BAYESIAN MULTIPLE TESTING UNDER SPARSITY FOR POLYNOMIAL-TAILED DISTRIBUTIONS

Xueying Tang[1], Ke Li[2] and Malay Ghosh[1]

[1]*University of Florida and* [2]*Southwestern University of Finance and Economics*

Abstract: This paper considers Bayesian multiple testing under sparsity for polynomial-tailed distributions satisfying a monotone likelihood ratio property. Included in this class of distributions are the Student's t, the Pareto, and many other distributions. We prove some general asymptotic optimality results under fixed and random thresholding. As examples of these general results, we establish the Bayesian asymptotic optimality of several multiple testing procedures in the literature for appropriately chosen false discovery rate levels. We also show by simulation that the Benjamini-Hochberg procedure with a false discovery rate level different from the asymptotically optimal one can lead to high Bayes risk.

Key words and phrases: Asymptotic optimality, Benjamini-Hochberg procedure, false discovery rate, Pareto distribution, Student's t distribution.

1. Introduction

Multiple testing has become a topic of growing importance in recent years. Its importance is particularly felt in the event of inference under sparsity, detecting a few signals in the midst of multiple noises. Applications abound, for example, in genetics, engineering, biology, and finance, just to name a few. A specific example is when one needs to identify a handful of genes attributable to a certain disease in the midst of thousands of others.

Currently, the most widely used approach for multiple testing is the one due to Benjamini and Hochberg (1995) that controls the false discovery rate (FDR). Since then, there are several noteworthy contributions in the area. Among others, we refer to Benjamini and Yekutieli (2001), Efron and Tibshirani (2002), Abramovich et al. (2006), Donoho and Jin (2006), Gavrilov, Benjamini and Sarkar (2009), Genovese and Wasserman (2002, 2004), Sarkar (2002), and Storey (2002).

Recently, Bogdan, Ghosh and Tokdar (2008) conducted an extensive simulation study to find closeness of the Benjamini-Hochberg procedure to an optimal Bayes procedure for multiple hypothesis testing under normality of the data.

Later, in Bogdan et al. (2011) and Frommlet and Bogdan (2013), it was shown that several multiple testing procedures, including the Benjamini-Hochberg procedure, asymptotically attained the Bayes oracle property under sparsity, once again under normality. As an extension of Bogdan et al. (2011), Neuvial and Roquain (2012) studied properties of FDR thresholding with observations coming from the Subbotin family that includes Laplace and normal distributions as special cases.

Here we consider distributions with polynomial tails. As we show later, the Bayes rules of the multiple testing problem of normal distributions and polynomial-tailed distributions produce quite different Bayes risks. Both type I and type II errors play a role in the limiting (as the number of tests goes to infinity) Bayes risk of the oracle multiple testing procedure for polynomial-tailed distributions while, as shown in Bogdan et al. (2011), the Bayes risk is asymptotically determined solely by type II errors for normal distributions. As indicated by Chi (2007), in multiple testing problems for polynomial-tailed distributions, controlling FDR under a certain threshold level α^* leads to asymptotically zero power. As a result, a vanishing FDR level is very unlikely to define an asymptotically optimal procedure in terms of Bayes risk. As this is not the case for normal distributions, we were motivated to study the performance of the Benjamini-Hochberg procedure and some other multiple testing procedures for polynomial-tailed distributions. We study the asymptotic optimality of multiple testing procedures for a general class of such distributions, including Student's t, Pareto, and many others.

Our framework follows that of Bogdan et al. (2011) and Neuvial and Roquain (2012), where the multiple testing problem is addressed in a decision theoretic framework. Suppose there are m independent observations, each of which comes from a mixture of two distributions in the scale family of a polynomial-tailed distribution. We are interested in testing simultaneously which distribution each observation comes from. We assume the loss of wrong decisions for the m tests is the sum of the losses of wrong decisions for individual tests (Lehmann (1957a,b)). For each test, nonzero loss occurs if and only if a type I or a type II error is made. All our results are obtained in an asymptotic framework that ensures that the limiting power of an individual test based on the Bayes oracle threshold converges to a constant between zero and one.

After finding the oracle Bayes risk, we define asymptotic Bayesian optimality under sparsity (ABOS) analogous to Bogdan et al. (2011). Under the asymptotic framework, a necessary and sufficient condition is provided for a fixed threshold-

ing procedure to be ABOS. A single constraint guarantees that the risks from the two error types converge to the optimal risk for our proposed class of distributions.

A more practically meaningful result is that we obtain a general sufficient condition for a random thresholding procedure to obtain ABOS. The condition requires comparison of a random threshold with a fixed ABOS threshold that is sometimes easier to work with than bounding the Bayes risk directly, as in Bogdan et al. (2011). Our general results show that the procedures controlling Bayesian false discovery rate (BFDR), the procedure of Genovese and Wasserman (2002), and the Benjamini-Hochberg procedure are all ABOS if the FDR level is chosen properly. On the other hand, it is shown via simulation that the Benjamini-Hochberg procedure with an FDR level different from the optimal one can lead to high Bayes risk.

The remaining sections of this article are organized as follows. In Section 2, we describe our asymptotic framework. The Bayes oracle rule and its Bayes risk are derived in this section, and we list a few important distributions with polynomial tails. Two general results about the conditions for fixed or random thresholding procedures to be ABOS are presented in Section 3. In Section 4, we provide conditions under which several procedures, including Benjamini-Hochberg, are ABOS by applying the results in Section 3. Section 5 contains numerical results suggesting the non-optimality of the Benjamini-Hochberg procedure. Some final remarks are made in Section 6. Proofs of theoretical results are provided in the supplementary material.

2. Oracle Bayes Risk and Asymptotic Framework

Suppose we have m independent observations $\mathbf{X} = (X_1, \ldots, X_m)$ from the same distribution. Let D and d be the cumulative distribution function (cdf) and the probability density function (pdf) with respect to Lebesgue measure of a distribution from a monotone polynomial tail (MPT) distribution family defined as follows.

