

This article was downloaded by: [Tel Aviv University]

On: 15 July 2012, At: 03:19

Publisher: Taylor & Francis

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Journal of the American Statistical Association

Publication details, including instructions for authors and subscription information:
<http://amstat.tandfonline.com/loi/uasa20>

Hierarchical False Discovery Rate-Controlling Methodology

Daniel Yekutieli

Daniel Yekutieli is Lecturer, Department of Statistics and Operations Research, School of Mathematical Sciences, Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, Israel . This work was supported by a grant from the Israeli Science Foundation and by NIH grant DA15087. The author thanks two anonymous referees, the associate editor, and joint editors for comments and suggestions that greatly improved the quality of the article.

Version of record first published: 01 Jan 2012

To cite this article: Daniel Yekutieli (2008): Hierarchical False Discovery Rate-Controlling Methodology, Journal of the American Statistical Association, 103:481, 309-316

To link to this article: <http://dx.doi.org/10.1198/016214507000001373>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://amstat.tandfonline.com/page/terms-and-conditions>

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae, and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand, or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

Hierarchical False Discovery Rate–Controlling Methodology

Daniel YEKUTIELI

We discuss methodology for controlling the false discovery rate (FDR) in complex large-scale studies that involve testing multiple families of hypotheses; the tested hypotheses are arranged in a tree of disjoint subfamilies, and the subfamilies of hypotheses are hierarchically tested by the Benjamini and Hochberg FDR-controlling (BH) procedure. We derive an approximation for the multiple family FDR for independently distributed test statistics: q , the level at which the BH procedure is applied, times the number of families tested plus the number of discoveries, divided by the number of discoveries plus 1. We provide a universal bound for the FDR of the discoveries in the new hierarchical testing approach, $2 \times 1.44 \times q$, and demonstrate in simulations that when the data has an hierarchical structure the new testing approach can be considerably more powerful than the BH procedure.

KEY WORDS: False-discovery rate; Hierarchical testing.

1. INTRODUCTION

The Benjamini and Hochberg (1995) false discovery rate (FDR)-controlling procedure (hereinafter the BH procedure) has been successfully applied in overcoming the multiplicity problem in many applications, but the task performed was limited: test a family of hypotheses, determined in advance, while controlling the FDR. Yekutieli et al. (2006) introduced a hierarchical testing methodology for controlling the FDR in multiple families of hypotheses and developed a general framework, based on the hierarchical approach, for controlling the FDR in complex experiments. This article formally defines the hierarchical FDR testing approach, studies its performance in simulations, and derives bounds and approximations for the FDR of the various types of discoveries produced by the hierarchical testing methodology. The exposition begins with a demonstration of use of this new testing approach in three different types of applications.

1.1 Analysis of Microarray Data

Yekutieli et al. (2006) applied hierarchical testing to a microarray experiment that included expression levels of 25,600 genes in 5 mice brain regions of 10 inbred mice strains. A 2-way ANOVA, with strain and brain region main effects, was fitted for each gene to identify genes with strain expression differences. The researchers also were interested in testing the interaction terms to locate areas in the brains and strains with abnormal expression levels. In the standard FDR approach, the two questions are addressed separately; when applied at level .05 to test the strain effect for the 25,600 genes, the BH procedure yielded 957 discoveries, but when applied to test the 1.2 million interaction terms, it yielded no discoveries. In the hierarchical approach, the discovery of genes with significant strain effects is considered the initial question for each gene, and localizing the effect to specific strains and brain regions are considered follow-up questions for genes with significant strain effects. Separately applying the .05 BH procedure in each of the 957 families of interactions corresponding to a strain effect discovery yielded a total of 170 discoveries. In this case

the researcher may be interested in four types of discoveries: the entire set of 1,127 discoveries, the 957 strain effect discoveries, the 170 interaction discoveries, and the most detailed information for each gene—either interaction discoveries, or strain effect discoveries—for which no interaction discoveries are found. Reiner et al. (2007) applied hierarchical testing to the same data, but in that study the expression of each gene with significant strain effect was tested for association with a series of 17 behavioral traits.

1.2 Quantitative Trait Loci Analysis

Quantitative trait loci (QTL) are genetic loci affecting quantitative traits. Weller, Song, Heyen, Lewin, and Ron (1998) suggested searching for QTL by applying the BH procedure to p values testing for linkage between a quantitative trait and a series of genetic markers. But because linkage is not specific to a genetic marker, FDR control over marker discoveries does not imply control over the proportion of false QTL discoveries (see the supplemental report available at http://www.amstat.org/publications/jasa/supplemental_materials). To control for the occurrence of false QTL discoveries, it is necessary to directly test for the existence of a QTL within a specified genomic region (Zheng 1994), preferably pinpointing the QTL to the smallest possible interval on the chromosome; however, due to a decreased power to discover QTLs and the increase in the number of hypotheses considered, searching at too high a resolution may result in failure to discover QTL.

The suggested solution is a hierarchical multiresolution search for QTL. In the first level, apply the BH procedure to test all chromosome-level hypotheses. On chromosomes in which a discovery has been made, search for QTL by testing the two half-chromosome-level hypotheses. If a half-chromosome discovery is made, then proceed to the quarter-chromosome level. As long as a discovery is made, continue searching at a higher resolution level. Adopting this testing strategy allows one to adaptively test hypotheses at the highest resolution possible. Because the discovery of high-resolution discoveries makes the initial low-resolution discoveries irrelevant, in this case, the researcher may be interested only in the highest-resolution discovery in each genomic region.

Daniel Yekutieli is Lecturer, Department of Statistics and Operations Research, School of Mathematical Sciences, Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, Israel (E-mail: yekutieli@post.tau.ac.il). This work was supported by a grant from the Israeli Science Foundation and by NIH grant DA15087. The author thanks two anonymous referees, the associate editor, and joint editors for comments and suggestions that greatly improved the quality of the article.

© 2008 American Statistical Association
Journal of the American Statistical Association
March 2008, Vol. 103, No. 481, Theory and Methods
DOI 10.1198/016214507000001373

1.3 Log-Linear Analysis of Mouse Behavior Data

Kafkafi, Yekutieli, Yarowsky, and Elmer (2006) presented an algorithm for discovering behavior patterns that differentiate between mutant and wild-type rats. A filmed session of exploratory behavior is divided into a series of 54,000 frames. Nine behavioral relevant endpoints are computed for each time frame, transformed into three five-level ordinal scales and summarized in a nine-way contingency table. The algorithm then scans these immense, sparse contingency tables for patterns with significant frequency differences between mutant and wild-type rats. The main challenge is to work at the highest resolution level and still avoid overfitting; thus Kafkafi et al. (2006) scanned only the 50,675 subsets of cells determined by combinations of up to 4 ordinal scales.

