement (d). The proof of part (d) combines the calculations of part (a) (where adaptivity is used) with part (c) (where reshaping is used). Define $f_g^{(m)} = \hat{k}^{(m)}$ and $c_g^{(m)} = \frac{w_g^{(m)} \alpha^{(m)}}{\hat{\pi}_{-g}^{(m)} G^{(m)}}$, where $\hat{\pi}_{-g}^{(m)}$ is defined as in equation (5.2) from part (a). Note that $c_g^{(m)}$ is no longer a constant, but nonetheless, proceeding as in part (a), we can calculate

$$\text{FDR}_u^{(m)} \leq \frac{\alpha^{(m)}}{G^{(m)}} \sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)} \mathbb{E} \left[\frac{\mathbf{1}\{P_g^{(m)} \leq c_g^{(m)} \cdot \beta^{(m)}(f_g^{(m)}(P))\}}{\hat{\pi}_{-g}^{(m)} \cdot c_g^{(m)} \cdot f_g^{(m)}(P)} \right].$$

Next, we condition on the p -values outside the group $A_g^{(m)}$:

$$\begin{aligned} \text{FDR}_u^{(m)} &\leq \frac{\alpha^{(m)}}{G^{(m)}} \sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)} \\ &\quad \times \mathbb{E} \left[\mathbb{E} \left[\frac{\mathbf{1}\{P_g^{(m)} \leq c_g^{(m)} \cdot \beta^{(m)}(f_g^{(m)}(P))\}}{\hat{\pi}_{-g}^{(m)} \cdot c_g^{(m)} \cdot f_g^{(m)}(P)} \mid P_{-A_g^{(m)}} \right] \right] \\ &= \frac{\alpha^{(m)}}{G^{(m)}} \sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)} \\ &\quad \times \mathbb{E} \left[\frac{1}{\hat{\pi}_{-g}^{(m)}} \cdot \mathbb{E} \left[\frac{\mathbf{1}\{P_g^{(m)} \leq c_g^{(m)} \cdot \beta^{(m)}(f_g^{(m)}(P))\}}{c_g^{(m)} \cdot f_g^{(m)}(P)} \mid P_{-A_g^{(m)}} \right] \right], \end{aligned}$$

where the last step holds since $\hat{\pi}_{-g}^{(m)}$ is a function of $P_{-A_g^{(m)}}$.

Finally, we apply Lemma 2(c) to show that each of these conditional expected values is ≤ 1 . Of course, the subtlety here is that we must condition on $P_{-A_g^{(m)}}$.

To do so, note that, after fixing $P_{-A_g^{(m)}}$, the function $f_g^{(m)}(P)$ can be regarded as a function of only the remaining unknowns (i.e., of $P_{A_g^{(m)}}$), and is still nonincreasing, the value $c_g^{(m)}$ is now a constant; and $P_g^{(m)} = T_g^{(m)}(P_{A_g^{(m)}})$ is indeed superuniform since, due to the independence of $P_{A_g^{(m)}}$ from $P_{-A_g^{(m)}}$, its distribution has not changed. Therefore, we can apply Lemma 2(c) (with the random vector $P_{A_g^{(m)}}$ in place of P , while $P_{-A_g^{(m)}}$ is treated as constant), to see that

$$\mathbb{E}\left[\frac{\mathbf{1}\{P_g^{(m)} \leq c_g^{(m)} \cdot \beta^{(m)}(f_g^{(m)}(P))\}}{c_g^{(m)} \cdot f_g^{(m)}(P)} \mid P_{-A_g^{(m)}}\right] \leq 1$$

and, therefore,

$$\text{FDR}_u^{(m)} \leq \frac{\alpha^{(m)}}{G^{(m)}} \sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)} \mathbb{E}\left[\frac{1}{\hat{\pi}_{-g}^{(m)}}\right].$$

Finally, we need to bound $\hat{\pi}_{-g}^{(m)}$. As in the proof of part (a), we see that the indicator variables $\mathbf{1}\{P_h^{(m)} > \lambda^{(m)}\}$ are independent, since $P_h^{(m)} = T_h^{(m)}(P_{A_h^{(m)}})$, and the sets of p -values $P_{A_h^{(m)}}$ are assumed to be independent from each other. Furthermore, since $T_h^{(m)}(P_{A_h^{(m)}})$ is assumed to be a valid p -value, that is, superuniform for any $h \in \mathcal{H}_0^{(m)}$, this means that the variable $\mathbf{1}\{P_h^{(m)} > \lambda^{(m)}\}$ is Bernoulli with chance $\geq 1 - \lambda^{(m)}$ of success. Therefore, the bound (5.4) calculated in the proof of part (a) holds here as well, and so

$$\text{FDR}_u^{(m)} \leq \frac{\alpha^{(m)}}{G^{(m)}} \sum_{g \in \mathcal{H}_0^{(m)}} u_g^{(m)} w_g^{(m)} \frac{G^{(m)}}{\sum_{h \in \mathcal{H}_0^{(m)}} u_h^{(m)} w_h^{(m)}} = \alpha^{(m)}.$$

This concludes the proof of all four parts of Theorem 1.

6. Proofs of supporting lemmas. In this section, we collect the proofs of some supporting lemmas.

6.1. Proof of superuniformity Lemma 1. Lemma 1 follows directly from earlier work [1, 10]. Statement (a) with inequality (but not with equality) follows as a special case of (b), since independence is a special case of positive dependence, and the distribution of a null P_i does not change on conditioning on an independent set of p -values. Statement (c) was proved also by Blanchard and Roquain [10]. We now prove the statements (a), (d).