Definition 1. *A distribution with cdf D and pdf d is said to be an MPT distribution if d is either an even function or a function whose support is the nonnegative real line, $d(x)x^{\gamma+1} \to C_d$ as $x \to \infty$ for some constant $\gamma > 0$ and $C_d > 0$, and for any $\theta > 1$, $d(x/\theta)/d(x)$ is a strictly increasing function in x for $x > 0$.*

The study of decision procedures for distributions with general monotone likelihood ratio (MLR) properties dates back to Karlin and Rubin (1956). The

Table 1. Some common distributions in the MPT family.

	pdf	C_d
Student's t	$d(x) = C_d(x^2 + \gamma)^{-(\gamma+1)/2}, \ -\infty < x < \infty$	$\dfrac{\gamma^{\gamma/2}\Gamma((\gamma+1)/2)}{\pi^{1/2}\Gamma(\gamma/2)}$
Pareto	$d(x) = C_d(x+1)^{-(\gamma+1)}, \ x > 0$	γ
Inverse Gamma	$d(x) = C_d x^{-\gamma-1}\exp(-1/x), \ x > 0$	$1/\Gamma(\gamma)$

MLR property here ensures a simple form for the Bayes rule. These distributions have polynomial tails with γ, the polynomial tail heaviness index, specifying the heaviness of the tail. By L'Hospital's rule, one has $x^\gamma\{1 - D(x)\} \to C_d/\gamma$. We focus on symmetric MPT distributions to make comparisons with normal distributions. For symmetric MPT, the scale family with pdf d has MLR property in $|x|$. The MPT family includes many important distributions. Some examples are given in Table 1.

Proposition 1. *A distribution in the MPT family satisfies*

    1. for $x > 0$, $g(x) = d(x)x^{\gamma+1}$ is a strictly increasing function;

    2. for $x > 0$, $h(x) = x^\gamma\{1 - D(x)\}$ is a strictly increasing function;

    3. if $\theta > 1$, $\{1 - D(x/\theta)\}/\{1 - D(x)\}$ is a strictly increasing function for $x > 0$;

    4. for $x > 0$, $f(x) = 2xd(x)/\gamma + 2D(x) - 1$ is a strictly increasing function.

We assume that the common distribution of the X's has cdf

$$pD(\frac{x}{\sigma_1}) + (1 - p)D(\frac{x}{\sigma_0}) \quad (0 < \sigma_0 < \sigma_1). \tag{2.1}$$

The mixture has the distribution of X_i determined by a latent Bernoulli random variable s_i with success probability p, and the latent variables $s_1, \ldots, s_m$ are mutually independent. The cdf of X_i is $D(x/\sigma_0)$ if $s_i = 0$ and $D(x/\sigma_1)$ if $s_i = 1$. We are interested in simultaneously testing

$$H_{0i} : s_i = 0 \text{ vs. } H_{Ai} : s_i = 1, \ i = 1, \ldots, m.$$

In a Bayesian hypothesis testing framework, the marginal distribution of the observations is often of the form (2.1). A specific example of our model is in stock selection, where $X_1, \ldots, X_m$ are the log returns of m stocks, often modeled by Student's t distributions. We can equivalently assume $X_i \,|\, \mu_i, \tau_i^2 \sim N(\mu_i, \tau_i^2)$ and $\tau_i^2 \sim \text{IG}(\gamma/2, \gamma/2)$. To test whether some stocks have extreme returns, if we assume

$$\mu_i \,|\, \tau_i^2 \sim (1 - p)N(0, \eta_0^2\tau_i^2) + pN(0, \eta_1^2\tau_i^2), \ i = 1, \ldots, m,$$

then the marginal distribution of X_i is a mixture of two Student's t distributions with γ degrees of freedom and $\sigma_j^2 = \eta_j^2 + 1$, $j = 0, 1$.

Let $r_i : \mathbb{R}^m \to \{0, 1\}$ be the decision rule used for the ith test. If $r_i(\mathbf{X}) = 1$, the null hypothesis is rejected; otherwise the null is not rejected. For each test, the loss is non-zero only if $r_i \neq s_i$, that is, only when a type I or type II error is made. Let δ_0 and δ_A denote the respective losses of making a type I and a type II error. We assume that the overall loss of the m tests is the sum of losses for individual tests. This additive loss structure is similar to the one in Lehmann (1957a,b) and in Bogdan et al. (2011). To simplify matters, we take

$$1 + u = \frac{\sigma_1^2}{\sigma_0^2}, \quad \delta = \frac{\delta_0}{\delta_A}, \quad f = \frac{1-p}{p}, \quad v = \delta f.$$

These parameters can vary with the number of tests m.

The Bayes risk of a multiple testing procedure is

$$R = p\delta_A \sum_{i=1}^{m} (vt_{1i} + t_{2i}), \tag{2.2}$$

where t_{1i} and t_{2i} denote the probabilities of type I and type II errors for the ith test; hereafter we call vt_{1i} and t_{2i} the type I and type II risk components, respectively. The Bayes rule minimizing the Bayes risk can be shown to reject H_{0i} if

$$\frac{(1+u)^{-1/2} d\left((1+u)^{-1/2} X_i/\sigma_0\right)}{d(X_i/\sigma_0)} > v, \ i = 1, \ldots, m.$$

Due to the MLR property, the Bayes rule rejects H_{0i} if $X_i^2/\sigma_0^2 \geq \omega_{\text{opt}}^2$, where ω_{opt} is the positive solution of the equation

$$\frac{(1+u)^{-1/2} d\left(\omega_{\text{opt}}(1+u)^{-1/2}\right)}{d(\omega_{\text{opt}})} = v. \tag{2.3}$$

As p is unknown in practice, we call ω_{opt} the oracle threshold.

We seek conditions for a multiple testing procedure to attain the Bayes oracle property under sparsity as $m \to \infty$. To impose sparsity, we assume $p \to 0$ as $m \to \infty$, and let $u \to \infty$ to ensure that the signals are strong enough to be discovered as $m \to \infty$. We assume that $\omega_{\text{opt}}^2/(1 + u) \to C \in (0, \infty)$ to avoid having the power of an individual test going to zero or one. By (2.3) and the definition of MPT family, this is equivalent to assuming $vu^{-\gamma/2} \to C_0$, where

$$C_0 = C^{(\gamma+1)/2} \frac{d(\sqrt{C})}{C_d} \tag{2.4}$$

is a strictly increasing function in C, according to Proposition 1. If $\delta \to \delta_\infty = 1$,

this can be simplified to $pu^{\gamma/2} \to C_0^{-1}$. Intuitively, it is more difficult to distinguish between signals and noises if the data have fewer signals (smaller p) and smaller signal-to-noise ratio (smaller u). The assumption that $pu^{\gamma/2}$ converges to a constant guarantees the signals are identifiable, while the magnitude of the constant C_0^{-1} indicates the intrinsic difficulty in identifying those signals. A larger C reflects more difficulties and we call C the difficulty index.