In the hierarchical approach, log-linear models are fitted for the mutant rat and wild-type rat contingency tables. The log-linear models are constructed hierarchically. At the first stage, the .05 BH procedure is applied to test the terms in the main-effects model; at the second stage, the .05 BH procedure is applied separately to test two-way interactions with each of the significant main effects from the first stage; in the third stage, the .05 BH procedure is applied separately to test three-way interactions with each of the significant two-way interaction; and so on. The tables are then scanned for the behavior patterns with the largest differences between the fitted values of the two log-linear models. In this case, the researcher is interested in controlling the FDR for all the terms in the models; however, the fact that hierarchical FDR model selection produces parsimonious models with small mean squared error may be a more relevant property.

1.4 Background

In their seminal article, Benjamini and Hochberg introduced the BH procedure and the FDR, a new measure for type I error in multiple testing. Since then, FDR methodology has become a very active field of research, but essentially, all of the methods apply only to a single family of hypotheses. Yekutieli and Benjamini (1999) expressed control of the FDR as an estimation problem; instead of comparing the sorted p values to a series of critical values determined by the FDR level q , estimate the FDR of a fixed rejection region test. Storey (2002, 2003) discussed a Bayesian setting for the fixed rejection region FDR and introduced the positive FDR and the q value with the following very appealing Bayesian interpretation: The conditional probability that a discovery is a false discovery given that its test statistic is in the rejection region. Efron, Tibshirani, Storey, and Tusher (2001) suggested empirical Bayes estimation of the FDR and even considered conditioning locally on the value of the test statistic, not just on the rejection region. Another estimation effort is the estimation of the proportion of true null hypotheses (Benjamini and Hochberg 2000; Storey, Taylor, and Siegmund 2004) to derive more powerful testing procedures. Genovese and Wasserman (2004) developed a framework in which the false discovery proportion (the number of false rejections divided by the number of rejections in a continuum of fixed rejection regions) is treated as a stochastic process. Benjamini and Yekutieli (2005a) generalized the FDR criterion to a measure for the validity of confidence intervals for parameters

after selection. This article presents a general framework for applying the BH procedure and controlling the FDR in a wide variety of settings not considered before.

Hierarchical modeling is widely used in the construction of complex statistical models, including the node-splitting decisions in CART (Brieman, Friedman, Olshen, and Stone 1984), and in the treatment of interaction terms in linear models. Separately correcting for multiplicity in several families of hypotheses also is not new. When the family-wise error rate (FWE) is separately controlled in several families of hypotheses, the FWE for the entire set of tested hypotheses is essentially the sum of the individual FWE levels; the usual solution is splitting the nominal FWE level, α , between the families of hypotheses (see Simonson and McIntyre 2004 for a QTL mapping example). In the FDR literature, applying the BH procedure, at level q , in each family of hypotheses has been suggested. Abramovich, Benjamini, Donoho, and Johnstone (1998) discussed the asymptotic theoretical properties of this practice. In the context of QTL mapping, Lee et al. (2002) suggested a separate FDR-controlled search for each quantitative trait. Yekutieli et al. (2006) suggested dividing the statistical analysis in complex studies into several main research directions and controlling the FDR separately in each research direction. In general, FDR control applied separately in several families does not necessarily imply FDR control for the entire study (Benjamini and Yekutieli 2005b). In this article, a bound for the FDR is computed when the BH is applied at level q to several families of hypotheses. The article's main message is that hierarchical application of the BH procedure inherently implies global FDR control for the entire set of discoveries.

Screening the null hypotheses before testing is a related method of alleviating the multiplicity practiced in recent microarray analyses (Pavlidis 2003; Letwin et al. 2006; Yekutieli et al. 2006; Reiner et al. 2007). Unlike hierarchical testing, any screening criterion can be applied to select the tested hypotheses, and all of the screened hypotheses are tested simultaneously; similar to hierarchical testing, screening must be independent of the hypotheses testing. Reiner et al. (2007) showed that the FDR is not controlled when the BH procedure is applied to test pairwise strain expression differences of genes screened by a 1-way ANOVA F test; they also showed in simulations that screening before testing offers less power than hierarchical testing.

Section 2 formally defines the hierarchical testing approach and summarizes the theoretical results. Section 3 describes the microarray analysis and signal denoising simulation studies. Section 4 is devoted to a discussion of the results. The Appendix presents the derived FDR bounds. The technical issues needed to derive the theoretical results and detailed accounts of the simulations are deferred to a supplemental report.

2. TESTING TREES OF HYPOTHESES

In the hierarchical approach, the set of tested hypotheses, H_1, \dots, H_m , is arranged in a tree with L levels. With the exception of hypotheses on the first level of the tree, which have no parent hypotheses, each hypothesis, H_i , on level $L(i) = 2, \dots, L$ is associated with a single-parent hypothesis, indexed by $Par(i)$, on level $L(i) - 1$. Let H_1, \dots, H_T denote the parent hypotheses; then we can divide the m hypotheses into $T + 1$

families: $\mathcal{T}_0 = \{H_i : L(i) = 1\}$ and $\mathcal{T}_t = \{H_i : Par(i) = t\}$ for $t = 1, \dots, T$. Here m_t and m_t^0 are the total number of hypotheses and the number of true null hypotheses in \mathcal{T}_t .

The hierarchical test of the tree of hypotheses has two elements: (a) Hypotheses in the same family are tested simultaneously, and (b) testing begins with \mathcal{T}_0 , and a family of hypotheses on higher levels of the tree is tested only if its parent hypothesis is rejected. Our concern in this work is with trees in which each family of hypotheses, \mathcal{T}_t , is tested by the BH procedure.

Definition 1. Level q BH procedure on \mathcal{T}_t :

1. Let $P_{(1)}^t \leq \dots \leq P_{(m_t)}^t$ denote the set of ordered p values corresponding to the hypotheses in \mathcal{T}_t .
2. Let $r_t = \max\{i : P_{(i)}^t \leq i \cdot q / m_t\}$.
3. If $r_t > 0$, then reject the r_t hypotheses corresponding to $P_{(1)}^t, \dots, P_{(r_t)}^t$.