Statement (a). We prove the first part of Lemma 1, under the assumptions that the function $P \mapsto f(P)$ satisfies the leave-one-out property with respect to in-

dex i , and that P_i is uniformly distributed and is independent of the remaining p -values. Since $\Pr\{P_i = 0\} = 0$, we ignore this possibility in the following calculations. Since f satisfies the LOOP condition, we have

$$\frac{\mathbf{1}\{P_i \leq f(P)\}}{f(P)} = \frac{\mathbf{1}\{P_i \leq f(P^{i \rightarrow 0})\}}{f(P^{i \rightarrow 0})}.$$

This can be seen by separately considering what happens when the numerator on the left-hand side is zero or one.

Since P^{-i} determines $f(P^{i \rightarrow 0})$, it immediately follows that

$$\begin{aligned} \mathbb{E}\left[\frac{\mathbf{1}\{P_i \leq f(P)\}}{f(P)} \mid P^{-i}\right] &= \mathbb{E}\left[\frac{\mathbf{1}\{P_i \leq f(P^{i \rightarrow 0})\}}{f(P^{i \rightarrow 0})} \mid P^{-i}\right] \\ &= \frac{\Pr\{P_i \leq f(P^{i \rightarrow 0}) \mid f(P^{i \rightarrow 0})\}}{f(P^{i \rightarrow 0})} \\ &= 1, \end{aligned}$$

where the last step follows since f has range $[0, 1]$, and P_i is uniformly distributed and is independent of $P^{i \rightarrow 0}$; therefore, we may deduce that $\Pr\{P_i \leq f(P^{i \rightarrow 0}) \mid f(P^{i \rightarrow 0})\} = f(P^{i \rightarrow 0})$. This concludes the proof of the superuniformity lemma under independence and uniformity.

Statement (d). For each $\ell = 1, \dots, m$, let ν_ℓ be a probability measure on $[0, \infty)$ chosen such that $\beta_\ell(k) = \beta_{\nu_\ell}(k) = \int_{x=0}^k x \, d\nu_\ell(x)$, as in the definition of a reshaping function. Let $X_\ell \sim \nu_\ell$ be drawn independently for each $\ell = 1, \dots, m$, and let ν be the probability measure on $[0, \infty)$ corresponding to the distribution of $Z = \prod_{\ell=1}^m X_\ell$. Then

$$\begin{aligned} c \cdot \prod_{\ell=1}^m \beta_\ell(f_\ell(P)) &= c \cdot \prod_{\ell=1}^m \left(\int_{x_\ell=0}^{f_\ell(P)} x_\ell \, d\nu_\ell(x_\ell) \right) \\ &= c \cdot \int_{x_1=0}^{\infty} \cdots \int_{x_m=0}^{\infty} \left(\prod_{\ell=1}^m x_\ell \cdot \mathbf{1}\{x_\ell \leq f_\ell(P)\} \right) d\nu_m(x_m) \cdots d\nu_1(x_1) \\ &= c \cdot \mathbb{E} \left[\prod_{\ell=1}^m (X_\ell \cdot \mathbf{1}\{X_\ell \leq f_\ell(P)\}) \right] \\ &= c \cdot \mathbb{E} [Z \cdot \mathbf{1}\{X_1 \leq f_1(P), \dots, X_m \leq f_m(P)\}] \\ &\leq c \cdot \mathbb{E} \left[Z \cdot \mathbf{1}\left\{Z \leq \prod_{\ell=1}^m f_\ell(P)\right\} \right] \\ &= c \cdot \int_{z=0}^{\prod_{\ell=1}^m f_\ell(P)} z \, d\nu(z) = c \cdot \beta_\nu \left(\prod_{\ell=1}^m f_\ell(P) \right). \end{aligned}$$

Therefore,

$$\mathbb{E} \left[\frac{\mathbf{1}\{P_i \leq c \cdot \prod_{\ell=1}^m \beta_\ell(f_\ell(P))\}}{c \cdot \prod_{\ell=1}^m f_\ell(P)} \right] \leq \mathbb{E} \left[\frac{\mathbf{1}\{P_i \leq c \cdot \beta_v(\prod_{\ell=1}^m f_\ell(P))\}}{c \cdot \prod_{\ell=1}^m f_\ell(P)} \right] \leq 1,$$

where the last step holds by Lemma 1(c).

6.2. Proof of group superuniformity Lemma 2. First, note that the proof of Lemma 2(c) is straightforward, by applying Lemma 1(c). More precisely, define an augmented vector $P' = (P_1, \dots, P_n, T(P_{A_g})) \in [0, 1]^{n+1}$, and define a function $f'(P') := f(P_1, \dots, P_n) = f(P)$. Since $T(P_{A_g})$ is assumed to be superuniform (since $g \in \mathcal{H}_0^{\text{grp}}$ is a null group), this means that $P'_{n+1} = T(P_{A_g})$ is superuniform, that is, index $n+1$ is a null p -value, in the augmented vector of p -values P' . Then applying Lemma 1(c), with P' and f' in place of P and f , and with index $i = n+1$, yields the desired bound.

Lemma 2(a) is simply a special case of Lemma 2(b) since independence is a special case of positive dependence, and conditioning on an independent set of p -values P_{-A_g} does not change the distribution of P_{A_g} .