To summarize, we study the properties of multiple testing procedures under the asymptotic framework (as $m \to \infty$)

$$p \to 0, \quad u \to \infty, \quad vu^{-\gamma/2} \to C_0, \tag{2.5}$$

whereas the asymptotic framework in Bogdan et al. (2011) has $p \to 0, u \to \infty, v \to \infty, 2\log(v)/u \to C$. Noticing that the second and the third assumptions in (2.5) imply $v \to \infty$, the only difference between the two frameworks is the relation between u and v. For normal distributions, the rate at which the signal strength u increases to infinity is the logarithm of v, while for polynomial-tailed distributions, it is a polynomial in v.

Proposition 2. *Let $C_1 = 2\sqrt{C}d(\sqrt{C})/\gamma$ and $C_2 = 2D(\sqrt{C}) - 1$. Under (2.5),*

$$\omega_{opt}^2 \sim C(\frac{v}{C_0})^{2/\gamma}, \quad t_{1i} = t_1 \sim C_1 v^{-1}, \quad t_{2i} = t_2 \sim C_2, \tag{2.6}$$

with C_0 as in (2.4). The corresponding Bayes risk is

$$R_{opt} \sim mp\delta_A(C_1 + C_2). \tag{2.7}$$

By Proposition 1, R_{opt} is a strictly increasing function in C, which agrees with the interpretation of the difficulty index C; a more difficult multiple testing task leads to a higher Bayes risk.

Since C_1 and C_2 are the limits of type I and type II risk components of the oracle procedure, we call them asymptotically optimal type I and type II risk components. In Bogdan et al. (2011) and Neuvial and Roquain (2012), the limiting Bayes risks of the oracle threshold are shown to depend solely on the type II risk component. For polynomial-tailed distributions, neither risk component of the oracle procedure is negligible as the number of tests goes to infinity. In both models, the oracle probability of type I errors goes to zero, but the probability decays at the rate of v^{-1} for polynomial-tailed models, while for the normal model, the rate is faster. Besides the need for stronger signals to ensure detectability, this is yet another effect of heavy-tailed signals and noises.

Definition 2. *A multiple testing rule is asymptotically Bayes optimal under sparsity (ABOS) under (2.5) if its Bayes risk R satisfies $R/R_{opt} \to 1$.*

3. Fixed and Random Thresholding Procedures

In this section, we consider multiple testing procedures that reject the ith null hypothesis if X_i^2/σ_0^2 is greater than or equal to a threshold, which can be either non-data dependent (fixed) or data dependent (random). To distinguish them, we let ω^2 denote the fixed threshold and $\widehat{\omega}^2$ denote the random threshold. For a fixed thresholding procedure, the events $\{r_i(\mathbf{X}) = 1 \mid s_i = 0\}$ and $\{r_i(\mathbf{X}) = 0 \mid s_i = 1\}$ are based only on the ith observation X_i with respective probabilities the same for each i. Therefore, the Bayes risk of a fixed thresholding procedure can be expressed as $R = mp\delta_A(vt_1 + t_2)$, where $t_1 = 2\{1 - D(\omega)\}$, and $t_2 = 2D(\omega(1 + u)^{-1/2}) - 1$. In contrast, the events $\{r_i(\mathbf{X}) = 1 \mid s_i = 0\}$ and $\{r_i(\mathbf{X}) = 0 \mid s_i = 1\}$ for a random thresholding procedure are potentially based on all m observations and the probabilities of type I and type II errors are not necessarily the same across different tests.

Theorem 1. *A fixed thresholding multiple testing procedure that rejects H_{0i} when $X_i^2/\sigma_0^2 \geq \omega^2, i = 1, \ldots, m$ is ABOS if and only if the threshold ω satisfies*

$$\frac{\omega^2}{\omega_{opt}^2} \to 1, \tag{3.1}$$

or, equivalently, with C_0 as in (2.4),

$$\omega^2 = C \left(\frac{v}{C_0}\right)^{2/\gamma} (1 + o(1)). \tag{3.2}$$

It may appear that even if the type I and type II risk components do not tend to the corresponding asymptotically optimal risk components, there is still a chance that $R \sim R_{\mathrm{opt}}$. However, the proof shows that the two components have to converge to the corresponding optimal risk components individually in order to achieve ABOS. This observation is also true for the normal distribution, but as shown by Theorem 3.2 in Bogdan et al. (2011), two conditions, one for the type II risk component, the other for the type I risk component, are needed to guarantee this, while in our case, only one condition is required. In Remark 3.1 of Bogdan et al. (2011), the authors argued the reason for the extra condition for type I error is that, for normal models, type I errors are more sensitive to changes in the critical value than type II errors. In their language, our Theorem 1 shows that for polynomial-tailed models, type I and type II errors are equally sensitive to changes in the critical value.

Theorem 2. *Under (2.5), a random thresholding multiple testing procedure that rejects H_{0i} if $X_i^2/\sigma_0^2 \geq \widehat{\omega}^2$ is ABOS if for all $\epsilon > 0$,*

$$\frac{1}{m} \sum_{i=1}^{m} P(|\widehat{\omega} - \omega_{opt}| > \epsilon v^{1/\gamma} | s_i = 0) = o(v^{-1}), \tag{3.3}$$

$$\frac{1}{m} \sum_{i=1}^{m} P(|\widehat{\omega} - \omega_{opt}| > \epsilon v^{1/\gamma} | s_i = 1) = o(1). \tag{3.4}$$

If δ does not converge to zero as $m \to \infty$, then a random thresholding procedure is ABOS if for all $\epsilon > 0$,

$$P(|\frac{\widehat{\omega}}{\omega_{opt}} - 1| > \epsilon) = o(v^{-1}). \tag{3.5}$$

An equivalent condition to (3.5) is that $P(|\widehat{\omega} - \omega_{\mathrm{opt}}| > \epsilon v^{1/\gamma}) = o(v^{-1})$. Theorem 2 continues to hold if the oracle threshold is replaced by the threshold of a fixed thresholding procedure that is ABOS. The left hand sides of (3.3) and (3.4) can be interpreted as the average departures of the probabilities of type I and type II errors of a random thresholding procedure from the corresponding errors of the Bayes oracle. Although less general than (3.3) and (3.4), condition (3.5) could be easier to verify in practice because of the symmetry of the distribution of m observations.

