

Genome-Wide Significance Levels and Weighted Hypothesis Testing

Kathryn Roeder and Larry Wasserman

Abstract. Genetic investigations often involve the testing of vast numbers of related hypotheses simultaneously. To control the overall error rate, a substantial penalty is required, making it difficult to detect signals of moderate strength. To improve the power in this setting, a number of authors have considered using weighted p -values, with the motivation often based upon the scientific plausibility of the hypotheses. We review this literature, derive optimal weights and show that the power is remarkably robust to misspecification of these weights. We consider two methods for choosing weights in practice. The first, external weighting, is based on prior information. The second, estimated weighting, uses the data to choose weights.

Key words and phrases: Bonferroni correction, multiple testing, weighted p -values.

1. INTRODUCTION

Testing for association between genetic variation and a complex disease typically requires scanning hundreds of thousands of genetic polymorphisms. In a multiple testing situation, such as a genome-wide association study (GWAS), the null hypothesis is rejected for any test that achieves a p -value less than a predetermined threshold (usually on the order of 10^{-8}). Data from these investigations has renewed interest in the multiple testing problem. The introduction of the false discovery rate and a procedure to control it by Benjamini and Hochberg (1995) inspired hope that this would be an effective way to control error while increasing power (Storey and Tibshirani, 2003; Sabatti, Service and Freimer, 2003). To further bolster power, recent statistical methods have been proposed that up-weight and down-weight hypotheses, based on prior likelihood of association with the phenotype (Genovese, Roeder and Wasserman, 2006; Roeder et al., 2006; Roeder, Wasserman and Devlin, 2007; Wang,

Li and Bucan, 2007). Such prior information is often available in practice.

Weighted procedures multiply the threshold by the weight w , for each test, raising the threshold when $w > 1$ and lowering it if $w < 1$. To control the overall rate of false positives, a budget must be imposed on the weighting scheme, so that the average weight is one. If the weights are informative, the procedure improves power substantially, but, if the weights are uninformative, the loss in power is usually small. Surprisingly, aside from this budget requirement, any set of nonnegative weights is valid (Genovese, Roeder and Wasserman, 2006). While desirable in some respects, this flexibility makes it difficult to select weights for a particular analysis.

The first such weighting scheme appears to be Holm (1979). Related ideas can be found in Benjamini and Hochberg (1997), Chen et al. (2000), Genovese, Roeder and Wasserman (2006), Kropf et al. (2004), Rosenthal and Rubin (1983), Schuster, Kropf and Roeder (2004), Westfall and Krishen (2001), Westfall, Kropf and Finos (2004), Blanchard and Roquain (2008) and Roquain and van de Wiel (2008), among others. Several of these approaches use data dependent weights and yet maintain familywise error control. There are, of course, other ways to improve power aside from weighting. Some notable recent approaches include Rubin, Dudoit and van der Laan (2006), Storey

Kathryn Roeder is Professor of Statistics, Department of Statistics, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA (e-mail: roeder@stat.cmu.edu).

Larry Wasserman is Professor of Statistics, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA (e-mail: wasserman@stat.cmu.edu).

(2007), Donoho and Jin (2004), Signoravitch (2006), Westfall, Krishen and Young (1998), Westfall and Soper (2001), Efron (2007) and Sun and Cai (2007). Of these, our approach is closest to Rubin, Dudoit and van der Laan (2006).

In some cases, the optimal weights can be estimated from the data. An approach developed by Westfall, Kropf and Finos (2004) utilizes quadratic forms to construct such weights; however, this approach assumes the individual measurements are normally distributed. This approach is suited to applications such as microarray data for which the observations are approximately normally distributed. We are interested in applications such as tests for genetic association. In this setting the individual observations are discrete, but the test statistics are approximately normally distributed.

In general, p -value weighting raises several important questions. How much power can we gain if we guess well in the weight assignment? How much power can we lose if we guess poorly? In this paper we show that the optimal weights have a simple parametric form and we investigate various approaches for estimating these weights. We also show the power is very robust to misspecification of the weights. In particular, in Section 3 we show that (i) sparse weights (few large weights and minimum weight close to 1) lead to huge power gains for well specified weights, but minute power loss for poorly specified weights; and (ii) in the nonsparse case, under weak conditions, the worst case power for poorly specified weights is typically better than the power obtained using equal weights.

We consider two methods for choosing the weights: (i) external weights, where prior information (based on scientific knowledge or prior data) singles out specific hypotheses (Section 4) and (ii) estimated weights where the data are used to construct weights (Section 5). External weights are prone to bias, while estimated weights are prone to variability. The two robustness properties reduce concerns about bias and variance.

To motivate this work consider an example (Figure 1) of external weighting that arises in genetic epidemiology. To identify variants of genes that induce greater susceptibility to disease, two types of studies (linkage and association) are often performed. Whole genome linkage analysis has been conducted for most major diseases. These data can be summarized by a linkage trace, a smooth stochastic process $\{Z(s) : s \in [0, L]\}$ where each s corresponds to a location on the genome. At points that correspond to a variant of a gene of interest, the mean of the process

$\mu(s) = E(Z(s))$ is a large positive value; however, due to extensive spatial correlation in the process, $\mu(s)$ is also nonzero in the vicinity of the variant. Tests for association between genetic polymorphisms and disease status for each of many genetic markers across the genome are also of interest. Like linkage analysis, the association statistics $\{T_j : j = 1, \dots, m\}$ map to spatial locations $\{s_j : j = 1, \dots, m\}$ on the genome. The number of tests m can be large, on the order of 1,000,000. Until recently, whole genome association analysis was prohibitively expensive, but technological advances have now made such studies feasible. Due to the multiple testing correction, it is difficult to achieve sufficient power to obtain definitive results in these studies. The linkage trace provides one obvious source of information from which the weights can be constructed; see Section 6 for further elaboration. Unlike linkage analysis, however, the spatial correlation in association tests is weak. For this reason, other choices such as genetic pathways could offer a more promising source for weights in the future.

2. BACKGROUND

2.1 Multiple Testing

Consider a multiple testing situation in which m tests are being performed. Suppose m_0 of the null hypotheses are true and $m_1 = m - m_0$ null hypotheses are false. We can categorize the m tests as in Table 1. In this notation F is the number of false positives. To control the familywise error rate, it is traditional to bound $P(F > 0)$ at α . When the tests are independent, the simplest way to control this probability is to reject only those tests for which the p -value is less than α/m ; this is called the Bonferroni procedure.

In 1995 Benjamini and Hochberg (BH) introduced a new approach to multiple hypothesis testing that controls the false discovery rate (FDR), defined as the expected fraction of false rejections among those hypotheses rejected. Let $P_{(1)} < \dots < P_{(m)}$ be the ordered p -values from m hypothesis tests, with $P_{(0)} \equiv 0$. Then, the BH procedure rejects any null hypothesis for which $P \leq T$ with

$$T = \max \left\{ P_{(i)} : P_{(i)} \leq \frac{\alpha i}{m} \right\}.$$

This quantity is of more scientific relevance than the overall type I error rate in GWAS. Also, the procedure is more powerful than the Bonferroni method. Adaptive variants of the procedure can increase power further at little additional computational expense; see

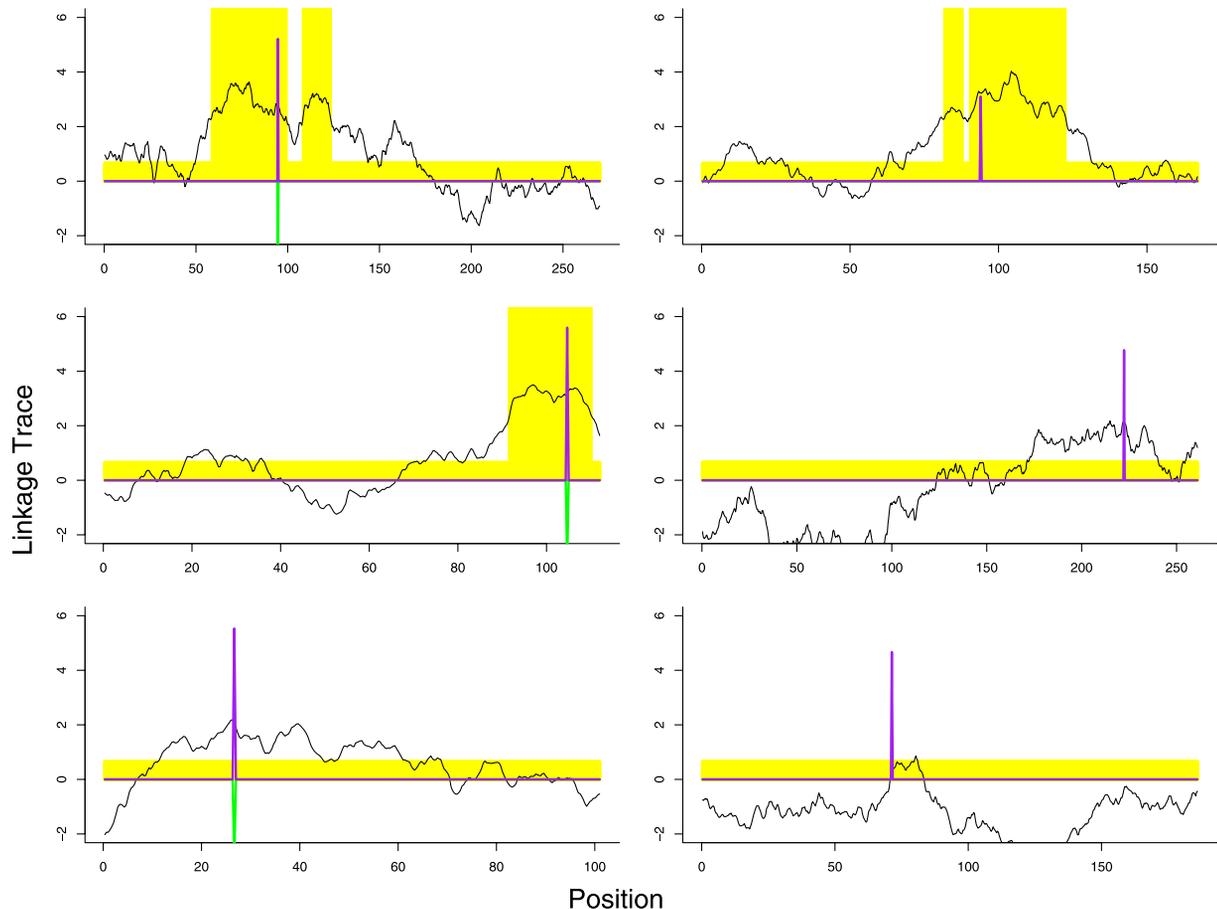


FIG. 1. Linkage trace and weights for 6 chromosomes. The trace is the linkage statistic plotted as a function of position on the chromosome. The shading indicates which p -values were up/down weighted. The upspike is the association test statistic. The 3 downspikes indicate tests that were rejected using the binary weights.