We assume that the p values are independently distributed; if H_j is a true null hypothesis, then $P_j \sim U[0, 1]$. For a false null hypothesis H_j , P_j satisfies the following condition.

Condition 1. For all $0 < \alpha_1 \leq \alpha_2 \leq 1$,

$$\alpha_1 / \alpha_2 \leq \Pr(P_j \leq \alpha_1 | P_j \leq \alpha_2).$$

Condition 1 states that the conditional marginal distribution of all of the p values is uniform, or stochastically smaller than uniform. It is satisfied if the cumulative distribution function of P_j is concave (as, e.g., under the monotone likelihood ratio condition), yet it is a somewhat stronger condition than stochastically smaller than $U[0, 1]$.

The final component of FDR trees is the specification of the set of discoveries that are of interest to the investigator. The FDR is then defined as the expected proportion of false “interesting” discoveries out of the total number of “interesting” discoveries made. (This proportion is set to 0 if no discoveries are made.) In this article, three types of FDR are discussed:

1. Full-tree FDR. Interest lies in the entire set of FDR tree discoveries.
2. Level-restricted FDR. The investigator is interested only in the discoveries on a specific level of the tree. The level-restricted FDR is also a model for the FDR of the entire study when the BH procedure is applied separately to several families of hypotheses (see Example A.2).
3. Outer-nodes FDR. Interest lies in discoveries that are not parents to other discoveries, for example, in QTL analysis, the highest-resolution discovery in each genomic region.

Figure 1 presents a schematic drawing of a tree of hypotheses, and the results of the hierarchical test. The tree includes 12 hypotheses in 6 families: $\mathcal{T}_0 = \{H_1, H_2\}$, $\mathcal{T}_1 = \{H_3, H_4\}$, $\mathcal{T}_2 = \{H_5, H_6\}$, $\mathcal{T}_3 = \{H_7\}$, $\mathcal{T}_4 = \{H_8, H_9\}$, and $\mathcal{T}_5 = \{H_{10}, H_{11}, H_{12}\}$. Six discoveries are considered in the full-tree FDR. H_1 and H_2 are level-1 discoveries, H_3 and H_5 are level-2 discoveries, and H_{10} and H_{12} are level-3 discoveries. H_3 , H_{10} , and H_{12} are the outer-nodes discoveries. \mathcal{T}_4 was not tested because H_4 was not rejected. Note that outer-node discoveries are not necessarily in the leaves of the tree; for instance, H_3 is an outer-node discovery, because none of its children (only H_7 in this example) were rejected in the procedure.

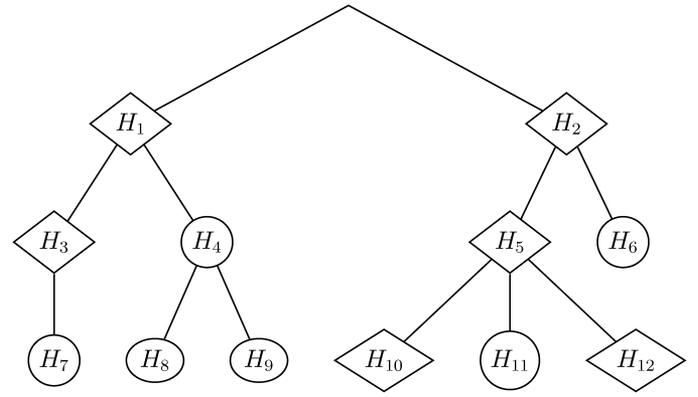


Figure 1. FDR tree schematic. Diamonds indicate null hypotheses rejected in the hierarchical testing scheme; circles are null hypotheses not rejected in the hierarchical testing scheme; and ovals are null hypotheses in families not tested in the hierarchical testing scheme.

2.1 Summary of the Results

The theoretical results in this article are derived under the assumption that the p value are independently distributed, true null hypotheses p values have $U[0, 1]$ distributions, and the false null hypotheses p values satisfy Condition 1. The FDR bounds are sums over $t = 0, \dots, T$ of FDR_t ,

$$FDR_t = E \left\{ I(\mathcal{T}_t \text{ is tested}, R > 0) \cdot \frac{V_t}{R} \right\},$$

where V_t is the number of false discoveries in \mathcal{T}_t and R is the total number of FDR tree discoveries.

To derive the bounds, FDR_t is expressed in terms of $R_t^{P_i=0}$ and $R^{P_i=0}$, the number of discoveries in \mathcal{T}_t and the total number of FDR tree discoveries given that a p value corresponding to H_i , a true null hypothesis in \mathcal{T}_t , is set to 0. (Setting $P_i = 0$ produces the conditional number of discoveries given that H_i is rejected.) Using this approach, Benjamini and Yekutieli (2001) proved that the FDR of the level- q BH procedure applied to a single family of independently distributed p values is $q \cdot m^0 / m$. Because it is not possible to aggregate the expression for FDR_t over multiple families of hypotheses, we substitute $R_t^{P_i=0}$ and $R^{P_i=0}$ with $R_t + 1$ and $R + 1$, the unconditional number of discoveries in \mathcal{T}_t and total number of FDR tree discoveries, plus 1. (Because true null hypotheses are rarely rejected in multiple-testing procedures, the conditional number of discoveries given that H_i is rejected can be approximated by the unconditional number of discoveries plus 1.) The main technical issue addressed in this article is finding a multiplicative factor, δ^* , as small as possible, yielding

$$E \frac{R_t^{P_i=0}}{R^{P_i=0}} / E \frac{R_t + 1}{R + 1} \leq \delta^*. \tag{1}$$

2.1.1 Assessment of δ^ .* We assess δ^* analytically and in simulations. At first, we consider the conditional distribution of $R_t^{P_i=1}$, the number of discoveries in \mathcal{T}_t for $P_i = 1$, given $R_t^{P_i=0}$. For $R_t^{P_i=0} = 1, \dots, m_t$, we examine the ratio in

$$R_t^{P_i=0} / E(R_t^{P_i=1} + 1 | R_t^{P_i=0}) \leq \delta^*. \tag{2}$$

We prove that for each value of $R_t^{P_i=0}$, the ratio is maximized when all of the p values have $U[0, 1]$ distributions, compute

the ratio for each value of $R_t^{P_i=0}$ under this assumption, and verify that under Condition 1, the maximal value of the ratio is slightly below 1.44 for $R_t^{P_i=0} = 4$. We then prove that (2) is a stronger condition than (1). Thus we prove that if the p values are independently distributed and Condition 1 holds, then $\delta^* = 1.44$ satisfies inequality (1).