As implied by Theorem 1, to obtain an ABOS fixed thresholding procedure, the fixed threshold itself is very likely to contain unknown parameters. In contrast, a random threshold consists of observed data only. For example, it could appear as an estimator of an ABOS fixed threshold. Therefore, a random thresholding procedure is naturally an implementable procedure, and, in this sense, Theorem 2 provides a more practical result.

4. ABOS of Several Special Procedures

4.1. Procedures controlling BFDR

Benjamini and Hochberg (1995) introduced FDR as a less stringent error measure than the familywise error rate, $\mathrm{FDR} = E(V/R)$, where R is the number of total rejections and V is the number of false rejections. Storey (2003) argued that the positive false discovery rate (pFDR), $\mathrm{pFDR} = E(V/R \mid R > 0)$, can overcome some of the concerns in Benjamini and Hochberg (1995) and, under certain conditions, it coincides with the Bayesian false discovery rate (BFDR) of Efron and Tibshirani (2002),

$$\mathrm{BFDR} = P(H_{0i} \text{ is true} \mid H_{0i} \text{ was rejected}) = \frac{(1-p)t_1}{(1-p)t_1 + p(1-t_2)}.$$

For a fixed thresholding procedure, the threshold ω and the BFDR level α are

linked by

$$\frac{(1-p)\{1-D(\omega)\}}{(1-p)\{1-D(\omega)\}+p\{1-D(\omega(1+u)^{-1/2})\}} = \alpha, \tag{4.1}$$

or equivalently,

$$\frac{1-D(\omega)}{1-D\left(\omega(1+u)^{-1/2}\right)} = \frac{r_\alpha}{f}, \tag{4.2}$$

where $r_\alpha = \alpha/(1-\alpha)$. Since we have already found a necessary and sufficient condition for the fixed thresholding procedure to be ABOS, by using (4.2), we are able to find conditions on α (depending on m) such that the BFDR controlling procedure is ABOS.

An alternative expression of the BFDR level α in (4.1) is that $\alpha = (1-p)\{1-p+pg(\omega)\}^{-1}$, where $g(\omega) = \{1-D(\omega)\}^{-1}\{1-D(\omega(1+u)^{-1/2})\}$. Property 3 of Proposition 1 shows that g is a strictly increasing function in ω, and, since $1-D(x) \sim x^{-\gamma}$ as $x \to \infty$, $g(\omega) \to (1+u)^{\gamma/2}$ as $\omega \to \infty$. As $g(\omega) = 1$ if $\omega = 0$, the BFDR of a finite fixed threshold procedure for a given m can only be controlled within the interval $I = (\beta^*, 1-p]$, where

$$\beta^* = \{1 + \frac{(1+u)^{\gamma/2}}{f}\}^{-1}. \tag{4.3}$$

Thus with a BFDR level less than β^*, the fixed threshold has to be infinite. In this case, none of the m null hypotheses is rejected and the power of an individual test is zero. Since our asymptotic framework requires the power of an individual test to go to a nonzero constant, we confine α to the interval I.

To distinguish from general fixed thresholds, $\omega_{B_\alpha}^2$ is used to denote the fixed threshold controlling BFDR under α, and the subscript α is omitted if there is no ambiguity.

Proposition 3. *A fixed thresholding procedure controlling* BFDR *under α is ABOS if and only if*

$$\delta r_\alpha \to \frac{C_1}{(1-C_2)}. \tag{4.4}$$

The threshold is of the form

$$\omega_B^2 = C_B \left(\frac{f}{r_\alpha}\right)^{2/\gamma} (1+o(1)), \tag{4.5}$$

where $C_B = [(C_d/\gamma)/\{1-D(\sqrt{C})\}]^{2/\gamma}$ and C_1, C_2 are as in Proposition 2.

Condition (4.4) implies that if either one of δ and α goes to a positive constant, the other is forced to converge to a positive constant as well. For example,

if δ converges to a positive constant δ_∞, then $\alpha \to \alpha_\infty$ where α_∞ is defined by

$$\alpha_\infty = \frac{1}{1 + \delta_\infty(1 - C_2)/C_1}. \tag{4.6}$$

Also, as more penalty is imposed for type II errors ($\delta \to 0$) as $m \to \infty$, then no control is made on BFDR since $\alpha \to 1$ as $m \to \infty$.

4.2. Genovese-Wasserman and Benjamini-Hochberg procedures

Let $Z_i = |X_i/\sigma_0|$ and $p_i = 2\{1 - D(Z_i)\}$ denote the p-values for the m tests, ordered as $p_{(1)} \le \cdots \le p_{(m)}$. The Benjamini-Hochberg procedure at FDR level α then looks for the largest k, denoted by $\widehat{k}$, that satisfies $p_{(k)} \le \alpha k/m$ and rejects all the tests whose p-values are less than or equal to $p_{(\widehat{k})}$. This is equivalent to rejecting the null hypothesis H_{0i} if $Z_i^2 \ge \omega_{\mathrm{BH}}^2$, where

$$\omega_{\mathrm{BH}} = \inf \left\{ y : \frac{2\{1 - D(y)\}}{1 - \widehat{F}(y)} \le \alpha \right\}, \tag{4.7}$$