For Lemma 2(b), our proof strategy is to reduce this statement into a form where Lemma 1(b) becomes applicable. (Note that we cannot simply take the approach of our proof of Lemma 2(c), because if we define an augmented vector of p -values $P' = (P_1, \dots, P_n, \text{Simes}_w(P_{A_g}))$, we do not know if this vector is positively dependent—specifically, whether P' is PRDS on entry $P'_{n+1} = \text{Simes}_w(P_{A_g})$.)

With this aim in mind, let $\hat{k}_g \in \{0, \dots, n_g\}$ be the number of discoveries made by the BH_w procedure when run on the p -values within group g at level $f(P)$. Using the connection between the Simes' test and the BH procedure, we may write

$$\mathbf{1}\{P_g \leq f(P)\} = \mathbf{1}\{\hat{k}_g > 0\} = \frac{\hat{k}_g}{\hat{k}_g} = \frac{\sum_{i \in A_g} \mathbf{1}\{P_i \leq \frac{w_i \hat{k}_g f(P)}{n_g}\}}{\hat{k}_g}$$

since for the BH_w procedure at level $f(P)$, the i th p -value P_i will be rejected if and only if $P_i \leq \frac{w_i \hat{k}_g f(P)}{n_g}$. Hence, we may conclude that

$$\frac{\mathbf{1}\{P_g \leq f(P)\}}{f(P)} = \frac{\sum_{i \in A_g} \mathbf{1}\{P_i \leq \frac{w_i \hat{k}_g f(P)}{n_g}\}}{\hat{k}_g f(P)} = \frac{1}{n_g} \sum_{i \in A_g} w_i \frac{\mathbf{1}\{P_i \leq \tilde{f}_g(P)\}}{\tilde{f}_g(P)},$$

where we have defined $\tilde{f}_g(P) := \frac{w_i \hat{k}_g f(P)}{n_g}$.

Taking expectations on both sides and applying Lemma 1(b) immediately proves Lemma 2(b). (Specifically, we know that $P \mapsto \hat{k}_g$ is a nonincreasing function of P , and $P \mapsto f(P)$ is also assumed to be nonincreasing; therefore, \tilde{f}_g is also nonincreasing in P .)

Given that Lemma 2(b) is proved, the proof of Lemma 2(d) follows exactly the same argument as above, except that in the very last equation, P_i is replaced by S_i , and Lemma 2(b) is invoked in place of Lemma 1(b).

6.3. Proof of inverse-binomial Lemma 3. The lower bound follows immediately from Jensen's inequality, since $\mathbb{E}[Z] = 1 + b \sum_{i=1}^d a_i$. We split the argument for the upper bound into three cases.

Case 1: Integer weights. First, suppose that all the weights a_i are integers, that is, $a_i \in \{0, 1\}$ for all i . In this case, we have $Z \sim 1 + \text{Binomial}(k, b)$, where k is the number of weights a_i that are equal to 1. A simple calculation shows that

$$\begin{aligned} \mathbb{E}\left[\frac{1}{1 + \text{Binomial}(k, b)}\right] &= \sum_{z=0}^k \frac{1}{1+z} \binom{k}{z} b^z (1-b)^{k-z} \\ &= \frac{1}{b(1+k)} \sum_{z=0}^k \binom{k+1}{z+1} b^{z+1} (1-b)^{(k+1)-(z+1)} \\ &= \frac{1}{b(1+k)} \cdot \Pr\{\text{Binomial}(k+1, b) \leq k\} \\ &\leq \frac{1}{b(1+k)} = \frac{1}{b(1 + \sum_i a_i)}. \end{aligned}$$

Case 2: One noninteger weight. Suppose that exactly one of the weights a_i is a noninteger. Without loss of generality, we can take $a_1 = \dots = a_k = 1$, $a_{k+1} = c$, $a_{k+2} = \dots = a_n = 0$, for some $k \in \{0, \dots, n-1\}$ and some $c \in (0, 1)$. Let $A = Z_1 + \dots + Z_{k+1} \sim \text{Binomial}(k+1, b)$, and $Y = Z_{k+1} \sim \text{Bernoulli}(b)$. Note that $\Pr\{Y = 1 \mid A\} = \frac{A}{1+k}$. Then

$$\begin{aligned} \mathbb{E}\left[\frac{1}{Z}\right] &= \mathbb{E}\left[\frac{1}{1 + A - (1-c)Y}\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\frac{1}{1 + A - (1-c)Y} \mid A\right]\right] \\ &= \mathbb{E}\left[\frac{1}{1+A} \cdot \Pr\{Y=0 \mid A\} + \frac{1}{c+A} \cdot \Pr\{Y=1 \mid A\}\right] \\ &= \mathbb{E}\left[\frac{1}{1+A} + \left(\frac{1}{c+A} - \frac{1}{1+A}\right) \cdot \Pr\{Y=1 \mid A\}\right] \\ &= \mathbb{E}\left[\frac{1}{1+A} + \frac{1-c}{(c+A)(1+A)} \cdot \frac{A}{1+k}\right] \\ &\leq \mathbb{E}\left[\frac{1}{1+A} + \frac{1-c}{(c+1+k)(1+A)} \cdot \frac{1+k}{1+k}\right], \end{aligned}$$

where the inequality holds since $\frac{A}{c+A} \leq \frac{1+k}{1+k+c}$ because $0 \leq A \leq k+1$. Simplifying, we get

$$\begin{aligned}\mathbb{E}\left[\frac{1}{Z}\right] &\leq \mathbb{E}\left[\frac{1}{1+A}\right] \cdot \frac{2+k}{1+k+c} \leq \frac{1}{b(2+k)} \cdot \frac{2+k}{1+k+c} \\ &= \frac{1}{b(1+k+c)} = \frac{1}{b(1+\sum_i a_i)},\end{aligned}$$

where the inequality uses the fact that $\mathbb{E}\left[\frac{1}{1+\text{Binomial}(k+1,b)}\right] \leq \frac{1}{b(2+k)}$ as calculated in Case 1.