Benjamini, Krieger and Yekutieli (2006) and Storey (2002).

BH controls the false discovery rate at level $\alpha m_0/m$, where m_0 is the number of true null hypotheses. With certain dependence assumptions on the p -values, this is true regardless of how many nulls are true and regardless of the distribution of the p -values under the alternatives (Benjamini and Yekutieli, 2001; Blanchard and Roquain, 2009; Sarkar, 2002). Under some distributional assumptions, Genovese and Wasserman (2002)

TABLE 1
 2×2 classification of m hypothesis tests

	H_0 rejected	H_0 not rejected	Total
H_0 true	F	$m_0 - F$	m_0
H_0 false	T	$m_1 - T$	m_1
Total	S	$m - S$	m

show that, asymptotically, the BH method corresponds to rejecting all p -values less than a particular p -value threshold u^* . Specifically, u^* is the solution to the equation $H(u) = \beta u$ and $\beta = (\frac{1}{\alpha} - A_0)/(1 - A_0)$, where $A_0 = m_0/m$ and H is the (common) distribution of the p -value under the alternative. The key result is that $\alpha/m \leq u^* \leq \alpha$, which shows that the BH method is intermediate between Bonferroni (corresponding to α/m) and uncorrected testing (corresponding to α). If A_0 is close to 0, however, as it usually is in GWA, then β is a very large quantity and the power of the FDR is not much improved over the Bonferroni procedure.

The power of the BH method can be improved with adaptations. Blanchard and Roquain (2008) have given numerical comparisons of different adaptive procedures under dependence. Romano, Shaikh and Wolf (2008) have considered improving the adaptive procedure of Benjamini, Krieger and Yekutieli (2006) using the bootstrap. Sarkar and Heller (2008) have noted that

the adaptive procedure of Benjamini et al. may not perform well compared to Storey’s (2002) procedure for certain parameter choices.

2.2 Weighted Multiple Testing

We are given hypotheses $H = (H_1, \dots, H_m)$ and standardized test statistics $T = (T_1, \dots, T_m)$, where $T_j \sim N(\xi_j, 1)$. Likewise, $T_j^2 \sim \chi_1^2(\xi_j^2)$. For a two-sided hypothesis, $H_j = 1$ if $\xi_j \neq 0$ and $H_j = 0$ otherwise. For the sake of parsimony, unless otherwise noted, results will be stated for a one-sided test where $H_j = 1$ if $\xi_j > 0$, although the results extend easily to the two-sided case. Let $\theta = (\xi_1, \dots, \xi_m)$ denote the vector of means.

The p -values associated with the tests are $P = (P_1, \dots, P_m)$, where $P_j = \overline{\Phi}(T_j)$, $\overline{\Phi} = 1 - \Phi$ and Φ denotes the standard Normal CDF. Let $P_{(1)} \leq \dots \leq P_{(m)}$ denote the sorted p -values and let $T_{(1)} \geq \dots \geq T_{(m)}$ denote the sorted test statistics.

A rejection set \mathcal{R} is a subset of $\{1, \dots, m\}$. Say that \mathcal{R} controls familywise error at level α if $\mathbb{P}(\mathcal{R} \cap \mathcal{H}_0) \leq \alpha$, where $\mathcal{H}_0 = \{j : H_j = 0\}$. The Bonferroni rejection set is $\mathcal{R} = \{j : P_j < \alpha/m\} = \{j : T_j > z_{\alpha/m}\}$ where we use the notation $z_\beta = \overline{\Phi}^{-1}(\beta)$.

The weighted Bonferroni procedure (Rosenthal and Rubin, 1983; Genovese, Roeder and Wasserman, 2006) is as follows. Specify nonnegative weights $w = (w_1, \dots, w_m)$ and reject hypothesis H_j if

$$(1) \quad j \in \mathcal{R} = \left\{ j : \frac{P_j}{w_j} \leq \frac{\alpha}{m} \right\}.$$

In the following lemma we show that as long as $m^{-1} \sum_j w_j \equiv \overline{w} = 1$, the rejection set \mathcal{R} controls familywise error at level α . The second lemma includes a simple modification that will be needed later.

LEMMA 2.1. *If $\overline{w} = 1$, then \mathcal{R} controls familywise error at level α .*

LEMMA 2.2. *Suppose that $W_j = g(V_j, c)$, $j = 1, \dots, m$, for some random variables V_1, \dots, V_m , some constant c and some function g . Further, suppose that V_j has a known distribution H whenever $j \in \mathcal{H}_0$ and that P_j is independent of V_j for all $j \in \mathcal{H}_0$. The rule that rejects when $P_j \leq \alpha W_j/m$ controls familywise error at level α if c is chosen to satisfy $\mathbb{E}_H(g(V_j, c)) \leq 1$.*

Genovese, Roeder and Wasserman (2006) also showed that false discovery methods benefit by weighting. Recall that the false discovery proportion (FDP) is

$$\text{FDP} = \frac{\text{number of false rejections}}{\text{number of rejections}} = \frac{|\mathcal{R} \cap \mathcal{H}_0|}{|\mathcal{R}|},$$

where the ratio is defined to be 0 if the denominator is 0. The false discovery rate (FDR) is $\text{FDR} = \mathbb{E}(\text{FDP})$. Benjamini and Hochberg (1995) proved $\text{FDR} \leq \alpha$ if $\mathcal{R} = \{j : P_{(j)} \leq T\}$ where $T = \max\{j : P_{(j)} \leq j\alpha/m\}$. Genovese, Roeder and Wasserman (2006) showed that $\text{FDR} \leq \alpha$ if the P_j ’s are replaced by $Q_j = P_j/w_j$ provided $\overline{w} = 1$. This paper focuses on familywise error using the weighted procedure (1). Similar results hold for FDR and other familywise controlling procedures such as Holm’s test.

3. POWER, ROBUSTNESS AND OPTIMALITY

The optimal weights, derived below, can be re-expressed as optimal cutoffs for testing. Specifically, rejecting when $P_j/w_j \leq \alpha/m$ is the same as rejection when $T_j > \xi_j/2 + c/\xi_j$. This result can be obtained from Spjøtvoll (1972) and is identical to the result in Rubin, Dudoit and van der Laan (2006) obtained independently. The remainder of the paper, which shows some good properties of the weighted method, can thus also be considered as providing support for their method for selecting test specific cutoffs. In particular, Rubin et al. (2006)’s simulations indicate that even poorly specified estimates of the cutoffs $\xi_j/2 + c/\xi_j$ can still perform well. In this section we provide insight into why this is true.

The power of a single, one-sided alternative in the unweighted case ($w_j = 1$) is

$$\pi(\xi_j, 1) = \mathbb{P}(T_j > z_{\alpha/m}) = \overline{\Phi}(z_{\alpha/m} - \xi_j).$$

The power¹ in the weighted case is

$$(2) \quad \begin{aligned} \pi(\xi_j, w_j) &= \mathbb{P}\left(P_j < \frac{\alpha w_j}{m}\right) \\ &= \mathbb{P}\left(T_j > \overline{\Phi}^{-1}\left(\frac{\alpha w_j}{m}\right)\right) \\ &= \overline{\Phi}(z_{\alpha w_j/m} - \xi_j). \end{aligned}$$

Weighting increases the power when $w_j > 1$ and decreases the power when $w_j < 1$ for the j th alternative.

Given $\theta = (\xi_1, \dots, \xi_m)$ and $w = (w_1, \dots, w_m)$, we define the average power

$$\frac{1}{m_1} \sum_{j=1}^m \pi(\xi_j, w_j) I(\xi_j > 0),$$

where $m_1 = \sum_{j=1}^m I(\xi_j > 0)$. More generally, if ξ is drawn from a distribution Q and $w = w(\xi)$ is a

¹For a two-sided alternative the power is

$$\pi(\xi_j, w_j) = \overline{\Phi}(z_{\alpha w_j/2m} - \xi_j) + \overline{\Phi}(z_{\alpha w_j/2m} + \xi_j).$$

weight function, we define the average power $\int \pi(\xi, w(\xi))I(\xi > 0) dQ(\xi) / \int I(\xi > 0) dQ(\xi)$. If we take Q to be the empirical distribution of (ξ_1, \dots, ξ_m) , then this reduces to the previous expression. In this formulation we require $w(\xi) \geq 0$ and $\int w(\xi) dQ(\xi) = 1$.

In the following theorem we see that the set of optimal weight functions form a one parameter family indexed by a constant c .

THEOREM 3.1. *Given $\theta = (\xi_1, \dots, \xi_m)$, the optimal weight vector $w = (w_1, \dots, w_m)$ that maximizes the average power subject to $w_j \geq 0$ and $\bar{w} = 1$ is $w = (\rho_c(\xi_1), \dots, \rho_c(\xi_m))$, where*

$$(3) \quad \rho_c(\xi) = \left(\frac{m}{\alpha}\right) \bar{\Phi}\left(\frac{\xi}{2} + \frac{c}{\xi}\right) I(\xi > 0),$$

and $c \equiv c(\theta)$ is defined by the condition

$$(4) \quad \frac{1}{m} \sum_{j=1}^m \rho_c(\xi_j) = 1.$$

The proof, essentially a special case of Spjøtvoll (1972), is in the Appendix. Figure 2 displays the func-

tion $\rho_c(\xi)$ for various values of c (the function is normalized to have maximum 1 for easier visualization). The result generalizes to the case where the alternative means are random variables with distribution Q in which case c is defined by $\int \rho_c(\xi) dQ(\xi) = 1$.

From (2) and (3) we have immediately:

LEMMA 3.2. *The power at an alternative with mean ξ under optimal weights is $\bar{\Phi}(c/\xi - \xi/2)$. The average power under optimal weights, which we call the oracle power, is*

$$\frac{1}{m_1} \sum_{j=1}^m \bar{\Phi}\left(\frac{c}{\xi_j} - \frac{\xi_j}{2}\right) I(\xi_j > 0),$$

where $m_1 = \sum_j I(\xi_j > 0)$.

The oracle power is not attainable since the optimal weights depend on $\theta = (\xi_1, \dots, \xi_m)$. In practice, the weights will either be chosen by prior information or by estimating the ξ 's. This raises the following question: how sensitive is the power to correct specification of the weights? Now we show that the power is very robust to weight misspecification.

Property I: Sparse weights (minimum weight close to 1) are highly robust. If most weights are less than 1 and the minimum weight is close to 1, then correct specification (large weights on alternatives) leads to large power gains but incorrect specification (large weights on nulls) leads to little power loss.