F being the common cdf of Z_i's and $1 - \widehat{F}(y) = \#\{Z_i \ge y\}/m$. Thus the Benjamini-Hochberg procedure is a random thresholding procedure. To study the ABOS of the Benjamini-Hochberg procedure via Theorem 2, we need to compare (4.7) with a fixed ABOS threshold. Genovese and Wasserman (2002) showed that the Benjamini-Hochberg procedure can be approximated by a fixed thresholding procedure whose threshold ω_{GW} is the solution of

$$\frac{1 - D(\omega_{\mathrm{GW}})}{(1 - p)\{1 - D(\omega_{\mathrm{GW}})\} + p\{1 - D(\omega_{\mathrm{GW}}(1 + u)^{-1/2})\}} = \alpha. \tag{4.8}$$

Proposition 4. *If $\alpha \not\to 1$, the rule that rejects the null hypothesis H_{0i} when $X_i^2/\sigma_0^2 \ge \omega_{GW}^2$ is ABOS if and only if (4.4) holds. In this case, with C_B as in Proposition 3,*

$$\omega_{GW}^2 = C_B \left(\frac{f}{r_\alpha} \right)^{2/\gamma} (1 + o(1)).$$

Theorem 3. *If*

$$p \propto 1/\log(m) \ or \ p \propto m^{-\kappa} \ for \ some \ 0 < \kappa < 1, \tag{4.9}$$

$$\delta \to \delta_\infty > 0, \tag{4.10}$$

the Benjamini-Hochberg procedure at FDR *level α is ABOS if $\alpha \to \alpha_\infty$, where α_∞ is at* (4.6).

The oracle Bayes rule balances type I and type II errors with the consideration of loss for each type of error. The optimal FDR level given in (4.6) is indeed

the result of balancing since it is determined by the limiting loss ratio δ_∞ and the asymptotically optimal risk components C_1, C_2.

The asymptotically optimal FDR level depends on the difficulty index C, which is usually an unknown parameter. Although not having a conclusive answer, we discuss how to find a practically usable FDR level in Section 6.

5. Simulation Results

We compared the performances of the Bayes oracle, the Benjamini-Hochberg procedure with the optimal FDR level, and the Benjamini-Hochberg procedure with FDR level $1/\log(m)$, through simulation studies. The FDR level $1/\log(m)$ was proved to be ABOS for the normal distributions by Bogdan et al. (2011). The simulation study in Ghosh et al. (2016) demonstrated its effectiveness in producing a misclassification probability curve similar to the one obtained from the oracle procedure. We considered this FDR level to illustrate the consequence of applying a multiple testing procedure regardless of the underlying distribution. We write the α-BH procedure for the Benjamini-Hochberg procedure with FDR level α.

The comparison of performances was done under two scenarios. In the first one, we took the sparsity parameter to vary with number of tests m, $p = m^{-0.5}$. We recorded the risks of the Bayes oracle and Benjamini-Hochberg procedures with different FDR levels. In the second scenario, with $m = 10^6$, we examined the behavior of the multiple testing procedures with changing values of p. In both scenarios, we fixed the parameters of the loss function, δ_A and δ_0, to be 1. We considered combinations of polynomial tail heaviness index γ and the difficulty index C, choosing from $\{3, 10\}$ and $\{0.1, 1, 10\}$, respectively. For each combination, 1,000 data sets of $X_i, i = 1, \ldots, m$ were generated from the mixture distribution (2.1) with $\sigma_0^2 = 1$, $\sigma_1^2 = \sigma_0^2(1 + u)$, where $u = (v/C_0)^{2/\gamma}$. Pareto and Student's t distributions in the MPT distribution family were considered in the simulation.

5.1. Results from scenario 1

The average risks based on 1,000 replicates were used to estimate the Bayes risks of the two procedures and to find the Bayes risk ratio of the BH procedure to the oracle. Panels (a) and (b) of Figures 1 and 2 show the plots of Bayes risk ratios against FDR level α. In the plots, the dashed vertical lines denote the asymptotically optimal FDR level α_∞ as at (4.6). When $m = 10^6$, the risk ratios at $\alpha = \alpha_\infty$ are close to 1 and almost reach the lowest point of the curve.

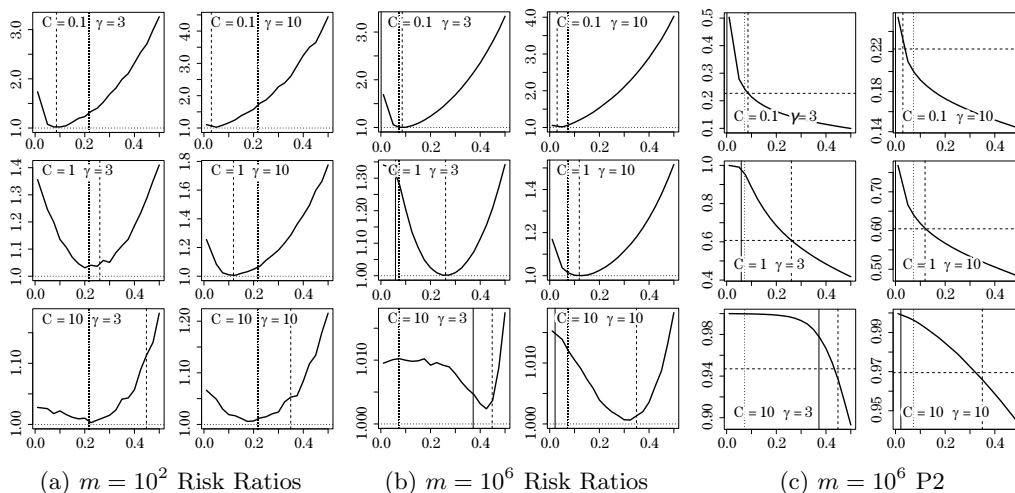(a) $m = 10^2$ Risk Ratios　　　(b) $m = 10^6$ Risk Ratios　　　(c) $m = 10^6$ P2

Figure 1. Plots of Bayes risk ratios (panels (a) and (b)) and probabilities of type II error (P2; panel (c)) against FDR level α for Student's T distributions. The dashed vertical lines denote the asymptotically optimal α and the dotted vertical lines denote $1/\log(m)$. The solid vertical lines in panels (b) and (c) denote β_∞^* defined in (6.1). The dashed horizontal lines in panel (c) represent P2 of the Bayes oracle. The ranges of vertical axes are different across plots.