In the first set of simulations assessing δ^* , we study the distribution of $R_t^{P_i=0}$ and the ratio in (2) for various p value distributions; in the second set of simulations, we directly assess the ratio in (1). The simulations reveal that for $q = .05$, the cases for which the ratio in (2) approaches its maximum, 1.44, are rare, and that in most cases $\delta^* \approx 1$ is sufficient to satisfy (1), whereas higher values of δ^* are needed only for testing several hundreds of hypotheses with a nearly uniform p value distribution.

2.1.2 False Discovery Rate Bounds and Approximations. With δ^* , we can phrase the main result of the article.

Proposition 1. For the three types of tree discoveries and δ^* satisfying inequality (1),

$$FDR_t \leq \delta^* \cdot q \cdot \frac{m_t^0}{m} \cdot E \left\{ I(\mathcal{T}_t \text{ is tested}) \cdot \frac{R_t + 1}{R + 1} \right\}. \quad (3)$$

The proof of this proposition, the lemmas leading to this result, and the assessment of δ^* are deferred to the supplemental report. The results given in the next paragraph are derived in the Appendix.

Summing (3) over $t = 0, \dots, T$ yields a bound for the FDR,

$$FDR \leq q \cdot \delta^* \cdot E \left\{ (\text{no. of discoveries} + \text{no. of families tested}) / (\text{no. of discoveries} + 1) \cdot \tilde{\pi}_0 \right\}, \quad (4)$$

where $\tilde{\pi}_0$ is mean of m_t^0/m_t , weighted proportionally to $R_t + 1$. Imposing the restriction of hierarchical testing yields the following universal bounds for the FDR:

1. The full tree FDR of any FDR tree is $< q \cdot \delta^* \cdot 2$.
2. The outer nodes FDR of any L -level FDR tree is $< L \cdot q \cdot \delta^* \cdot 2$.

Expression (4) also implies that if the number of discoveries greatly exceeds the number of families tested, then the three types of hierarchical FDR are approximately $q \cdot \delta^* \cdot \tilde{\pi}_0$. We suggest using the observed number of families tested and the observed number of discoveries to approximate the FDR:

$$FDR = q \cdot \delta^* \cdot (\text{observed no. of discoveries} + \text{observed no. of families tested}) / (\text{observed no. of discoveries} + 1). \quad (5)$$

Although the theoretical properties of the term on the right side of (5)—the FDR multiplier—are not clear, the simulations in Section 3.1 reveal that it can be used to approximate the FDR, especially for the level-restricted FDR, which has no universal bound. Thus, in the analysis of microarray data example discussed in the Introduction, the full-tree FDR for the entire set of 1,127 discoveries is approximately .092 [= .05 · (1,127 + 957 + 1)/(1,127 + 1)], and the FDR for the 170 interaction discoveries is approximately .330 [= .05 · (170 + 957)/(170 + 1)].

3. SIMULATION STUDIES

3.1 Hierarchical False Discovery Rate Analysis of Microarray Data

The simulations were run to model a 10,000-gene microarray experiment that compared a treatment and a control on 8 mouse strains. The first research question is to find genes whose expression level is affected by the treatment; the follow-up questions are comparisons of the treatment effect between the 28 pairs of mice strains. Gene expression levels are modeled as treatment effect, plus strain-specific treatment effect, plus correlated normal noise. In the direct approach, the 280,000 pairwise comparisons are tested simultaneously using the BH procedure. In the hierarchical approach, the data were analyzed with a two-level FDR tree tested by the .05 BH procedure. \mathcal{T}_0 includes the null hypotheses, for each gene, that the mean treatment effect (across all strains) is zero, and each hypothesis in \mathcal{T}_0 is parent to the 28 corresponding pairwise comparison hypotheses. In the hierarchical approach, four types of discoveries are considered: (a) Treatment-effect discoveries, the level-1 discoveries; (b) pairwise discoveries, the level-2 discoveries; (c) full-tree discoveries, the entire set of discoveries; and (d) outer-node discoveries, the pairwise differences plus the treatment-effect discoveries for genes for which no pairwise differences were found. Note that the addition of correlated noise and the use of pairwise comparisons induces dependence across the tree, yet each pairwise comparison is independent of the initial treatment effect test.

The results of the simulations are presented in Table 1. The mean number of discoveries is listed in column 5, and the FDR is listed in column 6, with standard errors in parentheses. The mean observed value of the FDR multiplier (the mean of the number of discoveries plus number of families tested divided by the number of discoveries plus 1) is given in column 7, with its mean absolute deviation (MAD) in parentheses.

Applied at level .05, the BH procedure yields 21–23 discoveries for sparse strain effect and dense strain effect configurations. In the sparse and small configuration, the BH procedure loses its power; at level .05, the mean number of discoveries drops to .1, and even for $q = .50$, the mean number of discoveries is only 10. For zero treatment effect and sparse strain effect, the power of the hierarchical approach was comparable to that of the direct approach, but for the dense strain-specific effect, the number of pairwise discoveries rose to 508. The performance of the hierarchical approach improved substantially for small and medium treatment effects. In the sparse and small strain-specific configuration with medium treatment effect, the hierarchical procedure yielded approximately 627 treatment discoveries and 128 pairwise discoveries—more than 12 times more discoveries than the level .50 BH procedure discoveries.