Case 3: General case. Now suppose that there are at least two noninteger weights, $0 < a_i \leq a_j < 1$. Let $C = \sum_{\ell \neq i,j} a_\ell Z_\ell$, then $Z = 1 + C + a_i Z_i + a_j Z_j$. Let $\alpha = \min\{a_i, 1 - a_j\} > 0$. Then

$$\begin{aligned}\mathbb{E}\left[\frac{1}{Z} \mid C\right] &= b^2 \cdot \frac{1}{1+C+a_i+a_j} + b(1-b) \cdot \frac{1}{1+C+a_i} \\ &\quad + b(1-b) \cdot \frac{1}{1+C+a_j} + (1-b)^2 \cdot \frac{1}{C} \\ &\leq b^2 \cdot \frac{1}{1+C+a_i+a_j} + b(1-b) \cdot \frac{1}{1+C+(a_i-\alpha)} \\ &\quad + b(1-b) \cdot \frac{1}{1+C+(a_j+\alpha)} + (1-b)^2 \cdot \frac{1}{C},\end{aligned}$$

where the inequality follows from a simple calculation using the assumption that $\alpha \leq a_i \leq a_j \leq 1 - \alpha$. Now, define a new vector of weights \tilde{a} where $\tilde{a}_i = a_i - \alpha$, $\tilde{a}_j = a_j + \alpha$ and $\tilde{a}_\ell = a_\ell$ if $\ell \notin \{i, j\}$. Defining $\tilde{Z} = 1 + \sum_\ell \tilde{a}_\ell Z_\ell$, the above calculation proves that $\mathbb{E}[\frac{1}{Z}] \leq \mathbb{E}[\frac{1}{\tilde{Z}}]$ (by marginalizing over C).

Note that $\sum_i a_i = \sum_i \tilde{a}_i$, but \tilde{a}_i has (at least) one fewer noninteger weight. Repeating this process inductively, we see that we can reduce to the case where there is at most one noninteger weight (i.e., Case 1 or Case 2). This proves the lemma.

7. Proof of propositions about **p-filter**.

7.1. Proof of “maximum-corner” Proposition 1. For each m , by definition of $\hat{k}^{(m)}$, there is some $k_1^{(m)}, \dots, k_{m-1}^{(m)}, k_{m+1}^{(m)}, \dots, k_M^{(m)}$ such that

$$(7.1) \quad (k_1^{(m)}, \dots, k_{m-1}^{(m)}, \hat{k}^{(m)}, k_{m+1}^{(m)}, \dots, k_M^{(m)}) \in \hat{\mathcal{K}}.$$

Thus, for each $m' \neq m$, $\hat{k}^{(m')} \geq k_{m'}^{(m)}$ by definition of $\hat{k}^{(m')}$. Then

$$\begin{aligned}\hat{\mathcal{S}}(k_1^{(m)}, \dots, k_{m-1}^{(m)}, \hat{k}^{(m)}, k_{m+1}^{(m)}, \dots, k_M^{(m)}) \\ \subseteq \hat{\mathcal{S}}(\hat{k}^{(1)}, \dots, \hat{k}_{m-1}, \hat{k}^{(m)}, \hat{k}_{m+1}, \dots, \hat{k}^{(M)}),\end{aligned}$$

because $\widehat{\mathcal{S}}(k^{(1)}, \dots, k^{(M)})$ is a nondecreasing function of $(k^{(1)}, \dots, k^{(M)})$, and this immediately implies that

$$\begin{aligned} & \widehat{\mathcal{S}}^{(m)}(k_1^{(m)}, \dots, k_{m-1}^{(m)}, \widehat{k}^{(m)}, k_{m+1}^{(m)}, \dots, k_M^{(m)}) \\ & \subseteq \widehat{\mathcal{S}}^{(m)}(\widehat{k}^{(1)}, \dots, \widehat{k}_{m-1}^{(m)}, \widehat{k}^{(m)}, \widehat{k}_{m+1}^{(m)}, \dots, \widehat{k}^{(M)}). \end{aligned}$$

Therefore, for each layer m ,

$$\sum_{g \in \widehat{\mathcal{S}}^{(m)}(\widehat{k}^{(1)}, \dots, \widehat{k}^{(m)})} u_g^{(m)} \geq \sum_{g \in \widehat{\mathcal{S}}^{(m)}(k_1^{(m)}, \dots, k_{m-1}^{(m)}, \widehat{k}^{(m)}, k_{m+1}^{(m)}, \dots, k_M^{(m)})} u_g^{(m)} \geq \widehat{k}^{(m)},$$

where the second inequality holds by observation (7.1), and by definition of $\widehat{\mathcal{K}}$ as the set of feasible vectors. Since this holds for all m , this proves that $(\widehat{k}^{(1)}, \dots, \widehat{k}^{(M)})$ is itself a feasible vector, and hence $(\widehat{k}^{(1)}, \dots, \widehat{k}^{(M)}) \in \widehat{\mathcal{K}}$.