Property II: Worst case analysis. Weighted hypothesis testing, even with poorly chosen weights, typically does as well or better than Bonferroni except when the alternative means are large, in which both have high power.

Let us now make these statements precise. Also, see Genovese, Roeder and Wasserman (2006) and Roeder et al. (2006) for other results on the effect of weight misspecification.

Property I. Consider first the case where the weights take two distinct values and the alternatives have a common mean ξ . Let ε denote the fraction of hypotheses given the larger of the two values of the weights B . Then, the weight vector w is proportional to

$$\underbrace{(B, \dots, B)}_{k \text{ terms}}, \underbrace{(1, \dots, 1)}_{m-k \text{ terms}},$$

where $k = \varepsilon m$ and $B > 1$, and, hence, the normalized weights are

$$w = \underbrace{(w_1, \dots, w_1)}_{k \text{ terms}}, \underbrace{(w_0, \dots, w_0)}_{m-k \text{ terms}},$$

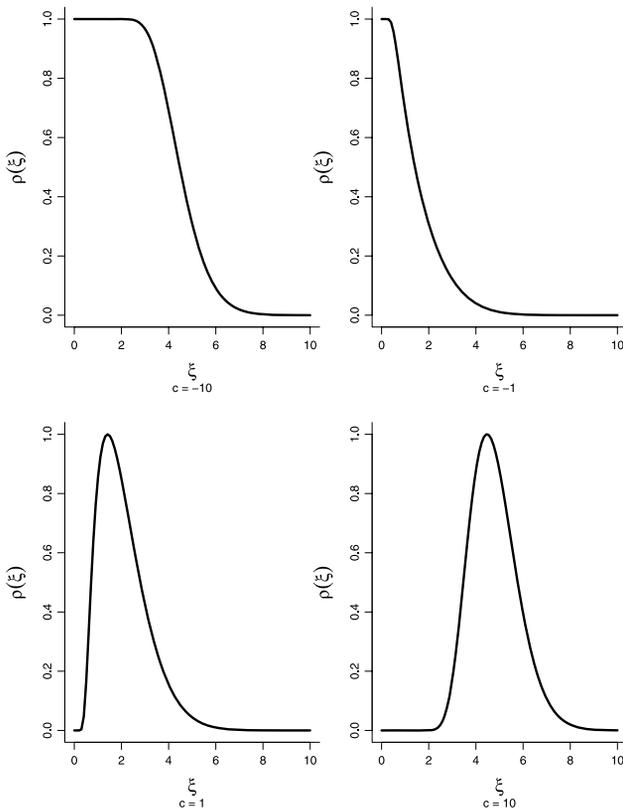


FIG. 2. Optimal weight function $\rho_c(\xi)$ for various c . In each case $m = 1000$ and $\alpha = 0.05$. The functions are normalized to have maximum 1.

where

$$w_1 = \frac{B}{\varepsilon B + (1 - \varepsilon)}, \quad w_0 = \frac{1}{\varepsilon B + (1 - \varepsilon)}.$$

We say that the weights are sparse if ε is small. Provided B is considerably less than $1/\varepsilon$, most weights are near 1 in the sparse case.

Rather than investigate the average power, we focus on a single alternative with mean ξ . The power gain by up-weighting this hypothesis is the power under weight w_1 minus the unweighted power $\pi(\xi, w_1) - \pi(\xi, 1)$. Similarly, the power loss for down-weighting is $\pi(\xi, 1) - \pi(\xi, w_0)$. The gain minus the loss, which we call the robustness function, is

$$\begin{aligned} R(B, \varepsilon) &\equiv (\pi(\xi, w_1) - \pi(\xi, 1)) \\ &\quad + (\pi(\xi, 1) - \pi(\xi, w_0)) \\ &= \bar{\Phi}(z_{\alpha w_1/m} - \xi) + \bar{\Phi}(z_{\alpha w_0/m} - \xi) \\ &\quad - 2\bar{\Phi}(z_{\alpha/m} - \xi). \end{aligned}$$

The gain outweighs the loss if and only if $R(B, \varepsilon) > 0$ (Figure 3).

In the sparse weighting scenario k is small and $w_0 \approx 1$ by assumption, consequently, an analysis of $R(B, \varepsilon)$ sheds light on the effect of weighting on power, without the added complications involved in a full analysis of average power.

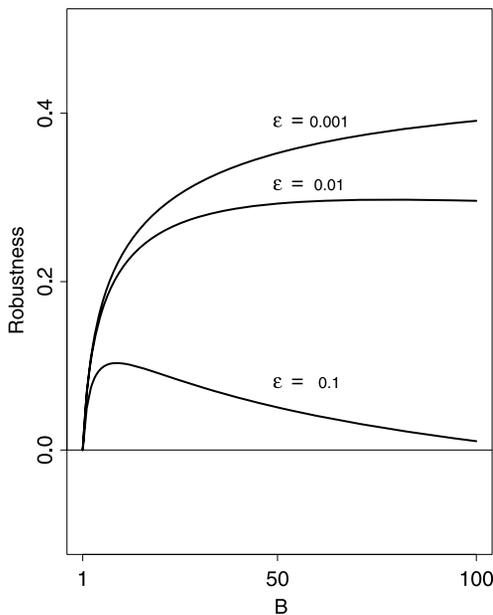


FIG. 3. Robustness function for $m = 1000$. In this example, $\xi = z_{\alpha/m}$ which has power $1/2$ without weighting. The gain of correct weighting far outweighs the loss for incorrect weighting as long as the fraction of large weights ε is small.

THEOREM 3.3. Fix $B > 1$. Then, $\lim_{\varepsilon \rightarrow 0} R(B, \varepsilon) > 0$. Moreover, there exists $\varepsilon^*(B) > 0$ such that $R(B, \varepsilon) > 0$ for all $\varepsilon < \varepsilon^*(B)$.

We can generalize this beyond the two-valued case as follows. Let w be any weight vector such that $\bar{w} = 1$. Now define the (worst case) robustness function

$$\begin{aligned} R(\xi) &\equiv \min_{\{j:w_j>1, H_j=1\}} \{\pi(\xi, w_j) - \pi(\xi, 1)\} \\ &\quad - \max_{\{j:w_j<1, H_j=1\}} \{\pi(\xi, 1) - \pi(\xi, w_j)\}. \end{aligned}$$

We will see that $R(\xi) > 0$ under weak conditions and that the maximal robustness is obtained for ξ near the Bonferroni cutoff $z_{\alpha/m}$.

THEOREM 3.4. A necessary and sufficient condition for $R(\xi) > 0$ is

$$\begin{aligned} (5) \quad R_{b,B}(\xi) &\equiv \Phi(z_{\alpha B/m} - \xi) + \Phi(z_{\alpha b/m} - \xi) \\ &\quad - 2\Phi(z_{\alpha/m} - \xi) \leq 0, \end{aligned}$$

where $B = \min\{w_j : w_j > 1\}$, $b = \min\{w_j\}$. Moreover,

$$R_{b,B}(\xi) = -\Delta(\xi) + O(1 - b),$$

where

$$\Delta(\xi) = (\Phi(z_{\alpha/m} - \xi) - \Phi(z_{\alpha B/m} - \xi)) > 0$$

and, as $b \rightarrow 1$, $\mu(\{\xi : R(\xi) < 0\}) \rightarrow 0$ and $\inf_{\xi} R(\xi) \rightarrow 0$.

Based on the theorem, we see that there is overwhelming robustness as long as the minimum weight is near 1. Even in the extreme case $b = 0$, there is still a safe zone, an interval of values of ξ over which $R(\xi) > 0$.

LEMMA 3.5. Suppose that $B \geq 2$. Then there exists $\xi_* > 0$ such that $R_{B,b}(\xi) > 0$ for all $0 \leq \xi \leq \xi_*$ and all b . An upper bound on ξ_* is $z_{\alpha/m} - 1/(z_{\alpha/m} - z_{B\alpha/m})$.

Property II. Even if the weights are not sparse, the power of the weighted test tends to be acceptable.

The result holds even though the weights themselves can be very sensitive to changes in θ . Consider the following example. Suppose that $\theta = (\xi_1, \dots, \xi_m)$ where each ξ_j is equal to either 0 or some fixed number ξ . The empirical distribution of the ξ_j 's is thus $Q = (1 - a)\delta_0 + a\delta_{\xi}$, where δ denotes a point mass and a is the fraction of nonzero means. The optimal weights are 0 for $\xi_j = 0$ and $1/a$ for $\xi_j = \xi$. Let $\tilde{Q} = (1 - a - \gamma)\delta_0 + \gamma\delta_u + a\delta_{\xi}$, where u is a small positive number. Since we have only moved the mass at 0 to u , and u is small, we would hope that $w(\xi)$

will not change much. But this is not the case. Set $\xi = A + \sqrt{A^2 - 2c}$, $u = B - \sqrt{B^2 - 2c}$, where

$$A = \bar{\Phi}^{-1}\left(\frac{\alpha}{(m(\gamma K + a))}\right),$$

$$B = \bar{\Phi}^{-1}\left(\frac{K\alpha}{(m(\gamma K + a))}\right).$$

This arrangement yields weights w_0 and w_1 on u and ξ such that $w_0/w_1 = K$. For example, if $m = 1000$, $\alpha = 0.05$, $a = 0.1$, $\gamma = 0.1$, $K = 1000$ and $c = 0.1$, then $u = 0.03$ and $\xi = 9.8$. The optimal weight on ξ under Q is 10 but under \tilde{Q} it is 0.00999 and so is reduced by a factor of 1001. More generally, we have the following result which shows that the weights are, in a certain sense, a discontinuous function of θ .

LEMMA 3.6. Fix α and m . For any $\delta > 0$ and $\varepsilon > 0$ there exists $Q = (1 - a)\delta_0 + a\delta_\xi$ and $\tilde{Q} = (1 - a - \gamma)\delta_0 + \gamma\delta_u + a\delta_\xi$ such that $d(Q, \tilde{Q}) < \delta$, and $\tilde{\rho}(\xi)/\rho(\xi) < \varepsilon$, where $a = \alpha/4$, $d(Q, \tilde{Q}) = \sup_\xi |Q(-\infty, \xi]|, \tilde{Q}(-\infty, \xi]|$ is the Kolmogorov–Smirnov distance, ρ is the optimal weight function for Q and $\tilde{\rho}$ is the optimal weight function for \tilde{Q} .

Fortunately, this feature of the weight function does not pose a serious hurdle in practice because it is possible to have high power even with poor weights. In Figure 4 the plots on the left show the power as a function of the alternative mean ξ . The dark solid line shows the lowest possible power assuming the weights were estimated as poorly as possible (under conditions specified below). The lighter solid line is the power of the unweighted (Bonferroni) method. The dotted line shows the power under theoretically optimal weights. The worst case weighted power is typically close to or larger than the Bonferroni power except for large ξ when they are both large.