The level-1 restricted FDR was $.05 \times m_0^0/m_0$. The full-tree and outer-nodes FDRs were less than the corresponding universal bounds, and even less than $.05 \cdot \delta^*$ times the FDR multiplier ($\delta^* \leq 1.1$ in the simulations). In the first four parameter configurations, the level-2 restricted FDR (which has no universal bound) also was $< .05 \cdot \delta^*$ times the FDR multiplier; but in the last configuration, the level-2 FDR was .358 and exceeded its approximation $.329 = .05 \cdot 1.1 \cdot 5.98$. These findings indicate

Table 1. Results of hierarchical analysis of microarray data simulation study

Treat. effect	Strain effect	Analysis method	Discovery type	Number of discoveries (standard error)	FDR (standard error)	FDR multiplier (MAD)		
Zero	Sparse	BH .05	Pairwise	21.1(.6)	.044(.002)	1		
			Hierarchical	Treatment	9.6(.4)	.044(.004)	1	
		Hierarchical	Pairwise	22.8(.9)	.045(.004)	1.38(.21)		
			Full tree	32.5(1.3)	.048(.003)	1.31(.11)		
			Outer node	27.6(1.1)	.055(.004)	1.36(.13)		
	Dense	BH .05	Pairwise	21.7(.7)	.044(.002)	1		
			Hierarchical	Treatment	404.5(4.7)	.039(.001)	1	
		Hierarchical	Pairwise	508.4(5.7)	.053(.001)	1.80(.08)		
			Full tree	912.9(10.2)	.047(0)	1.44(.03)		
			Outer node	777.2(8.7)	.055(0)	1.52(.03)		
	Small	Sparse	BH .05	Pairwise	22.1(.6)	.048(.003)	1	
				Hierarchical	Treatment	157.0(2.9)	.04(.001)	1
			Hierarchical	Pairwise	269.3(4.4)	.063(.001)	1.57(.09)	
				Full tree	426.3(7.2)	.055(.001)	1.36(.03)	
Outer node				362.3(6.2)	.064(.001)	1.43(.04)		
Dense		BH .05	Pairwise	22.9(.6)	.044(.002)	1		
			Hierarchical	Treatment	1,070(5)	.04(0)	1	
		Hierarchical	Pairwise	1,159(6)	.063(0)	1.92(.07)		
			Full tree	2,229(11)	.052(0)	1.48(.02)		
			Outer node	1,910(10)	.06(0)	1.56(.02)		
Medium and small		Sparse	BH .05	Pairwise	.11(.02)	.029(.007)	1	
				BH .35	Pairwise	2.73(.2)	.314(.017)	1
					Pairwise	10.1(.8)	.463(.016)	1
		Hierarchical	BH .50	Treatment	626.8(7.2)	.039(0)	1	
	Pairwise			128.3(1.7)	.358(.003)	5.98(.91)		
	Hierarchical		Full tree	755.2(8.6)	.093(.001)	1.83(.02)		
			Outer node	702.9(8.0)	.099(.001)	1.89(.02)		

that hierarchical FDRs may increase in response to dependence across the tree. Examining the MAD of the FDR multiplier reveals that in many of the configurations, the dispersion of the FDR multiplier was relatively small. This suggests that the observed FDR multiplier values (rather than the simulation mean FDR multiplier) can be used to approximate the FDR.

3.2 Hierarchical False Discovery Rate Signal Denoising

This section compares signal denoising through hard thresholding based on two-level hierarchical testing schemes with the FDR wavelet thresholding approach introduced by Abramovich and Benjamini (1996), and demonstrates the effectiveness of incorporating prior knowledge in the design of hierarchical testing procedures. The signal consisted of K segments of 128 observations: a segment of 128 observation de-measured Doppler signal, followed by $K - 1$ segments of 128 0's. In each run of the simulation, independent $N(0, 1)$ noise was added to the signal.

Abramovich and Benjamini (1996) FDR wavelet thresholding was performed by applying the BH procedure at level .05 to the p values corresponding to the $128 \times K - 1$ wavelet coefficient. For the hierarchical FDR approach, a wavelet transform was applied separately to each of the K segments. For

$k = 1, \dots, K$, \mathcal{T}_0 included a specific wavelet coefficient null hypothesis, H_k , to test the null hypothesis that segment k is pure noise; H_k was parent to the family of 126 remaining wavelet coefficient null hypotheses corresponding to segment k . We experimented with two choices of wavelet coefficients in \mathcal{T}_0 . The *default* choice, based on the assumption that in a nonnull signal, the effect size of the lowest-resolution wavelet coefficient is large, was the lowest resolution wavelet coefficient. The *optimal* choice, based on prior information on the type of signal used, was the largest wavelet coefficient in the wavelet transform of the Doppler signal.

The simulation study included 27 signal configurations: 3 levels of signal-to-noise ratio times 9 values of K ranging from 1 to 256. Each configuration was run 3,000 times. For each run, we recorded the observed proportions of false discoveries for single-stage FDR thresholding at $q = .05$ and for hierarchical FDR thresholding also at $q = .05$ were recorded, then the inverse wavelet transform was applied, and the SSE in relation to the true signal was computed.

The simulation FDR levels for the two hierarchical procedures was approximately $.05 \cdot \tilde{\pi}_0$. The SSE values of the three FDR denoising schemes are displayed in Figure 2. The optimal hierarchical procedure had the smallest SSE in all configurations—almost constant for all values of K . The SSE

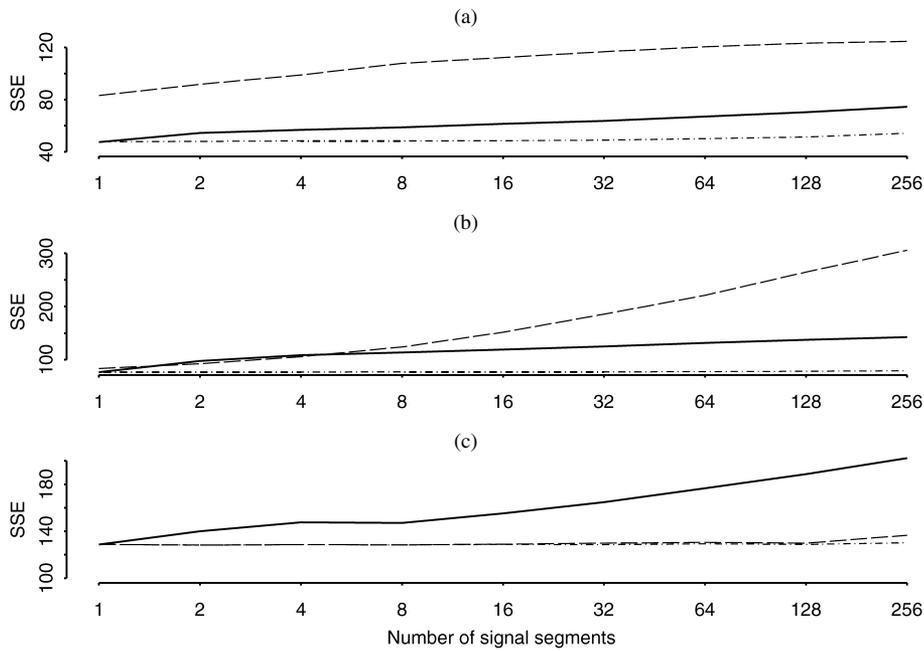


Figure 2. Simulation-based mean SSE level from the signal denoising simulation study (a) SNR = 1; (b) SNR = 2; (c) SNR = 4. The solid line represents BH level .05 FDR wavelet thresholding; the dashed line, default hierarchical level .05 FDR wavelet thresholding; the dot-and-dashed line, optimal hierarchical level .05 FDR wavelet thresholding.

of the BH procedure was the same for $K = 1$, but increased with K . For high SNR, the SSE of the default hierarchical procedure was almost as small as the SSE of the optimal procedure; for small signal-to-noise ratio, the poor choice of test statistic for H_1 made it difficult to recover the signal, resulting in high SSE.