When $m = 10^2$, the risk ratios at $\alpha = \alpha_\infty$ is not as close to the minimum as in the case $m = 10^6$, but the deviations are still moderate. This observation does not conflict with the asymptotic results we have established, but it suggests that the study of non-asymptotic results or the convergence rate of asymptotic results may be helpful to find a better α for smaller m. In the figure, the dotted vertical lines in the plots are $\alpha = 1/\log(m)$. In some situations, this choice of FDR level does a better job than α_∞, but it can also lead to a risk ratio away from 1 in other situations.

In the plots, the range of the Bayes risk ratios is narrower for a larger C. This is probably because the denominator of the ratio, the oracle Bayes risk, is an increasing function in C.

The optimal FDR level α_∞ (dashed vertical lines in Figures 1 and 2) increases as the difficulty index C increases. With a larger C, which signifies more difficulties in identifying signals from noises, the FDR can only be controlled at a higher level to achieve asymptotic Bayesian optimality. For both Student's t and Pareto distributions, when $C = 10$, α_∞ is close to 0.5, which could hardly provide satisfactory control of false discoveries in practice.
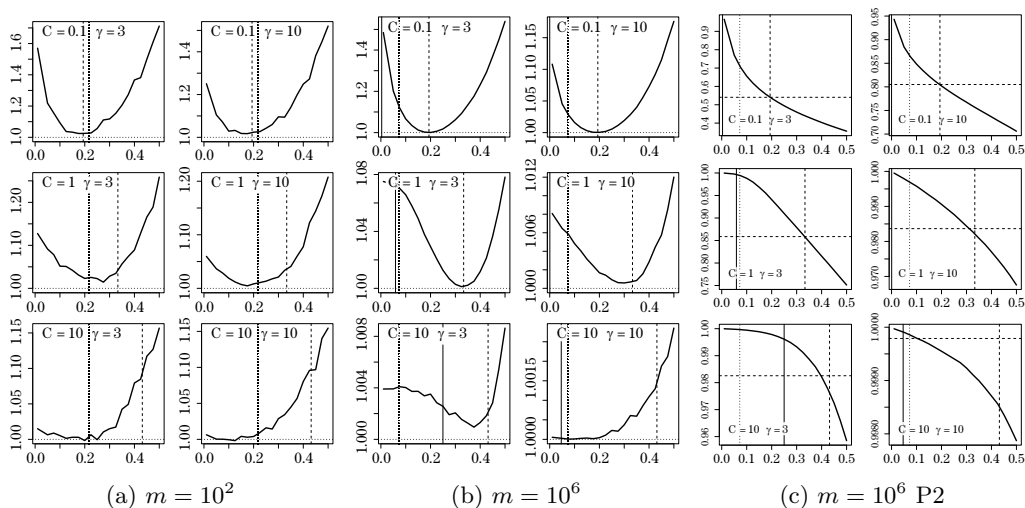
Figure 2. Plots of Bayes risk ratios (panels (a) and (b)) and probabilities of type II error (P2; panel (c)) against FDR level α for Pareto distributions. The dashed vertical lines denote the asymptotically optimal α and the dotted vertical lines denote $1/\log(m)$. The solid vertical lines in panels (b) and (c) denote β_∞^* defined in (6.1). The dashed horizontal lines in panel (c) represent P2 of the Bayes oracle. The ranges of vertical axes are different across plots.

5.2. Results from scenario 2

The average number of misclassified observations, type I, and type II errors based on 1,000 replicates were used to estimate the misclassification probability (MP), probability of type I errors (P1), and probability of type II errors (P2), respectively. Figures 3 and 4 display the plots of the three error measurements against p for Student's t and Pareto distributions, respectively. The solid, dashed, and dotted lines, respectively, represent the Bayes oracle, the α_∞-BH procedure, and the $1/\log(m)$-BH procedure. Here, the α_∞-BH procedure and the Bayes oracle behave similarly if p is small. The solid lines and the dashed lines are almost identical in most situations when $p < 0.1$. When p is large, the α_∞-BH procedure has lower P1 and higher P2 than the Bayes oracle, which suggests the former is conservative in identifying signals when they are abundant. Second, in terms of MP, for Student's t distributions, $1/\log(m)$-BH procedure works better when γ is larger since Student's t distributions with bigger degrees of freedom are closer to the normal distributions, for which the $1/\log(m)$ level is designed. In general, when applied to polynomial-tailed distributions, the $1/\log(m)$-BH procedure is more conservative in identifying signals than the α_∞-BH procedure. In the plots for $C = 1, \gamma = 3$, and $C = 10$, its P1 is in close vicinity of 0 and P2
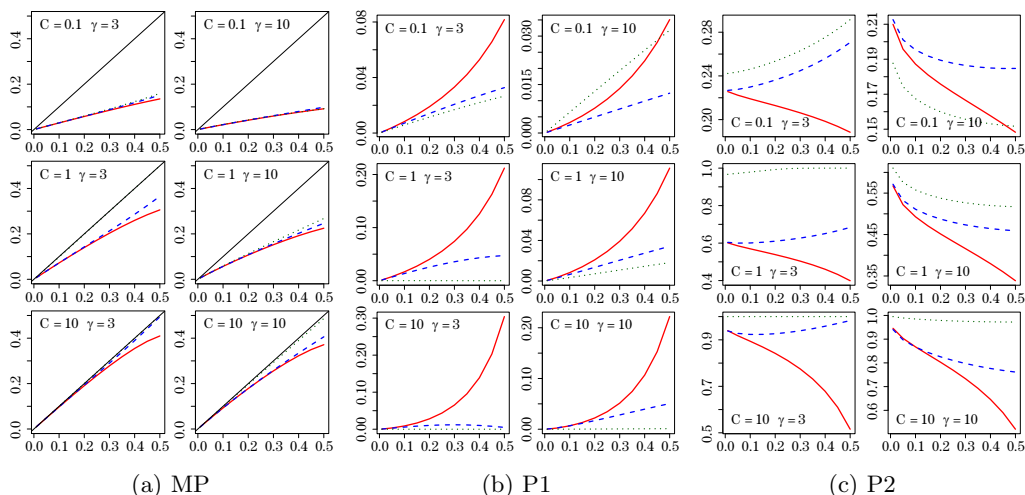
Figure 3. Plots of MP, P1, and P2 against sparsity parameter p for Student's T distributions. The red solid, blue dashed, and green dotted lines represent the Bayes oracle, the α_∞-BH procedure and the $1/\log(m)$-BH procedure respectively. In panel (a), the diagonal lines are the MP$=p$ line. The ranges of the vertical axes are different across the plots of P1 and P2.

is close to 1, which indicates the procedure identifies almost all observations as noises. In the MP panels of both figures, as C grows, the line corresponding to the Bayes oracle lies closer to the MP $= p$ line indicating increasing difficulty in multiple testing problem. This agrees with our findings in the first scenario.