To begin formal analysis, assume that each mean is either equal to 0 or ξ for some fixed $\xi > 0$. Thus, the empirical distribution is $Q = (1 - a)\delta_0 + a\delta_\xi$, where δ denotes a point mass and a is the fraction of nonzero ξ_j 's. The optimal weights are $1/a$ for hypotheses whose mean is ξ . To study the effect of misspecification error, consider the case where γm nulls are mistaken for alternatives with mean $u > 0$. This corresponds to misspecifying Q to be $\tilde{Q} = (1 - a - \gamma)\delta_0 + \gamma\delta_u + a\delta_\xi$. We will study the effect of varying u , so let $\pi(u)$ denote the power at the true alternative ξ as a function of u . Also, let π_{Bonf} denote the power using equal weights (Bonferroni). Note that changing $Q = (1 - a)\delta_0 + a\delta_\xi$ to $Q = (1 - a)\delta_0 + a\delta_{\xi'}$ for $\xi' \neq \xi$ does not change the weights.

As the weights are a function of c , we first need to find c as a function of u . The normalization condition (4) reduces to

$$(6) \quad \gamma \bar{\Phi}\left(\frac{u}{2} + \frac{c}{u}\right) + a \bar{\Phi}\left(\frac{\xi}{2} + \frac{c}{\xi}\right) = \frac{\alpha}{m},$$

which implicitly defines the function $c(u)$. First we consider what happens when u is restricted to be less than ξ .

THEOREM 3.7. Assume that $\alpha/m \leq \gamma + a \leq 1$. Let $Q = (1 - a)\delta_0 + a\delta_\xi$ and $\tilde{Q} = (1 - a - \gamma)\delta_0 + \gamma\delta_u + a\delta_\xi$ with $0 \leq u \leq \xi$. Let $C(\xi) = \sup_{0 \leq u \leq \xi} c(u)$ and define $\xi_0 = z_{\alpha/(m(\gamma+a))}$.

1. For $\xi \leq \xi_0$, $C(\xi) = \xi\xi_0 - \xi^2/2$. For $\xi > \xi_0$, $C(\xi)$ is the solution to

$$\gamma \bar{\Phi}(\sqrt{2c}) + a \bar{\Phi}\left(\frac{c}{\xi} + \frac{\xi}{2}\right) = \frac{\alpha}{m}.$$

In this case, $C(\xi) = z_{\alpha/(m\gamma)}^2/2 + O(a)$.

2. Let $\xi_* = z_{\alpha/m} + \sqrt{z_{\alpha/m}^2 - z_q^2}$, where $q = \alpha(1 - a)/(m\gamma)$. For $\xi < \xi_*$,

$$(7) \quad \inf_{0 < u < \xi} \pi(u) \geq \pi_{\text{Bonf}}.$$

For $\xi \geq \xi_*$ we have

$$(8) \quad \inf_{0 < u < \xi} \pi(u) \geq \bar{\Phi}\left(\frac{z_{\alpha/(m\gamma)}^2 - \xi_*^2}{2\xi_*}\right) - O(a)$$

$$(9) \quad \approx 1 - \bar{\Phi}\left(\sqrt{2 \log \frac{1-a}{\gamma}}\right) - O(a)$$

$$(10) \quad \geq 1 - \frac{\gamma}{1-a} - O(a).$$

The factor $\bar{\Phi}(\sqrt{2 \log(1-a)/\gamma}) \approx \gamma/(1-a)$ is the worst case power deficit due to misspecification. Now we drop the assumption that $u \leq \xi$.

THEOREM 3.8. Let $Q = (1 - a)\delta_0 + a\delta_\xi$ and let $Q_u \equiv (1 - a - \gamma)\delta_0 + \gamma\delta_u + a\delta_\xi$. Let π_u denote the power at ξ using the weights computed under Q_u .

1. The least favorable u is $u_* \equiv \arg \min_{u \geq 0} \pi_u = \sqrt{2c_*} = z_{\alpha/(m\gamma)} + O(a)$, where c_* solves

$$\gamma \bar{\Phi}(\sqrt{2c_*}) + a \bar{\Phi}\left(\frac{\xi}{2} + \frac{c_*}{\xi}\right) = \frac{\alpha}{m}$$

and $c_* = z_{\alpha/(m\gamma)}^2/2 + O(a)$.

2. The minimal power is

$$\inf_u \pi_u = \bar{\Phi}\left(\frac{c_*}{\xi} - \frac{\xi}{2}\right) = \bar{\Phi}\left(\frac{z_{\alpha/(m\gamma)}^2 - \xi^2}{2\xi}\right) + O(a).$$

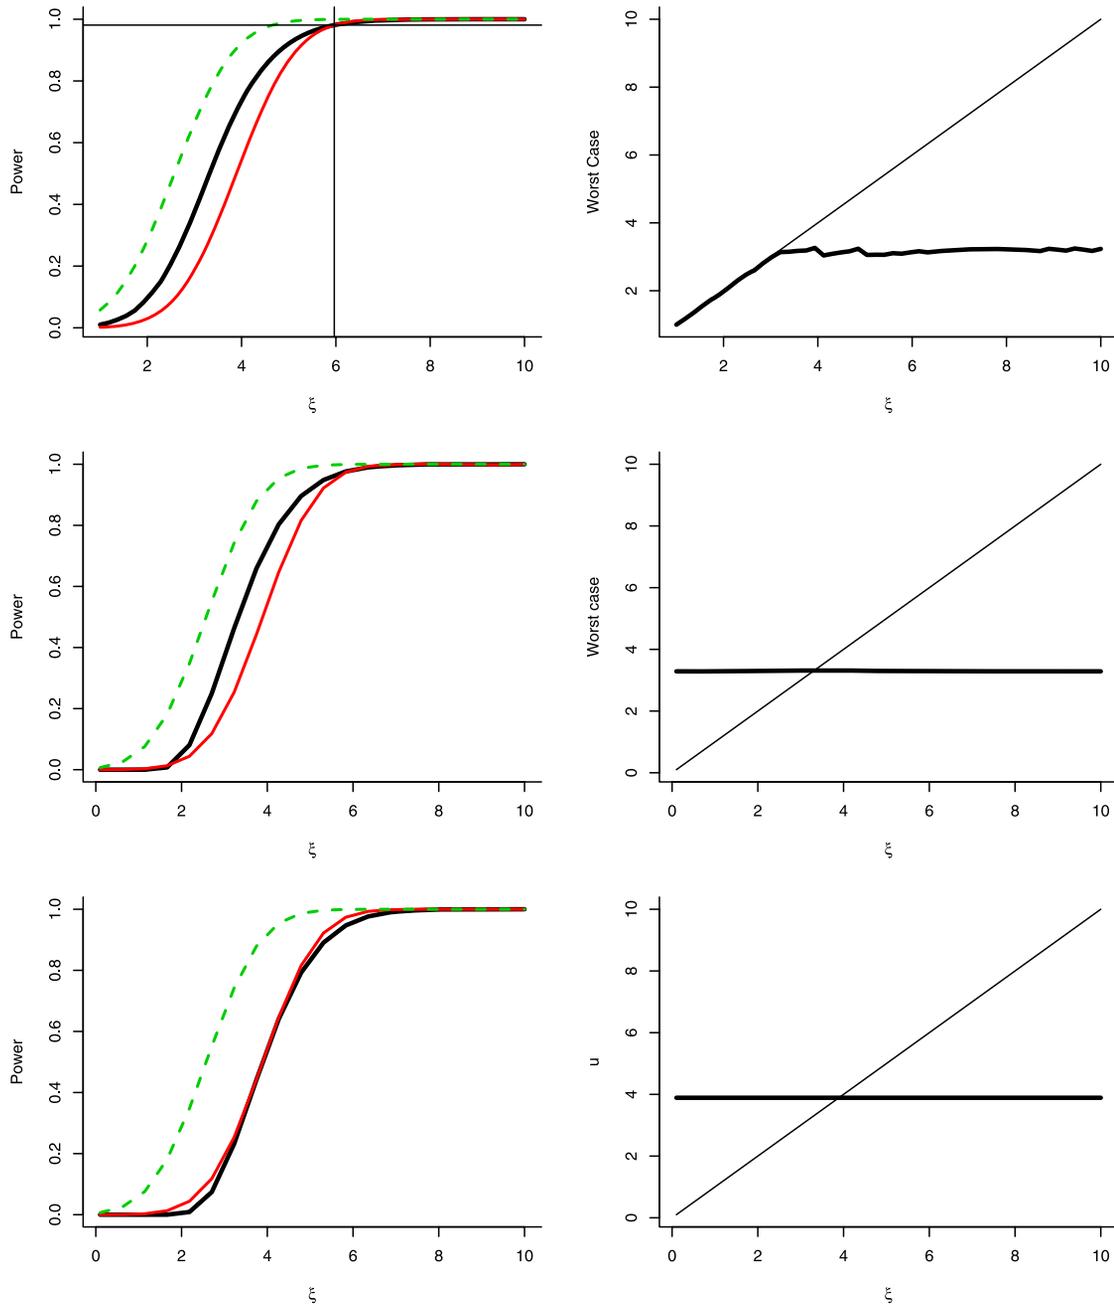


FIG. 4. Power as a function of the alternative mean ξ . In these plots, $a = 0.01$, $m = 1000$ and $\alpha = 0.05$. There are $(1 - a)m$ nulls and ma alternatives with mean ξ . The left plots show what happens when the weights are incorrectly computed assuming that a fraction γ of nulls are actually alternatives with mean u . In the top plot, we restrict $0 < u < \xi$. In the second and third plots, no restriction is placed on u . The top and middle plots have $\gamma = 0.1$, while the third plot has $\gamma = 1 - a$ (all nulls misspecified as alternatives). The dark solid line shows the lowest possible power assuming the weights were estimated as poorly as possible. The lighter solid line is the power of the unweighted (Bonferroni) method. The dotted line is the power under the optimal weights. The vertical line in the top plot is at ξ^* . The weighted method beats unweighted for all $\xi < \xi^*$. The right plot shows the least favorable u as a function of ξ . That is, mistaking γm nulls for alternatives with mean u leads to the worst power. Also shown is the line $u = \xi$.

3. A sufficient condition for $\inf_u \pi_u$ to be larger than the power of the Bonferroni method is $\xi \geq z_{\alpha/m} + \sqrt{z_{\alpha/m}^2 - z_{\alpha/(m\gamma)}^2} + O(a)$.