4. DISCUSSION

Hierarchical FDR methodology can be used to control the FDR in complex large-scale studies with multiple families of hypotheses. It also is considerably more powerful than the BH procedure in sparse testing problems, providing that the data have a hierarchical structure. In the examples considered in this article, the tested hypotheses were arranged in trees of homogeneous families: either families consisting of true null hypotheses with a true null hypothesis parent or families with a large proportion of false null hypotheses and a false null hypothesis parent. Thus hierarchical testing passed over the families of true null hypotheses and adaptively applied the BH procedure to the families with a high signal-to-noise ratio. Because many of the statistical methods are hierarchical, these testing schemes are intuitive and easily to set up. The p values tested in the examples discussed in the Introduction and in the simulation studies were computed through a series of nested linear models.

Throughout the article, the assumption is made that the p values are independently distributed. But it is important to make the distinction between dependence across the tree and dependence between a test statistic and its ancestors. In general, the validity of hypotheses testing is based on the assumption that the distribution of p values corresponding to true null hypotheses is $U[0, 1]$ (or stochastically larger than $U[0, 1]$). In the hierarchical FDR approach, this further implies that dependence between a p value and any of its ancestors should not be allowed.

Thus, when tested (i.e., given that all of its ancestor hypotheses were rejected), the distribution of a p value corresponding to a true null hypotheses is still $U[0, 1]$. (Recall that this property also was needed for screening before testing.) On the other hand, dependence across the tree can be allowed; it occurred in the examples presented in the Introduction, and its effect on the FDR was evaluated in the microarray analysis simulation study. The simulations indicate that whereas the universal FDR bounds apply for the dependent test statistics studied, dependence within the families of tested hypotheses and across the tree seemed to result in higher FDR values than expected for independent test statistics. Future work is planned to study the effect of dependence across the tree on the hierarchical FDR.

APPENDIX: FALSE DISCOVERY RATE COMPUTATIONS

Let D_j denote the event where H_j is rejected in the BH procedure on $\mathcal{T}_{Par(j)}$, and let D_j^{Par} denote the event where H_j and all of its ancestors are rejected in the BH procedure on their respective family, where $D_0^{Par} = \Omega$. We can now formally define the discovery of a null hypothesis, H_j , in the three tree testing schemes.

Definition A.1.

- a. H_j is a full-tree discovery, D_j^{Par} .
- b. H_j is an outer-nodes discovery, $D_j^{Par} \cap (\bigcap_{k \in \mathcal{T}_j} \overline{D_k})$.
- c. H_j is a level- l -restricted discovery, $D_j^{Par} \cap L(j) = l$.

Let R_t and V_t denote the total number of full-tree discoveries and the number of false full-tree discoveries in \mathcal{T}_t . R is the total number of discoveries in any of the three tree testing schemes. Proposition 1 can be expressed as

$$FDR_t \leq \delta^* \cdot q \cdot \frac{m_t^0}{m_t} \cdot E \left\{ I(D_t^{Par}) \cdot \frac{R_t + 1}{R + 1} \right\}. \quad (A.1)$$

We partition the sample space into disjoint events indexed by subtrees of hypotheses K ,

$$D_K^{Par} = \left(\bigcap_{t \in K} D_t \right) \cap \left(\bigcap_{t \notin K} \bar{D}_t \right),$$

and, for a given realization of the vector of m p values, denote the realized subtree by J . As the hypotheses are tested hierarchically, for $t = 0, \dots, T$,

$$t \notin J \Rightarrow D_J^{Par} \cap D_t^{Par} = \emptyset, \quad \text{whereas } t \in J \Rightarrow D_J^{Par} \subseteq D_t^{Par}. \quad (\text{A.2})$$

A.1 Full-Tree FDR

Let $R^{Full} = \sum_{t \in J} R_t$ denote the total number of full tree discoveries. Then full tree FDR can be expressed as

$$\begin{aligned} FDR^{Full} &= E \left\{ \sum_{t=0}^T I(D_t^{Par}, R_0 > 0) \cdot \frac{V_t}{R^{Full}} \right\} \\ &= \sum_{t=0}^T E \left\{ I(D_t^{Par}, R_0 > 0) \cdot \frac{V_t}{R^{Full}} \right\} \\ &\leq \sum_{t=0}^T \delta^* \cdot q \cdot \frac{m_t^0}{m_t} \cdot E \left\{ I(D_t^{Par}) \cdot \frac{R_t + 1}{R^{Full} + 1} \right\} \end{aligned} \quad (\text{A.3})$$

$$\begin{aligned} &= \delta^* \cdot q \cdot E \left\{ \sum_{t=0}^T I(D_t^{Par}) \cdot \frac{m_t^0}{m_t} \cdot \frac{R_t + 1}{R^{Full} + 1} \right\} \\ &= \delta^* \cdot q \cdot E \left\{ \sum_J I(D_J^{Par}) \cdot \sum_{t=0}^T I(D_t^{Par}, D_J^{Par}) \cdot \frac{m_t^0}{m_t} \cdot \frac{R_t + 1}{R^{Full} + 1} \right\} \end{aligned}$$

$$= \delta^* \cdot q \cdot E \left\{ \sum_J I(D_J^{Par}) \cdot \frac{\sum_{t \in J} (R_t + 1) \cdot m_t^0 / m_t}{R^{Full} + 1} \right\} \quad (\text{A.4})$$

$$\begin{aligned} &= \delta^* \cdot q \cdot E \left\{ \sum_J I(D_J^{Par}) \cdot \frac{R^{Full} + |J|}{R^{Full} + 1} \right. \\ &\quad \left. \times \frac{\sum_{t \in J} (R_t + 1) m_t^0 / m_t}{\sum_{t \in J} (R_t + 1)} \right\}, \end{aligned} \quad (\text{A.5})$$

where the inequality in (A.3) is due to (A.1) and the equality in (A.4) is due to (A.2).