7.2. Proof of “halting” Proposition 2. First, we introduce some notation: let $(k_{(s)}^{(1)}, \dots, k_{(s)}^{(M)})$ be the vector after the s th pass through the algorithm. We prove that $k_{(s)}^{(m)} \geq \widehat{k}^{(m)}$ for all m, s , by induction. At initialization, $k_{(0)}^{(m)} = G^{(m)} \geq \widehat{k}^{(m)}$ for all m . Now suppose that $k_{(s-1)}^{(m)} \geq \widehat{k}^{(m)}$ for all m ; we now show that $k_{(s)}^{(m)} \geq \widehat{k}^{(m)}$ for all m . To do this, consider the m th layer of the s th pass through the algorithm. Before this stage, we have vectors $k_{(s)}^{(1)}, \dots, k_{(s)}^{(s-1)}, k_{(s-1)}^{(m-1)}, k_{(s-1)}^{(m)}, k_{(s-1)}^{(m+1)}, \dots, k_{(s-1)}^{(M)}$, and we now update $k_{(s)}^{(m)}$. Applying induction also to this inner loop, and assuming that $k_{(s)}^{(m')} \geq \widehat{k}^{(m')}$ for all $m' = 1, \dots, m-1$, we can now prove that $k_{(s)}^{(m)} \geq \widehat{k}^{(m)}$. By definition of the algorithm,

$$(7.2) \quad \begin{aligned} k_{(s)}^{(m)} &= \max_{k'^{(m)} \in \{0, 1, \dots, G^{(m)}\}} \left\{ k'^{(m)} : \right. \\ & \left. \sum_{g \in \widehat{\mathcal{S}}^{(m)}(k_{(s)}^{(1)}, \dots, k_{(s)}^{(m-1)}, k'^{(m)}, k_{(s-1)}^{(m+1)}, \dots, k_{(s-1)}^{(M)})} u_g^{(m)} \geq k'^{(m)} \right\}. \end{aligned}$$

Since $k_{(s)}^{(m')} \geq \widehat{k}^{(m')}$ for all $m' = 1, \dots, m-1$, and $k_{(s-1)}^{(m')} \geq \widehat{k}^{(m')}$ for all $m' = m+1, \dots, M$, we have

$$\begin{aligned} & \sum_{g \in \widehat{\mathcal{S}}^{(m)}(k_{(s)}^{(1)}, \dots, k_{(s)}^{(m-1)}, \widehat{k}^{(m)}, k_{(s-1)}^{(m+1)}, \dots, k_{(s-1)}^{(M)})} u_g^{(m)} \\ & \geq \sum_{g \in \widehat{\mathcal{S}}^{(m)}(\widehat{k}^{(1)}, \dots, \widehat{k}^{(m-1)}, \widehat{k}^{(m)}, \widehat{k}^{(m+1)}, \dots, \widehat{k}^{(M)})} u_g^{(m)}, \end{aligned}$$

since $\widehat{\mathcal{S}}^{(m)}(\vec{k})$ is a nondecreasing function of \vec{k} by definition. The right-hand side of this expression is in turn $\geq \widehat{k}^{(m)}$ by definition of $(\widehat{k}^{(1)}, \dots, \widehat{k}^{(M)})$ being feasible.

Therefore, $\widehat{k}^{(m)}$ is in the feasible set for equation (7.2), and so we must have $k_{(s)}^{(m)} \geq \widehat{k}^{(m)}$. By induction, this is then true for all s, m , as desired.

Now suppose that the algorithm stabilizes at $(k_1^{(s)}, \dots, k_M^{(s)})$, after s full passes. After completing the m th layer of the last pass through the algorithm, we had vectors $k_{(s)}^{(1)}, \dots, k_{(s)}^{(m)}, k_{(s-1)}^{(m+1)}, \dots, k_{(s-1)}^{(M)}$; however, since the algorithm stops after the s th pass, this means that $k_{(s-1)}^{(m')} = k_{(s)}^{(m')}$ for all m' . Using this observation in the definition of $k_{(s)}^{(m)}$, we see that

$$\sum_{g \in \widehat{S}^{(m)}(k_{(s)}^{(1)}, \dots, k_{(s)}^{m-1}, k_{(s)}^{(m)}, k_{(s)}^{(m+1)}, \dots, k_{(s)}^{(M)})} u_g^{(m)} \geq k_{(s)}^{(m)}.$$

This means that $(k_{(s)}^{(1)}, \dots, k_{(s)}^{(M)}) \in \widehat{\mathcal{K}}$, and so $k_{(s)}^{(m)} \leq \widehat{k}^{(m)}$ for all m by the definition of $\widehat{k}^{(1)}, \dots, \widehat{k}^{(m)}$ and Proposition 1. But by the induction above, we also know that $k_{(s)}^{(m)} \geq \widehat{k}^{(m)}$ for all m, s , thus completing the proof.

8. Discussion and extensions. The procedures that we have analyzed and generalized do not fully cover the huge literature on FDR-controlling procedures. For example, *p*-filter is a generalized multidimensional step-up procedure, but much work has also been done on alternative styles of procedures, such as step-down, step-up-down and multistep methods. For example, Benjamini and Liu [8] propose step-down procedures that control FDR under independence. Also, procedures by Benjamini and Liu [7] and Romano and Shaikh [33] provably control FDR under arbitrary dependence, with Gavrilov et al. [15] extending them to adaptive control under independence. Two-step adaptive procedures have been analyzed in Benjamini et al. [6] under independence, and by Blanchard and Roquain [11] under dependence. Different methods of incorporating weights into such procedures have also been studied; cf. a different notion of the weighted Simes' *p*-value proposed by Benjamini and Hochberg [4].