4. CHOOSING EXTERNAL WEIGHTS

One approach to choosing external weights (or test statistic cutoffs) is to use empirical Bayes methods to model prior information while being careful to preserve error control as in Westfall and Soper (2001), for example. Here we consider a simple method that takes advantage of the robustness properties we have discussed. We will focus here on the two-valued case. Thus,

$$w = (\underbrace{w_1, \dots, w_1}_{k \text{ terms}}, \underbrace{w_0, \dots, w_0}_{m-k \text{ terms}}),$$

where $k = \varepsilon m$, $w_1 = B/(\varepsilon B + (1 - \varepsilon))$ and $w_0 = 1/(\varepsilon B + (1 - \varepsilon))$. In practice, we would typically have a fixed fraction of hypotheses ε that we want to give more weight to. The question is how to choose B . We will focus on choosing B to produce weights with good properties at interesting values of ξ . Now large values of ξ already have high power. Very small values of ξ have extremely low power and benefit little by weighting. This leads us to focus on constructing weights that are useful for a *marginal effect*, defined as the alternative ξ_0 that has power 1/2 when given weight 1. Thus, the marginal effect is $\xi_0 = z_{\alpha/m}$. In the rest of this section then we assume that all nonzero ξ_j 's are equal to ξ_0 . Of course, the validity of the procedure does not depend on this assumption being true.

Fix $0 < \varepsilon < 1$ and vary B . As we increase B , we will eventually reach a point $B_0(\varepsilon)$ where $R(B, \varepsilon) < 0$, which we call the turnaround point. Formally, $B_0(\varepsilon) = \sup\{B : R(B, \varepsilon) > 0\}$. The top panel in Figure 5 shows

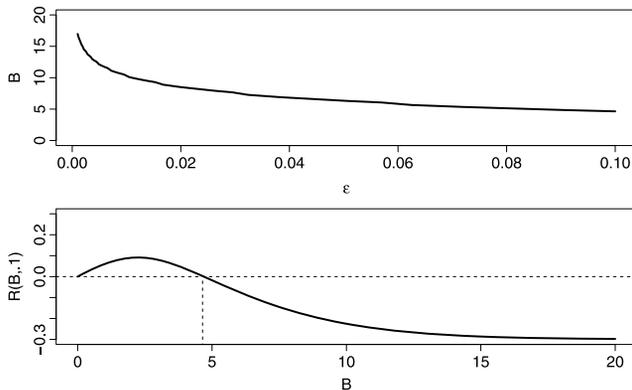


FIG. 5. Top plot: turnaround point $B_0(\varepsilon)$ versus ε . Bottom plot shows the robustness function $R(B, 0.1)$ versus B . The turnaround point $B_0(\varepsilon)$ is shown with a vertical dotted line.

$B_0(\varepsilon)$ versus ε , which shows that for small ε we can choose B large without loss of power. The bottom panel shows $R(B, \varepsilon)$ for $\varepsilon = 0.1$. Ideally, for a given ε , one chooses B near $B_*(\varepsilon)$, the value of B that maximizes $R(B, \varepsilon)$.

THEOREM 4.1. Fix $0 < \varepsilon < 1$. As a function of B , $R(B, \varepsilon)$ is unimodal and satisfies $R(1, \varepsilon) = 1$, $R'(1, \varepsilon) > 0$ and $R(\infty, \varepsilon) < 0$. Hence, $B_0(\varepsilon)$ exists and is unique. Also, $R(B, \varepsilon)$ has a unique maximum at some point $B^*(\varepsilon)$ and $R(B^*(\varepsilon), \varepsilon) > 0$.

When ε is very small, B can be large, provided $w_0 \approx 1$. For example, suppose we want to increase the chance of rejecting one particular hypothesis so that $\varepsilon = 1/m$. Then,

$$w_1 = \frac{mB}{B + m - 1} \approx B, \quad w_0 = \frac{1}{B + m - 1} \approx 1$$

and

$$\lim_{m \rightarrow \infty} \lim_{B \rightarrow \infty} \pi(\xi_j, w_1) = 1,$$

$$\text{while } \lim_{m \rightarrow \infty} \lim_{B \rightarrow \infty} \pi(\xi_j, w_0) = \frac{1}{2}.$$

The next results show that binary weighting schemes are optimal in a certain sense. Suppose we want to have at least a fraction ε with high power $1 - \beta$ and otherwise we want to maximize the minimum power.

THEOREM 4.2. Consider the following optimization problem: Given $0 < \varepsilon < 1$ and $0 < \beta < 1/2$, find a vector $w = (w_1, \dots, w_m)$ that maximizes $\min_j \pi(\xi_m, w_j)$ subject $\bar{w} = 1$, and $\#\{j : \pi(w_j, \xi_m) \geq 1 - \beta\}/m \geq \varepsilon$. The solution is given by $c = \bar{\Phi}(z_{\alpha/m} + z_{1-\beta})$, $B = cm(1 - \varepsilon)/(\alpha - \varepsilon cm)$, $w_1 = B/(\varepsilon B + (1 - \varepsilon))$, $w_0 = 1/(\varepsilon B + (1 - \varepsilon))$ and $k = \varepsilon m$.

If our goal is to maximize the number of alternatives with high power while maintaining a minimum power loss, the solution is given as follows.

THEOREM 4.3. Consider the following optimization problem: Given $0 < \beta < 1/2$, find a vector $w = (w_1, \dots, w_m)$ that maximizes $\#\{j : \pi(w_j, \xi_m) \geq 1 - \beta\}$ subject to $\bar{w} = 1$, and $\min_j \pi(w_j, \xi_m) \geq \delta$. The solution is

$$w_1 = \frac{m}{\alpha} \bar{\Phi}(z_{\alpha/m} + z_{1-\beta}), \quad w_0 = \frac{m}{\alpha} \bar{\Phi}(z_{\alpha/m} + z_\delta),$$

$$\varepsilon = \frac{1 - w_0}{w_1 - w_0}$$

and $k = m\varepsilon$.

A special case that falls under this theorem permits the minimum power to be 0. In this case $w_0 = 0$ and $\varepsilon = 1/w_1$.

5. ESTIMATED WEIGHTS

In practice, ξ_j is not known, so it must be estimated to utilize the weight function. A natural choice is to build on the two stage experimental design (Satagopan and Elston, 2003; Wang et al., 2006) and split the data into subsets, using one subset to estimate ξ_i , and hence $w(\xi_i)$, and the second to conduct a weighted test of the hypothesis (Rubin, Dudoit and van der Laan, 2006). This approach would arise naturally in an association test conducted in stages. It does lead to a gain in power relative to unweighted testing of stage 2 data; however, it is not better than simply using the full data set without weights for the analysis (Rubin, Dudoit and van der Laan, 2006). These results are corroborated by Skol et al. (2006) in a related context. They showed that it is better to use stages 1 and 2 jointly, rather than using stage 2 as an independent replication of stage 1.

To gain a strong advantage with data-based weights, prior information is needed. One option is to order the tests (Rubin, Dudoit and van der Laan, 2006), but with a large number of tests this can be challenging. The type of prior information readily available to investigators is often nonspecific. For instance, SNPs might naturally be grouped, based on features that make various candidates more promising for this disease under investigation. For a brain-disorder phenotype we might cross-classify SNPs by categorical variables such as functionality, brain expression and so forth. The SNPs in one group may seem most promising, a priori, while those in another seem least promising. Intermediate groups may be somewhat ambiguous. It is easy to imagine additional variables that further partition the SNPs into various classes that help to separate the more promising SNPs from the others. While this type of information lends itself to grouping SNPs, it does not lead directly to weights for the groups. Indeed, it might not even be possible to choose a natural ordering of the groups. What is needed is a way to use the data to determine the weights, once the groups are formed.

Until recently, methods for weighted multiple-testing required that prior weights be developed independently of the data under investigation (Genovese, Roeder and Wasserman, 2006; Roeder, Wasserman and Devlin, 2007). Here we provide a data-based estimate of weights based on results of grouped analysis. One way to implement this approach is to follow these steps:

1. Partition the tests into subsets $\mathcal{G}_1, \dots, \mathcal{G}_K$, with the k th group containing r_k elements, ensuring that r_k is at least 20–30.

2. Calculate the sample mean Y_k and variance S_k^2 for the test statistics in each group.

3. Label the i th test in group k , T_{ik} . At best, only a fraction of the elements in each group will have a signal, hence, we assume that for $i = 1, \dots, r_k$ the distribution of the test statistics is approximated by a mixture model

$$T_{ik} \sim (1 - \pi_k)N(0, 1) + \pi_k N(\xi_k, 1)$$

or

$$T_{ik} \sim (1 - \pi_k)\chi_1^2(0) + \pi_k \chi_1^2(\xi_k^2),$$

where ξ_k is the signal size for those tests with a signal in the k th group. This is an approximation because the signal is likely to vary across tests. The mixture of normals is only appropriate when the tests are one-sided. For two-sided alternatives, the χ^2 is the natural approach. This test squares the noncentrality parameter, effectively removing any ambiguity about the direction of the associations.

4. Estimate (π_k, ξ_k) using the method of moments estimator (for details see the Appendix). Because ξ_k has no meaning when $\pi_k = 0$, the $\hat{\xi}_k$ is set to 0 when $\hat{\pi}_k$ is close to zero. For the normal model the estimators are

$$(11) \quad \hat{\pi}_k = Y_k^2 / (Y_k^2 + S_k^2 - 1), \quad \hat{\xi}_k = Y_k / \pi_k,$$

provided $\hat{\pi}_k > 1/r_k$; otherwise $\hat{\xi}_k = 0$.

For the χ^2 model they are

$$(12) \quad \hat{\xi}_k^2 = \frac{(S_k^2 + Y_k^2 + 3)}{Y_k - 1}, \quad \hat{\pi}_k = \frac{Y_k - 1}{\hat{\xi}_k^2},$$

provided $Y_k > 1$ and $1/r_k < \hat{\pi}_k < (r_k - 1)/r_k$; otherwise $\hat{\xi}_k = 0$.

5. For each of the k groups, construct weights $w(\hat{\xi}_k)$. It is apparent in Figure 1 that if $|\hat{\xi}_k| < \delta$, for δ near 0, then $w(\hat{\xi}_k) \approx 0$ and it is unlikely that any tests in the k th group will be significant, regardless of the p -value. The stochastic quantity δ depends upon the relative values of $(\hat{\xi}_1, \dots, \hat{\xi}_K)$, and the number of elements in each group. For this reason we have found that smoothing the weights generally improves the power of the procedure. We suggest using a linear combination such as

$$\hat{w}_k = (1 - \gamma)w(\hat{\xi}_k) + \gamma K^{-1} \sum_k w(\hat{\xi}_k),$$

with $\gamma = 0.01$ or 0.05 . The larger the choice of γ , the more evenly distributed the weights across groups. Alternatively, one could smooth the weights by using a

Stein shrinkage estimator or bagging procedure to obtain a more robust estimator of (ξ_1, \dots, ξ_K) (Hastie, Tibshirani and Friedman, 2001). Regardless of how the weights are smoothed, one should renorm them to ensure the weights sum to m . Each test in group k receives the weight \hat{w}_k . Another effect of the smoothing is to ensure that each group gets a weight greater than 0.