Corollary A.1. The full-tree FDR is $\delta^* \cdot q \cdot 2$.

Proof. With the exception of \mathcal{T}_0 , each $t \in J$ corresponds to the full-tree discovery of H_t ; thus $R^{Full} \geq |J| - 1 \geq 0$, and, as $m_t^0 / m_t \leq 1$,

$$\begin{aligned} FDR^{Full} &\leq \delta^* \cdot q \cdot E \left\{ \sum_J I(D_J^{Par}) \cdot \frac{R^{Full} + |J|}{R^{Full} + 1} \right\} \\ &\leq \delta^* \cdot q \cdot E \left\{ \sum_J I(D_J^{Par}) \cdot \frac{|J| - 1 + |J|}{|J| - 1 + 1} \right\} \\ &< \delta^* \cdot q \cdot 2. \end{aligned}$$

Note that R^{Full} equals $|J| - 1$ plus the number of leaf discoveries. Therefore, the universal bound in Corollary A.1 can be approached only if there is a very small proportion of leaf discoveries and if m_t^0 / m_t in the families of tested hypotheses is close to 1. On the other hand, if $R^{Full} \gg |J|$, then the expression in (A.5) is approximately $\delta^* \cdot q \cdot \tilde{\pi}_0$, where

$$\tilde{\pi}_0 = E \left\{ \sum_J I(D_J^{Par}) \frac{\sum_{t \in J} (R_t + 1) \cdot m_t^0 / m_t}{\sum_{t \in J} (R_t + 1)} \right\}.$$

A.2 Outer-Nodes FDR

Let FDR^{Outer} , R^{Outer} , and V_t^{Outer} denote the outer-nodes FDR, the total number of outer-node discoveries, and the number of false outer-node discoveries in \mathcal{T}_t . Per definition, $R^{Outer} \leq R^{Full}$ and $V_t^{Outer} \leq V_t$. In the simulations presented in this article, true null hypotheses are parents to families of true null hypotheses, because families of true null hypotheses rarely yield discoveries $V_t^{Outer} \approx V_t$, whereas R^{Outer} usually is less than R^{Full} ; thus $FDR^{Outer} > FDR^{Full}$. The bounds for FDR^{Outer} are also greater than the bound for FDR^{Full} . But in general, FDR^{Outer} also can be smaller than or equal to FDR^{Full} . To derive the bound for FDR^{Outer} , we substitute V_t^{Outer} with V_t and apply Proposition 1:

$$\begin{aligned} FDR^{Outer} &= E \left\{ \sum_{t=0}^T I(D_t^{Par}, R_0 > 0) \cdot \frac{V_t^{Outer}}{R^{Outer}} \right\} \\ &\leq E \left\{ \sum_{t=0}^T I(D_t^{Par}, R_0 > 0) \cdot \frac{V_t}{R^{Outer}} \right\} \\ &\leq \delta^* \cdot q \cdot E \left\{ \sum_J I(D_J^{Par}) \cdot \frac{\sum_{t \in J} (R_t + 1) \cdot m_t^0 / m_t}{R^{Outer} + 1} \right\} \\ &= \delta^* \cdot q \cdot E \left\{ \sum_J I(D_J^{Par}) \cdot \frac{R^{Full} + |J|}{R^{Outer} + 1} \right. \\ &\quad \left. \times \frac{\sum_{t \in J} (R_t + 1) \cdot m_t^0 / m_t}{\sum_{t \in J} (R_t + 1)} \right\}. \end{aligned} \quad (\text{A.6})$$

Because J includes all indices of full-tree discoveries that are not outer-node discoveries, $R^{Full} \leq R^{Outer} + |J|$; thus if $R^{Full} \gg |J|$, then bound in (A.6) is approximately $\delta^* \cdot q \cdot \tilde{\pi}_0$.

For the universal bound, note that R^{Outer} is always greater than or equal to the number of full-tree discoveries at each level of the tree,

$$R^{Full} + |J| \leq 2 \cdot R^{Full} \leq 2 \cdot L \cdot R^{Outer}. \quad (\text{A.7})$$

As $m_t^0 / m_t \leq 1$, combining (A.7) and (A.6) yields the following result.

Corollary A.2. The outer-nodes FDR is $< \delta^* \cdot q \cdot 2 \cdot L$.

In the following example, we consider a very tall and narrow tree in which most discoveries are parent hypotheses. Thus we see that for large L , FDR^{Full} is close to $\delta^* \cdot q \cdot \tilde{\pi}_0$, whereas FDR^{Outer} is less than half of $\delta^* \cdot q \cdot \tilde{\pi}_0$.

Example A.1. $2 \cdot L$ null hypotheses are tested in a L -level tree. For $l = 1, \dots, L$, H_l are false null hypotheses with $P_l = 0$; for $l = L + 1, \dots, 2L$, H_l are true null hypotheses with iid $U[0, 1]$ p values; and for $l = 0, \dots, L - 1$, \mathcal{T}_l includes H_{l+1} and H_{L+l+1} . The tree of hypotheses is tested at level .05. Thus $J \equiv \{0, \dots, L - 1\}$; in each family $m_t^0 / m_t = 1/2$, and $\delta^* = 2/2.05$. In both testing schemes, $V \sim \text{Binom}(L, .05)$, whereas $R^{Full} = L + V$ and $R^{Outer} = 1 + V$. For $L = 2, 3, 6$, and 10 , FDR^{Full} is .033, .036, .042, and .044, and FDR^{Outer} is .049, .073, .138, and .216. As $L \rightarrow \infty$, $FDR^{Full} \rightarrow .0476$ and $FDR^{Outer} \rightarrow 1$.