6. Discussion

This paper establishes some asymptotic optimality properties of several multiple testing procedures in a Bayesian framework where the data are generated from distributions with polynomial tails. In particular, it is shown that some of the classical multiple testing procedures attain asymptotically the Bayes oracle property under sparsity. To the authors' knowledge, Theorem 2 is the first result clearly providing an approach to simplify the problem of finding an implementable random ABOS thresholding procedure to the construction of an appropriate estimator of a fixed ABOS threshold. Future work might extend results beyond polynomial-tailed distributions.

In Section 4.1, we show that, for a fixed m, the lower bound of the BFDR level that can be controlled is β^*, see (4.3), with

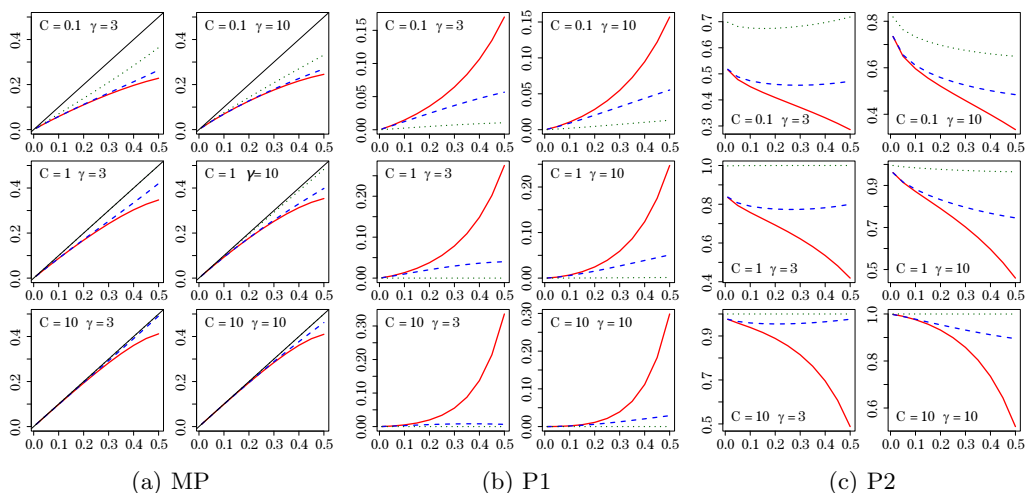$$\beta^*_\infty = \lim_{m \to \infty} \beta^* = (1 + \frac{\delta_\infty}{C_0})^{-1}. \tag{6.1}$$

Figure 4. Plots of MP, P1, and P2 against sparsity parameter p for Pareto distributions. The red solid, blue dashed, and green dotted lines represent the Bayes oracle, the α_∞-BH procedure and the $1/\log(m)$-BH procedure respectively. In panel (a), the diagonal lines are the MP=p line. The ranges of the vertical axes are different across the plots of P1 and P2.

In Section 3 of Chi (2007), it is shown that if both p and u do not vary with m and the cdf of the p-value is strictly concave, then as m grows to infinity, the BFDR is always bounded below by

$$\beta_* = \frac{1-p}{1-p+p\lim_{x\to\infty}\rho(x)}, \tag{6.2}$$

where $\rho(x)$ in our notation is $(1+u)^{-1/2}d(x/\sqrt{1+u})/d(x)$ with limit $(1+u)^{\gamma/2}$. Thus β^* in (4.3) and β_* in (6.2) have the same expression, although they are derived in different contexts. Chi (2007) also proved that, under his setting, there is a critical value $\alpha_* > 0$ for the target FDR control level α. If $0 < \alpha < \alpha_*$, the power of a multiple testing procedure decays to 0 as $m \to \infty$ and the BFDR converges to β_*. Under our setting, we believe that the criticality phenomenon still exists with both α_* and β_* replaced by β^*_∞ defined in (6.1). In panel (c) of Figures 1 and 2, we plot the probability of type II error (P2) of the Benjamini-Hochberg procedure against FDR level α when $m = 10^6$. The solid vertical lines represent β^*_∞. In the plots for $C = 1, \gamma = 3$, and $C = 10$, P2 is very close to 1, which suggests that the power is close to 0.

The asymptotic optimal FDR level of the Benjamini-Hochberg procedure, α_∞, depends on the difficulty index C that is usually unknown. For practical use, according to our simulation results, when γ is large, $1/\log(m)$ is a good surrogate

for α_∞ in terms of risk ratios and misclassification probabilities. Although in some situations $1/\log(m)$ is not close to α_∞, the risk ratio does not considerably deviate from 1. It is shown in Figures 1 and 2 that the risk ratios are less sensitive to the choice of α for a larger C. Therefore, a smaller value, say 0.1, is a safe guess for C when γ is small. There could be more delicate methods to estimate C. To illustrate an example, let m_1 and m_0 denote the number of observations with absolute values in intervals $(b, +\infty)$ and (a_1, a_2), respectively. If b is large enough and a_1, a_2 are relatively small, by an idea similar to the one used to estimate π_0 in Chapter 4.5 of Efron (2012), we could assume that almost all the m_1 observations are signals and almost all the m_0 observations are noises. Then $m_1/m \approx 2p\{1 - D(b/\sigma_1)\}$, and $m_0/m \approx 2(1 - p)\{D(a_2/\sigma_0) - D(a_1/\sigma_0)\}$. Taking the ratio of these two and using the polynomial tail equivalence of the MPT distribution, we have

$$\frac{m_0}{m_1} \approx \frac{1 - p}{p}\left(\frac{\sigma_0}{\sigma_1}\right)^\gamma \left[\left(\frac{b}{a_1}\right)^\gamma - \left(\frac{b}{a_2}\right)^\gamma\right] \approx \delta^{-1}vu^{-\gamma/2}\left[\left(\frac{b}{a_1}\right)^\gamma - \left(\frac{b}{a_2}\right)^\gamma\right].$$