This weighting scheme relies on data-based estimators of the optimal weights, but with a partition of the data sufficiently crude to preserve the control of family-wise error rate. The approach is an example of the “sieve principle” (Bickel et al., 1993). The sieve principle works because the number of parameters estimated is far less than the number of observations. Thus, many observations are used to estimate each parameter. Consequently, parameters are estimated with substantially less variability than if they were estimated using only the test statistics from the particular gene under investigation. Because the weights are determined by the size of the tests in the entire cluster, the probability of upweighting simply because a single test is large, due to chance, is small.

6. EXAMPLES

6.1 Binary Weights

In a study of nicotine dependence, Saccone et al. (2007) used binary weights in a candidate gene study. Their study involved 3713 genetic variants (single nucleotide polymorphisms or SNPs) encompassing 348 genes. The genes were divided into two types: 52 nicotinic and dopaminergic receptor genes; and 296 other candidate genes. Each SNP associated with a gene in the first group was allocated ten times the weight of a gene in the other category. Using a generous false discovery rate ($\alpha = 0.4$), they identified 39 SNPs; 78% of these were nicotine receptors, in contrast to the fraction of nicotine receptors overall (15%).

6.2 Independent Data Weights

For family-based study designs, tests of association are based on transmission data. In these studies, data are available from which one can compute the potential power to detect a signal at each SNP tested; see Ionita-Laza et al. (2007) for a detailed explanation of this unique feature of family-based data. Because the data used to calculate the power are independent of the test statistics for association, these data

are available for construction of the weights. Motivated by this possibility, Ionita-Laza et al. (2007) developed a weighting scheme. Using independent data, they ranked the SNPs from most to least promising, in terms of power. They then constructed an exponential weighting scheme, based on simulations of genetic models. The scheme results in a small number of SNPs receiving a top weight, successively more SNPs receiving correspondingly lower weights, and finally a large number receiving the lowest weight. In their simulations they found that the power of the test can often be doubled using this procedure. Using the FHS data, they apply the technique to a genome-wide association study with 116,204 SNPs and 923 participants. The phenotype of interest is height. Using their weighting scheme, they obtained one significant result with weights and none without weights.

6.3 Linkage Weights

Finding variation in the genetic code that increases the risk for complex diseases, such as Type II diabetes and schizophrenia, is critically important to the advancement of genetic epidemiology. In the [Introduction](#) we describe a means by which weights could be extracted from linkage data. Here we illustrate the idea with both data and simulations.

In the analysis of 955 cases and 1498 controls enrolled in a genome-wide association study, McQueen and colleagues (2008) used weights derived from published linkage results. They combined results from 11 linkage studies on bipolar disorder to obtain Z scores corresponding to the locations of each association test. From the linkage results they computed weighted p -values using the cumulative normal weight function (Roeder et al., 2006). Although none of their results were genome-wide significant, they obtained promising results in four regions. Three of these are obtained due to strong p -values in combination with a linkage peak. One signal did not correspond to a linkage peak, but continued to be in the top tier of p -values, after weights were applied.

To illustrate how binary weights could be derived from such linkage data, we present a realistic synthetic example. Using the methods described in Roeder et al. (2006), we create a linkage trace that captures many of the features found in actual linkage traces. In this simulation we generate a full genome (23 chromosomes) and place 20 disease variants at random, one per chromosome. The signals from these variants were designed to yield weak signals with broad peaks. Next, we simulated 100,000 normally distributed association

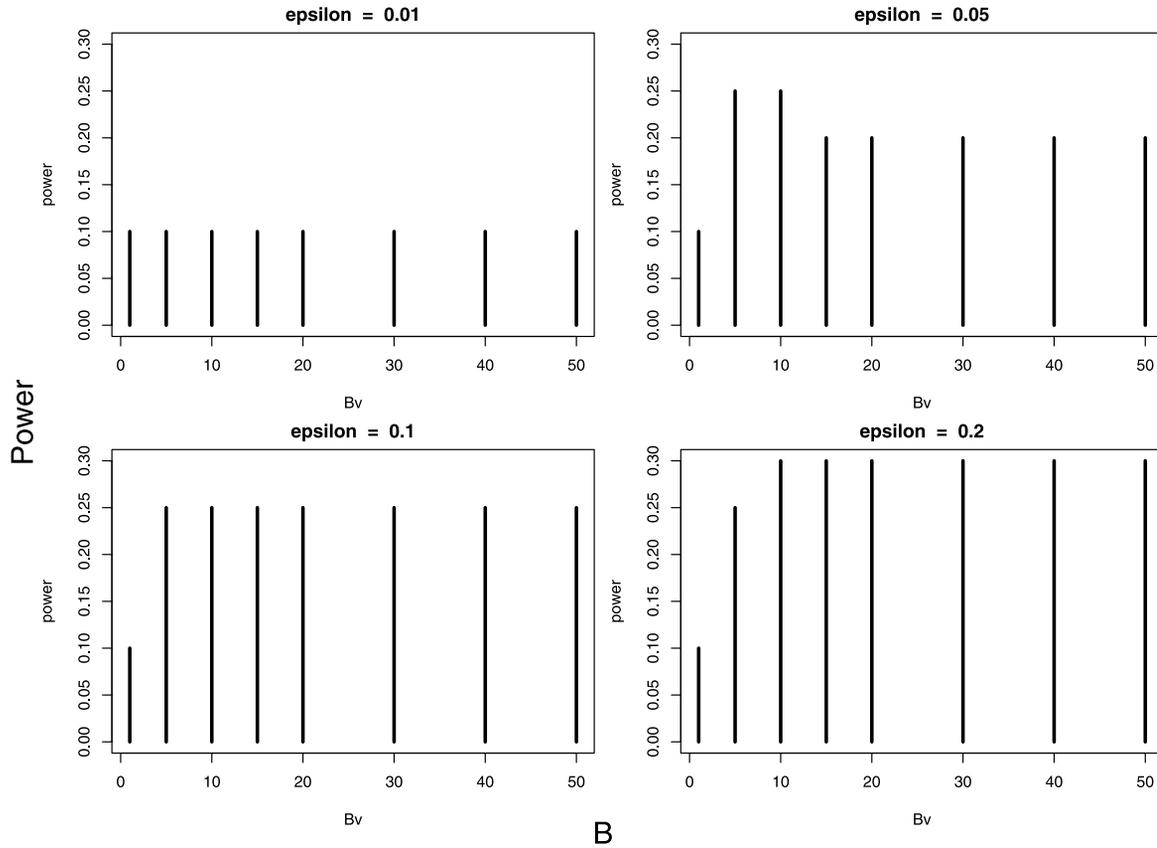


FIG. 6. Power as a function of B and ϵ .

test statistics mapped to the same genome. Again, 20 of these tests were generated under the alternative hypothesis of association. These signals were also weak.

To illustrate the synthetic data, six typical chromosomes are displayed in Figure 1. Each displayed chromosome has one true signal, with the association test statistic at that location indicated by an upspike; none of the association tests generated under the null hypothesis are plotted. Without weights, only 2 of the 20 signals could be detected using a Bonferonni correction. Using binary weights, as described above, with $\epsilon = 0.05$ and $B = 10$, we discover 5 of the 20 signals. In the left column of the figure all three signals were discovered, while in the right column none were discovered (indicated by presence of a down-spike). Comparing the top row, we see that both signals were up-weighted in the correct location, but the association signal was not strong enough in the top right chromosome to achieve significance. Alternatively, in the bottom left panel the association statistic was substantial enough to reject the null hypothesis without the benefit of up-weighting.

To examine the robustness of the procedure to choice of weights, we tried 4 choices of ϵ (0.01, 0.05, 0.1, 0.2) with $1 \leq B \leq 50$. We made no false discoveries with any of these choices. The power is displayed in Figure 6. To assist in the choice of parameters, we have found it helpful to examine the number of discoveries for each choice. In this example, the number of discoveries varied between 2 (unweighted, i.e., $B = 1$) to 6 ($\epsilon = 0.2$, $B \geq 10$). Five discoveries were made for a broad range of choices. In principle, choosing (ϵ, B) to maximize the number of discoveries can inflate the error rate. In our simulations we have found that searching within the family of weights defined by 1 or 2 parameters, such as this binary weight system based upon a linkage trace, tends to provide very close to nominal protection against false discoveries.

7. DISCUSSION

Several authors have explored the effect of weights on the power of multiple testing procedures [e.g., Westfall, Kropf and Finos (2004)]. These investigations show that the power of multiple testing procedures can

be increased by using weighted p -values. Here we derive the optimal weights for a commonly used family of tests and show that the power is remarkably robust to misspecification of these weights.

The same ideas used here can be applied to other testing methods to improve power. In particular, weights can be added to the FDR method, Holm’s step-down test, and the Donoho and Jin (2004) method. Weighting ideas can also be used for confidence intervals. Another open question is the connection with Bayesian methods which have already been developed to some extent in Efron et al. (2001).

GWAS for some phenotypes such as Type 1 diabetes have yielded exciting results (Todd et al., 2007), while results for other complex diseases have been much less successful. Presumably, many studies do not have sufficient power to detect the genetic variants associated with the phenotypes, even though thousands of cases and controls have been genotyped. To bolster power, we recommend up-weighting and down-weighting hypotheses, based on prior likelihood of association with the phenotype. For instance, Wang, Li and Bucan (2007) describe pathway-based approaches for the analysis of GWAS.

Multiple testing arises in GWAS analyses in other contexts as well. Frequently, multiple tests, assuming different genetic models, are applied to each genetic marker. Multiple markers in a neighborhood can be analyzed simultaneously to increase the signal, using haplotypes, multivariate models and fine-mapping techniques. Data are often collected in multiple stages of the experiment, and at each stage promising markers are tested for association. In summary, many questions concerning multiple testing remain open in the context of GWAS.