The superuniformity lemmas (Lemma 1 and, in the grouped setting, Lemma 2), can be used to quickly prove FDR control under dependence for many of the above procedures, and may be a useful tool for designing new multiple testing procedures in broader settings.⁴ For example, it has been used to derive a decentralized procedure for FDR control on sensor networks [27] and a sequential algorithm for FDR control on directed acyclic graphs [28]. This lemma was also used to derive a “post-selection BH procedure” [13]: if a set $S \subseteq [n]$ of hypotheses was selected by the user in an arbitrary monotone data-dependent manner (see footnote 2), one way to find a subset $T \subseteq S$ that controls the FDR is to run BH on S at level $\tilde{\alpha} := \alpha|S|/n$.

⁴An analog of the superuniformity lemma has also been discovered in the online FDR setting and has proved useful for designing new algorithms [21, 29, 30].

Indeed,

$$\begin{aligned} \text{FDR} &= \mathbb{E} \left[\frac{\sum_{i \in S \cap \mathcal{H}_0} \mathbf{1}\{P_i \leq \tilde{\alpha} \frac{|T|}{|S|}\}}{|T|} \right] \\ &\leq \sum_{i \in \mathcal{H}_0} \frac{\alpha}{n} \cdot \mathbb{E} \left[\frac{\mathbf{1}\{P_i \leq \alpha \frac{|T|}{n}\}}{\alpha \frac{|T|}{n}} \right] \leq \alpha \frac{|\mathcal{H}_0|}{n}. \end{aligned}$$

Notice that the post-selection BH procedure reduces to BH in the absence of selection, that is when $S = [n]$. As another particularly simple but striking example, consider the following novel “structured BH procedure.”⁵ Suppose we wish to insist that only certain subsets of $[n]$ are allowed to be rejected; let $\mathcal{K} \subseteq 2^{[n]}$ be the set of such allowed rejection sets (these could be determined by known logical constraints or structural requirements). Then, if the p -values are positively dependent, we may reject the largest set $T \in \mathcal{K}$ such that all its p -values are less than $\alpha|T|/n$. Completely equivalently, one can define $\widehat{\text{FDP}}(S) = \frac{n \cdot \max_{i \in S} P_i}{|S|}$ and reject the largest set $T \in \mathcal{K}$ such that $\widehat{\text{FDP}}(T) \leq \alpha$. This procedure controls FDR under positive dependence due to a trivial one-line proof using Lemma 1:

$$\begin{aligned} \text{FDR} &= \mathbb{E} \left[\frac{\sum_{i \in T \cap \mathcal{H}_0} \mathbf{1}\{P_i \leq \alpha \frac{|T|}{n}\}}{|T|} \right] \\ &\leq \sum_{i \in \mathcal{H}_0} \frac{\alpha}{n} \cdot \mathbb{E} \left[\frac{\mathbf{1}\{P_i \leq \alpha \frac{|T|}{n}\}}{\alpha \frac{|T|}{n}} \right] \leq \alpha \frac{|\mathcal{H}_0|}{n}. \end{aligned}$$

Again, notice that the structured BH procedure reduces to BH in the absence of structural constraints, that is, when $\mathcal{K} = 2^{[n]}$. Of course, except for special structured settings, the *largest* set $T \in \mathcal{K}$ may not be efficiently computable in general. In this case, any monotone set $T \in \mathcal{K}$ such that $\widehat{\text{FDP}}(T) \leq \alpha$ may be chosen, and FDR control will be maintained, and heuristics can be used to find large monotone sets. (In both the above examples, one may use reshaping to control for arbitrary dependence or for nonmonotonicity.)

While there exist works that can incorporate a single layer of groups [20], these often provide guarantees only for the finest partition. Alternative error metrics have been discussed by Benjamini and Bogomolov [2], who devise a way to take a single partition of groups into account and control a *selective* FDR. This idea has been extended by Peterson et al. [26] and Bogomolov et al. [12] to partitions that form a hierarchy (i.e., a tree). However, none of these aforementioned papers have been extended to handle arbitrary nonhierarchical partitions, leftover or overlapping groups, both sets of weights, adaptivity or reshaping. Recently, Katsevich and Sabatti [23] derived a knockoff p -filter that extended the work of Barber and

⁵This procedure was independently discovered recently by Katsevich et al. [24], along with several other substantial extensions.

Ramdas [1] in two ways: it allows the group p -values to be formed by procedures other than Simes' (like this paper), and it can use knockoff statistics instead of p -values. In both settings, it provides FDR control at a constant (between 1 and 2) times the target FDR. While their work can handle arbitrary nonhierarchical partitions (since it uses the same `p-filter` framework) along with knockoff statistics, it also does not handle null-proportion adaptivity, both sets of weights, reshaping, leftover or overlapping groups and so on. We believe that many of the algorithmic ideas and proof techniques (especially the lemmas) introduced here may generalize to these related works, and could be an avenue for future work.

Finally, as a last very general extension, it was recently noted by Katsevich and Ramdas [22] that the `p-filter` framework can arbitrarily “stack” together different layers, where each layer uses a different type of FDR-controlling algorithm (ordered testing, knockoffs, online algorithms, interactive algorithms and so on), and the `p-filter` framework can still be used as a wrapper to ensure internal consistency along with FDR control.