A.3 Level-Restricted False Discovery Rate

For $1 \leq l \leq L$, let $J_l = \{t : t \in J, L(t) = l - 1\}$, and denote $R^l = \sum_{t \in J_l} R_t$ the number of level- l discoveries. The level- l restricted FDR is

$$\begin{aligned} FDR^{Level=l} &= E \left\{ \sum_{t=0}^T I(D_t^{Par}, L(t) = l - 1, R_l > 0) \cdot \frac{V_t}{R^l} \right\} \end{aligned}$$

$$\begin{aligned}
&\leq \delta^* \cdot q \cdot E \left\{ \sum_J I(D_J^{Par}) \cdot \frac{\sum_{t \in J_l} (R_t + 1) \cdot m_t^0 / m_t}{\sum_{t \in J_l} R_t + 1} \right\} \\
&= \delta^* \cdot q \cdot E \left\{ \sum_J I(D_J^{Par}) \cdot \frac{R^l + |J_l|}{R^l + 1} \right. \\
&\quad \left. \times \frac{\sum_{t \in J_l} (R_t + 1) \cdot m_t^0 / m_t}{\sum_{t \in J_l} (R_t + 1)} \right\}. \tag{A.8}
\end{aligned}$$

There is no universal bound for $FDR^{Level=l}$, but if $R^l \gg |J_l|$, then the bound in (A.8) is approximately $\delta^* \cdot q \cdot \bar{\pi}_0$ (in this case $\bar{\pi}_0$ is the expected mean of m_t^0 / m_t over J_l).

Example A.2. A two level tree is tested. The level-1 hypotheses— H_1, \dots, H_T , are false null hypotheses with 0 p values. Each false null hypothesis H_t is a parent to a single true null hypothesis H_{T+t} . The number of false outer-node discoveries and false full-tree discoveries is $V \sim \text{Binom}(T, q)$; $R^{Full} = V + T$, and thus $FDR^{Full} \approx q / (1 + q)$, and $R^{Outer} = T$, and thus $FDR^{Outer} = q$. $FDR^{Level=1}$ is, by definition, 0. $FDR^{Level=2} = 1 - (1 - q)^T$ approaches 1 as $T \rightarrow \infty$.

[Received November 2005. Revised September 2007.]

REFERENCES

- Abramovich, F., and Benjamini, Y. (1996), "Adaptive Thresholding of Wavelet Coefficients," *Computational Statistics and Data Analysis*, 22, 351–361.
- Abramovich, F., Benjamini, Y., Donoho, D. L., and Johnstone, I. M. (1998), "The Amalgamation Challenge to Signal De-Noising," Research Paper RP-SOR-98-03, Tel Aviv University, Dept. of Statistics and Operations Research.
- Benjamini, Y., and Hochberg, Y. (1995), "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing," *Journal of the Royal Statistical Society*, Ser. B, 57, 289–300.
- (2000), "On the Adaptive Control of the False Discovery Rate in Multiple Testing With Independent Statistics," *Journal of Educational and Behavioral Statistics*, 25, 60–83.
- Benjamini, Y., and Yekutieli, D. (2001), "The Control of the False Discovery Rate in Multiple Testing Under Dependency," *The Annals of Statistics*, 29, 1165–1188.
- (2005a), "False Discovery Rate-Adjusted Multiple Confidence Intervals for Selected Parameters," *Journal of the American Statistical Association*, 100, 71–81.
- (2005b), "Quantitative Trait Loci Analysis Using the False Discovery Rate," *Genetics*, 171, 783–790.
- Brieman, L., Friedman, J. H., Olshen, R. A., and Stone, C. J. (1984), *Classification and Regression Trees*, Wadsworth Statistics/Probability Series (1st ed.), Wadsworth International Group/New York: Chapman & Hall.
- Efron, B., Tibshirani, R., Storey, J. D., and Tusher, V. (2001), "Empirical Bayes Analysis of a Microarray Experiment," *Journal of the American Statistical Association*, 96, 1151–1160.
- Genovese, C., and Wasserman, L. (2004), "A Stochastic Process Approach to False Discovery Control," *The Annals of Statistics*, 32, 1035–1061.
- Kafkafi, N., Yekutieli, D., Yarowsky, P., and Elmer, G., (2006), "Pattern Array: A Novel Approach for Mining Raw Behavioral Data," unpublished manuscript.
- Lee, H., Dekkers, J. C. M., Soller, M., Malek, M., Fernando, R. L., and Rothschild, M. F. (2002), "Application of the False Discovery Rate to Quantitative Trait Loci Interval Mapping With Multiple Traits," *Genetics*, 161, 905–914.
- Letwin, N. E., Kafkafi, N., Benjamini, Y., Mayo, C., Frank, B. C., Luu, T., Lee, N. H., and Elmer, G. I. (2006), "Combined Application of Behavior Genetics and Microarray Analysis to Identify Regional Expression Themes and Gene–Behavior Associations," *Journal of Neuroscience*, 26, 5277–5287.
- Pavlidis, P. (2003), "Using ANOVA for Gene Selection From Microarray Studies of the Nervous System," *Methods*, 31, 282–289.
- Reiner, A., Yekutieli, D., Letwin, N. E., Elmer, G. I., Lee, N. H., Kafkafi, N., and Benjamini, Y. (2007), "Associating Quantitative Behavioral Traits With Gene Expression in the Brain: Searching for Diamonds in the Hay," *Bioinformatics*, 23, 2239–2246.
- Simonson, K. L., and McIntyre, L. M. (2004), "Using Alpha Wisely: Improving Power to Detect Multiple QTL," *Statistical Applications in Genetics and Molecular Biology*, 3, issue 1, article 1.
- Storey, J. D. (2002), "A Direct Approach to False Discovery Rates," *Journal of the Royal Statistical Society*, Ser. B, 64, 479–498.
- (2003), "The Positive False Discovery Rate: A Bayesian Interpretation and the q Value," *The Annals of Statistics*, 31, 2013–2035.
- 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," *Journal of the Royal Statistical Society*, Ser. B, 66, 187–205.
- Weller, J. I., Song, J. Z., Heyen, D. W., Lewin, H. A., and Ron, M. (1998), "A New Approach to the Problem of Multiple Comparisons in the Genetic Dissection of Complex Traits," *Genetics*, 150, 1699–1706.
- Yekutieli, D., and Benjamini, Y. (1999), "Resampling-Based False Discovery Rate–Controlling Procedure for Dependent Test Statistics," *Journal of Statistical Planning and Inference*, 82, 171–196.
- Yekutieli, D., Reiner, A., Elmer, G. I., Kafkafi, N., Letwin, N. E., Lee, N. H., and Benjamini, Y. (2006), "Approaches to Multiplicity Issues in Complex Research in Microarray Analysis," *Statistica Neerlandica*, 60, 414–437.
- Zheng, Z. B. (1994), "Precision Mapping of Quantitative Trait Loci," *Genetics*, 136, 1457–1468.