APPENDIX

PROOF OF LEMMA 2.1. The familywise error is

$$\begin{aligned} \mathbb{P}((\mathcal{R} \cap \mathcal{H}_0) > 0) &= \mathbb{P}\left(P_j \leq \frac{\alpha w_j}{m} \text{ for some } j \in \mathcal{H}_0\right) \\ &\leq \sum_{j \in \mathcal{H}_0} \mathbb{P}\left(P_j \leq \frac{\alpha w_j}{m}\right) = \frac{\alpha}{m} \sum_{j \in \mathcal{H}_0} w_j \\ &\leq \alpha \bar{w} = \alpha. \end{aligned} \quad \square$$

PROOF OF LEMMA 2.2. The familywise error is

$$\begin{aligned} \mathbb{P}((\mathcal{R} \cap \mathcal{H}_0) > 0) \\ = \mathbb{P}\left(P_j \leq \frac{\alpha W_j}{m} \text{ for some } j \in \mathcal{H}_0\right) \end{aligned}$$

$$\begin{aligned} &\leq \sum_{j \in \mathcal{H}_0} \mathbb{P}\left(P_j \leq \frac{\alpha W_j}{m}\right) \\ &= \sum_{j \in \mathcal{H}_0} \mathbb{E}_H\left(\mathbb{P}\left(P_j \leq \frac{\alpha w_j}{m} \mid W_j = w_j\right)\right) \\ &= \sum_{j \in \mathcal{H}_0} \mathbb{E}_H(\alpha W_j / m) = \frac{\alpha}{m} \sum_{j \in \mathcal{H}_0} \mathbb{E}_H(W_j) \\ &\leq \frac{m_0 \alpha}{m} \leq \alpha. \end{aligned} \quad \square$$

PROOF OF THEOREM 3.1. Let C denote the set of hypotheses with $\xi_j > 0$. Power is optimized if $w_j = 0$ for $j \notin C$. The average power is

$$\frac{1}{m_1} \sum_{j \in C} \bar{\Phi}\left(\bar{\Phi}^{-1}\left(\frac{\alpha w_j}{m}\right) - \xi_j\right),$$

with constraint

$$\sum w_j = m.$$

Choose \underline{w} to maximize

$$\pi = \frac{1}{m_1} \sum_{j \in C} \bar{\Phi}\left(\bar{\Phi}^{-1}\left(\frac{\alpha w_j}{m}\right) - \xi_j\right) - \lambda(m - \sum w_i)$$

by setting the derivative to zero

$$\begin{aligned} \frac{\partial}{\partial w_i} \pi &= -\lambda + \frac{\phi(\bar{\Phi}^{-1}(\alpha w_j / m) - \xi_j) \alpha}{\phi(\bar{\Phi}^{-1}(\alpha w_j / m)) m} = 0, \\ \frac{m \lambda}{\alpha} &= \frac{\phi(\bar{\Phi}^{-1}(\alpha w_j / m) - \xi_j)}{\phi(\bar{\Phi}^{-1}(\alpha w_j / m))}. \end{aligned}$$

The \underline{w} that solves these equations is given in (3). Finally, solve for c such that $\sum_i w_i = m$. \square

PROOF OF THEOREM 3.4. The first statement follows easily by noting that the worst case corresponds to choosing weight B in the first term in $R(\xi)$ and choosing weight b in the second term in $R(\xi)$. The rest follows by Taylor expanding $R_{b,B}(\xi)$ around $b = 1$. \square

PROOF OF LEMMA 3.5. With $b = 0$, $R_{b,B}(\xi) \geq 0$ when

$$(13) \quad \bar{\Phi}(z_{B\alpha/m} - \xi) - 2\bar{\Phi}(z_{\alpha/m} - \xi) \geq 0.$$

With $B \geq 2$, (13) holds at $\xi = 0$. The left-hand side is increasing in ξ for ξ near 0, but (13) does not hold at $\xi = z_{\alpha/m}$. So (13) must hold in the interval $[0, \xi_*]$. Rewrite (13) as $\bar{\Phi}(z_{B\alpha/m} - \xi) - \bar{\Phi}(z_{\alpha/m} - \xi) \geq \bar{\Phi}(z_{\alpha/m} - \xi)$. We lower bound the left-hand side and upper bound the right-hand side. The left-hand side is

$\bar{\Phi}(z_{B\alpha/m} - \xi) - \bar{\Phi}(z_{\alpha/m} - \xi) = \int_{z_{B\alpha/m} - \xi}^{z_{\alpha/m} - \xi} \phi(u) du \geq (z_{\alpha/m} - z_{B\alpha/m})\phi(z_{\alpha/m} - \xi)$. The right-hand side can be bounded using Mill's ratio: $\bar{\Phi}(z_{\alpha/m} - \xi) \leq \phi(z_{\alpha/m} - \xi)/(z_{\alpha/m} - \xi)$. Set the lower bound greater than the upper bound to obtain the stated result. \square

PROOF OF LEMMA 3.6. Choose $K > 1$ such that $1/(K + 1) < 1/a - \varepsilon$. Choose $1 > \gamma > (2\alpha - a)/K$. Choose a small $c > 0$. Let $\xi = A + \sqrt{A^2 - 2c}$ and $u = B - \sqrt{B^2 - 2c}$, where

$$A = \bar{\Phi}^{-1}\left(\frac{\alpha}{(m(\gamma K + a))}\right),$$

$$B = \bar{\Phi}^{-1}\left(\frac{K\alpha}{(m(\gamma K + a))}\right).$$

Then $\rho(\xi) = 1/a$ and $\tilde{\rho}(\xi) = 1/(K + 1)$. Now $d(Q, \tilde{Q}) = \gamma$. Taking K sufficiently large and γ sufficiently close to $(2\alpha - a)/K$ makes $\gamma < \delta$. \square

It is convenient to prove Theorem 3.8 before proving Theorem 3.7.

PROOF OF THEOREM 3.8. Let c_* solve

$$(14) \quad \gamma \bar{\Phi}(\sqrt{2c_*}) + a \bar{\Phi}\left(\frac{\xi}{2} + \frac{c_*}{\xi}\right) = \frac{\alpha}{m}.$$

We claim first that, for any $c > c_*$, there is no u such that the weights average to 1. Fix $c > c_*$. The weights average to 1 if and only if

$$(15) \quad \gamma \bar{\Phi}\left(\frac{c}{u} + \frac{u}{2}\right) + a \bar{\Phi}\left(\frac{\xi}{2} + \frac{c}{\xi}\right) = \frac{\alpha}{m}.$$

Since $c > c_*$ and since the second term is decreasing in c , we must have

$$\bar{\Phi}\left(\frac{c}{u} + \frac{u}{2}\right) > \bar{\Phi}(\sqrt{2c_*}).$$

The function $r(u) = \bar{\Phi}(c/u + u/2)$ is maximized at $u = \sqrt{2c}$. So $r(\sqrt{2c}) \geq r(u)$. But $r(\sqrt{2c}) = \bar{\Phi}(\sqrt{2c})$. Hence, $\bar{\Phi}(\sqrt{2c}) \geq r(u) \geq \bar{\Phi}(\sqrt{2c_*})$. This implies $c < c_*$, which is a contradiction. This establishes that $\sup_u c(u) \leq c_*$. On the other hand, taking $c = c_*$ and $u = \sqrt{2c_*}$ solves equation (15). Thus, c_* is indeed the largest c that solves the equation which establishes the first claim. The second claim follows by noting that

$$\gamma \bar{\Phi}(\sqrt{2c_*}) + a \bar{\Phi}\left(\frac{\xi}{2} + \frac{c_*}{\xi}\right) = \gamma \bar{\Phi}(\sqrt{2c_*}) + O(a).$$

Now set this expression equal to α/m and solve. \square

PROOF OF THEOREM 3.7. Define c_* as in (14). If $u_* = \sqrt{2c_*} \leq \xi$, then the the proof proceeds as in the previous proof. So we first need to establish for

which values of ξ is this true. Let $r(c) = \gamma \bar{\Phi}(\sqrt{2c}) + a \bar{\Phi}(\xi/2 + c/\xi)$. We want to find out when the solution of $r(c) = \alpha/m$ is such that $\sqrt{2c} \leq \xi$, or, equivalently, $c \leq \xi^2/2$. Now r is decreasing in c . Since $\gamma + a \geq \alpha/m$, $r(-\infty) \geq \alpha/m$. Hence, there is a solution with $c \leq \xi^2/2$ if and only if $r(\xi^2/2) \leq \alpha/m$. But $r(\xi^2/2) = (\gamma + a)\bar{\Phi}(\xi)$, so we conclude that there is such a solution if and only if $(\gamma + a)\bar{\Phi}(\xi) \leq \alpha/m$, that is, $\xi \geq z_{\alpha/(m(\gamma+a))} = \xi_0$.

Now suppose that $\xi < \xi_0$. We need to find $u \leq \xi$ to make c as large as possible in the equation $v(u, c) \equiv \gamma \bar{\Phi}(u/2 + c/u) + a \bar{\Phi}(\xi/2 + c/\xi) = \alpha/m$. Let $u_* = \xi$ and $c_* = \xi z_{\alpha/(m(\gamma+a))} - \xi^2/2$. By direct substitution, $v(u_*, c_*) = \alpha/m$ for this choice of u and c and, clearly, $u_* \leq \xi$ as required. We claim that this is the largest possible c_* . To see this, note that $v(u, c) < v(u, c_*)$. For $\xi \leq \xi_0$, $v(u, c_*)$ is a decreasing function of u . Hence, $v(u, c) < v(u, c_*) \leq v(u_*, c_*) = \alpha/m$. This contradicts the fact that $v(u, c) = \alpha/m$.

For the second claim, note that the power of the weighted test beats the power of Bonferroni if and only if the weight $w = (m/\alpha)\bar{\Phi}(\xi/2 + C(\xi)/2) \geq 1$, which is equivalent to

$$(16) \quad C(\xi) \leq \xi z_{\alpha/m} - \xi^2/2.$$

When $\xi \leq \xi_0$, $C(\xi) = \xi \xi_0 - \xi^2/2$. By assumption, $\gamma + a \leq 1$ so that $z_{\alpha/(m(\gamma+a))} \leq z_{\alpha/m}$ and now suppose that $\xi_0 < \xi \leq \xi_*$. Then $C(\xi)$ is the solution to $r(c) = \gamma \bar{\Phi}(\sqrt{2c}) + a \bar{\Phi}(\xi/2 + c/\xi) = \alpha/m$. We claim that (16) still holds. Suppose not. Then, since $r(c)$ is decreasing in c , $r(\xi z_{\alpha/m} - \xi^2/2) > r(C(\xi)) = \alpha/m$. But, by direct calculation, $r(\xi z_{\alpha/m} - \xi^2/2) > \alpha/m$ implies that $\xi > \xi_*$, which is a contradiction. Thus, (7) holds.