As $vu^{-\gamma/2} \to C_0$, C_0 could be estimated by

$$\widehat{C}_0 = \left(\delta_\infty \frac{m_0}{m_1}\right)\left[\left(\frac{b}{a_1}\right)^\gamma - \left(\frac{b}{a_2}\right)^\gamma\right]^{-1}. \tag{6.3}$$

Since C_0 is an increasing function in C, the estimate of C can be solved analytically or numerically depending on the form of d. A problem of this method is how to choose a_1, a_2. We want a_1 and a_2 to be small enough so that nearly all the observations in intervals $(-a_2, -a_1)$ and (a_1, a_2) are noises. At the same time, a_1 and a_2 should be large enough so that the polynomial approximation of $D(a_i/\sigma_0), i = 1, 2$, is accurate. According to some simulations not shown here, there is no simple solution to this problem.

As far as we know, theories of multiple testing problems for polynomial-tailed distributions have not been well developed in literature. It is unclear whether some of the interesting results for normal distributions still exist for polynomial-tailed distributions. For example, Bogdan et al. (2011) mention that their assumption $u \propto -\log(p)$ (obtained from $2\log(v)/u \to C$ when $\log(\delta) = o(\log(p))$) can be related to asymptotically least-favorable configurations for $l_0[p]$ balls, discussed in Abramovich et al. (2006). To examine whether our assumption $vu^{-\gamma/2} \to C_0$ has similar connection with minimax estimation, a vital question to be answered is what the configurations for polynomial-tailed distributions look like.

Global-local shrinkage priors have received much attention recently in Bayesian

analysis. Ghosh et al. (2016) showed that a multiple testing procedure based on a group of global-local shrinkage priors can asymptotically achieve the oracle Bayes risk up to a multiplicative constant. In the same vein, we would like to examine, in future work, whether and how global-local shrinkage priors can be used for polynomial-tailed distributions.

Supplementary Materials

The online supplementary material includes proofs of our results.

Acknowledgment

References

Abramovich, F., Benjamini, Y., Donoho, D. L. and Johnstone, I. M. (2006). Special invited lecture: Adapting to unknown sparsity by controlling the false discovery rate. *The Annals of Statistics* **34**, 584–653.

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 B* **57**, 289–300.

Benjamini, Y. and Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *Annals of Statistics* **29**, 1165–1188.

Bogdan, M., Chakrabarti, A., Frommlet, F. and Ghosh, J. K. (2011). Asymptotic Bayes-optimality under sparsity of some multiple testing procedures. *The Annals of Statistics* **39**, 1551–1579.

Bogdan, M., Ghosh, J. K. and Tokdar, S. T. (2008). A comparison of the Benjamini-Hochberg procedure with some Bayesian rules for multiple testing. In *Beyond Parametrics in Interdisciplinary Research: Festschrift in Honor of Professor Pranab K. Sen*, 211–230. Institute of Mathematical Statistics.

Chi, Z. (2007). On the performance of FDR control: constraints and a partial solution. *The Annals of Statistics* **35**, 1409–1431.

Donoho, D. and Jin, J. (2006). Asymptotic minimaxity of false discovery rate thresholding for sparse exponential data. *The Annals of Statistics* **34**, 2980–3018.

Efron, B. (2012). *Large-scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction*, Volume 1. Cambridge University Press.

Efron, B. and Tibshirani, R. (2002). Empirical Bayes methods and false discovery rates for microarrays. *Genetic Epidemiology* **23**, 70–86.

Frommlet, F. and Bogdan, M. (2013). Some optimality properties of FDR controlling rules under sparsity. *Electronic Journal of Statistics* **7**, 1328–1368.

Gavrilov, Y., Benjamini, Y. and Sarkar, S. K. (2009). An adaptive step-down procedure with proven FDR control under independence. *The Annals of Statistics* **37**, 619–629.

Genovese, C. and Wasserman, L. (2002). Operating characteristics and extensions of the false discovery rate procedure. *Journal of the Royal Statistical Society B* **64**, 499–517.

Genovese, C. and Wasserman, L. (2004). A stochastic process approach to false discovery control. *Annals of Statistics* **32**, 1035–1061.

Ghosh, P., Tang, X., Ghosh, M. and Chakrabarti, A. (2016). Asymptotic properties of Bayes risk of a general class of shrinkage priors in multiple hypothesis testing under sparsity. *Bayesian Analysis* **11**, 753–796.

Karlin, S. and Rubin, H. (1956). The theory of decision procedures for distributions with monotone likelihood ratio. *The Annals of Mathematical Statistics* **27**, 272–299.

Lehmann, E. L. (1957a). A theory of some multiple decision problems, I. *The Annals of Mathematical Statistics* **28**, 1–25.

Lehmann, E. L. (1957b). A theory of some multiple decision problems. II. *The Annals of Mathematical Statistics* **28**, 547–572.

Neuvial, P. and Roquain, E. (2012). On false discovery rate thresholding for classification under sparsity. *The Annals of Statistics* **40**, 2572–2600.

Sarkar, S. K. (2002). Some results on false discovery rate in stepwise multiple testing procedures. *Annals of Statistics* **30**, 239–257.

Storey, J. D. (2002). A direct approach to false discovery rates. *Journal of the Royal Statistical Society B* **64**, 479–498.

Storey, J. D. (2003). The positive false discovery rate: A Bayesian interpretation and the q-value. *Annals of Statistics* **31**, 2013–2035.

Department of Statistics, University of Florida, Gainesville, Florida 32611-8545, U.S.A.

E-mail: xytang@stat.ufl.edu

School of Statistics & Center of Statistical Research, Southwestern University of Finance and Economics, Chengdu, Sichuan, China.

E-mail: likec@swufe.edu.cn

Department of Statistics, University of Florida, Gainesville, Florida 32611-8545, U.S.A.

E-mail: ghoshm@stat.ufl.edu