Acknowledgments. We thank Wenyu Chen for helping implement the new `p-filter` algorithm. We thank Eugene Katsevich, Etienne Roquain, Aditya Gun-tuboyina and Fanny Yang for relevant discussions.

SUPPLEMENTARY MATERIAL

Supplement to “A unified treatment of multiple testing with prior knowledge using the `p-filter`” (DOI: [10.1214/18-AOS1765SUPP](https://doi.org/10.1214/18-AOS1765SUPP); .pdf). Contains details on dotfractions, generalized Simes tests for the global null and the LOOP property.

REFERENCES

- [1] BARBER, R. F. and RAMDAS, A. (2017). The p -filter: Multilayer false discovery rate control for grouped hypotheses. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **79** 1247–1268. [MR3689317](#)
- [2] BENJAMINI, Y. and BOGOMOLOV, M. (2014). Selective inference on multiple families of hypotheses. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **76** 297–318. [MR3153943](#)
- [3] BENJAMINI, Y. and HOCHBERG, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. Ser. B* **57** 289–300. [MR1325392](#)
- [4] BENJAMINI, Y. and HOCHBERG, Y. (1997). Multiple hypotheses testing with weights. *Scand. J. Stat.* **24** 407–418. [MR1481424](#)
- [5] BENJAMINI, Y. and HOCHBERG, Y. (2000). On the adaptive control of the false discovery rate in multiple testing with independent statistics. *J. Educ. Behav. Stat.* **25** 60–83.
- [6] BENJAMINI, Y., KRIEGER, A. M. and YEKUTIELI, D. (2006). Adaptive linear step-up procedures that control the false discovery rate. *Biometrika* **93** 491–507. [MR2261438](#)
- [7] BENJAMINI, Y. and LIU, W. (1999). A distribution-free multiple test procedure that controls the false discovery rate. Technical report, Tel Aviv Univ.
- [8] BENJAMINI, Y. and LIU, W. (1999). A step-down multiple hypotheses testing procedure that controls the false discovery rate under independence. *J. Statist. Plann. Inference* **82** 163–170. [MR1736441](#)

- [9] BENJAMINI, Y. and YEKUTIELI, D. (2001). The control of the false discovery rate in multiple testing under dependency. *Ann. Statist.* **29** 1165–1188. [MR1869245](#)
- [10] BLANCHARD, G. and ROQUAIN, E. (2008). Two simple sufficient conditions for FDR control. *Electron. J. Stat.* **2** 963–992. [MR2448601](#)
- [11] BLANCHARD, G. and ROQUAIN, É. (2009). Adaptive false discovery rate control under independence and dependence. *J. Mach. Learn. Res.* **10** 2837–2871. [MR2579914](#)
- [12] BOGOMOLOV, M., PETERSON, C. B., BENJAMINI, Y. and SABATTI, C. (2017). Testing hypotheses on a tree: New error rates and controlling strategies. Preprint. Available at [arXiv:1705.07529](#).
- [13] BRZYSKI, D., PETERSON, C. B., SOBCZYK, P., CANDÈS, E. J., BOGDAN, M. and SABATTI, C. (2017). Controlling the rate of GWAS false discoveries. *Genetics* **205** 61–75.
- [14] GABRIEL, K. R. (1969). Simultaneous test procedures—Some theory of multiple comparisons. *Ann. Math. Stat.* **40** 224–250. [MR0240931](#)
- [15] GAVRILOV, Y., BENJAMINI, Y. and SARKAR, S. K. (2009). An adaptive step-down procedure with proven FDR control under independence. *Ann. Statist.* **37** 619–629. [MR2502645](#)
- [16] GENOVESE, C. R., ROEDER, K. and WASSERMAN, L. (2006). False discovery control with p -value weighting. *Biometrika* **93** 509–524. [MR2261439](#)
- [17] HEARD, N. A. and RUBIN-DELANCHY, P. (2018). Choosing between methods of combining p -values. *Biometrika* **105** 239–246. [MR3768879](#)
- [18] HOCHBERG, Y. and BENJAMINI, Y. (1990). More powerful procedures for multiple significance testing. *Stat. Med.* **9** 811–818.
- [19] HOCHBERG, Y. and LIBERMAN, U. (1994). An extended Simes’ test. *Statist. Probab. Lett.* **21** 101–105. [MR1309882](#)
- [20] HU, J. X., ZHAO, H. and ZHOU, H. H. (2010). False discovery rate control with groups. *J. Amer. Statist. Assoc.* **105** 1215–1227. [MR2752616](#)
- [21] JAVANMARD, A. and MONTANARI, A. (2018). Online rules for control of false discovery rate and false discovery exceedance. *Ann. Statist.* **46** 526–554. [MR3782376](#)
- [22] KATSEVICH, E. and RAMDAS, A. (2018). Towards “simultaneous selective inference”: Post-hoc bounds on the false discovery proportion. Preprint. Available at [arXiv:1803.06790](#).
- [23] KATSEVICH, E. and SABATTI, C. (2019). Multilayer knockoff filter: Controlled variable selection at multiple resolutions. *Ann. Appl. Stat.* **13** 1–33. [MR3937419](#)
- [24] KATSEVICH, E., SABATTI, C. and BOGOMOLOV, M. (2018). Controlling FDR while highlighting distinct discoveries. Preprint. Available at [arXiv:1809.01792](#).
- [25] LEHMANN, E. L. (1966). Some concepts of dependence. *Ann. Math. Stat.* **37** 1137–1153. [MR0202228](#)
- [26] PETERSON, C. B., BOGOMOLOV, M., BENJAMINI, Y. and SABATTI, C. (2016). Many phenotypes without many false discoveries: Error controlling strategies for multitrait association studies. *Genet. Epidemiol.* **40** 45–56.
- [27] RAMDAS, A., CHEN, J., WAINWRIGHT, M. J. and JORDAN, M. I. (2017). QuTE: Decentralized multiple testing on sensor networks with false discovery rate control. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)* 6415–6421. IEEE, New York.
- [28] RAMDAS, A., CHEN, J., WAINWRIGHT, M. J. and JORDAN, M. I. (2019). A sequential algorithm for false discovery rate control on directed acyclic graphs. *Biometrika* **106** 69–86. [MR3912384](#)
- [29] RAMDAS, A., YANG, F., WAINWRIGHT, M. J. and JORDAN, M. I. (2017). Online control of the false discovery rate with decaying memory. In *Advances in Neural Information Processing Systems* 5655–5664.
- [30] RAMDAS, A., ZRNIC, T., WAINWRIGHT, M. J. and JORDAN, M. I. (2018). SAFFRON: An adaptive algorithm for online control of the false discovery rate. In *Proceedings of the 35th International Conference on Machine Learning* 4286–4294.