Finally, we turn to (8). In this case, $C(\xi) = z_{\alpha/(m\gamma)}^2/2 + O(a)$. The worst case power is $\bar{\Phi}(C(\xi)/\xi - \xi/2) = \bar{\Phi}(z_{\alpha/(m\gamma)}^2/(2\xi) - \xi/2) + O(a)$. The latter is increasing in ξ and so is at least $\bar{\Phi}(z_{\alpha/(m\gamma)}^2/(2\xi_*) - \xi_*/2) + O(a) = \bar{\Phi}((z_{\alpha/(m\gamma)}^2/(2\xi_*) - \xi_*^2)/(2\xi_*)) + O(a)$, as claimed. The next two equations follow from standard tail approximations for Gaussians. Specifically, a Gaussian quantile $z_{\beta/m}$ can be written as $z_{\beta/m} = \sqrt{2 \log(mL_m/\beta)}$, where $L_m = c \log^a(m)$ for constants a and c [Donoho and Jin (2004)]. Inserting this into the previous expression yields the final expression. \square

PROOF OF THEOREM 4.2. Setting $\pi(w, \xi_m) = \bar{\Phi}(\bar{\Phi}^{-1}(w\alpha/m) - \xi_m)$ equal to $1 - \beta$ implies $w = (m/\alpha)\bar{\Phi}(z_{1-\beta} + z_{\alpha/m})$, which is equal to w_1 as stated in the theorem. The stated form of w_0 implies that

the weights average to 1. The stated solution thus satisfies the restriction that a fraction ε have power at least $1 - \beta$. Increasing the weight of any hypothesis whose weight is w_0 necessitates reducing the weight of another hypothesis. This either reduces the minimum power or forces a hypothesis with power $1 - \beta$ to fall below $1 - \beta$. Hence, the stated solution does in fact maximize the minimum power. \square

ACKNOWLEDGMENTS

The authors thank Jamie Robins for helping us to clarify several issues. Research supported in part by National Institute of Mental Health Grant MH057881 and NSF Grant AST 0434343.

REFERENCES

- 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](#)
- BENJAMINI, Y. and HOCHBERG, Y. (1997). Multiple hypotheses testing with weights. *Scand. J. Statist.* **24** 407–418. [MR1481424](#)
- 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](#)
- BENJAMINI, Y. and YEKUTIELI, D. (2001). The control of the false discovery rate in multiple testing under dependency. *Ann. Statist.* **29** 1165–1188. [MR1869245](#)
- BICKEL, P., KLAASSEN, C., RITOV, Y. and WELLNER, J. (1993). Efficient and adaptive estimation for semiparametric models. Technical report, Johns Hopkins Series in the Mathematical Statistics, Baltimore, Maryland. [MR1245941](#)
- BLANCHARD, G. and ROQUAIN, E. (2008). Two simple sufficient conditions for FDR control. *Electron. J. Stat.* **2** 963–992. [MR2448601](#)
- BLANCHARD, G. and ROQUAIN, E. (2009). Adaptive FDR control under independence and dependence. *J. Mach. Learn. Res.* To appear.
- CHEN, J. J., LIN, K. K., HUQUE, M. and ARANI, R. B. (2000). Weighted p -value adjustments for animal carcinogenicity trend test. *Biometrics* **56** 586–592.
- DONOHO, D. and JIN, J. (2004). Higher criticism for detecting sparse heterogeneous mixtures. *Ann. Statist.* **32** 962–994. [MR2065195](#)
- EFRON, B. (2007). Simultaneous inference: When should hypothesis testing problems be combined? *Ann. Appl. Statist.* **2** 197–223. [MR2415600](#)
- EFRON, B., TIBSHIRANI, R., STOREY, J. D. and TUSHER, V. (2001). Empirical Bayes analysis of a microarray experiment. *J. Amer. Statist. Assoc.* **96** 1151–1160. [MR1946571](#)
- GENOVESE, C. and WASSERMAN, L. (2002). Operating characteristics and extensions of the false discovery rate procedure. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **64** 499–517. [MR1924303](#)
- GENOVESE, C. R., ROEDER, K. and WASSERMAN, L. (2006). False discovery control with p -value weighting. *Biometrika* **93** 509–524. [MR2261439](#)
- HASTIE, T., TIBSHIRANI, R. and FRIEDMAN, J. (2001). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, New York. [MR1851606](#)
- HOLM, S. (1979). A simple sequentially rejective multiple test procedure. *Scand. J. Statist.* **6** 65–70. [MR0538597](#)
- IONITA-LAZA, I., MCQUEEN, M., LAIRD, N. and LANGE, C. (2007). Genomewide weighted hypothesis testing in family-based association studies, with an application to a 100K scan. *Am. J. Hum. Genet.* **81** 607–614.
- KROPP, S., LÄUTER, J., ESZLINGER, M., KROHN, K. and PASCHKE, R. (2004). Nonparametric multiple test procedures with data-driven order of hypotheses and with weighted hypotheses. *J. Statist. Plann. Inference* **125** 31–47. [MR2086887](#)
- MCQUEEN, M. and colleagues (2008). Personal communication.
- ROEDER, K., BACANU, S.-A., WASSERMAN, L. and DEVLIN, B. (2006). Using linkage genome scans to improve power of association in genome scans. *Am. J. Hum. Genet.* **78** 243–252.
- ROEDER, K., WASSERMAN, L. and DEVLIN, B. (2007). Improving power in genome-wide association studies: Weights tip the scale. *Genet. Epidemiol.* **31** 741–747.
- ROMANO, J. P., SHAIKH, A. M. and WOLF, M. (2008). Control of the false discovery rate under dependence using the bootstrap and subsampling. *TEST* **17** 417–442. [MR2470085](#)
- ROQUAIN, E. and VAN DE WIEL, M. (2008). Multi-weighting for FDR control. Available at [arXiv:0807.4081](#).
- ROSENTHAL, R. and RUBIN, D. (1983). Ensemble-adjusted p -values. *Psychol. Bull.* **94** 540–541.
- RUBIN, D., DUDOIT, S. and VAN DER LAAN, M. (2006). A method to increase the power of multiple testing procedures through sample splitting. *Stat. Appl. Genet. Mol. Biol.* **5**, Art. 19 (electronic). [MR2240850](#)
- SABATTI, C., SERVICE, S. and FREIMER, N. (2003). False discovery rate in linkage and association genome screens for complex disorders. *Genetics* **164** 829–833.
- SACCONI, S., HINRICHS, A. L., SACCONI, N., CHASE, G., KONVICKA, K., MADDEN, P., BRESLAU, N., JOHNSON, E., HATSUKAMI, D., POMERLEAU, O., SWAN, G., GOATE, A., RUTTER, J., BERTELSEN, S., FOX, L., FUGMAN, D., MARTIN, N., MONTGOMERY, G., WANG, J., BALLINGER, D., RICE, J. and BIERUT, L. (2007). Cholinergic nicotinic receptor genes implicated in a nicotine dependence association study targeting 348 candidate genes with 3713 SNPs. *Hum. Mol. Genet.* **16** 36–49.
- SARKAR, S. and HELLER, R. (2008). Comments on: Control of the false discovery rate under dependence using the bootstrap and subsampling. *TEST* **17** 450–455. [MR2470085](#)
- SARKAR, S. K. (2002). Some results on false discovery rate in stepwise multiple testing procedures. *Ann. Statist.* **30** 239–257. [MR1892663](#)
- SATAGOPAN, J. and ELSTON, R. (2003). Optimal two-stage genotyping in population-based association studies. *Genet. Epidemiol.* **25** 149–157.
- SCHUSTER, E., KROPP, S. and ROEDER, I. (2004). Micro array based gene expression analysis using parametric multivariate tests per gene—a generalized application of multiple procedures with data-driven order of hypotheses. *Biom. J.* **46** 687–698. [MR2108612](#)
- SIGNORAVITCH, J. (2006). Optimal multiple testing under the general linear model. Technical report, Harvard Biostatistics.
- SKOL, A., SCOTT, L., ABECASIS, G. and BOEHNKE, M. (2006). Joint analysis is more efficient than replication-based analysis

- for two-stage genome-wide association studies. *Nat. Genet.* **38** 390–394.
- SPJØTVOLL, E. (1972). On the optimality of some multiple comparison procedures. *Ann. Math. Statist.* **43** 398–411. [MR0301871](#)
- STOREY, J. D. (2002). A direct approach to false discovery rates. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **64** 479–498. [MR1924302](#)
- STOREY, J. D. (2007). The optimal discovery procedure: A new approach to simultaneous significance testing. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **69** 347–368. [MR2323757](#)
- STOREY, J. and TIBSHIRANI, R. (2003). Statistical significance for genome-wide studies. *Proc. Natl. Acad. Sci. USA* **100** 9440–9445. [MR1994856](#)
- SUN, W. and CAI, T. T. (2007). Oracle and adaptive compound decision rules for false discovery rate control. *J. Amer. Statist. Assoc.* **102** 901–912. [MR2411657](#)
- TODD, J., WALKER, N., COOPER, J., SMYTH, D., DOWNES, K., PLAGNOL, V., BAILEY, R., NEJENTSEV, S., FIELD, S., PAYNE, F., LOWE, C., SZESZKO, J., HAFLER, J., ZEITELS, L., YANG, J., VELLA, A., NUTLAND, S., STEVENS, H., SCHUILENBURG, H., COLEMAN, G., MAISURIA, M., MEADOWS, W., SMINK, L. J., HEALY, B., BURREN, O., LAM, A., OVERTON, N., ALLEN, J., ADLEM, E., LEUNG, H., WALLACE, C., HOWSON, J., GUJA, C., IONESCU-TIRGOVI, C., GENETICS OF TYPE 1 DIABETES IN FINLAND, SIMMONDS, M., HEWARD, J., GOUGH, S., WELLCOME TRUST CASE CONTROL CONSORTIUM, DUNGER, D., WICKER, L. and CLAYTON, D. (2007). Robust associations of four new chromosome regions from genome-wide analyses of type 1 diabetes. *Nat. Genet.* **39** 857–864.
- WANG, H., THOMAS, D., PE'ER, I. and STRAM, D. (2006). Optimal two-stage genotyping designs for genome-wide association scans. *Genet. Epidemiol.* **30** 356–368.
- WANG, K., LI, M. and BUCAN, M. (2007). Pathway-based approaches for analysis of genomewide association studies. *Am. J. Hum. Genet.* **81** 1278–1283.
- WESTFALL, P., KRISHEN, A. and YOUNG, S. (1998). Using prior information to allocate significance levels for multiple endpoints. *Stat. Med.* **17** 2107–2119.
- WESTFALL, P. H. and KRISHEN, A. (2001). Optimally weighted, fixed sequence and gatekeeper multiple testing procedures. *J. Statist. Plann. Inference* **99** 25–40. [MR1858708](#)
- WESTFALL, P. H., KROPF, S. and FINOS, L. (2004). Weighted FWE-controlling methods in high-dimensional situations. In *Recent Developments in Multiple Comparison Procedures. IMS Lecture Notes Monogr. Ser.* **47** 143–154. IMS, Beachwood, OH. [MR2118598](#)
- WESTFALL, P. H. and SOPER, K. A. (2001). Using priors to improve multiple animal carcinogenicity tests. *J. Amer. Statist. Assoc.* **96** 827–834. [MR1963409](#)