- [31] RAMDAS, A. K., BARBER, R. F., WAINWRIGHT, M. J. and JORDAN, M. I. (2019). Supplement to “A unified treatment of multiple testing with prior knowledge using the p-filter.” DOI:[10.1214/18-AOS1765SUPP](https://doi.org/10.1214/18-AOS1765SUPP).
- [32] ROMANO, J. P., SHAIKH, A. and WOLF, M. (2011). Consonance and the closure method in multiple testing. *Int. J. Biostat.* **7** Art. 12, 27. [MR2775079](#)
- [33] ROMANO, J. P. and SHAIKH, A. M. (2006). On stepdown control of the false discovery proportion. In *Optimality. Institute of Mathematical Statistics Lecture Notes—Monograph Series* **49** 33–50. IMS, Beachwood, OH. [MR2337829](#)
- [34] SARKAR, S. K. (2008). On methods controlling the false discovery rate. *Sankhyā* **70** 135–168. [MR2551809](#)
- [35] SARKAR, S. K. (2008). Two-stage stepup procedures controlling FDR. *J. Statist. Plann. Inference* **138** 1072–1084. [MR2384506](#)
- [36] SARKAR, T. K. (1969). Some lower bounds of reliability. Technical report, Stanford Univ. [MR2618992](#)
- [37] SEEGER, P. (1968). A note on a method for the analysis of significances en masse. *Technometrics* **10** 586–593.
- [38] SIMES, R. J. (1986). An improved Bonferroni procedure for multiple tests of significance. *Biometrika* **73** 751–754. [MR0897872](#)
- [39] SONNEMANN, E. (1982). Allgemeine Lösungen multipler Testprobleme. Institut für Mathematische Statistik und Versicherungslehre, Univ. Bern.
- [40] SONNEMANN, E. (2008). General solutions to multiple testing problems. *Biom. J.* **50** 641–656. [MR2542333](#)
- [41] SONNEMANN, E. and FINNER, H. (1988). Vollständigkeitssätze für multiple testprobleme. In *Multiple Hypothesenprüfung/Multiple Hypotheses Testing* 121–135. Springer, Berlin.
- [42] STOREY, J. D. (2002). A direct approach to false discovery rates. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **64** 479–498. [MR1924302](#)
- [43] STOREY, J. D., TAYLOR, J. E. and SIEGMUND, D. (2004). Strong control, conservative point estimation and simultaneous conservative consistency of false discovery rates: A unified approach. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **66** 187–205. [MR2035766](#)
- [44] STOUTER, S. A., SUCHMAN, E. A., DEVINNEY, L. C., STAR, S. A. and WILLIAMS, R. M. JR. (1949). *The American Soldier: Adjustment During Army Life*, Vol. 1. *Studies in Social Psychology in World War II*.
- [45] TUKEY, J. (1953). *The Problem of Multiple Comparisons: Introduction and Parts A, B, and C*. Princeton Univ., Princeton, NJ.
- [46] TUKEY, J. W. (1994). *The Collected Works of John W. Tukey. Vol. VIII: Multiple Comparisons: 1948–1983*. CRC Press, New York. [MR1263027](#)
- [47] VOVK, V. and WANG, R. (2018). Combining p-values via averaging. Preprint. Available at [arXiv:1212.4966v4](https://arxiv.org/abs/1212.4966v4).

A. K. RAMDAS
 DEPARTMENT OF STATISTICS AND DATA SCIENCE
 CARNEGIE MELLON UNIVERSITY
 5000 FORBES AVE
 PITTSBURGH, PENNSYLVANIA 15213
 USA
 E-MAIL: aramdas@cmu.edu

R. F. BARBER
 DEPARTMENT OF STATISTICS
 UNIVERSITY OF CHICAGO
 5801 S ELLIS AVENUE
 CHICAGO, ILLINOIS 60637
 USA
 E-MAIL: rina@uchicago.edu

M. J. WAINWRIGHT
M. I. JORDAN
DEPARTMENTS OF STATISTICS AND EECS
253 CORY HALL
UNIVERSITY OF CALIFORNIA, BERKELEY
BERKELEY, CALIFORNIA 94720-1770
USA
E-MAIL: wainwrig@berkeley.edu
jordan@stat.berkeley.edu